

# VLDB2014

40<sup>th</sup> International Conference on Very Large Data Bases, Hangzhou, China



## Proceedings of the VLDB Endowment

Volume 7, No. 7 – March 2014

**Proceedings of the 40th International Conference on  
Very Large Data Bases, Hangzhou, China**

Program Chairs and Editors-in-Chief:

**H. V. Jagadish, Aoying Zhou**

Associate Editors – Research and Innovative Systems Tracks:

**Shivnath Babu, Lei Chen, Graham Cormode, Bin Cui, Wynne Hsu, Martin Kersten,  
Donald Kossman, Elke Rundensteiner, Kyuseok Shim, Wang-Chiew Tan, Letizia Tanca, Jeffrey Yu**

Associate Editors – Experiments and Analysis Track:

**Gao Cong, Jens Dittrich**

Associate Editors – Vision Track:

**Zachary Ives**

Proceedings Chairs:

**Li Xiong, Cong Yu**

PVLDB – Proceedings of the VLDB Endowment

Volume 7, No. 7, March 2014.

The 40th International Conference on Very Large Data Bases, Hangzhou, China.

## **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 7, Number 7, March 2014: VLDB 2014

Pages ii - ix and 493 - 600

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 2014 Organization and Review Board .....	iv

### Letters

Letter from the Program Chairs of VLDB 2015.....	<i>Chen Li, Volker Markl</i>	ix
--	------------------------------	----

### Research Papers

Rank Join Queries in NoSQL Databases.....		493
.....	<i>Nikos Ntarmos, Ioannis Patlakas, Peter Triantafillou</i>	
Biperpedia: An Ontology for Search Applications .....		505
.....	<i>Rahul Gupta, Alon Halevy, Xuezhong Wang, Steven Euijong Whang, Fei Wu</i>	
GRAMI: Frequent Subgraph and Pattern Mining in a Single Large Graph .....		517
.....	<i>Mohammed Elseidy, Ehab Abdelhamid, Spiros Skiadopoulos, Panos Kalnis</i>	
Lightweight Indexing of Observational Data in Log-Structured Storage .....		529
.....	<i>Sheng Wang, David Maier, Beng Chin Ooi</i>	
epiC: an Extensible and Scalable System for Processing Big Data .....		541
.....	<i>Dawei Jiang, Gang Chen, Beng Chin Ooi, Kian-Lee Tan, Sai Wu</i>	
Hybrid Parallelization Strategies for Large-Scale Machine Learning in SystemML.....		553
.....	<i>Matthias Boehm, Shirish Tatikonda, Berthold Reinwald, Prithviraj Sen,</i> <i>Yuanyuan Tian, Douglas R. Burdick, Shivakumar Vaithyanathan</i>	
Schemaless and Structureless Graph Querying .....		565
.....	<i>Shengqi Yang, Yinghui Wu, Huan Sun, Xifeng Yan</i>	
Optimizing Graph Algorithms on Pregel-like Systems .....		577
.....	<i>Semih Salihoglu, Jennifer Widom</i>	
Toward Computational Fact-Checking .....		589
.....	<i>You Wu, Pankaj K. Agarwal, Chengkai Li, Jun Yang, Cong Yu</i>	

## **VLDB 2014 ORGANIZATION AND REVIEW BOARD**

### **Honorary Chair**

Yunhe Pan, Chinese Academy of Engineering

### **General Chairs**

Chun Chen, Zhejiang University

Sharad Mehrotra, University of California, Irvine

### **Program Chairs and Editors-in-Chief of PVLDB 7**

H. V. Jagadish, University of Michigan

Aoying Zhou, East Normal University, China

### **Research and Innovative Systems Tracks Associate Editors**

Shivnath Babu, Duke University

Lei Chen, Hong Kong University of Science and Technology

Graham Cormode, University of Warwick

Bin Cui, Peking University, China

Wynne Hsu, NUS

Martin Kersten, CWI

Donald Kossman, ETH

Elke Rundensteiner, WPI

Kyuseok Shim, Seoul National University

Wang-Chiew Tan, University of California, Santa Cruz

Letizia Tanca, Poli Milano

Jeffrey Yu, Chinese University of Hong Kong

### **Experiments and Analysis Track Associate Editors**

Gao Cong, Nanyang Technology University

Jens Dittrich, Saarland

### **Visions Track Associate Editor**

Zachary Ives, University of Pennsylvania

### **Industrial and Applications Track Associate Editors**

Umeshwar Dayal, HP

C. Mohan, IBM

Ge Yu, Northeastern University, China

### **Demonstration Chairs**

Mong-Li Lee, NUS  
Feifei Li, University of Utah  
Sunil Prabhakar, Purdue

### **Tutorial Chairs**

Xiaoyong Du, Renmin University  
Murat Kantarcioglu, University of Texas, Dallas  
Divesh Srivastava, AT&T Labs

### **Research Track Review Board**

Sibel Adali, Rensselaer Polytechnic Institute  
Foto Afrati, NTU Athens  
Yanif Ahmad, JHU  
Jose Luis Ambite, ISI - USC  
Walid Aref, Purdue University  
Claudia Bauzer Medeiros, University of Campinas  
Srikanta Bedathur, IIIT Delhi  
Michael Benedikt, Oxford University  
Sonia Bergamaschi, Universita Modena  
Laure Berti-Equille, IRD, France  
Leopoldo Bertossi, Carleton University, Ottawa  
Subhash Bhalla, University of Aizu, Japan  
Peter Boncz, CWI  
Angela Bonifati, University of Lille 1  
Rajesh Bordawekar, IBM Watson Research Center  
Omar Boucelma, Aix-Marseille University  
Nico Bruno, Microsoft Research  
Andrea Cali, University of London, Birkbeck College  
Malu Castellanos, HP Labs  
Badrish Chandramouli, Microsoft Research  
Adriane Chapman, MIT  
Gang Chen, Zhejiang University  
Yi Chen, New Jersey Institute of Technology  
James Cheng, CUHK  
Reynold Cheng, University of Hong Kong  
Brian Cooper, Google, USA

### **Workshop Chairs**

Anastasia Ailamaki, EPFL  
Kaushik Chakrabarti, Microsoft

### **Panel Chairs**

Hakan Hacigumus, NEC Labs  
Jignesh Patel, University of Wisconsin  
Xiaoyang Sean Wang, Fudan University

Philippe Cudré-Mauroux, University of Fribourg  
Carlo Curino, MIT  
Gautam Das, UT Arlington and QCRI  
Sudipto Das, Microsoft Research  
Anish Das Sarma  
Atish Das Sarma, eBay Research Labs  
Khuzaima Daudjee, University of Waterloo  
Antonios Deligiannakis, Technical University of Crete  
Daniel Deutch, Ben Gurion University  
Yanlei Diao, University of Massachusetts Amherst  
Xin (Luna) Dong, Google, USA  
Sameh Elnikety, Microsoft Research  
Mohamed Eltabakh, Worcester Polytechnic Institute  
Ihab F. Ilyas, QCRI  
Hakan Ferhatosmanoglu, Bilkent University  
Ada Wai-Chee Fu, Chinese University of Hong Kong  
Minos Garofalakis, Technical University of Crete  
Wolfgang Gatterbauer, Carnegie Mellon University  
Tingjian Ge, University of Massachusetts, Lowell  
Buğra Gedik, Bilkent University  
Rainer Gemulla, Max-Planck-Institut Saarbrücken  
Gabriel Ghinita, University of Massachusetts Boston  
Parke Godfrey, York University  
Lukasz Golab, University of Waterloo  
Sergio Greco, University of Calabria  
Le Gruenwald, University of Oklahoma

Giovanna Guerrini, Universita Genova  
 Krishna Gummadi, MPI-SWS  
 Rahul Gupta, Google Research  
 Rajeev Gupta, IBM Research  
 Shyam Gupta, IIT Delhi  
 Marios Hadjieleftheriou, AT&T labs  
 Wook-Shin Han, KNU, Korea  
 Kuno Harumi, HP Labs  
 Bingsheng He, NTU Singapore  
 Sven Helmer, Free University of Bozen-Bolzano  
 Jan Hidders, TUDelft  
 Wei Hong, Cisco System Inc.  
 Katja Hose, Aalborg University  
 Zi Huang, University of Queensland  
 Jeong-Hyon Hwang, SUNY - Albany  
 Seung-won Hwang, POSTECH, Korea  
 Stratos Idreos, CWI  
 Yoshiharu Ishikawa, Nagoya University  
 Zachary Ives, University of Pennsylvania  
 Ricardo Jimenez-Peris, Technical University of Madrid  
 Cheqing Jin, East China Normal University  
 Ruoming Jin, Kent State University  
 Alekh Jindal, Saarland University/MIT  
 Ryan Johnson, University of Toronto  
 Dmitri V Kalashnikov, UC Irvine  
 Panos Kalnis, KAUST, Saudi Arabia  
 Ben Kao, Hong Kong University  
 Panagiotis Karras, Rutgers University  
 Yiping Ke, Institute of High Performance Computing  
 Bettina Kemme, McGill University  
 Daniel Kifer, PSU  
 Benny Kimelfeld, IBM  
 Hideaki Kimura, Microsoft Jim Gray Systems Lab  
 George Kollios, Boston University  
 Christian König, Microsoft Research  
 Tim Kraska, Brown University  
 Laks V. S. Lakshmanan, University of British Columbia  
 Mounia Lalmas, Yahoo Inc.  
 Mong-Li Lee, National University of Singapore  
 Wolfgang Lehner, Technische University Dresden  
 Justin Levandoski, Microsoft Research  
 Chengkai Li, The University of Texas at Arlington  
 Cuiping Li, Renmin University of China  
 Feifei Li, University of Utah  
 Guoliang Li, Tsinghua University  
 Jianzhong Li, Harbin Institute of Technology  
 Yunyao Li, IBM Almaden  
 Zhanhuai Li, Northwestern Polytechnical University  
 Dan Lin, Missouri S&T, USA  
 Xuemin Lin, University of New South Wales  
 Bin Liu, NEC Labs America  
 Ziyang Liu, NEC Labs America  
 Eric Lo, The Hong Kong Polytechnic University  
 Qiong Luo, HKUST  
 Shuai Ma, Beihang University  
 Ashwin Machanavajjhala, Duke University  
 Brad Malin, Duke University  
 Nikos Mamoulis, University of Hong Kong  
 Stefan Manegold, CWI  
 Murali Mani, University of Michigan  
 Ioana Manolescu, INRIA, France  
 Amélie Marian, Rutgers University  
 Volker Markl, TU Berlin  
 Marta Mattoso, Federal University of Rio de Janeiro  
 Frank McSherry, Microsoft  
 Alexandra Meliou, Umass Amherst  
 Marco Mesiti, University of Milano  
 Dan Miranker, The University of Texas at Austin  
 Mohamed Mokbel, University of Minnesota  
 Bongki Moon, Seoul National University  
 Yasuhiko Morimoto, Hiroshima University  
 Mirella Moro, Universidade Federal de Minas Gerais  
 Kyriakos Mouratidis, SMU, Singapore  
 Karin Murthy, IBM India  
 Arnab Nandi, Ohio State University  
 Wolfgang Nejdl, University of Hannover

Thomas Neumann, Technology University Munchen  
Boris Novikov, St Petersburg University  
Dan Olteanu, Oxford University  
Gultekin Ozsoyoglu, Case Western Reserve University  
Tamer Ozsu, University of Waterloo  
Esther Pacitti, University of Montpellier  
Themis Palpanas, University of Trento  
Ippokratis Pandis, IBM Almaden  
Stelios Pappas, Microsoft Research  
Aditya Parameswaran, Stanford University  
Srinivasan Parthasarathy, The Ohio State University  
Jignesh Patel, University of Wisconsin  
Andrew Pavlo, Brown University  
Peter Pietzuch, Imperial College London  
Neoklis Polyzotis, University of California - Santa Cruz  
Cecilia M. Procopiuc, AT&T Labs  
Li Qian, University of Michigan  
Jorge Quiané-Ruiz, QCRI  
Elisa Quintarelli, Politecnico di Milano  
Maya Ramanath, IIT Delhi  
Louiqa Raschid, University of Maryland  
Vibhar Rastogi, Yahoo  
Matthias Renz, University of Munich  
Kenneth Ross, Columbia University  
Sourav S Bhowmick, NTU, Singapore  
Dimitris Sacharidis, IMIS Athena, Greece  
Kenneth Salem, University of Waterloo  
Maria Sapino, University of Torino  
Kai-Uwe Sattler, TU Ilmenau  
Monica Scannapieco, ISTAT  
Bernhard Seeger, University of Marburg  
Lidan Shou, Zhejiang University  
Adam Silberstein, Trifacta  
Lisa Singh, Georgetown University  
Radu Sion, Stony Brook University  
Yufei Tao, Chinese University of Hong Kong

Nesime Tatbul, ETH Zurich  
Arash Termehchy, Oregon State University  
Evimaria Terzi, University of Boston  
Martin Theobald, Max Planck Institute, Germany  
Srikanta Tirthapura, Iowa State University  
Riccardo Torlone, Roma Tre University  
Anthony Tung, National University of Singapore  
Kostas Tzoumas, Technical University of Berlin  
Sergei Vassilvitskii, Google Research  
Marcos Vaz Salles, University of Copenhagen (DIKU)  
Stratis Viglas, University of Edinburgh  
Hoang Tam Vo, National University of Singapore  
Daisy Zhe Wang, University of Florida  
Haixun Wang, Microsoft Research Asia  
Ke Wang, Simon Fraser University  
Wei Wang, University of New South Wales  
Xiaoling Wang, East China Normal University  
Ingmar Weber, Yahoo  
Raymond Chi Wing Wong, HKUST  
Sai Wu, Zhejiang University  
Yuqing Wu, Indiana University  
Xiaokui Xiao, NTU  
Dong Xin, Google  
Jianliang Xu, Hong Kong Baptist University  
Jun (Jim) Xu, Georgia Institute of Technology  
Xifeng Yan, University of Santa Barbara  
Xiaoyan Yang, Advanced Digital Science Center  
Ke Yi, HKUST  
Ge Yu, Northeastern University, China  
Hwanjo Yu, POSTECH, Korea  
Meihui Zhang, National University of Singapore  
Wenjie Zhang, The University of New South Wales  
Ying Zhang, The University of New South Wales  
Zhenjie Zhang, Advanced Digital Science Center  
Wenzhao Zhou, Georgetown University  
Xiaofang Zhou, University of Queensland

**PhD Workshop Chairs**

Erich Neuhold, University of Vienna  
Yun Yao Li, IBM

**Sponsorship Chairs**

Mike Carey, University of California, Irvine  
Lizhu Zhou, Tsinghua University

**Local Organization Chair**

Lidan Shou, Zhejiang University

**Web Management Chair**

Sai Wu, Zhejiang University

**Conference and Registration Chairs**

Ke Chen, Zhejiang University  
Cuiping Li, Renmin University

**Publicity Chairs**

Vasilis Vassalos, AUEB, Greece  
Dunlu Peng, USST, China

**Proceedings Chairs**

Li Xiong, Emory University  
Cong Yu, Google Research

**Treasury Chair**

Li (Eric) Qian, University of Michigan

**VLDB Endowment Liaison**

Kyu-Young Whang, KAIST

**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

**Logo Design**

Guanmin Guo



## LETTER FROM THE PROGRAM CHAIRS OF VLDB 2015

As the program chairs of VLDB 2015 and Editors in Chief of the upcoming Volume 8 of PVLDB, we are looking forward to your contributions to shape the next year of database research. VLDB 2015 will continue the well-established process of monthly submissions using Microsoft's conference management tool. We have assembled a program committee of world-class researchers, who will evaluate your contributions. Meanwhile, we have made minor changes to the paper-submission process for PVLDB Volume 8, which may be of interest to authors:

- We changed the monthly submission deadline to the last day of the month at 23:59 US Pacific Time (instead of the previous deadline, which was the 1st day of each month at 5 AM pacific time). One reason is that we want the monthly deadline to be within that month, not the first day of the next month. We also wish people working on the west coast can go to sleep earlier :-). We will open CMT for submissions of papers to PVLDB Volume 8 in March 2014. The last submission date for consideration by the 2015 committee will be February 28, 2015. Note that the cutoff date for presentation at VLDB 2015 is full acceptance by June 1st, 2015, irrespective of submission date.
- We dropped the "Vision" track, as we believe that all submissions should have strong innovative and visionary aspects. We therefore did not see the need to specifically single out vision papers in the proceedings or the conference.
- We included the "Innovative Systems Track" and "Experiments and Analysis (E&A) Track" into the Research Track Call for Papers, in order to make clear that both "innovative systems papers" and "experiments and analysis papers" are full research papers in their own right. However, we strongly endorse both of these paper types and continue to maintain the special features and reviewing requirements of these papers, which are described on the Web site of VLDB 2015, following the corresponding links on the Call for Papers page.

We look forward towards to your submissions and creating an exciting program for VLDB 2015. (Hey, it will be in Hawaii!)

---

Chen Li, UC Irvine and SRCH2, USA  
Volker Markl, TU Berlin, Germany  
Program Chairs, VLDB 2015