

AKSHITHA SRIRAMAN

Curriculum Vitae — September 2021

Carnegie Mellon University
Department of Electrical and Computer Engineering
4720 Forbes Avenue (Office 4118), Pittsburgh, PA 15213

Phone: +1 (215) 796-3381
Email: akshitha@cmu.edu
Web: <https://akshithasriraman.eecs.umich.edu>

BRIEF BIOGRAPHY

My work bridges computer architecture and software systems, demonstrating the importance of that bridge in realizing efficient web services via solutions that span the systems stack.

Modern web services require data centers that scale to hundreds of thousands of servers, i.e., *hyperscale*. The key challenge in enabling hyperscale web services arises from two important trends. First, over the past few years, there has been a radical shift in hyperscale computing due to an unprecedented growth in data, users, and web service software functionality. Second, modern hardware can no longer support this growth in hyperscale trends due to a steady decline in hardware performance scaling. To enable the new hyperscale era, hardware architects must become more aware of hyperscale software requirements and software researchers can no longer expect unlimited hardware performance scaling. In short, systems researchers can no longer follow the traditional approach of building each systems stack layer separately. Instead, they must rethink the synergy between the software and hardware worlds from the ground up. My work establishes such a synergy to enable futuristic hyperscale web services. Some of my solutions have been deployed on real hyperscale systems and serve billions of users.

I am a Facebook fellow, a Rising Stars in EECS Workshop participant, and a recipient of the Rackham Graduate Fellowship. My work has been recognized with an IEEE Micro Top Picks distinction and has appeared in top architecture and systems venues like OSDI, ISCA, ASPLOS, MICRO, and HPCA.

EDUCATION

Ph.D., Computer Science and Engineering University of Michigan <i>Advisor:</i> Prof. Thomas F. Wenisch <i>Dissertation title:</i> Enabling Hyperscale Web Services	2015 - 2021
M.S., Embedded Systems University of Pennsylvania <i>Advisor:</i> Prof. Joseph Devietti	2013 - 2015
B.E., Electronics and Communication Visvesvaraya Technological University	2008 - 2012

AWARDS AND HONORS

<input type="checkbox"/> Facebook Fellowship \$200,000 towards tuition, stipend, and travel	2020 - 2022
<input type="checkbox"/> Accelerometer selected as an IEEE Micro Top Pick Awarded to the top 12 computer architecture papers of 2020	2020
<input type="checkbox"/> ASPLOS Best Presentation Award Best presentation out of 86 presentations	2020
<input type="checkbox"/> Selected to attend the Heidelberg Laureate Forum	2020
<input type="checkbox"/> CSE Graduate Student Honors (University of Michigan) Won the “best student research” award in the hardware discipline	2020

- **Selected to attend the Rising Stars in EECS Workshop** 2019
- **Facebook Fellowship Finalist** 2019
Recognized as the first runner-up
- **Best of Wild & Crazy Ideas (ASPLOS)** 2019
Chair’s Choice Award
- **Cross-layer Computing Summer School Student Scholarship** 2018
20 winners nation-wide
- **Anita Borg Grace Hopper Scholarship** 2017
- **Rackham Merit Ph.D. Fellowship** 2015
\$140,000 towards tuition, stipend, and travel
- **CIS Full Tuition Scholarship (University of Pennsylvania)** 2014
\$55,000 towards tuition, stipend, and travel
- **Award for academic excellence (Visvesvaraya Technological University)** 2012
Ranked 5th (out of ~10,000 students) in the state
- **“Power Player” awards at Microsoft (India)** 2012

PEER-REVIEWED CONFERENCE/JOURNAL PUBLICATIONS

- Tanvir Ahmed Khan, Nathan Brown, **Akshitha Sriraman**, Niranjana Soundararajan, Rakesh Kumar, Joseph Devietti, Sreenivas Subramoney, Gilles A Pokam, Heiner Litz, Baris Kasikci
Twig: Profile-Guided BTB Prefetching for Data Center Applications. In proceedings of the 54th IEEE/ACM International Symposium on Microarchitecture (**MICRO 2021**). Oct 2021.
Acceptance rate: 94/430 = 21.8%
First instruction prefetching technique to achieve near-ideal BTB performance by using program context information to inform BTB prefetching decisions
- **Akshitha Sriraman**, Abhishek Dhanotia
Understanding Acceleration Opportunities at Hyperscale. In **IEEE Micro**, May-June 2021.
Issue: **Top Picks** in Computer Architecture from Conferences in 2020.
Acceptance: Top 12 computer architecture papers in 2020
Provides key insights on which hyperscale overheads are worth accelerating, and analytically models hardware acceleration benefits to help make well-informed hyperscale hardware investments
- Tanvir Ahmed Khan, Dexin Zhang, **Akshitha Sriraman**, Joseph Devietti, Gilles A Pokam, Heiner Litz, Baris Kasikci
Ripple: Profile-Guided Instruction Cache Replacement for Data Center Applications. In proceedings of the 48th International Symposium on Computer Architecture (**ISCA 2021**). Jun 2021.
Acceptance rate: 76/406 = 18.7%
A novel profile-guided technique that uses program context information to inform the underlying I-cache replacement policy about efficient replacement decisions
- Tanvir Ahmed Khan, **Akshitha Sriraman**, Joseph Devietti, Gilles Pokam, Heiner Litz, Baris Kasikci
I-SPY: Context-Driven Conditional Instruction Prefetching with Coalescing. In proceedings of the 53rd IEEE/ACM International Symposium on Microarchitecture (**MICRO 2020**). Oct 2020.
Acceptance rate: 66/422 = 15.6%
First instruction prefetching technique to achieve near-ideal I-cache performance by conditionally prefetching instructions only when the program context is known to lead to misses
- **Akshitha Sriraman**, Abhishek Dhanotia
Accelerometer: Understanding Acceleration Opportunities for Data Center Overheads at Hyperscale.

In proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS 2020**). Mar 2020. Acceptance rate: 86/476 = 18.1%

IEEE Micro Top Picks

Best Presentation Award

Received the “Artifact Available” and “Artifact Functional” ACM badges

*Analytically models hardware acceleration benefits at hyperscale, and is currently used by hyperscale enterprises (e.g., with developing their encryption accelerator) to make well-informed hardware investments; the acceleration opportunities identified **brought about the new Intel IPU***

- **Akshitha Sriraman**, Abhishek Dhanotia, Thomas F. Wenisch
SoftSKU: Optimizing Server Architectures for Microservice Diversity @Scale. In proceedings of the 46th International Symposium on Computer Architecture (**ISCA 2019**). Jun 2019.
Acceptance rate: 62/365 = 16.9%
Approach and automated tool to improve hyperscale microservice performance on cheap commodity hardware; SoftSKU has been deployed across real-world global data centers, serving billions of users, and has influenced changes in Intel’s Alder Lake (Golden Cove and beyond) server designs
- Amirhossein Mirhosseini, **Akshitha Sriraman**, Thomas F. Wenisch
Enhancing Server Efficiency in the Face of Killer Microseconds. In proceedings of the 25th International Symposium on High-Performance Computer Architecture (**HPCA 2019**). Feb 2019.
Acceptance rate: 46/233 = 19.7%
Solves the infamous “killer microsecond” problem via a new server architecture that schedules latency-insensitive batch threads when a latency-critical microservice awaits a microsecond-scale access
- **Akshitha Sriraman**, Thomas F. Wenisch
μTune: Auto-Tuned Threading for OLDI Microservices. In proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (**OSDI 2018**). Oct 2018.
Acceptance rate: 47/264 = 17.8%
Makes the important observation that no microservice threading model is best across all loads, paving the way for an automatic load adaptation system that tunes threading models to improve efficiency
- **Akshitha Sriraman**, Thomas F. Wenisch
μSuite: A Benchmark Suite for Microservices. In proceedings of the 13th International Symposium on Workload Characterization (**IISWC 2018**). Sep - Oct 2018. Acceptance rate: 17/47 = 36.1%
First benchmark suite of end-to-end web services composed of microservices, facilitating future microservice research
- Liang Luo, **Akshitha Sriraman**, Brooke Fugate, Shiliang Hu, Gilles Pokam, Chris J. Newburn, Joseph Devietti
LASER: Light, Accurate Sharing dEtection and Repair. In proceedings of the 22nd International Symposium on High Performance Computer Architecture (**HPCA 2016**). Mar 2016.
Acceptance rate: 53/240 = 22.0%
Novel low-overhead run-time tool that detects cache contention-induced performance bugs and mitigates them using dynamic binary re-writing

PEER-REVIEWED WORKSHOP PUBLICATIONS & POSTERS

- Lillian Pentecost, Marco Donato, **Akshitha Sriraman**, Gu-Yeon Wei, David Brooks
Analytically Modeling NVM Design Trade-Offs. Non-Volatile Memories Workshop (**NVMW**). Mar 2020.
Analytically models the large design space of emerging memories and their efficiency implications
- Radhika Ghoshal, Yu-Shun Hsiao, **Akshitha Sriraman**, David Brooks
Efficient Event Notification Paradigms for Hyperscale Microservices. Young Architect Workshop held in association with the International Conference on Architectural Support for Programming

Languages and Operating Systems (**YArch - ASPLOS**). Mar 2020.

Analyzes why existing I/O event notification mechanisms fall short at hyperscale

- **Akshitha Sriraman**, Abhishek Dhanotia, Thomas F. Wenisch
Optimizing Server Architectures for Microservice Diversity. Career Workshop for Women and Minorities in Computer Architecture held in association with the International Symposium on Microarchitecture (**CWWMCA - MICRO**). Oct 2019.
Tool to automatically improve microservice performance- and cost-efficiency by customizing hardware and OS knobs in commodity servers
- **Akshitha Sriraman**
Unfair Data Centers for Fun and Profit. In proceedings of the Wild And Crazy Ideas session held in association with the International Conference on Architectural Support for Programming Languages and Operating Systems (**WACI - ASPLOS**). Apr 2019.
Received the Best of WACI, Chair's Choice Award
Proposes investigating the societal and ethical implications of correlating user traits and their acceptable wait times to design user-specific microservice Service Level Objectives
- **Akshitha Sriraman**, Thomas F. Wenisch
Performance-Efficient Notification Paradigms for Disaggregated OLDI Microservices. In proceedings of the Workshop on Resource Disaggregation held in association with the International Conference on Architectural Support for Programming Languages and Operating Systems (**WORD - ASPLOS**). Apr 2019.
New hardware-assisted I/O notification paradigm to achieve near-constant time notification
- Amirhossein Mirhosseini, **Akshitha Sriraman**, Thomas F. Wenisch
Hiding the Microsecond-Scale Latency of Storage-Class Memories with Duplexity. Non-Volatile Memories Workshop (**NVMW**). Mar 2019.
First server architecture to improve server utilization in the presence of microsecond-scale stalls, without sacrificing the Quality-of-Service of a microservice
- **Akshitha Sriraman**, Thomas F. Wenisch
Auto-Tuned Threading for OLDI Microservices. Career Workshop for Women and Minorities in Computer Architecture held in association with the International Symposium on Microarchitecture (**CWWMCA - MICRO**). Oct 2018.
Automatic load adaptation system that tunes threading models to minimize microservice tail latency
- **Akshitha Sriraman**, Thomas F. Wenisch
A Benchmark Suite for Microservices. Workshop on Architectures and Systems for Big Data held in association with the International Symposium on Computer Architecture (**ASBD - ISCA**). Jun 2018.
Suggests how μ Suite can be used by researchers to facilitate future research
- **Akshitha Sriraman**, Thomas F. Wenisch
Performance Characterization of a Taxonomy of Threading Models. Career Workshop for Women and Minorities in Computer Architecture held in association with the International Symposium on Microarchitecture (**CWWMCA - MICRO**). Oct 2017.
Makes the important observation that inherent latency trade-offs between threading models can be exploited at system run-time to minimize microservice tail latency
- **Akshitha Sriraman**, Sihang Liu, Sinan Gunbay, Shan Su, Thomas F. Wenisch
Deconstructing the Tail at Scale Effect Across Network Protocols. Workshop on Duplicating, Deconstructing and Debunking, held in association with the International Symposium on Computer Architecture (**WDDD - ISCA**). Jun 2016.
Establishes that widely-used network protocol stacks can significantly degrade microservice tail latency

DISSERTATION

Akshitha Sriraman. “Enabling Hyperscale Web Services”. PhD dissertation. University of Michigan.

OPEN-SOURCE TOOLS AND INFRASTRUCTURE

Accelerometer: Analytical Model for Hardware Acceleration

Author: Akshitha Sriraman

An analytical model built using C++ for projecting speedup from hardware acceleration for microservice functionalities. (ASPLOS 2020).

Code repository: <https://github.com/akshithasriraman/Accelerometer> & <https://doi.org/10.5281/zenodo.3612797>

μ Tune: A Framework that Auto-Tunes Threading for OLDI Microservices

Author: Akshitha Sriraman

A C++ framework that uses an event-based technique to detect offered load and seamlessly switch between threading models and scale thread pool sizes depending on the load; μ Tune builds on gRPC to abstract complicated threading details from user-level application code. (OSDI 2018).

Code repository: <https://github.com/wenischlab/MicroTune>

μ Suite: A Benchmark Suite for OLDI Microservices

Author: Akshitha Sriraman

The first open-source benchmark suite of end-to-end OLDI services composed of microservices. (IISWC 2018).

Code repository: <https://github.com/wenischlab/MicroSuite>

SELECTED PRESS

- Intel IT Peer Network* Aug 2021
The IPU: A New, Strategic Resource for Cloud Service Providers
<https://itpeernetwork.intel.com/ipu-cloud/gs.bwz5m5>
- Facebook Engineering* May 2020
Accelerometer & SoftSKU: Improving hardware platform performance for diverse microservices
<https://engineering.fb.com/data-center-engineering/accelerometer-and-softsku/>
- Engineering Jobs* May 2020
Accelerometer, SoftSKU for diverse microservices
<https://engineeringjobs4u.co.uk/accelerometer-softsku-for-diverse-microservices>
- TechXplore* Apr 2020
Analytical model predicts exactly how much a piece of hardware will speed up data centers
<https://techxplore.com/news/2020-04-analytical-piece-hardware-centers.html>
- The Michigan Engineer News Center* Apr 2020
Analytical model predicts exactly how much a piece of hardware will speed up data centers
<https://news.engin.umich.edu/2020/04/analytical-model-predicts-exactly-how-much-a-piece-of-hardware-will-speed-up-data-centers/>
- Debug Lies News* Apr 2020
Researchers from Facebook have designed a way to measure exactly how much a hardware accelerator would speed up a datacenter
<https://debuglies.com/2020/04/08/researchers-from-facebook-has-designed-a-way-to-measure-exactly-how-much-a-hardware-accelerator-would-speed-up-a-datacenter/>
- The Michigan Engineer News Center* Jan 2020
Facebook Fellowship for improving high-demand web services
<https://cse.engin.umich.edu/stories/facebook-fellowship-for-improving-high-demand-web-services>

- *The Michigan Engineer News Center* Oct 2019
Two CSE grad students selected for Rising Stars in EECS Workshop
<https://cse.engin.umich.edu/stories/two-cse-grad-students-selected-for-rising-stars-in-eeecs-workshop>
- *Real World Technologies* June 2019
Facebook Workload Analysis
<https://www.realworldtech.com/forum/?threadid=185536curpostid=185539>

PROFESSIONAL EXPERIENCE

- Assistant Professor, **Carnegie Mellon University**, Pittsburgh, PA Jan 2022 -
- Visiting Researcher, **Google**, Pittsburgh, PA Sep 2021 - Dec 2021
Manager: Dr. Hank Levy
Identifying new hardware design and optimization opportunities for hyperscale web services
- Visiting Research Fellow, **University of British Columbia**, Canada May 2020 - Dec 2020
Advisor: Prof. Margo Seltzer
Developing a generic hardware-software interface for diverse hardware accelerators
- Visiting Research Fellow, **Harvard University**, Cambridge, MA Sep 2019 - Apr 2020
Advisor: Prof. David Brooks
Designing future hardware systems for data centers
- Research Scientist, **Facebook Research**, Cambridge, MA Sep 2019 - Apr 2020
Manager: Vijay Balakrishnan
Designing custom hardware for diverse microservice functionalities
- Research Intern, **Facebook Research**, Menlo Park, CA May - Aug 2019
Manager: Abhishek Dhanotia
Developed Accelerometer, an analytical model for hardware acceleration
- Research Engineer, **Facebook Research**, Ann Arbor, MI Sep - Dec 2018
Manager: Murray Stokely
Developed SoftSKU, a strategy to maintain hardware fungibility despite microservice diversity
- Research Intern, **Facebook Research**, Menlo Park, CA May - Aug 2018
Manager: Abhishek Dhanotia
Characterized Facebook's production microservices' system-level and architectural bottlenecks
- Research Intern, **Microsoft Research**, Redmond, WA May - Aug 2017
Manager: Dr. Ed Nightingale
Developed a bare-metal hypervisor from scratch (including a virtualized MMU) to serve as a defense-in-depth security mechanism for Microsoft Azure Sphere; demonstrated two security attacks and defenses
- Research Intern, **Intel Labs**, Santa Clara, CA Jun - Aug 2015
Manager: Dr. Gilles Pokam
Low-overhead run-time tool to detect and mitigate different kinds of cache misses
- Performance Engineer, **Microsoft**, India Jul 2012 - Jun 2013
Manager: Tajdar Salam
Performance analysis of Windows server platforms
- Research Intern, **Hindustan Aeronautics Limited**, India Jan - Mar 2012
Manager: Mohan Rao
Real-time "rotation-per-minute"-based flight warning system for military helicopters/airplanes

TEACHING

University of Pennsylvania , Teaching Assistant with Prof. Joseph Devietti <i>Computer Architecture, Graduate</i>	Spring 2015
University of Pennsylvania , Teaching Assistant with Prof. Camillo J. Taylor <i>Introduction to Computer Systems, Undergraduate</i>	Fall 2014
Invited guest lecture on hyperscale computing <i>CS 146/246: Computer Architecture at Harvard University, Cambridge, MA</i>	Nov 2019
Invited guest lecture on cache coherence protocols <i>CIS 501: Computer Architecture at the University of Pennsylvania, Philadelphia, PA</i>	Apr 2015

RESEARCH ADVISING

Mihailo Rancic (Ph.D. student at Carnegie Mellon University) <i>Identifying new hardware design opportunities and optimizations for hyperscale web services</i>	2021 - Present
Vishwanath Seshagiri (Ph.D. student at Emory University) <i>Dynamically predicting efficient threading design decisions for hyperscale web services</i>	2021 - Present
Shixin Song (Undergraduate student at U. Michigan) <i>Using execution context to improve BTB overheads in data center applications</i>	2021 - Present
Deepanjali Mishra (Undergraduate student at VTU, India) <i>Developing virtual memory optimizations for taint-tracking applications</i>	2021 - Present
Tanvir Ahmed Khan (Ph.D. student at U. Michigan) <i>Using execution context to improve instruction footprint-induced overheads in data center services (MICRO 2020, ISCA 2021, MICRO 2021)</i>	2020 - Present
Sara Mahdizadeh Shahri (Ph.D. student at U. Michigan) <i>Analyzing web service evolution patterns to predict future service behaviors (CWIDCA 2021)</i>	2020 - Present
Dexin Zhang (Undergraduate student at USTC, China) <i>Using program context to design I-cache replacement for data center services (ISCA 2021)</i>	2020 - 2021
Radhika Ghoshal (Ph.D. student at Harvard) <i>Efficient Event Notification Paradigms for Hyperscale Microservices (YArch 2020)</i>	2019 - 2020
Yu-Shun Hsiao (Ph.D. student at Harvard) <i>Efficient Event Notification Paradigms for Hyperscale Microservices (YArch 2020)</i>	2019 - 2020
Lillian Pentecost (Ph.D. student at Harvard) <i>Analytically Modeling NVM Design Trade-Offs (NVMW Poster 2020)</i>	2019 - 2020
Mark Wilkening (Ph.D. student at Harvard) <i>Using service clustering to inform server selection at hyperscale (work under submission)</i>	2019 - 2020

INVITED SEMINAR TALKS

Enabling Hyperscale Web Services

<input type="checkbox"/> Google, Seattle, WA	Jul 2021
<input type="checkbox"/> University of Washington (CSE), Seattle, WA	Apr 2021
<input type="checkbox"/> University of Michigan (CSE), Ann Arbor, MI	Apr 2021
<input type="checkbox"/> University of Illinois Urbana-Champaign (CS), Champaign, IL	Apr 2021

<input type="checkbox"/> <i>University of British Columbia (CS)</i> , Vancouver, Canada	Apr 2021
<input type="checkbox"/> <i>Microsoft Research</i> , Bangalore, India	Apr 2021
<input type="checkbox"/> <i>University of Wisconsin-Madison (CS)</i> , Madison, WI	Mar 2021
<input type="checkbox"/> <i>Brown University (CIS)</i> , Providence, RI	Mar 2021
<input type="checkbox"/> <i>Georgia Institute of Technology (CS & ECE)</i> , Atlanta, GA	Mar 2021
<input type="checkbox"/> <i>École polytechnique fédérale de Lausanne (EPFL) (CS)</i> , Lausanne, Switzerland	Mar 2021
<input type="checkbox"/> <i>University of Chicago (CS)</i> , Chicago, IL	Mar 2021
<input type="checkbox"/> <i>University of Toronto (CS)</i> , Toronto, Canada	Mar 2021
<input type="checkbox"/> <i>Microsoft Research</i> , Redmond, WA	Mar 2021
<input type="checkbox"/> <i>Carnegie Mellon University (ECE)</i> , Pittsburgh, PA	Feb 2021
<input type="checkbox"/> <i>Cornell University (CIS)</i> , Ithaca, NY	Feb 2021
<input type="checkbox"/> <i>University of Pennsylvania (CIS)</i> , Philadelphia, PA	Feb 2021
<input type="checkbox"/> <i>University of Texas, Austin (CS & ECE)</i> , Austin, TX	Feb 2021
<input type="checkbox"/> <i>University of California, Los Angeles (CS)</i> , Los Angeles, CA	Jan 2021
<input type="checkbox"/> <i>University of Waterloo (CS)</i> , Waterloo, Canada	Jan 2021

Understanding Hyperscale Web Services

<input type="checkbox"/> <i>New York University</i> , NYC, NY	Dec 2020
<input type="checkbox"/> <i>Cornell University</i> , Ithaca, NY	June 2020
<input type="checkbox"/> <i>École polytechnique fédérale de Lausanne (EPFL)</i> , Switzerland	May 2020
<input type="checkbox"/> <i>University of Wisconsin</i> , Madison, WI	Mar 2020
<input type="checkbox"/> <i>Google</i> , Madison, WI	Mar 2020
<input type="checkbox"/> <i>Yale University</i> , New Haven, CT	Jan 2020
<input type="checkbox"/> <i>Harvard University</i> , Cambridge, MA	Dec 2019
<input type="checkbox"/> <i>University of Pennsylvania</i> , Philadelphia, PA	Dec 2019
<input type="checkbox"/> <i>Google</i> , Sunnyvale, CA	Jul 2019
<input type="checkbox"/> <i>Brown University</i> , Providence, RI	Apr 2019
<input type="checkbox"/> <i>University of Rhode Island</i> , Kingston, RI	Apr 2019

μ Suite & μ Tune: Auto-Tuned Threading for OLDI Microservices

<input type="checkbox"/> <i>University of California Los Angeles</i> , Los Angeles, CA	Mar 2019
<input type="checkbox"/> <i>Indian Institute of Science</i> , Bangalore, India	Jan 2019
<input type="checkbox"/> <i>Microsoft Research</i> , Bangalore, India	Jan 2019
<input type="checkbox"/> <i>Intel Labs</i> , Bangalore, India	Jan 2019
<input type="checkbox"/> <i>University of California San Diego</i> , San Diego, CA	Oct 2018
<input type="checkbox"/> <i>University of Southern California</i> , Los Angeles, CA	Oct 2018
<input type="checkbox"/> <i>University of Texas, Austin</i> , Austin, TX	Sep 2018

OTHER SELECTED TALKS

- Accelerometer: Understanding Acceleration Opportunities for Data Center Overheads at Hyperscale**
International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Lausanne, CN Mar 2020
- Optimizing Server Architectures for Microservice Diversity**
Career Workshop for Women & Minorities in Comp. Arch. (CWWMCA), Columbus, OH Oct 2019
- Understanding Acceleration Opportunities for Data Center Overheads**
Facebook HQ, Menlo Park, CA Aug 2019
- SoftSKU: Optimizing Server Architectures for Microservice Diversity @Scale**
International Symposium on Computer Architecture (ISCA), Phoenix, AZ Jun 2019
- Unfair Data Centers for Fun and Profit**
Workshop on Wild and Crazy Ideas (WACI), Providence, RI Apr 2019
- Performance-Efficient Notification Paradigms for Disaggregated OLDI Microservices**
Workshop on Resource Disaggregation (WORD), Providence, RI Apr 2019
- Optimizing Server Architectures for Microservice Diversity**
Facebook HQ, Menlo Park, CA Dec 2018
- μ Tune: Auto-Tuned Threading for OLDI Microservices**
Symposium on Operating Systems Design and Implementation (OSDI), Carlsbad, CA Oct 2018
- μ Suite: A Benchmark Suite for Microservices**
International Symposium on Workload Characterization (IISWC), Raleigh, NC Oct 2018
- Auto-Tuned Threading for OLDI Microservices**
Career Workshop for Women & Minorities in Computer Architecture (CWWMCA), Japan Oct 2018
- A Comprehensive Characterization of Facebook's Heavy Hitter Microservices**
Facebook HQ, Menlo Park, CA Aug 2018
- A Benchmark Suite of Microservices**
Workshop on Architectures and Systems for Big Data (ASBD), Los Angeles, CA June 2018
- A Benchmark Suite of Microservices**
Intel VEC retreat, Ann Arbor, MI June 2018
- Characterization of a Taxonomy of Threading Models**
Career Workshop for Women & Minorities in Comp. Arch. (CWWMCA), Boston, MA Oct 2017
- Hypervisor-Based Defense-In-Depth for Microsoft Azure Sphere**
Microsoft Research, Redmond, WA Aug 2017
- A Case Study Characterizing Bottlenecks in High Dimensional Search**
CRA-Women Grad Cohort Workshop, Washington D.C. Apr 2017
- Data Center-Scale System Support for Encyclopedic Recognition**
Intel VEC retreat, Santa Clara, CA Dec 2016
- Imagen: Custom Scaled-Out High Dimensional Search**
ARM retreat, Ann Arbor, MI Nov 2016
- Deconstructing the Tail at Scale Effect Across Network Protocols**
Workshop on Duplicating, Deconstructing and Debunking (WDDD), Seoul, Korea Jun 2016

4C'sONS Haswell: 4C's - ONline cache profiling on Server platforms*Intel Labs, Santa Clara, CA*

Aug 2015

Crash Prevention System in a Helicopter*Hindustan Aeronautics Limited, Bangalore, India*

Jan 2012

PROFESSIONAL SERVICE

Program Committee Member

- IEEE Micro Top Picks 2022
- Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2022
- EuroSys 2022
- Eurosys Doctoral Workshop (EuroDW) 2021
- ACM Symposium on Cloud Computing (SoCC) 2020
- Young Architect Workshop (YArch-ASPLOS) 2020 - 2021

External Review Committee Member

- International Symposium on Microarchitecture (MICRO) 2021
- Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2021

Artifact Evaluation Committee Member

- Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2020
- Symposium on Operating Systems Principles (SOSP) 2019

Conference Shadow Program Committee Member

- EuroSys 2018 - 2019
- Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2017

Invited Reviewer

- ACM SIGMETRICS 2019
- International Symposium on Microarchitecture (MICRO) 2016

Journal Reviewer

- ACM Transactions on Computer Systems (TOCS) 2021 - Present
- ACM Transactions on Architecture and Code Optimization (TACO) 2018 - 2020

Workshop Co-organizer

- Young Architect Workshop (YArch) at ASPLOS 2022, 2020
- JOBS Workshop at MICRO 2020
- Career Workshop for Women & Minorities in Computer Architecture at MICRO 2019

Technical Blog Editor

- SIGOPS Blog 2020 - Present

Session Chair

- ACM Symposium on Cloud Computing (SoCC) 2020

Web Chair

- International Symposium on Low Power Electronics and Design (ISLPED) 2018 - 2020

Publicity Chair

- Architectural Support for Programming Languages & Operating Systems (ASPLOS) 2022, 2020
- Young Architect Workshop (YArch-ASPLOS) 2020
- International Symposium on Microarchitecture (MICRO) 2019

Faculty Candidate Hiring Committee (Student-Run Interviews)

- University of Michigan, Computer Science Department 2020

Ph.D. Admissions Committee

- Carnegie Mellon University, Electrical & Computer Engineering Department 2021 - 2022
- University of Michigan, Computer Science Department 2019

Student Organizer

- IEEE Micro Top Picks 2018
- University of Michigan Ph.D. prospective student visit day 2018
- Explore Grad Studies in CSE Workshop, University of Michigan 2016

OUTREACH ACTIVITIES

Female Mentoring

- Vidushi Goyal, Ph.D. student (U. Michigan) 2016 - Present
- Harini Muthukrishnan, Ph.D. student (U. Michigan) 2016 - Present
- Hiwot Tadese Kassa, Ph.D. student (U. Michigan) 2017 - Present
- Amani Alkayyali, Ph.D. student (U. Michigan) 2019 - Present
- Katie Lim, Ph.D. student (U. Washington) 2019 - Present
- Lillian Pentecost, Ph.D. student (Harvard University) 2019 - Present
- Sara Mahdizadeh Shahri, Ph.D. student (Penn State) 2019 - Present
- Aninda Manocha, Ph.D. student (Princeton) 2020 - Present
- Ketaki Joshi, Ph.D. student (Yale) 2020 - Present
- Shixin Song, Undergraduate student (U. Michigan) 2021 - Present
- Deepanjali Mishra, Undergraduate student (VTU) 2021 - Present
- Katia Flores, Undergraduate student (U. Michigan) 2018 - 2019
- Linh Le, Undergraduate student (U. Michigan) 2018 - 2019

Women In Computer Architecture (WICArch) Webinar Series Lead 2018 - Present
Organizing monthly webinars for women in computer architecture

WICArch Mentoring Series Co-organizer 2018 - Present
Organizing a mentorship program for female students in computer architecture

Middle School Outreach Co-organizer, Ann Arbor, MI 2017 - 2019
Created the “Middle School Outreach” program to get middle school students from underserved groups interested in CS early on; designed curriculum, trained and hired instructors, secured funding, etc

Middle School Teacher, Scarlett Middle School, Ann Arbor, MI 2018 - 2019
Taught computer science basics to middle school students from underserved groups

Ensemble of CSE Ladies Officer, University of Michigan, Ann Arbor, MI 2018 - 2019
Co-ordinated activities for a female graduate student support organization

CS Kickstart Hackathon Co-organizer, University of Michigan, Ann Arbor, MI Sep 2016
Workshop aimed at improving gender diversity in CSE through increased female enrollments

Girls Encoded Co-organizer (along with Prof. Reetuparna Das), Ann Arbor, MI Mar 2016
Workshop aimed at getting high school female students interested in computer science