Axel Feldmann

(917)-691-5910 • axelf@csail.mit.edu • feldmann.nyc

GOAL

Looking for computer architecture and/or performance engineering roles in industry. I'm specifically interested in roles that involve working across the hardware/software interface.

EDUCATION

Massachusetts Institute of Technology, Cambridge, Massachusetts
Ph.D. in Electrical Engineering and Computer Science — *anticipated* in December 2024
- Advisor: Daniel Sanchez

S.M. in Electrical Engineering and Computer Science — May 2021
- GPA: 5.0 / 5.0

Carnegie Mellon University, Pittsburgh, Pennsylvania
BS in Computer Science — May 2019
- GPA: 3.9 / 4.0

Selected coursework: computer architecture, operating systems, compilers, algorithms, computer networks, numerical algorithms

SKILLS

Computer architecture research, hardware prototyping/development, parallel programming models, analytical and simulation-based performance modeling, hardware-software codesign.
C++, C, Python, CUDA, x86/64 assembly, PyTorch, Triton, Pandas, Pallas, Linux/Unix

HONORS & AWARDS

Mathworks Fellowship — 2024
IEEE Micro Top Picks — 2022

EXPERIENCE

Research Assistant — September 2019 – present
MIT CSAIL - Cambridge, MA
- Currently working on designing hardware accelerators for sparse linear algebra, specifically focused on scientific computing.
- Designed *Spatula* and *Azul*, hardware accelerators for sparse linear solvers that outperform state-of-the-art GPU baselines by gmean $47\times$ and gmean $242\times$ respectively.
  - Python and C++ cycle-level simulation, Design Compiler area and power estimates, Xilinx FPGA prototyping
- Contributed custom high-performance GPU kernels to multiple collaborators working on machine learning.
  - CUDA, Triton, Pallas, PyTorch
- Worked on designing hardware accelerators for Fully Homomorphic Encryption (FHE), a computationally expensive cryptographic system that enables arithmetic operations on encrypted data.
- Designed *F1*, the first proposed ASIC accelerator capable of executing entire FHE programs. *F1* was selected for IEEE Micro Top Picks 2022.
- Developed compilation techniques that enabled *F1* and *CraterLake* (our second-generation accelerator design) to achieve $5000\times$ speedups over CPU baselines.
  - Python and C++ compiler stack, C++ cycle-level simulator, Design Compiler area and power estimates

Intern — Summer 2022
Cerebras Systems - Sunnyvale, CA
- Worked on the CSL (Cerebras Systems Language) team to designing new abstractions for programming the Cerebras Wafer Scale Engine.
- Designed and built a complete MPI-style collectives library for CSL programmers.
  - Project entirely in CSL, syntactically similar to Zig

Undergraduate Research Assistant — 2018 – 2019
Computer Organization Research Group - Pittsburgh, PA
- Worked with Prof. Nathan Beckmann on *Livia*, a system architecture for data centric computing.
- Created zsim-based simulation infrastructure (C++) in to evaluate our proposed architecture.
- Wrote applications to effectively utilize *Livia*'s novel hardware features.

Systems Software Intern — Summer 2018
NVIDIA - Santa Clara, CA
- Improved display driver performance for existing and upcoming Tegra SoCs.
- Reduced kernel test time by 30% via improved thread synchronization.

Software Engineering Intern Summer 2017
Yahoo, Flurry Analytics - Sunnyvale, CA
- Created webapp to help users design metrics API queries.
- Re-engineering User Acquisition Analysis (UAA) features on the Flurry data platform.

PAPERS

Azul: An Accelerator for Sparse Iterative Solvers Leveraging Distributed On-Chip Memory
Axel Feldmann, Courtney Golden, Yifan Yang, Joel S. Emer, Daniel Sanchez, *to appear at MICRO 2024*.

MiniFold: Simple, Fast and Accurate Protein Structure Prediction
Jeremy Wohlwend, Mateo Reveiz, Matt McPartlon, **Axel Feldmann**, Wengong Jin, Regina Barzilay, *under submission to Neurips 2024*.

FeatUp: A model-agnostic framework for features at any resolution
Stephanie Fu, Mark Hamilton, Laura Brandt, **Axel Feldmann**, Zhoutong Zhang, and William T. Freeman, *ICLR 2024*.

DsDm: Model-Aware Dataset Selection with Datamodels
Logan Engstrom, **Axel Feldmann**, and Aleksander Madry, *ICML 2024*.

Spatula: A Hardware Accelerator for Sparse Matrix Factorization
Axel Feldmann, Daniel Sanchez, *MICRO 2023*.

An Architecture to Accelerate Computation on Encrypted Data
Axel Feldmann*, Nikola Samardzic*, Aleksander Krastev, Srinivas Devadas, Ronald Dreslinski, Chris Peikert, Daniel Sanchez, *IEEE Micro Top Picks 2022*.

CraterLake: A Hardware Accelerator for Efficient Unbounded Computation on Encrypted Data
Nikola Samardzic, **Axel Feldmann**, Aleksander Krastev, Nathan Manohar, Nicholas Genise, Srinivas Devadas, Chris Peikert, Daniel Sanchez, *ISCA 2022*.

F1: A Fast and Programmable Accelerator for Fully Homomorphic Encryption
Axel Feldmann*, Nikola Samardzic*, Aleksander Krastev, Srinivas Devadas, Ron Dreslinski, Chris Peikert, Daniel Sanchez, *MICRO 2021*.

Livia: Data-Centric Computing Throughout the Memory Hierarchy
Eliot Lockerman, **Axel Feldmann**, Mohammad Bakhshalipour, Alexandru Stanescu, Shashwat Gupta, Daniel Sanchez, and Nathan Beckmann, *ASPLOS 2020*.

TEACHING ASSISTANT

MIT 6.5900: Computer Systems Architecture	Fall 2024 (ongoing)
CMU 15-410: Operating System Design and Implementation	2018-2019
CMU 15-213: Introduction to Computer Systems	2017-2018

[*Resume compiled on 2024-08-19*]