Amir Yazdanbakhsh

+1 (608)335-6884

https://cc.gatech.edu/~ayazdanb

a.yazdanbakhsh@gatech.edu

Objective

My objective is threefold: (1) secure a research position where I can fully challenge my knowledge and research skills, (2) collaborate with elite researchers in a vibrant environment, and (3) train myself to conduct world-class machine learning research.

Short Bio

Amir is an experienced researcher with a particular interest on designing efficient specialized hardware for machine learning applications. He is also interested in exploring the interplay between machine learning techniques and efficient computing system design. He published his research in multiple well-recognized peer-reviewed conferences and journals. His work has been recognized by multiple prestigious fellowships including Qualcomm Innovation Fellowship and Microsoft Graduate Research Fellowship.

Coursework

Machine Learning (Coursera)
Game Theory (Coursera)
Algorithm
Programming Language
Data Structure
Tools for Optimization
Advanced Computer Architecture
Advanced Database

Skills

Verilog
Python
C\C++
FPGA Design
CUDA
TensorFlow

Interests

Computer Architecture • Approximate Computing • Architecture Design for FPGA • Deep (Reinforcement) Learning • Analog-Digital Hybrid Computing

Education

2018 Ph.D. candidate in Computer Science — Georgia Institute of Technology
Computer Architecture and Systems

2013 M.Sc. in Electrical Engineering — University of Wisconsin-Madison
Algorithms and CAD

2010 M.Sc. Computer Engineering — University of Tehran
Computer Architecture and Systems

2007 B.Sc. in Computer Engineering — Shiraz University
Hardware Design and Computer Architecture

Awards

2016 Microsoft Research PhD Fellowship (two-year fully covered)
2016 Best Poster Award ARCHITECH (co-located with ASPLOS)
2016 Distinguished Paper Award in HPCA
2014 Qualcomm Innovation Fellowship, Qualcomm, USA

Publications

2017 A Unified SIMD-MIMD FPGA Acceleration for GANs
2017 Predictive Early Activation for Reducing Computation in CNNs
2016 AxBench: A Multi-Platform Benchmark Suite for Approximate Computing
2016 Controlling Quality Tradeoffs in Approximate Acceleration
2015 Neural Acceleration for GPU Throughput Processors
2015 Axilog: Language Support for Approximate Hardware Design
2015 RFVP: Rollback-Free Value Prediction with Approximate Loads
2014 General-Purpose Acceleration with Limited-Precision Analog Computation

Experience

2017 Research Intern — Nvidia, MA
Designed and implemented microarchitectural units for a particular spatial array architecture that aims to accelerate a variety of machine learning applications.

2016 Research Intern — Microsoft Research, WA
Developed a distributed version of Stochastic Gradient Descent (SGD) on COSMOS— Microsoft internal data storage/query system for enormous amounts of data— with large scalability.

2015 Research Intern — Rambus, CA
Developed a power-aware correcting code with partial correction capability. Developed an approximate memory design for hard failures.

2013 Visiting Researcher — Information Science Institute, VA
Developed a fully object-oriented program (C++) to analyze the effect of radiation on CMOS chips. Successfully verified the tool on novel residue-based arithmetic codes. The results show 99% detection of the single event upsets.

Open Source

2016 FPGA Accelerator Generator for ML Algorithms — http://act-lab.org/artifacts/tabla
2016 AxBench: Approximate Computing Benchmark — http://axbench.org
2015 NPiler: NPU Compiler — http://act-lab.org/artifacts/npiler

Patents

2017 Predictive Early Activation for Reducing Computation in DNNs — In Progress
2017 In-DRAM Near-Data Approximate Acceleration — In Progress
2016 Computer System Predicting Memory Failure — US 20160148707
2012 High-Performance and Low-Power Wireless Sensor Network — IR 79179