

VLDB2018

44th International Conference on Very Large Data Bases, Rio de Janeiro, Brazil



Proceedings of the VLDB Endowment

Volume 11, No. 11 – July 2018

**Proceedings of the 44th International Conference on
Very Large Data Bases, Rio de Janeiro, Brazil**

Program Chairs:

Sihem Amer-Yahia and Jian Pei

Associate Editors – Research Track:

Luc Bouganim, Juliana Freire, Johannes Gehrke, Wook-Shin Han, Chris Jermaine, Jimmy Lin, Ioana Manolescu, Renee Miller, Mohamed Mokbel, Felix Naumann, Srinivasan Parthasarathy, Andrew Pavlo, S. Sudarshan, Jens Teubner, Yuanyuan Tian, Jianliang Xu, Meihui Zhang, Xiaodong Zhang

Proceedings Chairs:

Sourav Bhowmick, Ricardo Torres

PVLDB – Proceedings of the VLDB Endowment

Volume 11, No. 11, July 2018.

The 44th International Conference on Very Large Data Bases, Rio de Janeiro, Brazil.

Copyright 2018 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.

Volume 11, Number 11, July 2018: VLDB 2018

Pages i – viii and 1317 - 1741

ISSN 2150-8097

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

TABLE OF CONTENTS

Front Matter

Copyright Notice	i
Table of Contents	ii
VLDB 2018 Organization and Review Board	v

Research Papers

In-RDBMS Hardware Acceleration of Advanced Analytics.....Divya Mahajan, Joon Kyung Kim, Jacob Sacks, Adel Ardalan, Arun Kumar, Hadi Esmaeilzadeh	1317
Join Query Optimization Techniques for Complex Event Processing Applications	1332
Efficient Adaptive Detection of Complex Event Patterns. Join Query Optimization Techniques for Complex Event Processing Applications	1346
Robustness Metrics for Relational Query Execution Plans	1360
Question Answering Over Knowledge Graphs: Question Understanding Via Template Decomposition Weiguo Zheng, Jeffrey Xu Yu, Lei Zou, Hong Cheng	1373
Explaining Repaired Data with CFDs	1387
AIDA - Abstraction for Advanced In-Database Analytics	1400
RHEEM: Enabling Cross-Platform Data Processing -- May The Big Data Be With You! --..... Divy Agrawal, Sanjay Chawla, Bertty Contreras-Rojas, Ahmed Elmagarmid, Yasser Idris, Zoi Kaoudi, Sebastian Kruse, Ji Lucas, Essam Mansour, Mourad Ouzzani, Paolo Papotti, Jorge-Arnulfo Quiane-Ruiz, Nan Tang, Saravanan Thirumuruganathan, Anis Troudi	1414
An Experimental Evaluation of Task Assignment in Spatial Crowdsourcing	1428
2SCENT: An Efficient Algorithm to Enumerate All Simple Temporal Cycles	1441
Distributed Representations of Tuples for Entity Resolution	1454
Efficient Construction of Approximate Ad-Hoc ML models Through Materialization and Reuse	1468
Axiomatic Foundations and Algorithms for Deciding Semantic Equivalences of SQL Queries	1482

HomeRun: Scalable Sparse-Spectrum Reconstruction of Aggregated Historical Data	Faisal Almutairi, Fan Yang, Hyun Ah Song, Christos Faloutsos, Nicholas Sidiropoulos, Vladimir Zadorozhny	1496
Differentially Private Hierarchical Count-of-Counts Histograms	Yu-Hsuan Kuo, Cho-Chun Chiu, Daniel Kifer, Michael Hay, Ashwin Machanavajjhala	1509
Efficient Document Analytics on Compressed Data: Method, Challenges, Algorithms, Insights	Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, Wenguang Chen	1522
You Say `What', I Hear `Where' and `Why'? (Mis-)Interpreting SQL to Derive Fine-Grained Provenance	Tobias Müller, Benjamin Dietrich, Torsten Grust	1536
An Eight-Dimensional Systematic Evaluation of Optimized Search Algorithms on Modern Processors	Lars-Christian Schulz, David Broneske, Gunter Saake	1550
Vocalizing Large Time Series Efficiently	Immanuel Trummer, Mark Bryan, Ramya Narasimha	1563
Filter Before You Parse: Faster Analytics on Raw Data with Sparser	Shoumik Palkar, Firas Abuzaid, Peter Bailis, Matei Zaharia	1576
Streaming Graph Partitioning: An Experimental Study	Zainab Abbas, Vasiliki Kalavri, Paris Carbone, Vladimir Vlassov	1590
Efficient Distributed Memory Management with RDMA and Caching	Qingchao Cai, Wentian Guo, Hao Zhang, Divyakant Agrawal, Gang Chen, Beng Chin Ooi, Kian-Lee Tan, Yong Meng Teo, Sheng Wang	1604
Causal Consistency and Latency Optimality: Friend or Foe?	Diego Didona, Rachid Guerraoui, Jingjing Wang, Willy Zwaenepoel	1618
A Unified Approach to Route Planning for Shared Mobility	Yongxin Tong, Yuxiang Zeng, Zimu Zhou, Lei Chen, Jieping Ye, Ke Xu	1633
Moment-Based Quantile Sketches for Efficient High Cardinality Aggregation Queries	Edward Gan, Jialin Ding, Kai Sheng Tai, Vatsal Sharan, Peter Bailis	1647
How Good Are Modern Spatial Analytics Systems?	Varun Pandey, Andreas Kipf, Thomas Neumann, Alfons Kemper	1661
Locality-Sensitive Hashing for Earthquake Detection: A Case Study Scaling Data-Driven Science ..	Kexin Rong, Clara E. Yoon, Karianne J. Bergen, Hashem Elezabi, Peter Bailis, Philip Levis, Gregory C. Beroza	1674
ModelarDB: Modular Model-Based Time Series Management with Spark and Cassandra	Søren Jensen, Torben Pedersen, Christian Thomsen	1688
Exploiting Coroutines to Attack the "Killer Nanoseconds"	Christopher Jonathan, Umar Farooq Minhas, James Hunter, Justin Levandoski, Gor Nishanov	1702
The Tao of Inference in Privacy-Protected Databases	Vincent Bindschaedler, Paul Grubbs, David Cash, Thomas Ristenpart, Vitaly Shmatikov	1715

Efficient Searchable Encryption Through Compression
..... Ioannis Demertzis, Charalampos Papamanthou, Rajdeep Talapatra 1729

VLDB 2018 ORGANIZATION AND REVIEW BOARD

General Chairs

Alberto Laender, Universidade Federal de Minas Gerais
Fabio Porto, LNCC
Marco Antonio Casanova, PUC Rio

Honorary Chair

Antonio Furtado, PUC Rio
Nivio Ziviani, UFMG

Organization Committee Chair

Stephan Günemann, TUM
Alfons Kemper, TUM
Thomas Neumann, TUM

Program Chairs and Editors in Chief of PVLDB 11

Jian Pei, Simon Fraser University
Sihem Amer-Yahia, University of Grenoble Alpes, CNRS

Associate Editors of PVLDB 11

Andrew Pavlo, Carnegie Mellon University
Chris Jermaine, Rice University
Felix Naumann, Hasso Plattner Institute
Ioana Manolescu, INRIA Saclay
Jens Teubner, TU Dortmund
Jianliang Xu, Hong Kong Baptist U.
Jimmy Lin, University of Waterloo
Johannes Gehrke, Microsoft
Juliana Freire, New York University
Luc Bouganim, INRIA
Meihui Zhang, SUTD
Mohamed Mokbel, University of Minnesota
Renee Miller, University of Toronto
Srinivasan Parthasarathy, Ohio State University
S. Sudarshan, IIT Bombay
Wook-Shin Han, Postech
Xiaodong Zhang, Ohio State University
Yuanyuan Tian, IBM Almaden

VLDB Endowment Representative

Divesh Srivastava, AT&T Labs-research

Sponsorship Committee Chairs

Artur Ziviani, LNCC
Anand Deshpande, Persistent
Mike Carey, University of California, UCI
Patrick Valduriez, INRIA

Publicity Committee Chair

Carmem Hara, Universidade Federal do Paraná
Mahashweta Das, Visa Research

Tutorial Chairs

Nick Koudas, University of Toronto
Sergio Lifschitz, PUC Rio

Industrial Chairs

Karin Breitman
Rakesh Agrawal, Data Insight Laboratories

Demonstration Chairs

Ming Hua, Facebook
Vanessa Braganholo, Universidade Federal Fluminense

Panel Chairs

Letizia Tanca, Politecnico di Milano
Mario Nascimento, Univ. Alberta

Workshop Chairs

Mirella Moro, Universidade Federal Minas Gerais
Xuemin Lin, University of New South Wales

PhD Workshop Chairs

Altigran Silva, Universidade Federal do Amazonas
Senjuti Basu Roy, New Jersey Institute of Technology

Proceedings Chairs

Ricardo Torres, University of Campinas
Sourav Bhowmick, Nanyang Technological University

Website Chair

Daniel De Oliveira, Universidade Federal Fluminense
Enver Anibal Choque Cayo, LNCC

PVLDB Managing Editor

Divesh Srivastava, AT&T Labs-research

PVLDB Advisory Committee

Juliana Freire, Jayant Haritsa, Wolfgang Lehner, Chen Li,
Renée J. Miller, Tova Milo, M. Tamer Özsu, Divesh
Srivastava, Kian-Lee Tan

Research Track Review Board

Abdul Quamar, IBM Almaden
Aijun An, York University, Canada
Alan Fekete, University of Sydney
Alex Thomo, University of Victoria
Alexandra Meliou, University of Massachusetts Amherst
Allison Holloway, Oracle
Anastasia Ailamaki, EPFL
Andrea Cali, Birkbeck Univ. of London, UK
Andrew Pavlo, Carnegie Mellon University
Anja Gruenheid, Google Research
Anshumali Shrivastava, Rice University
Antonios Deligiannakis, Technical University of Crete
Arbee Chen, Asia University, Taiwan
Aristides Gionis, Aalto University
Arnab Bhattacharya, IIT Kanpur
Arun Kumar, University of California, San Diego
Ashraf Aboulmaga, Qatar Computing Research Institute
Ashwin Machanavajjhala, Duke
Asterios Katsifodimos, TU Berlin, Germany
Atsuyuki Morishima, University of Tsukuba
Avrilia Floratou, Microsoft
Azza Abouzied, New York University Abu Dhabi
Baihua Zheng, Singapore Management University
Barzan Mozafari, University of Michigan
Beng Chin Ooi, NUS
Bernd Amann, Université Pierre et Marie Curie, France
Bolin Ding, Microsoft Research
Byron Choi, HKBU
Carlo Curino, Microsoft
Ce Zhang, ETH
Chee-Yong Chan, National University of Singapore
Chi Wang, Microsoft Research
Chris Jermaine, Rice University
Christopher Re, Stanford University
Chuan Xiao, Nagoya Univ., Japan
Cristian Bizer, University of Mannheim
Da Yan, University of Alabama
Daisy Zhe Wang, University of Florida
Dan Ports, Univ. of Washington USA
Dario Colazzo, U. Paris Dauphine
David Koop, U Mass Dartmouth
De-Nian Yang, Academia Sinica
Denilson Barbosa, University of Alberta - Canada
Divesh Srivastava, AT&T Labs Research
Elke Rundensteiner, WPI
Emmanuel Muller, Hasso Plattner Institute
Essam Mansour, QCRI
Fatma Ozcan, IBM Almaden
Feida Zhu, Singapore Management University
Feifei Li, University of Utah
Felix Naumann, Hasso Plattner Institute
Fernando Chirigati, NYU
Florian Rusu, University of California, Merced
Floris Geerts, University of Antwerp, USA

Francesco Bonchi, Yahoo Labs Barcelona, Spain
François Goasdoué, U. Rennes 1
Gao Cong, Nanyang Technological University
George Fletcher, Eindhoven University of Technology the Netherlands
George Kollios, Boston University
Georgia Koutrika, ATHENA Research Center, Greece
Gillian Dobbie, Univ. of Auckland, New Zealand
Guoliang Li, Tsinghua University
H. Jagadish, University of Michigan
Haibo Hu, Hong Kong Polytechnic
Haixun Wang, Facebook
Hakan Ferhatosmanoglu, Bilkent University
Harish Doraiswamy, NYU
Haryadi Gunawi, Univ. of Chicago, USA
Heiko Mueller, NYU
Herodotos Herodotou, Cyprus University of Technology
Holger Pirk, MIT
Hong Cheng, Chinese University of Hong Kong
Huy Vo, CUNY
Ihab Ilyas, University of Waterloo
Immanuel Trummer, EPFL, Switzerland
Ioana Manolescu, INRIA Saclay, France
Ira Assent, University of Aarhus
Isabel Cruz, University of Illinois at Chicago, USA
Jaideep Vaidya, Rutgers University
James Cheng, Chinese University of Hong Kong, Hong Kong
Jeffrey Xu Yu, Chinese University of Hong Kong, Hong Kong
Jens Dittrich, Saarland University
Jens Teubner, TU Dortmund
Jianliang Xu, Hong Kong Baptist U.
Jiannan Wang, SFU
Jignesh Patel, UW - Madison
Jimmy Lin, University of Waterloo
Jing Gao, State University of New York at Buffalo
Johannes Gehrke, Microsoft
Jonathan Goldstein, Microsoft
Josep Domingo-Ferrer, Universitat Rovira i Virgili, Catalonia
Ju Fan, Remin University
Juliana Freire, New York University
Justin Levandoski, Microsoft Research
Karthik Ramachandra, Microsoft Gray Systems Lab
Ke Wang, SFU
Ken Barker, University of Calgary, Canada
Khuzaima Daudjee, University of Waterloo
Kyuseok Shim, Seoul National University
Laks V.S. Lakshmanan, The University of British Columbia
Lei Chen, Hong Kong University of Science and Technology
Letizia Tanca, Politecnico di Milano, Italy
Li Xiong, Emory University
Luc Bouganim, INRIA
Lukasz Golab, University of Waterloo

Luna Dong, Amazon.com
Magdalena Balanziska, University of Washington
Marco Serafini, Qatar Computing Research Institute, Qatar
Maria Damiani, University of Milano, Italy
Martin Theobald, University of Luxembourg
Masaru Kitsuregawa, University of Tokyo, Japan
Matteo Golfarelli, University of Bologna, Italy
Matthias Boehm, IBM Almaden
Matthias Renz, George Mason Univ., USA
Matthieu Latapy, LIP6, France
Meihui Zhang, SUTD
Melanie Herschel, Universität Stuttgart
Michael Cafarella, University of Michigan
Mohamed Mokbel, University of Minnesota
Mohamed Sharaf, University of Queensland, Australia
Mohamed Sarwat, ASU
Mohammad Sadoghi, Purdue University
Nesime Tatbul, Intel Labs and MIT
Nicolas Ancaux, INRIA
Niketani Pansare, IBM Almaden
Nikos Mamoulis, The University of Hong Kong
Ninghui Li, Purdue University
Norman May, SAP Research
Oktie Hassanzadeh, IBM Research
Olga Papaemmanouil, Brandeis University
Oliver Kennedy, University of Buffalo
Panagiotis Bours, Aarhus University
Panagiotis Karras, Aalborg University
Panagiotis Papapetrou, Stockholm University
Pankaj Agarwal, Duke University
Paolo Papotti, Arizona State University
Peter Bailis, MIT and Stanford
Philippe Lamarre, INSA-Lyon France
Pinar Tozun, IBM Research
Quanquan Gu, University of Virginia
Rainer Gemulla, Universität Mannheim
Rajeev Rastogi, Amazon
Ralf Schenkel, University of Trier
Rebecca Taft, MIT
Rene Mueller, IBM Research - Almaden
Renee Miller, University of Toronto
Reynold Cheng, The University of Hong Kong, China
Reza Akbarinia, INRIA
Ryan Johnson, Logic Blox USA
S. Sudarshan, IIT Bombay
Sandeep Tata, Google, USA
Sang Won Lee, Sungkyunkwan University, Korea
Sara Cohen, The Hebrew University of Jerusalem
Selçuk Candan, Arizona State University
Sergey Melnik, Google, USA
Sergio Greco, Univ. of Calabria, Italy
Shady Elbassuoni, AUB

Shahram Ghandeharizadeh, UCS
Shel Finkelstein, University of California at Santa Cruz, USA
Sourav S Bhowmick, Nanyang Technological University
Spyros Blanas, Ohio State University
Srikanta Bedathur, IBM India
Srikanta Tirthapura, Iowa State Univ., USA
Stéphane Bressan, National University of Singapore, Singapore
Steven Whang, Google, USA
Sudipto Das, Microsoft Research
Sungpack Hong, Oracle
Theodore Johnson, AT&T Labs, USA
Thomas Neumann, TU Munich
Tiark Rompf, Purdue University
Tim Kraska, Brown University
Tingjian Ge, University of Massachusetts at Lowell, USA
Ulf Leser, Humboldt-Universität zu Berlin
Umar Farooq Minhas, Microsoft Research
Venkatesan Chakaravarthy, IBM Research, India
Vijayshankar Raman, IBM Research - Almaden
Viktor Leis, Technische Universität München
Vincent Leroy, University Grenoble-Alps, CNRS, LIG, France
Vineet Chaoji, Amazon
Walid Aref, Purdue Univ., USA
Wang-Chien Lee, Pennsylvania State University, USA
Wei Lu, Renmin University of China
Wei Wang, National University of Singapore
Willis Lang, Microsoft Gray Systems
Wook-Shin Han, Postech
Wynne Hsu, National University of Singapore
Xiaochun Yang, Northeast University
Xiaodong Zhang, Ohio State University
Xiaofang Zhou, University of Queensland
Xuan Liu, Baidu
Xuemin Lin, University of New
Yael Amsterdamer, Tel-aviv Univ., Israel
Yizhou Sun, UCLA
Yongxin Tong, Beihang University, China
Yoshiharu Ishikawa, Nagoya University
Yuanyuan Tian, IBM Almaden
Yufei Tao, Univ. of Queensland
Zhenhui Li, Penn State University
Zhifeng Bao, RMIT University
Ziawasch Abedjan, TU Berlin
Zoi Kaoudi, Qatar Computing Research Institute, Qatar

LETTER FROM THE PROGRAM CHAIRS

The Proceedings of the VLDB Endowment (PVLDB) provides a high-quality publication service to the data management research community. Each volume offers twelve monthly submission deadlines on the first of each month and a quick, six week, reviewing cycle. This publication model was pioneered by PVLDB and combines a journal-style reviewing process, which includes a three month revision cycle, with the agility and visibility provided by rapid on-line publication and presentation at the annual VLDB conference.

This is the eleventh issue of the eleventh volume of the PVLDB.

PVLDB attracts many submissions, and the PVLDB reviewing process is implemented by a large team. The Review Board of PVLDB Volume 11 consists of 201 expert researchers, and reviewing is coordinated by 18 Associate Editors. Board members provide timely (within a 4-week deadline) high-quality reviews, and participate in online discussions with the Associate Editors and the other reviewers assigned to each paper. In order to honor these efforts, this year PVLDB will be recognizing outstanding reviewers. For this purpose, we are gathering feedback from the Associate Editors as well as from the authors, who in a feedback phase are asked to judge the quality of the reviews. These outstanding reviewers will be honored at the 44th VLDB conference, to be held in Rio de Janeiro, Brazil, on August 27-31, 2018. Here also the large majority of the PVLDB Volume 11 papers will be presented.

Issue 11 is a good example of the variety of topics covered in PVLDB, ranging from advanced analytics to query optimization, cross-platform data processing, data cleaning and entity resolution, provenance, databases for ML, graphs and knowledge graphs, data consistency, memory management, spatial data management, etc. We hope that the selected papers will provide insight to the readers and create impact by inspiring additional novel contributions.

Sihem Amer-Yahia and Jian Pei
PVLDB Volume 11 Editors in Chief
VLDB 2018 Program Committee Chairs