



Proceedings of the VLDB Endowment

Volume 12, No. 11 – July 2019

Editors in Chief:

Lei Chen and Fatma Özcan

Associate Editors:

Azza Abouzied, Selcuk Candan, Surajit Chaudhuri, Amol Desphande, Johann-Christoph Freytag, Rainer Gemulla, Nick Koudas, Georgia Koutrika, Yunyao Li, Alexandra Meliou, Arnab Nandi, M. Tamer Özsu, Themis Palpanas, Alkis Polyzotis, Kyuseok Shim, Xiaokui Xiao, Meihui Zhang

Publication Editors:

Abdul Quamar, Yongxin Tong

PVLDB – Proceedings of the VLDB Endowment

Volume 12, No. 11, July 2019.

All papers published in this issue will be presented at the 45th International Conference on Very Large Data Bases, Los Angeles, California, 2019.

Copyright 2019 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 12, Number 11, July 2019

Pages i – viii and 1235 - 1777

ISSN 2150-8097

Available at: <http://www.pvldb.org> and <https://dl.acm.org>.

TABLE OF CONTENTS

Front Matter

| | |
|---|----|
| Copyright Notice | i |
| Table of Contents | ii |
| PVLDB Organization and Review Board – Vol. 12 | v |

Research Papers

| | |
|---|------|
| Online Template Induction for Machine-Generated Emails..... | 1235 |
| <i>Michael J. Whittaker, Nick Edmonds, Sandeep Tata, James B. Wendt, Marc Najork</i> | |
| Querying Shortest Paths on Time Dependent Road Networks | 1249 |
| <i>Yong Wang, Guoliang Li, Nan Tang</i> | |
| Example-Driven Query Intent Discovery: Abductive Reasoning using Semantic Similarity | 1262 |
| <i>Anna Fariha, Alexandra Meliou</i> | |
| Automated Verification of Query Equivalence Using Satisfiability Modulo Theories | 1276 |
| <i>Qi Zhou, Joy Arulraj, Shamkant Navathe, William Harris, Dong Xu</i> | |
| Towards a Unified Framework for String Similarity Joins | 1289 |
| <i>Pengfei Xu, Jiaheng Lu</i> | |
| NETS: Extremely Fast Outlier Detection from a Data Stream via Set-Based Processing..... | 1303 |
| <i>Susik Yoon, Jae-Gil Lee, Byung Suk Lee</i> | |
| STAR: Scaling Transactions through Asymmetric Replication | 1316 |
| <i>Yi Lu, Xiangyao Yu, Samuel Madden</i> | |
| Subjective Databases..... | 1330 |
| <i>Yuliang Li, Aaron Feng, Jinfeng Li, Saran Mumick, Alon Halevy, Vivian Li, Wang-Chiew Tan</i> | |
| Fast and Robust Distributed Subgraph Enumeration | 1344 |
| <i>Xuguang Ren, Junhu Wang, Wook-Shin Han, Jeffrey Xu Yu</i> | |
| An Experimental Evaluation of Large Scale GBDT Systems | 1357 |
| <i>Fangcheng Fu, Jiawei Jiang, Yingxia Shao, Bin Cui</i> | |
| PrivateSQL: A Differentially Private SQL Query Engine | 1371 |
| <i>Ios Kotsogiannis, Yuchao Tao, Xi He, Maryam Fanaeepour, Ashwin Machanavajjhala, Michael Hay, Jerome Miklau</i> | |
| CAPER: A Cross-Application Permissioned Blockchain..... | 1385 |
| <i>Mohammad Javad Amiri, Divyakant Agrawal, Amr El Abbadi</i> | |
| Crossbow: Scaling Deep Learning with Small Batch Sizes on Multi-GPU Servers..... | 1399 |
| <i>Alexandros Koliouisis, Pijika Watcharapichat, Matthias Weidlich, Luo Mai, Paolo Costa, Peter Pietzuch</i> | |
| Finding Attribute-Aware Similar Region for Data Analysis | 1414 |
| <i>Kaiyu Feng, Gao Cong, Christian S. Jensen, Tao Guo</i> | |
| Intermittent Query Processing | 1427 |
| <i>Dixin Tang, Zechao Shang, Aaron J. Elmore, Sanjay Krishnan, Michael J. Franklin</i> | |

| | |
|--|------|
| Hillview: A trillion-cell spreadsheet for big data | 1442 |
| <i>Mihai Budiu, Parikshit Gopalan, Lalith Suresh, Udi Wieder, Han Krueger, Marcos K. Aguilera</i> | |
| Embedded Functional Dependencies and Data-completeness Tailored Database Design..... | 1458 |
| <i>Ziheng Wei, Sebastian Link</i> | |
| Ocean Vista: Gossip-Based Visibility Control for Speedy Geo-Distributed Transactions..... | 1471 |
| <i>Hua Fan, Wojciech Golab</i> | |
| An IDEA: An Ingestion Framework for Data Enrichment in AsterixDB | 1485 |
| <i>Xikui Wang, Michael Carey</i> | |
| DimmStore: Memory Power Optimization for Database Systems | 1499 |
| <i>Alexey Karyakin, Kenneth Salem</i> | |
| Generating Application-specific Data Layouts for In-memory Databases..... | 1513 |
| <i>Cong Yan, Alvin Cheung</i> | |
| Rewriting of Plain SO Tgds into Nested Tgds | 1526 |
| <i>Rihan Hai, Christoph Quix</i> | |
| Blockchain Meets Database: Design and Implementation of a Blockchain Relational Database..... | 1539 |
| <i>Senthil Nathan, Chander Govindarajan, Adarsh Saraf, Manish Sethi, Praveen Jayachandran</i> | |
| An Intermediate Representation for Optimizing Machine Learning Pipelines..... | 1553 |
| <i>Andreas Kuntz, Asterios Katsifodimos, Sebastian Schelter, Sebastian Breß, Tilmann Rabl, Volker Markl</i> | |
| Accelerating Raw Data Analysis with the ACCORDA Software and Hardware Architecture..... | 1568 |
| <i>Yuanwei Fang, Chen Zou, Andrew Chien</i> | |
| Comparing Synopsis Techniques for Approximate Spatial Data Analysis | 1583 |
| <i>A. B. Siddique, Ahmed Eldawy, Vagelis Hristidis</i> | |
| BlockchainDB - A Shared Database on Blockchains | 1597 |
| <i>Muhammad El-Hindi, Carsten Binnig, Arvind Arasu, Donald Kossmann, Ravi Ramamurthy</i> | |
| Efficient Task-Specific Data Valuation for Nearest Neighbor Algorithms | 1610 |
| <i>Ruoxi Jia, David Dao, Boxin Wang, Frances Ann Hubis, Nezihe Merve Gürel, Bo Li, Ce Zhang, Costas J. Spanos, Dawn Song</i> | |
| Distributed Implementations of Dependency Discovery Algorithms | 1624 |
| <i>Hemant Saxena, Lukasz Golab, Ihab F. Ilyas</i> | |
| Rethinking Database High Availability with RDMA Networks | 1637 |
| <i>Erfan Zamanian, Xiangyao Yu, Michael Stonebraker, Tim Kraska</i> | |
| Motivo: Fast Motif Counting via Succinct Color Coding and Adaptive Sampling..... | 1651 |
| <i>Marco Bressan, Stefano Leucci, Alessandro Panconesi</i> | |
| Arx: An Encrypted Database using Semantically Secure Encryption..... | 1664 |
| <i>Rishabh Poddar, Tobias Boelter, Raluca Ada Popa</i> | |
| Efficient Knowledge Graph Accuracy Evaluation | 1679 |
| <i>Junyang Gao, Xian Li, Yifan Ethan Xu, Bunyamin Sisman, Xin Luna Dong, Jun Yang</i> | |
| Optimizing Subgraph Queries by Combining Binary and Worst-Case Optimal Joins..... | 1692 |

Amine Mhedhbi, Semih Salihoglu

| | |
|---|------|
| Neo: A Learned Query Optimizer | 1705 |
| <i>Ryan C. Marcus, Parimarjan Negi, Hongzi Mao, Chi Zhang, Mohammad Alizadeh, Tim Kraska, Olga Papaemmanouil, Nesime Tatbul</i> | |
| Efficient Algorithms for Densest Subgraph Discovery | 1719 |
| <i>Yixiang Fang, Kaiqiang Yu, Reynold Cheng, Laks V.s. Lakshmanan, Xuemin Lin</i> | |
| Plan-Structured Deep Neural Network Models for Query Performance Prediction | 1733 |
| <i>Ryan C. Marcus, Olga Papaemmanouil</i> | |
| SLOG: Serializable, Low-latency, Geo-replicated Transactions..... | 1747 |
| <i>Kun Ren, Dennis Li, Daniel J. Abadi</i> | |
| GRAIL: Efficient Time-Series Representation Learning..... | 1762 |
| <i>John Paparrizos, Michael Franklin</i> | |

PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 12

Editors in Chief of PVLDB

Lei Chen, HKUST
Fatma Özcan, IBM Research - Almaden

Associate Editors of PVLDB

Azza Abouzied, NYU Abu Dhabi
Selcuk Candan, Arizona State University
Surajit Chaudhuri, Microsoft Research
Amol Desphande, University of Maryland
Johann-Christoph Freytag, HU Berlin
Rainer Gemulla, University of Mannheim
Nick Koudas, University of Toronto
Georgia Koutrika, Athena Research Center
Yunyao Li, IBM Research - Almaden
Alexandra Meliou, University of Massachusetts
Arnab Nandi, Ohio State University
M. Tamer Özsu, University of Waterloo
Themis Palpanas, French University Institute
Alkis Polyzotis, Google
Kyuseok Shim, Seoul National University

Xiaokui Xiao, National University of Singapore
Meihui Zhang, Beijing Institute of Technology

Publication Editors

Abdul Quamar, IBM Research - Almaden
Yongxin Tong, Beihang University

PVLDB Managing Editor

Wolfgang Lehner, TU Dresden

PVLDB Advisory Committee

Sihem Amer-Yahia, CNRS
Peter Boncz, Centrum Wiskunde & Informatica
Xin Luna Dong, Amazon
Juliana Freire, New York University
Wolfgang Lehner, TU Dresden
Renée J. Miller, University of Toronto
Tova Milo, Tel Aviv University
M. Tamer Özsu, University of Waterloo
Divesh Srivastava, AT&T Labs-Research

Review Board

Abdul Quamar, IBM Research - Almaden
Ada Waichee Fu, Chinese University of Hong Kong
Ahmet Erdem Sariyuca, University at Buffalo
Alan Fekete, University of Sydney
Alkis Simitsis, Hewlett Packard Labs
Ambuj Singh, UC Santa Barbara
Andrew Pavlo, Carnegie Mellon University
Angela Bonifati, University of Lyon
Arijit Khan, Nanyang Technological University
Arnab Bhattacharya, IIT Kanpur
Arun Kumar, University of California, San Diego
Arvind Arasu, Microsoft Research
Ashraf Aboulnaga, QCRI
Ashwin Machanavajhala, Duke University
Avrilia Floratou, Microsoft Research
Azade Nazi, Microsoft Research
Badrish Chandramouli, Microsoft Research
Barzan Mozafari, University of Michigan
Beng Chin OOI, National University of Singapore
Berthold Reinwald, IBM Research - Almaden
Bin Cui, Peiking University
Bobbie Cochrane, IBM
Bolin Ding, Alibaba
Boris Glavic, Illinois Institute of Technology
Bugra Gedik, Bilkent University
Byron Choi, Hong Kong Baptist University
Carlo Curino, Microsoft Research
Chee-Yong Chan, National University of Singapore
Chen Li, University of California, Irvine
Chengkai Li, UT Arlington
Chuan Lei, IBM Research - Almaden
Cong Yu, Google
Curtis Dyreson, Utah State University
Danica Probic, Oracle
Daniel Kifer, Penn State University
Davide Mottin, Hasso-Plattner Institute
Demetrios Zeinalipour-Yazti, University of Cyprus
Dimitris Papadias, HKUST
Diptikalyan Saha, IBM Research - India
Divyakant Agrawal, UC Santa Barbara
Donald Kossmann, Microsoft Research
Egemen Tanin, University of Melbourne
Eser Kandogan, IBM Research - Almaden
Essam M. Mansour, QCRI
Fabio Porto, LNCC
Fei Chiang, McMaster University
Feifei Li, University of Utah
Florin Rusu, University of California, Merced
Floris Geerts, University of Antwerp
George Papadakis, University of Athens
Goetz Graefe, Google
Guoliang Li, Tsinghua University
H. V. Jagadish, University of Michigan
Hakan Ferhatosmanoglu, Bilkent University
Hakan Hacigumus, Google
Hanghang Tong, Arizona State University
Helen Huang, University of Queensland

Heng Tao Shen, UESTC
Hong Cheng, Chinese University of Hong Kong
Hongzhi Yin, University of Queensland
Hua Lu, Aalborg University
Huiping Cao, New Mexico State University
Ilaria Bartolini, University of Bologna
Ilkay Altintas, San Diego Supercomputing Center
Immanuel Trummer, Cornell University
Ioana Manulescu, INRIA
Ismail Sengor Altingovde, METU
James Cheng, Chinese University of Hong Kong
Jens Dittrich, University of Saarland
Jens Teubner, TU Dortmund
Jianliang Xu, Hong Kong Baptist University
Jignesh Patel, University of Wisconsin - Madison
Jinyang Gao, National University of Singapore
Johann Gamper, Free University of Bozen-Bolzano
Jun Yang, Duke University
Junjie Yao, East China Normal University
Kai Zheng, UESTC
Karthik Sankaranarayanan, IBM Research - India
Katja Hose, Aalborg University
Khuzaima Daudjee, University of Waterloo
Kostas Stefanidis, University of Tampere
Kostas Zoumpatianos, Harvard University
Letizia Tanca, Polytechnic University of Milan
Lucian Popa, IBM Research - Almaden
Luna Dong, Amazon
Manos Karpathiotakis, Facebook London
Maria Luisa Sapino, University of Torino
Mario Nascimento, University of Alberta
Martin Theobald, University of Luxemburg
Mary Roth, IBM Research - Almaden
Matthias Boehm, Graz University of Technology
Matthias Renz, George Mason University
Maya Ramanath, Indian Institute of Technology Delhi
Melanie Herschel, University of Stuttgart
Michael Böhlen, University of Zurich
Michael Hay, Colgate University
Michael Mathioudakis, University of Helsinki
Min Li, JD.com
Mirek Riedewald, Northeastern University
Mirella Moro, Universidade Federal de Minas Gerais
Mohamed Eltabakh, Worcester Polytechnic Institute
Mohamed Mokbel, QCRI
Mohamed Sarwat, Arizona State University
Murat Kantarcioglu, University of Texas at Dallas
Nan Tang, QCRI
Nicolas Ancaux, INRIA
Nikolaus Augsten, University of Salzburg
Oktie Hassanzadeh, IBM Research - Yorktown
Olga Papaemmanouil, Brandeis University
Paolo Papotti, EURECOM
Parth Nagarkar, New Mexico State University
Pelin Angin, Middle East Technical University
Philip Bernstein, Microsoft Research
Philippe Bonnet, IT University of Copenhagen
Pinar Karagoz, Middle East Technical University
Pinar Tozun, IT University of Copenhagen

Raymond Ng, University of British Columbia
Sai Wu, Zhejiang University
Sang Kyun Cha, Seoul National University
Sebastian Breß, DFKI and TU Berlin
Semih Salihoglu, University of Waterloo
Senjuti Basu Roy, New Jersey Institute of Technology
Seung-Won Hwang, Yonsei University
Shaoxu Song, Tsinghua University
Shuo Shang, IIAI
Spyros Blanas, Ohio State University
Stefan Mangeold, Centrum Wiskunde & Informatica
Stefano Paraboschi, University of Bergamo
Steffen Zeuch, DFKI and TU Berlin
Stratis Viglas, University of Edinburgh
Sudip Roy, Google
Tingjian Ge, University of Massachusetts Lowell
Tyson Condie, University of California, Los Angeles
Umar Farooq Minhas, Microsoft Research
Vijayshankar Raman, Google
Viktor Leis, University of Jena
Vincent Oria, New Jersey Institute of Technology
Vivek Narasayya, Microsoft Research
Wenjie Zhang, University of New South Wales
Wook-Shin Han, POSTECH

Xiang Lian, Kent State University
Xiangmin Zhou, RMIT
Xiaochun Yang, Northeastern University
Xiaofang Zhou, University of Queensland
Li Xiong, Emory University
Xu Chu, Georgia Institute of Technology
Xuemin Lin, University of New Southwales
Yael Amsterdamer, Bar-Ilan University
Yannis Velegarakis, Utrecht University
Yanyan Shen, Shanghai Jiao Tong University
Yi Chen, New Jersey Institute of Technology
Ying Zhang, University of Technology Sydney
Yinghui Wu, Washington State University
Yingjun Wu, IBM Research - Almaden
Yingxia Shao, Peking University
Yongxin Tong, Beihang University
Yoshiharu Ishikawa, Nagoay University
Ye Yuan, Northeastern University
Yuanyuan Tian, IBM Research - Almaden
Yucel Saygin, Sabanci University
Yunjun Gao, Zhejiang University
Zhiguo Gong, University of Macau

LETTER FROM THE EDITORS IN CHIEF

The Proceedings of the VLDB Endowment (PVLDB) provides a high-quality journal publication service to the data management research community. Each volume offers twelve monthly submission deadlines on the first day of each month and a quick, six weeks, 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.

PVLDB attracts many submissions spanning diverse data management topics, and the PVLDB reviewing process is implemented by a large team of dedicated researchers. The Review Board of PVLDB Volume 12 consists of 166 expert researchers, and reviewing is coordinated by 17 Associate Editors. Review Board members provide timely (within a 4-week deadline) high-quality reviews, and participate actively in online discussions led by the Associate Editors for each paper.

This is the eleventh issue of the twelfth volume of the PVLDB. There are thirty-nine papers accepted in this issue that will be presented at the 45th International Conference on Very Large Data Bases (VLDB 2019), to be held in Los Angeles, California during August 26 to August 30, 2019.

For the tenth issue of PVLDB Volume 12, the review board has selected contributions proposing advances to topics on traditional database management problems, such as schema design, transaction management, data availability, memory power optimization, serializability in distributed databases, and query processing techniques, such as query by example, intermittent query processing, equivalent query identification, nearest neighbor and similarity search. This issue also contains papers on advanced data analytics and mining, including shortest path query, synopses comparison, subgraph search and enumeration, dense subgraph discovery, fast motif discovery and time series representation. Moreover, this issue covers advanced topics including knowledge bases, such as information extraction, schema mapping and knowledge base accuracy estimation, encrypted databases, private SQL, raw data analytics and trillion-cell spreadsheet visualization, and blockchain, such as permission-less blockchain and blockchain databases. Last, but not the least, this issue covers machine learning and deep learning topics in the context of data analytics, including evaluation of GBDTs, scaling deep learning, deep learning for query optimization, and intermediate representation for machine learning. We hope that the readers will find the selected papers engaging, and thought provoking. We also hope that the selected papers will provide valuable insights and inspire novel systems contributions and follow-up research.

Lei Chen and Fatma Özcan

PVLDB Volume 12 Editors in Chief