Maleen Abeydeera

https://www.maleen-abeydeera.com
Email: maleen@csail.mit.edu Phone: (617) 800 9924

Skills : C, C++, Java, Python, Matlab, Verilog, FPGA, CUDA

Academic/ Research

PhD. Computer Science (Massachusetts Institute of Technology) (2015 - 2021)
Advisor: Daniel Sanchez, **Expected graduation** - June 2021 GPA: 5.0/5.0
Thesis: Architectural Techniques for Efficient Speculative Parallelism

- Developed C++ based architectural simulators to evaluate novel architectural features [2].
- Benchmarked the performance of emerging applications on novel architectures [3, 4].
- Built FPGA accelerators for several hard-to-parallelize applications in graph analytics and simulation [1].

BSc. Engineering (University of Moratuwa, Sri Lanka) (2009 - 2014)
Thesis: Accelerating Video Compression on FPGA GPA: 4.0/4.2

- Designed the system architecture for an FPGA-based HEVC (H.265) video decoder [5].
- Implemented the entropy decoder and inverse transform modules in RTL.
- Proposed a novel architecture for high-throughput inverse transform by making use of coefficient sparsity.

Industry Experience

Google – Software Engineering Intern (2020 Jun - Aug)

- Modelled the behavior of a security controller chip within a QEMU emulation framework. This model communicates with an external chip simulator, with a QEMU-side caching mechanism to reduce RPC calls.
- Ultimately, this work would accelerate new hardware development within Google.

Microsoft Research – Research Intern (2017 Jun - Aug)

- Demonstrated that, by reorganizing coefficient storage order, Convolutional Neural Networks (CNN) can be efficiently accelerated by an architecture that only supports pure vector operations.
- This work was ultimately adopted into the Microsoft Brainwave deep-learning platform.

Wave Computing – Consultant on behalf of Paraqum Technologies (2014 - 2015)

- Developed new benchmarks to characterize Wave's CGRA architecture and verify its compiler toolchain.

Selected Publications

[1] Chronos: Efficient Speculative Parallelism for Accelerators (paper, talk, video)
M. Abeydeera, D. Sanchez (ASPLOS 2020)

[2] SAM: Optimizing Multithreaded Cores for Speculative Parallelism (paper, talk)
M. Abeydeera, S. Subramanian, M. Jeffrey, J. Emer, D. Sanchez (PACT 2017)

[3] Fractal: An Execution Model for Fine-Grain Nested Speculative Parallelism (paper, talk)
S. Subramanian, M. Jeffrey, **M. Abeydeera**, H.R. Lee, V. Ying, J. Emer, D. Sanchez (ISCA 2017)

[4] Data-Centric Execution of Speculative Parallel Programs (paper, talk)
M. Jeffrey, S. Subramanian, **M. Abeydeera**, J. Emer, D. Sanchez (MICRO 2016)

[5] 4K Real-Time HEVC Decoder on an FPGA (paper)
M. Abeydeera, M. Karunaratne, G. Karunaratne, K. De Silva, A. Pasqual (IEEE TCSVT Jan 2016)