

Jonathan C. Beard

Captain Dr. ``Whittle-Beard" • soldier → scientist → silicon architect

System Performance Architect & Technical Lead / Manager • Google • Austin & Sunnyvale

jonathan.c.beard@gmail.com | jonathanbeard.io | github.com/jonathan-beard | linkedin.com/in/jonathanbeard | Google Scholar

A hardware-architecture and systems-performance leader with **18+ years** spanning the U.S. Army, academia, and hyperscale silicon. I turn architecture, memory, interconnect, power, cost, and real workload behavior into data-driven platform decisions — and I like building the teams that make those calls stick.

BY THE NUMBERS

18+ years in systems & silicon

29 granted U.S. patents

26 academic publications

40 engineers led at peak

3 university degrees

EDUCATION

Ph.D., Computer Science

Washington University in St. Louis · 2015

M.S., Bioinformatics

Johns Hopkins University · 2010

B.S. Biological Sciences

B.A. International Studies

Louisiana State University · 2005

WHAT I WORK ON

AI accelerators; CPU/GPU microarchitecture; memory systems (HBM/DDR, CXL, disaggregation); interconnects; performance modeling & benchmarking; hardware/software co-design; TCO.

Tools: C/C++, Python, Go, CUDA, gem5, SST, Verilog, Linux/eBPF/perf.

OPEN SOURCE

RaftLib — the C++ stream-processing DSL I authored and still maintain. ≈969 GitHub stars and front page of Hacker News *twice*. raftlib.io

NOW --- GOOGLE (2022--PRESENT)

- Own **system-wide Perf/TCO** for Google's internal SoC program (gSoC) across multiple silicon generations and product lines.
- Lead the **New Platform Introduction Performance** team assessing the forward-looking Arm-based platform roadmap.
- Drive ML-guided SoC tuning and Google-specific accelerator IP; manage and technically lead orgs of up to 40 engineers.

THE ROAD HERE

- **Arm (2015–2022), to Principal System Architect.** CHI memory-copy enhancements, the **LS64** architecture + gather-hint instruction, Project-38 (joint DoD/DoE late-Moore research), GenZ/CXL subcommittee; funded 12+ students at Barcelona Supercomputing Center, UT Austin, and Georgia Tech.
- **FastData.io** — startup technical advisor.
- **Washington University** — PhD research in the Stream-Based Supercomputing Lab.
- **U.S. Army veteran** — 2001–2010.

RESEARCH & PHD THESIS

- **Thesis:** *Online Modeling and Tuning of Parallel Stream Processing Systems* (2015) — queueing theory, machine learning, and control theory to auto-tune dataflow systems. It grew into RaftLib.
- **25+ papers** across IJHPCA, IPDPS, ICPP, PACT, ICS, Euro-Par, MEMSYS, and HPEC.

COMMUNITY & SPEAKING

- Supercomputing (SC) Technical Program Committee 2016–2022, **Architecture & Networks Co-Chair (SC19)**; TPC for ISCA, MICRO, and HPCA; ICPP Architecture Track Co-Chair; **General Chair, Arm Research Summit 2019**.
- Talks & panels at C++Now (RaftLib), MEMSYS, and SC; guest on **CppCast** (Ep. 50).

Ask me about: stream processing & post-Moore architecture, the leap from Army officer to silicon architect --- or, yes, the beard.