Sizhuo ZHANG

(617)306-1938
Room 2031A, 235 Albany Street, Cambridge, MA, 02139
szzhang@mit.edu
http://people.csail.mit.edu/szzhang/

OVERVIEW

- 5+ years of research and development experience in computer architecture and hardware design.
- First author of multiple publications in computer architecture top conferences.
- Programming languages: C, C++, Python, Verilog, System Verilog, Bluespec System Verilog, MATLAB, Bash, Tcl.
- Passionate about computer/system architecture and hardware design; interested in other topics as well.

EDUCATION

Massachusetts Institute of Technology (MIT) Cambridge, MA
Department of Electrical Engineering and Computer Science 09/2013 – 06/2019
- **PhD** candidate in Computer Science, **Advisor:** Prof. Arvind GPA: 4.9 / 5.0
- **Honor:** Irwin Mark Jacobs and Joan Klein Jacobs Presidential Fellowship, 2013.9

Tsinghua University Beijing, China
Department of Microelectronics and Nanoelectronics 09/2009 – 06/2013
- **B.E.** in Microelectronics (Ranking: 1 / 54) GPA: 94.4 / 100
- **Honors:** Outstanding graduate of Tsinghua University, Outstanding thesis of Tsinghua University

WORKING EXPERIENCE

FPGA-accelerated Hardware Design Verification 05/2018 – 08/2018
Research Intern, Microsoft
- Developed a random testing framework and a coverage-driven testing framework on Microsoft's cloud FPGA.
- Evaluated the effectiveness of the frameworks in bug finding using an internal design that just entered the verification process. The frameworks were able to detect all the bugs found by a commercial model checking tool, potentially saving the cost of purchasing the commercial tool.

Workload Scheduling on Multicore 06/2017 – 08/2017
Engineering Intern, Waymo (Google Self-Driving Car Project)
- Analyzed the resource utilization of each major software module on the multicore processor in the self-driving car, and discovered congestion in the mapping from software modules to computing resources.
- Alleviated the congestion and improved software performance by adjusting the allocation of computing resources.

Messaging Framework for Hardware Accelerators 06/2015 – 08/2015
Engineering Intern, Microsoft
- Developed a messaging framework, HGum, which automatically generates serialization and deserialization logic from a given message schema to support software-to-hardware, hardware-to-software and hardware-to-hardware messaging.
- Evaluated HGum by applying it to a real hardware accelerator for feature extraction in the Bing ranking algorithm. HGum can save 70% hand-written codes with merely 5% performance overhead on average.

RESEARCH EXPERIENCE

Composable Building Blocks for Out-of-Order Multiprocessor 12/2015 – Present

Research Assistant, MIT

- Designed and Implemented a set of composable building blocks for out-of-order multiprocessors. The building blocks can be refined individually without affecting the correctness of the overall processor design, enabling fast exploration of microarchitectures.
- Composed a two-way superscalar out-of-order cache-coherent multiprocessor using the building blocks. The processor uses the RISC-V instruction set, can boot Linux on FPGA, and has comparable performance to other state-of-the-art academic out-of-order processors.

Weak Memory Consistency Models 10/2014 – 11/2017

Research Assistant, MIT

- Designed two weak memory consistency models, *WMM* and *GAM*, for multicores.
- WMM has a much simpler definition than any other weak memory model by forbidding one type of load/store reordering which has little impact on performance.
- GAM allows more flexibility in implementations than WMM in that it allows all 4 types of load/store reorderings.
- Developed axiomatic and operational definitions for WMM and GAM to cater for different types of program analysis.

Effects of Using Simplified Processor Models in Architectural Studies 09/2013 – 09/2014

Research Assistant, MIT

- Conducted architectural experiments on an FPGA-based multicore simulator to show that the influence of using a simplified core model instead of an accurate in-order core model is negligible.
- Improved the simulator for the study by using time-multiplexing to simulate more cores and adding a non-blocking last-level cache.

PUBLICATIONS

- **Sizhuo Zhang**, Andrew Wright, Thomas Bourgeat, Arvind, "Composable Building Blocks to Open up Processor Design", The 51st IEEE/ACM International Symposium on Microarchitecture (MICRO), October, 2018. (To appear.)
- **Sizhuo Zhang**, Muralidaran Vijayaraghavan, Andrew Wright, Mehdi Alipour, Arvind, "Constructing a Weak Memory Model", The 45th International Symposium on Computer Architecture (ISCA), June, 2018. DOI: 10.1109/ISCA.2018.00021
- Sang Woo Jun, Andrew Wright, **Sizhuo Zhang**, Shuotao Xu, Arvind, "GraFBoost: Accelerated Flash Storage for External Graph Analytics", The 45th International Symposium on Computer Architecture (ISCA), June, 2018. DOI: 10.1109/ISCA.2018.00042
- **Sizhuo Zhang**, Hari Angepat, Derek Chiou, "HGum: Messaging Framework for Hardware Accelerators", 2017 International Conference on ReConFigurable Computing and FPGAs (ReConFig), December, 2017. DOI: 10.1109/RECONFIG.2017.8279799
- **Sizhuo Zhang**, Muralidaran Vijayaraghavan, Arvind, "Weak Memory Models: Balancing Definitional Simplicity and Implementation Flexibility", The 26th International Conference on Parallel Architectures and Compilation Techniques (PACT), September, 2017. DOI: 10.1109/PACT.2017.29