

# Curriculum Vitae

## Carole-Jean Wu

Assistant Professor,  
School of Computing, Informatics, and Decision Systems Engineering,  
Arizona State University

Office address: 660 W. 6<sup>th</sup> St. Tempe, AZ 85281

Website: <http://faculty.engineering.asu.edu/carolewu/>

Email: [carole-jean.wu@asu.edu](mailto:carole-jean.wu@asu.edu)

Office telephone: +1 480 965 1494, Cellphone: +1 607 592 6397

Last update: August 1, 2017

### **BIOGRAPHICAL NOTE**

Carole-Jean Wu is an Assistant Professor of Computer Science and Engineering in the School of Computing, Informatics, and Decision Systems Engineering in Arizona State University (ASU). She is the director of the Energy Efficient Computer Architecture Laboratory (EECALab) and the associate site director of the NSF I/UCRC Center for Embedded Systems (CES). She holds an affiliated faculty appointment in EECE at ASU. Prof. Wu receives her M.A. and Ph.D. in Electrical Engineering from Princeton University in 2008 and 2012, respectively. She receives a B.Sc. degree in Electrical and Computer Engineering from Cornell University. Before joining ASU, Prof. Wu held a number of industrial internship positions with Intel, IBM, and Google. She is an IEEE senior member.

Prof. Wu works in the area of Computer and System Architectures. In particular, her research interests include high-performance and energy-efficiency computer architectures through hardware heterogeneity, energy harvesting techniques for emerging computing devices, temperature and energy management for portable electronics, performance characterization, analysis and prediction, and memory subsystem designs. She has published 6 journal papers and 20 conference papers. Prof. Wu is the recipient of the 2017 NSF CAREER Award, the 2017 IEEE Young Engineer of the Year Award, the 2014 Best of CAL Award, the 2013 Science Foundation of Arizona Bisgrove Early Career Award, and the 2011-12 Intel Ph.D. Fellowship Award. She was the recipient of the 2014 IEEE Best of Computer Architecture Letter award. Her research has been supported by both industry sources and the National Science Foundation to a level over \$1.8 million.

### **RESEARCH INTERESTS**

- High-performance and energy-efficient heterogeneous architecture
- Energy harvesting for emerging computing devices
- Energy management for portable electronics
- Temperature-aware computing
- Advanced processor and system cooling designs and management
- Performance characterization, analysis and prediction
- High-performance and power-efficient memory designs

## **EDUCATION**

**Princeton University**, Princeton, NJ  
Doctor of Philosophy, Electrical Engineering April 2012  
Thesis: Dynamic Techniques for Mitigating Inter- and Intra-Application Cache Interference

**Princeton University**, Princeton, NJ  
Master of Art, Electrical Engineering May 2008

**Cornell University**, Ithaca, NY  
Bachelor of Science, Electrical and Computer Engineering May 2006  
Honor Thesis: Predictability of Microprocessor L2 Cache Misses

## **ACADEMIC EMPLOYMENT**

**Assistant Professor**, August 2012 – present

Arizona State University, Tempe, AZ

- Faculty member, School of Computing, Informatics, and Decision Systems Engineering (CIDSE)
- Director: Energy Efficient Computer Architecture Laboratory (EECALab)
- Associate Director: NSF I/UCRC Center for Embedded Systems (CES)
- Graduate faculty member, School of Electrical, Computer and Energy Engineering (ECEE)
- Honors Faculty, Honors College at ASU

## **PRIOR PROFESSIONAL EMPLOYMENT**

**Intel Architecture Group**, Hudson, MA

*Graduate Technical Intern* 9/2010 - 5/2011

High-Performance Cache Designs

- Designed novel high performing memory management techniques to improve cache efficiency for multimedia, games, enterprise servers, and scientific workloads.

**Google Platform Development**, Mountain View, CA

*Graduate Technical Intern* 6/2009 - 8/2009

Power Utilization Prediction for Data Center Workloads

- Developed a portfolio of application-level power prediction algorithms for data centers. This portfolio includes machine learning algorithms and linear regression techniques that exploit application CPU/power utilization history.

**Intel Labs**, Hillsboro, OR

*Graduate Technical Intern* 6/2008 - 8/2008

Performance and Scalability Characterization for Many-Core Systems

- Analyzed performance and scalability issues related to System Management Mode (SMM) on Intel's many-core and Tera-scale platforms.

**IBM TJ Watson Research Center**, Hawthorne, NY

*Graduate Technical Intern* 6/2007 - 8/2007

Visualization Software Development for Large-Scale Multicore System

- Created the visualization software, TuningFork, for large-scale multicore systems.

**Intel Software Solution Group**, Chandler, AZ

*Undergraduate Technical Intern* 8/2004 - 12/2004; 5/2005 - 8/2005

Performance Characterization and Analysis

- Generated multi-threaded traces for IA32 pre-Si simulations for various applications, validated traces were representative, and performed pre-Si architectural analysis.

## **HIGHLIGHTS OF ACHIEVEMENTS**

### Research

- NSF Early CAREER Award – 2017.
- IEEE Young Engineer of the Year Award – 2017.
- IEEE Best of Computer Architecture Letters – 2014.
- Science Foundation Arizona Bisgrove Early Career Scholar – 2013.
- 26 peer-reviewed, published journal (6) and conference (20) papers.
  - 2 papers in the International Symposium on Computer Architecture (ISCA).
  - 2 papers in the International Symposium on Microarchitecture (MICRO).
  - 1 paper in the International Symposium on High Performance Computer Architecture (HPCA).
  - 1 paper in the International Conference on Parallel Architectures and Compilation Techniques (PACT).
  - 5 journal and 17 conference papers are carried out during Prof. Wu's ASU employment.
  - 4 journal and 17 conference papers are carried out with ASU PhD/MS students.
- Four NSF awards and a number of industrial research grants, total to a level over **\$1.8 million**.

### Teaching

- Currently supervising a total of 4 PhD students, 1 post-doctoral researcher, 1 summer REU student, and 2 high school students.
- Graduated a total of 3 PhD students, 2 MS students, and 1 undergraduate honors student as the Thesis Chair/Co-Chair.
- Have taught a total of three courses.
  - Regularly teach a core undergraduate course of Computer Systems and Engineering – CSE420: Computer Architecture I
  - Regularly teach a graduate course – CSE520: Computer Architecture II
  - Designed and taught a new graduate seminar course – CSE591: Advanced Memory Systems and Architecture
  - A total of 311 students

### Service

- Served on the Program Committees of the top-tier, major conferences in computer architecture
  - International Symposium on Computer Architecture (ISCA), 2017.
  - International Symposium on Computer Architecture (ISCA), 2016.
  - International Symposium on Computer Architecture (ISCA), 2015.
  - International Symposium on Computer Architecture (ISCA), 2014.
  - International Symposium on Microarchitecture (MICRO), 2017.
  - International Symposium on Microarchitecture (MICRO), 2016.
  - International Symposium on Microarchitecture (MICRO), 2014.
  - International Symposium on High Performance Computer Architecture (HPCA), 2017.
  - International Symposium on High Performance Computer Architecture (HPCA), 2016.
  - International Symposium on High Performance Computer Architecture (HPCA), 2014.
- Served as a reviewer for the major journals in computer architecture
  - IEEE Transactions on Computers
  - ACM Transactions on Architecture and Code Optimization
  - ACM Transactions on Parallel and Distributed Systems
- Organized and have been leading the ASU Grace Hopper Celebration for Women in Computing Scholarship Program.

- 100, 50, 50 ASU scholarship students attended the conference for professional career development (75% of which are undergraduate students) in 2014, 2015, and 2016, respectively.
- We were the number-one attending university in 2014.
- Since the inception of the scholarship program and with other BRAID-related activities, our CIDSE female vs. male undergraduate student ratio has improved from 10% to 15% from 2012 to 2016.
- Served on the Engineering Dean's New Faculty Cohort Steering Committee.
  - Faculty peer mentors, 2013-16.
- Participated at the Engineering E2 Camps 2012, 2014-16.
- Served on the Faculty Search Committees of Next Generation Computing, Big Data Architecture, and Trustworthy Internet-of-Things, 2012-14, 2014-15, 2016-17.
- Served on the Computer Engineering admission committee, 2013-16.
- Served as the course coordinator for CSE420: Computer Architecture I, 2014-17.

## **HONORS AND AWARDS**

- **NSF Early CAREER Award**, National Science Foundation, 2017.

The National Science Foundation’s most prestigious award in supporting junior tenure-track faculty who exemplify the role of teach-scholars through outstanding research, excellent education and the integration of education and research.

- **IEEE Young Engineer of the Year Award**, Institute of Electrical and Electronics Engineers, Phoenix Section, 2017.

A prestigious award given to recognize “[the PI’s] contributions to Computer Architecture, with an emphasis on high performance and energy efficient computing systems”.

- **Best of CAL 2014**, IEEE Computer Architecture Letters, February 2015, for “Architectural Thermal Energy Harvesting Opportunities for Sustainable Computing,” (Wu, C.-J.).

- **SFAz Bisgrove Early Career Scholar**, Science Foundation of Arizona 2013-15.  
<http://www.sfaz.org/bisgrove-scholars/>

A prestigious award given to outstanding early career researchers in the state of Arizona.

- **Intel Ph.D. Graduate Fellow**, 2011-12.  
<http://www.intel.com/content/www/us/en/education/university/intel-phd-fellowship-infographic2012.html>

A prestigious award given to 21 outstanding Ph.D. students in semiconductor research.

- **Best Paper Nominee**, IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), April 2011, for “Characterization and Dynamic Mitigation of Intra-Application Cache Interference,” (Wu, C.-J. and Martonosi, M.).
- **Excellence in Leadership Award** for organizing Computer Architecture Day @ Princeton 2009.

## **IN THE NEWS**

***“Effort to produce more energy-efficient computing earns science foundation’s support”***

ASU News and ASU Full Circle

<https://fullcircle.asu.edu/research/effort-to-produce-more-energy-efficient-computing-earns-science-foundations-support/>

***“Making energy-saving strides in sustainable computing”***

ASU News, ASU Full Circle, and ASU Research Matters

<https://researchmatters.asu.edu/stories/asu-engineer-makes-strides-sustainable-computing-3820>

***“Energy harvesting: Creating useful power out of processor waste heat”***

Embedded Computing Design

<http://embedded-computing.com/articles/energy-processor-waste-heat/>

***“Breaking barriers to diversity in Computer Science”***

ASU News and ASU Full Circle

<https://fullcircle.asu.edu/students/breaking-barriers-diversity-computer-science/>

***“Women’s 101: ASU faculty discuss strides for Women’s History Month”***

ASU Now: Access, Excellence, and Impact

<https://asunow.asu.edu/20170330-sun-devil-life-womens-101-asu-faculty-discuss-strides-womens-history-month>

***“User satisfaction: The new framework for smartphone performance optimization”***

ASU News and ASU Full Circle

<https://fullcircle.asu.edu/faculty/user-satisfaction-new-framework-smartphone-performance-optimization/>

***“ASU wins record 14 NSF career awards”***

ASU Now: Access, Excellence, and Impact

<https://asunow.asu.edu/20170428-asu-wins-record-14-nsf-career-awards>

## **GRANTS**

### **Active/Funded research programs<sup>1</sup>**

[AP1]      **NSF:SHF:CAREER—Maximizing Energy Efficiency with Statistical Performance and Skin Temperature QoS Guarantee for Handhelds.** C.-J. Wu (PI); NSF-CISE-#1652132. \$450,000.

**Brief Description:**

Smartphone designs must achieve balance between performance, temperature and energy management in a coordinated manner to maximize user satisfaction. Their designs must offer high performance with guarantees of high energy efficiency, along with temperature requirements, namely, in device surface and processor junction temperature limits. Performance as measured by mobile application execution time has long been assumed to be a deterministic quantity. In reality, execution times vary substantially, depending on the characteristics of data inputs and the varying states of the system. Furthermore, the device surface temperature (or skin temperature) is a unique constraint for handheld devices. It is tightly coupled with the location of the major heat-generating source within a smartphone. Thus, a smartphone’s energy efficiency is intricately related to both performance quality and skin temperature management, making optimization for system energy efficiency a complex task for battery-powered platforms. This proposal tackles the problem of performance, temperature and energy efficiency co-optimization from the vantage point of managing the hardware resources. It builds on the PI’s prior work in system and hardware architecture by proposing an optimizing user satisfaction (OPUS) framework for holistic management and coordination of performance quality, skin temperature, and energy efficiency for handhelds.

[AP2]      **NSF:SHF—Latency Tolerance Aware Runtime Optimization for General-Purpose GPU Architectures.** C.-J. Wu (PI); NSF-CISE-SHF #1618039. \$330,000.

**Brief Description:**

The availability of programming and architecture support for executing general-purpose computations on GPUs enables the acceleration of a diverse set of important and large problems in scientific computing, environmental and disaster modeling and prediction, real-time image-based learning and prediction, medical applications, and many other domains. The algorithms and the computation patterns behind these problems are inherently different from one another, let alone being entirely different from graphics workloads. For example, some of the graph-based algorithms used to solve these problems exhibit a high degree of computation and memory access divergence. This presents a significant challenge for GPUs and leads to low computation and memory resources utilization for GPU architectures. On average, the pipeline utilization of running GPGPU applications hovers between 20-40% only. A natural, straight-forward approach to improve GPGPU resource utilization is to port and adapt existing CPU-centric optimization techniques to GPU architectures, in the hope to accelerate these general-purpose workloads. However, such approaches lead to sub-optimal performance improvement, if any. Thus, the main objective of this proposal is (1) to explore and identify important GPU characteristics that state-of-the-art optimization techniques must

---

<sup>1</sup> **NSF:** National Science Foundation. **CES:** Center for Embedded Systems. This is an NSF Industry University Collaborative Research Center at ASU. My research has been funded by the member companies, Marvell and Qualcomm Research. **SFAz:** Science Foundation Arizona.

consider to improve performance acceleration opportunities for GPGPU workloads, and (2) to design, implement, and evaluate a collection of novel architectural techniques in the LATTE runtime optimization framework that increases GPGPU performance by taking into account the key characteristics of GPU architectures—the latency tolerance feature.

[AP3] **Runtime Energy Management and Optimization.** C.-J. Wu (PI); VMware gift. \$10,000.

[AP4] **CES: Designing Advanced Thermal Solutions and Intelligent Control Mechanisms for High Performance Embedded Systems.** C.-J. Wu (PI); Qualcomm and Marvell (NSF I/UCRC). \$33,000.

**Brief Description:**

Modern high performance embedded electronics, fueling smartphones, self-driving automobiles, and augmented reality wearable, are multifunctional to provide real-time direction navigation, high definition imaging and video quality entertainment, messaging, real-time sensing and control. One of the major performance limiting factors in these platforms is the poorly designed thermal solution that is triggered to prevent overheating in electronics—the processor transistor junction and the electronics surface-level temperature. The former limits the maximum operating temperature of transistors to guarantee electronics reliability and lifetime, whereas the latter limits the maximum temperature of platform surface to guarantee user ergonomic comfort.

To overcome the thermal issue, we propose to investigate a hybrid cooling solution that combines the application of thin-film thermoelectric cooler (TFTEC) with advanced heat spreading and heat storage materials. TFTEC is suitable for site-specific, on-demand localized cooling of high heat-flux regions. We hypothesize that, with the proposed hybrid TFTEC-PCM approach (when the power source is not severely constrained), the electronics junction- and surface-level temperature can be controlled in a more energy efficient manner, leading to higher electronics performance and better user experience.

[AP5] **CES: Dynamic Memory Interference-Aware Technique for Mobile Web Browsing.** C.-J. Wu (PI); Qualcomm and Marvell (NSF I/UCRC). \$32,000.

**Brief Description:**

Mobile web browsing is the most important application accessed in smartphones. Thus, optimizing the performance of web browsing is pivotal to user satisfaction. It is also important to maximize smartphone energy efficiency to ensure the most efficient usage of the limited energy source on smartphones. An important aspect that has been neglected in prior work is the presence of memory interference incurred by background processes and co-scheduled applications. As more cores are packed into the mobile processor and more programmable accelerators are integrated, the shared memory hierarchy is becoming increasingly contentious, leading to significant performance degradation for real-time, user-centric web browsing and sub-optimal energy efficiency for smartphones.

To address the performance variability and the energy efficiency sub-optimality problems, this project proposes to develop models to (1) predict the performance of web browsing in the presence of background processes and co-scheduled applications and (2) estimate the power consumption of smartphones, by considering workload characteristics and dynamically-varying architectural and system conditions. By allowing the mobile processor to operate at an optimal energy-efficient condition, the battery lifetime of smartphones is

expected to increase significantly.

- [AP6] **NSF:SHF—Novel Architecture Energy Harvesting for Sustainable Spot Cooling and Energy Management.** C.-J. Wu (PI); NSF-CISE-SHF #1525462. \$419,066.

**Brief Description:**

Increased power dissipation in computing devices has led to a sharp rise in thermal hot spots on computer chips, creating a vicious cycle—higher temperatures brings higher leakage power, and higher power dissipation increases temperature, thereby leading to thermal avalanche. To reduce the additional power dissipation and reliability concerns caused by high temperature, current heat management approaches apply cooling mechanisms to remove heat aggressively as well as devise dynamic management techniques that avoid thermal emergencies by slowing down heat generation of processors. This proposal addresses the heat management problem using an innovative approach—rather than removing heat or slowing down heat generation, the proposed work transforms the waste heat into reusable energy for new applications such as self-powered spot-cooling. The main objective of this proposal is to design and implement a novel architectural Thermal Energy Harvesting framework (TEHar) to *create the mechanisms, policies, and system support that allow waste heat generated by computing devices to be harvested efficiently, to achieve better energy utilization efficiency.* This will be achieved by exploiting the thermal characteristics of modern computing nodes and by leveraging thermoelectric and pyroelectric energy harvesting materials.

- [AP7] **CES: A high-performance, QoS-aware memory scheduling approach for highly parallel, heterogeneous smartphone SoCs.** C.-J. Wu (PI); Qualcomm and Marvell (NSF I/UCRC). \$47,500.

**Brief Description:**

This proposal aims to design a high-performance, quality of service (QoS)-aware memory scheduler for the main memory shared by multiple processors on the mobile platform. We first study the degree of memory interference in emerging mobile phone applications to quantify the impact of memory interference caused by the various processing elements on application QoS and system throughput. Based on the characterization results, we design a memory scheduling algorithm that dynamically determines the memory request scheduling order to take advantage of the heterogeneous access patterns. The scheduling algorithm will enforce QoS constraints of real-time applications while improving the performance of co-scheduled applications to maximize system throughput.

- [AP8] **NSF:CNS:EAGER—Characterization and Modeling for Architectural Thermal Energy Harvesting.** C.-J. Wu (PI); NSF-CISE-CSR #1358805. \$200,000.

**Brief Description:**

To reduce the ever-increasing power dissipation caused by high temperature in computing systems, the current approaches seek to apply cooling mechanisms to remove heat aggressively, as well as devising management techniques to avoid thermal emergencies by slowing down heat generation. Complementary to existing techniques, this proposal attempts to address the heat management problem using a fundamentally different approach -- rather than removing the heat or slowing down heat generation, we transform this heat into reusable energy by using thermoelectric materials. A framework is proposed here that will allow heat energy generated by computing devices to be recovered, transformed, and harvested efficiently, to achieve better energy efficiency. The framework is designed based

on the interesting implication of thermal energy distribution of computing platforms: the temperature differences between the hottest and the coldest components can be more than tens of degrees, creating a steep spatial thermal gradient. We discover that, by leveraging the thermoelectric effects at the architectural level, the varying spatial thermal gradients created as a result of computations can be exploited to transform heat into reusable energy. Therefore, the heat generated by the circuitry of the computing devices is not wasted but is rather harvested for reuse. This research explores possibilities in thermal energy harvesting techniques at the architectural level. Overall, this proposal explores the potential for energy harvestability, particularly in the steep thermal gradients commonly observed in computing systems, while also investigating applications that can reuse this recovered energy.

### **Completed research programs**

- [CP1]      *CES: Energy-aware Application Scheduling for Heterogeneous and Parallel Smartphone Architectures.* C.-J. Wu (PI); Qualcomm and Marvell (NSF I/UCRC). \$36,000.
  
- [CP2]      *SFAz Bisgrove—Architectural Thermal Energy Recovery, Transformation, and Harvesting for Sustainable Computing.* C.-J. Wu (PI); Science Foundation of Arizona: Bisgrove Early Career Scholarship. \$200,000.
  
- [CP3]      *CES: Achieving Energy Efficient Mobile Computing Through Explicit Data Communication and Global Power Management.* C.-J. Wu (PI); Qualcomm and Marvell (NSF I/UCRC). \$43,408.

## **PUBLICATIONS**

All of my publications are available online at <http://faculty.engineering.asu.edu/carolewu/publications/> through links to the publishers and/or to preprint PDF files.

Convention for name ordering: Student names usually appear first. Among both student authors and non-student authors, the ordering of the names is based on the effort invested in writing the paper. The advisor of the major contributing author is typically listed as the last author.

### **Legend:**

- (\*) Corresponding Author
- Bold Font:** ASU Ph.D. Student
- Underline: ASU Master's Student
- (#) ASU Undergraduate Student
- (∞) Other/Visiting Undergraduate Student
- (✕) ASU Postdoctoral Researcher
- ‡ High School Student
- (+) Equal Contributions (if not equal include % of participation)

### **Impact Assessment for Publication Venues in Computer Architecture:**

The table below summarizes the impact assessment collected with Google Scholar Top Publications in Engineering and Computer Science

([https://scholar.google.com/citations?view\\_op=top\\_venues&hl=en&vq=eng](https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng)).

In Computer Architecture and many computer science disciplines, the impact factors (h-index) of top-tier conferences, such as the International Symposium on Computer Architecture (ISCA), the International Symposium on Microarchitecture (MICRO), and the International Symposium on High Performance Computer Architecture (HPCA) is on-par or higher than the best IEEE/ACM journal proceedings, such as the IEEE Transactions on Computers (TOC) or the ACM Transactions on Architecture and Code Optimization (TACO).

Papers published in the top-tier computer architecture conferences, such as the International Symposium on Computer Architecture (ISCA), the International Symposium on Microarchitecture (MICRO), and the International Symposium on High Performance Computer Architecture (HPCA), are typically double-column, single-spaced 12-page papers. Each paper typically receives 4-5 detailed technical reviews, 3 of which come from the Program Committee and the remaining 2 reviews can either come from the Program Committee or external reviewers of specific areas of expertise.

	<b>Publication</b>	<b>h5-median (Higher is better)</b>	<b>h5-index (Higher is better)</b>
	IEEE Transactions on Computers (TOC)	66	45
<b>Journal</b>	ACM Transactions on Architecture and Code Optimization (TACO)	35	23
	ACM Transactions on Embedded Computing Systems (TECS)	33	23

	25	16
IEEE Computer Architecture Letters		
	21	16
ASME Journal of Electronic Packaging		
	79	50
International Symposium on Computer Architecture (ISCA) <i>Typical acceptance rate: 15-20%</i>		
	71	41
International Symposium on Microarchitecture (MICRO) <i>Typical acceptance rate: 20%</i>		
	61	42
International Symposium on High Performance Computer Architecture (HPCA) <i>Typical acceptance rate: 20%</i>		
<b>Conferences</b>	58	43
	Design Automation Conference (DAC) <i>Typical acceptance rate: 20%</i>	
	45	30
	International Conference on Parallel Architectures and Compilation Techniques (PACT) <i>Typical acceptance rate: 25%</i>	
	28	19
	IEEE International Conference on Computer Design (ICCD) <i>Typical acceptance rate: 30%</i>	

\*h5-median for a publication is the median number of citations for the articles that make up its h5-index.

\*h5-index is the h-index for articles published in the last 5 complete years. It is the largest number h such that h articles published in 2011-2015 have at least h citations each.

## Peer-Reviewed, Referred Journal Publications

- A. Arunkumar\***, E. Bolotin, B. Cho, U. Milic, E. Ebrahimi, O. Villa, A. Jaleel, C.-J. Wu, and D. Nellans, “MCM-GPU: Multi-Chip-Module GPUs for Continued Performance Scalability,” Accepted and To Appear in *ACM SIGARCH Computer Architecture News*. This is the journal version of the ISCA paper (C1).
- [J1] **H.-M. Lee\***, S. Jeloka, **A. Arunkumar**, D. Blaauw, C.-J. Wu, T. Mudge, and C. Chakrabarti, “Using Low Cost Erasure and Error Correction Schemes to Improve Reliability of Commodity DRAM Systems,” In *IEEE Transactions on Computers (TOC)*, Vol. 65 Number 12, December 2016, 14 pages.
- [J2] **H.-M. Lee\***, C.-J. Wu, T. Mudge, and C. Chakrabarti, “RATT-ECC: Rate Adaptive Two-Tiered Error Correction Codes for Reliable 3D Die-Stacked Memory,” *ACM Transactions on Architecture and Code Optimization (TACO)*, Vol. 13, Issue 3, Article No. 24, September 2016, 24 pages.
- [J3] **S. Lee\***, P. E. Phelan, and C.-J. Wu, “Hot Spot Cooling and Harvesting CPU Waste Heat Using Thermoelectric Modules,” *Journal of Electronic Packaging*, Vol. 137, Issue 3, September 2015.
- S.-Y. Lee\***, **A. Arunkumar**, and C.-J. Wu, “CAWA: Coordinated Warp Scheduling and Cache Prioritization for Critical Warp Acceleration for GPGPU Workloads,” In *ACM SIGARCH Computer Architecture News*, Vol. 43, Issue 3, June 2015, pp. 515-527. This is the journal version of the ISCA paper (C9).
- [J4] **V. Hanumaiah\***, **D. Desai**, **B. Gaudette**, C.-J. Wu, and S. Vrudhula, “STEAM: A Smart Temperature and Energy Aware Multicore Controller,” In *ACM Transactions on Embedded Computing Systems (TECS)*, Vol. 13, Issue 5s, November 2014, pp. 151:1-25.
- [J5] C.-J. Wu\*, “Architectural Thermal Energy Harvesting Opportunities for Sustainable Computing,” In *IEEE Computer Architecture Letter*, Vol. 13, No. 2, July-December 2014, pp. 65-68.
- \*Received the Best of IEEE Computer Architecture Letter 2014 (CAL) Award**, February 2015.
- [J6] C.-J. Wu\* and M. Martonosi, “Adaptive Timekeeping Replacement: Fine-Grained Capacity Management for Shared CMP Caches,” In *ACM Transactions on Architecture and Code Optimization (TACO)*, Vol. 8, No. 1, Article 3, February 2011, pp. 3:1-26.

## Peer-Reviewed, Referred Conference Publications

Number of reviewers: typically 4 or more, minimum 3 (in rare cases).

- [C1] **S.-Y. Lee\*** and C.-J. Wu, “Performance Characterization, Prediction, and Optimization for Heterogeneous Systems with Multi-Level Memory Interference,” To Appear in *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Seattle, WA, October 2017.
- [C2] **A. Arunkumar\***, E. Bolotin, B. Cho, U. Milic, E. Ebrahimi, O. Villa, A. Jaleel, C.-J. Wu, and D. Nellans, “MCM-GPU: Multi-Chip-Module GPUs for Continued Performance Scalability,” To Appear in *Proceedings of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, Toronto, Canada, June 2017.

- [C3] **S.-Y. Lee\*** and C.-J. Wu, “Ctrl-C: Control Loop Based Adaptive Cache Bypassing for GPUs,” In *Proceedings of the International Conference on Computer Design (ICCD)*, Scottsdale, AZ, October 2016, pp. 133-140.
- [C4] **A. Arunkumar\***, **S.-Y. Lee**, and C.-J. Wu, “ID-Cache: Instruction and Memory Divergence Based Management for GPGPU Caches” In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Providence, RI, September 2016, pp. 158-167.
- [C5] **S. Lee\***, D. Pandiyan, J.-S. Seo, P. Phelan, C.-J. Wu, “Thermoelectric-based Sustainable Self-Cooling for Fine-Grained Processor Hot Spots,” In *Proceedings of the Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM)*, May 2016.
- [C6] **B. Gaudette\***, C.-J. Wu, and S. Vrudhula, “Improving Smartphone/Mobile User Experience by Balancing Performance and Energy with Probabilistic QoS Guarantee,” In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA)*, March 2016.
- [C7] D. Shingari\*, **A. Arunkumar**, and C.-J. Wu, “Characterization and Throttling-based Mitigation of Memory Interference for Heterogeneous Smartphones,” In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Atlanta, GA, October 2015, pp. 22-33.
- [C8] **H.-M. Chen\***, **A. Arunkumar**, C.-J. Wu, T. Mudge, and C. Chakrabarti, “E-ECC: Low Power Erasure and Error Correction Schemes for Increasing Reliability of Commodity DRAM Systems,” In *Proceedings of International Symposium on Memory Systems (MEMSYS)*, Washington, DC, October 2015, pp. 60-70.
- [C9] **S. Lee\***, N. Singh, P. E. Phelan, and C.-J. Wu, “Harvesting CPU Waste Heat through Pyroelectric Materials,” In *Proceedings of the ASME International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems (InterPACK)*, San Francisco, CA, July 2015 (10 pages).
- [C10] **S.-Y. Lee\***, **A. Arunkumar**, and C.-J. Wu, “CAWA: Coordinated Warp Scheduling and Cache Prioritization for Critical Warp Acceleration for GPGPU Workloads,” In *Proceedings of the ACM/IEEE International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2015, pp. 515-527.
- [C11] C. Gao\*, A. Gutierrez, M. Rajan, R. G. Dreslinski, T. Mudge, and C.-J. Wu, “A Study of Mobile Device Utilization,” In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Philadelphia, PA, March 2015, pp. 225-234.
- [C12] **S. Lee\***, P. E. Phelan, and C.-J. Wu, “Hot Spot Cooling and Harvesting CPU Waste Heat Using Thermoelectric Modules,” In *Proceedings of the ASME 2014 International Mechanical Engineering Congress & Exposition (IMECE)*, Montreal, Canada, November 2014 (10 pages).
- [C13] **A. Arunkumar\*** and C.-J. Wu, “ReMAP: Reuse and Memory Access Cost Aware Eviction Policy for Last Level Cache Management,” In *Proceedings of the International Conference on Computer Design (ICCD)*, Seoul, Korea, October 2014, pp. 110-117.
- [C14] D. Pandiyan\* and C.-J. Wu, “Quantifying the Energy Cost of Data Movement for Emerging Smart Phone Workloads on Mobile Platforms,” In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Raleigh, NC, October 2014, pp. 171-180.

- [C15] **S.-Y. Lee\*** and C.-J. Wu, “CAWS: Criticality-Aware Warp Scheduling for GPGPU Workloads,” In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Edmonton, Canada, August 2014, pp. 175-186.
- [C16] A. Shrivastava\*, A. Rhisjeekeesan, R. Jeyapaul, and C.-J. Wu, “Quantitative Analysis of Control Flow Checking Mechanisms for Soft Errors,” In *Proceedings of the Design Automation Conference (DAC)*, San Francisco, CA, June 2014, pp. 13:1-6.
- [C17] **D. Pandiyan\***, **S.-Y. Lee**, and C.-J. Wu, “Performance, Energy Characterizations and Architectural Implications of An Emerging Mobile Platform Benchmark Suite – MobileBench,” In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, Portland, OR, September 2013, pp. 133-142.
- [C18] C.-J. Wu\*, A. Jaleel, M. Martonosi, S. Steely Jr., and J. Emer, “PACMan: Prefetch-Aware Cache Management for High Performance Caching,” In *Proceedings of the 44<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO)*, Porto Alegre, Brazil, December 2011, pp. 442-453.
- [C19] C.-J. Wu\*, A. Jaleel, W. Hasenplaugh, M. Martonosi, S. Steely Jr., and J. Emer, “SHiP: Signature-Based Hit Predictor for High Performance Caching,” In *Proceedings of the 44<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO)*, Porto Alegre, Brazil, December 2011, pp. 430-441.
- [C20] C.-J. Wu\* and M. Martonosi. “Characterization and Dynamic Mitigation of Intra-Application Cache Interference,” In *Proceedings of IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Austin, TX, April 2011, pp. 2-11. [Best Paper Award Nominee]

## Other Peer-Reviewed, Referred Publications

This category includes: (1) indexed conference and workshop tutorials and abstracts and (2) workshop papers.

- [W1] Y.-J. Yu<sup>x\*</sup>, and C.-J. Wu, “Understanding the Thermal Challenges of High-Performance Mobile Devices Through An Accurate and Detailed Platform Level Temperature Model,” to Appear in *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC)*, October 2017 (Poster Abstract).
- [W2] **S.-Y. Lee\*** and C.-J. Wu, “Characterizing the Latency Hiding Ability of GPUs,” In *Proceedings of IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, March 2013 (Poster Abstract).
- [W3] C.-J. Wu\*, M. Martonosi, “A Comparison of Capacity Management Schemes for Shared CMP Caches,” In *Proceedings of the 7th Annual Workshop on Duplicating, Deconstructing, and Debunking (WDDD)* in conjunction with ISCA-35, Beijing, China, June 2008 (9 pages).

## Patents

- [IP1] *Sustainable Self-Cooling for CPU Hot Spots Using Thermoelectric Materials (Invention ID: M14-257P)*. C.-J. Wu, P. E. Phelan, and S. Lee. Provisional Invention Disclosure filed with ASU AzTE. Utility-application filed.
- [IP2] *Signature Based Hit-Predicting Cache (Application Number: 20140006717)*. S. C. Steely, W. C. Hasenplaugh, A. Jaleel, J. Emer, and C.-J. Wu. Filed in June 2012. Issued in January 2014.

## **PRESENTATIONS**

- [T1] **Holistic Management and Coordination of Performance Quality, Temperature, and Energy Efficiency.**  
*University of Wisconsin – Madison, Georgia Institute of Technology.* April 2017.
- [T2] **Delivering Accelerated Performance Improvement with High Performance GPUs.**  
*University of Rochester.* April 2017.
- [T3] **Designing Smartphones for Performance Quality of Service and High Energy Efficiency Delivery.**  
*Samsung Austin Research Center.* February 2017.
- [T4] **Processor Thermal Management: From Energy Harvesting to Advanced Active Cooling.**  
*MICRO NOPE --- Workshop on Negative Outcomes Post-mortems, and Expectations.* October 2016.
- [T5] **Designing Smartphone System and Architectures for Performance and Energy Efficiency.**  
*ISCA MobiTools --- Tutorial on Infrastructure and Tools for Mobile Computer Architecture Research with an Emphasis on Real System Measurement.* June 2016.
- [T6] **CAWA: Coordinated Warp Scheduling and Cache Prioritization for Critical Warp Acceleration for GPGPU Workloads.**  
*Cornell University, Harvard University, nVIDIA Research, and National Taiwan University.* May—July 2015.
- [T7] **Architectural Thermal Energy Harvesting Opportunities for Sustainable Computing.**  
*HPCA-Best of CAL session.* February 2015.
- [T8] **Sustainable Self-Cooling for CPU Hot Spot.**  
*Intel Thermal and Fluids Lab.* November 2014.
- [T9] **Energy-Efficient Smartphone Architecture.**  
*Qualcomm Research and Science Foundation of Arizona.* July/September 2014.
- [T10] **Energy-Efficient Heterogeneous Architecture – Criticality-Aware Warp Scheduling for General-Purpose GPU Workloads.**  
*AMD and Intel-National Taiwan University.* June/July 2014.
- [T11] **Performance, Energy Characterizations and Architectural Implications of An Emerging Mobile Platform Benchmark Suite – MobileBench.**  
*Marvell.* March 2014.
- [T12] **Energy Efficient Multicore Architecture.**  
*Intel Labs.* September 2013.
- [T13] **Dynamic Techniques for Mitigating Inter- and Intra-Application Cache Interference.**  
*Brown University, Arizona State University, and University of Texas at San Antonio.* February 2012.

- [T14] **Characterizing and Improving Last-Level Cache Management Using Signature-based and Prefetch-aware Approaches.**  
*GigaScale Systems Research Center (GSRC)*. November 2011.

## **TEACHING EXPERIENCE**

### CSE 420 Computer Architecture I: Spring 2013, Fall 2014, and Fall 2016.

Some student comments over the years:

- “[The course is] conceptual and beneficial if you want a career in microprocessor industry. I liked the course because I understood the basics and gained [an] in-depth knowledge of Computer Architecture very clearly and the labs were really helpful to see the implementation of what we were studying in class.”
- “I enjoyed that the professor has a lot of practical experience and related the course to current implementations. Often times we work on 10-15 year old implementations which are of little use but in this course we worked on current (less than 5 year) implementations.”
- “The course content and the professor's clarity in concepts. Dr. Wu is very knowledgeable about the subject and eager to share her knowledge with students.”
- “The assignments were challenging. Also the organization of course is really good.”

### CSE 520 Computer Architecture II: Fall 2012, Fall 2013, Spring 2015, Fall 2015, and Spring 2017.

Some student comments over the years:

- “The structure of the class was good. Emphasis was given on the concepts and keeping up to date by reading recent papers on important topics.”
- “Highly relevant for anyone hoping to pursue industry careers in system level software / hardware jobs.”
- “The instructor was excited about new research area related to the course and motivated students to study and form novel ideas”
- “[What I liked the most about this course is the] Group Research Project and Panel Discussions. Second half was very good compared to part I (This doesn't mean that part I was bad). Professor spent a good amount of time to review each and every project in details. She timely reviewed and gave proper suggestions for the project. Assignment1 was very easy, and Assignment2 was time consuming, it is but much needed. Clarity of the professor on all the topics was excellent, [and] she discussed doubts raised by students in fair manner.”

### CSE 591 Advanced Memory Systems and Architecture: Spring 2015 and Spring 2017.

Some student comments over the years:

- “This is a great course to learn many state-of-the-art designs and a few important legacy designs in the area of memory system design, in one course. This is also a great introductory course to research, for students who are starting/considering to do research.”

## **STUDENT MENTORING**

### *Current Postdoctoral Researchers/Students*

- **Postdoctoral Researcher**

1. Ying-Ju Yu (PhD, Carnegie Mellon University; December 2016 – present)

- **PhD Candidates**

- **Chair**

2. Akhil Arunkumar (January 2013 – present; Expected graduation date: December 15, 2017)
3. Srinath Dasari (August 2016 – present)
4. Jhe-Yu Liou (August 2015 – present)

- **Co-Chair**

5. Benjamin Gaudette (August 2012 – present; Co-advised with Prof. Vrudhula; Expected graduation date: August 31, 2017)

- **Undergraduate Students**

- **Summer REU**

6. TJ Smith (July 2017 – September 2017; REU student from Princeton University)

- **High School Students**

7. Viraj Wadhwa (BASIS Chandler Primary)
8. Katherine Hann (Xavier College Preparatory High School)

### *Alumni*

- **PhD**

- **Chair**

1. Shin-Ying Lee (2017)
  - Dissertation— Intelligent Scheduling and Memory Management for Modern GPU Architectures

- **Co-Chair**

2. Amrit Panda (2014; co-advised with Prof. Chatha)
  - Dissertation— StreamWorks: An Energy-efficient Embedded Co-processor for Stream Computing
  - First employment: Qualcomm Research
3. Soochan Lee (2015; co-advised with Prof. Phelan)
  - Dissertation— A Study of Latent Heat of Vaporization in Aqueous Nanofluids
  - First employment: LG Electronics

- **MS**
  - **Chair**
    4. Dhinakaran Pandiyan (2014)
      - Dissertation— Data Movement Energy Characterization of Emerging Smartphone Workloads for Mobile Platform
      - Best of Computer Engineering/Computer Science Outstanding Graduate Student Award
      - First employment: Intel
    5. Davesh Shingari (2016)
      - Dissertation— Memory Interference Characterization and Mitigation for Heterogeneous Smartphones
      - First employment: Marvell
- **Undergraduate**
  - **Co-Chair**
    6. Kody Stribny (2017, co-advised with Prof. Vrudhula)
      - Dissertation— Mobile Waterway Monitor
      - First employment: Amazon
  - **Fulton Undergraduate Research Initiative**
    7. Ryan Brazones (2014)
      - First employment: Intel
  - **Barret Honors College Contract**
    8. Michael Storto (2014)
      - First employment: General Dynamics

*Thesis Committees*

- **PhD**
  1. Zahra Abbasi – Sustainable Cloud Computing, 2014.
  2. Madhi Hamzeh – Compiler and Architecture Design for Coarse-Grained Programmable Accelerators, 2015.
  3. Priyanka Bagade – Evidence-based Trustworthy Development of Mobile Medical Apps, 2015.
  4. Yooseong Kim – Schedulability-driven Instruction Scratchpad Memory Management for Hard Real-time Systems. School of Computing, Informatics and Decision Systems Engineering, 2017.
  5. Hsing-Min Chen – Using Low Cost Error Correction Schemes to Improve Memory Reliability, 2017.
  6. Jian Cai – Memory Management in Software Managed Multicore Architectures, Current.
- **MS**
  1. Haeseung Lee – Dynamic Scheduling of Stream Programs on Embedded Multi-core Processors, 2013.

2. Digant Pareshkumar Desai – Towards Energy Efficient Computing with Linux: Enabling Task Level Power Awareness and Support for Energy Efficient Accelerator, 2013.
3. Abhishek Rhisheekesan – Control Flow Based Soft Error Protection Mechanism, 2013.
4. Dipal Saluja – Register File Organization for Coarse-Grained Reconfigurable Architectures, 2014.
5. Girish Rao Bulusu – Asymmetric Multiprocessing Real Time Operating System on Multicore Platforms, 2014.
6. Jeevan Prasath – Android Application Context Aware I/O Scheduler, 2014.
7. James Welch – Performance Comparison between Hypervisor- and Container-based Virtualization for Latency-Sensitive Workloads, 2015.
8. Srinivas Karthik Tanikella – GemV: A Validated Micro-architectural Vulnerability Estimation Tool, 2015.
9. Raymond Robles, Analysis and Design of Native File System Enhancements for Storage Class Memory, 2016.

## **SERVICE AND PROFESSIONAL ACTIVITIES**

### **UNIVERSITY AND DEPARTMENT SERVICE**

**Arizona State University** Tempe, AZ, USA

- Building Recruiting And Inclusion for Diversity (BRAID) Organization Committee  
Fall 2014 – Present
- Organization of ASU Booth at the Anita Borg Institute Grace Hopper Celebration of Women in Computing Conference (Record-breaking attendance for ASU, ~100 undergrad/graduate students for GHC-2014, ~50 undergrad/graduate students for GHC-2015, and ~50 undergrad/graduate students for GHC-2016)  
Fall 2014 – Present

#### The ASU Grace Hopper Celebration of Women in Computing Scholarship

I took a leadership role in organizing the annual ASU scholarship program to support female undergraduate and graduate students in computing-related majors to attend the Grace Hopper Celebration of Women in Computing (2014 – Present). Under my organization, we brought around 100, 50, 50 ASU scholarship students to the conference (75% of which are undergraduate students) for professional career development in 2014, 2015, and 2016, respectively.

- We were the number-one attending university in 2014.
- Since the inception of the scholarship program and with other BRAID-related activities, our CIDSE female vs. male undergraduate student ratio has improved from 10% to 15% from 2012 to 2016.
- 89% of the participating students agreed on that attending GHC reinforces their confidence in succeeding in Computer Science.
- 87% of the participating students agreed on that attending GHC reinforces their confidence in solving important world problems.
- 62% of the participating students became more interested in pursuing advanced degree and 53% of the participating students became more interested in pursuing advanced degrees in ASU.

An impact report summarizing the program efforts and outcomes can be found [here](#).

- Next Generation Computing/Trustworthy IoT Faculty Search Committee      Fall 2016 – Spring 2017
- Course Coordinator for CSE420: Computer Architecture      Fall 2014 – Present
- K-14 Girls Make-A-Thon      Fall 2016
- Big Data Architecture Search Committee      Fall 2014 – Spring 2015
- Participation at FSE294 Fulton Undergraduate Research Experience      Fall 2014
- Computer Engineering Graduate Admission Committee      Fall 2013 – 2016
- Fulton Undergraduate Research Intern (FURI) Mentor      Fall 2013 – Spring 2014
- Faculty Peer Mentor of Heni Ben Amor (CIDSE)      Fall 2015 – Spring 2016
- Faculty Peer Mentor of Prof. Jingrui He (CIDSE)      Fall 2014 – Spring 2015
- Faculty Peer Mentor of Prof. Shimeng Yu (CIDSE)      Fall 2013 – Spring 2014
- Organizer of the Smartphone App camp (part of the Best Engineering Science and Technology (BEST) program).      Summer 2013
- Participation at the Ira A. Fulton Schools of Engineering E2 Camp      Summers 2012, 2014-16
- New Faculty Cohort Steering Committee (NFAC) under FSE      Fall 2012 – Spring 2015
- Next Generation Computing Faculty Search Committee      Fall 2012 – Spring 2014
- Participation at the Society of Women Engineers (SWE) Dinner with the Professors      Fall 2012 – Present

### **SCIENTIFIC COMMUNITY**

**IEEE Technical Committee on Computer Architecture (TCCA) Executive Committee, 2017-19.**

### **Best Paper Award Selection Committees**

- The Best Paper Award Selection Committee for the 2016 IEEE International Symposium on High Performance Computer Architecture (HPCA).
- The Best Paper Award Selection Committee for the 2015 IEEE International Symposium on Workload Characterization (IISWC).

### **Conference Technical Program Committees**

The typical review load of a Program Committee member in the top-tier computer architecture conferences, such as the International Symposium on Computer Architecture (ISCA), the International Symposium on Microarchitecture (MICRO), and the International Symposium on High Performance Computer Architecture (HPCA), is 14-22 papers of 11-13 page, double-column, single-space submissions. Each paper typically receives 4-5 reviews, 3 of which come from the Program Committee.

#### **2017**

- The 44<sup>th</sup> ACM/IEEE International Symposium on Computer Architecture (ISCA).
- The 50<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO).
- The 22<sup>nd</sup> ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)\*.
- The 23<sup>rd</sup> IEEE International Symposium on High Performance Computer Architecture (HPCA).
- The 37<sup>th</sup> IEEE International Conference on Distributed Computing Systems (ICDCS).
- The Journal of Instruction-Level Parallelism—the 2<sup>nd</sup> Cache Replacement Championship Workshop (JILP-CRC).
- The IEEE International Symposium on Workload Characterization (IISWC).
- The 14<sup>th</sup> Annual Workshop on Duplicating, Deconstructing and Debunking (WDDD).

#### **2016**

- The 43<sup>rd</sup> ACM/IEEE International Symposium on Computer Architecture (ISCA).
- The 49<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO).
- The 22<sup>nd</sup> IEEE International Symposium on High Performance Computer Architecture (HPCA).
- The 25<sup>th</sup> International Conference on Parallel Architectures and Compilation Techniques (PACT)\*.
- The IEEE International Symposium on Workload Characterization (IISWC).
- The 9<sup>th</sup> General Purpose GPU (GPGPU) Workshop.

#### **2015**

- The 42<sup>nd</sup> ACM/IEEE International Symposium on Computer Architecture (ISCA).
- The 48<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO)\*.
- The 21<sup>st</sup> IEEE International Symposium on High Performance Computer Architecture (HPCA)\*.
- The 33<sup>rd</sup> IEEE International Conference on Computer Design (ICCD).
- The IEEE International Symposium on Workload Characterization (IISWC).
- The 27<sup>th</sup> IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD).
- The ACM/IEEE International Conference on High Performance Computing, Networking, Storage and Analysis (SC15) Technical Posters and ACM Student Research Competition.

#### **2014**

- The 41<sup>st</sup> ACM/IEEE International Symposium on Computer Architecture (ISCA).
- The 47<sup>th</sup> ACM/IEEE International Symposium on Microarchitecture (MICRO).
- The 20<sup>th</sup> IEEE International Symposium on High Performance Computer Architecture (HPCA).
- The IEEE International Symposium on Workload Characterization (IISWC).
- The International Parallel and Distributed Processing Symposium (IPDPS).
- The ACM International Conference on Computing Frontiers (CF).

**2013**

- The 19<sup>th</sup> IEEE International Conference Symposium on High Performance Computer Architecture (HPCA)\*.
- The 18<sup>th</sup> ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)\*.
- The IEEE International Symposium on Workload Characterization (IISWC).
- The International Parallel and Distributed Processing Symposium (IPDPS).

\*External Program Committee Member

**Conference Technical Organizing Committee Memberships**

- Finance Chair for the International Symposium on Workload Characterization (IISWC-2017).
- Student Travel Grant Chair for the International Conference on Parallel Architectures and Compilation Techniques (PACT-2017).
- Poster Chair for the International Symposium on Performance Analysis of Systems and Software (ISPASS-2017).
- Registration Chair for the International Symposium on Code Generation and Optimization (CGO-2017).
- Student Travel Grant Chair for the International Symposium on Microarchitecture (MICRO-2016).
- Workshops/Tutorials Chair for the International Symposium on High Performance Computer Architecture (HPCA-2016).
- Finance Chair for Career Workshop for Women and Minorities in Computer Architecture (CWWMCA) 2014, held in conjunction with the ACM/IEEE International Symposium on Microarchitecture (MICRO-2014).
- Finance Chair for the International Symposium on Workload Characterization (IISWC-2014).
- Student Travel Grant Chair for the International Symposium on Performance Analysis of Systems and Software (ISPASS-2014).
- Publicity and Local Arrangement Chair for the International Symposium on Networks-on-Chip (NOCS-2013).

**Conference Session Chair**

- ISCA-2017 Power and Energy Session
- HPCA-2017 Novel Architectures Session
- MICRO-2016 Multicore Session
- ISCA-2016 Cache Session
- HPCA-2016 GPUs Session
- ISCA-2015 Accelerator Session
- IISWC-2014 GPU Session
- ISCA-2014 Circuits and Architecture Session
- IPDPS-2014 GPU Session
- ISPASS-2014 Cache and Memory Systems Session
- IISWC-2013 Work-in-Progress Session

### **Journal Reviewing**

- ACM Transactions on Architecture and Code Optimization (TACO)
- ACM Transactions on Information and System Security (TISSEC)
- ACM Transactions on Computer Systems (TOCS)
- ACM Transactions on Design Automation of Electronic Systems (TODAES)
- IEEE Transactions on Computers (TOC)
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
- IEEE Computer Architecture Letters (CAL)
- Elsevier Journal of Parallel and Distributed Computing (JPDC)

### **OTHER SERVICE AND OUTREACH ACTIVITIES**

- Served on DOE SBIR Panel 2016
- Served on NSF CCF Panel 2016
- Served on DOE SBIR Panel 2015
- ASU Math Club – “Energy-Efficient Heterogeneous Architecture” Fall 2014
- Served on NSF CCF Panel 2014
- Served as a Grand Award Judge for Intel’s International Science and Engineering Fair (the world’s largest international pre-college science competition) Spring 2013
- Organized and participated the NYC High School Girls Computer Science and Engineering Conference to teach high school girls Android App coding Spring 2012
- Co-founded Princeton’s Computer Architecture Group 2007 – 2012
- Actively involved in the nation-wide Computing Research Association for Women (CRA-W) mentoring program 2007 – 2010
- General Co-Chair of the First Computer Architecture Day @ Princeton 2009