

BENJAMIN LAMOTHE

Software Engineer and System Architect: Databases and Distributed Systems

@ blamothe@alum.mit.edu

📞 917-601-6567

🌐 zonotope

📍 Brooklyn, NY

19 years building databases and distributed systems, specializing in query engines and optimization, data storage and indexing, and consensus protocols. Most recently, hybrid graph, vector, and full-text retrieval for AI workloads. Twice a venture-backed co-founder (YC S12). B.S. in theoretical mathematics from MIT.

EXPERIENCE

Chief Architect

Fluree

📅 January 2024 – Present

📍 Brooklyn, NY (Remote)

- Set the architectural direction of the Fluree Core Database, a distributed system with independently scaled query peers
- Built a federated query layer presenting graph, vector, and full-text indexes across heterogeneous local and remote storage back ends as a single logical database, providing the hybrid retrieval substrate for AI workloads
- Designed a cost-based query optimizer that reorders and rewrites user queries using cardinality estimates derived from index statistics
- Owned query execution, optimization, and indexing for a ground-up rewrite of the database in Rust, undertaken for throughput and fine-grained memory control
- Added Raft-based replication on openraft, implementing the log storage, state machine, and snapshot layers against the engine's storage
- Restructured the system into a modular, fault-tolerant architecture separating storage, indexing, and query execution

Technical Lead, Core Database

Fluree

📅 September 2022 – January 2024

📍 Brooklyn, NY (Remote)

- Led a team of four database engineers with mentorship through weekly 1 on 1 meetings and code review
- Built the core query processing engine and declarative query language on top of the streaming index layer, using transducers to process datasets larger than available memory

Senior Software Engineer

Fluree

📅 September 2020 – September 2022

📍 Brooklyn, NY (Remote)

- Designed and implemented a streaming re-indexing system that merges persisted indexes with in-memory novelty in constant memory, independent of database size
- Wrote an SQL parser and transpiler allowing users already familiar with SQL to query Fluree databases

Co-founder and CTO

Quorum Control

📅 April 2018 – September 2020

📍 Brooklyn, NY (Remote)

- Led a fully remote development team building a distributed platform using a novel Byzantine fault tolerant consensus algorithm that supported a throughput of 200 transactions per second with 144 millisecond finality on 21 nodes
- Wrote SDKs in Go, JavaScript, and WebAssembly allowing third party developers to build decentralized applications on the platform

EDUCATION

B.S. in Mathematics (Theoretical)

Massachusetts Institute of Technology

📅 June 2007

📍 Cambridge, MA

AWARDS



Y Combinator Founder

Summer 2012



Forbes 30 under 30

2012 and 2013

SYSTEMS

Query Optimization Storage Engines

Indexing Distributed Consensus

Vector Search Full-Text Search

Knowledge Graphs Streaming

SQL

AWS

Hadoop

Vertica

WebAssembly

CODE IN PRODUCTION

📦 Clojure

🐹 Go

☕ Java

🦀 Rust

📄 JavaScript

🔗 Ruby

EXPERIENCE (CONTINUED)

Senior Software Engineer

DemocracyWorks

📅 January 2017 – April 2018

📍 Brooklyn, NY (Remote)

- Built Clojure micro-services to send millions of user notifications for United States voter registration and election information on peak days
 - Mentored multiple junior developers on the team through pair programming, code review, and regular 1 on 1 discussions
-

Independent Software Contractor

📅 August 2014 – January 2017

📍 New York, NY

- Migrated a legacy CakePHP healthcare web application to Ruby on Rails, then built and managed a HIPAA compliant and fault-tolerant AWS infrastructure for the new application (Spreemo)
 - Performed code quality audits and security reviews across multiple client engagements
 - Developed a declarative AWS deployment infrastructure (Vitals)
 - Designed a candidate interview process and recruited two full-time engineers (Spremo)
-

Co-Founder and Lead Software Engineer

Amicus (Y Combinator S12)

📅 February 2012 – July 2014

📍 New York, NY

- Developed a Ruby on Rails and Backbone.js web application for non-profits and political organizations to better engage their supporters
 - Developed a JRuby Object-Graph Mapping for property graph databases in Rails applications using the TinkerPop API
 - Built a Ruby console that allowed for fully automated deployment and provisioning of Amazon EC2 instances.
-

Software Engineer

Visible Measures

📅 September 2007 – March 2011

📍 Boston, MA

- Developed and maintained a high-performance, scalable, and distributed analytics engine with Hadoop on 60+ node cloud compute clusters to process dozens of Terabytes of video stream viewership data daily
 - Implemented and managed a sharded data warehouse based on MySQL and later Vertica database management systems
-

Research Assistant

M.I.T. Computer Science and Artificial Intelligence Laboratory

📅 January 2006 – May 2007

📍 Cambridge, MA

- Worked under Professor Patrick Winston to build a natural language understanding system that parsed English sentences into syntactic representations, extracted structured knowledge through pattern recognition, and generated visual animations of described events
- Surveyed knowledge representation schemes in the Cognitive Linguistics literature and developed novel knowledge representations through analysis of English sentence structure