

Michael Ferdman

Associate Professor, Computer Science, Stony Brook University
Director, Computer Architecture at Stony Brook (COMPAS) Laboratory

Curriculum Vitae - August 2022

<http://compas.cs.stonybrook.edu/~mferdman/>
mferdman@cs.stonybrook.edu
+1 (631) 632-8449

Department of Computer Science
343 New Computer Science
Stony Brook, NY 11794-2424

Research Interests

Computer architecture, with particular emphasis on the design of efficient server systems. Most recently, my main focus has been on Machine Learning Accelerators, developing hardware techniques to enable fast and efficient implementations of deep learning, and making FPGA-based accelerators more practical and easier to program. More broadly, my work seeks to understand the fundamental properties and interactions of application software, operating systems, networks, processor microarchitecture, and datacenter dynamics, to enable software and hardware co-design of high-performance, power-efficient, and compact servers.

Education

Carnegie Mellon University	Pittsburgh, PA
<i>Ph.D. in Electrical and Computer Engineering</i>	June 2012
<i>M.S. in Electrical and Computer Engineering</i>	December 2002
<i>B.S. in Electrical and Computer Engineering</i>	December 2002
<i>B.S. in Computer Science</i>	May 2002

Honors and Awards

2020 Young Academic Inventor's Award from the National Academy of Inventors SBU Chapter
Undergraduate Teaching Award (2019)
David R. Smith Young Scholar in Computer Science Award (2016)
NSF CAREER Award (2015)
Graduate Teaching Award (2014)
Best Paper Award at the 11th International Conference on Virtual Execution Environment (VEE) for "A Comprehensive Implementation and Evaluation of Direct Interrupt Delivery."
IEEE Micro Top Picks from Computer Architecture Conferences of 2013, "A Case for Specialized Processors for Scale-Out Workloads."
Best Paper Award at the 17th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) for "Clearing the Clouds: A Study of Emerging Scale-out Workloads on Modern Hardware."
Best Paper Finalist at the 17th International Symposium on High-Performance Computer Architecture (HPCA) for "Cuckoo Directory: A Scalable Directory for Many-Core Systems."
Paper Award from the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC) for "Cuckoo Directory: A Scalable Directory for Many-Core Systems."
IEEE Micro Top Picks from Computer Architecture Conferences of 2009, "R-NUCA: Data Placement in Distributed Shared Caches."
IEEE Micro Top Picks from Computer Architecture Conferences of 2009, "Practical Off-chip Meta-data for Temporal Memory Streaming."
2005 DARPA Grand Challenge driverless desert race, 2nd and 3d place autonomous vehicles for RedTeam.

Publications

Peer-reviewed Conference Papers

- [2022] *Domains Do Change Their Spots: Quantifying Potential Abuse of Residual Trust*
J. So, N. Miramirkhani, M. Ferdman, N. Nikiforakis, In *2022 IEEE Symposium on Security and Privacy (SP), IEEE Computer Society, volume , 2022. [doi]*
- [2022] *AppBastion: Protection from Untrusted Apps and OSes on ARM*
Darius Suci, Michael Ferdman, Radu Sion, In *Computer Security - ESORICS 2022 - 27th European Symposium on Research in Computer Security, Copenhagen, Denmark, September 26-30, 2022.*
- [2021] *Leveraging FPGA Layout to Minimize Jitter in Statistical Time-to-Digital Converters*
Farid Samandi, Tianchu Ji, Shengsun Cho, Michael Ferdman, Peter Milder, In *29th IEEE International Symposium on*

Field-Programmable Custom Computing Machines (FCCM, withdrawn due to IP dispute), 2021.

- [2021] *On the Distribution, Sparsity, and Inference-time Quantization of Attention Values in Transformers*
Tianchu Ji, Shraddhan Jain, Michael Ferdman, Peter Milder, H. Andrew Schwartz, Niranjan Balasubramanian, In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, Association for Computational Linguistics, 2021.
- [2020] *A Scheduling Approach to Incremental Maintenance of Datalog Programs*
Shikha Singh, Sergey Madaminov, Geoffrey Washburn, Ryan Johnson, Hung Ngo, Dung Nguyen, Soeren Olesen, Kurt Stirewalt, Michael A. Bender, Michael Ferdman, Benjamin Moseley., In *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2020.
- [2020] *Flick: Fast and Lightweight ISA-Crossing Call for Heterogeneous-ISA Environments*
Shenghsun Cho, Han Chen, Sergey Madaminov, Michael Ferdman, Peter Milder, In *47th International Symposium on Computer Architecture (ISCA)*, 2020.
- [2020] *FPGA-Accelerated Samplesort For Large Data Sets*
Han Chen, Sergey Madaminov, Michael Ferdman, Peter Milder, In *2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2020.
- [2019] *x86-64 Instruction Usage Among C/C++ Applications*
Amogh Akshintala, Bhushan Jain, Chia-Che Tsai, Michael Ferdman, Donald E. Porter, In *12th ACM International Conference on Systems and Storage (SYSTOR)*, 2019.
- [2019] *Swarm Model Checking on the GPU*
Richard DeFrancisco, Shenghsun Cho, Michael Ferdman, Scott A. Smolka, In *26th International SPIN Symposium on Model Checking of Software*, Springer International Publishing, 2019. [doi]
- [2019] *An Efficient, Scalable and Exact Representation of High-Dimensional Color Information Enabled via de Bruijn Graph Search*
Fateme Almodaresi, Prashant Pandey, Michael Ferdman, Rob Johnson, Rob Patro, In *22nd Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, 2019.
- [2019] *Runtime-Programmable Pipelines for Model Checkers on FPGAs*
Mrunal Patel, Shenghsun Cho, Michael Ferdman, Peter Milder, In *29th International Conference on Field Programmable Logic and Applications (FPL)*, 2019. (nominated for the Best Paper award)
- [2018] *Panning for gold.com: Understanding the dynamics of domain dropcatching*
Najmeh Miramirkhani, Timothy Barron, Michael Ferdman, Nick Nikiforakis, In *Web Conference (WWW)*, 2018.
- [2018] *Mantis: A Fast, Small, and Exact Large-Scale Sequence Search Index*
Prashant Pandey, Fateme Almodaresi, Michael A. Bender, Michael Ferdman, Rob Johnson, Rob Patro, In *21st Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, 2018.
- [2018] *Taming the Killer Microsecond*
Shenghsun Cho, Amoghavarsha Suresh, Tapti Palit, Michael Ferdman, Nima Honarmand, In *51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2018.
- [2018] *Impact of Device Parameters on Internet-based Mobile Applications*
Malleham Dasari, Santiago Vargas, Arani Bhattacharya, Aruna Balasubramanian, Samir Das, Mike Ferdman, In *2018 Conference on Internet Measurement Conference (IMC)*, 2018.
- [2018] *Medusa: A Scalable Memory Interconnect for Many-Port DNN Accelerators and Wide DRAM Controller Interfaces*
Yongming Shen, Tianchu Ji, Michael Ferdman, Peter Milder, In *28th International Conference on Field Programmable Logic and Applications (FPL)*, 2018.
- [2018] *FPGASwarm: High Throughput Model Checking Using FPGAs*
Shenghsun Cho, Michael Ferdman, Peter Milder, In *28th International Conference on Field Programmable Logic and Applications (FPL)*, 2018.
- [2018] *A Full-System VM-HDL Co-Simulation Framework for Servers with PCIe-Connected FPGAs*
Shenghsun Cho, Mrunal Patel, Han Chen, Michael Ferdman, Peter Milder, In *2018 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2018.
- [2017] *Maximizing CNN Accelerator Efficiency Through Resource Partitioning*
Yongming Shen, Michael Ferdman, Peter Milder, In *44th International Symposium on Computer Architecture (ISCA)*, 2017.
- [2017] *Escher: A CNN Accelerator with Flexible Buffering to Minimize Off-Chip Transfer*
Yongming Shen, Michael Ferdman, Peter Milder, In *25th IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2017.
- [2016] *Fused-Layer CNN Accelerators*
Manoj Alwani, Han Chen, Michael Ferdman, Peter Milder, In *49th Annual IEEE/ACM International Symposium on*

Microarchitecture (MICRO), 2016.

- [2016] *Demystifying Cloud Benchmarking*
Tapti Palit, Yongming Shen, Michael Ferdman, In *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2016.
- [2016] *Overcoming Resource Underutilization in Spatial CNN Accelerators*
Yongming Shen, Michael Ferdman, Peter Milder, In *26th International Conference on Field Programmable Logic and Applications (FPL)*, 2016.
- [2015] *A Comprehensive Implementation and Evaluation of Direct Interrupt Delivery*
Cheng-Chun Tu, Michael Ferdman, Chao-tung Lee, Tzi-cker Chiueh, In *Processors of the 11th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environment (VEE)*, 2015. (recognized as Best Paper by the program committee)
- [2015] *Architectural Support for Dynamic Linking*
Varun Agrawal, Abhiroop Dabral, Tapti Palit, Yongming Shen, Michael Ferdman, In *20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2015.
- [2014] *DIMMER: A case for turning off DIMMs in clouds*
Dongli Zhang, Moussa Ehsan, Michael Ferdman, Radu Sion, In *ACM Symposium on Cloud Computing (SOCC)*, 2014.
- [2014] *Temporal Stream Branch Predictor*
Yongming Shen, Michael Ferdman, In *JWAC-4: Championship Branch Prediction workshop (in conjunction with ISCA'14)*, 2014.
- [2012] *Scale-Out Processors*
Pejman Lotfi-Kamran, Boris Grot, Michael Ferdman, Stavros Volos, Onur Kocberber, Javier Picorel, Almutaz Adileh, Djordje Jevdjic, Sachin Idgunji, Emre Ozer, Babak Falsafi, In *39th International Symposium on Computer Architecture (ISCA)*, 2012.
- [2012] *Clearing the Clouds: A Study of Emerging Scale-out Workloads on Modern Hardware*
Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafae, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, Babak Falsafi, In *17th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2012. (recognized as Best Paper by the program committee and recognized as Top Pick of 2013 by IEEE Micro)
- [2011] *Proactive Instruction Fetch*
Michael Ferdman, Cansu Kaynak, Babak Falsafi, In *44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2011.
- [2011] *Cuckoo Directory: A Scalable Directory for Many-Core Systems*
Michael Ferdman, Pejman Lotfi-Kamran, Ken Balet, Babak Falsafi, In *17th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2011. (selected by the program committee for Best Student Papers session)
- [2010] *TurboTag: lookup filtering to reduce coherence directory power*
Pejman Lotfi-Kamran, Michael Ferdman, Daniel Crisan, Babak Falsafi, In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2010.
- [2009] *Reactive NUCA: near-optimal block placement and replication in distributed caches*
Nikos Hardavellas, Michael Ferdman, Babak Falsafi, Anastasia Ailamaki, In *36th International Symposium on Computer Architecture (ISCA)*, 2009. (recognized as Top Pick of 2009 by IEEE Micro)
- [2009] *Practical Off-Chip Meta-Data for Temporal Memory Streaming*
Thomas F. Wenisch, Michael Ferdman, Anastasia Ailamaki, Babak Falsafi, Andreas Moshovos, In *15th International Symposium on High Performance Computer Architecture (HPCA)*, 2009. (recognized as Top Pick of 2009 by IEEE Micro)
- [2008] *Temporal Instruction Fetch Streaming*
Michael Ferdman, Thomas F. Wenisch, Anastasia Ailamaki, Babak Falsafi, Andreas Moshovos, In *41st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2008.
- [2008] *Cache bursts: A new approach for eliminating dead blocks and increasing cache efficiency*
Haiming Liu, Michael Ferdman, Jaehyuk Huh, Doug Burger, In *41st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2008.
- [2008] *Temporal Streams in Commercial Server Applications*
Thomas F. Wenisch, Michael Ferdman, Anastasia Ailamaki, Babak Falsafi, Andreas Moshovos, In *2008 IEEE International Symposium on Workload Characterization (IISWC)*, 2008.
- [2007] *Last-Touch Correlated Data Streaming*
Michael Ferdman, Babak Falsafi, In *2007 IEEE International Symposium on Performance Analysis of Systems and Software*

(ISPASS), 2007.

- [2007] *Mitigating multi-bit soft errors in L1 caches using last-store prediction*
Brian T Gold, Michael Ferdman, Babak Falsafi, Ken Mai, In *Workshop on Architectural Support for Gigascale Integration (ASGI)*, 2007.
- [2003] *Analysis of IC Manufacturing Process Deformations: An automated approach using SRAM bit fail maps*
Thomas Zanon, Michael Ferdman, Kambiz Komeyli, Wojciech P. Maly, In *29th International Symposium for Testing and Failure Analysis*, 2003.

Journal Articles

- [2022] *An Incrementally Updatable and Scalable System for Large-Scale Sequence Search using the Bentley-Saxe Transformation*
Fateme Almodaresi, Jamshed Khan, Sergey Madaminov, Michael Ferdman, Rob Johnson, Prashant Pandey, Rob Patro, In *Bioinformatics*, 2022.
- [2021] *Practical Model Checking on FPGAs*
Shengsun Cho, Mrunal Patel, Michael Ferdman, Peter Milder, In *ACM Trans. Reconfigurable Technol. Syst.*, Association for Computing Machinery, volume 14, 2021. [doi]
- [2020] *Swarm Model Checking on the GPU*
Richard DeFrancisco, Shengsun Cho, Michael Ferdman, Scott A. Smolka, In *International Journal on Software Tools for Technology Transfer (Thomas Given-Wilson, Axel Legay, eds.)*, 2020.
- [2020] *An Efficient, Scalable, and Exact Representation of High-Dimensional Color Information Enabled Using de Bruijn Graph Search*
Fateme Almodaresi, Prashant Pandey, Michael Ferdman, Rob Johnson, Rob Patro, In *Journal of Computational Biology*, volume 27, 2020. [doi]
- [2019] *Massively Parallel Server Processors*
Varun Agrawal, Mina Abbasi Dinani, Yuxuan Shui, Michael Ferdman, Nima Honarmand, In *Computer Architecture Letters (CAL)*, 2019.
- [2019] *Argus: an End-to-End Framework for Accelerating CNNs on FPGAs*
Yongming Shen, Tianchu Ji, Michael Ferdman, Peter Milder, In *IEEE Micro*, volume , 2019. [doi]
- [2014] *A Case for Specialized Processors for Scale-Out Workloads*
Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafae, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, Babak Falsafi, In *IEEE Micro's Top Picks*, 2014. (original at ASPLOS'12)
- [2012] *Quantifying the Mismatch between Emerging Scale-Out Applications and Modern Processors*
Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafae, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, Babak Falsafi, In *ACM Trans. Comput. Syst.*, ACM, volume 30, 2012.
- [2011] *Toward Dark Silicon in Servers*
Nikos Hardavellas, Michael Ferdman, Babak Falsafi, Anastasia Ailamaki, In *IEEE Micro*, volume 31, 2011.
- [2011] *Spatial Memory Streaming*
Stephen Somogyi, Thomas F. Wenisch, Michael Ferdman, Babak Falsafi, In *Journal of Instruction-Level Parallelism (JILP)*, volume 13, 2011.
- [2010] *Near-Optimal Cache Block Placement with Reactive Nonuniform Cache Architectures*
Nikos Hardavellas, Michael Ferdman, Babak Falsafi, Anastasia Ailamaki, In *IEEE Micro's Top Picks*, volume 30, 2010. (original at ISCA'09)
- [2010] *Making Address-Related Prefetching Practical*
Thomas F. Wenisch, Michael Ferdman, Anastasia Ailamaki, Babak Falsafi, Andreas Moshovos, In *IEEE Micro's Top Picks*, volume 30, 2010. (original at HPCA'09)
- [2006] *SimFlex: Statistical Sampling of Computer System Simulation*
Thomas F. Wenisch, Roland E. Wunderlich, Michael Ferdman, Anastasia Ailamaki, Babak Falsafi, James C. Hoe, In *IEEE Micro*, volume 26, 2006.

Technical Reports

- [2017] *A VM-HDL Co-Simulation Framework for Systems with PCIe-Connected FPGAs*
Shengsun Cho, Mrunal Patel, Basavaraj Kaladagi, Han Chen, Tapti Palit, Michael Ferdman, Peter Milder, Technical report #839, Stony Brook CEAS, 2017.
- [2017] *Mantis: A Fast, Small, and Exact Large-Scale Sequence Search Index*
Prashant Pandey, Fateme Almodaresi, Michael A. Bender, Michael Ferdman, Rob Johnson, Rob Patro, Cold Spring Harbor Laboratory, 2017.

Patents

US 10,726,330 B2: System, Method, and Accelerator to Process Convolutional Neural Network Layers
Michael Ferdman, Peter Milder, Manoj Alwani.

Academic Employment

<i>Microsoft Corporation</i>	<i>Redmond, WA</i>
<i>Visiting Researcher (academic sabbatical)</i>	<i>June 2020-June 2021</i>
<i>Stony Brook University</i>	<i>Stony Brook, NY</i>
<i>Associate Professor</i>	<i>September 2018-present</i>
<i>Assistant Professor</i>	<i>August 2012-August 2018</i>
<i>Ecole Polytechnique Fédérale de Lausanne</i>	<i>Lausanne, Switzerland</i>
<i>Research Assistant</i>	<i>July 2008-June 2012</i>
<i>Carnegie Mellon University</i>	<i>Pittsburgh, PA</i>
<i>Research Assistant</i>	<i>August 2004-June 2012</i>

Teaching

<i>Stony Brook University</i>	<i>Stony Brook, NY</i>
<i>CSE 502 - Graduate Computer Architecture</i>	<i>Spring 2020</i>
<i>CSE 356 - Cloud Computing</i>	<i>Fall 2019</i>
<i>CSE 356 - Cloud Computing</i>	<i>Spring 2019</i>
<i>CSE 502 - Graduate Computer Architecture</i>	<i>Spring 2019</i>
<i>CSE 356 - Cloud Computing</i>	<i>Spring 2018</i>
<i>CSE 506 - Graduate Operating Systems</i>	<i>Fall 2017</i>
<i>CSE 502 - Graduate Computer Architecture</i>	<i>Spring 2017</i>
<i>CSE 356 - Cloud Computing</i>	<i>Spring 2017</i>
<i>CSE 506 - Graduate Operating Systems</i>	<i>Fall 2015</i>
<i>CSE 356(391) - Cloud Computing</i>	<i>Fall 2015</i>
<i>CSE 506 - Graduate Operating Systems</i>	<i>Spring 2015</i>
<i>CSE 602 - Graduate Advanced Computer Architecture</i>	<i>Fall 2014</i>
<i>CSE 502 - Graduate Computer Architecture</i>	<i>Spring 2014</i>
<i>CSE 506 - Graduate Operating Systems</i>	<i>Fall 2013</i>
<i>CSE 502 - Graduate Computer Architecture</i>	<i>Spring 2013</i>
<i>CSE 602 - Graduate Advanced Computer Architecture</i>	<i>Fall 2012</i>
<i>Ecole Polytechnique Fédérale de Lausanne</i>	<i>Lausanne, Switzerland</i>
<i>TA - Advanced Topics on Memory Systems (graduate)</i>	<i>Spring 2009 (Babak Falsafi)</i>
<i>TA - Multiprocessor Architecture (graduate)</i>	<i>Fall 2008 (Babak Falsafi)</i>
<i>Carnegie Mellon University</i>	<i>Pittsburgh, PA</i>
<i>TA - Multiprocessor Architecture (graduate)</i>	<i>Spring 2006 (Babak Falsafi)</i>
<i>TA - Advanced Techniques in Microprocessors (PhD)</i>	<i>Fall 2005 (Babak Falsafi)</i>
<i>TA - Operating Systems (undergraduate)</i>	<i>Fall 2001 (Gregory Kesden)</i>
<i>TA - Embedded Systems (undergraduate)</i>	<i>Fall 2001 (Raj Rajkumar)</i>

Industry Employment

<i>Microsoft Corporation</i>	<i>Redmond, WA</i>
<i>Visiting Researcher</i>	<i>2020-2021</i>
<i>Telinta, Inc.</i>	<i>Springfield, NJ</i>
<i>Chief Technology Officer</i>	<i>2002-present</i>
<i>Cadence Design Systems</i>	<i>Pittsburgh, PA</i>
<i>Software Engineer</i>	<i>April 2004-August 2007</i>
<i>Neoliner, Inc. (startup acquired by Cadence)</i>	<i>Pittsburgh, PA</i>
<i>Software Engineer</i>	<i>March 2003-April 2004</i>
<i>Automatika, Inc.</i>	<i>Pittsburgh, PA</i>
<i>Independent Contractor</i>	<i>September 2002-January 2003</i>

Professional Service

Organizing committees: ASPLOS'22 (co-general chair), ISPASS'21 (general chair), ISPASS'20 (program committee chair), HPCA'20 (registration chair), DPC'19 (program committee chair), ISCA'17 (finance chair), IISWC'17 (travel grant chair), HPCA'17 (workshops & tutorials chair), ISPASS'17 (workshops & tutorials chair), ISPASS'16 (publication chair), ACM SRC at CGO'15 (local organizer), ISPASS'15 (publication chair), MICRO'14 (publication chair), ISPASS'14 (web chair)

Steering Committees: Architectural Support for Programming Languages and Operating Systems (ASPLOS), International Symposium on Performance Analysis of Systems and Software (ISPASS)

Program committees: SIGMETRICS'23, SIGMETRICS'22, MICRO'22, ISCA'21, IEEE MICRO's Top Picks'21, MICRO'20, IPC'20, USENIX ATC'20, IISWC'19, ISCA'19, HPCA'19, MICRO'18, ICCD'18, IISWC'18, ISCA'18, DAC'18, GLSVLSI'18, HPCA'18, MICRO'17, ISCA'17, HPCA'17, CRC'17, ISCA'16, IISWC'16, ISPASS'16, MICRO'15, IISWC'15, ISCA'15, CGO'15, MICRO'14, ICS'14, ICPP'14, HiPEAC'14, ICCD'13, WIVOSCA'13, DATE'13, CCGrid'13, ISPASS'13, IPDPS'13

University Service: Graduate Council Representative for CS (F'21), College of Engineering and Applied Sciences Executive Committee (S'21), CS Executive Committee (S'19-), CS Operations Committee (S'17, F'16, S'16), Graduate Committee (S'17, F'16, S'16), Undergraduate Committee (S'16), Open House Chair (S'17, S'16, S'15), Graduate Admission Committee (S'20, F'19, S'19, F'18, S'18, F'17, S'17, F'16, S'16, F'15, S'15, S'14, F'14, F'13, S'13, F'12), Faculty Recruitment Committee (S'17, S'14, F'14), Department Orientation Organizer (F'16)

NSF invited workshops: Workshop on Sustainable Data Centers '15, XPS Workshop on Exploiting Parallelism and Scalability '15

External reviewer: HPCA'23, MICRO'21, ASPLOS'21, SIGMETRICS'21, ISCA'20, CAL'19, IEEE Micro'19, TECS'19, MICRO'19, CAL'19, TACO'19, CAL'18, TOCS'18, TACO'18, CAL'17, IEEE Micro'17, ACM TACO'17, ACM TOS'16, MICRO'16, ACM TACO'16, HPCA'16, ACM TACO'15, CAL'15, HPCA'15, ASPLOS'15, CF'14, ISCA'14, TC'14, HPCA'14, PPOPP'14, CAL'13, DAC'13, HPCA'13, JCST'13, MICRO'12, IISWC'12, CAL'12, HPCA'12, IISWC'11, MICPRO/DSD'11, ICS'11, ISCA'11, HPCA'11, HiPEAC'10, ISCA'10, HPCA'10, JPDC'09

NSF service: 2022 (panelist x2), 2021 (panelist), 2020 (panelist), 2019 (panelist), 2016 (panelist, reviewer), 2014 (panelist)

Invited Lectures and Talks: Cloud Computing course at HiPEAC ACACES'17, Keynote at RAPIDO'13

PhD committees: Prashant Pandey (Fast and Space-Efficient Maps for Large Data Sets), Weicheng Liu (Low Voltage Clocking Methodologies for Nanoscale ICs), Tan Li (Harness Multicore Parallelism for High Performance Data Replication), Fatima Zarinni (Understanding and Improving Performance in Next-Generation WiFi and Cellular Networks), Mingwei Zhang (Static Binary Instrumentation with Applications to COTS Software Security), Niranjan Hasabnis (Infrastructure for Architecture-independent Binary Analysis and Transformation), Vasily Tarasov (Multi-dimensional Analysis of I/O Workloads for Modern Storage Systems), Zhichao Li (GreenDM: A Versatile Tiering Hybrid Drive for the Trade-Off Evaluation of Performance, Energy, and Endurance), Cheng-Chun (William) Tu (Memory-based Rack-area Network)

MS committees: Bharath Kumar Reddy Vangoor (To FUSE or not to FUSE?), Kavita Agarwal (A Study of Virtualization Overheads), Arun Olappamanna Vasudevan (Finding the right balance - Security vs Performance with Network Storage Systems)

Co-developer of CloudSuite, a benchmark suite for scale-out workloads.

Co-developer of FLEXUS, a scalable, full-system, cycle-accurate multi-processor and multi-core simulation framework between 2005 and 2012.

SIMFLEX and ProtoFlex: Fast, Accurate, and Flexible Simulation of Computer Systems Tutorial at

- 2010 IEEE International Symposium on Workload Characterization (IISWC). Atlanta, GA, December 2010 with Eric Chung, Pejman Lotfi-Kamran, and Michael Papamichael.
- 42st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO). New York, NY, December 2009 with Eric Chung and Michael Papamichael.
- 17th International Conference on Parallel Architectures and Compilation Techniques (PACT), Toronto, Canada, October 2008 with Eric Chung and Nikos Hardavellas.

Organizer of the Fall 2009 weekly seminar of the Systems Labs at Ecole Polytechnique Fédérale de Lausanne.

Organizer of the Fall 2007 weekly seminar of the Computer Architecture Lab at Carnegie Mellon (CALCM).

Senior Member: IEEE, Member: ACM SIGARCH, ACM SIGMICRO, ACM SIGOPS, HiPEAC Associate.

Funding

Stony Brook University-Brookhaven National Lab Seed Grant - Exploration of FPGAs for Real-Time ML-Based Data Compression in sPHENIX

Co-PI, \$37,000, 8/1/2022 - 3/1/2023
National Science Foundation - SHF: Small: Massively Parallel Server Processors
PI, \$600,000, 6/15/2022 - 5/31/2025
OVPR Revise&Resubmit - Distributed Protocol Offload using FPGA SmartNICs
Co-PI, \$13,000, 4/22/2022
AMD Corporation - FPGA Hardware for research
Donation, equipment (\$13,195), 3/9/2022
Ampere Computing - ARM Server Hardware for research
Donation, equipment (\$32,000), 5/28/2021
National Science Foundation - SHF: Small: Sparsity-Aware Hardware Accelerators for Natural Language Processing with Transformers
Co-PI, \$500,000, 10/01/2020 - 9/30/2023
Stony Brook University-Brookhaven National Lab Seed Grant - Software Assist for Hardware-Managed Virtual Memory on FPGA Accelerators
PI, \$37,000, 8/1/2020 - 3/1/2022
Xilinx Corporation - FPGA Hardware for research
Donation, equipment (\$6,495), 6/16/2020
Xilinx Corporation - FPGA Hardware for research
Donation, equipment (\$6,495), 2/20/2020
National Science Foundation - MRI: Acquisition of Heterogeneous Computer System for Machine Learning
Co-PI, \$561,000, 7/15/2019 - 7/14/2022
National Science Foundation / Intel Corporation - FoMR: IPC Improvement through Hardware Memorization
PI, \$400,000, 8/1/2019 - 7/31/2022
Telluric Labs LLC / Department of Energy - IP Access Gateway
PI, \$168,000, 4/8/2019 - 4/7/2020
National Science Foundation - CSR: Medium: Approximate Membership Query Data Structures in Computational Biology and Storage
Co-PI, \$1,200,000, 8/15/2018 - 8/14/2022
Intel Corporation - FPGA Hardware for research
Donation, equipment (\$5,500), 6/22/2018
National Science Foundation - SPX: Harnessing the Power of High-Bandwidth Memory via Provably Efficient Parallel Algorithms
PI, \$750,000 (\$500,000 SBU, \$250,000 WUSTL), 9/15/2017 - 8/14/2021
Xilinx Corporation - FPGA Hardware for research
Donation, equipment (\$7,000), 7/28/2017
Samsung - SSD Hardware for research
Donation, equipment (\$2,000), 7/6/2017
National Science Foundation - Domestic student travel grant funding for IISWC
PI, \$15,000, 6/01/2017 - 12/31/2017
National Science Foundation - EAGER: Measuring the Stability of Web Links
Co-PI, \$89,200, 4/15/2017 - 10/15/2017
National Science Foundation - Research Experiences for Undergraduates: Secure and Efficient Cloud Infrastructure and Accessibility Services
PI, \$21,900, 8/10/2016 - 8/9/2017
National Science Foundation - EAGER: Preliminary Study to Demonstrate Feasibility and Advantages of Massively Parallel Server Processors
Co-PI, \$146,000, 10/1/2016 - 9/30/2017
Oracle Labs - Exploring Custom Graph Algorithms with PGX and Green-Marl
Gift, \$55,000, 8/17/2016
Google - Taming the Killer Microsecond
Gift, \$58,500, 9/2/2016
National Science Foundation - XPS: FPGA Cloud Platform for Deep Learning, Applications in Computer Vision
PI, \$875,000 (\$574,000 SBU, \$301,000 UNC), 9/1/2015 - 8/31/2019
Intel Corporation - Hardware for research
Donation, equipment (\$21,600), 8/6/2015
National Science Foundation - CAREER: Leveraging temporal streams for micro-architectural innovation in data center servers
PI, \$500,000, 2/15/2015 - 1/31/2020
National Science Foundation - EAGER: Preliminary Study to Demonstrate the Performance and Power Advantages of FPGAs for Deep Learning in Computer Vision
PI, \$95,000, 8/1/2014 - 7/31/2016

Altera Corporation - FPGA Hardware for research

Donation, equipment (\$16,000), 10/22/2014

Cavium - Support of research activities

Gift, \$34,400 + equipment, 7/17/2014

National Science Foundation - CRI: Secure and Efficient Cloud Infrastructure and Accessibility Services

PI, \$200,000, 9/1/2014 - 8/31/2017

Semiconductor Research Corporation - Flexible Hardware Acceleration of the Network Stack for Performance and Energy Efficiency

PI, \$300,000, 1/1/2014 - 1/31/2017

Advisees

PhD (4 students)

Sergey Madaminov, 2016-present

Reza Ali Modadi, 2018-present

Farid Samandi, 2019-present

Zavosh Mottahedeh, 2019-present

Graduated PhDs (3 students)

Yongming Shen, 2014-2021

Shenghsun Cho, 2014-2021

Varun Agrawal, 2013-2020

MS Thesis (5 students)

Raveendra Soori, 2018-2020, NUCALloc: Fine-Grained Block Placement in Hashed Last-Level NUCA Caches

Mina Abbasi Dinani, 2016-2020, Massively Parallel Server Processors

Raghav Dogra, 2017-2018, Predictive Batching for Deep Learning Inference Servers

Manoj Alwani, 2015-2016, Fused Convolutional Neural Network Accelerators

Tapti Palit, 2014-2015, Benchmarking Network-Intensive Applications

MS Advanced Project (84 students)

Aakshintala Amogh, Aavuty Rajesh, Abhiroop Dabral, Abhishek Chauhan, Abhishek Patil, Aditi Singh, Ajay Gopal Krishna, Ajay Paddayuru Shreepathi Akshay Kale, Alok Thatikunta, Anirudh Mallick, Ankit Dewan, Anurag Porripreddi, Arjun Mathew Dan, Arnabjyoti Kalita, Balaji Srinivasan, Basavaraj Kaladagi, Bhavya Agarwal, Bhuvnesh Kumar Biswaranjan Panda, Chaitanya Chakka Krishna, Chidambaram Ramanathan, Devashish Thakur, Dhruva Kumar Devineni, Gangabarani Balakrishnan, Harinath Kanchu, Heta Ashit Saraiya, Jacob Samuel Harder, Jerrin Shaji George, Jiemin Zeng, Jihyu Yang, Kai-Chieh Huang, Karan Pugla, Komal Agrawal, Kushal Dhar, Malvika Modi, Mandar Naik, Michael Laurence Anderson, Mingchen Zhang, Mitesh Kumar Singh, Neeraj Dixit, Palit Tapti, Parag Gupta, Parikshit Bhattacharjee, Paul Mathew, Pratik Sudarshan Pantode, Priyanka Nath, Pruthvi Kumar Madugundu, Rahul Gadi, Rajendra Kumar Raghupatruni, Ravikiran Patil, Ravi Prakash Pandey, Rohit Chandramohanan, Romeyo Dsouza, Sahil Parmar, Sai Madan Mohan Reddy Patlolla, Saptarshi Sen, Saraj Munjal, Sayli Yogesh Karnik, Scott Harvey, Shantanu Potdar, Shanuj Shekhar, Shreyas Prabhu Binnamangala, Shubham Pandurang Zope, Shubhanga Narasimha, Shyam Sundar Chandrasekaran, Sneha Pathrose, Sri Krishna Jayadev Peddibhotla, Srinath Battula Yagna Reddy, Subin Mathew, Sumeeth Kyathanahalli, Sunad S Bhandary, Tamilmani Manoharan, Toby Babu, Vaibhav Srivastav Bandikatla, Vamshi Muthineni, Varsha Venkatesh, Vertika Vaid Vinay Krishnamurthy, Vinay Shetty, Vishal Nayak, Yigong Wang, Yu-Jiun Kao, Yuxuan Shui

BS Honors Project (4 students)

David Pokryvailo, Lise Ho, Mehrab Hoque, Yang Sheng Fang

BS Research (12 students)

Benjamin Michalowicz, Brandon Paradiso, Edgar Samudio, Eniola Abdul, Janet Vorobyeva Jessica Chan, Lawrence Lin, Matthew Ng, Minqi (Jason) Shi, Mrunal Patel, Qi Zhang, Roman Scher