

## Objective

Full-time research or development position working on compilers or low-level systems. I seek to build performant and robust platforms at the bottom of the software stack, because (1) it's critical infrastructure that must be done well, and (2) that's where the fun is.

## Full-time Experience

**IBM Research** Yorktown Heights, NY 2023.06 - Present

*Research scientist*

- Compiler for AI accelerator - graph optimizations
- Confidential computing
- 1 pending patent application (20230195860)

**Pindrop** Atlanta, GA 2017.01 - 2021.02

*Voice-based fraud detection, VC-funded start-up*

*Senior software engineer*

- Designed, led development, and delivered a messaging layer in C for the company's flagship product
- Adopted and righted underperforming in-memory service that collates call details and predicted features from numerous sources
- Guided requirements and design for product-wide logging and team-wide CI system
- Wrote a new identity access management system to support single sign-on and replace the original, insecure version for the cloud product
- Mentored summer intern's development of much-needed, rudimentary in-house libraries
- Contributed to a larger team effort for constructing a new platform for our product: ripping out vagrant; unifying a build system hodgepodge; repackaging deps as rpms; moving all development, build, and product environments onto Red Hat; decoupling services; etc.

**Qualcomm Research** San Diego, CA 2014.07 - 2016.12

*Senior software engineer*

- Co-pioneered Qualcomm's first profile-guided optimization for the modem, improving L2I\$ miss rate by 33%, branch target prediction by 36%, memory b/w by 25%, and cycles per instruction packet by 5%
- Optimized the modem's runtime compression algorithms for speed and size, implementing techniques such as partial hashing and zero run-length encoding; prototyped in C and hand-crafted in assembly; team contributions led to 5% size improvement, 1.9x compression speedup, and 8x decompression speedup
- Co-led investigations into clock boost latency, identifying 25% speedup on critical paths
- Boosted modem throughput by contributing to an asynchronous compression library
- Collaborated to shrink modem image in first 3 months by 60% more than prior MB/month reduction avg
- Contributed to the machine learning group's tooling, including CMake infrastructure and CI/CD system
- Co-authored the 4th-place (top 10%) software paper in Qualcomm's internal "2016 QTech Paper Series"
- Earned 3 employee recognition awards in first 3 review periods; promoted in under 30 months
- New-hire mentor
- 3 patents (10198362, 10169246, 10061698)

## Other Work Experience

- **IBM Research** *Intern* Yorktown Heights, NY 2021.05 - 2021.08
- **Georgia Institute of Technology** *Teaching assistant* Atlanta, GA, 2011.01 - 2019.12
- **Qualcomm** *Intern, Software engineer* San Diego, CA, 2013.05 - 2013.08
- **Doubleshot Technologies** *Software engineer* Atlanta, GA, 2013.01 - 2013.05
- **Georgia Tech Compute for Good Lab** *Software engineer* Atlanta, GA, 2012.05 - 2012.09
- **Georgia Tech Research Institute** *Research assistant* Atlanta, GA, 2012.01 - 2012.08

## Education

**Georgia Institute of Technology** Atlanta, GA 2017.08 - 2023.05

*PhD in Computer Science*

*Presidential Fellowship*

- Compiler and ML techniques for on-the-fly call set prediction and software debloating
- Multi-GPU and many-core scheduling via compiler assistance

**Georgia Institute of Technology** Atlanta, GA 2010.05 - 2014.05

*BS and MS in Computer Science*

*Highest Honors*

- Master's project: Compiler-assisted prediction for VM scheduling
- Systems, architecture, and networking

## Publications & Conferences

- **Tackling ML-based Dynamic Mispredictions using Statically Computed Invariants for Attack Surface Reduction.** Porter, Khan, Ni, Pande. *Under review.*
- **From Confidential Computing to Zero Trust, Come Along for The (Bumpy?) Ride.** Ye, Koteshwara, Dunn, Franke, Porter, Feldman-Fitzthum, Ruocco, Buono, Carvalho. *HASP@MICRO 2024.*
- **Going Beyond Confidential Attestation with Trustee.** Feldman-Fitzthum, Carvalho, Porter, Dubey, Buono. *Confidential Computing @LPC 2024.*
- **Aligning Confidential Computing with Cloud-native ML Platforms.** Ruocco, Porter, Carvalho, Buono, Dunn, Franke, Bottomley, Silva, Ye, Dubey, Feldman-Fitzthum. *AISCC@NDSS 2024.*
- **Beacons: End-to-End Compiler Framework for Predicting and Utilizing Dynamic Loop Characteristics.** Mururu, Khan, Chatterjee, Chen, Porter, Gavrilovska, Pande. *OOPSLA 2023.*
- **Decker: Attack Surface Reduction via On-demand Code Mapping.** Porter, Khan, Pande. *ASPLOS 2023.*
- **Compiler-Assisted Scheduling for Multi-Instance GPUs.** Porter, Chen, Pande. *GPGPU 2022.*
- **CASE: A Compiler-Assisted Scheduling Framework for Multi-GPU Systems.** Chen, Porter, Pande. *PPoPP 2022.*
- **BlankIt library debloating: getting what you want instead of cutting what you don't.** Porter, Mururu, Barua, Pande. *PLDI 2020.*

## Toolset

- **Languages:** C/++, Python, Rust, Bash, Q6 assembly
- **Platforms, frameworks, libs:** Linux, LLVM, gcc, gdb, Make, git, pthreads, valgrind, Pin, Nebula VPN, perf, SNP/TDX, CMake, CUDA, AWS, OpenMP, MPI, KVM/QEMU, Kata, Kubernetes, node.js, Yii, Android, Apache, nginx, Flask

## Coursework

Algorithms, Combinatorics, Compilers, Distributed computing, Embedded systems, High performance computer architecture, High performance parallel computing, Theory of computation, Data structures, Databases, Operating systems, Processor design, Programming languages, Networking