VLDB2017
43rd International Conference on Very Large Data Bases, Munich, Germany

Proceedings of the VLDB Endowment

Volume 10, No. 13 – September 2017

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

PVLDB – Proceedings of the VLDB Endowment

Volume 10, No. 13, September 2017.

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

Copyright 2017 VLDB Endowment

Volume 10, Number 13, September 2017: VLDB 2017

Pages i – vii and 2025 – 2096

ISSN 2150-8097

Additional copies only online at: portal.acm.org, arxiv.org/corr, and www.vldb.org

TABLE OF CONTENTS

Front Matter

Research Papers

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
Ansger 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 Sidirourgos, 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 Ozsu, 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 Velegrakis, 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

Letter from the VLDB 2017 Program Chairs

The 43rd International Conference on Very Large Data Bases was held in Munich, from 28 August to 1 September 2017, organized locally by general chairs Alfons Kemper and Thomas Neumann. The conference was very successful, drawing 1018 registered attendees from 46 countries.

The technical program of VLDB 2017 included presentations of 131 research papers from Issues 1-11 of Volume 10 of PVLDB, plus 8 research papers from Issue 14 of Volume 9. The papers in this final issue (Issue 13) of Volume 10 will be "rolled over" for presentation at VLDB 2018. The technical program also included 20 industrial papers, 34 demonstrations, and 8 tutorials, all of which are collected in Issue 12 of Volume 10.

Each year, the VLDB conference confers a best paper award. For VLDB 2017, we compiled a shortlist of best paper candidates based on reviews and on nominations from the Volume 10 Associate Editors. The papers on this shortlist were carefully reviewed by a Best Paper Selection Committee consisting of three distinguished members of the PVLDB Review Board: Alan Fekete, Meichun Hsu, and S. Sudarshan. The Committee unanimously selected the following paper, from Volume 10, Issue 5, as the winner of the 2017 Best Paper award:

Provenence for Natural Language Queries
Daniel Deutch, Nave Frost, and Amir Gilad

In the words of the Committee: *The award is conferred on this paper for its innovation and potential for impact, in the area of natural language interfaces to databases. A novel aspect of this work is its focus on giving meaningful answers in natural language, along with explanations, to a natural language query. Two key insights of the paper are to generate answers by leveraging the structure from the original NL queries, and to provide simpler explanations by factorizing and summarizing provenance. User studies and performance results in the paper show the value and practicality of the proposed approach.*

The VLDB conference also confers an annual award for best demonstration, chosen by an Demo Track awards committee after viewing the demos at the conference. For 2017, two demonstrations were selected to share the Best Demo award:

GRAPE: Parallelizing Sequential Graph Computations
Wenfei Fan, Jingbo Xu, Yinghui Wu, Wenyuan Yu, Jiaxin Jiang

Interactive Navigation of Open Data Linkages
Erkang Zhu, Ken Pu, Fatemeh Nargesian, Renee Miller

Volume 10 of PVLDB, and the VLDB 2017 conference, would not have been possible without the hard work of the Associate Editors and the Review Board. This year, we established a process to identify and recognize particularly effective members of the Review Board. These reviewers consistently produce fair, informative, and timely reviews, and participate constructively in the discussion and decision-making process. To help identify top reviewers, we collected review feedback from authors, and also solicited nominations from the Associated Editors. Ultimately, we chose 9 members of the Review Board (about 5%) to receive this recognition. As announced at the conference, they are:

PVLDB 2017 Top Reviewers

Alan Fekete
Johann Gamper
Joseph Gonzales
Allison Holloway
Meichun Hsu
Konstantinos Karanasos
Jeff Pound
S. Sudarshan
Xiaodong Zhang

We would like to thank all of the members of the Review Board, as well as the Associated Editors, for their efforts.

Peter Boncz and Ken Salem
PVLDB Volume 10 Editors in Chief
VLDB 2017 Program Committee Chairs