

Matthew Butrovich

Last updated: April, 2024

✉ mbutrovich@gmail.com

☎ (xxx) xxx-xxxx

🌐 mattbutrovi.ch

🔗 github.com/mbutrovich

Summary

Database management systems (DBMS) expert with a deep understanding of high-performance transaction processing—from when queries arrive over the network until their results are durable on disk. Strong background in transactional and analytical DBMS internals (query execution, storage management, concurrency control), Linux internals (eBPF, networking), storage systems (distributed, devices, file systems), performance engineering (benchmarking, observability), and machine learning for systems. Committed to lifelong learning, peer mentorship, and contributing to scientific and open-source communities.

Education

2019 – 2024	Carnegie Mellon University Doctor of Philosophy (Ph.D.), Computer Science <ul style="list-style-type: none">Thesis: “On Embedding Database Management System Logic in Operating Systems via Restricted Programming Environments”Advisor: Andy Pavlo	Pittsburgh, PA
2017 – 2018	Carnegie Mellon University Master of Science (M.Sc.), Computer Science	Pittsburgh, PA
2013 – 2017	University of California, Irvine Bachelor of Science (B.Sc.), <i>cum laude</i> , Computer Science and Engineering	Irvine, CA

Professional Experience

2009 – 2017	Western Digital Corporation Lead Technician V <ul style="list-style-type: none">Led a team of 17 technicians verifying firmware for HDD and SSD storage devices.Managed a regression testing environment supporting over 50 test developers.Wrote Bash scripts to increase test throughput and reduce maintenance overhead.Assisted leadership with design of 10,000 sq. ft. lab containing over 3,000 test systems.Coordinated test efforts across teams in the United States, India, Thailand, and Malaysia.Promoted four times from initial role as Technician II while achieving B.Sc. degree.	Irvine, CA
-------------	---	------------

Selected Publications



2023	Matthew Butrovich , Karthik Ramanathan, John Rollinson, Wan Shen Lim, William Zhang, Justine Sherry, and Andrew Pavlo. “Tigger: A Database Proxy That Bounces With User-Bypass”. <i>Proc. VLDB Endow.</i>
2023	Wan Shen Lim, Matthew Butrovich , William Zhang, Andrew Crotty, Lin Ma, Peijing Xu, Johannes Gehrke, and Andrew Pavlo. “Database Gyms”. <i>CIDR</i> .
2022	Matthew Butrovich , Wan Shen Lim, Lin Ma, John Rollinson, William Zhang, Yu Xia, and Andrew Pavlo. “Tastes Great! Less Filling! High Performance and Accurate Training Data Collection for Self-Driving Database Management Systems”. <i>SIGMOD Conference</i> .
2020	Tianyu Li, Matthew Butrovich , Amadou Ngom, Wan Shen Lim, Wes McKinney, and Andrew Pavlo. “Mainlining Databases: Supporting Fast Transactional Workloads on Universal Columnar Data File Formats”. <i>Proc. VLDB Endow.</i>

Skills




Programming Languages: C, C++17, eBPF, Python, Java

Development Tools: CMake, GDB/LLDB, Git, Google Test, Make, perf, pytest

Projects

- 2024 **BPF-DB**
 ACID-compliant embedded DBMS for kernel-space eBPF applications.
 Role: Sole developer. Wrote eBPF programs in C for storage management (variable length key-value store), transaction management (serializable isolation level), and write-ahead logging. Implemented two sample applications built on BPF-DB: (1) Redis-protocol transactional key-value store that matches state-of-the-art approaches (Dragonfly), and (2) Voter workload comparing serializable transaction performance against VoltDB, achieving 43% higher throughput.
- 2023 **Tigger** 
 PostgreSQL-compatible DBMS proxy based on PgBouncer that bypasses user-space via eBPF.
 Role: Sole developer. Wrote eBPF programs in C that attach at the socket and Traffic Control layers of the Linux networking stack to apply PostgreSQL protocol logic. This design reduced CPU utilization by up to 42%, and lowered transaction latencies by up to 29%.
- 2018 – 2022 **NoisePage** 
 PostgreSQL-compatible in-memory DBMS designed for autonomous operation.
 Role: Founding developer responsible for HyPer-style MVCC over Apache Arrow-compatible memory format, PostgreSQL protocol (Simple and Extended Query), and training data collection. Wrote a C++ version of TPC-C to exercise the storage layer for performance and correctness, ported the Bw-Tree from [Peloton](#), and contributed to system catalog design.

Teaching Experience

- 2022 **15-799: Special Topics in Database Systems** 
 Head teaching assistant
- 2021 **15-445/645: Database Systems** 
 Teaching assistant
- 2020 **15-721: Advanced Database Systems** 
 Head teaching assistant

Academic Service

To the University

- 2022 – 2024 Diversity, Equity and Inclusion Committee, Computer Science Department, *Carnegie Mellon University*.
- 2020 – 2024 Doctoral Review Committee, Computer Science Department, *Carnegie Mellon University*.
- 2019 – 2024 MSCS Admissions Committee, Computer Science Department, *Carnegie Mellon University*.
- 2023 Dean's Welcome Panel, School of Computer Science, *Carnegie Mellon University*.

To the Profession

- 2023 Reviewer, *eBPF Summit Program Committee*. 
- 2023 Reviewer, *ACM SIGMOD Artifacts & Reproducibility Committee*. 
- 2023 Organizer, *ML \rightleftharpoons DB Seminar Series*. 
- 2022 Organizer, *jDatabases! - A Database Seminar Series*. 
- 2022 Organizer, *Vaccination Database Tech Talks - Booster*. 
- 2021 Organizer, *Vaccination Database Tech Talks - Second Dose*. 
- 2021 Organizer, *Vaccination Database Tech Talks - First Dose*. 
- 2020 Organizer, *Quarantine Database Tech Talks*. 