

Amir Yazdanbakhsh

<http://www.cc.gatech.edu/~ayazdanb/>
a.yazdanbakhsh@gatech.edu | (608).335.6884

RESEARCH INTERESTS

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

EDUCATION

■ Georgia Institute of Technology

PhD Candidate in Computer Science | August 2018
Area | Computer Architecture and Systems
Advisor | [Dr. Hadi Esmaeilzadeh](#)
Topic | Neuro-General Computing: An Acceleration-Approximation Approach

■ University of Wisconsin-Madison

MSc in Electrical Engineering | Dec. 2013
Area | Algorithms and CAD
Advisor | [Prof. Azadeh Davoodi](#)
Topic | Operand-Aware Fault-Tolerant Design Techniques

■ University of Tehran

MSc in Computer Engineering | Sep. 2010
Area | Computer Architecture and Systems
Advisor | [Prof. Sied Mehdi Fakhraie](#)
Topic | Integrated Framework for Modeling Extensible Processor Architectures

■ Shiraz University

BSc in Computer Engineering | Sep. 2007
Area | Hardware Design and Computer Architecture
Advisor | [Prof. Hooman Tahayori](#)

HONORS AND AWARDS

- **Gold Medal in ACM Student Research Competition (SRC)** | 2018
 - GANAX: A Unified SIMD-MIMD Acceleration for Generative Adversarial Networks @ASPLOS'18
 - Awarded a \$500 cash prize
- **Microsoft Research PhD Fellowship (two-year fully covered tuition and stipend)** | 2016-2018
12 winners out of 211 applications nation-wide
- **Best Poster Award in ARCHITECTH (co-located with ASPLOS)** | 2016
"Neural Acceleration for GPU Throughput Processors"
- **Honorable Mention in IEEE Micro Top Picks** | 2016
"General-Purpose Code Acceleration with Limited-Precision Analog Computation"
- **Distinguished Paper Award in IEEE Symposium on High Performance Computer Architecture** | 2016
"TABLA: A Unified Template-based Framework for Accelerating Statistical Machine Learning"

- **Qualcomm Innovation Fellowship, Qualcomm, USA** | 2014-2015
Together with Bradley Thwaites. 9 teams out of 137 teams nation-wide
- **Nominated for a CACM Research Highlights** | 2014
"General-Purpose Code Acceleration with Limited-Precision Analog Computation"
- **UW Chancellor's Opportunity Scholarship, UW-Madison, USA** | 2011

PUBLICATIONS

■ Peer-Reviewed Conference Papers

- 2018
- [P24] **A. Yazdanbakhsh**, C. Song, J. Sacks, P. Lotfi-Kamran, N. Kim, and H. Esmailzadeh, "In-DRAM Near-Data Approximate Acceleration for GPUs," in International Conference on Parallel Architectures and Compilation Techniques (PACT).
- [P23] **A. Yazdanbakhsh**, M. Brzozowski, B. Khaleghi, S. Ghodrati, K. Samadi, H. Esmailzadeh, and N. Sung Kim, "FlexiGAN: An End-to-End Solution for FPGA Acceleration of Generative Adversarial Networks," in IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM).
- [P22] **A. Yazdanbakhsh**, H. Falahati, P. J. Wolfe, K. Samadi, H. Esmailzadeh, and N. Sung Kim, "GANAX: A Unified SIMD-MIMD Acceleration for Generative Adversarial Network," in ACM/IEEE International Symposium on Computer Architecture (ISCA).
- [P21] V. Akhlaghi[†], **A. Yazdanbakhsh**[†], K. Samadi, H. Esmailzadeh, and R. K. Gupta, "SnaPEA: Predictive Early Activation for Reducing Computation in Deep Convolutional Neural Networks," in ACM/IEEE International Symposium on Computer Architecture (ISCA). [†](Equal contribution)
- [P20] **A. Yazdanbakhsh** and H. Esmailzadeh "GANAX: A Unified SIMD-MIMD Acceleration for Unsupervised Learning with Generative Adversarial Networks," in Student Research Competition collocated with the International Conference on Architectural Support for Programming Languages and Operating Systems (SRC@ASPLOS'18).
- [P19] **A. Yazdanbakhsh**, K. Samadi, H. Esmailzadeh, and N. S. Kim "A SIMD-MIMD Acceleration with Access-Execute Decoupling for Generative Adversarial Networks," in SysML Conference.
- 2016
- [P18] D. Mahajan, **A. Yazdanbakhsh**, J. Park, B. Thwaites, and H. Esmailzadeh, "Towards Statistical Guarantees in Controlling Quality Tradeoffs in Approximate Acceleration," in ACM/IEEE International Symposium on Computer Architecture (ISCA).
- [P17] A. Lotfi, A. Rahimi, **A. Yazdanbakhsh**, H. Esmailzadeh, and R. K. Gupta, "GRATER: An Approximation Workflow for Exploiting Data-Level Parallelism in FPGA Acceleration," in Design, Automation & Test in Europe (DATE).
- [P16] D. Mahajan, J. Park, E. Amaro, H. Sharma, **A. Yazdanbakhsh**, J. Kim, and Hadi Esmailzadeh, "TABLA: A Unified Template-based Framework for Accelerating Statistical Machine Learning," in IEEE Symposium on High Performance Computer Architecture (HPCA). (Distinguished Paper Award)
- [P15] **A. Yazdanbakhsh**, G. Pekhimenko, B. Thwaites, H. Esmailzadeh, O. Mutlu, and T. Mowry, "RFVP: Rollback-Free Value Prediction with Approximate Loads," in High Performance Embedded Architectures and Compiler (HiPEAC).
- 2015
- [P14] D. Mahajan, **A. Yazdanbakhsh**, J. Park, B. Thwaites, and Hadi Esmailzadeh, "MITHRA: Controlling Quality Tradeoffs in Approximate Computing," in TechCon.
- [P13] **A. Yazdanbakhsh**, J. Park, H. Sharma, P. Lotfi-Kamran, and H. Esmailzadeh, "Neural Acceleration for GPU Throughput Processors," in International Symposium on Microarchitecture (MICRO).
- [P12] **A. Yazdanbakhsh**, D. Palframan, A. Davoodi, N. S. Kim, and M. Lipasti, "Online and Operand-Aware Detection of Failures Utilizing False Alarm Vectors," in ACM Great Lakes Symposium on VLSI (GLSVLSI).
- [P11] **A. Yazdanbakhsh**, D. Mahajan, B. Thwaites, J. Park, A. Nagendrakumar, S. Sethuraman, K. Ramkrishnan, N. Ravindran, R. Jariwala, A. Rahimi, H. Esmailzadeh, and Kia Bazargan, "Axilog: Language Support for Approximate Hardware Design," in Design, Automation & Test in Europe (DATE).

..... 2014

[P10] B. Thwaites, G. Pekhimenko, **A. Yazdanbakhsh**, J. Park, G. Mururu, H. Esmailzadeh, O. Mutlu, and T. Mowry, **"Rollback-Free Value Prediction with Approximate Loads,"** in International Conference on Parallel Architecture and Compiler Techniques (PACT).

[P9] R. St. Amant, **A. Yazdanbakhsh**, J. Park, B. Thwaites, H. Esmailzadeh, A. Hassibi, L. Ceze, and D. Burger, **"General-Purpose Code Acceleration with Limited-Precision Analog Computation,"** in International Symposium on Computer Architecture (ISCA). **(Honorable Mention in IEEE Micro Top Picks & Nominated for a CACM Research Highlights).**

[P8] M. D. Sika, **A. Yazdanbakhsh**, B. Kiddie, J. Ahlbin, M. Bajura, M. Fritze, J. Damoulakis, and J. Granacki **"Low Energy Hardening of Combinatorial Logic using Standard Cells and Residue Codes,"** in Government Microcircuit Applications and Critical Technology (GOMACTech).

..... 2013

[P7] M. D. Sika, **A. Yazdanbakhsh**, B. Kiddie, J. Ahlbin, M. Bajura, M. Fritze, J. Damoulakis, and J. Granacki **"Applying Residue Arithmetic Codes to Combinational Logic to Reduce Single Event Upset,"** in International Conference on Radiation Effects on Components and Systems (RADECS).

..... 2011

[P6] F. Firouzi, **A. Yazdanbakhsh**, H. Dorosti and S. M. Fakhraie **"Dynamic Soft Error Hardening via Joint Body Biasing and Dynamic Voltage Scaling,"** in Euromicro Conference on Digital System Design (DSD).

..... 2010

[P5] **A. Yazdanbakhsh**, M. E. Salehi, S. Safari, and S. M. Fakhraie **"Locality Considerations in Exploring Custom Instruction Selection Algorithms,"** in Asia Symposium on Quality Electronic Design (ASQED).

[P4] **A. Yazdanbakhsh**, M. Kamal, M. E. Salehi, H. Noori, and S. M. Fakhraie **"Energy-Aware Design Space Exploration of Register File for Extensible Processors,"** in International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS).

[P3] M. E. Salehi, A. Azarpeyvand, F. Firouzi, and **A. Yazdanbakhsh**, **"Reliability Analysis of Embedded Applications in Non-Uniform Fault Tolerant Processors,"** in International Conference on Future Information Technology (FutureTech).

[P2] **A. Yazdanbakhsh**, M. E. Salehi, and S. M. Fakhraie, **"Architecture-Aware Graph-Covering Algorithm for Custom Instruction Engines,"** in International Conference on Future Information Technology (FutureTech).

[P1] A. Azarpeyvand, M. E. Salehi, F. Firouzi, **A. Yazdanbakhsh**, and S. M. Fakhraie, **"Instruction Reliability Analysis for Embedded Processors,"** in IEEE Symposium on Design and Diagnostics of Electronic Circuits and Systems (DDECS).

■ **Refereed Journal Articles**

..... 2018

[J10] Z. Liu, **A. Yazdanbakhsh**, T. Park, H. Esmailzadeh, and N. S. Kim, **"SiMul: An Algorithm-Driven Approximate Multiplier Design for Machine Learning,"** in IEEE Micro, special issue on Approximate Computing.

..... 2016

[J9] **A. Yazdanbakhsh**, D. Mahajan, P. Lotfi-Kamran, and H. Esmailzadeh, **"AxBench: A Multi-Platform Benchmark Suite for Approximate Computing,"** in IEEE Design and Test, special issue on Computing in the Dark Silicon Era.

..... 2015

[J8] **A. Yazdanbakhsh**, G. Pekhimenko, B. Thwaites, H. Esmailzadeh, O. Mutlu, and T. Mowry, **"RFVP: Rollback-Free Value Prediction with Approximate Loads,"** in ACM Transactions on Architecture and Code Optimization (TACO).

[J7] **A. Yazdanbakhsh**, G. Pekhimenko, B. Thwaites, H. Esmailzadeh, O. Mutlu, and T. Mowry, **"Mitigating the Bandwidth Bottleneck with Approximate Load Value Prediction,"** in IEEE Design and Test Journal.

[J6] D. Mahajan, K. Ramkrishnan, R. Jariwala, **A. Yazdanbakhsh**, J. Park, B. Thwaites, A. Rahimi, H. Esmailzadeh, and Kia Bazargan, **"Axilog: Abstractions for Approximate Hardware Design and Reuse,"** in IEEE Micro, special issue on Alternative Computing Designs and Technologies.

..... 2014

[J5] **A. Yazdanbakhsh**, R. Balasubramanian, T. Nowatzki, and K. Sankaralingam, **"Comprehensive Circuit Failure Prediction for Logic and SRAM using Virtual Aging,"** in IEEE Micro, special issue on Harsh Chips.

[J4] **A. Yazdanbakhsh**, M. Kamal, S. M. Fakhraie, A. A. Kusha, S. Safari, and M. Pedram **"Implementation-Aware Selection of The custom Instruction Set for Extensible Processors,"** in Journal of Microprocessors and Microsystems: Embedded Hardware Design (MICPRO).

[J3] **A. Yazdanbakhsh**, M. E. Salehi, and S. M. Fakhraie **"Customized Pipeline and Instruction Set Architecture for Embedded Processing Engines,"** in Journal of Supercomputing.
 2013

[J2] M. Kamal, **A. Yazdanbakhsh**, H. Noori, A. Afzali-Kusha, and M. Pedram **"A New Merit Function for Custom Instruction Selection Under an Area Budget Constraint,"** in Journal of Design Automation for Embedded Systems (DAEM).
 2012

[J1] M. E. Salehi, S. M. Fakhraie, and **A. Yazdanbakhsh**, **"Instruction Set Architectural Guidelines for Embedded Packet-Processing Engines,"** in Journal of System Architecture.

■ **Book Chapters**

..... 2018

[B2] **A. Yazdanbakhsh**, G. Pekhimenko, H. Esmailzadeh, O. Mutlu, and T. C. Mowry, **"Mitigation of Memory Bottleneck with Approximate Load Value Prediction,"** Approximate Circuits and Methodologies Book, Springer.
 2016

[B1] **A. Yazdanbakhsh**, J. Park, H. Sharma, P. Lotfi-Kamran, and H. Esmailzadeh, **"Neural Acceleration for GPU Throughput Processors,"** Advances in GPU Research and Practice, Elsevier.

■ **Workshop Papers**

..... 2016

[W5] **A. Yazdanbakhsh**, J. Sacks, Ch. Song, P. Lotfi-Kamran, H. Esmailzadeh, and N. Sung-Kim, **"NAX: Near-Data Approximate Computing,"** in Workshop on Approximate Computing (AC) @ESWEEK.
 2015

[W4] D. Mahajan, **A. Yazdanbakhsh**, J. Park, B. Thwaites, and H. Esmailzadeh, **"Prediction-Based Quality Control for Approximate Accelerators,"** in Workshop on Approximate Computing Across System Stack (WACAS) @ASPLOS.
 2014

[W3] **A. Yazdanbakhsh**, R. St. Amant, B. Thwaites, J. Park, H. Esmailzadeh, A. Hassibi, L. Ceze, and Doug Burger, **"Toward General-Purpose Code Acceleration with Analog Computation,"** in Workshop on Approximate Computing Across System Stack (WACAS) @ASPLOS.

[W2] **A. Yazdanbakhsh**, B. Thwaites, J. Park, and H. Esmailzadeh, **"Methodical Approximate Hardware Design and Reuse,"** in Workshop on Approximate Computing Across System Stack (WACAS) @ASPLOS.

[W1] B. Thwaites, **A. Yazdanbakhsh**, J. Park, and H. Esmailzadeh, **"Bio-Accelerators: Bridging Biology and Silicon for General-Purpose Computing,"** in Wild and Crazy Ideas (WACI) @ASPLOS.

PROFESSIONAL EXPERIENCE

■ **NVIDIA RESEARCH | RESEARCHER INTERNSHIP**

WINTER 2017 | WESTFORD, MA

TEAM: COMPUTER ARCHITECTURE

MENTORS: [Dr. Michael Pellauer](#) and [Dr. Joel Emer](#) • MANAGER: [Dr. Steve Keckler](#)

- Modified and improved a simulation infrastructure for a particular spatial array architecture that aims to accelerate a variety of machine learning applications. Designed and implemented microarchitectural units for a particular spatial array architecture that aims to accelerate a variety of machine learning applications.

■ **MICROSOFT RESEARCH | RESEARCHER INTERNSHIP**

SUMMER 2016 | REDMOND, WA

TEAM: RESEARCH IN SOFTWARE ENGINEERING (RISE)

MENTOR: [Dr. Saeed Maleki](#) • MANAGER: [Dr. Ben Zorn](#)

- 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.

■ **RAMBUS LABS | RESEARCHER INTERNSHIP**

SUMMER 2015 | SUNNYVALE, CA

TEAM: APPROXIMATE COMPUTING

MENTOR: [Dr. David Stork](#) • MANAGER: [Dr. Craig Hampel](#)

- Developed a power-aware correcting code with partial correction capability. Developed an approximate memory design for hard failures.

■ **INFORMATION SCIENCE INSTITUTE | VISITING RESEARCHER**

SUMMER 2013 | ARLINGTON, VA

TEAM: ADVANCED ELECTRONICS / DISRUPTIVE ELECTRONICS

MENTOR: [Michel Sika](#) • MANAGER: [Dr. Michael Fritze](#)

- 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 (published @RADECS 2013).
- Synthesized the residue-based arithmetic codes with the characterized library for low-power operation. The results show **1.42x** less delay overhead and **1.57x** less energy per bit compared to the equivalent TMR implementations (published @GOMACTech 2014).

PROJECTS

■ **A UNIFIED SIMD-MIMD FPGA ACCELERATION FOR GENERATIVE ADVERSARIAL NETWORKS**

ADVISOR: [Dr. Hadi Esmaeilzadeh](#) and [Dr. Nam Sung Kim](#)

FEBRUARY 2017 – DECEMBER 2017

Although Generative Adversarial Networks (GANs) are gaining prominence in different fields, there are no accelerators for these new models. GANs leverage a new operator, called the transposed convolution that exposes new challenges for hardware acceleration. This operator first inserts zeros within the multidimensional input and then convolves a kernel over this expanded array to add information to the embedded zeros. The inserted zeros lead to underutilization of the compute resources when a traditional convolution accelerator is used. We propose the GANAX architecture to alleviate these sources of inefficiencies, making the first GAN accelerator design possible. Evaluations shows **3.5x** speedup and **3.1x** energy savings over Eyeriss without compromising the efficiency of traditional convolution. These benefits come with only $\approx 7.8\%$ area increase.

PUBLICATION:

- International Symposium on Computer Architecture (ISCA) 2018.
- International Symposium on Field-Programmable Custom Computing Machines (FCCM) 2018.

■ **PREDICTIVE EARLY ACTIVATION FOR REDUCING COMPUTATION IN CNNs**

ADVISORS: [Dr. Hadi Esmaeilzadeh](#) and [Prof. Rajesh K. Gupta](#)

FEBRUARY 2017 – DECEMBER 2017

In modern CNNs, the output of compute-heavy convolution operations are fed to activation units that output zero if their input is negative. By exploiting this unique algorithmic property of CNNs, we propose a predictive early activation technique. The proposed technique, called SNAPEA, cuts the computation of convolution operations short if it determines that the output will be negative. With only **1%** loss in classification accuracy, SNAPEA, on average, yields **2.2x** speedup and **1.9x** energy reductions. These benefits grow to **3.0x** speedup and **2.3x** energy reductions if the acceptable loss in classification accuracy is set to **3%**. Compared to the static pruning approaches, which are complimentary to the dynamic approach of SNAPEA, our proposed technique offers **1.3x** speedup and **1.1x** energy reductions with no loss in classification accuracy.

PUBLICATION:

- International Symposium on Computer Architecture (ISCA) 2018.

■ NEAR DATA APPROXIMATE COMPUTING

ADVISOR: [Dr. Hadi Esmaeilzadeh](#)

MARCH 2016 – DECEMBER 2016

Designed a novel DRAM architecture to enable near-data approximate acceleration in GPUs. Results show on average, **2.6x** speedup and **13.4x** energy reduction over a baseline GPU with no acceleration. These benefits are achieved while reducing overall DRAM system power by **26%** with an area cost of merely **2.1%**.

PUBLICATIONS:

- Workshop on Approximate Computing (AC) co-located with ESWEK 2016.

■ TEMPLATE-BASED FRAMEWORK FOR ACCELERATING STATISTICAL MACHINE LEARNING

ADVISOR: [Dr. Hadi Esmaeilzadeh](#)

APRIL 2015 – DECEMBER 2015

Helped to design TABLA, a framework that generates accelerators for a class of machine learning algorithms. TABLA leverages the insight that many learning algorithms can be expressed as a stochastic optimization problem. TABLA-generated accelerators provide **19.4x** and **2.9x** average speedup over the ARM and Xeon processors, respectively. These accelerators provide **17.57x**, **20.2x**, and **33.4x** higher Performance-per-Watt in comparison to Tegra, GTX 650 Ti and Tesla, respectively.

PUBLICATIONS/AWARDS:

- IEEE Symposium on High-Performance Computer Architecture (HPCA), 2016 (distinguished paper award).

■ NEURAL ACCELERATION FOR GPU THROUGHPUT PROCESSORS

ADVISORS: [Dr. Hadi Esmaeilzadeh](#) and [Dr. Pejman Lotfi-Kamran](#)

FEBRUARY 2015 – MAY 2015

Designed and integrated an approximate form of acceleration, neural acceleration, within the GPGPU architectures. The architecture design for the neurally accelerated Streaming Multiprocessors (SMs), called **GNPU**, provides significant performance and energy efficiency benefits while providing reasonably low hardware overhead (**1.2%** area overhead per SM). The results show on average **2.4x** speedup and **2.8x** energy reduction with **10%** quality loss across a diverse set of GPGPU benchmarks.

PUBLICATIONS/AWARDS:

- Best Poster Award in ARCHITECT co-located with ASPLOS, 2016.
- Book Chapter in Advances in GPU Research and Practice, Elsevier, 2016.
- Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), 2015.

■ AN APPROXIMATION WORKFLOW FOR EXPLOITING DATA-LEVEL PARALLELISM IN FPGA ACCELERATION

ADVISORS: [Dr. Hadi Esmaeilzadeh](#) and [Prof. Rajesh K. Gupta](#)

APRIL 2015 – OCTOBER 2015

Helped to devise an automated design workflow for FPGA accelerators that leverages imprecise computation to increase data-level parallelism and achieve higher computational throughput. By selectively reducing the precision of the data and operation, the required area to synthesize the kernels on the FPGA decreases allowing to integrate a larger number of operations and parallel kernels in the fixed area of the FPGA. The synthesis result on a modern Altera FPGA with a diverse set of data-intensive OpenCL benchmarks shows that our approximation workflow yields **1.4x - 3.0x** higher throughput with less than **1%** quality loss.

PUBLICATIONS:

- IEEE Design, Automation and Test in Europe (DATE), 2016.

■ MITHRA: CONTROLLING QUALITY TRADEOFFS IN APPROXIMATE ACCELERATION

ADVISOR: [Dr. Hadi Esmaeilzadeh](#)

SEPTEMBER 2014 – MAY 2015

Introduced a hardware software mechanism, called MITHRA, that controls the tradeoffs between performance, energy efficiency and quality loss in approximate accelerators. MITHRA uses a neural predictor to identify inputs that are likely to lead to relatively large quality losses, and instructs the core to run the original precise function, instead of invoking the approximate accelerator. The results show on average **2.7x** speedup and **3.2x** energy reduction while effectively restricting the quality loss below **5%**.

PUBLICATIONS:

- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2016.
- Workshop on Approximate Computing Across System Stack (WACAS) co-located with ASPLOS, 2015.

■ **AXILOG: LANGUAGE SUPPORT FOR APPROXIMATE HARDWARE DESIGN**

ADVISOR: [Dr. Hadi Esmaeilzadeh](#)

APRIL 2014 – SEPTEMBER 2015

Defined a language extension to Verilog HDL that provides the syntax and the semantics necessary for approximate hardware design and reuse. Successfully realized the introduced language extension with two different synthesis approaches. The results show on average **1.8x** energy reduction.

PUBLICATIONS:

- IEEE Design, Automation and Test in Europe (DATE), 2015.
- IEEE Micro, Alternative Computing Designs and Technologies, 2015.

■ **BRIDGING ANALOG NEURAL AND DIGITAL COMPUTING**

NOV 2013 – JUNE 2015

Used an algorithmic transformation that automatically converts regions of code from a von Neumann model to an analog neural model, which can be offloaded to a special analog accelerator. Introduced a customized training algorithm for limited-precision analog neural processing units. Results show application speedup of **3.7x** and energy savings of **6.3x**.

PUBLICATIONS/AWARDS:

- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2014.
- Qualcomm Innovation Fellowship, 2014.
- Honorable Mention in IEEE Micro Top Picks 2015.
- Nominated for a CACM Research Highlights.

■ **MITIGATE THE MEMORY BANDWIDTH WALL IN GPGPU**

ADVISOR: [Dr. Hadi Esmaeilzadeh](#)

MAY 2014 – SEPTEMBER 2015

Introduced a new approximation technique in GPGPUs by dropping certain fraction of performance-critical loads. Results show, on average, **40%** speedup and **31%** energy reduction.

PUBLICATIONS:

- High Performance Embedded Architectures and Compilers (HiPEAC), 2016.
- ACM Journal of Architecture and Code Optimization (TACO), 2015.
- IEEE International Conference on Parallel Architecture and Compilation Techniques, 2014.

■ **PREDICTING THE WEAR-OUT FAILURES IN MANY-CORE PROCESSORS**

ADVISOR: [Prof. Karu Sankaralingam](#)

AUGUST 2013 – APRIL 2014

DESCRIPTION:

Designed a unified, yet simple microarchitecture to predict and detect logic and SRAM wear-out failures in the field for many-core processors. The proposed system can predict and detect wear-out failures within **0.4 days** for logic and within **milliseconds** for SRAM-based modules after their appearance.

PUBLICATIONS:

- IEEE Micro, Harsh Chips, 2014.

■ ONLINE AND OPERAND-AWARE DETECTION OF FAILURES UTILIZING FALSE ALARM VECTORS

ADVISOR: [Prof. Azadeh Davoodi](#)

AUGUST 2011 – MAY 2012

DESCRIPTION:

Developed a framework which detects *online* and at *operand level* of granularity *all* the vectors which excite a set of diagnosed failures in combinational modules. Also, designed a checker unit which gets programmed after identifying one or more failures within combinational modules. For a given failure, if an operand yields to an observable fault, it will be detected by our checker and a recovery flag will be activated at runtime. The main contribution of this work is to significantly *minimize* the number of stored test cubes by inserting only a few but carefully-selected “false-alarm” vectors.

PUBLICATIONS:

- IEEE Great Lakes Symposium on VLSI (GVLSI), 2015.
- International Workshop on Logic and Synthesis (IWL) co-located with DAC, 2013.

ARTIFACTS

- GANAX: A Unified MIMD-SIMD Acceleration for Generative Adversarial Networks | <http://ganax.org>
- Axilog: Language Support for Approximate Hardware Design | <http://act-lab.org/artifacts/axilog>
- AxBench: Approximate Computing Benchmark | <http://axbench.org>
- NPiler: NPU Compiler | <http://act-lab.org/artifacts/npiler>

PATENTS

- Amir Yazdanbakhsh, Hardik Sharma, and Hadi Esmaeilzadeh, **Bit-Level Software Techniques for Executing Variable Precision Operations on General Purpose Processors**, Pending.
- Amir Yazdanbakhsh, Hadi Esmaeilzadeh, and Nam Sung Kim, **An Acceleration for Training and Inference of Generative Adversarial Networks**, Pending.
- Amir Yazdanbakhsh, Hadi Esmaeilzadeh, and Nam Sung Kim, **In-DRAM Near-Data Approximate Acceleration**, Pending.
- Amir Yazdanbakhsh, and Hadi Esmaeilzadeh, **Predictive Early Activation for Reducing Computation in Deep Neural Networks**, Pending.
- Amir Yazdanbakhsh, Raghuraman Balasubramanian, Anthony Nowatzki, Karthikeyan Sankaralingam, **Computer System Predicting Memory Failure**, US 20160148707.
- H. Dorosti, M. E. Salehi, A. Mazraei-Farahani, **A. Yazdanbakhsh**, and S. M. Fakhraie, **“Customized High-Performance and Low-Power Wireless Sensor Network,”** IRAN Patent 79179.

TALKS

- GANAX: A Unified SIMD-MIMD Acceleration for Generative Adversarial Network, ISCA, Los Angeles, CA, 2018.
- SnaPEA: Predictive Early Activation for Reducing Computation in Deep Convolutional Neural Networks, ISCA, Los Angeles, CA, 2018.
- GANAX: A Unified SIMD-MIMD Acceleration for Generative Adversarial Networks, SRC Competition (co-located with ASPLOS), Williamsburg, VA, 2018.

- **Neural Acceleration for GPU Throughput Processors**, MICRO, Waikiki, HI, 2015.
- **Approximate Computing: Cross-Stack Solutions for Efficient Computing**, AMD, Sunnyvale, CA, 2015.
- **Approximate Computing: Cross-Stack Solutions for Efficient Computing**, USC, Los Angeles, CA, 2015.
- **Bridging Analog Neuromorphic and Digital Von Neumann Computing**, Qualcomm, San Diego, CA, 2015.
- **General-Purpose Code Acceleration with Limited-Precision Analog Computation**, ISCA, Minneapolis, MN, 2014.
- **Energy-Aware Design Space Exploration of Register File for Extensible Processors**, SAMOS, Samos, Greece, 2010.

SUPERVISED STUDENTS

- **Hajar Falahati** | 2018
Research Project | Heterogeneous Architecture for Machine Learning
- **Marc Marone** | 2018
Research Project | RLArch: Optimized Design of Neural Network Acceleration using Reinforcement Learning
- **Michael Brzozowski (Winner of Georgia Tech President's Undergraduate Research Award)** | 2018
Research Project | RLArch: Optimized Design of Neural Network Acceleration using Reinforcement Learning
| A Unified SIMD-MIMD Acceleration for Generative Adversarial Networks

MENTORED STUDENTS

- **Ecclesia Morain** | 2017
Research Project | A Unified SIMD-MIMD Acceleration for Generative Adversarial Networks
- **Vahideh Akhlaghi** | 2017
Research Project | Predictive Early Activation for Reducing Computation in CNN
- **Hardik Sharma** | 2015
Research Project | Neural Acceleration for GPU Throughput Processors
- **Manali Kumar** | 2015
Research Project | Customized Training Algorithm for Analog Neural Processing Unit
First Employment | IBM, Austin, TX
- **Shashank Phalke** | 2015
Research Project | Customized Training Algorithm for Analog Neural Processing Unit
- **Priyank Deshpande, Advait Sakhalkar, and Ameya Ghadi** | 2014-2015
Research Project | Axilog Compiler and Synthesis framework
- **Aniket Patwardhan** | 2014-2015
Research Project | Axilog Compiler and Synthesis framework
First Employment | Intel Corp., Hillsboro, OR
- **Oleg Filatov** | 2013-2015
Research Project | Neural Transformation for Conventional Programming Languages
- **Anandhavel Nagendrakumar** | 2013-2014
Research Project | Approximate Hardware Description Language
First Employment | IM Flash Technologies, Lehi, UT
- **Sindhuja Sethuraman** | 2013-2014
Research Project | Approximate Hardware Description Language
First Employment | Apple Inc., Orlando, FL

PROFESSIONAL ACTIVITIES AND SERVICE

- **PROGRAM AND ORGANIZATION COMMITTEES**
 - ArchiTech | 2016
- **CONFERENCE REVIEWS**
 - IISWC, MICRO, ISCA, ISPASS, HPCA | 2016

■ **JOURNAL REVIEWS**

- IEEE Computer Architecture Letters (CAL), Journal on Emerging Technologies in Computing Systems | 2018

■ **IEEE K-12 OUTREACH**

- Family Science Nights | February 2016

Location | Eastvalley Elementary School

Description | Taught the students how to write a very simple computer program

■ **IEEE WOMEN IN ENGINEERING**

- An Introduction to Engineering for the Female Scholars (National Engineering Week) | February 2016

Location | Ivy Preparatory Academy

Description | Led two science sessions for female students in 6/7th and 7/8th grade

■ **PROFESSIONAL MEMBERSHIPS**

- IEEE Student Member | Since 2015

- ACM SIGARCH Online Membership | Since 2015