

Lawrence Benson

Email mail@lawben.com

Website lawben.com

Languages German (native), English (native)

Phone +49 157 73001118

LinkedIn linkedin.com/in/lawben

GitHub github.com/lawben

Summary

I am a postdoc database researcher with a strong focus on modern hardware. I am passionate about efficiently leveraging hardware in novel system designs. Currently, I am working on SIMD vectorization and efficient data storage layouts, with multiple published papers at top venues (VLDB, SIGMOD, CIDR). I've supervised various student projects, multiple of which have been published.

Work Experience

- Mar 2024 - present** **Technical University of Munich**, Germany
Postdoctoral Researcher in database group of Viktor Leis.
Working on portable SIMD code with LLVM, efficient B-Tree layouts, and cloud databases.
- Nov 2019 - Nov 2023** **Hasso Plattner Institute**, Potsdam, Germany
Research Associate and head teaching assistant.
Research focus on persistent memory, stream processing engines, and SIMD code.
- Jan 2019 - Oct 2019** **bakdata GmbH**, Berlin, Germany
Working Student Software Engineer
Developed Kafka Streams application to de-duplicate document streams at terabyte-scale.
- Jul 2018 - Oct 2018** **Google Inc.**, New York, NY, USA
Software Engineering Intern
Designed and implemented disaster recovery for petabyte-scale distributed streaming service.
- Oct 2017 - Jul 2018** **Hasso Plattner Institute**, Potsdam, Germany
Student Research Assistant
Developed core functionality of columnar in-memory research database.
- Oct 2016 - Jan 2017** **Google Inc.**, Mountain View, CA, USA
Software Engineering Intern
Extended server fuzz-testing framework with traffic log invariants to increase efficiency.

Education

- Nov 2019 - Nov 2023** **PhD Student Computer Science** - Hasso Plattner Institute, Potsdam, Germany
Focus on data management with modern hardware. Supervisor: Tilmann Rabl
with distinction | Thesis: Efficient State Management with Persistent Memory
- Apr 2017 - Nov 2019** **M.Sc. IT-Systems Engineering** - Hasso Plattner Institute, Potsdam, Germany
with distinction | Focus on scalable systems, databases, and stream processing
- Oct 2013 - Jul 2016** **B.Sc. IT-Systems Engineering** - Hasso Plattner Institute, Potsdam, Germany

Selected Publications (full list at <https://lawben.com>)

- ADMS 2023** **Evaluating SIMD Compiler-Intrinsics for Database Systems**
Lawrence Benson, Richard Ebeling, Tilmann Rabl
- VLDB 2023** **Analyzing Vectorized Hash Tables Across CPU Architectures**
Maximilian Böther, Lawrence Benson, Ana Klimovic, Tilmann Rabl
- VLDB 2022** **PerMA-Bench: Benchmarking Persistent Memory Access**
Lawrence Benson, Leon Papke, Tilmann Rabl
- CIDR 2022** **Darwin: Scale-In Stream Processing**
Lawrence Benson, Tilmann Rabl
- VLDB 2021** **Viper: An Efficient Hybrid PMem-DRAM Key-Value Store**
Lawrence Benson, Hendrik Makait, Tilmann Rabl
- SIGMOD 2021** **Maximizing Persistent Memory Bandwidth Utilization for OLAP Workloads**
Björn Daase, Lars Jonas Bollmeier, Lawrence Benson, Tilmann Rabl

Open Source Contributions (selection, full overview at github.com/lawben)

- 2023/24** **LLVM Compiler Project** - Acquired Push/Merge Rights
Notable Contributions (focused on SIMD code):
- Introduced new `@llvm.vector.compress` intrinsic (#92289)
 - Introduced `__builtin_vectorelements()` function in Clang (#69010)
 - Improved instruction selection for SIMD vector shuffles on ARM (#146212)
 - Improved instruction selection for SIMD vector to scalar bitcast on ARM (#145301)

Teaching & Supervision (selection, full list at lawben.com)

- Theses**
- **A Data Generator for Realtime Streaming Analytics Benchmarks** (Master Thesis, 2023)
 - **Query-Sharing in Query-Compiled Stream Processing Engines** (Master Thesis, 2023)
 - **Column-Oriented Stream Processing** (Master Thesis, 2022)
 - **Efficient Network-aware Data Stream Aggregation** (Master Thesis, 2022)
 - **Designing a CPU-aware Hash Table for Streaming Hash Joins** (Master Thesis, 2022)
 - **Efficient Data Ingestion for Stream Processing Systems**, (Master Thesis, 2022)
 - **R-Tree Data Placement on Persistent Memory** (Master Thesis, 2021)
- Seminars**
- **Hardware-Conscious Data Processing** (2024)
 - **Data Management on Modern Storage Technologies** (2021)
 - **Open Source Data Processing** (2021)
 - **Data Processing on Modern Hardware** (2020)
- Lectures**
- **Hardware-Conscious Data Processing** (2023, 2022)
 - **Big Data Systems** (2021, 2019)
 - **Database Systems I** (2020)

Invited Talks

- 2024** - Compiler SIMD @ TUMuchData
- 2023** - Compiler SIMD @ LLVM Meetup Berlin, TU München, Simon Fraser University
- Hardware-Conscious Databases @ Oracle
- 2022** - PerMA-Bench @ TU München, FAU Erlangen, TU Darmstadt, Intel
- Darwin @ TU Berlin
- 2021** - Persistent Memory in Databases @ ITU Copenhagen
- Viper @ Intel, TU Darmstadt

Program Committee Member / Reviewer

- 2024** PVLDB 18, DaMoN, VLDB Journal, ADMS, DEBS
- 2023** PVLDB 17, VLDB Journal
- 2022** Datenbankspektrum