

# VLDB2014

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



## Proceedings of the VLDB Endowment

Volume 7, No. 6 – February 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. 6, February 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 6, February 2014: VLDB 2014

Pages ii - ix and 429 - 492

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

Planet of the APEs.....	<i>Graham Cormode</i>	ix
-------------------------	-----------------------	----

### Research Papers

Shared Workload Optimization .....	429
.....	<i>Georgios Giannikis, Darko Makreshanski, Gustavo Alonso, Donald Kossmann</i>
Scalable and Adaptive Online Joins.....	441
.....	<i>Mohammed Elseidy, Abdallah Elguindy, Aleksandar Vitorovic, Christoph Koch</i>
Support the Data Enthusiast: Challenges for Next-Generation Data-Analysis Systems.....	453
.....	<i>Kristi Morton, Magdalena Balazinska, Dan Grossman, Jock Mackinlay</i>
A Provenance Framework for Data-Dependent Process Analysis.....	457
.....	<i>Daniel Deutch, Yuval Moskovitch, Val Tannen</i>
Tracking Entities in the Dynamic World: A Fast Algorithm for Matching Temporal Records.....	469
.....	<i>Yueh-Hsuan Chiang, AnHai Doan, Jeffrey F. Naughton</i>
Edelweiss: Automatic Storage Reclamation for Distributed Programming .....	481
.....	<i>Neil Conway, Peter Alvaro, Emily Andrews, Joseph M. Hellerstein</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, Mitre  
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 Paparizos, 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, Univesity 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



## PLANET OF THE APES

As an experimental hybrid between a conference and a journal, VLDB employs associate editors who operate differently from similar roles in more traditional venues. The VLDB associate editor is akin to an associate editor of a journal, and a meta-reviewer, or area chair in a conference, but different from each. The exact actions have varied over time, as the PVLDB workflow is fine-tuned, and different structures have been tried. The main tasks of the Associate PVLDB Editor (or APE, for short), are to assist in the allocation of reviewers to papers, and to oversee the path from reviews to decision.

Finding a good set of reviewers for a paper means choosing reviewers who are expert on the subject matter, or have a good understanding of related topics, so that they can comment on the novelty and importance of the submission. The compressed timescale of PVLDB denies us the luxury of inviting individual experts, or allowing a large PC the time to bid for papers. The current workflow includes a very small time window for the VLDB PC members to enter preferences (in fact a majority do, up to 75%).

An initial allocation of papers to reviewers is then made by the conference management system, based on the preferences, and the categories that have been used to tag the papers and reviewers. Matching based on categories is quite haphazard, due to the subjectivity of labeling people and papers. The APE, who is assigned a set of papers to oversee, can manually modify this assignment, under some constraints, such as a limit on the number of papers to assign any reviewer each month (typically two). Not every PC member is "on call" each month, so there is sometimes a paper with an "ideal" reviewer from the PC list who cannot be assigned. Due to these constraints, authors who strategize over the reviewer list should realize that there is no guarantee that their paper will reach the reviewers they think are best suited to it.

After reviews are in at the end of the month, the APE looks over them. In some cases, the decision is easy -- three accepts or three rejects will inevitably yield the consensus outcome. When the reviews are mixed, the APE tries to stimulate the discussion, much as in typical database conference, but on a much shorter timescale. For the most part, the APE acts as a facilitator of the discussion, trying to help identify on what points there is agreement or dissent, and trying to reach a consensus discussion. The APE is not a "fourth reviewer", and does not read the paper in as much detail as the assigned reviewers. The option of inviting a revision is always on the table, but with the warning that this is not to be used in order to "kick the can down the road": APEs are encouraged to reach firm decisions in the first round if possible. There is still a range of opinions among PVLDB participants over whether to interpret the revision option as equivalent to a minor revision or major revision in the journal context. In my view, it is a different beast. If a revision is requested, the reviews should indicate the main points that need to be addressed to allow acceptance, but then it is up to the authors to decide if they wish to submit a revision within the (again, quite tight) time allowed. Due to the high standards of PVLDB, revision can be a minority outcome: in my sample for the current iteration, about 20-25% of papers received a revision response, while close to 75% were rejected in the first round (acceptances of first submissions are rare).

The last step for the APE is to enter a decision, and prepare a "meta-review" (terminology inherited from the Conference Management Tool used to run the process). The meta-review is usually short, and allows the APE to identify which points from reviews were most significant in reaching the decision, or to highlight where to focus attention if a revision is requested. It can also summarize points raised in the (private) discussion phase that may not have been mentioned in the reviews.

---

Graham Cormode, University of Warwick  
Research Track Associated Editor, VLDB 2014