



# Proceedings of the VLDB Endowment

Volume 16, No. 13 – September 2023

**Editors in Chief:**

Georgia Koutrika and Jun Yang

**Associate Editors:**

Alkis Simitsis, Amol Deshpande, Angela Bonifati, Ashwin Machanavajjhala, Badrish Chandramouli,  
Boris Glavic, Ce Zhang, Cyrus Shahabi, Dan Olteanu, Eric Lo, Evaggelia Pitoura, Evimaria Terzi,  
Gustavo Alonso, Helen (Zi) Huang, Hong Cheng, Kenneth Ross, Khuzaima Daudjee, Kyuseok Shim,  
Letizia Tanca, Lucian Popa, Magdalena Balazinska, Meihui Zhang, Neoklis Polyzotis,  
Nesime Tatbul, Nikos Mamoulis, Rachel Pottinger, Wenjie Zhang, Wolfgang Gatterbauer,  
Wook-Shin Han, Xiaokui Xiao, Yannis Velegrakis, Yanyan Shen, Yi Chen, Yongxin Tong, Zhifeng Bao

**Publication Editors:**

Manos Athanassoulis, Kostas Stefanidis, Ju Fan

PVLDB – Proceedings of the VLDB Endowment

Volume 16, No. 13, September 2023.

All papers published in this issue will be presented at the 49th International Conference on Very Large Data Bases, Vancouver, Canada, 2023.

## Copyright 2023 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 16, Number 13, September 2023

Pages i – vii and 4160 - 4352

ISSN 2150-8097

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

## TABLE OF CONTENTS

### Front Matter

Copyright Notice .....	i
Table of Contents .....	ii
PVLDB Organization and Review Board – Vol. 16 .....	iv

### Research Papers

Doquet: Differentially Oblivious Range and Join Queries with Private Data Structures .....	4160
<i>Lina Qiu, Georgios Kellaris, Nikos Mamoulis, Kobbi Nissim, George Kollios</i>	
AMNES: Accelerating the computation of data correlation using FPGAs.....	4174
<i>Monica Chiosa, Thomas B. Preußner, Michaela Blott, Gustavo Alonso</i>	
VOCALExplore: Pay-as-You-Go Video Data Exploration and Model Building .....	4188
<i>Maureen Daum, Enhao Zhang, Dong He, Stephen Mussmann, Brandon Haynes, Ranjay Krishna, Magdalena Balazinska</i>	
Flexible Resource Allocation for Relational Database-as-a-Service .....	4202
<i>Pankaj Arora, Surajit Chaudhuri, Sudipto Das, Junfeng Dong, Cyril George, Ajay Kalhan, Arnd Christian König, Willis Lang, Changsong Li, Feng Li, Jiaqi Liu, Lukas M. Maas, Akshay Mata, Ishai Menache, Justin Moeller, Vivek Narasayya, Matthaios Olma, Morgan Oslake, Elnaz Rezai, Yi Shan, Manoj Syamala, Shize Xu, Vasileios Zois</i>	
ShadowAQP: Efficient Approximate Group-by and Join Query via Attribute-oriented Sample Size Allocation and Data Generation.....	4216
<i>Rong Gu, Han Li, Haipeng Dai, Wenjie Huang, Jie Xue, Meng Li, Jiaqi Zheng, Haoran Cai, Yihua Huang, Guihai Chen</i>	
Optimizing Data Pipelines for Machine Learning in Feature Stores .....	4230
<i>Rui Liu, Kwanghyun Park, Fotis Psallidas, Xiaoyong Zhu, Jinghui Mo, Rathijit Sen, Matteo Interlandi, Konstantinos Karanasos, Yuan Yuan Tian, Jesús Camacho-Rodríguez</i>	
SparqLog: A System for Efficient Evaluation of SPARQL 1.1 Queries via Datalog .....	4240
<i>Renzo Angles, Georg Gottlob, Aleksandar Pavlović, Reinhard Pichler, Emanuel Sallinger</i>	
Solver-In-The-Loop Cluster Resource Management for Database-as-a-Service .....	4254
<i>Arnd Christian König, Yi Shan, Karan Newatia, Luke Marshall, Vivek Narasayya</i>	
RTIndex: Exploiting Hardware-Accelerated GPU Raytracing for Database Indexing .....	4268
<i>Justus Henneberg, Felix Schuhknecht</i>	
ContTune: Continuous Tuning by Conservative Bayesian Optimization for Distributed Stream Data Processing Systems .....	4282
<i>Jinqing Lian, Xinyi Zhang, Yingxia Shao, Zenglin Pu, Qingfeng Xiang, Yawen Li, Bin Cui</i>	
Single Update Sketch with Variable Counter Structure .....	4296
<i>Dimitrios Melissourgios, Haibo Wang, Shigang Chen, Chaoyi Ma, Shiping Chen</i>	
Can Large Language Models Predict Data Correlations from Column Names?.....	4310
<i>Immanuel Trummer</i>	
GraphOS: Towards Oblivious Graph Processing .....	4324
<i>Javad Ghareh Chamani, Ioannis Demertzis, Dimitrios Papadopoulos, Charalampos Papamanthou, Rasool Jalili</i>	

Catalyst: Optimizing Cache Management for Large In-memory Key-value Systems..... 4339  
*Kefei Wang, Feng Chen*

## PVLDB ORGANIZATION AND REVIEW BOARD - Vol. 16

### Editors in Chief of PVLDB

Georgia Koutrika (Athena Research Center)  
Jun Yang (Duke University)

### Associate Editors of PVLDB

Alkis Simitsis (Athena Research Center)  
Amol Deshpande (University of Maryland at College Park)  
Angela Bonifati (Lyon 1 University)  
Ashwin Machanavajjhala (Duke University/Tumult Labs)  
Badrish Chandramouli (Microsoft Research)  
Boris Glavic (Illinois Institute of Technology)  
Ce Zhang (ETH Zurich)  
Cyrus Shahabi (University of Southern California)  
Dan Olteanu (University of Zurich)  
Eric Lo (The Chinese University of Hong Kong)  
Evaggelia Pitoura (University of Ioannina)  
Evmiria Terzi (Boston University)  
Gustavo Alonso (ETH Zurich)  
Helen (Zi) Huang (University of Queensland)  
Hong Cheng (The Chinese University of Hong Kong)  
Kenneth Ross (Columbia University)  
Khuzaima Daudjee (University of Waterloo)  
Kyuseok Shim (Seoul National University)  
Letizia Tanca (Politecnico di Milano)  
Lucian Popa (IBM Research - Almaden)  
Magdalena Balazinska (University of Washington)  
Meihui Zhang (Beijing Institute of Technology)  
Neoklis Polyzotis (Databricks)  
Nesime Tatbul (Intel Labs and MIT)  
Nikos Mamoulis (University of Ioannina)  
Rachel Pottinger (University of British Columbia)  
Wenjie Zhang (University of New South Wales)  
Wolfgang Gatterbauer (Northeastern University)

Wook-Shin Han (Pohang University of Science and Technology)  
Xiaokui Xiao (National University of Singapore)  
Yannis Velegrakis (University of Trento and Utrecht University)  
Yanyan Shen (Shanghai Jiao Tong University)  
Yi Chen (New Jersey Institute of Technology)  
Yongxin Tong (Beihang University)  
Zhifeng Bao (RMIT University)

### Publication Editors

Manos Athanassoulis (Boston University)  
Kostas Stefanidis (Tampere University)  
Ju Fan (Renmin University of China)

### PVLDB Managing Editor

Wolfgang Lehner (Dresden University of Technology)

### PVLDB Advisory Board

Vanessa Braganholo (Universidade Federal Fluminense)  
Sourav S Bhowmick (Nanyang Technological University)  
Chris Jermaine (Rice University)  
Peter Triantafillou (University of Warwick)  
Xin Luna Dong (Facebook)  
Fatma Ozcan (Google)  
Lei Chen (Hong Kong University of S&T)  
Juliana Freire (New York University)  
Graham Cormode (University of Warwick)  
Divesh Srivastava (AT&T Labs-Research)  
Wolfgang Lehner (Dresden University of Technology)  
Felix Naumann (HPI)  
Xuemin Lin (University of New South Wales)  
Georgia Koutrika (Athena Research Center)  
Jun Yang (Duke University)

## Review Board

Abolfazl Asudeh (University of Illinois at Chicago)  
Alexander J Ratner (University of Washington)  
Alexandra Meliou (University of Massachusetts Amherst)  
Amelie Marian (Rutgers University)  
Amir Gilad (Duke University)  
Amir Shaikhha (University of Edinburgh)  
Amrita Roy Chowdhury (University of Wisconsin-Madison)  
Anastasios Kementsietsidis (Google Research)  
Andrew Crotty (Carnegie Mellon University)  
Anna Fariha (Microsoft)  
Anton Dignös (Free University of Bozen-Bolzano)  
Antonios Deligiannakis (Technical University of Crete)  
Arijit Khan (Nanyang Technological University)  
Ashraf Aboulnaga (Qatar Computing Research Institute, HBKU)  
Asterios Katsifodimos (Delft University of Technology)  
Baihua Zheng (Singapore Management University)  
Bin Cui (Peking University)  
Bingsheng He (National University of Singapore)  
Binhang Yuan (ETH Zurich)  
Bogdan Cautis (University Paris-Saclay)  
Bojan Karlas (ETH Zurich)  
Bolin Ding (Data Analytics and Intelligence Lab, Alibaba Group)  
Bolong Zheng (Huazhong University of Science and Technology)  
Bongki Moon (Seoul National University)  
Botong Huang (Alibaba)  
Brad Glasbergen (University of Waterloo)  
Brandon Haynes (Microsoft Gray Systems Lab)  
Cedric Renggli (ETH Zurich)  
Chao Zhang (Lyon 1 University)  
Chen Li (UC Irvine)  
Chengfei Liu (Swinburne University of Technology)  
Chengkai Li (The University of Texas at Arlington)  
Chengliang Chai (Tsinghua University)  
Chong Wang (Amazon)  
Cristian Riveros (PUC Chile)  
Daichi Amagata (Osaka University)  
Dan Kifer (Pennsylvania State University)  
Daniel Deutch (Tel Aviv University)  
Daniel Kang (Stanford University)  
Diego Calvanese (Free University of Bozen)  
Dimitrios Skoutas (Athena Research Center)  
Dimitris Sacharidis (ULB)  
Dirk Habich (TU Dresden)  
Dong Deng (Rutgers University - New Brunswick)  
Dong Wen (University of New South Wales)  
Dong Xie (Pennsylvania State University)  
Dongxiang Zhang (Zhejiang University)  
Dumitrel Loghin (National University of Singapore)  
Elena Ferrari (University of Insubria, Varese)  
Eleni Tzirita Zacharatou (IT University of Copenhagen)  
Essam Mansour (Concordia University)  
Faisal Nawab (University of California at Irvine)  
Fan Zhang (Guangzhou University)  
Fatemeh Nargesian (University of Rochester)  
Fei Chiang (McMaster University)  
Floris Geerts (University of Antwerp)  
Gao Cong (Nanyang Technological University)

George Fakas (Uppsala University)  
George Fletcher (Eindhoven University of Technology the Netherlands)  
George Papadakis (University of Athens)  
George Papastefanatos (Athena Research Center)  
Giovanni Simonini (University of Modena and Reggio Emilia)  
Graham Cormode (University of Warwick)  
Guna Prasaad (Meta Platforms Inc.)  
Guoliang Li (Tsinghua University)  
Guoren Wang (Beijing Institute of Technology)  
Haibo Hu (Hong Kong Polytechnic University)  
Hannes Voigt (Neo4j)  
Haridimos Kondylakis (FORTH-ICS)  
Holger Pirk (Imperial College)  
Huanchen Zhang (Tsinghua University)  
Ibrahim Sabek (MIT)  
Immanuel Trummer (Cornell University)  
Ingo Müller (Google)  
James Cheng (The Chinese University of Hong Kong)  
Jeffrey Xu Yu (The Chinese University of Hong Kong)  
Jens Teubner (TU Dortmund University)  
Jia Yu (Washington State University)  
Jian Lou (Xidian University)  
Jianguo Wang (Purdue University)  
Jianliang Xu (Hong Kong Baptist University)  
Jianxin Li (Deakin University)  
Jiawei Jiang (ETH Zurich)  
Jieming Shi (Hong Kong Polytechnic University)  
Jinfei Liu (Zhejiang University)  
Jing Tang (Hong Kong University of Science and Technology)  
John Liagouris (Boston University)  
John Paparrizos (University of Chicago)  
Joseph Near (University of Vermont)  
Junhao Gan (University of Melbourne)  
K. Selcuk Candan (Arizona State University)  
Kai Wang (University of New South Wales)  
Karima Echiabi (Mohammed VI Polytechnic University)  
Kartik Nayak (Duke University)  
Katja Hose (Aalborg University)  
Kexin Rong (Stanford University)  
Kun Qian (Amazon)  
Kunsoo Park (Seoul National University)  
Kyriakos Mouratidis (Singapore Management University)  
Laks Lakshmanan (University of British Columbia)  
Laurel Orr (Stanford University)  
Lei Cao (MIT)  
Lei Chen (Hong Kong University of Science and Technology)  
Lei Li (Hong Kong University of Science and Technology, Guang Zhou)  
Lijun Chang (The University of Sydney)  
Lin Ma (Carnegie Mellon University)  
Long Yuan (Nanjing University of Science and Technology)  
Lu Qin (UTS)  
Lucas Lersch (Amazon Web Services)  
Lukasz Golab (University of Waterloo)  
Matteo Interlandi (Microsoft)  
Matteo Lissandrini (Aalborg University)  
Matthias Renz (University of Kiel)  
Matthias Weidlich (Humboldt University of Berlin)

Michael Abebe (University of Waterloo)  
 Michael H Boehlen (University of Zurich)  
 Michael Hay (Colgate University/Tumult Labs)  
 Michael Mathioudakis (University of Helsinki)  
 Michal Friedman (ETH)  
 Milos Nikolic (University of Edinburgh)  
 Mirek Riedewald (Northeastern University)  
 Mohamed Sharaf (United Arab Emirates University)  
 Mohammad Sadoghi (University of California, Davis)  
 Mostafa Milani (The University of Western Ontario)  
 Nick Koudas (University of Toronto)  
 Nikolaos Tziavelis (Northeastern University)  
 Nikolay Yakovets (Eindhoven University of Technology)  
 Ning Wang (Beijing Jiaotong University)  
 Oliver A Kennedy (University at Buffalo, SUNY)  
 Panagiotis Bouros (Johannes Gutenberg University Mainz)  
 Panos Kalnis (King Abdullah University of Science and Technology)  
 Panos Vassiliadis (University of Ioannina)  
 Paolo Papotti (EURECOM, France)  
 Periklis Andritsos (University of Toronto)  
 Prashant Pandey (University of Utah)  
 Raghav Kaushik (Microsoft)  
 Rainer Gemulla (University of Mannheim)  
 Raul Castro Fernandez (University of Chicago)  
 Raymond Chi-Wing Wong (Hong Kong University of Science and Technology)  
 Renata Borovica-Gajic (University of Melbourne)  
 Reynold Cheng (The University of Hong Kong, China)  
 Riccardo Torlone (Roma Tre University)  
 Ronghua Li (Beijing Institute of Technology)  
 Ryan C Marcus (MIT)  
 Ryan Stutsman (University of Utah)  
 Sai Wu (Zhejiang Univ)  
 Sairam Gurajada (Apple)  
 Sebastian Link (University of Auckland)  
 Senjuti Basu Roy (New Jersey Institute of Technology)  
 Seokki Lee (University of Cincinnati)  
 Shantanu Sharma (New Jersey Institute of Technology)  
 Shaoxu Song (Tsinghua University)  
 Shiyu Yang (Guangzhou University)  
 Shuai Ma (Beihang University)  
 Sibo Wang (The Chinese University of Hong Kong)  
 Siqiang Luo (Nanyang Technological University)  
 Sourav S Bhowmick (Nanyang Technological University)  
 Spyros Blanas (The Ohio State University)  
 Srikanta Bedathur (IIT Delhi)  
 Stefania Dumbrava (ENSIIE)  
 Stefano Paraboschi (Universita' degli Studi di Bergamo)  
 Steffen Zeuch (DFKI Berlin)  
 Steven E Whang (KAIST)  
 Stijn Vansummeren (Hasselt University)  
 Sudipto Das (Amazon Web Services)  
 Tarique Siddiqui (Microsoft Research)  
 Theodore Dalamagas (Athena Research Center)  
 Thomas Neumann (TU Munich)  
 Tian Li (Carnegie Mellon University)  
 Tianhao Wang (University of Virginia)  
 Tianzheng Wang (Simon Fraser University)  
 Tien Tuan Anh Dinh (Singapore University of Technology and Design)  
 Torben Bach Pedersen (Aalborg University)  
 Utku Sirin (Harvard University)  
 Vasiliki Kalavri (Boston University)  
 Vassilios S Verykios (Hellenic Open University)  
 Walid G Aref (Purdue)  
 Wang-Chiew Tan (Facebook AI)  
 Weiguo Zheng (Fudan University)  
 Wendy Hui Wang (Stevens Institute of Technology)  
 Wentao Wu (Microsoft Research)  
 Xi He (University of Waterloo)  
 Xiang Lian (Kent State University)  
 Xiangmin Zhou (RMIT University)  
 Xiangyao Yu (University of Wisconsin-Madison)  
 Xiaochun Yang (Northeastern University)  
 Xiaofei Zhang (University of Memphis)  
 Xiaoyang Wang (University of New South Wales)  
 Xin Cao (University of New South Wales)  
 Xin Huang (Hong Kong Baptist University)  
 Yan Zhao (Aalborg University)  
 Yang Cao (Kyoto University)  
 Yao Lu (Microsoft Research)  
 Ye Yuan (Beijing Institute of Technology)  
 Yeye He (Microsoft Research)  
 Yi Yu (NII)  
 Yinghui Wu (Case Western Reserve University)  
 Yingxia Shao (BUPT)  
 Yixiang Fang (The Chinese University of Hong Kong, Shenzhen)  
 Yongluan Zhou (University of Copenhagen)  
 You Peng (University of New South Wales)  
 You Wu (Google)  
 Yufei Tao (The Chinese University of Hong Kong)  
 Yuncheng Wu (National University of Singapore)  
 Yuyu Luo (Tsinghua University)  
 Zeke Wang (Zhejiang University)  
 Zhiwei Zhang (Beijing Institute of Technology)  
 Zhongle Xie (Zhejiang University)  
 Zhuoyue Zhao (University at Buffalo - SUNY)  
 Ziawasch Abedjan (Leibniz University Hannover)  
 Zimu Zhou (Singapore Management University)

## LETTER FROM THE EDITORS IN CHIEF

It is our pleasure to present the thirteenth and final issue of PVLDB (Proceedings of the VLDB) Volume 16. This issue includes research papers that will be presented at VLDB 2024 in Guangzhou, China, together with those in PVLDB Volume 17.

PVLDB presents original research papers on a broad range of topics related to all aspects of data and information management, spanning from theoretical foundations, system architectures, models and techniques, to novel applications as well as large-scale deployment and evaluation. There are four equally important categories of papers in the research track: (a) regular research, (b) scalable data science, (c) experiment, analysis & benchmark, and (d) vision. Each paper is evaluated by at least three reviewers and an Associate Editor, who summarizes in a meta-review all reviews and the results of a three-week discussion phase during which the reviewers exchange their viewpoints and converge to a joint decision.

This issue includes 14 papers, spanning the topics of: Data Mining and Analytics; Data Privacy and Security; Database Engines; Database Performance and Manageability; Distributed Database Systems; Graph and Network Data; Information Integration and Data Quality; Languages; Machine Learning, AI, and Databases; Novel Database Architectures (these cover all but two top-level topics of interest defined by our Call for Contributions). The most popular topics in this issue are:

- Database Engines (8 papers);
- Machine Learning, AI, and Databases (5 papers).

The breakdown according to the paper categories is as follows:

- Scalable data science (1 paper);
- Experiment, analysis & benchmark (2 papers);
- Regular research papers (11 papers).

Out of the 14 papers, 10 were accepted after revision, and 4 were accepted after revision plus shepherding.

This issue is the result of all the work put in by the authors as well as the great commitment and effort of our associate editors and reviewers as well as our proceedings chairs.

Georgia Koutrika and Jun Yang  
Editors-in-Chief of PVLDB Vol. 16  
Program Chairs for VLDB 2023