



PVLDB

Proceedings of the VLDB Endowment

Editor-in-Chief:

H. V. Jagadish

Steering Committee:

Serge Abiteboul, Peter Apers, Philip Bernstein, Elisa Bertino, Peter Buneman, Martin Kersten, Z. Meral Ozsoyuglu



36th International Conference on

Very Large Data Bases

13th to 17th Sep 2010

Singapore, Singapore

Volume 3, No. 1,2 Sep 2010.

**Proceedings of the 36th International
Conference on Very Large Data Bases,
Singapore, Singapore**

Guest Editors:

Elisa Bertino, Paolo Atzeni, Kian Lee Tan, Yi Chen, Y. C. Tay

SAP
EXHIBIT 1009

PVLDB

Proceedings of the VLDB Endowment

Volume 3, No. 1, September 2010.

Proceedings of the 36th International Conference on Very Large Data Bases, Singapore.

Copyright 2010 VLDB Endowment

Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for components of this work owned by others than VLDB Endowment must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists requires prior specific permission and/or a fee. Request permission to republish from PVLDB under email: info@vldb.org.

Volume 3, Number 1: Full Papers

Pages iii - I and 1 - 1381

ISSN 2150-8097, September 2010.

Additional copies only online at: portal.acm.org and www.vldb.org

PVLDB

Proceedings of the VLDB Endowment

Volume 3, No. 2, September 2010.

Proceedings of the 36th International Conference on Very Large Data Bases, Singapore.

Copyright 2010 VLDB Endowment

Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for components of this work owned by others than VLDB Endowment must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists requires prior specific permission and/or a fee. Request permission to republish from PVLDB under email: info@vldb.org.

Volume 3, Number 2: Additional Conference Contributions

Pages 1382 - 1665

ISSN 2150-8097, September 2010.

Additional copies only online at: portal.acm.org and www.vldb.org

WELCOME MESSAGE FROM GENERAL CHAIRS

On behalf of the conference committee for VLDB 2010, it is our pleasure to welcome you to Singapore for the 36th International Conference on Very Large Data Bases. This is the second time that this conference is being held in Singapore; the first time was 26 years ago in 1984. The conference is jointly organized by the National University of Singapore, Nanyang Technological University, Singapore Management University, and the VLDB Endowment.

VLDB is a premier annual international forum for data management and database researchers, vendors, practitioners, application developers, and users. The conference will feature research talks, industry presentations, panels, tutorials, demonstrations, and 11 workshops. It will cover current issues in data management, database and information systems research. Data management and databases remain among the main technological cornerstones of emerging applications of the twenty-first century.

Singapore is situated at the southern tip of the Malaysian Peninsula in South-East Asia. It is a small and prosperous cosmopolitan nation that offers visitors the cultural diversity of its four main ethnic groups, namely, Chinese, Malays, Indians and Eurasians. Singapore is a vibrant Science & Technology hub where several universities and numerous public and corporate research institutes perform high-quality research over a broad range of disciplinary and cross-disciplinary areas. Singapore benefits from a knowledge-based economy with a national R&D spending target of at least 3% of GDP by 2010. With museums, theme parks, beaches, nature parks and two new integrated resorts as well as bustling shopping and dining havens, Singapore is an excellent place to unwind.

The conference would not have been possible if not for the efforts of many people. Thanks are due to the Technical Program Chair - Dr. Elisa BERTINO, and her team of track program chairs – Dr. Kian Lee TAN, Dr. Paolo ATZENI, Dr. Philip S. YU, and Dr. Masaru KITSUREGAWA and their PC members for producing an exciting technical programme. Thanks also to the efforts of the workshop co-chairs – Dr. Amol DESHPANDE, Dr. Zachary G. IVES, and Dr. Anthony Kum Hoe TUNG, workshops’ organizers and their committee members, the VLDB Endowment Liaison – Dr. David LOMET, the Local Arrangement Co-Chairs – Dr. Sourav S BHOWMICK, and Dr. Mong Li LEE, the Sponsorship Co-Chairs – Dr. Dan SUCIU and Dr. Limsoon WONG, the Publicity co-chairs – Dr. Stéphane BRESSAN, Dr. Umeshwar DAYAL, and Dr. Christian S. JENSEN, the proceedings co-chairs – Dr. Y.C. TAY and Dr. Yi CHEN, the Registration Chair – Dr. Chee Yong CHAN, the Treasurer – Dr. Wynne HSU, webmaster – Mr. Zaini MOHD, the Conference Secretary – Ms. Siew Foong HO, and other committee chairs, co-chairs, and members.

We are grateful for the generous support of our sponsors, Platinum Sponsors - Microsoft and Oracle; Gold Sponsors - Google, IBM, SAP, and Sybase; Silver Sponsors - ADSC, Yahoo!, TPC, and HP; Supporter - Sun Microsystems. We also wish to thank the following supporting organizations - A*Star and Singapore Tourism Board; supporting societies - Korean Institute of Information Scientists and Engineers, COMAD (India), The Database Society of Japan, and the China Computer Federation Database Technical Society.

Last but not least, our sincere thanks go to the authors of the papers, the speakers, and all the participants of VLDB 2010 who have made this conference a resounding success.

Welcome and enjoy the conference and have a good time in Singapore!

Tok Wang LING, National University of Singapore, Singapore
Krithi RAMAMRITHAM, Indian Institute of Technology Bombay, India
VLDB 2010 General Chairs

MESSAGE FROM THE PROGRAM CHAIRS

A warm welcome to the 36th International Conference on Very Large Data Bases (VLDB) in Singapore. The annual VLDB Conference is the premier international venue for the dissemination and exchange of research advances and ideas in the broad area of data management, both in the development and application of new technology. We hope you find both the conference and its location enjoyable.

This year's edition of VLDB received 686 submissions overall. For the research tracks: The Core Database Technology track received 280 submissions, of these 19 were rejected without review because of formatting violations and 48 were accepted. The acceptance rate computed with respect to the reviewed papers is 18.4%.

The Infrastructure for Information Systems track received 215 submissions, of these 11 were rejected without review because of formatting violations and 33 were accepted. The acceptance rate computed with respect to the reviewed papers is 16.1%.

In addition to the traditional program committees, papers were also reviewed through a "journal track". 119 submissions were received in this track during this year, of which 30 were accepted. In addition, there were 6 acceptances from other submission periods.

In addition, the Industrial Applications and Experience track received 51 submissions, of which 12 (23.5 %) were accepted.

The Demonstrations track received 89 submissions, of which 32 (36 %) were accepted.

The conference program also includes two keynote talks. We are very fortunate that two outstanding researchers have accepted our invitation to serve as keynote speakers. Divesh Srivastava will give the first keynote talk, focusing on real time data analysis, whereas Paul Matsudaira will give the second keynote talk, focusing on the challenges of managing very large 3D+ datasets resulting from high-end biological images.

The conference is complemented by a strong program of 11 workshops selected by Amol Deshpande, University of Maryland, Zachary G. Ives, University of Pennsylvania, and Anthony Kum Hoe Tung, National University of Singapore, Singapore. Of these, the VLDB PhD workshop continues to be encouraged and embraced by the VLDB Program Chairs. It was chaired this year by John Mylopoulos, University of Toronto, Lizhu Zhou, Tsinghua University, and Xiaofang Zhou, University of Queensland.

The Tutorial Program Chairs, Guozhu Dong, Wright State University, and Hwee Hwa Pang, Singapore Management University, selected 6 tutorials covering a wide variety of topics ranging from well established topics to more novel topics, such as large-scale data in cloud computing.

The Panel Chairs, Marianne Winslett, University of Illinois at Urbana-Champaign, and Jeffrey Xu Yu, The Chinese University of Hong Kong, selected two panels addressing exciting questions in our field. This year's edition of the conference had two Awards Committees.

The Best Paper Award Committee consisted of Paolo Atzeni, Università Roma tre, Elisa Bertino, Purdue University, Christian S. Jensen, Aalborg University, Meral Özsoyoglu, Case Western Reserve University, Kian-Lee Tan, National University of Singapore, Yufei Tao, Chinese University of Hong Kong.

The 10-Ten-Year Best Paper Award Committee consisted of Stefano Ceri, Politecnico di Milano, M. Tamer Özsu, University of Waterloo, Canada, and Kyu-Young Whang, KAIST. The 10-year Best Paper award this year is given to two papers: "Don't be lazy, be consistent: Postgres-R, a New Way to Implement Database Replication" by Bettina Kemme and Gustavo Alonso, and "Evolution of the Web and Implications for an Incremental Crawler" by Junghoo Cho and Hector Garcia-Molina.

Assembling the program of VLDB 2010 has been an effort that has involved a large number of members of our community.

We would like to thank first of all the authors who submitted top quality papers addressing many exciting and important research topics. After all, it is their work which makes the conference program. We are grateful to the PC members and external reviewers for the outstanding job with paper review. We appreciated the quality of the reviews, with many suggestions and comments given to the authors, and the active participation to the paper discussions.

WELCOME TO VLDB 2010

We would also like to deeply thank the proceedings editors, Yi Chen, Arizona State University, USA and Y.C. Tay, National University of Singapore, for all their hard work in assembling the proceedings and making sure that all papers were formatted according to the guidelines.

Last but not least, we would like to thank the VLDB General Co-Chairs, Tok Wang Ling, National University of Singapore, and Krithi Ramamritham, Indian Institute of Technology Bombay, for the constant advice and supervision.

Enjoy VLDB 2010!!

Elisa Bertino, Purdue University, USA
Technical Program Chair

Kian Lee Tan, National University of Singapore, Singapore
Core Database Technology Track Chair

Paolo Atzeni, Università Roma Tre, Italy
Infrastructure for Information Systems Track Chair

Masaru Kitsuregawa, University of Tokyo, Japan
Experiments and Analyses Track Chair

Philip S. Yu, University of Illinois at Chicago, USA
Industrial Applications and Experience Track Chair

Ralf Hartmut Guting, University of Hagen, Germany
Wang-Chien Lee, The Pennsylvania State University, USA
Ee-Peng Lim, Singapore Management University, Singapore
Xiaofeng Meng, Remin University, China
Demonstrations Track Co-Chairs

H.V. Jagadish, Univ. Michigan
Editor-in-Chief, Proceedings of the VLDB Endowment

ORGANIZING COMMITTEE

General Co-Chairs

Tok Wang LING
National University of Singapore, Singapore

Krithi RAMAMRITHAM
Indian Institute of Technology Bombay, India

Technical Program Chair

Elisa BERTINO
Purdue University, USA

Core Database Technology Program Chair

Kian Lee TAN
National University of Singapore, Singapore

Infrastructure for Information Systems Program Chair

Paolo ATZENI
Università Roma Tre, Italy

Industrial Applications and Experience Program Chair

Philip S. YU
University of Illinois at Chicago, USA

Experiments and Analyses Program Chair

Masaru KITSUREGAWA
University of Tokyo, Japan

Demonstrations Program Co-Chair

Ralf Hartmut GÜTING
University of Hagen, Germany

Wang-Chien LEE
The Pennsylvania State University, USA

Ee-Peng LIM
Singapore Management University, Singapore

Xiaofeng MENG
Remin University, China

PhD Program Co-chairs

John MYLOPOULOS
University of Toronto, Canada

Lizhu ZHOU
Tsinghua University, China

Xiaofang ZHOU
University of Queensland, Australia

Workshop Program Co-Chairs

Amol DESHPANDE
University of Maryland, USA

Zachary G. IVES
University of Pennsylvania, USA

Anthony Kum Hoe TUNG
National University of Singapore, Singapore

Tutorial Program Co-Chairs

Guozhu DONG
Wright State University, USA

Hwee Hwa PANG
Singapore Management University, Singapore

Panel Program Co-Chairs

Marianne WINSLETT
University of Illinois at Urbana-Champaign, USA

Jeffrey Xu YU
The Chinese University of Hong Kong, China

Ten-Year Best Paper Award Committee

Stefano CERI
Politecnico di Milano, Italy

M. Tamer ÖZSU
University of Waterloo, Canada (Chair)

Kyu-Young WHANG
KAIST, South Korea

Best Paper Award Committee

Paolo ATZENI
Università Roma Tre, Italy

Elisa BERTINO
Purdue University, USA (Chair)

Christian S. JENSEN
Aalborg University, Denmark

Meral ÖZSOYOĞLU

Case Western Reserve University, USA

Kian-Lee TAN

National University of Singapore, Singapore

Yufei TAO

Chinese University of Hong Kong, China

VLDB Endowment Liaison

David LOMET

Microsoft Research, USA

Geographical Area Coordinators

(Americas)

Jiawei HAN

University of Illinois at Urbana-Champaign, USA

(Asia-Pacific)

Kyu-Young WHANG

KAIST, South Korea

(Europe, Mideast, and Africa)

Timos SELLIS

National Tech. University of Athens, Greece

National Database Society Liaisons, Asia Pacific

(CCF-DBS)

Jianzhong LI

Harbin Institute of Technology, China,

(DB India)

Srinath SRINIVASA

IIIT Bangalore, India

(The Database Society of Japan (DBSJ))

Haruo YOKOTA

Tokyo Institute of Technology, Japan

(KIISE (Korean Institute of Information Scientists and Engineers) Database Society)

Bong Hee HONG

Pusan National University, South Korea

Local Arrangements Co-Chairs

Sourav S BHOWMICK

Nanyang Technological University, Singapore

Mong Li LEE
National University of Singapore, Singapore

Sponsorship Co-Chairs

Limsoon WONG
National University of Singapore, Singapore

Dan SUCIU
University of Washington, USA

Publicity Co-Chairs

Stéphane BRESSAN
National University of Singapore, Singapore

Umeshwar DAYAL
HP Labs, USA

Christian S. JENSEN
Aalborg University, Denmark

Proceedings Co-Chairs

Yi CHEN
Arizona State University, USA

Y.C. TAY
National University of Singapore, Singapore

Registration Chair

Chee Yong CHAN
National University of Singapore, Singapore

Treasurer

Wynne HSU
National University of Singapore, Singapore

Web Management

Stéphane BRESSAN
National University of Singapore, Singapore

Zaini MOHD
National University of Singapore, Singapore

Conference Secretariat

Siew Foong HO
National University of Singapore, Singapore

PROGRAM COMMITTEES

Core Database Technology

Kian Lee TAN, **Program Chair** (National University of Singapore, Singapore)

Daniel ABADI (Yale University, USA)

Ashraf ABOULNAGA (University of Waterloo, Canada)

Laurent AMSALEG (IRISA-CNRS, France)

Walid AREF (Purdue University, USA)

Sourav BHOWMICK (Nanyang Technological University, Singapore)

Angela BONIFATI (Icar-CNR, Italy)

Peter BONCZ (CWI, The Netherlands)

Philippe BONNET (University of Copenhagen, Denmark)

Luc BOUGANIM (INRIA, France)

Stéphane BRESSAN (National University of Singapore, Singapore)

Nicolas BRUNO (Microsoft Research, USA)

Barbara CATANIA (Università di Genova, Italy)

Chee Yong CHAN (National University of Singapore, Singapore)

Ugur CETINTEMEL (Brown University, USA)

Lei CHEN (Hong Kong University of Science and Technology, China)

Shimin CHEN (Intel Research Pittsburgh, USA)

Reynold CHENG (University of Hong Kong, China)

Brian COOPER (Yahoo! Research, USA)

Bin CUI (Peking University, China)

Gautam DAS (University of Texas, Arlington, USA)

Amol DESHPANDE (University of Maryland, USA)

Alin DEUTSCH (University of California, San Diego, USA)

Yanlei DIAO (University of Massachusetts, Amherst, USA)

Wenfei FAN (University of Edinburgh, UK)

Alan FEKETE (University of Sydney, Australia)

Elena FERRARI (University of Insubria, Italy)

Shel FINKELSTEIN (SAP, USA)

Peter FISCHER (ETH Zurich, Switzerland)

Minos GAROFALAKIS (Technical University of Crete, Greece)

Johannes GEHRKE (Cornell University, USA)

Gabriel GHINITA (Purdue University)

Leo GIAKOUMAKIS (Microsoft, USA)

Torsten GRUST (Technische Universität München, Germany)

Jarek GRYZ (York University, USA)

Wook-Shin HAN (Kyungpook National University, Korea)

Jayant HARITSA (Indian Institute of Science, India)

Stavros HARIZOPOULOS (HP Labs, Palo Alto, CA)

Zachary IVES (University of Pennsylvania, USA)

Christian S. JENSEN (Aalborg University, Denmark)

Chris JEMAIN (University of Florida / Rice University, USA)

Panos KALNIS (KAUST, Saudi Arabia)

Panagiotis KARRAS (National University of Singapore)

Murat KANTARCIOGLU (University of Texas, Dallas, USA)

Alfons KEMPER (Technische Universität München, Germany)

George KOLLIOS (Boston University, USA)

Yannis KOTIDIS (Athens University of Economics and Business, Greece)

Ramamohanarao KOTAGIRI (University of Melbourne, Australia)

Nick KOUDAS (University of Toronto, Canada)

Laks LAKSHMANAN (The University of British Columbia, Canada)

Chen LI (University of California, Irvine & BiMaple)

Jianzhong LI (Harbin Institute of Technology, China)

Lipyeow LIM (IBM T J Watson Research Center, USA)

David LOMET (Microsoft Research, USA)

Boon Thau LOO (University of Pennsylvania, USA)

Qiong LUO (Hong Kong University of Science and Technology, China)

Sharad MEHROTRA (University of California, Irvine, USA)

Sergey MELNIK (Google, USA)

Gerome MIKLAU (University of Massachusetts, Amherst, USA)

Guido MOERKOTTE (University of Mannheim, Germany)

Bongki MOON (University of Arizona, USA)

Mario NASCIMENTO (University of Alberta, Canada)

Wolfgang NEJDL (University of Hannover, Germany)

Christopher OLSTON (Yahoo! Research, USA)

M. Tamer ÖZSU (University of Waterloo, Canada)

HweeHwa PANG (Singapore Management University, Singapore)

Dimitris PAPADIAS (Hong Kong University of Science and Technology, China)

Evaggelia PITOURA (University of Ioannina, Greece)

Neoklis POLYZOTIS (University of California, Santa Cruz, USA)

Rachel POTTINGER (The University of British Columbia, Canada)
Sunil PRABHAKAR (Purdue University, USA)
Rajeev RASTOGI (Yahoo! Research, India)
Tore RISCH (Uppsala University, Sweden)
Kenneth ROSS (Columbia University, USA)
Elke RUNDENSTEINER (Worcester Polytechnic Institute, USA)
Kai-Uwe SATTLER (Technische Universität Ilmenau, Germany)
Ralf SCHENKEL (Max Planck Institute, Germany)
Eugene SHEKITA (IBM Almaden Research Center, USA)
Hengtao SHEN (University of Queensland, Australia)
Kyuseok SHIM (Seoul National University, Korea)
Richard SNODGRASS (University of Arizona, USA)
Divesh SRIVASTAVA (AT&T, USA)
Ioana STANOI (IBM T.J. Watson Research Center, USA)
Dan SUCIU (University of Washington, USA)
S. SUDARSHAN (IIT Bombay, India)
Wang-Chiew TAN (University of California, Santa Cruz, USA)
Yufei TAO (Chinese University of Hong Kong, China)
Victor VIANU (University of California, San Diego, USA)
Gerhard WEIKUM (Max-Planck Institute, Germany)
Jun YANG (Duke University, USA)
Carlo ZANIOLO (University of California, Los Angeles, USA)
Donghui ZHANG (Microsoft, USA)
Sai WU (National University of Singapore, Singapore)
Aoying ZHOU (East China Normal University, China)
Yongluan ZHOU (University of Southern Denmark, Denmark)

Infrastructure for Information Systems

Paolo ATZENI, **Program Chair** (Università Roma Tre, Italy)

Karl ABERER (EPFL, Switzerland)

Foto AFRATI (NTUA, Greece)

Rakesh AGRAWAL (Microsoft, USA)

Suad ALAGIC (University of Southern Maine, USA)

Mike ATALLAH (Purdue University, USA)

Ricardo BAEZA-YATES (Yahoo! Research, Spain)

Elena BARALIS (Politecnico di Torino, Italy)

Denilson BARBOSA (University of Alberta, Canada)

Leopoldo BERTOSSI (Carleton University, Canada)

Boualem BENATALLAH (University of New South Wales, Australia)

Jan Van den BUSSCHE (Universiteit Hasselt, Belgium)

Jose BLAKELEY (Microsoft, USA)

Mike CAFARELLA (University of Washington, USA)

Michael CAREY (University of California, Irvine, USA)

Sharma CHAKRAVARTHY (The University of Texas at Arlington, USA)

Kevin CHANG (University of Illinois at Urbana-Champaign, USA)

Arbee CHEN (National Chengchi University, Taiwan)

Laura CHITICARIU (IBM Research, USA)

Jan CHOMICKI (University at Buffalo, USA)

Chris CLIFTON (Purdue University, USA)

Susan DAVIDSON (University of Pennsylvania, USA)

Umeshwar DAYAL (HP Labs, USA)

Lois DELCAMBRE (Portland State University, USA)

AnHai DOAN (University of Wisconsin, USA)

Jens DITTRICH (Saarland University, Germany)

Michael FRANKLIN (University of California, Berkeley, USA)

Johann-Christoph FREYTAG (Humboldt-Universität zu Berlin, Germany)

Aristides GIONIS (Yahoo! Research, Spain)

Jonathan GOLDSTEIN (Microsoft, USA)

Philip GOLLE (PARC, USA)

Luis GRAVANO (Columbia University, USA)

Ralf Hartmut GÜTING (Fernuniversität Hagen)

Alon HALEVY (Google, Inc., USA)

Mauricio HERNANDEZ (IBM Research, USA)
Howard HO (IBM Research, USA)
Rick HULL (IBM Research)
Panos IPEIROTIS (New York University, USA)
Anastasios KEMENTSIETSIDIS (IBM Research, USA)
Maurice van KEULEN (University of Twente, The Netherlands)
Christoph KOCH (Cornell University, USA)
Hank KORTH (Lehigh University, USA)
Donald KOSSMANN (ETH Zürich, Switzerland)
Georgia KOUTRIKA (Stanford University, USA)
Peer KROEGER (Ludwig-Maximilians-Universitaet München, Germany)
Alexandros LABRINIDIS (University of Pittsburgh, USA)
Alberto LAENDER (Federal University of Minas Gerais, Brazil)
Maurizio LENZERINI (Sapienza, Università di Roma, Italy)
Frederick LOCHOVSKY (HKUST, Hong Kong)
Jiaheng LU (Renmin University of China)
Ioana MANOLESCU (INRIA, France)
Yannis MANOLOPOULOS (Aristotle University, Thessaloniki, Greece)
Giansalvatore MECCA (Università della Basilicata, Italy)
Paolo MERIALDO (Università Roma Tre, Italy)
Moirra NORRIE (ETH Zuerich, Switzerland)
Z. Meral ÖZSOYOĞLU (Case Western Reserve University, USA)
Themis PALPANAS (Università di Trento, Italy)
Jan PAREDAENS (University of Antwerp, Belgium)
Norman PATON (University of Manchester, UK)
Hamid PIRAHESH (IBM Research, USA)
Lucian POPA (IBM Research, USA)
Alexandra POULOVASSILIS (Birkbeck, University of London, UK)
Erhard RAHM (Universität Leipzig, Germany)
Matthias RENZ (Ludwig-Maximilians-Universität München, Germany)
Marc SCHOLL (Universität Konstanz)
Ning SHANG (Purdue University, USA)
Jayavel SHANMUGASUNDARAM (Yahoo! Research, USA)
Dennis SHASHA (New York University, USA)
Oded SHMUELI (Technion, Israel)
Jianwen SU (University of California, Santa Barbara, USA)

Val TANNEN (University of Pennsylvania, USA)

Jens TEUBNER (ETH, Switzerland)

Yannis VELEGRAKIS (Università di Trento, Italy)

Gottfried VOSSEN (WWU Muenster, Germany)

Ouri WOLFSON (University of Illinois at Chicago)

Peter WOOD (Birkbeck, University of London, UK)

Cong YU (Yahoo! Research, USA)

Stan ZDONIK (Brown University, USA)

Pavel ZEZULA (Masaryk University, Czech Republic)

Zhenjie ZHANG (National University of Singapore, Singapore)

Industrial Applications and Experience

Philip S. YU, **Program Chair** (University of Illinois at Chicago, USA)

Henrique ANDRADE (IBM USA)

Michael BRODIE (Verizon, USA)

Malu CASTELLANOS (HP Labs, USA)

Edward CHANG (Google, China)

Yuan-chi CHANG (IBM, USA)

Ming-Syan CHEN (NTU, Taiwan)

Yun CHI (NEC USA)

Graham CORMODE (AT&T Labs, USA)

Takeshi FUKUDA (IBM, Japan)

Bugra GEDIK (IBM, USA)

Alex GORELIK (Exeros, USA)

Goetz GRAEFE (HP Labs, USA)

Hui I HSIAO (IBM USA)

Meichun HSU (HP Labs, China)

Dean JACOBS (SAP, Germany)

Chiao LIU (Microsoft, USA)

Bo LONG (Yahoo, USA)

Wei-Ying MA (Microsoft, China)

Ravi MURTHY (Oracle, USA)

Makoto ONIZUKA (NTT Labs, Japan)

Alex SZALAY (JHU, USA)

Ming-Chien SHAN (SAP, China)

Evimaria TERZI (IBM USA)

Michale VLACHOS (IBM, Switzerland)

Haixun WANG (Microsoft, China)

Xiaoxin YIN (Microsoft USA)

Experiments and Analyses

Masaru KITSUREGAWA, **Program Chair** (University of Tokyo, Japan)

Kasim CANDAN (ASU, USA)

Umeshwar DAYAL (HP Labs, USA)

H JAGADISH (MU, USA)

Hiroyuki KITAGAWA (Tsukuba University, Japan)

Paul LARSON (Microsoft Research, USA)

Wen-Syan LI (SAP, China)

Xuemin LIN (U of New South Wales .Australia)

Ling LIU (George Tech. , USA)

Sanjay MADRIA (UMR, USA)

Volker MARKL (Tu-Nerlin, Germany)

Beng Chin OOI (NUS, Singapore)

Cyrus SHAHABI (USC, USA)

Kyu-Young WHANG (KAIST Korea)

Kung Lung WU (IBM research Watson, USA)

Jefferey YU (CUHK, HongKong)

Demo PC members

Jidong CHEN (EMC China Lab, China)

Lei CHEN (Hong Kong University of Science and Technology, China)

Xin DONG (AT&T Labs Research, USA)

Bugra GEDIK (IBM Watson Research Center, USA)

Fosca GIANNOTTI (Istituto di Scienza e Tecnologie dell'Informazione "Alessandro Faedo", Italy)

Hyoil HAN (LeMoyne-Owen College, USA)

Wook-Shin HAN (Kyung-Pook National University, Korea)

San-Yih HWANG (National Sun Yat Sen University, Taiwan)

Daxing JIAN (Microsoft Research Asia, China)

Vana KALOGERAKI (Univ. of California, Riverside, USA)

Daniel KIFER (Pennsylvania State University, USA)

Ulf LESER (Humboldt University, Berlin, Germany)

Weiyi MENG (Binghamton University, USA)

Mohamed F. MOKBEL (University of Minnesota, USA)

Felix NAUMANN (Hasso Plattner Institut, Germany)

Frank NEVEN (Hasselt University, Belgium)

Daniela NICKLAS (Uni Oldenburg, Germany)

Alex NTOULAS (Microsoft Search Labs, USA)

Yasushi SAKURAI (NTT Communication Science Laboratories, Japan)

Peter SCHEUERMANN (Northwestern University, USA)

Katsumi TANAKA (Kyoto University, Japan)

Nesime TATBUL (ETH Zuerich, Switzerland)

Jens TEUBNER (ETH Zürich Switzerland)

Jianyong WANG (Tsinghua University, China)

Li XIONG (Emory University, USA)

Denian YANG (Academic Sinica, Taiwan)

Ge YU (Northeastern University, China)

Baihua ZHENG (Singapore Management University, Singapore)

PhD Workshop

John MYLOPOULOS, **Program Co-Chair** (University of Toronto, Canada)

Lizhu ZHOU, **Program Co-Chair** (TsinghuaUniversity, China)

Xiaofang ZHOU, **Program Co-Chair** (University of Queensland, Australia)

Lei CHEN (Hong Kong University of Science and Technology, China)

Ke DENG (University of Queensland, Australia)

Alan FEKETE (University of Sydney, Australia)

Ihab ILYAS (University of Waterloo, Canada)

Peer KRÖGER(Ludwig-Maximilians-Universität München, Germany)

Xuemin LIN (University of New South Wales, Australia)

Amélie MARIAN (RutgersUniversity, USA)

Miyuki NAKANO (University of Tokyo, Japan)

Yoshiharu ISHIKAWA (Nagoya University, Japan)

Jian PEI (Simon Fraser University, Canada)

Jagan SANKARANARAYANAN (University of Maryland, USA)

Divesh SRIVASTAVA (AT&T, USA)

Fabio SCHREIBER (Politecnico di Milano, Italy)

Mohamed SHARAF (University of Queensland, Australia)

Kyuseok SHIM (National Seoul University, Korea)

Anthony TUNG (National University of Singapore, Singapore)

Yannis VELEGRAKIS (University of Trento, Italy)

Kyu-Young WHANG (KAIST, Korea)

Ke YI (Hong Kong University of Science and Technology, China)

Jeffrey X. YU (Chinese University of Hong Kong, China)

Aoying ZHOU (East China Normal University, China)

EXTERNAL REVIEWERS

Core Database Technology

Azza Abouzied (Yale University, USA)	Feifei Li (Florida State University, USA)
Bogdan Alexe (UC Santa Cruz, USA)	Guoliang Li (Tsinghua University, China)
Sattam Alsubaiee (UC Irvine, USA)	Ninghui Li (Purdue University, USA)
Kamil Bajda-Pawlikowski (Yale University, USA)	Xiang Lian (HKUST, China)
Alexander Behm (UC Irvine, USA)	Changbin Liu (University of Pennsylvania, USA)
Vinayak Borkar (UC Irvine, USA)	Alexandra Meliou (University of Washington, USA)
Mustafa Canim (University of Texas at Dallas, USA)	Chaitanya Mishra (Facebook, USA)
Barbara Carminati (University of Insubria, Italy)	Hamid Mousavi (UCLA/CSD, USA)
Lei Chen (HKUST, China)	Barzan Mozafari (UCLA, USA)
Jiefeng Cheng (The University of Hong Kong, China)	Kunal Mukerjee (Microsoft, USA)
Chenyun Dai (Purdue University, USA)	Rimma Nehme (Microsoft Research, USA)
Christos Doukeridis (NTNU, Norway)	Przemyslaw Pawluk (York University, USA)
Campbell Fraser (Microsoft, USA)	Weixiong Rao (HKUST, China)
Craig Freedman (Microsoft, USA)	Vibhor Rastogi (University of Washington, USA)
Sébastien Gambs (Irisa, France)	Nikos Sarkas (University of Toronto, Canada)
Wolfgang Gatterbauer (University of Washington, USA)	Karl Schnaitter (UC Santa Cruz, USA)
Nikos Giatrakos (University of Piraeus, Greece)	Warren Shen (Google, USA)
Parke Godfrey (York University, Canada)	Balder ten-Cate (UC Santa Cruz, USA)
Ravindra Guravannavar (IIT Hyderabad, India)	dimitris Tsirogiannis (University of Toronto, Canada)
Michael Hay (University of Massachusetts, USA)	Kostas Tzoumas (Aalborg University, Denmark)
Jiewen Huang (Yale University, USA)	Rares Vernica (UC Irvine, USA)
Xuegang Huang (Danske Bank, Denmark)	Akrivi Vlachou (NTNU, Norway)
Ali Inan (University of Texas at Dallas, USA)	Anduo Wang (University of Pennsylvania, USA)
Abhay Jha (University of Washington, USA)	Dingming Wu (Aalborg University, Denmark)
Shengyue Ji (UC Irvine, USA)	Xiaokui Xiao (NTU, Singapore)
Dmitri Kalashnikov (UC Irvine, USA)	Mingxuan Yuan (HKUST, China)
Ben Kao (The University of Hong Kong, China)	Zhuoyao Zhang (University of Pennsylvania, USA)
Nikolay Laptev (UCLA, USA)	Wenchao Zhou (University of Pennsylvania, USA)

Infrastructure for Information Systems

Raman Adaikkalavan (Indiana University South Bend, USA)

Jose Aguilar Saborit (IBM Canada)

Alsayed Algergawy (Otto-von-Guericke-Universität Magdeburg, Germany)

Lory Al Moakar (University of Pittsburgh, USA)

Alexandru Arion (EPFL, Switzerland)

Artin Avanes (Humboldt-Universität zu Berlin)

Daniel Ayala (University of Illinois at Chicago, USA)

Zhuowei Bao (University of Pennsylvania, USA)

Michal Batko (Masaryk University, Czech Republic)

Hila Becker (Columbia University, USA)

Alexander Behm (University of California, Irvine, USA)

Thomas Bernecker (Ludwig-Maximilians-University Munich, Germany)

Vasudha Bhatnagar (University of Delhi, India)

Joel Booth (University of Illinois at Chicago, USA)

Vinayak Borkar (University of California, Irvine, USA)

Loreto Bravo (University of Concepcion, Chile)

David Briggs (University of Southern Maine)

Yuanzhe Cai (The University of Texas at Arlington, USA)

Yu Cao (National University of Singapore)

Xiaoyong Chai (University of Wisconsin, Madison, USA)

Bee-Chung Chen (Yahoo! Research, USA)

Ting Chen (University of Wisconsin, Madison, USA)

Xinxing Chen (Renmin University of China)

En Cheng (Case Western Reserve University, USA)

Tao Cheng (University of Illinois at UrbanaChampaign, USA)

Paolo Ciaccia (Università di Bologna, Italy)

Alexander Connor (University of Pittsburgh, USA)

Altigran S. da Silva (Universidade Federal do Amazonas, Brazil)

Alin Deutsch (University of California, San Diego)

Vlastislav Dohnal (Masaryk University, Czech Republic)

Tobias Emrich (Ludwig-Maximilians-University Munich, Germany)

Stefan Endrullis (Universität Leipzig, Germany)

Yao-Chung Fan (National Tsing Hua University, Taiwan)

Roxana Gheorghiu (University of Pittsburgh, USA)

George Giannakopoulos (Università di Trento, Italy)

Chaitanya Gokhale (University of Wisconsin, Madison, USA)

Marcos A. Gonçalves (Federal University of Minas Gerais, Brazil)

Hector Gonzalez (Google, Inc., USA)

Franz Graf (Ludwig-Maximilians-University Munich, Germany)

Todd J. Green (University of Pennsylvania)

Chetan Gupta (HP Labs, USA)

Till Haselmann (WWU Muenster, Germany)

Ruwen Hess (Microsoft Corporation, USA)

Stefan Hildenbrand (ETH Zürich, Switzerland)

Marie Jacob (University of Pennsylvania)

HoYoung Jeung (EPFL, Switzerland)

Dustin Jiang (The University of Texas at Arlington, USA)

Wei Jiang (Missouri University of Science & Technology, USA)

Alekh Jindal (Saarland University, Germany)

Satyen Kale (Yahoo! Research, USA)

Grigoris Karvounarakis (University of Pennsylvania)

Verena Kantere (EPFL, Switzerland)

Jia-Ling Ko (National Taiwan Normal University, Taiwan)

Hanna Koepcke (Universität Leipzig, Germany)

Jens Lechtenbörger (WWU Muenster, Germany)

Guanling Lee (National Dong-Hwa University, Taiwan)

Chengkai Li (The University of Texas at Arlington, USA)

Rui Li (University of Illinois at Urbana-Champaign, USA)

Yunyao Li (IBM Research, USA)

Chen-Yi Lin (National Tsing Hua University, Taiwan)

Chunbin Lin (Renmin University of China, China)

Ying Lu (Renmin University of China, China)

Sol Ma (University of Illinois at Chicago, USA)

Ashwin Machanavajjhala (Cornell University, USA)

Jayant Madhavan (Google, Inc., USA)

Sabine Massmann (Universität Leipzig, Germany)

Marcus Fontoura (Yahoo! Research, USA)

Sharad Mehrotra (University of California, Irvine, USA)

Svilen Mihaylov (University of Pennsylvania, USA)

Zoltan Miklos (EPFL, Switzerland)

Faruck Morcos Gonzalez (University of Notre Dame, USA)

Mirella M. Moro (Federal University of Minas Gerais, Brazil)

Kyriakos Mouratidis (Singapore Management University, Singapore)

Panayiotis Neophytou (University of Pittsburgh, USA)

Ahmet Erhan Nergiz (Purdue University, USA)

David Novak (Masaryk University, Czech Republic)

Irene Ntoutsis (Ludwig-Maximilians-University Munich, Germany)

Tekin Ozsoyoglu (Case Western Reserve University, USA)

Paolo Papotti (Università Roma Tre, Italy)

Marco Patella (Università di Bologna, Italy)

Thao Pham (University of Pittsburgh, USA)

Jorge-Arnulfo Quiané-Ruiz (Saarland University, Germany)

Flavio Rizzolo (Carleton University, Canada)

Jose Antonio Rodriguez Mondejar

Royi Ronen (Technion, Israel)

George Roussos (Birkbeck, University of London, UK)

Sudeepa Roy (University of Pennsylvania, USA)

Yasushi Sakurai (NTT, Japan)

Mark Sandler (Google, USA)

Saket Sathe (EPFL, Switzerland)

Jörg Schad (Saarland University, Germany)

Erich Schubert (Ludwig-Maximilians-University Munich, Germany)

Matthias Schubert (Ludwig-Maximilians-University Munich, Germany)

Mohamed Shehab (University of North Carolina at Charlotte, USA)

Warren Shen (Google, Inc., USA)

Lila Shnaiderman (Technion, Israel)

Slawek Staworko (University Lille, France)

Julia Stoyanovich (University of Pennsylvania, USA)

Weifeng Su (United International College, Hong Kong)

Gianvito Summa (Università della Basilicata, Italy)

Piotr Szczurek (University of Illinois at Chicago, USA)

Aditya Telang (The University of Texas at Arlington, USA)

Gunnar Thies (WWU Muenster, Germany)

Marisa Thoma (Ludwig-Maximilians-University Munich, Germany)

Ba-Quy Vuong (University of Wisconsin, Madison, USA)

Xiaoheng Xie (HKUST, Hong Kong)

Bo Xu (University of Illinois at Chicago, USA)

En Tzu Wang (National Tsing Hua University, Taiwan)

Jiying Wang (City University of Hong Kong)

Song Wang (HP Labs, USA)

Yi-Hung Wu (Chung Yuan Christian University, Taiwan)

Lei Yang (Case Western Reserve University, USA)

Mianwei Zhou (University of Illinois at UrbanaChampaign, USA)

Hengzhi Zhong (University of Illinois at UrbanaChampaign, USA)

Andreas Züfle (Ludwig-Maximilians-University Munich, Germany)

Industrial Applications and Experience

Longbin Chen (Yahoo! Labs, USA)

Bin He (IBM, USA)

Lan Nie (Yahoo! Labs, USA)

Michalis Potamias (Boston University, USA)

Yuval Sherman (HP, Israel)

Hans Zeller (HP, USA)

Daniel Zhi (Microsoft Research, USA)

Dong Wang (Google, USA)

Robin Sun (Google, USA)

Xiance Si (Google, USA)

Experiments and Analyses

Muhammad Aamir Cheema (The University of New South Wales, Australia)

Bhuvan Bamba (Georgia Tech, USA)

Farnoush BanaeiKashani (USC, USA)

Qinyuan Feng (Georgia Tech, USA)

Bingsheng He (Chinese University of Hong Kong, Hong Kong, China)

Ali Khodaei (USC, USA)

Ali Jeff Khoshgozaran (USC, USA)

Shicong Meng (Georgia Tech, USA)

Gong Zhang (Georgia Tech, USA)

Ke Zhu (The University of New South Wales, Australia)

Demo

Xian Li (Binghamton University, USA)

Yu Jiang (Binghamton University, USA)

Dustin Lange (Hasso Plattner Institute, Potsdam, Germany)

Mohammed AbuJarour (Hasso Plattner Institute, Potsdam, Germany)

Christoph Böhm (Hasso Plattner Institute, Potsdam, Germany)

Johannes Lorey (Hasso Plattner Institute, Potsdam, Germany)

Hang Guo (EMC Research China, China)

Ziye Yang (EMC Research China, China)

Changliang Wang (VMware, China)

Wei Shen (Tsinghua University, China)

Xu Pu (Tsinghua University, China)

Qingyan Yang (Tsinghua University, China)

Lili Jiang (Lanzhou University, China)

Chuancong Gao (Tsinghua University, China)

Luciano Barbosa (AT&T Labs-Research, USA)

VLDB Endowment Board of Trustees

The VLDB Endowment now has a board of 21 elected trustees, who are the legal guardians of the Endowment's charter and activities. Trustees are elected for a six-year period; the election procedure is documented in the code of regulations available on the Endowment's Web site. The board is continuously renewed, with one third of its members being replaced every two years. The trustees are elected among internationally distinguished researchers and professionals in the field of database and information systems who have contributed to the objectives of the Endowment with dedication and distinction. All trustees provide voluntary service.

Executive

President:
Renée J. Miller

Vice-President:
Paolo Atzeni

Treasurer:
Michael J. Franklin

Secretary:
Beng Chin Ooi

Current trustees (as of Jan 2010):

Gustavo Alonso
Paolo Atzeni
Susan B. Davidson
Michael J. Franklin
Donald Kossmann
David B. Lomet
Renée J. Miller
Tova Milo
Beng Chin Ooi
Z. Meral Özsoyoğlu
Sunita Sarawagi
Anastasia Ailamaki
Sihem Amer-Yahia
Michael H. Böhlen
Michael J. Carey
Surajit Chaudhuri
Alon Y. Halevy
Volker Markl
Divesh Srivastava
S. Sudarshan
Kyu-Young Whang

Trustees Emeriti

Rakesh Agrawal
Peter Apers
Phil A. Bernstein
Elisa Bertino
Giampio Bracchi
Michael L. Brodie
Janis A. Bubenko
Peter Buneman
Alfonso Cardenas
Stefano Ceri
Sophie Cluet
Umesh Dayal
David J. DeWitt
Antonio L. Furtado
Hector Garcia-Molina
Georges Gardarin
Laura Haas
David K. Hsiao
Yannis E. Ioannidis
H. V. Jagadish
Christian S. Jensen
Raghu Ramakrishnan
Krithi Ramamritham
Gerhard Weikum
Stan Zdonik
Keith G Jeffery
Martin L. Kersten
Masaru Kitsuregawa
Peter C. Lockemann
Vincent Y. Lum
Stuart Madnick
Dennis McLeod
John Mylopoulos
Shamkant B. Navathe
Erich J. Neuhold
Maria E. Orlowska

M. Tamer Özsu
Alain Pirotte
Nick Roussopoulos
Ron Sacks-Davis
Hans-Jörg Schek
Gunter Schlageter
Timos Sellis
Arie Shoshani
Arne Solvberg
Stanley Su
Dennis Tsichritzis
Patrick Valduriez
Yannis Vassiliou
Herbert Weber
Jennifer Widom
S. Bing Yao
Amr El Abbadi
Johann-Christoph Freytag
Shojiro NISHIO

Table of Contents

ii Copyright Notice	xi Program Committees
iii Welcome from the General Chairs	xxi External Reviewers
v Message from the Program Chairs	xxviii VLDB Endowment
vii Organizing Committee	xxix Table of Contents

Keynotes

1

Enabling Real Time Data Analysis.

Divesh Srivastava (AT&T Labs Research, United States of America), Lukasz Golab (AT&T Labs Research, United States of America), Rick Greer (AT&T Labs Research, United States of America), Theodore Johnson (AT&T Labs Research, United States of America), Joseph Seidel (AT&T Labs Research, United States of America), Vladislav Shkapenyuk (AT&T Labs Research, United States of America), Oliver Spatscheck (AT&T Labs Research, United States of America), Jennifer Yates (AT&T Labs Research, United States of America).

3

High-End Biological Imaging Generates Very Large 3D+ and Dynamic Datasets

Paul Matsudaira (National University of Singapore, Republic of Singapore).

10-Year Best Paper Awards

Keynote Sessions

4

Dealing with Web Data: History and Look ahead

Junghoo Cho (University of California Los Angeles, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

5

Database Replication: a Tale of Research across Communities

Bettina Kemme (McGill University, Canada), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland).

Research Sessions

Database Security

13

Building Disclosure Risk Aware Query Optimizers for Relational Databases

Mustafa Canim (University of Texas at Dallas, United States of America), Murat Kantarcioglu (University of Texas at Dallas, United States of America), Bijit Hore (University of California Irvine, United States of America), Sharad Mehrotra (University of California Irvine, United States of America).

25

Secure Personal Data Servers: a Vision Paper

Tristan Allard (University of Versailles, France), Nicolas Anciaux (Institut National de Recherche en Informatique et Automatique, France), Luc Bouganim (Institut National de Recherche en Informatique et Automatique, France), Yanli Guo (Institut National de Recherche en Informatique et Automatique, France), Lionel Le Folgoc (Institut National de Recherche en Informatique et Automatique, France), Benjamin Nguyen (Institut National de Recherche en Informatique et Automatique, France), Philippe Pucheral (Institut National

de Recherche en Informatique et Automatique, France), Indrajit Ray (Colorado State University, United States of America), Indrakshi Ray (Colorado State University, United States of America), Shaoyi Yin (Institut National de Recherche en Informatique et Automatique, France).

36

PolicyReplay: Misconfiguration-Response Queries for Data Breach Reporting

Daniel Fabbri (University of Michigan, United States of America), Kristen LeFevre (University of Michigan, United States of America), Qiang Zhu (University of Michigan, United States of America).

Parallel and Distributed Databases

48

Schism: a Workload-Driven Approach to Database Replication and Partitioning

Carlo Curino (Massachusetts Institute of Technology, United States of America), Yang Zhang (Massachusetts Institute of Technology, United States of America), Evan Jones (Massachusetts Institute of Technology, United States of America), Samuel Madden (Massachusetts Institute of Technology, United States of America).

58

Ten Thousand SQLs: Parallel Keyword Queries Computing

Lu Qin (The Chinese University of Hong Kong, People's Republic of China), Jefferey Yu (The Chinese University of Hong Kong, People's Republic of China), Lijun Chang (The Chinese University of Hong Kong, People's Republic of China).

70

The Case for Determinism in Database Systems

Alexander Thomson (Yale University, United States of America), Daniel Abadi (Yale University, United States of America).

Data Exchange

81

MapMerge: Correlating Independent Schema Mappings

Bogdan Alexe (University of California Santa Cruz, United States of America), Mauricio Hernández (IBM Research, United States of America), Lucian Popa (IBM Almaden Research Center, United States of America), Wang-Chiew Tan (University of California Santa Cruz, United States of America).

93

Chase Termination: A Constraints Rewriting Approach

Francesca Spezzano (Università della Calabria, Italy), Sergio Greco (Università della Calabria, Italy).

105

Scalable Data Exchange with Functional Dependencies

Bruno Marnette (University of Oxford, United Kingdom), Giansalvatore Mecca (Università della Basilicata, Italy), Paolo Papotti (Università Roma Tre, Italy).

Database Services and Applications

117

Interactive Route Search in the Presence of Order Constraints

Roy Levin (Technion-Israel Institute of Technology, Israel), Yaron Kanza (Technion-Israel Institute of Technology, Israel), Eliyahu Safra, Yehoshua Sagiv (Hebrew University of Jerusalem, Israel).

129

Energy Management for MapReduce Clusters

Willis Lang (University of Wisconsin-Madison,

United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

140

Toward Scalable Keyword Search over Relational Data

Akanksha Baid (University of Wisconsin-Madison, United States of America), Ian Rae (University of Wisconsin-Madison, United States of America), Jiexing Li (University of Wisconsin-Madison, United States of America), AnHai Doan (University of Wisconsin-Madison, United States of America), Jeffrey Naughton (University of Wisconsin-Madison, United States of America).

Data Models and Languages

150

From Regular Expressions to Nested Words: Unifying Languages and Query Execution for Relational and XML Sequences

Barzan Mozafari (University of California Los Angeles, United States of America), Kai Zeng (University of California Los Angeles, United States of America), Carlo Zaniolo (University of California Los Angeles, United States of America).

162

Avalanche-Safe LINQ Compilation

Torsten Grust (Universität Tübingen, Germany), Jan Rittinger (Universität Tübingen, Germany), Tom Schreiber (Universität Tübingen, Germany).

173

Towards Certain Fixes with Editing Rules and Master Data

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Nan Tang (University of Edinburgh,

United Kingdom), Wenyan Yu (University of Edinburgh, United Kingdom).

Semantics

185

Explaining Missing Answers to SPJUA Queries

Melanie Herschel (Universität Tübingen, Germany), Mauricio Hernández (IBM Research, United States of America).

197

Sampling the Repairs of Functional Dependency Violations under Hard Constraints

George Beskales (University of Waterloo, Canada), Ihab Ilyas (University of Waterloo, Canada), Lukasz Golab (AT&T Labs - Research, United States of America).

208

Evaluating Entity Resolution Results

David Menestrina (Stanford University, United States of America), Steven Whang (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

Stream Databases

220

High-Performance Dynamic Pattern Matching over Disordered Streams

Badrish Chandramouli (Microsoft Research, United States of America), Jonathan Goldstein (Microsoft Research, United States of America), David Maier (Portland State University, United States of America).

232

SECRET: A Model for Analysis of the Execution Semantics of Stream Processing Systems

Irina Botan (Eidgenössische Technische Hochschule Zürich, Switzerland), Roozbeh Derakhshan (Eidgenössische Technische Hochschule

Zürich, Switzerland), Nihal Dindar (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research Center, United States of America), Renée Miller (University of Toronto, Canada), Nesime Tatbul (Eidgenössische Technische Hochschule Zürich, Switzerland).

244

Recognizing Patterns in Streams with Imprecise Timestamps

Haopeng Zhang (University of Massachusetts Amherst, United States of America), Yanlei Diao (University of Massachusetts Amherst, United States of America), Neil Immerman (University of Massachusetts Amherst, United States of America).

RDF and Graphs

256

x-RDF-3X: Fast Querying, High Update Rates, and Consistency for RDF Databases

Thomas Neuman (Technische Universität München, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

264

Graph Pattern Matching: From Intractable to Polynomial Time

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Nan Tang (University of Edinburgh, United Kingdom), Yinghui Wu (University of Edinburgh, United Kingdom), Yunpeng Wu (National University of Defense Technology, People's Republic of China).

276

GRAIL: Scalable Reachability Index for Large Graphs

Hilmi Yildirim (Rensselaer Polytechnic Institute, United States of America), Vineet Chaoji

(Yahoo! Research Labs, United States of America), Mohammed Zaki (Rensselaer Polytechnic Institute, United States of America).

Middleware Platforms for Data Management

285

HaLoop: Efficient Iterative Data Processing on Large Clusters

Yingyi Bu (University of Washington, United States of America), Bill Howe (University of Washington, United States of America), Magdalena Balazinska (University of Washington, United States of America), Michael Ernst (University of Washington, United States of America).

297

The Impact of Virtual Views on Containment

Michael Benedikt (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom).

309

Updatable and Evolvable Transforms for Virtual Databases

James Terwilliger (Microsoft Corporation, United States of America), Lois Delcambre (Portland State University, United States of America), David Maier (Portland State University, United States of America), Jeremy Steinhauer (Portland State University, United States of America), Scott Britell (Portland State University, United States of America).

Novel/Advanced Applications

320

Navigating in Complex Mashed-Up Applications

Daniel Deutch (Tel Aviv University, Israel), Ohad Greenspan (Tel Aviv University, Israel), Tova Milo (Tel Aviv University, Israel).

330

Dremel: Interactive Analysis of Web-Scale Datasets

Sergey Melnik (Google, United States of America), Andrey Gubarev (Google, United States of America), Jing Jing Long (Google, United States of America), Geoffrey Romer (Google, United States of America), Shiva Shivakumar (Google, United States of America), Matt Tolton (Google, United States of America), Theo Vassilakis (Google, United States of America).

340

On Graph Query Optimization in Large Networks

Peixiang Zhao (University of Illinois at Urbana-Champaign, United States of America), Jiawei Han (University of Illinois at Urbana-Champaign, United States of America).

Ranking Queries

352

Proximity Rank Join

Davide Martinenghi (Politecnico di Milano, Italy), Marco Tagliasacchi (Politecnico di Milano, Italy).

364

Identifying the Most Influential Data Objects with Reverse Top-k Queries

Akrivi Vlachou (Norwegian University of Science and Technology, Norway), Christos Doukolidis (Norwegian University of Science and Technology, Norway), Kjetil Nørvag (Norwegian University of Science and Technology, Norway), Yannis Kotidis (Athens University of Economics and Business, Greece).

373

Retrieving Top-k Prestige-Based Relevant Spatial Web Objects

Xin Cao (Nanyang Technological University, Republic of Singapore), Gao Cong (Nanyang Technological University, Republic of Singapore), Christian Jensen (Aarhus University, Denmark).

Spatial and Temporal Databases

385

Parsimonious Linear Fingerprinting for Time Series

Lei Li (Carnegie Mellon University, United States of America), B. Aditya Prakash (Carnegie Mellon University, United States of America), Christos Faloutsos (Carnegie Mellon University, United States of America).

397

The HV-tree: a Memory Hierarchy Aware Version Index

Rui Zhang (University of Melbourne, Australia), Martin Stradling (University of Melbourne, Australia).

409

Transforming Range Queries To Equivalent Box Queries To Optimize Page Access

Sakti Pramanik (Michigan State University, United States of America), Alok Watve (Michigan State University, United States of America), Chad Meiners (Michigan State University, United States of America), Alex Liu (Michigan State University, United States of America).

Record Linkage

417

Record Linkage with Uniqueness Constraints and Erroneous Values

Songtao Guo (AT&T Interactive, United States of America), Xin Dong (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America), Remi Zajac (AT&T Interactive, United States of America).

429

On-the-Fly Entity-Aware Query Processing in the Presence of Linkage

Ekaterini Ioannou (L3S Research Center, Germany), Wolfgang Nejdl (L3S Research Center, Germany), Claudia Niederée (L3S Research Center, Germany), Yannis Velegrakis (University of Trento, Italy).

439

Behavior Based Record Linkage

Mohamed Yakout (Purdue University, United States of America), Ahmed Elmagarmid (Purdue University, United States of America), Hazem Elmeleegy (Purdue University, United States of America), Mourad Ouzzani (Purdue University, United States of America), Alan Qi (Purdue University, United States of America).

Experimental Analysis and Performance

449

iGraph: A Framework for Comparisons of Disk-Based Graph Indexing Techniques

Wook-Shin Han (Kyungpook National University, Republic of Korea), Jinsoo Lee (Kyungpook National University, Republic of Korea), Minh-Duc Pham (Kyungpook National University, Republic of Korea), Jeffrey Yu (The Chinese University of Hong Kong, People's Republic of China).

460

Runtime Measurements in the Cloud: Observing, Analyzing, and Reducing Variance

Jörg Schäd (Saarland University, Germany), Jens Dittrich (Saarland University, Germany), Jorge-Arnulfo Quiané-Ruiz (Saarland University, Germany).

472

The Performance of MapReduce: An In-depth Study

Dawei Jiang (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Lei Shi (National University of Singapore, Republic of Singapore), Sai Wu (National University of Singapore, Republic of Singapore).

484

Evaluation of entity resolution approaches on real-world match problems

Hanna Köpcke (University of Leipzig, Germany), Andreas Thor (University of Leipzig, Germany), Erhard Rahm (University of Leipzig, Germany).

Cloud Computing

494

MRShare: Sharing Across Multiple Queries in MapReduce

Tomasz Nykiel (University of Toronto, Canada), Michalis Potamias (Boston University, United States of America), Chaitanya Mishra (Facebook, United States of America), George Kollios (Boston University, United States of America), Nick Koudas (University of Toronto, Canada).

506

Towards Elastic Transactional Cloud Storage with Range Query Support

Hoang Tam Vo (National University of Singapore, Republic of Singapore), Chun Chen (Zhejiang University, People's Republic of China), Beng Chin Ooi (National University of Singapore, Republic of Singapore).

518

Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)

Jens Dittrich (Saarland University, Germany), Jorge-Arnulfo Quiané-Ruiz (Saarland University, Germany), Alekh Jindal (Saarland University/ International Max Planck Research

School for Computer Science, Germany), Yagiz Kargin (International Max Planck Research School for Computer Science, Germany), Vinay Setty (International Max-Planck Institute for Informatics, Germany), Jörg Schäd (Saarland University, Germany).

Query Processing I

530

Slicing Long-Running Queries

Nicolas Bruno (Microsoft Research, United States of America), Vivek Narasayya (Microsoft Research, United States of America), Ravi Ramamurthy (Microsoft Research, United States of America).

542

Sharing-Aware Horizontal Partitioning for Exploiting Correlations During Query Processing

Kostas Tzoumas (Aalborg University, Denmark), Amol Deshpande (University of Maryland College Park, United States of America), Christian Jensen (Aalborg University, Denmark).

554

Advanced Processing for Ontological Queries

Andrea Cali (University of Oxford, United Kingdom), Georg Gottlob (University of Oxford, United Kingdom), Andreas Pieris (Brunel University, United Kingdom).

Data Extraction

566

Towards The Web of Concepts: Extracting Concepts from Large Datasets

Aditya Parameswaran (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America), Anand Rajaraman (Kosmix Corporation, United States of America).

578

Exploiting Content Redundancy for Web Information Extraction

Pankaj Gulhane (Yahoo! Research Labs, Bangalore, India), Rajeev Rastogi (Yahoo! Research Labs, India), Srinivasan Sengamedu (Yahoo! Research Labs, Bangalore, India), Ashwin Tengli (Microsoft India Development Center, Bangalore, India).

588

Automatic Rule Refinement for Information Extraction

Bin Liu (University of Michigan, United States of America), Laura Chiticariu (IBM Almaden Research Center, United States of America), Vivian Chu (IBM Almaden Research Center, United States of America), H. Jagadish (University of Michigan, United States of America), Frederick Reiss (IBM Almaden Research Center, United States of America).

Privacy

598

Embellishing Text Search Queries To Protect User Privacy

HweeHwa Pang (Singapore Management University, Republic of Singapore), Xuhua Ding (Singapore Management University, Republic of Singapore), Xiaokui Xiao (Nanyang Technological University, Republic of Singapore).

608

Small Domain Randomization: Same Privacy, More Utility

Rhonda Chaytor (Simon Fraser University, Canada), Ke Wang (Simon Fraser University, Canada).

619

Nearest Neighbor Search with Strong Location Privacy

Stavros Papadopoulos (The Chinese University of Hong Kong, People's Republic of China), Spiridon Bakiras (City University of New York, United States of America), Dimitris Papadias (The Hong Kong University of Science and Technology, People's Republic of China).

Probabilistic and Uncertain Databases

630

UPI: A Primary Index for Uncertain Databases

Hideaki Kimura (Brown University, United States of America), Samuel Madden (Massachusetts Institute of Technology, United States of America), Stan Zdonik (Brown University, United States of America).

638

Ranking Continuous Probabilistic Datasets

Jian Li (University of Maryland College Park, United States of America), Amol Deshpande (University of Maryland College Park, United States of America).

650

Set Similarity Join on Probabilistic Data

Xiang Lian (The Hong Kong University of Science and Technology, People's Republic of China), Lei Chen (The Hong Kong University of Science and Technology, People's Republic of China).

Databases on Modern Hardware

660

Complex Event Detection at Wire Speed with FPGAs

Louis Woods (Eidgenössische Technische Hochschule Zürich, Switzerland), Jens Teubner (Eidgenössische Technische Hochschule Zürich, Switzerland), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland).

670

Database Compression on Graphics Processors

Wenbin Fang (The Hong Kong University of Science and Technology, People's Republic of China), Bingsheng He (Nanyang Technological University, Republic of Singapore), Qiong Luo (The Hong Kong University of Science and Technology, People's Republic of China).

681

Aether: A Scalable Approach to Logging

Ryan Johnson (Carnegie Mellon University, United States of America), Ippokratis Pandis (Carnegie Mellon University, United States of America), Radu Stoica (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Manos Athanassoulis (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland).

Data Mining

693

Scalable Discovery of Best Clusters on Large Graphs

Kathy Macropol (University of California Santa Barbara, United States of America), Ambuj Singh (University of California Santa Barbara, United States of America).

703

An Architecture for Parallel Topic Models

Alexander Smola (Yahoo! Research Labs, United States of America), Shravan Narayana-murthy (Yahoo! Research Labs, India).

711

Keyword++: A Framework to Improve Keyword Search Over Entity Databases

Dong Xin (Microsoft Research, United States of America), Yeye He (University of Wisconsin-Madison, United States of America), Venkatesh Ganti (Google, United States of America).

Moving Object Databases

723

Swarm: Mining Relaxed Temporal Moving Object Clusters

Zhenhui Li (University of Illinois at Urbana-Champaign, United States of America), Bolin Ding (University of Illinois at Urbana-Champaign, United States of America), Jiawei Han (University of Illinois at Urbana-Champaign, United States of America), Roland Kays (New York State Museum, United States of America).

735

An Adaptive Updating Protocol for Reducing Moving Object Databases Workload

Su Chen (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Zhenjie Zhang (National University of Singapore, Republic of Singapore).

747

Shortest Path Computation on Air Indexes

Georgios Kellaris (Singapore Management University, Republic of Singapore), Kyriakos Mouratidis (Singapore Management University, Republic of Singapore).

Probabilistic Data

758

Efficient and Effective Similarity Search over Probabilistic Data based on Earth Mover's Distance

Jia Xu (Northeastern University, People's Republic of China), Zhenjie Zhang (National University of Singapore, Republic of Singapore), Anthony Tung (National University of Singapore, Republic of Singapore), Ge Yu (Northeastern University, People's Republic of China).

770

Probabilistic XML via Markov Chains

Michael Benedikt (University of Oxford, United Kingdom), Evgeny Kharlamov (Free University of Bozen-Bolzano, Italy), Dan Olteanu (University of Oxford, United Kingdom), Pierre Senellart (Institut Télécom ParisTech, France).

Fuzzy, Probabilistic and Approximate Databases

782

MCDB-R: Risk Analysis in the Database

Peter Haas (IBM Research, United States of America), Christopher Jermaine (Rice University, United States of America), Subi Arumugam (University of Florida Gainesville, United States of America), Fei Xu (Microsoft Corporation, United States of America), Luis Perez (Rice University, United States of America), Ravi Jampani (University of Florida Gainesville, United States of America).

794

Scalable Probabilistic Databases with Factor Graphs and MCMC

Michael Wick (University of Massachusetts Amherst, United States of America), Andrew McCallum (University of Massachusetts Amherst, United States of America), Gerome Miklau (University of Massachusetts Amherst, United States of America).

Discovery and Exploration

805

On Multi-Column Foreign Key Discovery

Meihui Zhang (National University of Singapore, Republic of Singapore), Marios Hadjieleftheriou (AT&T Labs - Research, United States of America), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Cecilia Procopiuc (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

815

Explore or Exploit? Effective Strategies for Disambiguating Large Databases

Reynold Cheng (The University of Hong Kong, People's Republic of China), Eric Lo (Hong Kong Polytechnic University, People's Republic of China), Xuan Yang (The University of Hong Kong, People's Republic of China), Ming-Hay Luk (Hong Kong Polytechnic University, People's Republic of China), Xiang Li (The University of Hong Kong, People's Republic of China), Xike Xie (The University of Hong Kong, People's Republic of China).

Information Filtering and Dissemination

826

Building Ranked Mashups of Unstructured Sources with Uncertain Information

Mohamed Soliman (University of Waterloo, Canada), Ihab Ilyas (University of Waterloo, Canada), Mina Saleeb (University of Waterloo, Canada).

838

Computing Closed Skycubes

Chedy Raïssi (Institut National de Recherche en Informatique et Automatique, France), Jian Pei (Simon Fraser University, Canada), Thomas Kister (Laboratoire d'Informatique de Robotique et de Microélectronique de Montpellier).

Query Processing II

848

Generating Databases for Query Workloads

Eric Lo (Hong Kong Polytechnic University, People's Republic of China), Nick Cheng (Hong Kong Polytechnic University, People's Republic of China), Wing-Kai Hon (National Tsing Hua University, Taiwan).

860

Processing Top-k Join Queries

Minji Wu (Rutgers University, United States of America), Laure Berti-Équille (University of Rennes 1, France), Amélie Marian (Rutgers University, United States of America), Cecilia Procopiuc (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

871

Two-way Replacement Selection

Xavier Martinez-Palau (Universitat Politècnica de Catalunya, Spain), David Dominguez-Sal (Universitat Politècnica de Catalunya, Spain), Josep Lluís Larriba-Pey (Universitat Politècnica de Catalunya, Spain).

XML Data

882

XPath Whole Query Optimization

Sebastian Maneth (National Information and Communication Technologies Australia / University of New South Wales, Australia), Kim Nguyen (National Information and Communication Technologies Australia, Australia).

894

Fast Optimal Twig Joins

Nils Grimsmo (Norwegian University of Science and Technology, Norway), Truls Bjørklund (Norwegian University of Science and Technology, Norway), Magnus Hetland (Norwegian University of Science and Technology, Norway).

906

Destabilizers and Independence of XML Updates

Michael Benedikt (University of Oxford, United Kingdom), James Cheney (University of Edinburgh, United Kingdom).

Workflows, Transactions and Business Processes

918

Searching Workflows with Hierarchical Views

Ziyang Liu (Arizona State University, United States of America), Qihong Shao (Arizona State University, United States of America), Yi Chen (Arizona State University, United States of America).

928

Data-Oriented Transaction Execution

Ippokratis Pandis (Carnegie Mellon University, United States of America), Ryan Johnson (Carnegie Mellon University, United States of America), Nikos Hardavellas (Northwestern University, United States of America), Anastasia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland).

940

Optimal Top-K Query Evaluation for Weighted Business Processes

Daniel Deutch (Tel Aviv University, Israel), Tova Milo (Tel Aviv University, Israel), Neoklis Polyzotis (University of California Santa Cruz, United States of America), Tom Yam (Tel Aviv University, Israel).

Scientific databases

952

Behavioral Simulations in MapReduce

Guozhang Wang (Cornell University, United States of America), Marcos Vaz Salles (Cornell University, United States of America), Benjamin Sowell (Cornell University, United States of America), Xun Wang (Cornell University, United States of America), Tuan Cao (Cornell University, United States of America), Alan Demers (Cornell University, United States of America), Johannes Gehrke (Cornell University, United States of America), Walker White (Cornell University, United States of America).

964

A-tree: A Structure for Storage and Modeling of Uncertain Multidimensional Arrays*

Tingjian Ge (University of Kentucky, United States of America), Stan Zdonik (Brown University, United States of America).

975

On Dense Pattern Mining in Graph Streams

Charu Aggarwal (IBM Thomas J. Watson Research Center, United States of America), Yao Li (University of Illinois at Chicago, United States of America), Philip Yu (University of Illinois at Chicago, United States of America), Ruoming Jin (Kent State University, United States of America).

Mobility and Spatial Queries

985

Efficient Proximity Detection among Mobile Users via Self-Tuning Policies

Man Lung Yiu (Hong Kong Polytechnic University, People's Republic of China), Leong Hou U (The University of Hong Kong, People's Republic of China), Simonas Šaltenis (Aalborg University, Denmark), Kostas Tzoumas (Aalborg University, Denmark).

997

k-Nearest Neighbors in Uncertain Graphs

Michalis Potamias (Boston University, United States of America), Francesco Bonchi (Yahoo! Research Labs, United States of America), Aristides Gionis (Yahoo! Research Labs, United States of America), George Kollios (Boston University, United States of America).

1009

Mining Significant Semantic Locations From GPS Data

Xin Cao (Nanyang Technological University, Republic of Singapore), Gao Cong (Nanyang

Technological University, Republic of Singapore), Christian Jensen (Aarhus University, Denmark).

Data Anonymization Techniques

1021

Boosting the Accuracy of Differentially Private Histograms Through Consistency

Michael Hay (University of Massachusetts Amherst, United States of America), Vibhor Rastogi (University of Washington, United States of America), Gerome Miklau (University of Massachusetts Amherst, United States of America), Dan Suciu (University of Washington, United States of America).

1033

rho-uncertainty: Inference-Proof Transaction Anonymization

Jianneng Cao (National University of Singapore, Republic of Singapore), Panagiotis Karras (National University of Singapore, Republic of Singapore), Chedy Raïssi (Institut National de Recherche en Informatique et Automatique, France), Kian-Lee Tan (National University of Singapore, Republic of Singapore).

1045

Minimizing Minimality and Maximizing Utility: Analyzing Method-based attacks on Anonymized Data

Graham Cormode (AT&T Labs - Research, United States of America), Ninghui Li (Purdue University, United States of America), Tiancheng Li (Purdue University, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

Querying and Integrating Probabilistic Databases

1057

Querying Probabilistic Information Extraction

Daisy Zhe Wang (University of California

Berkeley, United States of America), Michael Franklin (University of California Berkeley, United States of America), Minos Garofalakis (Technical University of Crete, Greece), Joseph Hellerstein (University of California Berkeley, United States of America).

1068

Read-Once Functions and Query Evaluation in Probabilistic Databases

Prithviraj Sen (Yahoo! Research Labs, United States of America), Amol Deshpande (University of Maryland College Park, United States of America), Lise Getoor (University of Maryland College Park, United States of America).

1080

Foundations of Uncertain-Data Integration

Parag Agrawal (Stanford University, United States of America), Anish Das Sarma (Yahoo! Research Labs, United States of America), Jeffrey Ullman (Stanford University, United States of America), Jennifer Widom (Stanford University, United States of America).

Database Design

1091

Identifying, Attributing and Describing Spatial Bursts

Michael Mathioudakis (University of Toronto, Canada), Nilesh Bansal (University of Toronto, Canada), Nick Koudas (University of Toronto, Canada).

1103

CORADD: Correlation Aware Database Designer for Materialized Views and Indexes

Hideaki Kimura (Brown University, United States of America), George Huo (Google, United States of America), Alexander Rasin (Brown University, United States of America), Samuel Madden (Massachusetts Institute

of Technology, United States of America), Stanley Zdonik (Brown University, United States of America).

1114

Regret-Minimizing Representative Databases

Danupon Nanongkai (Georgia Institute of Technology, United States of America), Atish Das Sarma (Georgia Institute of Technology, United States of America), Ashwin Lall (Georgia Institute of Technology, United States of America), Richard Lipton (Georgia Institute of Technology, United States of America), Jun Xu (Georgia Institute of Technology, United States of America).

Query Optimization

1125

An Access Cost-Aware Approach for Object Retrieval over Multiple Sources

Benjamin Arai (University of California Riverside, United States of America), Gautam Das (University of Texas at Arlington, United States of America), Dimitrios Gunopulos (University of Athens, Greece), Vagelis Hristidis (Florida International University, United States of America), Nick Koudas (University of Toronto, Canada).

1137

On the Stability of Plan Costs and the Costs of Plan Stability

M. Abhirama (Indian Institute of Science, India), Sourjya Bhaumik (Indian Institute of Science, India), Atreyee Dey (Indian Institute of Science, India), Harsh Shrimal (Indian Institute of Science, India), Jayant Haritsa (Indian Institute of Science, India).

1149

Xplus: A SQL-Tuning-Aware Query Optimizer

Herodotos Herodotou (Duke University, United States of America), Shivnath Babu (Duke University, United States of America).

Graph and Pattern Matching

1161

Graph Homomorphism Revisited for Graph Matching

Wenfei Fan (University of Edinburgh, United Kingdom), Jianzhong Li (Harbin Institute of Technology, People's Republic of China), Shuai Ma (University of Edinburgh, United Kingdom), Hongzhi Wang (Harbin Institute of Technology, People's Republic of China), Yinghui Wu (University of Edinburgh, United Kingdom).

1173

SigMatch: Fast and Scalable Multi-Pattern Matching

Ramakrishnan Kandhan (University of Wisconsin-Madison, United States of America), Nikhil Teletia (University of Wisconsin-Madison, United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

1185

SAPPER: Subgraph Indexing and Approximate Matching in Large Graphs

Shijie Zhang (Case Western Reserve University, United States of America), Jiong Yang (Case Western Reserve University, United States of America), Wei Jin (Case Western Reserve University, United States of America).

Indexing Techniques

1195

Tree Indexing on Solid State Drives

Yinan Li (University of Wisconsin-Madison, United States of America), Bingsheng He (The Hong Kong University of Science and Technology, People's Republic of China), Robin Jun Yang (The Hong Kong University of Science and Technology, People's Republic of China), Qiong Luo (The Hong Kong University of Science and Technology, People's Republic of China), Ke

Yi (The Hong Kong University of Science and Technology, People's Republic of China).

1207

Efficient B-tree Based Indexing for Cloud Data Processing

Sai Wu (National University of Singapore, Republic of Singapore), Dawei Jiang (National University of Singapore, Republic of Singapore), Beng Chin Ooi (National University of Singapore, Republic of Singapore), Kun-Lung Wu (IBM Thomas J. Watson Research Center, United States of America).

1219

Trie-Join: Efficient Trie-based String Similarity Joins with Edit-Distance Constraints

Jiannan Wang (Tsinghua University, People's Republic of China), Guoliang Li (Tsinghua University, People's Republic of China), Jianhua Feng (Tsinghua University, People's Republic of China).

Query Processing III

1231

VoR-Tree: R-trees with Voronoi Diagrams for Efficient Processing of Spatial Nearest Neighbor Queries

Mehdi Sharifzadeh (Google, United States of America), Cyrus Shahabi (University of Southern California, United States of America).

1243

Efficient RkNN Retrieval with Arbitrary Non-Metric Similarity Measures

Deepak P (IBM India Research Lab, India), Prasad Deshpande (IBM India Research Lab, India).

1255

Efficient Skyline Evaluation over Partially Ordered Domains

Shiming Zhang (The University of Hong Kong, People's Republic of China), Nikos Mamoulis (The University of Hong Kong, People's Republic of China), Ben Kao (The University of Hong Kong, People's Republic of China), David Cheung (The University of Hong Kong, People's Republic of China).

Streaming and Sensor Data

1267

Achieving High Output Quality under Limited Resources through Structure-based Spilling in XML Streams

Mingzhu Wei (Worcester Polytechnic Institute, United States of America), Elke Rundensteiner (Worcester Polytechnic Institute, United States of America), Murali Mani (University of Michigan at Flint, United States of America).

1279

Dynamic Join Optimization in Multi-Hop Wireless Sensor Networks

Svilen Mihaylov (University of Pennsylvania, United States of America), Marie Jacob (University of Pennsylvania, United States of America), Zachary Ives (University of Pennsylvania, United States of America), Sudipto Guha (University of Pennsylvania, United States of America).

1291

Database-support for Continuous Prediction Queries over Streaming Data

Mert Akdere (Brown University, United States of America), Uğur Çetintemel (Brown University, United States of America), Eli Upfal (Brown University, United States of America).

1302

Conditioning and Aggregating Uncertain Data Streams: Going Beyond Expectations

Thanh Tran (University of Massachusetts Amherst, United States of America), Andrew McGregor (University of Massachusetts

Amherst, United States of America), Yanlei Diao (University of Massachusetts Amherst, United States of America), Liping Peng (University of Massachusetts Amherst, United States of America), Anna Liu (University of Massachusetts Amherst, United States of America).

Information Integration and Retrieval

1314

TRAMP: Understanding the Behavior of Schema Mappings through Provenance

Boris Glavic (University of Zurich, Switzerland), Gustavo Alonso (Eidgenössische Technische Hochschule Zürich, Switzerland), Renée Miller (University of Toronto, Canada), Laura Haas (IBM Research, United States of America).

1326

Entity Resolution with Evolving Rules

Steven Whang (Stanford University, United States of America), Hector Garcia-Molina (Stanford University, United States of America).

1338

Annotating and Searching Web Tables Using Entities, Types and Relationships

Girija Limaye (Indian Institute of Technology - Bombay, India), Sunita Sarawagi (Indian Institute of Technology - Bombay, India), Soumen Chakrabarti (Indian Institute of Technology - Bombay, India).

Data Mining, Copy Detection and Data Publication

1348

Interesting-Phrase Mining for Ad-Hoc Text Analytics

Srikanta Bedathur (Max-Planck Institute for

Informatics, Germany), Klaus Berberich (Max-Planck Institute for Informatics, Germany), Jens Dittrich (Saarland University, Germany), Nikos Mamoulis (Max-Planck Institute for Informatics, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

1358

Global Detection of Complex Copying Relationships Between Sources

Xin Dong (AT&T Labs - Research, United States of America), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

1370

Fragments and Loose Associations: Respecting Privacy in Data Publishing

Sabrina De Capitani di Vimercati (Università degli Studi di Milano, Italy), Sara Foresti (Università degli Studi di Milano, Italy), Sushil Jajodia (George Mason University Fairfax, United States of America), Stefano Paraboschi (University of Bergamo, Italy), Pierangela Samarati (Università degli Studi di Milano, Italy).

Industrial Applications and Experience Sessions

Real-Time and Stream Processing

1382

Net-Fli: On-the-fly Compression, Archiving and Indexing of Streaming Network Traffic

Francesco Fusco (IBM Research - Zurich, Switzerland), Marc Stoecklin (IBM Research - Zurich, Switzerland), Michael Vlachos (IBM Research - Zurich, Switzerland).

1394

From a Stream of Relational Queries to Distributed Stream Processing

Qiong Zou (IBM China Research Lab, People's Republic of China), Huayong Wang (IBM China Research Lab, People's Republic of China), Robert Soulé (New York University, United States of America), Martin Hirzel (IBM Thomas J. Watson Research Center, United States of America), Henrique Andrade (IBM Thomas J. Watson Research Center, United States of America), Buğra Gedik (IBM Thomas J. Watson Research Center, United States of America), Kun-Lung Wu (IBM Thomas J. Watson Research Center, United States of America).

1406

*UASMA*s (Universal Automated SNP Mapping Algorithms): a set of algorithms to instantaneously map SNPs in real time to aid functional SNP discovery

James Mah (Institute for Infocomm Research, Republic of Singapore), Danny Poo (National University of Singapore, Republic of Singapore), Shaojiang Cai (National University of Singapore, Republic of Singapore).

Advances in Database Systems

1414

FlashStore: High Throughput Persistent Key-Value Store

Biplob Debnath (University of Minnesota, United States of America), Sudipta Sengupta (Microsoft Research - Redmond, United States of America), Jin Li (Microsoft Research - Redmond, United States of America).

1426

MEET DB2: Automated Database Migration Evaluation

Reynold Xin (University of California Berkeley, United States of America), Patrick Dantresangle (IBM Research, United Kingdom), Sam Lightstone (IBM Research, Canada), William McLaren (IBM, United States of America), Steve Schormann (IBM Research, Canada), Maria Schwenger (IBM, United States of America).

1435

SSD Bufferpool Extensions for Database Systems

Mustafa Canim (University of Texas at Dallas, United States of America), George Mihaila (IBM Thomas J. Watson Research Center, United States of America), Bishwaranjan Bhattacharjee (IBM Thomas J. Watson Research Center, United States of America), Kenneth Ross (Columbia University, United States of America), Christian Lang (IBM Thomas J. Watson Research Center, United States of America).

Data Warehousing and Mining

1447

DataGarage: Warehousing Massive Performance Data on Commodity Servers

Charles Loboz (Microsoft Corporation, United States of America), Slawek Smyl (Microsoft Corporation, United States of America), Suman Nath (Microsoft Research, United States of America).

1459

Cheetah: A High Performance, Custom Data Warehouse on Top of MapReduce

Songting Chen (Turn Incorporated, United States of America).

1469

Distance-Based Outlier Detection: Consolidation and Renewed Bearing

Gustavo Orair (Universidade Federal de Minas Gerais, Brazil), Carlos Teixeira (Universidade Federal de Minas Gerais, Brazil), Ye Wang (The Ohio State University, United States of America), Wagner Meira (Universidade Federal de Minas Gerais, Brazil), Srinivasan Parthasarathy (The Ohio State University, United States of America).

Mobile and Social Applications

1481

Adaptive Logging for Mobile Device

Young-Seok Kim (Samsung Advanced Institute of Technology, Republic of Korea), Heegy Jin (Samsung Advanced Institute of Technology, Republic of Korea), Kyoung-Gu Woo (Samsung Advanced Institute of Technology, Republic of Korea).

1493

RoadTrack: Scaling Location Updates for Mobile Clients on Road Networks with Query Awareness

Peter Pesti (Georgia Institute of Technology, United States of America), Ling Liu (Georgia Institute of Technology, United States of America), Bhuvan Bamba (Georgia Institute of Technology, United States of America), Arun Iyengar (IBM Thomas J. Watson Research Center, United States of America), Matt Weber (Georgia Institute of Technology, United States of America).

1505

Confucius and Its Intelligent Disciples: Integrating Social with Search

Xiance Si (Google Research, People's Republic of China), Edward Chang (Google Research, People's Republic of China), Zoltán Gyöngyi (Google Research, United States of America), Maosong Sun (Tsinghua University, People's Republic of China).

Demonstrations

Core Database Technologies (12 Demos)

1517

The Picasso Database Query Optimizer Visualizer

Jayant Haritsa (Indian Institute of Science, India).

1521

CODS: Evolving Data Efficiently and Scalably in Column Oriented Databases

Ziyang Liu (Arizona State University, United States of America), Sivaramakrishnan Natarajan (Arizona State University, United States of America), Bin He (IBM Almaden Research Center, United States of America), Hui-I Hsiao (IBM Almaden Research Center, United States of America), Yi Chen (Arizona State University, United States of America).

1525

Efficient Event Processing through Reconfigurable Hardware for Algorithmic Trading

Mohammad Sadoghi (University of Toronto, Canada), Hans-Arno Jacobsen (University of Toronto, Canada), Martin Labrecque (University of Toronto, Canada), Warren Shum (University of Toronto, Canada), Harsh Singh (University of Toronto, Canada).

1529

CareDB: A Context and Preference-Aware Location-Based Database System

Justin Levandoski (University of Minnesota, United States of America), Mohamed Mokbel (University of Minnesota, United States of America), Mohamed Khalefa (University of Minnesota, United States of America).

1533

Cloudy: A Modular Cloud Storage System

Donald Kossmann (Eidgenössische Technische Hochschule Zürich, Switzerland), Tim Kraska (University of California Berkeley, United States of America), Simon Loesing (Eidgenössische Technische Hochschule Zürich, Switzerland), Stephan Merkli (Eidgenössische Technische Hochschule Zürich, Switzerland), Raman Mittal (Eidgenössische Technische Hochschule Zürich, Switzerland), Flavio Pfaffhauser (Eidgenössische Technische Hochschule Zürich, Switzerland).

1537

Geospatial Stream Query Processing using Microsoft SQL Server StreamInsight

Seyed Jalal Kazemitabar (University of Southern California, United States of America), Ugur Demiryurek (University of Southern California, United States of America), Mohamed Ali (Microsoft Corporation, United States of America), Afsin Akdogan (University of Southern California, United States of America), Cyrus Shahabi (University of Southern California, United States of America).

1541

Using XMorph to Transform XML Data

Curtis Dyreson (Utah State University, United States of America), Sourav Bhowmick (Nanyang Technological University, Republic of Singapore), Kirankanth Mallampalli (Utah State University, United States of America).

1545

Active Complex Event Processing: Applications in Real-Time Health Care

Di Wang (Worcester Polytechnic Institute, United States of America), Elke Rundensteiner (Worcester Polytechnic Institute, United States of America), Richard Ellison (University of Massachusetts Medical School, United States of America), Han Wang (Worcester Polytechnic Institute, United States of America).

1549

Thirteen New Players in the Team: A Ferry-based LINQ to SQL Provider

Tom Schreiber (Universität Tübingen, Germany), Simone Bonetti (Universität Tübingen, Germany), Torsten Grust (Universität Tübingen, Germany), Manuel Mayr (Universität Tübingen, Germany), Jan Rittinger (Universität Tübingen, Germany).

1553

AXART - Enabling Collaborative Work with AXML Artifacts

Bogdan Marinoiu (SAP Labs, France), Serge Abiteboul (Institut National de Recherche en Informatique et en Automatique, France), Pierre Bourhis (Institut National de Recherche en Informatique et en Automatique, France), Alban Galland (Institut National de Recherche en Informatique et en Automatique, France).

1557

iFlow: An Approach for Fast and Reliable Internet-Scale Stream Processing Utilizing Detouring and Replication

Christopher McConnell (State University of New York University at Albany, United States of America), Fan Ping (State University of New York University at Albany, United States of America), Jeong-Hyon Hwang (State University of New York University at Albany, United States of America).

1561

Peer coordination through distributed triggers

Verena Kantere (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Maher Manoubi (University of Ottawa, Canada), Iluju Kiringa (University of Ottawa, Canada), Timos Sellis (National Technical University of Athens, Greece), John Mylopoulos (University of Toronto, Canada).

Web Data Management and Information Retrieval Systems (12 Demos)

1565

Seaform: Search-As-You-Type in Forms

Hao Wu (Tsinghua University, People's Republic of China), Guoliang Li (Tsinghua University, People's Republic of China), Chen Li (University of California Irvine, United States of America), Lizhu Zhou (Tsinghua University, People's Republic of China).

1569

TimeTrails: A System for Exploring Spatio-Temporal Information in Documents

Jannik Strötgen (University of Heidelberg, Germany), Michael Gertz (University of Heidelberg, Germany).

1573

QUICK: Expressive and Flexible Search over Knowledge Bases and Text Collections

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

1577

Transforming XML Documents as Schemas Evolve

Jarek Gryz (York University, Canada), Marcin Kwietniewski (York University, Canada), Stephanie Hazlewood (IBM, Canada), Paul Van Run (IBM, Canada).

1581

XSACT: A Comparison Tool for Structured Search Results

Ziyang Liu (Arizona State University, United States of America), Sivaramakrishnan Nataraajan (Arizona State University, United States of America), Peng Sun (Arizona State University, United States of America), Stephen Booher (Arizona State University, United States of America), Tim Meehan (Arizona State University, United States of America), Robert Winkler (Arizona State University, United States of America), Yi Chen (Arizona State University, United States of America).

1585

ObjectRunner: Lightweight, Targeted Extraction and Querying of Structured Web Data

Talel Abdessalem (Télécom ParisTech, France),

Bogdan Cautis (Télécom ParisTech, France), Nora Derouiche (Télécom ParisTech, France).

1589

ROXXI: Reviving witness dOcuments to eXplore eXtracted Information

Shady Elbassuoni (Max-Planck Institute for Informatics, Germany), Katja Hose (Max-Planck Institute for Informatics, Germany), Steffen Metzger (Max-Planck Institute for Informatics, Germany), Ralf Schenkel (Max-Planck Institute for Informatics, Germany).

1593

EXTRACT: Using Deep Structural Information in XML Keyword Search

Arash Termehchy (University of Illinois at Urbana-Champaign, United States of America), Marianne Winslett (University of Illinois at Urbana-Champaign, United States of America).

1597

SQL QueRIE Recommendations

Javad Akbarnejad (San Jose State University, United States of America), Gloria Chatzopoulou (University of California Riverside, United States of America), Magdalini Eirinaki (San Jose State University, United States of America), Suju Koshy (San Jose State University, United States of America), Sarika Mittal (San Jose State University, United States of America), Duc On (San Jose State University, United States of America), Neoklis Polyzotis (University of California Santa Cruz, United States of America), Jothi Swarubini Vindhiya Varman (San Jose State University, United States of America).

1601

P2PDocTagger: Content management through automated P2P collaborative tagging

Hock Hee Ang (Nanyang Technological University, Republic of Singapore), Vivekanand

Gopalkrishnan (Nanyang Technological University, Republic of Singapore), Wee Keong Ng (Nanyang Technological University, Republic of Singapore), Steven C.H. Hoi (Nanyang Technological University, Republic of Singapore).

1605

InZeit: Efficiently Identifying Insightful Time Points

Vinay Setty (Max-Planck Institute for Informatics, Germany), Srikanta Bedathur (Max-Planck Institute for Informatics, Germany), Klaus Berberich (Max-Planck Institute for Informatics, Germany), Gerhard Weikum (Max-Planck Institute for Informatics, Germany).

1609

iAVATAR: An Interactive Tool for Finding and Visualizing Visual-Representative Tags in Image Search

Aixin Sun (Nanyang Technological University, Republic of Singapore), Sourav Bhowmick (Nanyang Technological University, Republic of Singapore), Yao Liu (Nanyang Technological University, Republic of Singapore).

Data Extraction, Integration and Mining (8 Demos)

1613

Deep Web Integration with VisQI

Thomas Kabisch (Humboldt Universität zu Berlin, Germany), Eduard Dragut (University of Illinois at Chicago, United States of America), Clement Yu (University of Illinois at Chicago, United States of America), Ulf Leser (Humboldt Universität zu Berlin, Germany).

1617

SOLOMON: Seeking the Truth Via Copying Detection

Xin Dong (AT&T Labs - Research, United States of America), Laure Berti-Equille (University of Rennes 1, France), Yifan Hu (AT&T

Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

1621

Just-in-time Data Integration in Action

Martin Hentschel (Eidgenössische Technische Hochschule Zürich, Switzerland), Laura Haas (IBM Almaden Research, United States of America), Renée Miller (University of Toronto, Canada).

1625

Massively Parallel Data Analysis with PACTs on Nephele

Alexander Alexandrov (Technische Universität Berlin, Germany), Dominic Battré (Technische Universität Berlin, Germany), Stephan Ewen (Technische Universität Berlin, Germany), Max Heibel (Technische Universität Berlin, Germany), Fabian Hueske (Technische Universität Berlin, Germany), Odej Kao (Technische Universität Berlin, Germany), Volker Markl (Technische Universität Berlin, Germany), Erik Nijkamp (Technische Universität Berlin, Germany), Daniel Warneke (Technische Universität Berlin, Germany).

1629

Using Sentinel Technology in the TARGIT BI Suite

Morten Middelfart (TARGIT A/S, Denmark), Torben Bach Pedersen (Aalborg University, Denmark).

1633

CoDA: Interactive Cluster Based Concept Discovery

Stephan Günnemann (RWTH Aachen University, Germany), Ines Färber (RWTH Aachen University, Germany), Hardy Kremer (RWTH Aachen University, Germany), Thomas Seidl (RWTH Aachen University, Germany).

1637

Keymantic: Semantic Keyword-based Searching in Data Integration Systems

Sonia Bergamaschi (University of Modena and Reggio Emilia, Italy), Elton Domnori (University of Modena and Reggio Emilia, Italy), Francesco Guerra (University of Modena and Reggio Emilia, Italy), Mirko Orsini (University of Modena and Reggio Emilia, Italy), Raquel Trillo Lado (University of Zaragoza, Spain), Yannis Velegrakis (University of Trento, Italy).

1641

Data Auditor: Exploring Data Quality and Semantics using Pattern Tableaux

Lukasz Golab (AT&T Labs - Research, United States of America), Howard Karloff (AT&T Labs - Research, United States of America), Flip Korn (AT&T Labs - Research, United States of America), Divesh Srivastava (AT&T Labs - Research, United States of America).

Tutorials

1645

Distributed Caching Platforms

Anil Nori (Microsoft Corporation, United States of America).

1647

Big Data and Cloud Computing: New Wine or just New Bottles?

Divyakant Agrawal (University of California Santa Barbara, United States of America), Sudipto Das (University of California Santa Barbara, United States of America), Amr El Abbadi (University of California Santa Barbara, United States of America).

1649

Techniques for Similarity Searching in Multimedia Databases

Hanan Samet (University of Maryland College Park, United States of America).

1651

Event Processing – past, present and future

Opher Etzion (IBM Haifa Research Lab, Israel).

1653

Similarity Search and Mining in Uncertain Databases

Matthias Renz (Ludwig Maximilian University of Munich, Germany), Reynold Cheng (The University of Hong Kong, People's Republic of China), Hans-Peter Kriegel (Ludwig Maximilian University of Munich, Germany).

1655

Data Management and Mining in Internet Ad Systems

S. Muthukrishnan (Rutgers U., and Google Inc., United States of America).

Panels

Keynote Sessions

1657

Cloud Databases: What's New?

Daniel Abadi (Yale University, United States of America), Michael Carey (University of California Irvine, United States of America), Surajit Chaudhuri (Microsoft Research, United States of America), Hector Garcia-Molina (Stanford University, United States of America), Jignesh Patel (University of Wisconsin-Madison, United States of America).

1658

Time for Our Field to Grow Up

Anastassia Ailamaki (Ecole Polytechnique Fédérale de Lausanne, Switzerland), Laura Haas (IBM Almaden Research Center, United

States of America), HV Jagadish (University of Michigan, United States of America), David

Maier (Portland State University, United States of America), Tamer Özsu (University of Waterloo, Canada).

On Multi-Column Foreign Key Discovery

Meihui Zhang
National University of
Singapore
mhzhang@comp.nus.edu.sg

Marios Hadjieleftheriou
AT&T Labs - Research
marioh@research.att.com

Beng Chin Ooi
National University of
Singapore
ooibc@comp.nus.edu.sg

Cecilia M. Procopiuc
AT&T Labs - Research
magda@research.att.com

Divesh Srivastava
AT&T Labs - Research
divesh@research.att.com

ABSTRACT

A foreign/primary key relationship between relational tables is one of the most important constraints in a database. From a data analysis perspective, discovering foreign keys is a crucial step in understanding and working with the data. Nevertheless, more often than not, foreign key constraints are not specified in the data, for various reasons; e.g., some associations are not known to designers but are inherent in the data, while others become invalid due to data inconsistencies. This work proposes a robust algorithm for discovering single-column and multi-column foreign keys. Previous work concentrated mostly on discovering single-column foreign keys using a variety of rules, like inclusion dependencies, column names, and minimum/maximum values. We first propose a general rule, termed *Randomness*, that subsumes a variety of other rules. We then develop efficient approximation algorithms for evaluating randomness, using only two passes over the data. Finally, we validate our approach via extensive experiments using real and synthetic datasets.

1. INTRODUCTION

A foreign/primary key relationship between relational tables is one of the most important constraints in a database. From a data analysis perspective, discovering foreign keys is a crucial step in understanding and working with the data. For that reason, database systems allow the explicit specification of foreign key constraints in the database schema. Nevertheless, in practice, database designers frequently fail to specify such constraints for various reasons, including: they are not aware of implicit relationships inherent in the data; such relationships might hold across multiple databases; it is not feasible to specify the constraints due to data inconsistencies (e.g., those arising from data integration or from database evolution over time); or because of performance considerations. When this happens in enterprise databases, which often contain hundreds of tables, thousands of columns and insufficient (or missing) documentation, even expert users have a difficult time identifying foreign key constraints.

In this paper, we propose a novel approach for discovering foreign/primary key (fk/pk) relationships between single or multiple

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were presented at The 36th International Conference on Very Large Data Bases, September 13-17, 2010, Singapore.

Proceedings of the VLDB Endowment, Vol. 3, No. 1

Copyright 2010 VLDB Endowment 2150-8097/10/09... \$ 10.00.




Figure 1: A small subset of the TPC-E schema with one multi-column and several single-column foreign keys.

columns in relational databases. Surprisingly, little previous work deals with the case of discovering multi-column foreign keys [14]. Even for single-column keys, existing work is limited and focuses mainly on identifying inclusion dependencies, since the only formal requirement for specifying a foreign key constraint is that the foreign key be a subset of the primary key [1, 14]. However, checking only for inclusion can lead to a large number of false positives.

For example, Figure 1 shows a portion of the benchmark TPC-E schema, which represents a stock transaction system. It has information about customer accounts, companies, brokers, stock trades, etc. Column Trade.TID contains all integers in the interval [1, 10000], while column Broker.BID, which is unrelated to TID, contains all integers in [1, 100]. A simple inclusion test would incorrectly report (Broker.BID, Trade.TID) as a foreign/primary key pair. This scenario arises frequently in practice because of auto-increment fields. Of course, one could adapt the test so that it discards pairs in which one column is a consecutive subset (e.g., a prefix or a suffix) of the other. However, that is not sufficient. Notice that the values in column Customer Account.BID, which is a foreign key of column Broker.BID, are a *random subset of a prefix* of Trade.TID. Hence, the inclusion test adapted as above would still incorrectly report (Customer Account.BID, Trade.TID) as a foreign/primary key pair. To complicate matters further, this problem is not limited to numerical attributes. It arises with date-time fields that may contain consecutive values, or even alphanumeric fields composed of letters followed by a number (e.g., A-1, A-2). The same is true for multi-column keys. For example, Holding.(CID, SMB) is a two-column foreign key of Holding Summary.(CID, SMB). However, Broker.(BID, STID) is not a valid foreign key of Trade History.(TID, STID), even though column-wise inclusion is satisfied.

Reducing the number of false positives is a critical requirement

in order to make the identification of useful relationships feasible. As we show in the experimental section the number of false positives (i.e., pairs of columns that satisfy inclusion but are not valid fk/pk constraints) can be in the order of hundreds. Even for domain experts the task of sifting through and manually validating candidates is overwhelming. Previous work has proposed heuristic rules to reduce the number of false positives by identifying important properties that a good foreign key should satisfy. A comprehensive list of such properties, compiled based on extensive experimentation, appears in Rostin et al. [17]. Some of the most important rules are: 1. A foreign key should have significant cardinality; 2. A foreign key should have good coverage of the primary key; 3. A foreign key should not be at the same time a primary key for too many other foreign keys; 4. The set of values of a foreign key should not be a subset of too many primary keys; 5. The average length of the values in foreign/primary key columns should be similar (mostly for strings); 6. The primary key should have only a small percentage of values outside the range of the foreign key; 7. The column names of foreign/primary keys should be similar. Indeed, a lot of previous work, especially in the realm of schema matching, has used similar rules to find associations between columns [10].

It is important to note that counter-examples exist for any rule one tries to devise. One can easily come up with such examples for rules 1 and 2. A counter-example for rules 3 and 4 is when social security numbers or telephone numbers are used in a database as primary keys. Such keys are expected to appear in a large number of tables. Kang and Naughton [10] give several counter-examples for rule 7, i.e., columns with no meaningful association but very similar names.

In this work, we propose a novel method for measuring the likelihood that a pair of columns that satisfy inclusion is a useful fk/pk constraint. Our approach subsumes a variety of previous rules, and, as we show in Section 5, is both highly scalable and accurate.

Consider the set of values in a primary key P , ordered by the natural order in the underlying domain (i.e., numerical order for numeric attributes, lexicographic order for strings). We conjecture that in most cases the values in a foreign key column F form a (nearly) uniform random sample of the values in P . For example, consider columns Broker.BID and Customer Account.BID from Figure 1. The broker ids that appear in Customer Account are expected to be “sprinkled” uniformly throughout the ordered set of all broker ids. This is because we have no reason to expect a correlation between the semantics of the foreign key constraint (“this broker works with these customers”) and the mechanism through which the broker ids are generated. For example, ids may be consecutive numbers between 1 and 100 generated via auto-increment. By contrast, the subset of broker ids that appear in Customer Account may reflect, say, those brokers with great reputation. It is highly unlikely that a database instance is designed such that a foreign key is a biased sample of the respective primary key (e.g., a prefix or a suffix in the ranked order). Even if this is the case, for dynamic databases the distribution of fk/pk values is expected to change over time, eventually eliminating such bias. The closer a column F is to being a uniform random sample of a primary key P , the higher we consider the likelihood that (F, P) is a useful fk/pk constraint.

Randomness is a strong requirement that implies rules 1-6: If F is a uniform random sample of P , rule 2 (and by extension rule 1 relative to the cardinality of P) is satisfied. Similarly for rule 6. If the underlying distribution of column F is the same as column P , and F is a random sample of P , then the probability that a substantial number of columns F' are random samples of F , without any real correlation between F and F' , is very small (rule 3). Similarly,

if F is a random sample of P , and F is a random sample of some other column P' with the same underlying distribution as P , then P and P' are clearly highly correlated. First, it is unlikely that a large number of such correlated columns P' exist (rule 4). Second, any such association (F, P') has high confidence if (F, P) has high confidence, so it is equally valid. Finally, if F is a random sample of P rule 5 is straightforwardly satisfied.

Nevertheless, one can come up with counter-examples for the randomness rule as well. Consider a data warehouse that contains a table P with all historical transactions, and a table F that references only the last month of transactions (for the purpose of efficiently answering queries on the latest data). If transaction identifiers are assigned using an auto-increment field, then the transaction id field in table F is a foreign key to table P and the transaction id values in F form a suffix of the ids in P . Note that this example also invalidates rule 6. Unless a foreign key constraint is specified in the schema, no formal method can decide *with 100% certainty* whether a column F is a foreign key with respect to primary key P . As mentioned above, useful fk/pk relationships are often data-dependent, and may not be specified in the schema. Hence, we cannot expect to find a solution with 100% precision/recall. However, as we show via extensive experiments over a large number of real databases, the randomness rule eliminates a very large number of false positives in practice.

In this paper we show that *Randomness* efficiently discovers meaningful foreign keys, including multi-column foreign-keys (which have not been considered by previous work). Our experiments show that our approach has higher accuracy than previously proposed methods, scales to very large datasets, and does not require any prior knowledge of the data (in contrast with the method in [17]). Our contributions are as follows:

- We define a distance measure between distributions, which allows us to quantify *Randomness*. This leads to a novel foreign key discovery rule that prunes a large number of false positives.
- We design fast approximate algorithms for evaluating randomness over a large set of columns, using quantile summaries.
- We design an I/O efficient algorithm for discovering single and multi-column foreign keys, which requires only two linear scans of the data. It outputs a list of fk/pk pairs, in descending order of their randomness scores. The score reflects the likelihood that the pair is a useful foreign key constraint.
- We present a comprehensive experimental validation of our approach using a large number of real and synthetic datasets.

2. PRELIMINARIES

Let $\mathbf{T}$ be a collection of relational tables, possibly from multiple databases. Let $\mathbf{C}$ be the set of all columns in $\mathbf{T}$. For the remainder of this paper, we use $F(P)$ to refer to both a single and multi-column foreign (primary) key. Abusing notation, F and P refer both to the names of the columns (or multi-columns) and to their respective set of distinct values (or tuples). In general, F and P are multi-sets. Let $|X|$ be the number of *distinct* values in multi-set X (or distinct tuples if X is multi-column).

We assume that the single and multi-column primary keys in $\mathbf{T}$ are known, either from schema specification or from a preprocessing phase. In the latter case, Gordian [19] can be used to compute them. Let $\mathbf{P}_s$ ($\mathbf{P}_m$) denote the set of single-column (multi-column) primary keys and $\mathbf{P} = \mathbf{P}_s \cup \mathbf{P}_m$.

Similar to previous work, in order to cope with data inconsistencies, we relax the inclusion property that a foreign key must satisfy.




Figure 2: A good foreign key F is a set of random values from the primary key. Column F' fails the randomness test.

More precisely, we require that

$$\sigma(F, P) = \frac{|F \cap P|}{|F|} \geq \theta,$$

where $\sigma(F, P)$ is the *inclusion coefficient* and θ is user-defined. In our experiments, we use $\theta = 0.9$, i.e., partial inclusion is satisfied if at least 90% of the values in F are also contained in P . We use the notation $F \subset_{\theta} P$ to denote that $\sigma(F, P) \geq \theta$.

Computing $\sigma(F, P)$ is very expensive, especially when considering the potentially very large number of multi-column candidate fk/pk pairs. Therefore, we estimate all inclusion coefficients by computing a bottom- k sketch [5] for each column (refer to Appendix C for details on bottom- k sketches). We then use the SCS estimator from [5], which estimates the Jaccard coefficient $\rho(F, P) = \frac{|F \cap P|}{|F \cup P|}$. Since $\sigma(F, P) = \frac{\rho(F, P)}{\rho(F \cup P, F)}$, we estimate $\sigma(F, P)$ by dividing the estimators for the two Jaccards.¹ Section 4 provides details on how to efficiently compute bottom- k sketches for both single and multi-column candidate keys.

3. RANDOMNESS

In this section we assume that the inclusion coefficients between all pairs of single/multi-column pks and columns in C have been computed, and pairs that do not satisfy partial inclusion have been discarded. As mentioned in Section 1, we conjecture that randomness is a strong indicator of the quality of an fk/pk pair. Formally:

DEFINITION 1 (RANDOMNESS TEST). *Given two sets of values (tuples) F and P , test the statistical hypothesis that the distinct values (tuples) in F have the same underlying distribution as the distinct values (tuples) in P .*

Figure 2 shows an example of a two-column primary key and two candidate foreign keys. Set F is a good fk, since it appears to be a random subset of values from the pk. Set F' is a contiguous subset of the pk and does not pass the randomness test.

Domain Order. The randomness test requires the existence of an underlying order over the domain of the primary and foreign keys. To see this, consider the example in Figure 3. If the values are sorted numerically, then the candidate column F is a prefix of the primary key. However, when the same values are sorted lexicographically, F falsely appears to be a random sample of the primary key. To handle this issue, we adopt the following natural convention: numeric values are sorted numerically, and strings are sorted lexicographically. The implicit assumption is that it is very rare that a column containing only numeric values is a foreign key for a primary key that contains strings (in which case it should have been sorted lexicographically, rather than numerically). When columns contain both numeric, alphanumeric, and string values, we use a combination sort (same as the Unix “sort -n” command). For multi-column keys we define an order along each dimension, as above.

¹In Section 5 we discuss two alternative estimators we considered. Each had a significant drawback. By contrast, this estimator proved highly accurate.




Figure 3: A column containing numeric values might falsely appear to be a random sample of a primary key based on lexicographic sorting of values.




Figure 4: The Wilcoxon test: 1. Sort values in multi-set $F \cup P$; 2. Assign ranks; 3. Compute the rank-sum of values in F (13.5 in this example).

Randomness measure. A standard, non-parametric statistical test for randomness is the Wilcoxon rank-sum test [18]. Assume that F, P are single-column candidate keys. Sort the values in the multi-set union $F \cup P$ and rank them. Since $F \subset_{\theta} P$, the majority of values in F appear in P , so there are duplicate values. Assign the mean rank for duplicate values (i.e., if a duplicate value is 3rd and 4th in the sorted order, it is assigned rank 3.5). Finally, compute the sum of ranks for values in F . This rank-sum is an indication of whether F and P are drawn from the same distribution. Intuitively, if the rank-sum is too small, then most values in F are contained in a prefix of P , and if the rank-sum is too large then most values in F are contained in a suffix of P .

The Wilcoxon test is straightforward for univariate distributions but does not generalize to multivariate distributions, so it cannot be used for multi-column keys. Attempting to apply the Wilcoxon test separately for each dimension of a multi-column key results in false negatives. For example, consider the multi-column key F in Figure 2. Even though F appears to be a uniform random sample of P , the projection of F in either dimension is not a uniform sample due to the multiplicity of some of the values (two points project into the same value in both dimensions). An independent Wilcoxon test in either dimension would dismiss F .

We now propose a novel approach for deciding whether two multi-dimensional sets are drawn from the same distribution. Our method computes a value that reflects how close the distributions of the two sets are. We start by defining a probability distribution for each set, so that the total probability mass is 1 (this step is detailed later in the section). A standard distance measure between two probability distributions is the Earth Mover’s Distance (EMD) [16]. The smaller the value $EMD(F, P)$ is, the closer the distributions of F and P are. The output of our algorithm is the list of (F, P) pairs, in increasing order of their (normalized) EMD values.

Intuitively, EMD measures the amount of work needed to convert the set of values of the foreign key into the set of values of the primary key. If we regard each distribution as piles of dirt spread over some space, EMD is the least amount of effort needed to convert the first set of piles into the second. The effort is the amount of dirt that needs to be moved times the distance it has to travel. Figure 5 illustrates the computation of EMD for pairs (F, P) and (F', P) from Figure 2. In this example, all points in a set have equal probability and the sum in each set is equal to 1. To convert F into P , a probability mass of 0.1 needs to be moved from each




Figure 5: EMD quantifies the amount of work required to convert one set of values into another.

point $p \in F$ to the nearest point $np(p) \in P \setminus F$. Similarly for F' and P . Since the points of F are uniformly distributed over P the average distance between p and $np(p)$ is smaller than the average distance for F' . Hence, the amount of work needed to convert distribution F to P is smaller than the one to convert F' to P .

While the definition of EMD applies to single and multi-dimensional sets, it has a crucial restriction: unlike the Wilcoxon test, EMD requires a metric distance between the values of the two distributions. A metric distance can be used only when both columns F and P contain numeric values, but not when they contain strings. Even for numeric values, using the underlying distance is undesirable because we need to be able to compare EMD values between different candidate pairs for sorting pairs according to confidence. However, given distinct F, F', P, P' , if F and P have larger ranges of values than F' and P' , then $EMD(F, P)$ will generally tend to be larger than $EMD(F', P')$, even if F is a “more random” subset of P than F' is for P' . Therefore, a uniform way of defining a distance function for numeric and string columns is needed, which is independent of the range of values in any column.

We propose using the distance between the ranks of the values in the pk column. For single-column F and P , rank all values in P in the underlying ordered space, then define the rank distance between two values in F or P to be the (absolute) difference between their ranks in P . For multi-columns F and P , define the rank distance to be the sum of single-dimensional rank distances (i.e., the Manhattan distance). However, the rank distance will still introduce bias when comparing $EMD(F, P)$ and $EMD(F', P')$ if the number of values in $|P|$ is much larger than the number of values in $|P'|$. Therefore, the rank distance is normalized by the number of values, in effect replacing ranks by quantiles:

DEFINITION 2 (QUANTILE DISTANCE). Given a multi-column set X consisting of n columns, a total order in each column X_i , a function $q_i(x)$ that returns the quantile order of value x in column X_i , and two tuples $v, w \in X$, the quantile distance is

$$d(v, w) = \sum_{1 \leq i \leq n} |q_i(v) - q_i(w)|.$$

Notice that the quantile distance is independent of the type of values in X as long as a total ordering of the values in each dimension is defined. We refer to the EMD measure using the quantile distance as *Quantile-EMD*. A final normalization is needed to compare (F, P) and (F', P') when they have different dimensionality. Let $EMD_n(F, P) = EMD(F, P)/n$, where n is the dimensionality of F and P .

Computing Quantile-EMD. We now consider the problem of efficiently approximating $EMD(F, P)$ for all pairs of candidate keys (F, P) . The first step is to define a probability distribution for F and P . The easiest choice is to let each value in F have a probability mass of $1/|F|$, and each value in P have a probability of $1/|P|$. Computing EMD is equivalent to the well-known transportation problem and can be solved by the Hungarian algorithm [7]. However, the Hungarian algorithm has cubic complexity and is very inefficient over large F and P . For our purposes, it




Figure 6: Constructing a 2-dimensional 4-quantile histogram for primary key P .

is sufficient to compute EMD on coarser probability distributions. More precisely, we use a quantile histogram to define the probability distribution in the primary key, since quantiles best approximate the original distribution w.r.t. the quantile distance. The probability distribution in the candidate foreign key is then defined with respect to the quantiles of the primary key.

For every single/multi-column key $P \in \mathcal{P}$ construct a *quantile histogram* based on the ℓ -quantiles of P (for some constant ℓ). In one dimension, the histogram is equi-depth. In multiple dimensions, compute quantiles separately on each dimension (over the distinct values in that dimension) and construct a grid based on the quantiles in each dimension. An example 2-dimensional 4-quantile histogram is shown in Figure 6. Notice that in this particular example there exists a three point tie in each dimension. After projecting the points in either dimension there are only 8 distinct values left. Hence, the 1st 4-quantile is the point with rank $8 \cdot 1/4$, the 2nd is the one with rank $8 \cdot 2/4$, etc. The probability distribution of P based on the corresponding histogram is defined as:

DEFINITION 3 (QUANTILE HISTOGRAM). Given a multi-column primary key P consisting of n columns, let $Q_i = \{q_1^i, \dots, q_{\ell}^i\}$ be the ℓ_i -quantiles of P in column i (different columns may have different number of quantiles). Let $G_P = Q_1 \times \dots \times Q_n$ be the corresponding n -dimensional quantile grid. The quantile histogram $\bar{P}$ is defined as the number of values of P within each grid cell of G_P . The total number of grid cells is $|G_P| = \ell_1 \times \dots \times \ell_n$. The probability distribution over P is defined as the normalized $\bar{P}$; i.e., the count in each cell is divided by $|P|$.

For a candidate multi-column F , the probability distribution histogram based on the quantile grid G_P of P is defined as:

DEFINITION 4 (DISTRIBUTION HISTOGRAM). Given a candidate pair (F, P) , the distribution histogram $\bar{F}_P$ of F with respect to P is defined as the number of distinct values of F within each grid cell of the quantile grid G_P . The probability distribution over F is defined as the normalized $\bar{F}_P$; i.e., the count in each cell is divided by $|F|$.

We now describe how to approximate $EMD(F, P)$ using the quantile histograms. Assume that the probability mass of a grid cell is concentrated in its upper right corner. Therefore, the distance between two grid cells is defined as the quantile distance between the upper right corners of the cells. For example, in Figure 6, the distance between grid cells A and B is $(3/4 - 2/4) + (2/4 - 1/4) = 0.5$. As before, the Hungarian algorithm is used to compute the EMD between the two distributions. The input size is now $|G_P| = \ell^n \ll |P|$ (usually $1 \leq n \leq 4$ and ℓ is small). Once the normalized histograms are computed, the method requires no additional access to the raw data. Note that the value ℓ need not be the same for all primary keys. Since Quantile-EMD uses the quantile distance we can compare $EMD(F, P)$ and $EMD(F', P')$ even if the quantile histograms were computed for different values of ℓ . This is important, since some primary keys may have only a few values. On the other hand, a larger ℓ for larger primary keys will improve accuracy.

Now we can bound the approximation error for the Quantile-

EMD in the grid space G_P versus the initial space P (the proof is in Appendix D).

LEMMA 1. *Let n be the space dimensionality, $1/\ell$ be the side length of the cells in G_P (in every dimension without loss of generality), and $EMD_{n,P}$, EMD_{n,G_P} be the normalized EMD in the primary space P and reduced space G_P respectively. Then*

$$|EMD_{n,P} - EMD_{n,G_P}| \leq 2/\ell.$$

4. DISCOVERING FOREIGN KEYS

Throughout this section, we use the notations from Section 2 (see also Table 5 in the Appendix). To discover foreign keys, we first compute inclusion between all pairs of primary keys and columns in $\mathbf{C}$ and then evaluate randomness only on the pairs satisfying inclusion. We accomplish this by computing bottom- k sketches and quantile histograms with two passes over the data. A pseudocode of the algorithm described below appears in Appendix B.

For single-column candidate foreign keys and single/multi-column primary keys, the bottom- k sketches can be computed in one linear scan of the database. Multi-column candidate foreign keys are challenging due to their potentially large number. However, if P is a primary key consisting of columns $(C_1, \dots, C_n)$ and F is a candidate foreign key consisting of columns $(C'_1, \dots, C'_n)$, then $\sigma(F, P) \leq \min_{i=1}^n \sigma(C'_i, C_i)$. Hence, it is sufficient to consider only candidates F such that $\sigma(C'_i, C_i) \geq \theta$, for all i , where all C'_i s belong to the same table. We expect only a small number of pairs (F, P) to have these properties (and confirm this experimentally). For such pairs (F, P) , we compute the bottom- k sketch of F and estimate $\sigma(F, P)$, with a second pass over only the relevant columns in F (recall that the sketch of P has been computed during the first pass).

Quantile histograms for a single-column primary key P can be computed exactly in one linear scan if there exists an index on P – as is usually the case for pks – by reading P in sorted order and computing the quantiles incrementally (this requires knowledge of $|P|$ which can be found from table statistics). If an index does not exist we can approximate the quantiles in linear time using quantile summaries [9]. The distribution histograms of single-column candidate fks can be trivially computed in linear time after the quantile histograms of all pks have been computed and stored in memory; for an fk F we simultaneously compute all histograms w.r.t. all pks P for which (F, P) passes inclusion. In our experiments, the average number of such pairs, for a fixed F , was less than 10.

For a multi-column primary key P , the quantile histogram requires two passes over the data. In the first pass, we compute the quantile grid for each column $C \in P$ (either using an index or a quantile summary) and construct a multi-dimensional quantile grid. A subtle point here is that a column C of a multi-column pk might not be a pk itself; estimating quantiles on the distinct values in C requires using duplicate insensitive quantile summaries [6]. In the second pass, we scan P and populate the quantile grid. We also scan each multi-column fk F and simultaneously compute all relevant distribution histograms. In our experiments, the average number of such pairs, for a fixed F , was less than 5.

We now summarize each of the two linear scans:

Phase 1. Read all columns in table-wise order (i.e., row by row) and build bottom- k sketches for all columns in $\mathbf{C}$, as well as for all multi-column primary keys in $\mathbf{P}_m$. Also build quantile grids for all single/multi-column primary keys. All structures are stored in two hash tables B (for bottom- k sketches) and Q (for quantile/distribution histograms), using the name of the column(s) as the hash key. Evaluate all (single-column) inclusions between $F \in \mathbf{C}$ and $P \in \mathbf{P}_s$ and store candidate pairs in $\mathbf{F}_s$. Finally, evaluate

(single-column) inclusions between $C \in \mathbf{C}$ and $C_i \in P$, $P \in \mathbf{P}_m$ and store candidates in $S(C_i)$.

Phase 2. For each multi-column pk $P = (C_1, \dots, C_n)$, consider the n sets $S(C_i) = \{C \in \mathbf{C} \mid \sigma(C, C_i) \geq \theta, 1 \leq i \leq n\}$ computed in Phase 1. Then $F = (C'_1, \dots, C'_n) \in S(C_1) \times \dots \times S(C_n)$ is a candidate fk for P if there exists a table T s.t. $\forall 1 \leq i \leq n : C'_i \in T$. Compute, for each multi-column pk P the set of its candidate foreign keys and insert pairs (F, P) in $\mathbf{F}_m$. This requires access only to the sets $S(C_1), \dots, S(C_n)$, which are stored in memory.

Next, for each $(F, P) \in \mathbf{F}_s$, scan each single-column F and compute its distribution histograms w.r.t. all relevant primary keys P . Compute $EMD_1(F, P)$ and store it in memory. For each multi-column candidate $(F, P) \in \mathbf{F}_m$ scan F and compute its multi-column bottom- k sketch as well as its distribution histograms w.r.t. all P . For each such P , verify whether (F, P) satisfies inclusion (recall that the multi-column bottom- k sketch of P was computed in Phase 1). If (F, P) does not pass the test, discard the distribution histogram $\bar{F}_P$ and remove (F, P) from $\mathbf{F}_m$. Finally, compute $EMD_n(F, P)$ for all $(F, P) \in \mathbf{F}_m$ and return $\mathbf{F} = \mathbf{F}_s \cup \mathbf{F}_m$, sorted in increasing order of EMD_n values.

5. EXPERIMENTAL EVALUATION

We evaluate our algorithm for discovering fk/pk constraints on two benchmark synthetic databases (TPC-E and TPC-H), as well as on two real datasets: a Wikipedia (WP) snapshot from March 2008 and an IMDB snapshot from January 2010. All these datasets come with a schema specification (see Appendix E for details). We implemented our algorithms in C++, and performed the experiments on an Intel Core2 Duo 2.33 GHz CPU with 4GB RAM running MySQL. We use three standard accuracy measures to evaluate our method: *precision*, *recall* and *F-measure* (the harmonic mean of precision and recall). Each measure is applied to two sets of fk/pk constraints: the “golden standard” set specified in the schema, and the top- $X\%$ constraints reported by our algorithm (for various X). We start by evaluating the accuracy of EMD computation, then evaluate the overall algorithm for both accuracy and scalability. In addition, we discuss how the results change if we also take into account the similarity of column names. Finally, we compare our results with the machine learning approach of Rostin et al. [17], which uses the 7 rules discussed in Section 1. For completeness, we also include experiments on bottom- k sketches, which show that the partial inclusion estimator is highly accurate and does not influence the overall results.

5.1 EMD Computation

By Lemma 1, the error in the computation of $EMD(F, P)$ is bounded by $2/\ell$, where ℓ is the number of quantiles in each dimension. In practice, this error is negligible even for small ℓ . Table 1 shows the average difference between the EMD computed using $\ell \in \{4, \dots, 1024\}$ quantiles and 2048 quantiles. The averages are over all columns in the TPC-H database that pass inclusion. The reason we do not compute the differences to the actual EMD value (i.e., using all quantiles) is that EMD has cubic complexity, which is very expensive to compute when the primary key has $\Omega(10^4)$ values. Clearly, the EMD values converge very quickly as ℓ increases. Therefore, small values of ℓ are sufficient.

The last column of Table 1 shows that, when using 256 approximate quantiles, the difference is only 10^{-4} bigger than for 256 exact quantiles. This is not surprising, since the quantile histograms of foreign keys should remain roughly the same over small shifts in the quantile grids. The main difference between approximate versus exact quantiles is computation time. To our surprise, com-

ℓ		4	16	64	256	1024	256(A)
Diff		0.06	0.009	0.002	$3 \cdot 10^{-4}$	$4 \cdot 10^{-5}$	$4 \cdot 10^{-4}$
Top-25%	R	1	1	1	1	1	1
	P	0.9	0.9	0.9	0.9	0.9	0.9
Top-20%	R	0.78	0.78	0.89	0.89	0.89	0.78
	P	0.88	0.88	1	1	1	0.88
Top-15%	R	0.56	0.67	0.67	0.67	0.67	0.56
	P	0.83	1	1	1	1	0.83

Table 1: EMD accuracy for different quantile grid sizes; Diff=EMD_{n,G_ℓ} − EMD_{n,G₂₀₄₈}.

Dataset	TPC-H		TPC-E		WP		IMDB	
SC-FK	9		44		10		8	
MC-FK	1		1		0		0	
$\theta =$	0.9	1	0.9	1	0.9	1	0.9	1
Cand. SC	38	34	304	214	12	8	24	24
Cand. MC	1	1	4	3	0	0	0	0

Table 2: Number of candidate pairs that satisfy inclusion; SC=single-column, MC=multi-column.

puting exact quantiles was about 20% faster! The reason is that the vast majority of primary keys are indexed, so using an ORDER BY SQL query had the complexity of a linear scan. By contrast, computing approximate quantiles incurred the overhead of maintaining the quantile summaries [9]. Henceforth, all reported results are for exact quantiles.

Since normalized EMD values determine the order in which we report fk/pk pairs, we also measured the effect of quantile grid sizes on the precision and recall of the final results. Smaller grid sizes, as well as approximate quantiles, do impact both precision (P) and recall (R) for the top-15% and top-20% reported pairs (see last 4 rows in Table 1). However, they have no effect for the top-25% pairs. The reason is that EMD values are very close to each other within the top-25% pairs, so even small changes impact the order of results. However, pairs below top-25% have EMD values at least one order of magnitude larger. As we discuss below, we use precisely this jump in EMD values to determine the best set of constraints to present to a user. For TPC-H, the best set is the top-25%, which remains unchanged for all grid sizes.

5.2 Overall Algorithm

For all experiments in this section, we use $\ell = 256$ quantiles for single-column primary keys, and $\ell = 16$ quantiles per dimension for multi-column primary keys. Table 2 shows the number of candidate fk/pk pairs that pass the inclusion test for each dataset. This illustrates the large number of false positives our algorithm must eliminate. For example, 217 pairs pass inclusion for TPC-E (214 single-column pairs and 3 multi-column pairs). Of these, only 45 are specified in the schema.

To evaluate the utility of our method, we measure the precision, recall and F-measure after selecting the top- $X\%$ results as the answer set, and comparing them with the golden standard specified in the schema. We report the results in two groups, for reasons that we explain later in this section.

TPC-H, WP and IMDB: Figure 7 shows the results for these three databases. A larger X (i.e., a bigger answer set) implies a larger number of false positives, hence lower precision. On the other hand, a smaller X has more false negatives (i.e., undiscovered fk/pk constraints), and thus worse recall. There is a sweet spot, which depends on each dataset, that balances precision and recall. That spot also corresponds to a big jump in the respective EMD values; we illustrate this by plotting the EMD values on the same graph. For TPC-H the sweet spot occurs at $X = 25$, for IMDB at $X = 35$, and for WP at $X = 80$. Thus, by examining




Figure 8: Utility measures on TPC-E using the golden standard and extended constraints.

Foreign Key		Primary Key	
Column	Values	Column	Values
Exchange.ex_ad_id	1-4	Address.ad_id	1-7504
Company.co_ad_id	5-2504	Address.ad_id	1-7504
Customer.c_ad_id	2505-7504	Address.ad_id	1-7504
Trade.t_st_id	B	Status.type.st_id	A, . . . , E
Broker.b_st_id	A	Status.type.st_id	A, . . . , E
Company.co_st_id	A	Status.type.st_id	A, . . . , E
Customer.c_st_id	A	Status.type.st_id	A, . . . , E
Security.s_st_id	A	Status.type.st_id	A, . . . , E

Table 3: False negatives in TPC-E (A=Active, B=Completed, C=Canceled, D=Pending, E=Submitted).

the significant jumps in EMD values, the algorithm automatically proposes one or more answer sets deemed “relevant” (note that for WP, an EMD jump also occurs for $X = 60$). An answer set can then be verified either experimentally (e.g., by running queries and testing if the results are meaningful), or by a domain expert.

For all three datasets in Figure 7, we achieve F-measure above 0.8 at their respective sweet spot X . For TPC-H, the answer set (top-25%) has only one false positive: the pair of columns Part-Supp.PS.AVAILQTY and Supplier.S.SUPPKEY, the first containing values from 1 to 9999 and second from 1 to 10000. Clearly, only a supervised algorithm would be able to recognize this false positive. For IMDB, the answer set is exactly the golden standard.

For WP, the loss in recall originates from an unlikely source. The schema specifies that ImageLinks.il_to is a foreign key to the primary key Image.img_name. However, its inclusion coefficient is only 51%. Since we set the inclusion threshold at 90%, this results in a false negative. (We note that three other specified fk/pk pairs have inclusion coefficient below 1, but higher than 0.9). Such data inconsistencies verify the intuition that real data does not always follow the ideal rules from database theory.

TPC-E: For this database, we report two sets of experiments. In the first set (Figure 8(a)) we measure precision/recall with respect to the golden standard specified in the schema. Notice that the F-measure is below 0.6 across the board. A careful analysis of the data reveals that the reported accuracy is misleading: many false positives occur either because of *symmetry* or *transitivity*. Symmetry refers to the case when a pair of primary key columns (A, B) is specified in the schema as fk/pk in one direction, but not in the other. However, (B, A) is clearly a valid constraint in this case. Transitivity occurs when, for three columns A, B and C , with B and C being primary keys in their respective tables, the constraints (A, B) and (B, C) are specified in the schema, while (A, C) is not (although it is clearly valid). By applying symmetry and transitivity rules, we extend the set of valid constraints against which we test our results. Reporting precision/recall with respect to this augmented set of constraints improves results significantly; see Figure 8(b).

Nevertheless, none of the measures reaches 1. We attribute this




Figure 7: Utility measures on TPC-H, Wikipedia and IMDB.




Figure 9: Scalability results.

to the data generation process itself. A total of 15 false positive pairs (some pairs are counted in both directions) are between only eight columns. These columns belong to seven different tables and contain exactly the same number of rows and number of distinct values: the numbers 1 to 5000. Clearly, only a domain expert can label them as false positives (in our case, we used the extensive TPC-E documentation). The algorithm also fails to discover 8 out of 45 true constraints and 28 implied constraints; the pairs are shown in Table 3. Five of these foreign keys contain only one distinct value (either the status ‘Completed’ or ‘Active’). Clearly, the generator assigns a default value for this column for every row in the table, since not all trades in Trade can be completed, while all trades in Broker are active. One column contains address identifiers 1 to 4 even though the corresponding primary key contains 7504 distinct addresses. Finally, the other two address columns are (almost) a prefix and a suffix of the primary key and constitute a counter-example for the randomness rule.

5.3 Scalability

We tested the scalability of our method on TPC-H, for which it is easy to generate instances of progressively larger sizes. We used five instances with sizes 1MB, 10MB, 100MB, 1GB and 10GB. The running times for each of the two phases, as well as the total time are shown in Figure 9. For readability, we use a logarithmic scale on both axes. As expected, each phase takes linear time. The second phase is faster because we only have to scan the columns that satisfy inclusion (while in the first case, we scan all columns). For the 10GB instance, the total running time is less than 2.5 hours, making our method applicable to enterprise-scale datasets.

5.4 Column Names

So far our foreign key discovery process has been a data-driven approach. However, it is easy to enhance it by considering the orthogonal approach of looking at the column names (rule 7 in Section 1). As shown in [10], comparing column names is not necessarily straightforward, and can lead to false conclusions. For example, in TPC-E columns that form valid fk/pk constraints have very different names, because they contain an abbreviation of the table name as a prefix (e.g., columns Trade History.th.Lid and Trade.Lid are an fk/pk pair; the prefixes ‘th’ and ‘t’ in the column names stand

	Precision	Recall	F-measure
TPC-H	1	1	1
TPC-E	0.57	0.82	0.67
TPC-E Ext.	1	0.89	0.94
IMDB	1	1	1

Table 4: Results after eliminating non-matching column names.

for Trade History and Trade respectively). Fortunately, TPC-E has extensive documentation that explains the naming conventions, so we can delete these prefixes and compare the remaining strings. The resulting names are identical only if the pair is a valid constraint.

We are not aware of any method for automatically determining which string similarity measure to use for any given schema. In Table 4, we report our results using string identity (for TPC-E, we apply this to column names after deleting their table prefixes). The results are generated as follows: First, we compute for each database the most relevant answer set, i.e., the top- $X\%$ for the best value X . We then delete all pairs from these sets whose column names are not identical, and compute the precision/recall on the resulting answer set. For TPC-E, we also report results using the extended set of valid constraints. For WP there is no single pair with identical column names, hence we exclude it from this experiment.

5.5 Comparison With Alternatives

The algorithm of Rostin et al. [17] uses a learning phase to train four different classifiers that are then used to discover single-column keys only. Each classifier uses a training set consisting of known fk/pk pairs from four out of five different databases. The goal is to learn the relative importance of rules 1-7 stated in Section 1, then apply them to the fifth database. No classifier was consistently the best across all datasets.

We compare our results over TPC-H using the results already reported in [17]. As reported in that paper, the best classifier (J48) for the TPC-H dataset results in F-measure equal to 0.95, with the average value over all classifiers being 0.915. The success of J48 for TPC-H can be largely attributed to the use of rule 7 (matching column names), making TPC-H an easy target. Our method achieves an F-measure of 1 for TPC-H when using column names (even without column names, F-measure is 0.95).

5.6 Inclusion Estimators

We considered two alternative estimators for the inclusion coefficient $\sigma(F, P)$:

1. The estimator proposed in [3], which is unbiased. However, it is defined over sketches whose sizes are a user-defined fraction of the size of the original column. This is generally too large for practical purposes (e.g., the Wikipedia database has size $O(10^9)$, so the size of 1%-sketches is $O(10^7)$). Note




Figure 10: Accuracy of bottom- k estimators for the inclusion coefficient, as a function of k .

that the size of sketches impacts not just storage requirements, but more importantly, the running time for computing all inclusion coefficients.

2. Divide the estimated value of $|F \cap P|$ by the estimated value of $|F|$. Both estimated values are computed using the estimators proposed in [2]. The advantage is that these estimators work over bottom- k sketches which have constant size. Figure 10 shows an experimental comparison (over TPC-E) of this estimator, denoted Estimator2, to the one described in Section 2, which we denote as Estimator1. Clearly, Estimator1 is significantly more accurate, and it also uses bottom- k sketches. Therefore, all our experiments use Estimator1 for the inclusion coefficient. We set $k = 256$.

6. RELATED WORK

The state of the art for efficient, automatic discovery of single and multi-column primary keys is Gordian [19]. Gordian formulates the problem as a cube computation that corresponds to the computation of the entity counts of all possible column projections.

Surprisingly, very little work has dealt with discovery of foreign keys. Most work focuses on computing inclusion dependencies only [1, 4, 12–15]. Recently, Bauckmann et al. [1] proposed SPIDER for efficiently detecting single-column inclusion dependencies. The algorithm first sorts the distinct values in all columns in C and then uses a parallel merge-sort like algorithm to compute all inclusions simultaneously. SPIDER computes inclusions exactly, but the cost is super-linear to the size of the data. A similar approach was proposed by Marchi et al. [14], using a linear pass over the data to compute an inverted index over each data type (e.g., strings, floats, integers). Subsequent passes over the index can discover single/multi-column inclusions. Marchi and Petit [15] propose a hybrid technique based on association rule mining to find low-dimensional inclusions and an optimistic exploration of high-dimensional inclusions using clique-finding. Koeller and Rundensteiner [12] utilize clique-finding for discovering high-dimensional inclusions. Partial inclusion is not addressed in these works. Lopes et al. [13] use a query workload based approach to discover foreign key relationships based on the assumption that SQL join queries use fk/pk keys. This approach is based on the availability of a query workload. Furthermore, inclusion is not a sufficient condition for foreign keys, resulting in a large number of spurious keys. Rostin et al. [17] introduced a machine learning approach for discovering foreign keys that is based not only on inclusion, but on a variety of other properties of good foreign keys. The authors use the SPIDER algorithm to discover all inclusion dependencies. Most importantly, the algorithm requires a learning step, which implies the availability of datasets with known fk/pk keys. The quality of the training dataset affects performance significantly. Finally, multi-column keys are not addressed in that work. Dasu et al. [8] proposed using minhash sketches to find potential associations between columns (or sets of columns) as a function of Jaccard

coefficient. However, Jaccard is not a good indicator of inclusion coefficient when set sizes differ substantially.

7. CONCLUSION

In this work we introduced the notion of *Randomness* and showed that it can be used effectively to reduce a large number of false positive pairs produced by partial inclusion. We also provided an efficient approximation algorithm for evaluating randomness between pairs of multi-column candidate keys. We presented a combined algorithm that can discover good single/multi-column foreign keys with only two linear scans over the data. Finally, we performed a comprehensive experimental evaluation showing the efficacy of our techniques.

8. ACKNOWLEDGEMENT

The work of Meihui Zhang and Beng Chin Ooi were in part supported by Singapore NRF grant R-252-000-376-279.

9. REFERENCES

- [1] J. Bauckmann, U. Leser, F. Naumann, and V. Tietz. Efficiently detecting inclusion dependencies. In *ICDE*, pages 1448–1450, 2007.
- [2] K. Beyer, P. J. Haas, B. Reinwald, Y. Sismanis, and R. Gemulla. On synopsis for distinct-value estimation under multiset operations. In *SIGMOD*, pages 199–210, 2007.
- [3] A. Broder. On the resemblance and containment of documents. In *SEQUENCES*, pages 21–30, 1997.
- [4] M. A. Casanova, R. Fagin, and C. H. Papadimitriou. Inclusion dependencies and their interaction with functional dependencies. In *PODS*, pages 171–176, 1982.
- [5] E. Cohen and H. Kaplan. Leveraging discarded samples for tighter estimation of multiple-set aggregates. *Joint Intl. Conf. on Measurement and Modeling of Comp. Syst.*, pages 251–262, 2009.
- [6] J. Considine, M. Hadjieleftheriou, F. Li, J. W. Byers, and G. Kollios. Robust approximate aggregation in sensor data management systems. *TODS*, 34(1), 2009.
- [7] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2001.
- [8] T. Dasu, T. Johnson, S. Muthukrishnan, and V. Shkapenyuk. Mining database structure; or, how to build a data quality browser. In *SIGMOD*, pages 240–251, 2002.
- [9] M. Greenwald and S. Khanna. Space-efficient online computation of quantile summaries. In *SIGMOD*, pages 58–66, 2001.
- [10] J. Kang and J. F. Naughton. On schema matching with opaque column names and data values. In *SIGMOD*, pages 205–216, 2003.
- [11] D. Kifer, S. Ben-David, and J. Gehrke. Detecting change in data streams. In *VLDB*, pages 180–191, 2004.
- [12] A. Koeller and E. A. Rundensteiner. Discovery of high-dimensional inclusion dependencies. In *ICDE*, pages 683–685, 2003.
- [13] S. Lopes, J.-M. Petit, and F. Toumani. Discovering interesting inclusion dependencies: application to logical database tuning. *Information Systems*, 27(1):1–19, 2002.
- [14] F. D. Marchi, S. Lopes, and J.-M. Petit. Unary and n-ary inclusion dependency discovery in relational databases. *Journal of Intelligent Information Systems*, 32(1):53–73, 2009.
- [15] F. D. Marchi and J.-M. Petit. Zigzag: a new algorithm for mining large inclusion dependencies in databases. In *ICDM*, pages 27–34, 2003.
- [16] S. Peleg, M. Werman, and H. Rom. A unified approach to the change of resolution: space and gray-level. *TPAMI*, 11(7):739–742, 1989.
- [17] A. Rostin, O. Albrecht, J. Bauckmann, F. Naumann, and U. Leser. A machine learning approach to foreign key discovery. In *WebDB*, 2009.
- [18] S. Siegel and N. Castellan. *Nonparametric statistics for the behavioral sciences*. McGraw-Hill, Inc., second edition, 1988.
- [19] Y. Sismanis, P. Brown, P. J. Haas, and B. Reinwald. Gordian: Efficient and scalable discovery of composite keys. In *VLDB*, pages 691–702, 2006.

APPENDIX

A. NOTATIONS

Table 5 summarizes the notations used in the paper.

Symbol	Description
T	Set of tables
C	Set of columns
P_s	Single-column primary keys
P_m	Multi-column primary keys
F_s	Single-column candidate foreign keys
F_m	Multi-column candidate foreign keys
B	Hash table for bottom- k sketches
Q	Hash table for quantile/distribution histograms
F	Single/multi-column candidate foreign key
P	Single/multi-column primary key
$\hat{C}$	Bottom- k sketch of C
$\hat{F}$	Bottom- k sketch of F
$\hat{P}$	Bottom- k sketch of P
$\bar{P}$	Quantile histogram of P
$\bar{F}_P$	Distribution histogram of F with respect to P
$\sigma(F, P)$	Inclusion coefficient
θ	User-defined threshold for inclusion coefficient

Table 5: Notation used throughout the paper.

B. ALGORITHM

Algorithm B.1: DISCOVER FOREIGN KEYS(C, P_s, P_m, θ)

Phase 1.

```

 $F_s \leftarrow \emptyset, F_m \leftarrow \emptyset, B \leftarrow \emptyset, Q \leftarrow \emptyset, S \leftarrow \emptyset$ 
for all  $C \in C : B[C] \leftarrow \hat{C}$ 
for all  $P = \{C_1, \dots, C_n\} \in (P_s \cup P_m)$ 
  for  $p \leftarrow 1$  to  $n$ 
    for all  $C_f \in C$ 
      do
        if  $\sigma(\hat{C}_f, \hat{C}_p) \geq \theta \wedge \hat{C}_f, \hat{C}_p \in B$ 
          then
            if  $n = 1 : F_s \leftarrow (C_p, C_f)$ 
            if  $n > 1 : S[P, C_p] \leftarrow C_f$ 
        if  $n > 1 : B[P] \leftarrow \hat{P}$ 
        /*For  $n = 1, \hat{P}$  already in  $B$ */
         $Q[P] \leftarrow \bar{P}$ 

```

Phase 2.

```

for all  $P = \{C_1, \dots, C_n\} \in P_m$ 
  for all  $T \in T$ 
     $F_m \leftarrow (\{C'_1, \dots, C'_n\} \mid C'_i \in S[P, C_i] \cap T, P)$ 
    for all  $F = (\{C'_1, \dots, C'_n\}, P) \in F_m$ 
      Build  $\hat{F}$ 
      if  $\sigma(\hat{F}, \hat{P}) \geq \theta \wedge \hat{P} \in B$ 
        then
           $Q[P] \leftarrow \bar{P}$ 
           $Q[F] \leftarrow \bar{F}_P$ 
        else Remove  $(F, P)$  from  $F_m$ 
  for all  $(F, P) \in F_s : Q[F] \leftarrow \bar{F}_P$ 

```

```

for all  $(F, P) \in (F_s \cup F_m)$ 
  Compute  $EMD_n(F, P)$  /*Using  $Q$ */
Output  $F = F_s \cup F_m$  in increasing order of  $EMD_n$ 

```

C. BOTTOM- K SKETCHES

Given a set F , a bottom- k sketch $\hat{F}$ for F is computed as follows: Assign ranks to all values in F uniformly at random, and let




Figure 11: Constructing a Bottom- k sketch.

$\hat{F}$ be the set of k values with the smallest ranks. In practice, to compute the rank assignment we choose a hash function h , hash each value in F , and keep the k values corresponding to the smallest k hash values. If F is a set of tuples, rather than simple values, we first concatenate all values in a tuple using a predefined field separator and hash the resulting string as a whole. Figure 11 shows an example bottom-1 sketch for a set of tuples. Clearly, a bottom- k sketch can be computed in one pass over F .

Bottom- k sketches have been used to estimate various measures, such as the Jaccard coefficient $\rho(F, P) = \frac{|F \cap P|}{|F \cup P|}$ (see [5]) or the intersection size $|F \cap P|$ (see [2]). The estimators require that the same hash function h be used for computing both bottom- k sketches $\hat{F}$ and $\hat{P}$ (hence, the sketches are called *coordinated*).

D. PROOF OF LEMMA 1

PROOF. Let $EMD_P = n \cdot EMD_{n,P}$ and $EMD_{G_P} = n \cdot EMD_{n,G_P}$ be the unnormalized EMD values in the primary space P and the reduced space G_P respectively. Consider a movement of mass m that EMD_P executes in the primary space P , from a point p to a point q . Its cost is $m \cdot d(p, q)$. Let a and b be the upper right corners of the cells that contain p , respectively q . Then we can define a valid movement of mass m in the space G_P , between a and b . The cost of this mass movement is $m \cdot d(a, b) \leq m(d(a, p) + d(p, q) + d(q, b)) \leq md(p, q) + m \frac{2n}{\ell}$. Making this transformation for all mass movements in EMD_P , we obtain a valid mass movement in G_P , of cost at most $EMD_P + 2 \frac{n}{\ell} \sum m \leq EMD_P + 2 \frac{n}{\ell}$ (the sum is over all the mass moved in EMD_P). Since EMD_{G_P} is the *minimum* cost movement in G_P , we deduce $EMD_{G_P} \leq EMD_P + 2n/\ell$. A similar argument holds for the other inequality, by transforming mass movements from EMD_{G_P} into valid mass movements in P . We deduce that $|EMD_P - EMD_{G_P}| \leq \frac{2n}{\ell}$. $\square$

E. DATASET DESCRIPTIONS

The datasets can be downloaded from the following sites: TPC-H from <http://www.tpc.org/tpch>, TPC-E from <http://www.tpc.org/tpce>, WP from <http://www.archive.org/details/enwiki-20080312>, IMDB from <http://www.imdb.com/interfaces>.

When generating instances for the synthetic datasets, we use the following parameter settings: For TPC-H we use scale factor 1. For TPC-E we use 1000 customers, 20 trading days, and scale factor 1000. The characteristics of all datasets are given in Table 6, where $|T|$ is the number of non-empty tables, $|C_T|$ and $\max |C_T|$ are the average and maximum number of columns per table, and $|R_T|$ and $\max |R_T|$ are the average and maximum number of rows per table.

All these datasets come with a schema specification. Table 7 summarizes the single/multi-column foreign/primary keys explicitly stated in each schema. Notice that, e.g., TPC-E specifies nine 2-column primary keys but only one 2-column foreign key.

F. SCHEMA AND DATA UPDATES

Our methods can easily handle insertions and deletions of new tables and columns given the existing bottom- k sketches and quantile/distribution histograms. Let the new set of columns be C' .

	$ \mathbf{T} $	$ \mathbf{C}_T $	$\max \mathbf{C}_T $	$ \mathbf{R}_T $	$\max \mathbf{R}_T $
TPC-H	8	8	16	1082504	6000003
TPC-E	32	6	24	171127	4469625
WP	15	6	16	24356005	227867141
IMDB	9	2	2	1136607	5107802

Table 6: Datasets characteristics.

	TPC-H		TPC-E		WP		IMDB	
	PK	FK	PK	FK	PK	FK	PK	FK
SC	6	9	20	44	5	10	5	8
MC	2	2	1	9	1	7	—	4
	3	—	—	2	—	3	—	—
	4	—	—	1	—	—	—	—
Total	8	10	32	45	15	10	9	8

Table 7: Foreign/primary keys according to schema specifications.

First, identify new primary keys and insert them in $\mathbf{P}_s, \mathbf{P}_m$. Then, re-run Algorithm B.1 on $\mathbf{C}', \mathbf{P}_s, \mathbf{P}_m, \theta$, building only the new bottom- k and quantile/distribution histograms, as necessary.

Handling data insertions and deletions on existing columns is a little harder. Existing bottom- k sketches can easily be updated under insertions only. The new values are simply hashed and inserted in the corresponding sketches if necessary. However, deletions are not straightforward: if a deleted value was part of the bottom- k sketch, a rescan of the corresponding column is needed in order to identify the new k -th minimum hash value. One way to handle deletions without rescanning the data is to maintain larger bottom- k sketches (e.g., twice as large as needed). That way, we only rescan the data infrequently. Since in practice we expect a balanced insertions and deletions workload, this simple strategy is likely to obviate the need of a rescan in most settings.

Updating the quantile/distribution histograms is generally hard, both under insertions and deletions. A small number of insertions or deletions can be accommodated by identifying the histogram cells that contain the respective tuples, and incrementing or decrementing their counters. However, if a large amount of data is inserted or deleted, the distribution of the underlying columns is likely to change. As a result, the quantile grids on each dimension also change. This requires rescanning the data in order to compute a new quantile grid and a new histogram. A simple way of reducing the cost of updates is to use the existing quantile grid for a batch of updates and rebuild it only after a certain number of updates. Depending on the application and the underlying data we can also use well known techniques to detect a change in the distribution and trigger a rebuild [11].