+1 (734) 773-5372
tk3070@columbia.edu
web.eecs.umich.edu/~takh

Tanvir Ahmed Khan

Assistant Professor, Columbia University in the City of New York

Research Interests

I build systems techniques to enable efficient data center processing. Efficient data center processing is challenging due to rapid growth in data and software complexity along with the ongoing slowdown in hardware performance scaling. As a consequence, micro-architectural structures (*e.g.*, instruction and data cache, branch predictors) can no longer meet the demand of data center applications. Combining insights from computer architecture, compilers, and operating systems, I design techniques to enable near-ideal micro-architectural structures via profile-guided optimizations. Consequently, Intel and ARM have adopted a couple of my techniques. Bridging hardware and software, my work appears in venues like ISCA, MICRO, ASPLOS, OSDI, PLDI, FAST, and EuroSys. My research has also received 2024 ACM SIGARCH/IEEE CS TCCA Outstanding Dissertation Award, IEEE Micro Top Picks 2023 distinction, and MICRO 2022 Best Paper Award.

Education

2017-2023 **Ph.D.**, *University of Michigan*, Ann Arbor, Michigan, USA
- Computer Science and Engineering
- Thesis: Rescuing Data Center Processors

2017 **M.Sc.**, *Bangladesh University of Engineering and Technology*, Dhaka, Bangladesh
- Computer Science and Engineering

2014 **B.Sc.**, *Bangladesh University of Engineering and Technology*, Dhaka, Bangladesh
- Computer Science and Engineering
- Class Rank: 1/153

Awards and Honors

2024 ACM SIGARCH/IEEE CS TCCA Outstanding Dissertation Award
2023 Richard F. and Eleanor A. Towner Prize for Outstanding Ph.D. Research ($2,500), University of Michigan
2023 IEEE Micro Top Pick (awarded to the top 12 computer architecture papers of 2022)
2023 DATE Best Paper Award Nomination
2022-2023 Rackham Predoctoral Fellowship ($100,000 towards tuition, stipend, and insurance), University of Michigan
2022 MICRO Best Paper Award (top two of 83 papers, top 1% of 348 submissions)
2022 Qualcomm Innovation Fellowship Finalist (one of 46 finalists)
2021 Graduate Student Honors (Best PhD Research Award in the Michigan Systems Lab), University of Michigan
2020 Facebook Fellowship Finalist (top 4% of 1800 applicants)
2017-2018 Rollin M. Gerstacker Foundation Fellowship ($100,000 towards tuition, stipend, and insurance)
2014 Crest of Honor, Highest CGPA in the department, presented by BUET alumni association
2009-2014 University Merit Scholarship, Bangladesh University of Engineering and Technology
2009-2014 Dean's List Scholarship, Bangladesh University of Engineering and Technology

Mentee Awards and Honors

2022 First place in ACM undergraduate student research competition (MICRO'22), Kan Zhu
2021 CRA outstanding undergraduate researcher award honorable mention, Shixin Song
2021 First place in ACM undergraduate student research competition (MICRO'21), Shixin Song

2020 First place in ACM undergraduate student research competition (CGO'20), Nathan Brown

Peer-Reviewed Conference Publications

HPCA'25 Kan Zhu, Yilong Zhao, Yufei Gao, Peter Braun, **Tanvir Ahmed Khan**, Heiner Litz, Baris Kasikci, and Shuwen Deng
From Optimal to Practical: Efficient Micro-op Cache Replacement Policies for Data Center Applications.
In proceedings of the 31st High-Performance Computer Architecture (**HPCA**), 2025.
The first micro-op cache replacement policy.

ISCA'24 Surim Oh, Mingsheng Xu, **Tanvir Ahmed Khan**, Baris Kasikci, and Heiner Litz
UDP: Utility-Driven Fetch Directed Instruction Prefetching.
In proceedings of the 51th International Symposium on Computer Architecture (**ISCA**), 2024.
First utility-based mechanism for instruction prefetching that ensures both accuracy and timeliness

ASPLOS'24 Yuxuan Zhang, Nathan Sobotka, Soyoon Park, Saba Jamilan, **Tanvir Ahmed Khan**, Baris Kasikci, Gilles A Pokam, Heiner Litz, and Joseph Devietti
RPG^2: Robust Profile-Guided Runtime Prefetch Generation.
In proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS**), 2024.
An online system that injects and tunes prefetch instructions to make prefetching more robust optimization

DATE'23 Yuhan Chen[G], Alireza Khadem, Xin He, Nishil Talati, **Tanvir Ahmed Khan**, and Trevor Mudge
PEDAL: A Power Efficient GCN Accelerator with Multiple DAtafLows.
In proceedings of the 26th Design, Automation, and Test in Europe (**DATE**) conference, 2023.
Best Paper Award Nomination
An accelerator for graph convolutional neural networks that automatically selects optimized dataflow and phase ordering

MICRO'22 **Tanvir Ahmed Khan**, Muhammed Ugur[U], Krishnendra Nathella, Dam Sunwoo, Heiner Litz, Daniel A. Jiménez, and Baris Kasikci
Whisper: Profile-Guided Branch Misprediction Elimination for Data Center Applications.
In proceedings of the 55th IEEE/ACM International Symposium on Microarchitecture (**MICRO**), IEEE, 2022.
Best Paper Award
First branch prediction technique that identifies precise program contexts leading to branch mispredictions and encodes corresponding hard-to-predict correlations in branch history efficiently using Boolean formulas

MICRO'22 Yuxuan Zhang[G], **Tanvir Ahmed Khan**, Gilles Pokam, Baris Kasikci, Heiner Litz, and Joseph Devietti
OCOLOS: Online COde Layout OptimizationS.
In proceedings of the 55th IEEE/ACM International Symposium on Microarchitecture (**MICRO**), IEEE, 2022.
IEEE Micro Top Pick
First online system that enables profile-guided optimizations of unmanaged applications on a running process

ISCA'22 Shixin Song[U], **Tanvir Ahmed Khan**, Sara Mahdizadeh Shahri[G], Akshitha Sriraman, Niranjan K Soundararajan, Sreenivas Subramoney, Daniel A. Jiménez, Heiner Litz, and Baris Kasikci
Thermometer: Profile-Guided BTB Replacement for Data Center Applications.
In proceedings of the 49th International Symposium on Computer Architecture (**ISCA**), ACM, 2022.
First replacement technique that realizes applications' holistic behavior to avoid branch target mispredictions

EuroSys'22 Saba Jamilan[G], **Tanvir Ahmed Khan**, Grant Ayers, Baris Kasikci, and Heiner Litz
APT-GET: Profile-Guided Timely Software Prefetching.
In proceedings of the 17th European Conference on Computer Systems (**EuroSys**), ACM, 2022.
First data prefetching technique that ensures prefetch timeliness using Intel LBR's cycle information

MICRO'21 **Tanvir Ahmed Khan**, Nathan Brown[U], Akshitha Sriraman, Niranjan Soundararajan, Rakesh Kumar, Joseph Devietti, Sreenivas Subramoney, Gilles Pokam, Heiner Litz, and Baris Kasikci
Twig: Profile-Guided BTB Prefetching for Data Center Applications.
In proceedings of the 54th IEEE/ACM International Symposium on Microarchitecture (**MICRO**), IEEE, 2021.
Semiconductor Research Corporation (SRC) Best Paper, Q3'2021
First prefetching technique that enables near-ideal branch predecoding to avoid branch target mispredictions

MICRO'21 Niranjan Soundararajan, Peter Braun[G], **Tanvir Ahmed Khan**, Baris Kasikci, Heiner Litz, and Sreenivas Subramoney
PDede: Partitioned, Deduplicated, Delta Branch Target Buffer.
In proceedings of the 54th IEEE/ACM International Symposium on Microarchitecture (**MICRO**), IEEE, 2021.
First architectural design that stores branch targets efficiently by removing redundancies among branches

OSDI'21 **Tanvir Ahmed Khan**, Ian Neal, Gilles Pokam, Barzan Mozafari, and Baris Kasikci
DMon: Efficient Detection and Correction of Data Locality Problems using Selective Profiling.
In proceedings of the 15th USENIX Symposium on Operating Systems Design and Implementation (**OSDI**), USENIX Association, 2021.
Received the "Artifact Available" USENIX badge
First profiling technique that enables in-production profiling without any overhead, guides optimizations to make real-world workloads twice as faster, and has been adopted in the Arm Neoverse N1 Core

ISCA'21 **Tanvir Ahmed Khan**, Dexin Zhang[U], Akshitha Sriraman, Joseph Devietti, Gilles Pokam, Heiner Litz, and Baris Kasikci
Ripple: Profile-Guided Instruction Cache Replacement for Data Center Applications.
In proceedings of the 48th International Symposium on Computer Architecture (**ISCA**), IEEE, 2021.
First cache replacement technique that uses program context to make efficient code replacement decisions

FAST'21 Ian Neal, Gefei Zuo, Eric Shiple, **Tanvir Ahmed Khan**, Youngjin Kwon, Simon Peter, and Baris Kasikci
Rethinking File Mapping for Persistent Memory.
In proceedings of the 19th USENIX Conference on File and Storage Technologies (**FAST**), USENIX, 2021.
Optimizes the file offsets to persistent memory mapping operation, providing up to 45% performance gains

MICRO'20 **Tanvir Ahmed Khan**, Akshitha Sriraman, Joseph Devietti, Gilles Pokam, Heiner Litz, and Baris Kasikci
I-SPY: Context-Driven Conditional Instruction Prefetching with Coalescing.
In proceedings of the 53rd IEEE/ACM International Symposium on Microarchitecture (**MICRO**), IEEE, 2020.
Semiconductor Research Corporation (SRC) Best Paper, Q3'2020
First instruction prefetching technique that enables conditional prefetching only when program context leads to instruction cache misses; I-SPY achieves near-ideal cache performance and has been adopted by Intel

IISWC'20 Yuhan Chen[U], Jingyuan Zhu[U], **Tanvir Ahmed Khan**, and Baris Kasikci
CPU Microarchitectural Performance Characterization of Cloud Video Transcoding.
In proceedings of the IEEE International Symposium on Workload Characterization (**IISWC**), IEEE, 2020.
Finds key bottlenecks in instruction cache, data cache, and branch predictor for video transcoding workloads

PLDI'19 **Tanvir Ahmed Khan**, Yifan Zhao[U], Gilles Pokam, Barzan Mozafari, and Baris Kasikci
Huron: Hybrid False Detection and Repair.
In proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (**PLDI**), ACM, 2019.
Received the "Artifacts Available" and "Artifacts Functional" ACM badges
Semiconductor Research Corporation (SRC) Best Paper, Q2'2019
A system that makes parallel programs up to 8× faster by automatically detecting and repairing false sharing

Peer-Reviewed Journal/Workshop Publications and Posters

Underlined authors are undergraduate[U] and graduate[G] students mentored by me.

IEEE Micro'23 — Yuxuan Zhang[G], **Tanvir Ahmed Khan**, Gilles Pokam, Baris Kasikci, Heiner Litz, and Joseph Devietti
Online COde Layout OptimizationS via OCOLOS.
IEEE Micro, IEEE, 2023.

OSR'22 — Muhammed Ugur[U], Cheng Jiang[U], Alex Erf[U], **Tanvir Ahmed Khan**, and Baris Kasikci
One Profile Fits All: Profile-Guided Linux Kernel Optimizations for Data Center Applications.
ACM Special Interest Group on Operating Systems (SIGOPS) Operating Systems Review, ACM, 2022.

WoSC'22 — Truls Asheim[G], **Tanvir Ahmed Khan**, Baris Kasikci, and Rakesh Kumar
Impact of Microarchitectural State Reuse on Serverless Functions.
In proceedings of the 8th International Workshop on Serverless Computing (WoSC), ACM, 2022.

Wireless Networks'19 — **Tanvir Ahmed Khan** and A. B. M. Alim Al Islam
Enhancing Throughput in Multi-Radio Cognitive Radio Networks.
Wireless Networks, Springer, 2019.

MobiSys'16 — **Tanvir Ahmed Khan** and A. B. M. Alim Al Islam
Poster: Overcoming Throughput Degradation in Multi-Radio Cognitive Radio Networks.
In proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys) Companion, ACM, 2016.

WiMob'15 — **Tanvir Ahmed Khan**, Chowdhury Sayeed Hyder, and A. B. M. Alim Al Islam
Towards Exploiting a Synergy between Cognitive and Multi-Radio Networking.
In proceedings of the 11th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), IEEE, 2015.

Employment

2024-present — **Columbia University in the City of New York**, New York, USA
- Assistant Professor, Electrical Engineering Department
- Fu Foundation School of Engineering and Applied Science

2017-2023 — **University of Michigan**, Ann Arbor, Michigan, USA
- Research Assistant, Electrical Engineering and Computer Science
- Advisor: Baris Kasikci

Summer 2021 — **ARM**, Austin, Texas, USA
- Research Intern, Microarchitecture Research Group
- Mentor: Krishnendra Nathella and Dam Sunwoo

Summer 2020 — **Facebook**, Menlo Park, California, USA
- Software Engineer Intern, Binary Optimization and Layout Tool (BOLT) Team
- Mentor: Maksim Panchenko

Summer 2019 — **Microsoft**, Redmond, Washington, USA
- Research Intern, Azure Hardware Research Group
- Mentor: Gagan Gupta and Rathijit Sen

2014-2017 — **Bangladesh University of Engineering and Technology**, Dhaka, Bangladesh
- Lecturer, Department of Computer Science and Engineering

Teaching

Fall 2024 **Columbia University in the City of New York**, EECS 6894
- **Primary Instructor** for the course, Hardware/Software Co-Design for Data Center Processing

Spring 2024 **Columbia University in the City of New York**, CSEE 4824
- **Primary Instructor** for the course, Computer Architecture

Winter 2022 **University of Michigan**, EECS 582
- **Primary Instructor** for the graduate course, Advanced Operating Systems

Spring 2022 **Carnegie Mellon University**, 18-847C
- Invited guest lecture on profile-guided compiler optimizations for Data Center Computing

Winter 2019 **University of Michigan**, EECS 570
- Teaching Assistant with Prof. Thomas Wenisch for the graduate course, Parallel Computer Architecture

2014-2017 **Bangladesh University of Engineering and Technology**, CSE 305, CSE 313, CSE 309
- **Primary Instructor** for the undergraduate course, Computer Architecture
- **Primary Instructor** for the undergraduate course, Operating Systems
- **Primary Instructor** for the undergraduate course, Compilers

Research Mentoring

2024-present Mingkai Li (Columbia PhD)
Tiered Memory Safety

2024-present Matt Weingarten (Columbia PhD)
Profile-Guided Memory Layout Optimizations

2024-present Mohammad Tawhid Bhuiyan (Columbia PhD)
Stale Profile Inference

2024-present Chenyang Zhou (Columbia MSc)
Compiler-Driven Simulation of Chained Data Center Accelerators

2024-present Zhixuan Zhang (Columbia MSc)
Optimizing BTB Hierarchy for Data Center Applications

2024-present Hiba Altaf (Columbia BSc)
False Sharing Bug Detection and Correction with Large Language Models

2023-present Sumya Hoque (Columbia BSc)
Profile-Guided Optimizations for Microservices

2022-present Pooneh Safayenikoo (UCSC PhD)
Data-driven far memory allocation, prefetching, and replacement

2022-2023 Yiwei Yang (UCSC PhD)
Emulating memory disaggregation for data center applications

2022-2023 Woojin Jung (Cranbrook High School)
Performance characterization and optimization of distributed gradient boosting workloads

2022-2023 Kan Zhu (UM BSc → University of Washington PhD)
Architectural implications of Google's data center applications
First place in ACM undergraduate student research competition (MICRO'22)

2022-2023 Diane Chiang (UM BSc → University of Washington MSc)
Visualization of thread interleavings among Google's data center applications

2022-2023 Fangjia Shen (Purdue PhD)
Performance characterization and acceleration of Apache Spark workloads

2021-2023 Saba Jamilan (UCSC PhD)
Profile-guided timely software prefetching and replacement
First author of APT-GET [EuroSys'22]

2021-2023 Surim Oh (UCSC PhD)
Precise and timely wrong-path prefetching

2021-2023 Ali Ansari (EPFL PhD)
Architectural and systems-level implications of CloudSuite-4.0 applications

2021-2023 Shanqing Lin (EPFL PhD)
Systems-level architectural simulation of CloudSuite-4.0 applications

2021-2023 Truls Asheim (NTNU PhD)
Analyzing microarchitectural behaviour of serverless functions
First author of the WoSC'22 paper

2020-2023 Peter Braun (UCSC PhD)
Performance characterization and optimization of modern processor's frontend for data center applications
Co-author of PDede [MICRO'21]

2020-2023 Yuxuan Zhang (UPenn PhD)
Online profile-guided optimizations for C/C++ applications
First author of OCOLOS [MICRO'22]

2020-2023 Yuhan Chen (UM BSc → UM PhD)
Performance characterization and acceleration of irregular workloads
First author of the IISWC'20 paper
First author of the DATE'23 paper

2020-2022 Sara Mahdizadeh Shahri (UM PhD → CMU PhD)
Proxy-web: a proxy app suite for production web services
Co-author of Thermometer [ISCA'22]

2021-2022 Shixin Song (UM BSc → MIT PhD)
Profile-guided BTB replacement for data center applications
First author of Thermometer [ISCA'22]
First place in ACM undergraduate student research competition (MICRO'21)
CRA outstanding undergraduate researcher award honorable mention, 2022

2020-2022 Muhammed Ugur (UM BSc, MSc → Yale PhD)
Profile-guided Linux Kernel optimizations for data center applications
Co-author of Whisper [MICRO'22]
First author of One Profile Fits All [OSR'22]

2022 Zhenhang He (UM BSc → Pinterest)
Enabling microarchitectural simulations of data center applications

2021 Scott Hadley (UM BSc → ARM)
Enabling microarchitectural simulations of managed workloads

2019-2021 Nathan Brown (UM BSc, MSc → ARM)
Profile-guided instruction cache and BTB prefetching for data center applications
Co-author of Twig [MICRO'21]
First place in ACM undergraduate student research competition (CGO'20)

2020 Dexin Zhang (USTC BSc → USTC PhD)
Profile-guided instruction cache replacement for data center applications
Co-author of Ripple [ISCA'21]

2020 Ashfaqur Rahaman (BUET BSc → Utah PhD)
Load-time code layout optimizations

2020 Yineng Yan (UM BSc → UT Austin PhD)
Record and replay debugging for arbitrary GPU programs

2019 Shariq Hafeez (UM BSc → Citadel)
Optimizing data locality for stash databases

2019 Zhiqi Chen (UM BSc → CMU MSc)
Optimizing data locality for MySQL

2019 Xiaohe Cheng (HKUST BSc → Google)
Optimizing data locality for Memcached, Redis, SQLite, Graph500, and XStream

2018-2019 Yifan Zhao (UM BSc → UIUC PhD)
Hybrid false sharing detection and repair
Co-author of Huron [PLDI'19]

Grants and Gifts

2024 **Google gift, 35K USD**
A Data Center Simulator Enabling Large Design Space Exploration of Distributed Accelerators
Principal Investigators: Tanvir Ahmed Khan

2023 **Google gift, 300K USD**
Rethinking Micro-Architectural and Operating Systems Abstractions for Data Center Applications
Principal Investigators: Tanvir Ahmed Khan and Baris Kasikci

2022 **Intel TSA grant, 600K USD**
Data-Driven Processor Design for Datacenter Applications
Principal Investigators: Baris Kasikci, Daniel A. Jiménez, Heiner Litz, and Akshitha Sriraman

2021 **Semiconductor Research Corporation (SRC) realignment grant, 406K USD**
Proxy-Web: A Proxy App Suite for Production Web Services
Principal Investigators: Baris Kasikci, Timothy Rogers, and David Brooks

2020 **NSF/Intel Partnership on Foundational Microarchitecture Research (FoMR), 360K USD**
Taming the Instruction Bottleneck in Modern Datacenter Applications
Principal Investigators: Baris Kasikci and Joseph Devietti

2020 **Semiconductor Research Corporation (SRC) seed grant 150K USD**
Proxy-Web: A Proxy App Suite for Production Web Services
Principal Investigators: Baris Kasikci, Timothy Rogers, and David Brooks

Open-Source Software and Tools

September 2022 **OCOLOS**
The artifact implements the first online code layout optimization system (MICRO 2022) for applications written in C/C++ languages and allows profile-guided optimization to be performed on a running process. https://github.com/upenn-acg/Ocolos-MySQL

July 2022 **One Profile Fits All**
The artifact consists of software and source code of Linux Kernel's profile-guided optimizations (OSR 2022 and Linux Plumbers Conference 2021) and evaluation scripts for eight open-source data center applications (Apache, Nginx, Redis, Memcached, Leveldb, Rocksdb, MySQL, and PostgreSQL).
https://github.com/efeslab/lk-profile and https://github.com/facebookincubator/BOLT

June 2022 **Thermometer**
The artifact consists of Thermometer's (ISCA 2022) software and source code, evaluation, and instructions.
https://github.com/efeslab/thermometer-artifact

July 2021 **DMon**
The artifact contains the prototype of DMon (OSDI 2021) implementing selective profiling, a technique that locates data locality problems with low-enough overhead that is suitable for production use.
https://github.com/efeslab/dmon-ae

June 2021 **I-SPY** and **Ripple**
The artifact consists of software and source code of I-SPY (MICRO 2020) and Ripple (ISCA 2021), a VirtualBox image with pre-installed data center applications, and these applications' traces.
https://github.com/efeslab/ispy-ripple

February 2021 **NVM File Indexing**
The open-source file system is optimized for file mapping on persistent/non-volatile memory (FAST 2021).
https://github.com/efeslab/nvm-file-indexing

May 2019 **Huron**
The artifact consists of a VirtualBox image containing the Huron (PLDI 2019) software and source code, evaluation data, instructions, and all open source evaluation applications and input files.
https://github.com/efeslab/huron and https://doi.org/10.1145/3325966

Selected Press

March 2022 **Tanvir Ahmed Khan awarded Rackham Predoctoral Fellowship**
The Michigan Engineer News Center
https://cse.engin.umich.edu/stories/tanvir-ahmed-khan-awarded-rackham-predoctoral-fellowship

November 2021 **Outstanding research recognized at Graduate Honors Competition**
The Michigan Engineer News Center
https://cse.engin.umich.edu/stories/outstanding-research-recognized-at-graduate-honors-competition

October 2021 **Arm Neoverse N1 – Performance Analysis Methodology to Tune Production Systems and Application Code**
Arm Community Blogs
https://community.arm.com/arm-community-blogs/b/tools-software-ides-blog/posts/arm-neoverse-n1-performance-analysis-methodology

September 2021 **Facebook Has Been Working On BOLT'ing The Linux Kernel For Greater Performance**
Phoronix
https://www.phoronix.com/news/Facebook-BOLTing-The-Kernel

April 2020 **Undergraduate research on speeding up data centers earns ACM first prize**
The Michigan Engineer News Center
https://cse.engin.umich.edu/stories/undergraduate-research-on-speeding-up-data-centers-earns-acm-first-prize

August 2019 **Automated tool optimizes complex programs better than humans**
The Michigan Engineer News Center
https://cse.engin.umich.edu/stories/automated-tool-optimizes-complex-programs-better-than-humans

February 2019 **Speeding up code with clever data manipulation**
The Michigan Engineer News Center
https://cse.engin.umich.edu/stories/speeding-up-code-with-clever-data-manipulation

Invited Talks

Filling the RPC Gap Between Google and Open-Source Microservices

April 2024 Google Simulation Summit (Sunnyvale)
Host: Victor Lee

Rescuing Data Center Processors

September 2023 ACCESS-CEDA Seminar Series, Hong Kong University of Science and Technology (HKUST)
Host: Zhiyao Xie

August 2023 Azure Systems Research group
Host: Esha Choukse and Daniel S. Berger

August 2023 Samsung Tech. Seminar
Host: Diane Jung

July 2023 Shanghai Jiao Tong University
Host: Zhengwei Qi

April 2023 Cornell University (Department of Computer Science)
Host: Giulia Guidi

April 2023 Stony Brook University (Department of Computer Science)
Host: Dongyoon Lee

April 2023 Texas A&M University (Department of Computer Science and Engineering)
Host: Daniel A. Jiménez

March 2023 University of Waterloo (Department of Electrical and Computer Engineering)
Host: Hiren Patel

March 2023 University of Waterloo (David R. Cheriton School of Computer Science)
Host: Martin Karsten

March 2023 Rutgers University (Department of Computer Science)
Host: Richard P. Martin

March 2023 École Polytechnique Fédérale de Lausanne (School of Computer and Communication Sciences)
Host: Sanidhya Kashyap

March 2023 University of Minnesota Twin Cities (Department of Electrical and Computer Engineering)
Host: Ulya Karpuzcu

March 2023 University of North Carolina at Chapel Hill (Department of Computer Science)
Host: Shahriar Nirjon

March 2023 University of Toronto (Department of Computer Science)
Host: Gennady Pekhimenko

March 2023 Columbia University (Electrical Engineering Department)
Host: Asaf Cidon

February 2023 University of Illinois Urbana-Champaign (Electrical and Computer Engineering Department and Department of Computer Science)
Hosts: Nam Sung Kim and Tianyin Xu

February 2023 Northeastern University (College of Engineering-Electrical and Computer Engineering)
Host: Yanzhi Wang

February 2023 Purdue University (Department of Computer Science)
Host: Xuehai Qian

February 2023 New York University (NYU) Courant Institute of Mathematical Sciences (Computer Science Department)
Host: Lakshminarayanan Subramanian

February 2023 — Boston University (Department of Computer Science)
Host: Jonathan Appavoo

February 2023 — Boston University (Department of Electrical and Computer Engineering)
Host: Manuel Egele

December 2022 — Columbia University
Host: Baishakhi Ray

October 2022 — Intel Labs ArchFest
Host: Zeshan Chishti

January 2022 — University of California, Riverside
Host: Zhiyun Qian

January 2022 — University of California, Irvine
Host: Sangeetha Abdu Jyothi

November 2021 — Applications Driving Architectures (ADA) Fall Symposium
Session chair: Ada Gavrilovska

November 2021 — University of Michigan Graduate Honors Competition
Host: Emily Mower Provost

November 2021 — University of California, Los Angeles
Host: Jens Palsberg

November 2021 — University of California, San Diego
Host: Dean Tullsen

October 2021 — University of Michigan Systems Lab
Host: Max S. New

October 2021 — University of California, Santa Cruz
Hosts: Lindsey Kuper and Tyler Sorensen

October 2021 — Texas A&M University
Host: Daniel A. Jiménez

September 2021 — University of Rochester
Host: Sreepathi Pai

How Can We Help Students Learn to Review Papers?

June 2023 — Revisiting the Review Processes Workshop (Co-located with ISCA 2023)
Orlando, FL, USA

Whisper: Profile-Guided Branch Misprediction Elimination for Data Center Applications

December 2022 — Apple Computer Architecture Reading Group
Host: Muawya Al-Otoom and Tyler Huberty

November 2022 — AMD Tech Talk
Host: Jagadish Kotra

November 2022 — ARM Austin CPU Group
Host: Dam Sunwoo

October 2022 — IEEE/ACM International Symposium on Microarchitecture (MICRO)
Session chairs: Ulya Karpuzcu and Boris Grot

Thermometer: Profile-Guided BTB Replacement for Data Center Applications

June 2022 — IEEE/ACM International Symposium on Computer Architecture (ISCA)
Session chair: Esha Choukse

Twig: Profile-Guided BTB Prefetching for Data Center Applications

May 2022 Applications Driving Architectures (ADA) Annual Symposium
Session chair: Zachary Tatlock

October 2021 IEEE/ACM International Symposium on Microarchitecture (MICRO)
Session chair: Pedro Trancoso

Locality Optimizations for Data Center Applications

September 2021 Semiconductor Research Corporation (SRC) TECHCON
Session chair: Lisa Green

July 2021 ARM Research
Host: Alex Rico

May 2021 Applications Driving Architectures (ADA) Annual Symposium
Session chair: Timothy G. Rogers

April 2021 EuroSys Doctoral Workshop
Session chair: Irene Zhang

December 2020 Students' Forum Speaker, International Conference on Networking, Systems and Security
Session chair: Mahmuda Naznin

DMon: Efficient Detection and Correction of Data Locality Problems Using Selective Profiling

July 2021 USENIX Symposium on Operating Systems Design and Implementation (OSDI)
Session chairs: Deniz Altinbüken and Rashmi Vinayak

Ripple: Profile-Guided Instruction Cache Replacement for Data Center Applications

June 2021 IEEE/ACM International Symposium on Computer Architecture (ISCA)
Session chair: Esha Choukse

I-SPY: Context-Driven Conditional Instruction Prefetching with Coalescing

October 2020 IEEE/ACM International Symposium on Microarchitecture (MICRO)
Session chair: Trevor E. Carlson

Huron: Hybrid False Sharing Detection and Repair

August 2019 Microsoft C++ Compiler Team
Host: Mo Di

June 2019 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)
Session chair: Veselin Raychev

May 2019 Azure Hardware Research Group
Host: Lisa Hsu

Overcoming Throughput Degradation in Multi-Radio Cognitive Radio Networks

May 2018 Intel Labs Wireless Networking Research Group
Host: Satish C. Jha

Selected Professional Services

2025 Program Committee Member for International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'25)

2024 Program Committee Member for International Symposium on Computer Architecture (ISCA'24)

2024 Program Committee Member for USENIX Annual Technical Conference (ATC'24)

2024 Program Committee Member for European Conference on Computer Systems (EuroSys'24)

2024 Program Committee Member for International Symposium on Code Generation and Optimization (CGO'24)

2023 Program Committee Member for International Symposium on Code Generation and Optimization (CGO'23)

2023 Scholarship Selection Committee Member for ACM Symposium on Operating Systems Principles (SOSP'23)

2022 External Review Committee Member for International Symposium on Microarchitecture (MICRO'22)

2022 Publicity Chair for Young Architect Workshop (YArch'22)

2022 External Review Committee Member for International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'22)

2021 Shadow Program Committee Member for European Conference on Computer Systems (EuroSys'21)

2021 External Review Committee Member for International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'21)

2020 Artifact Evaluation Committee Member for SIGPLAN Conference on Programming Language Design and Implementation (PLDI'20)

2020 Artifact Evaluation Committee Member for Third Conference on Machine Learning and Systems (MLSys'20)

2020 Artifact Evaluation Committee Member for International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'20)

2019 Artifact Evaluation Committee Member for ACM Symposium on Operating Systems Principles (SOSP'19)

2019 Student Volunteer for SIGPLAN Conference on Programming Language Design and Implementation (PLDI'19)

Outreach Activities

2022-present **High School Student Mentoring**

- Woojin Jung (Cranbrook High School)

2018-present **Mentoring Students from Underserved Groups**

- Hiba Altaf (Columbia BSc)
- Zhixuan Zhang (Columbia MSc)
- Sumya Hoque (Columbia BSc)
- Diane Chiang (UM BSc)
- Yiwei Yang (UCSC PhD)
- Pooneh Safayenikoo (UCSC PhD)
- Saba Jamilan (UCSC PhD)
- Surim Oh (UCSC PhD)
- Yuxuan Zhang (UPenn PhD)
- Sara Mahdizadeh Shahri (UM PhD)
- Shixin Song (UM BSc)
- Ashfaqur Rahaman (BSc in Naval Engineering)
- Xiaohe Cheng (HKUST BSc)

2022 **Data Analytics for Detroit Digital Inclusion**

- Eliminating digital inequality in the greater Detroit area
- As a volunteer, advising and mentoring undergraduate and master's student research assistants

2021 **"My CS PhD" Information Session Panelist**

- Attending as a panelist in the info session for undergraduates about getting a PhD in computer science