Anthony Cabrera

HETEROGENEOUS COMPUTING RESEARCH SCIENTIST

✉ cabreraam AT ieee DOT org | ⌂ cabreraam.github.io

Summary

Experienced heterogeneous computing researcher and lifelong learner with a penchant for exploring questions around the hardware-software interface. Adept at quickly learning skills, languages, or tools necessary to answer research questions. Strives to communicate complex ideas effectively and accessibly.

Research and Work Experience

Architectures and Performance Group @ Oak Ridge National Laboratory
Oak Ridge, TN (Remote)

RESEARCH SCIENTIST — *Jan 2022 - Present*
SOFTWARE ENGINEER — *Aug 2020 - December 2021*

- Created an MLIR pass to enable qubit-wise commutativity optimizations in the Xanadu Catalyst MLIR quantum compiler
- Deployed TF/PyTorch models on AMD/Xilinx Versal ACAP VCK190 platform for scientific application edge computing
- Upstreamed MLIR dialect and lowering features for Flang – the LLVM front-end for Fortran
- Led a multi-institution performance and portability evaluation comparing Intel and Xilinx FPGA OpenCL kernels
- Developed Hexagon DSP kernels for the Qualcomm Snapdragon chip, as part of DARPA's DSSoC project
- Created performance models for CXL-based GPU-FPGA collaboration on HPC mini applications as part of the DoE Exascale Compute Project

Computer Science and Engineering Department @ Washington University in St. Louis
St. Louis, MO

ADJUNCT INSTRUCTOR — *August 2024 - Present*
VISITING RESEARCH SCIENTIST — *May 2022 - Present*

- Teaching a graduate level course called Accelerating Algorithms in Reconfigurable Logic
- Led seminar presentations and contributed to research proposals in the areas of hardware and compilers

Stream Based Supercomputing Laboratory @ Washington University in St. Louis
St. Louis, MO

GRADUATE RESEARCH ASSISTANT — *July 2016 - July 2020*

- Thesis: Domain Specific Computing in Tightly Coupled Heterogeneous Systems
- Evaluated the Intel HARPv2 CPU+FPGA platform as a domain specific compute solution
- Created a benchmark suite of data integration applications (DIBS) to identify opportunities for hardware acceleration
- Ported Needleman-Wunsch OpenCL Kernels to the Intel HARPv2 CPU+FPGA platform to analyze kernel design, performance, and portability
- Architected and optimized hardware for DIBS applications using OpenCL targeting the Intel HARPv2

The MITRE Corporation
Shiloh, IL

GRADUATE PROTOTYPING AND SOFTWARE ENGINEER — *May 2019 - August 2019*

- Created a neural network to detect cars from wireless iPhone camera stream targeting the NVIDIA Jetson Nano
- Deployed containers on GPU-enabled HPC resources to train convolutional neural networks
- Mentored undergraduate intern project on hyperparameter performance analysis
- Maintained GitLab repository to document work and enable continued development of project
- Selected as one of four interns across all of MITRE's sites to deliver company-wide presentation on project

Arm Holdings
Austin, TX

GRADUATE RESEARCH INTERN — *May. 2018 - Aug. 2018*

- Quantified spatial and temporal locality by creating a novel technique based on reuse distance
- Developed dynamic binary instrumentation clients to profile memory subsystem characteristics
- Identified strategies around data layout transformations and paging to improve memory subsystem performance

Advanced Sensors Research Laboratory @ Washington University in St. Louis
St. Louis, MO

UNDERGRADUATE AND GRADUATE RESEARCH ASSISTANT — *May 2014 - July 2016*

- Developed the software/UI for filter alignment of NIR fluorescence imagers
- Assisted medical researchers with NIR fluorescence and polarization imaging studies
- Aided in the design and fabrication of a custom PCB around an ultra low-noise imaging sensor

Education

Washington University in St. Louis   *St. Louis, MO*
PHD COMPUTER ENGINEERING   *August 2020*
MS COMPUTER SCIENCE   *August 2018*
BS COMPUTER ENGINEERING, SECOND MAJOR COMPUTER SCIENCE   *May 2015*
BSAS ELECTRICAL ENGINEERING

Hendrix College   *Conway, AR*
BA CHEMICAL PHYSICS, MINOR MUSIC   *May 2013*

Peer Reviewed Publications (Chronological)

1. **AM Cabrera**, S Afrose, D Claudino, TS Humble "Toward Exploiting Qubit-Wise Commutativity Using an MLIR Approach", *In Preparation*.

2. **AM Cabrera**, DE Bernholdt, C Zimmer "A Flang Plugin for Fortran Feature Characterization", Accepted for publication in *IEEE/ACM Eighth Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC '24)*.

3. E Wong, VL Ortega, D Claudino, S Johnson, S Afrose, M Gowrishankar, **AM Cabrera**, TS Humble "A Cross-Platform Execution Engine for the Quantum Intermediate Representation", *arXiv preprint arXiv:2404.14299*.

4. **AM Cabrera**, YA Yucesan, FY Liu, W Blokland, JS Vetter "Errant Beam Detection Using the AMD Versal ACAP and Vitis AI", *IEEE High Performance Extreme Computing Conference (HPEC '23)*.

5. CJ Faber, SD Harris, Z Xiao, RD Chamberlain, **AM Cabrera** "Challenges Designing for FPGAs Using High-Level Synthesis", *IEEE High Performance Extreme Computing Conference (HPEC '22)*.

6. NR Miniskar, AR Young, FY Liu, **AM Cabrera**, JS Vetter, "Ultra Low Latency Machine Learning for Scientific Edge Applications", *IEEE International Conference on Field Programmable Logic and Applications (FPL '22)*.

7. AR Young*, **AM Cabrera***, JS Vetter, "Design and Analysis of CXL Performance Models for Tightly-Coupled Heterogeneous Computing", *ACM International Workshop on Extreme Heterogeneity Solutions (ExHET '22 @ PPoPP '22)*.

8. CJ Faber, T Plano, S Kodali, Z Xiao, A Dwaraki, JD Buhler, RD Chamberlain, **AM Cabrera**, "Platform Agnostic Streaming Data Application Performance Models", *ACM/IEEE Redefining Scalability for Diversely Heterogeneous Architectures (RSDHA '21 @ SC '21)*.

9. Zhili Xiao, RD Chamberlain, **AM Cabrera**. "HLS Portability from Intel to Xilinx: A Case Study", *IEEE High Performance Extreme Computing Conference (HPEC '21)*. [Paper] [Slides]

10. **AM Cabrera**, S Hitefield, J Kim, S Lee, NR Miniskar, JS Vetter, "Toward Performance Portable Programming for Heterogeneous System-on-Chips: Case Study with Qualcomm Snapdragon SoC", *IEEE High Performance Extreme Computing Conference (HPEC '21)*. [Paper] [Slides]

11. **AM Cabrera**, AR Young, J Lambert, Z Xiao, A An, S Lee, Z Jin, J Kim, J Buhler, RD Chamberlain, JS Vetter, "Toward Evaluating High-Level Synthesis Portability and Performance between Intel and Xilinx FPGAs", *ACM International Workshop on OpenCL (IWOCL '21)*. [Paper] [Slides] [Video]

12. **AM Cabrera**, RD Chamberlain, "Design and Performance Evaluation of Optimizations for OpenCL FPGA Kernels", *IEEE High Performance Extreme Computing Conference (HPEC '20)*. [Paper] [Slides]

13. **AM Cabrera**, RD Chamberlain, "Designing Domain Specific Computing Systems", *IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM '20)*. [Paper] [Slides] [Video]

14. **AM Cabrera**, RD Chamberlain, JC Beard, "Multi-spectral Reuse Distance: Divining Spatial Information from Temporal Data", *IEEE High Performance Extreme Computing Conference (HPEC '19)*. [Paper] [Slides]

15. **AM Cabrera**, RD Chamberlain, "Exploring Portability and Performance of OpenCL FPGA Kernels on Intel HARPv2", *ACM International Workshop on OpenCL (IWOCL '19)* **Best Presentation Award**. [Paper] [Slides]

16. CJ Faber, **AM Cabrera**, O Booker, G Maayan, RD Chamberlain, "Data Integration Tasks on Heterogeneous Systems Using OpenCL", *ACM International Workshop on OpenCL (IWOCL '19)*. [Paper]

17. **AM Cabrera**, CJ Faber, K Cepeda, R Derber, C Epstein, J Zheng, RK Cytron, RD Chamberlain, "DIBS: A Data Integration Benchmark Suite", *ACM/SPEC International Conference on Performance Engineering (ICPE '18)*. [Paper] [Slides]

*Denotes equal contribution.

PhD Dissertation

Domain Specific Computing in Tightly-Coupled Heterogeneous Systems [Text] [Slides]

Honors and Awards

2020	**SC20 Early Career Program**, Supercomputing 2020	*Atlanta, GA*
2020	**Honors Designation for PhD Progress Review (Top 15-20% of students)**, CSE Department @ WUSTL	*St. Louis, MO*
2020	**Engineering PhD Student Commencement Marshal**, WUSTL	*St. Louis, MO*
2019	**Best Presentation Award**, International Workshop on OpenCL	*Boston, MA*
2019	**Graduate Student Ambassador**, Intel Corporation	*St. Louis, MO*
2019	**Travel Grant**, Supercomputing 2019	*Denver, CO*
2018	**Travel Grant**, Supercomputing 2018	*Dallas, TX*
2017	**Travel Grant**, Supercomputing 2017	*Denver, CO*
2015	**Graduate Danforth Scholar**, WUSTL	*St. Louis, MO*
2013	**Harold P. Brown Engineering Fellowship**, McKelvey School of Engineering @ WUSTL	*St. Louis, MO*
2013	**Hendrix College Chamber Orchestra Award**, Hendrix College	*Conway, AR*
2012	**Transamerica Employer Solution & Pension Scholarship Award**, Transamerica Corporation	*Little Rock, AR*
2011	**Hendrix College Chamber Orchestra Award**, Hendrix College	*Conway, AR*

Teaching Experience

FL24	**WUSTL CSE 565M Accelerating Algorithms in Reconfigurable Logic**, Instructor of Record	*St. Louis, MO*
FL17, FL18	**WUSTL CSE 560M Computer Systems Architecture I**, Graduate Teaching Assistant	*St. Louis, MO*
SU18	**WUSTL CSE 566S High Performance Computing**, Graduate Teaching Assistant	
SP16	**WUSTL CSE {4,5}63M Digital Integrated Circuit Design and Architecture**, Graduate Teaching Assistant	
FL14, SP15	**WUSTL CSE 200 Scientific Computing**, Undergraduate Teaching Assistant	

Press

2021	**SCTV Interview**, SC21 Inclusion and Diversity with AJ Lauer and Anthony Cabrera	*St. Louis*
2021	**SCTV Interview**, Promoting the SCALE students program at SC	*Virtual*

Professional Service

2022	**SC22 Inclusivity Committee**, Supercomputing 2022	*Dallas, TX*
2022	**Program Committee**, International Workshop on OpenCL and SYCLcon	*Remote*
2021	**SC21 Inclusivity Committee**, Supercomputing 2021	*St. Louis, MO*
2021	**Program Committee**, International Workshop on OpenCL and SYCLcon	*Remote*
2019	**Lead Student Volunteer: Communications Committee Press Liaison**, Supercomputing 2019	*Denver, CO*
2019	**Student Volunteer**, Supercomputing 2018	*Dallas, TX*
2017	**Student Volunteer**, Supercomputing 2017	*Denver, CO*

Skills

Languages	Bash, C, C++, Python
Frameworks	CMake, Git, Intel HLS, LLVM, MLIR, OpenCL, Xilinx HLS