



# Proceedings of the VLDB Endowment

Volume 13, No. 8 – April 2020

Editors in Chief:

**Magdalena Balazinska and Xiaofang Zhou**

Associate Editors:

**Azza Abouzied, Amr El Abbadi, Phil Bernstein, Xin Luna Dong, Zi (Helen) Huang,  
Nick Koudas, Georgia Koutrika, Guoliang Li, Alexandra Meliou, Felix Naumann,  
Dan Olteanu, M. Tamer Özsu, Aditya Parameswaran, Andy Pavlo,  
Xiaokui Xiao, Jeffrey Xu Yu, Meihui Zhang, Jingren Zhou**

Publication Editors:

**Hiroaki Shiokawa and Sen Wang**

PVLDB – Proceedings of the VLDB Endowment

Volume 13, No. 8, April 2020.

PVLDB is indexed in Scopus (Elsevier) as well as covered by the following Clarivate Analytics services:

- Science Citation Index Expanded (also known as SciSearch®)
- Journal Citation Reports/Science Edition, and
- Current Contents®/Engineering Computing and Technology

All papers published in this issue will be presented at the 46th International Conference on Very Large Data Bases, Tokyo, Japan, 2020.

## **Copyright 2020 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 13, Number 8, April 2020

Pages i – vi and 1147 - 1318

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. 13 .....	iii

### Research Papers

Dash: Scalable Hashing on Persistent Memory.....	1147
<i>Baotong Lu, Xiangpeng Hao, Tianzheng Wang, Eric Lo</i>	
The PGM-index: a fully-dynamic compressed learned index with provable worst-case bounds .....	1162
<i>Paolo Ferragina, Giorgio Vinciguerra</i>	
Diagnosing Root Causes of Intermittent Slow Queries in Cloud Databases.....	1176
<i>Minghua Ma, Zheng Yin, Shenglin Zhang, Sheng Wang, Christopher Zheng, Xinhao Jiang, Hanwen Hu, Cheng Luo, Yilin Li, Nengjun Qiu, Feifei Li, Changcheng Chen, Dan Pei</i>	
Pangolin: An Efficient and Flexible Graph Mining System on CPU and GPU .....	1190
<i>Xuhao Chen, Roshan Dathathri, Gurbinder S Gill, Keshav Pingali</i>	
Quantifying TPC-H Choke Points and Their Optimizations.....	1206
<i>Markus Dreseler, Martin Boissier, Tilmann Rabl, Matthias Uflacker</i>	
Efficient Algorithms for Crowd-Aided Categorization.....	1221
<i>Yuanbing Li, Xian Wu, Yifei Jin, Jian Li, Guoliang Li</i>	
Set-valued Data Publication with Local Privacy: Tight Error Bounds and Efficient Mechanisms .....	1234
<i>Shaowei Wang, Yuqiu Qian, Jiachun Du, Wei Yang, Liusheng Huang, Hongli Xu</i>	
Translation of Array-Based Loops to Distributed Data-Parallel Programs.....	1248
<i>Leonidas Fegaras, Md Hasanuzzaman Noor</i>	
Incrementalization of Graph Partitioning Algorithms.....	1261
<i>Wenfei Fan, Muyang Liu, Chao Tian, Ruiqi Xu, Jingren Zhou</i>	
Optimizing Item and Subgroup Configurations for Social-Aware VR Shopping.....	1275
<i>Shao-Heng Ko, Hsu-Chao Lai, Hong-Han Shuai, Wang-Chien Lee, Philip Yu, De-Nian Yang</i>	
Efficient Confidentiality-Preserving Data Analytics over Symmetrically Encrypted Datasets .....	1290
<i>Savvas Savvides, Darshika Khandelwal, Patrick Eugster</i>	
Single Machine Graph Analytics on Massive Datasets Using Intel Optane DC Persistent Memory ..	1304
<i>Gurbinder S Gill, Roshan Dathathri, Loc Hoang, Ramesh Peri, Keshav Pingali</i>	

## **PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 13**

### **Editors in Chief of PVLDB**

Magdalena Balazinska (University of Washington, USA)  
Xiaofang Zhou (University of Queensland, Australia)

Meihui Zhang (Beijing Institute of Technology, China)  
Jingren Zhou (Alibaba Group, China)

### **Associate Editors of PVLDB**

Azza Abouzied (New York University Abu Dhabi, UAE)  
Amr El Abbadi (University of California, Santa Barbara, USA)  
Phil Bernstein (Microsoft Research, USA)  
Xin Luna Dong (Amazon, USA)  
Zi (Helen) Huang (University of Queensland, Australia)  
Nick Koudas (University of Toronto, Canada)  
Georgia Koutrika (Athena Research Center, Greece)  
Guoliang Li (Tsinghua University, China)  
Alexandra Meliou (University of Massachusetts, Amherst, USA)  
Felix Naumann (Hasso Plattner Institute, University of Potsdam, Germany)  
Dan Olteanu (University of Oxford, United Kingdom)  
M. Tamer Özsu (University of Waterloo, Canada)  
Aditya Parameswaran (University of California, Berkeley, USA)  
Andy Pavlo (Carnegie Mellon University, USA)  
Xiaokui Xiao (National University of Singapore, Singapore)  
Jeffrey Xu Yu (The Chinese University of Hong Kong, China)

### **Publication Editors**

Hiroaki Shiokawa (University of Tsukuba, Japan)  
Sen Wang (University of Queensland, Australia)

### **PVLDB Managing Editor**

Wolfgang Lehner (TU Dresden, Germany)

### **PVLDB Advisory Committee**

Divesh Srivastava (AT&T Labs-Research, USA)  
M. Tamer Özsu (University of Waterloo, Canada)  
Juliana Freire (New York University, USA)  
Xin Luna Dong (Amazon, USA)  
Peter Boncz (CWI, Netherlands)  
Xiaofang Zhou (University of Queensland, Australia)  
Magdalena Balazinska (University of Washington, USA)  
Lei Chen (Hong Kong University of Science and Technology, China)  
Fatma Ozcan (IBM Almaden, USA)  
Graham Cormode (University of Warwick, United Kingdom)  
Felix Naumann (HPI, Germany)

## Review Board

Ziawasch Abedjan (TU Berlin, Germany)  
Ashraf Aboulnaga (Qatar Computing Research Institute, Qatar)  
Pelin Angin (Middle East Technical University, Turkey)  
Arvind Arasu (Microsoft Research, USA)  
Joy Arulraj (Georgia Tech, USA)  
Manos Athanassoulis (Boston University, USA)  
Zhifeng Bao (RMIT University, Australia)  
Iliaria Bartolini (University of Bologna, Italy)  
Leilani Battle (University of Maryland, USA)  
Kaustubh Beedkar (TU Berlin, Germany)  
Arnab Bhattacharya (IIT Kanpur, India)  
Sourav S Bhowmick (Nanyang Technological University, Singapore)  
Carsten Binnig (TU Darmstadt, Germany)  
Spyros Blanas (The Ohio State University, USA)  
Matthias Boehm (Graz University of Technology, Austria)  
Alexander Böhm (SAP SE, Germany)  
Michael Böhlen (University of Zürich, Switzerland)  
Peter Boncz (Centrum Wiskunde & Informatica, Netherlands)  
Angela Bonifati (Lyon 1 University, France)  
Philippe Bonnet (IT University of Copenhagen, Denmark)  
Renata Borovica-Gajic (University of Melbourne, Australia)  
Huiping Cao (New Mexico State University, USA)  
Lei Cao (MIT, USA)  
Lijun Chang (The University of Sydney, Australia)  
Surajit Chaudhuri (Microsoft Research, USA)  
Lei Chen (Hong Kong University of Science and Technology, China)  
Hong Cheng (The Chinese University of Hong Kong, China)  
Reynold Cheng (The University of Hong Kong, China)  
Fei Chiang (McMaster University, Canada)  
Xu Chu (Georgia Tech, USA)  
Bobbie Cochrane (IBM, USA)  
Gao Cong (Nanyang Technological University, Singapore)  
Brian Cooper (Google, USA)  
Natacha Crooks (University of Texas at Austin, USA)  
Andrew Crotty (Brown University, USA)  
Bin Cui (Peking University, China)  
Sudipto Das (Amazon Web Service, USA)  
Akash Das Sarma (Facebook, USA)  
Khuzaima Daudjee (University of Waterloo, Canada)  
Niv Dayan (Harvard University, USA)  
Dong Deng (Rutgers University, USA)  
Bailu Ding (Microsoft Research, USA)  
Bolin Ding (Alibaba Group, China)  
Jens Dittrich (Saarland University, Germany)  
Harish Doraiswamy (New York University, USA)  
Eduard C. Dragut (Temple University, USA)  
Curtis Dyreson (Utah State University, USA)  
Mohamed Y. Eltabakh (Teradata Labs, USA)

Jose M. Faleiro (Microsoft Research, USA)  
Ju Fan (Renmin University of China, China)  
Raul Castro Fernandez (The University of Chicago, USA)  
Avrilia Floratou (Microsoft Research, USA)  
Avigdor Gal (Technion, Israel)  
Alex Galakatos (Brown University, USA)  
Johann Gamper (Free University of Bozen-Bolzano, Italy)  
Jing Gao (University at Buffalo, USA)  
Yunjun Gao (Zhejiang University, China)  
Tingjian Ge (University of Massachusetts, Lowell, USA)  
Floris Geerts (University of Antwerp, Belgium)  
Johannes Gehrke (Microsoft Research, USA)  
Jonathan Goldstein (Microsoft Research, USA)  
Torsten Grust (University of Tübingen, Germany)  
Wook-Shin Han (POSTECH, South Korea)  
Takahiro Hara (Osaka University, Japan)  
Oktie Hassanzadeh (IBM Research, USA)  
Michael Hay (Colgate University, USA)  
Xi He (University of Waterloo, Canada)  
Melanie Herschel (University of Stuttgart, Germany)  
Katja Hose (Aalborg University, Denmark)  
Wen Hua (The University of Queensland, Australia)  
Xin Huang (Hong Kong Baptist University, China)  
Yan Huang (University of North Texas, USA)  
Seung-won Hwang (Yonsei University, South Korea)  
Christopher Jermaine (Rice University, USA)  
Ruoming Jin (Kent State University, USA)  
Eser Kandogan (Megagon Labs, USA)  
Murat Kantarcioglu (University of Texas at Dallas, USA)  
Verena Kantere (University of Ottawa, Canada)  
Pinar Karagoz (Middle East Technical University, Turkey)  
Manos Karpathiotakis (Facebook, United Kingdom)  
Batya Kenig (University of Washington, USA)  
Oliver Kennedy (University at Buffalo, USA)  
Arijit Khan (Nanyang Technological University, Singapore)  
Daniel Kifer (Pennsylvania State University, USA)  
Hideaki Kimura (Oracle, USA)  
Sanjay Krishnan (University of Chicago, USA)  
Arun Kumar (University of California, San Diego, USA)  
Chuan Lei (IBM Research - Almaden, USA)  
Viktor Leis (Technical University of Munich, Germany)  
Ulf Leser (Humboldt-Universität zu Berlin, Germany)  
Chengkai Li (The University of Texas at Arlington, USA)  
Feifei Li (University of Utah, USA)  
Rong-Hua Li (Beijing Institute of Technology, China)  
Sebastian Link (The University of Auckland, New Zealand)  
Chengfei Liu (Swinburne University of Technology, Australia)  
Hua Lu (Aalborg University, Denmark)  
Jiaheng Lu (University of Helsinki, Finland)  
Wei Lu (Renmin University of China, China)  
Shuai Ma (Beihang University, China)  
Nikos Mamoulis (University of Ioannina, Greece)  
Ioana Manolescu (INRIA, France)  
Essam Mansour (Concordia University, Canada)

Ryan Marcus (MIT, USA)  
 Sergey Melnik (Google, USA)  
 Mohamed Mokbel (Qatar Computing Research Institute, Qatar)  
 Mirella Moura Moro (Universidade Federal de Minas Gerais, Brazil)  
 Davide Mottin (Aarhus University, Denmark)  
 Parth Nagarkar (New Mexico State University, USA)  
 Faisal Nawab (University of California, Santa Cruz, USA)  
 Thomas Neumann (Technical University of Munich, Germany)  
 Milos Nikolic (The University of Edinburgh, United Kingdom)  
 Beng Chin Ooi (National University of Singapore, Singapore)  
 Ismail Oukid (SAP SE, USA)  
 Mourad Ouzzani (Qatar Computing Research Institute, Qatar)  
 Themis Palpanas (Paris Descartes University, France)  
 George Papadakis (University of Athens, Greece)  
 Olga Papaemmanouil (Brandeis University, USA)  
 Thorsten Papenbrock (Hasso Plattner Institute, Germany)  
 Paolo Papotti (EURECOM, France)  
 Stefano Paraboschi (Universita' degli Studi di Bergamo, Italy)  
 Yongjoo Park (University of Michigan, USA)  
 Jignesh M. Patel (University of Wisconsin-Madison, USA)  
 Peter Pietzuch (Imperial College London, United Kingdom)  
 Holger Pirk (Imperial College London, United Kingdom)  
 Fábio Porto (National Laboratory for Scientific Computing (LNCC), Brazil)  
 Dan R. K. Ports (Microsoft Research, USA)  
 Lu Qin (University of Technology Sydney, Australia)  
 Abdul H. Quamar (IBM Research – Almaden, USA)  
 Tilmann Rabl (TU Berlin, Germany)  
 Karthik Ramachandra (Microsoft Research, USA)  
 Maya Ramanath (IIT Delhi, India)  
 Berthold Reinwald (IBM Research, USA)  
 Theodoros Rekatsinas (University of Wisconsin-Madison, USA)  
 Uwe Roehm (The University of Sydney, Australia)  
 Jennie Rogers (Northwestern University, USA)  
 Florin Rusu (University of California, Merced, USA)  
 Diptikalyan Saha (IBM Research AI India, India)  
 Ken Salem (University of Waterloo, Canada)  
 Semih Salihoglu (University of Waterloo, Canada)  
 Maria Luisa Sapino (University of Torino, Italy)  
 A. Erdem Sariyuce (University at Buffalo, USA)  
 Mohamed Sarwat (Arizona State University, USA)  
 Maximilian Schleich (University of Oxford, United Kingdom)  
 Mohamed Sharaf (University of Queensland, Australia)  
 Yanyan Shen (Shanghai Jiao Tong University, China)  
 Kyuseok Shim (Seoul National University, South Korea)  
 Prashant Shiralkar (Amazon, USA)  
 Alkis Simitsis (Hewlett Packard Enterprise, USA)  
 Kostas Stefanidis (Tampere University, Finland)  
 Rebecca Taft (Cockroach Labs, USA)  
 Nan Tang (Qatar Computing Research Institute, Qatar)  
 Yufei Tao (The Chinese University of Hong Kong, China)  
 Jens Teubner (TU Dortmund, Germany)  
 Andreas Thor (University of Applied Sciences for Telecommunications Leipzig, Germany)  
 Yongxin Tong (Beihang University, China)  
 Anthony K. H. Tung (National University of Singapore, Singapore)  
 Yannis Velegarakis (Utrecht University, Netherlands)  
 Stratis Viglas (University of Edinburgh, United Kingdom)  
 Daisy Zhe Wang (University of Florida, USA)  
 Guoren Wang (Beijing Institute of Technology, China)  
 Jiannan Wang (Simon Fraser University, USA)  
 Junhu Wang (Griffith University, Australia)  
 Sibow Wang (The Chinese University of Hong Kong, China)  
 Eugene Wu (Columbia University, USA)  
 Yingjun Wu (Amazon Web Service, USA)  
 Yinglong Xia (Huawei Research America, USA)  
 Chuan Xiao (Osaka University, Japan)  
 Yanghua Xiao (Fudan University, China)  
 Li Xiong (Emory University, USA)  
 Jianliang Xu (Hong Kong Baptist University, China)  
 Xiaochun Yang (Northeastern University, China)  
 Junjie Yao (East China Normal University, China)  
 Hongzhi Yin (The University of Queensland, Australia)  
 Man Lung Yiu (Hong Kong Polytechnic University, China)  
 Haruo Yokota (Tokyo Institute of Technology, Japan)  
 Masatoshi Yoshikawa (Kyoto University, Japan)  
 Xiangyao Yu (University of Wisconsin-Madison, USA)  
 Demetrios Zeinalipour-Yazti (University of Cyprus, Cyprus)  
 Baihua Zheng (Singapore Management University, Singapore)  
 Rui Zhang (University of Melbourne, Australia)  
 Wenjie Zhang (University of New South Wales, Australia)  
 Xiaofei Zhang (The University of Memphis, USA)  
 Ying Zhang (University of Technology Sydney, Australia)  
 Yuanyuan Zhu (Wuhan University, China)  
 Lei Zou (Peking University, China)  
 Kostas Zoumpatianos (Harvard University, USA)

## 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 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.

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 13 consists of 186 expert researchers, and reviewing is coordinated by 18 Associate Editors. Review Board members provide timely (within a 3-week deadline) high-quality reviews, and participate actively in online discussions led by the Associate Editors for each paper. When needed, the Associate Editors together with the Editors-in-Chief solicit additional reviews from external experts. We give special thanks here to those additional reviewers who in most cases need to complete their expert reviews on a very short notice.

Most of the accepted papers go through a revision process that requires a second round of reviews after the authors have addressed an initial set of issues and concerns raised by the reviewers during the first round. Some papers are further accepted with shepherding, which means that one of the reviewers works with the authors to address a final set of comments.

This is the eighth issue of the thirteenth volume of PVLDB. There are twelve accepted papers in this issue. These papers will be presented at the 46th International Conference on Very Large Data Bases (PVLDB 2020), originally planned to be held in Tokyo, Japan during August 31 to September 4, 2020 but now will be rescheduled due to the coronavirus pandemic. The final format and time for PVLDB2020 have not been decided yet.

This issue includes multiple papers on solving new data management and data processing issues arising from the use of persistent memory, CPU and GPU, cloud, and data-parallel programming. Data privacy issues have been addressed in two papers, and graph processing remains to be a hot topic with three papers in this issue on this topic. Other papers in this issue report new advances on traditional database topics such as indexing, database performance benchmarking and optimization, as well as new topics related to machine learning and crowdsourcing. There is also one very interesting paper to deal with novel applications of VR shopping.

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.

---

Magdalena Balazinska and Xiaofang Zhou

PVLDB Volume 13 Editors in Chief