Giorgis Georgakoudis

Principal Compiler Engineer | LLVM/MLIR, OpenMP, GPU Runtime Systems

Pleasanton, CA | georgakoudis1@llnl.gov | +1 925-495-9544 | ggeorgakoudis | ggeorgakoudis

Summary

Principal compiler and systems engineer building **LLVM/MLIR**, **OpenMP**, and **GPU runtime** infrastructure for production scientific computing. Lead developer of **Proteus** and co-lead of **PyOMP**, with hands-on ownership of **CUDA/HIP** code generation, JIT optimization, Python compiler integration through **Numba**, and OpenMP standards leadership.

Core Skills

Compilers: LLVM/Clang, **MLIR**, code generation, **JIT optimization**, compiler-guided specialization

Runtime Systems: OpenMP, **GPU offload**, runtime systems, **performance portability**, autotuning

Languages: C++, C, Python, **CUDA**, **HIP**

Systems: Linux, profiling, benchmarking, **performance analysis**, scientific software

Experience

Principal Computer Scientist, Lawrence Livermore National Laboratory – Livermore, CA — 2018 – present (7 years 5 months)

Lead compiler, runtime, and performance engineering efforts for large-scale scientific software and GPU-accelerated HPC applications.

- Lead development of **Proteus**, a programmable **LLVM/MLIR** JIT optimization framework for host and GPU kernels targeting **CUDA** and **HIP**, used in production HPC physics applications.
- Co-lead **PyOMP**, an OpenMP-for-Python system built on **Numba**, extending parallel and offload support in a major Python compiler ecosystem.
- Contribute to **Mneme**, a GPU kernel record/replay and autotuning workflow built around LLVM IR artifacts and Proteus-based optimization experiments.
- Contributed to OpenMP compilation and the OpenMP offloading runtime.
- Serve as technical lead for a **$2.7M / 3-year** just-in-time optimization program for HPC applications.
- Chair the **OpenMP Python Language Subcommittee**, helping shape a de-facto industry standard for parallel programming and Python-facing OpenMP support.

Selected Technical Publications

Proteus: Portable Runtime Optimization of GPU Kernel Execution with Just-In-Time Compilation — 2025
G. Georgakoudis, K. Parasyris, D. Beckingsale
10.1145/3696443.3708939 (CGO 2025)

Programming GPUs with OpenMP and Python — 2025
G. Georgakoudis, T. A. Anderson, S. Archibald, B. R. de Supinski, T. G. Mattson
10.1007/978-3-032-06343-4_14 (IWOMP 2025, Best Paper Award)

Recognition

- Best paper awards (IWOMP 2025, IWOMP 2021, IWOMP 2020, DATE 2020, COSMIC 2013)
- LLNL Director's Excellence in Publication (2024); LLNL Achievement Award (2024, 2020)
- SC'23 Best Paper Finalist; SC'21 Best Reproducibility Award

Education

University of Thessaly, PhD in Computer Engineering – Greece — 2017