



# Proceedings of the VLDB Endowment

Volume 8, No. 3 – November 2014

**Proceedings of the 41st International Conference on  
Very Large Data Bases, Kohala Coast, Hawaii**

Program Chairs and Editors-in-Chief:

**Chen Li and Volker Markl**

Associate Editors – Research and Innovative Systems Tracks:

**Kevin Chang, Shivnath Babu, Magdalena Balazinska, Felix Naumann, Stefan Manegold, Yi Chen, Fatma Ozcan,  
Jignesh Patel**

Associate Editors – Experiments and Analysis Track:

**Rainer Gemulla**

Proceedings Chairs:

**Tyson Condie, Daisy Zhe Wang**

PVLDB – Proceedings of the VLDB Endowment

Volume 8, No. 3, November 2014.

The 41st International Conference on Very Large Data Bases, Kahola Coast, Hawaii.

## Copyright 2014 VLDB Endowment

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>. Obtain permission prior to any use beyond those covered by the license. Contact copyright holder by emailing [info@vldb.org](mailto:info@vldb.org).

Volume 8, Number 3, November 2014: VLDB 2015

Pages ii - ix and 185 - 328

ISSN 2150-8097

Additional copies only online at: [portal.acm.org](http://portal.acm.org), [arxiv.org/corr](http://arxiv.org/corr), and [www.vldb.org](http://www.vldb.org)

## TABLE OF CONTENTS

### Front Matter

Copyright Notice .....	ii
Table of Contents .....	iii
VLDB 2015 Organization and Review Board .....	iv

### Letters

Letter from the Managing Editor and Advisory Committee .... <i>Divesh Srivastava and S. Sudarshan</i>	ix
---	----

### Research Papers

Coordination Avoidance in Database Systems .....	185
..... <i>Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Hellerstein, Ion Stoica</i>	
QuickFOIL: Scalable Inductive Logic Programming.....	197
..... <i>Qiang Zeng, Jignesh M. Patel, David Page</i>	
Staring into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores .....	209
..... <i>Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, Michael Stonebraker</i>	
Multi-Objective Parametric Query Optimization.....	221
..... <i>Immanuel Trummer, Christoph Koch</i>	
Deployment of Query Plans on Multicores.....	233
..... <i>Jana Giceva, Gustavo Alonso, Timothy Roscoe, Tim Harris</i>	
E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing .....	245
<i>Rebecca Taft, Essam Mansour, Marco Serafini, Jennie Duggan, Aaron J. Elmore, Ashraf Aboulnaga, Andrew Pavlo, Michael Stonebraker</i> .....	
Beyond Itemsets: Mining Frequent Featuresets over Structured Items .....	257
..... <i>Saravanan Thirumuruganathan, Habibur Rahman, Sofiane Abbar, Gautam Das</i>	
Inferring Continuous Dynamic Social Influence and Personal Preference for Temporal Behavior Prediction .....	269
..... <i>Jun Zhang, Chaokun Wang, Jianmin Wang, Jeffrey Xu Yu</i>	
Large-Scale Distributed Graph Computing Systems: An Experimental Evaluation .....	281
..... <i>Yi Lu, James Cheng, Da Yan, Huanhuan Wu</i>	
Faster Set Intersection with SIMD instructions by Reducing Branch Mispredictions .....	293
..... <i>Hiroshi Inoue, Moriyoshi Ohara, Kenjiro Taura</i>	
Scalable Topical Phrase Mining from Text Corpora.....	305
..... <i>Ahmed El-Kishky, Yanglei Song, Chi Wang, Clare R. Voss, Jiawei Han</i>	
Efficient Top-K SimRank-based Similarity Join .....	317
..... <i>Wenbo Tao, Minghe Yu, Guoliang Li</i>	

## VLDB 2015 ORGANIZATION AND REVIEW BOARD

### General Chairs

Michael J. Carey, University of California, Irvine

### Program Chairs and Editors-in-Chief of PVLDB 8

Chen Li, University of California, Irvine

Volker Markl, TU Berlin

### Research and Innovative Systems Tracks Associate Editors

Kevin Chang, U Illinois

Shivnath Babu, Duke University

Felix Naumann, Hasso Plattner Institute

Stefan Manegold, CWI Amsterdam

Yi Chen, NJIT

Fatma Ozcan, IBM Research Almaden

Jignesh Patel, University of Wisconsin, Madison

Magdalena Balazinska, University of Washington

### Experiments and Analysis Track Associate Editors

Rainer Gemulla, MPI Saarbrücken, Germany

### Industrial, Applications and Experience Track Associate Editors

Anhai Doan, University of Wisconsin, Madison

Prasan Roy, Sclera

Gregor Hackenbroich, SAP

### Demonstration Chair

Alfons Kemper, TU München

### Tutorial Chairs

Tova Milo, Tel Aviv University

Pierre Senellart, Telecom Paris Tech, France

### Panel Chair

Joseph M. Hellerstein, University of California, Berkeley

### PhD Workshop Chairs

Rachel Pottinger (UBC)

### Proceedings Chairs

Daisy Zhe Wang, University of Florida

Tyson Condie, University of California, Los Angeles

### Sponsorship Chairs

Michael Franklin, University of California, Berkeley

Edward Change, HTC

Patrick Valduriez, INRIA

### Local Organization Chair

Lipyeow Lim, University of Hawaii

### Conference and Registration Chairs

Ke Chen, Zhejiang University

Cuiping Li, Renmin University

### Publicity and Web Management Chair

Rada Chirkova, North Carolina State

### Treasury Chair

Malu Castellanos, HP Labs

### VLDB Endowment Liaison

Paul Larson, Microsoft Research

### PVLDB Managing Editor

Divesh Srivastava, AT&T Labs

### PVLDB Information Director

Gerald Weber, University of Auckland

### PVLDB Advisory Committee

Philip Bernstein, Michael Böhlen, Peter Buneman,  
Susan Davidson, Z. Meral Ozsoyoglu, S. Sudarshan,  
Gerhard Weikum

## Research Track Review Board

Daniel Abadi, Yale University  
Alberto Abello, UPC Barcelona  
Ashraf Aboulnaga, Qatar Computing Research Institute  
Foto Afrati, NTU Athens  
Sihem Amer Yahia, CNRS LIG  
Aijun An, York University  
Arvind Arasu, Microsoft Research  
Walid Aref, Purdue University  
Paolo Atzeni, Roma Tre University  
Denilson Barbosa, University of Alberta  
Srikanta Bedathur, IBM Research  
Philip Bernstein, Microsoft Research  
Michael Böhlen, University of Zürich  
Peter Boncz, CWI Amsterdam  
Angela Bonifati, Lille 1 U and INRIA  
Philippe Bonnet, IT U of Copenhagen  
Nico Bruno, Microsoft  
Alex Buchmann, TU Darmstadt  
Mike Cafarella, University of Michigan  
K. Selcuk Candan, Arizona State University  
Malu Castellanos, HP Labs  
Kaushik Chakrabarti, Microsoft Research  
Lei Chen, Hong Kong U of Science and Technology  
Fei Chiang, McMaster University  
Byron Choi, Hong Kong Baptist University  
Philippe Cudre Mauroux, Fribourg University  
Mahashweta Das, HP Labs  
Sudipto Das, Microsoft Research  
Amol Desphande, University of Maryland  
Stefan Dessloch, TU Kaiserslautern  
Jens Dittrich, Saarland University  
Alin Dobra, University of Florida  
Xin Luna Dong, Google  
Jennie Duggan, MIT  
Wenfei Fan, University of Edinburgh  
Alan Fekete, University of Sydney Australia

Peter Fischer, Universität Freiburg  
Avrilia Floratou, IBM  
Avigdor Gal, Technion  
Minos Garofalakis, Technical U of Crete  
Wolfgang Gatterbauer, Carnegie Mellon U  
Tingjian Ge, University of Massachusetts Lowell  
Floris Geerts, University of Antwerp  
Lukasz Golab, University of Waterloo  
Torsten Grust, Universität Tübingen  
Jarek Gryz, York University  
Dimitrios Gunopulos, University of Athens  
Hakan Hacigumus, NEC Labs  
Wook Shin Han, POSTECH  
Seif Haridi, KTH Stockholm  
Oktie Hassanzadeh, IBM Research  
Bingsheng He, Nanyang Technological University  
Jeffrey Heer, University of Washington  
Herodotos Herodotou, Microsoft Research  
Katja Hose, Aalborg University  
Vagelis Hristidis, University of California, Riverside  
Jeong-Hyon Hwang, State University of New York at Albany  
Stratos Idreos, Harvard University  
Yannis Ioannidis, University Of Athens  
Zachary Ives, University of Pennsylvania  
Christopher Jermaine, Rice University  
Ruoming Jin, Kent State University  
Alekh Jindal, Massachusetts Institute of Technology  
Ryan Johnson, University of Toronto  
Eser Kandogan, IBM Research Almaden  
Gjergji Kasneci, Hasso Plattner Institute  
Asterios Katsifodimos, TU Berlin  
Yannis Katsis, University of California, San Diego  
Daniel Keim, Universität Konstanz  
Bettina Kemme, McGill University  
Eamonn Keogh, University of California, Riverside  
Martin Kersten, CWI Amsterdam  
Daniel Kifer, Penn State University  
Hideaki Kimura, Hewlett Packard  
George Kollios, Boston University

Donald Kossman, ETH Zurich  
Nick Koudas, University of Toronto  
Georgia Koutrika, HP Labs  
Tim Kraska, Brown University  
Harumi Kuno, HP Labs  
Laks Lakshmanan, University of British Columbia  
Paul Larson, Microsoft Research  
Hongrae Lee, Google  
Wolfgang Lehner, TU Dresden  
Alberto Lerner, New York University  
Ulf Leser, Humboldt Universität zu Berlin  
Justin Levandoski, Microsoft Research  
Feifei Li, University of Utah  
Guoliang Li, Tsinghua University  
Jianzhong Li, Harbin Institute of Technology  
Yunyao Li, IBM Research Almaden  
Erietta Liarou, EPF Lausanne  
Xuemin Lin, University of New South Wales  
Ziyang Liu, NEC Labs America  
Eric Lo, Polytecnic University of Hong Kong  
Guy Lohman, IBM Research Almaden  
Jiaheng Lu, Renmin University of China  
Qiong Luo, Hong Kong University of Science and  
Technology  
Jayant Madhavan, Google  
Ioana Manolescu, INRIA  
Patrick Marcel, University of Tours  
Marta Mattoso, Federal University of Rio de Janeiro  
Alexandra Meliou, University of Massachusetts  
Amherst  
Sergey Melnik, Google  
Weiyi Meng, Binghamton University  
Sebastian Michel, Saarland University  
Iris Miliaraki, Yahoo Labs Barcelona  
Renee Miller, University of Toronto  
Zhou Minqi, East China Normal University  
Prasenjit Mitra, Penn State University  
Bernhard Mitschang, Universität Stuttgart  
Mohamed Mokbel, Northeastern University  
Barzan Mozafari, University of Michigan

Hannes Mühleisen, CWI Amsterdam  
Arnab Nandi, Ohio State University  
Vivek Narasayya, Microsoft Research  
Jeffrey Naughton, University of Wisconsin Madison  
Rimma Nehme, Microsoft  
Thomas Neumann, TU Munich  
Raymond Ng, University of British Columbia  
Christopher Olston, Google  
Dan Olteanu, Oxford University & LogicBlox  
Beng Chin Ooi, National University of Singapore  
M. Tamer Özsu, University of Waterloo  
Themis Palpanas, Paris Descartes University  
Ippokratis Pandis, IBM Research Almaden  
Dimitris Papadias, Hong Kong U of Science and  
Technology  
Paolo Papotti, Qatar Computing Research Institute  
Andy Pavlo, Carnegie Mellon University  
Torben Bach Pedersen, Aalborg University  
Jian Pei, Simon Fraser University  
Peter Pietzuch, Imperial College London  
Evaggelia Pitoura, University of Ioannina  
Alkis Polyzotis, Google and University of California,  
Santa Cruz  
Fabio Porto, LNCC Brazil  
Li Qian, Facebook  
Jorge Arnulfo Quiane Ruiz, Qatar Computing Research  
Institute  
Tilman Rabl, University of Toronto  
Erhard Rahm, Universität Leipzig  
Krithi Ramamritham, IIT Bombay  
Ravi Ramamurthy, Microsoft Research  
Vijayshankar Raman, IBM Research Almaden  
Mirek Riedewald, Northeastern University  
Tore Risch, Uppsala Universitet  
Kenneth Ross, Columbia University  
Elke Rundensteiner, Worcester Polytechnic Institute  
US  
Barna Saha, AT&T Labs Research  
Kenneth Salem, University of Waterloo  
Simonas Saltenis, Aalborg University

Kai Uwe Sattler, TU Ilmenau  
Eric Sedlar, Oracle Labs  
Bernhard Seeger, Philipps Universität Marburg  
Kyuseok Shim, Seoul National University  
Jerome Simeon, IBM Research Watson  
Alkis Simitsis, HP Labs  
Michael Stonebraker, Massachusetts Institute of  
Technology  
Julia Stoyanovich, Drexel University  
Nan Tang, Qatar Computing Research Institute  
Yufei Tao, The Chinese University of Hong Kong  
Sandeep Tata, Google  
Nesime Tatbul, Intel Labs and MIT  
Arash Termehchy, Oregon State University  
Evimaria Terzi, Boston University  
Jens Teubner, TU Dortmund  
Martin Theobald, University of Antwerp  
Andreas Thor, University of Leipzig  
Yuanyuan Tian, IBM Research Almaden  
Anthony Tung, National University of Singapore  
Vasilis Vassalos, Athens University  
Yannis Vassiliou, NTU Athens  
Yannis Velegarakis, University of Trento

Rares Vernica, HP Labs  
Gottfried Vossen, WWU Münster  
Florian Waas, Datometry Inc.  
Daisy Zhe Wang, University of Florida  
Haixun Wang, Google  
Wei Wang, University of New South Wales  
Gerhard Weikum, Max Planck Institut für Informatik  
Till Westmann, Oracle Labs  
Steven Whang, Google  
Kyu Young Whang, KAIST  
Raymond Chi Wing Wong, Hong Kong University of  
Science and Technology  
Xiaokui Xiao, Nanyang Technological University  
Xifeng Yan, University of California, Santa Barbara  
Jun Yang, Duke University  
Xiaochun Yang, Northeastern University China  
Cong Yu, Google Research  
Ge Yu, Northeastern University China  
Jeffrey Xu Yu, Chinese University of Hong Kong  
Xiaofang Zhou, The University of Queensland  
Esteban Zimanyi, UL Brussels  
Marcin Zukowski, Snowflake Computing

## LETTER FROM THE MANAGING EDITOR AND ADVISORY COMMITTEE

To err is human, to allow for corrections is the new PVLDB policy.

Every one of us who has written programs (or got them written by students or employees) has surely found bugs in code. Some of these bugs result in errors in results presented in papers, which we discover later. Even theoreticians are not immune to errors, and several of us have found errors in our own published papers or in papers published by others. Based on the personal experience of one of the two authors of this article, easily a few percent of published papers have one or more such issues that are detected only after publication. Till now PVLDB did not have a formal way to correct mistakes.

We now have a new policy on errata. PVLDB will henceforth accept Errata Notes, which are basically very short (1 or 2 page) papers reporting errors in papers published earlier in PVLDB/VLDB. We expect many of these will come from the authors of the original paper when they find errors, but some will come from others who have found bugs in earlier papers. Errata Notes will be subjected to a basic level of reviewing if submitted by original authors, and more detailed reviewing plus feedback/rebuttal from original authors if submitted by someone other than the original authors. Errata Notes will also be explicitly identified by including the keywords "Errata for" in the title.

In case of serious errors that severely compromise the results in a paper, we will also allow retraction of an already published paper. We do not expect to see very many retractions, and hope to see none at all, but are working on policies to handle such cases if and when they arise. Typically we expect a retraction to be either initiated by or approved by the original authors, but in case of issues such as plagiarism/self-plagiarism, the PVLDB editors-in-chief may unilaterally make the decision to retract an already published paper. Policies for how to handle retractions are being finalized, but in case a paper is retracted, PVLDB will replace the paper online by a page that explains why the paper was retracted and by whom.

---

Divesh Srivastava (PVLDB Managing Editor)  
S. Sudarshan (Member, PVLDB Advisory Committee)