

# VLDB2017

43<sup>rd</sup> International Conference on Very Large Data Bases, Munich, Germany



## Proceedings of the VLDB Endowment

Volume 10, No. 4 – December 2016

**Proceedings of the 43rd International Conference on  
Very Large Data Bases, Munich, Germany**

Program Chairs:

**Peter Boncz and Ken Salem**

Associate Editors – Research Track:

**Ashraf Aboulnaga, Shimin Chen, Gautam Das, Amol Deshpande, Zack Ives, Qiong Luo, Stefan Manegold, Ioana Manolescu, Sharad Mehrotra, Fatma Ozcan, Themis Palpanas, Rachel Pottinger, Ken Ross, Gerhard Weikum**

Proceedings Chairs:

**Alvin Cheung, Aaron Elmore**

PVLDB – Proceedings of the VLDB Endowment

Volume 10, No. 4, December 2016.

The 43rd International Conference on Very Large Data Bases, Munich, Germany.

## **Copyright 2016 VLDB Endowment**

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org).

Volume 10, Number 4, December 2016: VLDB 2017

Pages i – vi and 289 – 468

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 2016 Organization and Review Board .....	iv

### Research Papers

Finding Persistent Items in Data Streams .....	
.....Haipeng Dai, Muhammad Shahzad, Alex X. Liu, Yuankun Zhong	289
BlueCache: A Scalable Distributed Flash-based Key-value Store .....	
.....Shuotao Xu, Sungjin Lee, Sang-Woo Jun, Ming Liu, Jamey Hicks, Arvind	301
A General and Parallel Platform for Mining Co-Movement Patterns over Large-scale Trajectories .....	
.....Qi Fan, Dongxiang Zhang, Huayu Wu, Kian-Lee Tan	313
VIP-Tree: An Effective Index for Indoor Spatial Queries .....	
.....Zhou Shao, Muhammad Cheema, David Taniar, Hua Lu	325
Write-Behind Logging .....	
.....Joy Arulraj, Matthew Perron, Andrew Pavlo	337
The TileDB Array Data Storage Manager .....	
.....Stavros Papadopoulos, Kushal Datta, Samuel Madden, Timothy Mattson	349
DOCS: Domain-Aware Crowdsourcing System .....	
.....Yudian Zheng, Guoliang Li, Reynold Cheng	361
Lifting the Haze off the Cloud: A Consumer-Centric Market for Database Computation in the Cloud .....	
.....Yue Wang, Alexandra Meliou, Gerome Miklau	373
Two Birds, One Stone: A Fast, yet Lightweight, Indexing Scheme for Modern Database Systems .....	
.....Jia Yu, Mohamed Sarwat	385
History is a mirror to the future: Best-effort approximate complex event matching with insufficient resources .....	
.....Zheng Li, Tingjian Ge	397
PHYTM: Persistent Hybrid Transactional Memory .....	
.....Hillel Avni, Trevor Brown	409
Skipping-oriented Partitioning for Columnar Layouts .....	
.....Liwen Sun, Michael Franklin, Jiannan Wang, Eugene Wu	421
Estimating Quantiles from the Union of Historical and Streaming Data .....	
.....Sneha Singh, Divesh Srivastava, Srikanta Tirthapura	433
Clay: Fine-Grained Adaptive Partitioning for General Database Schemas .....	
Marco Serafini, Rebecca Taft, Aaron Elmore, Andrew Pavlo, Ashraf Aboulnaga, Michael Stonebraker	445

Effortless Data Exploration with zenvisage: An Expressive and Interactive Visual Analytics System  
.....Tarique Ashraf Siddiqui, Albert Kim, John Lee, Karrie Karahalios, Aditya Parameswaran 457

## VLDB 2017 ORGANIZATION AND REVIEW BOARD

### General Chairs

Alfons Kemper, TUM  
Thomas Neumann, TUM

### Honorary Chair

Johann-Christoph Freytag, HU Berlin

### Organization Committee Chair

Stephan Günnemann, TUM  
Alfons Kemper, TUM  
Thomas Neumann, TUM

### Program Chairs and Editors in Chief of PVLDB 10

Peter Boncz, CWI  
Ken Salem, University of Waterloo

### Associate Editors of PVLDB 10

Ashraf Aboulnaga, Qatar Computing Research Institute  
Shimin Chen, Chinese Academy of Sciences  
Gautam Das, University of Texas at Arlington  
Amol Deshpande, University of Maryland  
Zack Ives, University of Pennsylvania  
Qiong Luo, HKUST  
Stefan Manegold, CWI  
Ioana Manolescu, INRIA  
Sharad Mehrotra, UC Irvine  
Fatma Ozcan, IBM Research  
Themis Palpanas, Paris Descartes University  
Rachel Pottinger, University of British Columbia  
Ken Ross, Columbia University  
Gerhard Weikum, MPI

### VLDB Endowment Representative

Volker Markl, TU Berlin

### Sponsorship Committee Chairs

Mike Carey, UC Irvine  
Sang Kyun Cha, University of Seoul  
Wolfgang Lehner, TU Dresden

### Publicity Committee Chair

Jens Dittrich, Saarland University

### Tutorial Chairs

Vagelis Hristidis, UC Riverside  
Aristides Gionis, Aalto University

### Industrial Chairs

Felix Naumann, HPI  
Jonathan Goldstein, Microsoft Research  
Jingren Zhou, Alibaba

### Demonstration Chairs

Martin Theobald, Ulm University  
Bingsheng He, NUS  
Reynold Xin, Databricks

### Panel Chairs

Stratos Idreos, Harvard  
Michael Brodie, MIT

### Workshop Chairs

Christian Jensen, Aalborg University

### PhD Workshop Chairs

Erhard Rahm, University of Leipzig  
Peter Christen, ANU  
Bettina Kemme, McGill University

### Proceedings Chairs

Aaron Elmore, University of Chicago  
Alvin Cheung, University of Washington

### Website Chair

Felix Martin Schuhknecht, Saarland University

### PVLDB Managing Editor

Divesh Srivastava, AT&T Labs

### PVLDB Information Director

Gerald Weber, University of Auckland

### PVLDB Advisory Committee

H.V. Jagadish, Tan Kian Lee, Renee Miller, S. Sudarshan,  
Juliana Freire, Tamer Ozsu, Chen Li, Wolfgang Lehner

## Research Track Review Board

Alan Fekete, Sydney University  
Alekh Jindal, Microsoft  
Alexander Löser, Beuth University of Applied Sciences Berlin  
Alexandros Labrinidis, University of Pittsburgh  
Allison Holloway, Oracle  
Angela Bonifati, Université Lyon 1  
Ansgar Scherp, Kiel University  
Anthony Tung, NUS Singapore  
Aris Anagnostopoulos, Sapienza University of Rome  
Arnab Nandi, Ohio State University  
Arvind Arasu, Microsoft Research  
Asterios Katsifodimos, TU Berlin  
Atsuyuki Morishima, University of Tsukuba  
Avrilia Floratou, IBM Research Almaden  
Azza Abouzied, NYU Abu Dhabi  
Barzan Mozafari, University of Michigan  
Bernhard Seeger, University of Marburg  
Berthold Reinwald, IBM Research Almaden  
Bin Cui, Peking University  
Bingsheng He, Nanyang Technological University  
Bolin Ding, Microsoft Research  
Bongki Moon, SNU  
Boris Glavic, Illinois Institute of Technology  
Carmem Hara, Universidade Federal do Parana  
Chee-Yong Chan, National University of Singapore  
Chengkai Li, University of Texas at Arlington  
Chi Wang, Microsoft Research  
Chris Jermaine, Rice University  
Christian König, Microsoft Research  
Christina Lioma, Copenhagen University  
Cong Yu, Google  
Curtis Dyreson, Utah State University  
Cyrus Shahabi, University of Southern California  
Daisy Zhe, Wang, University of Florida  
Dan Olteanu, University of Oxford  
Daniel de Oliveira, Universidade Federal Fluminense  
David Koop, University of Massachusetts  
Davide Mottin, HPI  
Dmitri Kalashnikov, AT&T Labs Research  
Eli Cortez, Microsoft  
Elisa Bertino, Purdue University  
Eric Lo, Hong Kong Polytechnic University  
Essam Mansour, Qatar Computing Research Institute  
Eugene Wu, Columbia University  
Fabrizio Silvestri, Yahoo Research London  
Fei Chiang, McMaster University  
Feifei Li, University of Utah  
Florent Masseglia, INRIA  
Florian Kerschbaum, SAP  
George Papadakis, University of Athens  
Georgia Koutrika, Hewlett Packard Labs  
Giansalvatore Mecca, University Basilicata  
Goetz Graefe, Hewlett Packard Labs  
Guoliang Li, Tsinghua University  
Hakan Ferhatosmanoglu, Bilkent University  
Hannes Voigt, TU Dresden

Hannes Mühleisen, CWI  
Harumi Kuno, Hewlett Packard Labs  
Henrik Muehe, Google  
Holger Pirk, MIT  
Huy Vo, CUNY-CCNY  
Ihab Ilyas, University of Waterloo  
Indrakshi Ray, Colorado State University  
Ingmar Weber, Qatar Computing Research Institute  
Ippokratis Pandis, Amazon Web Services  
Ira Assent, Aarhus University  
Jaewoo Kang, Korea University  
James Cheng, Chinese University of Hong Kong  
Jeff Pound, SAP  
Jeffrey Yu, Chinese University of Hong Kong  
Jennie Duggan, Northwestern University  
Jens Teubner, TU Dortmund  
Jiaheng Lu, University of Helsinki  
Jianliang Xu, Hong Kong Baptist University  
Jignesh Patel, University of Wisconsin  
Johann Gamper, Free University of Bozen-Bolzano  
Joseph Gonzalez, UC Berkeley  
Julia Stoyanovich, Drexel University  
Julien Leblay, AIST, Japan  
Kai-Uwe Sattler, TU Ilmenau  
Karthik Sankaranarayanan, IBM Research India  
Katja Hose, Aalborg University  
Khuzaima Daudjee, University of Waterloo  
Konstantinos Karanasos, Microsoft  
Kostis Kyzirakos, CWI  
Lee Mong Li, NUS Singapore  
Lefteris Sidorouros, CWI  
Lei Zou, Peking University  
Li Xiong, Emory University  
Luc Bouganim, INRIA  
Luciano Barbosa, IBM Research Brazil  
Lucja Kot, Cornell University  
Mahashweta Das, Hewlett Packard Labs  
Marco Serafini, Qatar Computing Research Institute  
Martin Kersten, CWI  
Masatoshi Yoshikawa, Kyoto University  
Maurice Van Keulen, TU Twente  
Maya Ramanath, IIT Delhi  
Meichun Hsu, Hewlett Packard Labs  
Meikel Poess, Oracle  
Melanie Herschel, University of Stuttgart  
Michael Benedikt, Oxford University  
Michael Bohlen, University of Zurich  
Michael Hay, Colgate University  
Michael Grossniklaus, University of Konstanz  
Mirella Moro, Universidade Federal de Minas Gerais  
Mohamed Eltabakh, Worcester Polytechnic Institute  
Mohamed Mokbel, University of Minnesota  
Mohamed Sarwat, Arizona State University  
Mohammad Sadoghi, IBM Research T.J. Watson  
Mourad Ouzzani, Qatar Computing Research Institute  
Murat Kantarcioglu, UT Dallas  
Nan Zhang, George Washington University  
Nick Koudas, University of Toronto  
Nicolas Bruno, Microsoft Research

Nikolaus Augsten, University of Salzburg  
Nikos Mamoulis, Hong Kong University  
Norman Paton, University of Manchester  
Oliver Kennedy, University at Buffalo  
Panagiotis Papapetrou, Stockholm University  
Panos Kalnis, KAUST  
Panos Chrysanthis, University of Pittsburgh  
Paolo Merialdo, Roma Tre University  
Paris Koutris, University of Wisconsin-Madison  
Patricia Arocena, University of Toronto  
Peter Fischer, Universität Freiburg  
Peter Bailis, Stanford University  
Peter Alvaro, University of California, Santa Cruz  
Philippe Cudre-Mauroux, University of Fribourg  
Pierangela Samarati, University of Milan  
Pinar Tozun, IBM Research  
Raghav Kaushik, Microsoft Research  
Raluca Ada Popa, UC Berkeley  
Raymond Ng, University of British Columbia  
Reynold Cheng, Hong Kong University  
Ricardo Torres, UNICAMP Brazil  
S. Sudarshan, IIT Bombay  
Sai Wu, Zhejiang University  
Sebastian Michel, University of Kaiserslautern  
Selcuk Candan, Arizona State University  
Semih Salihoglu, University of Waterloo  
Senjuti Basu Roy, University of Washington Tacoma  
Seung-won Hwang, Yonsei University  
Sourav Bhowmick, Nanyang Technological University  
Spyros Blanas, Ohio State University  
Srikanta Bedathur, IBM Research India  
Stavros Papadopoulos, Intel Labs and MIT  
Stefanie Scherzinger, Ostbayerische Technische Hochschule Regensburg  
Stratis Viglas, University of Edinburgh  
Sudeepa Roy, Duke University  
Sudipto Das, Microsoft Research  
Sven Helmer, Free University of Bozen-Bolzano  
Tamer Ozsü, University of Waterloo  
Theodoros Rekatsinas, Stanford University  
Thomas Heinis, Imperial College  
Todd Green, Logicblox  
Torsten Grust, University of Tuebingen  
Tyson Condie, UCLA  
Umar Farooq Minhas, IBM Research  
Uwe Röhm, University of Sydney  
Verena Kantere, University of Geneva  
Viktor Leis, TU München  
Vivek Narasayya, Microsoft Research  
Wei Wang, University of New South Wales  
Wenchao Zhou, Georgetown University  
Wendy Wang, Stevens Institute of Technology  
Xiaochun Yang, Northeastern University, China  
Xiaodong Zhang, Ohio State University  
Xiaofang Zhou, University of Queensland  
Xiaohui Yu, York University  
Xiaoyang Wang, Fudan University  
Xin Luna Dong, Google

Yannis Manolopoulos, Aristotle University of Thessaloniki  
Yannis Velegarakis, University of Trento  
Yeye He, Microsoft Research  
Yi Chen, New Jersey Institute of Technology  
Yinan Li, Microsoft Research  
Yizhou Sun, Northeastern University  
Yoshiharu Ishikawa, Nagoya University  
Yuanyuan Tian, IBM Research (Almaden)  
Yufei Tao, University of Queensland  
Zhao Cao, Beijing Institute of Technology  
Zhifeng Bao, RMIT University  
Zoi Kaoudi, Qatar Computing Research Institute