

ABHISHEK BHATTACHARJEE

Associate Professor of Computer Science
Yale University
abhishek@cs.yale.edu
<http://www.cs.yale.edu/homes/abhishek>



BRIEF BIOGRAPHY

I am an Associate Professor of Computer Science (with tenure) at Yale. I direct the Systems Architecture Group within Yale's Computer Systems Lab. We work at the boundary between computer architecture and systems software.

My group proposed coalesced TLBs, which have been adopted on AMD's Zen architecture; superpage migration, which we have integrated in Linux from the 4.14 kernel series onwards; the POSIX-compliant GPU system call interface now hosted under the Radeon Open Compute platform for ultrascale computing; and the compilation stack for PsyNeuLink, a brain simulator for cognitive neuroscience.

PROFESSIONAL EXPERIENCE

- 2019-Present **Yale University, New Haven, Connecticut, USA.**
Associate Professor (with tenure), Department of Computer Science.
- 2010-2018 **Rutgers, The State University of New Jersey, New Brunswick, New Jersey, USA.**
2016-2018 Associate Professor (with tenure), Department of Computer Science.
2010-2016 Assistant Professor, Department of Computer Science.
- 2017-2019 **Princeton University, Princeton, New Jersey, USA.**
CV Starr Fellow, Princeton Neuroscience Institute.
- 2013 **Columbia University, New York, New York, USA.**
Visiting Research Scientist (on sabbatical), Department of Computer Science.
- 2007/2008 **Intel Corporation, Hudson, Massachusetts, USA.**
Research Intern, VSSAD Group.
Research Intern, Strategic CAD Labs.

ACADEMIC HISTORY

- 2005-2010 **Ph.D., Electrical Engineering, Princeton University, Princeton, New Jersey, USA.**
2007 **M.A., Electrical Engineering, Princeton University, Princeton, New Jersey, USA.**
Advisor: Margaret R. Martonosi.
Thesis: Thread Criticality and TLB Enhancement Techniques for Chip Multiprocessors.
- 2018 **M.A., Privatum, Yale University, New Haven, Connecticut, USA.**
- 2001-2005 **B.Eng, Honours Electrical Engineering, McGill University, Montréal, Québec, Canada.**

AWARDS AND HONORS

- [1] **The Rutgers-New Brunswick Chancellor's Award for Faculty Excellence in Research:** Awarded annually to one recently tenured faculty member. Conferred in “recognition of outstanding contributions to the entire computer architecture research community, impact on the computer hardware industry, and exploration of a new research direction that holds great promise in improving our understanding of the human brain”, 2017.
- [2] **CV Starr Fellowship in Neuroscience:** Princeton Neuroscience Institute, Princeton University, 2017-2019.
- [3] **IEEE Micro's Top Picks in Computer Architecture:** ASPLOS '17, ASPLOS '16, ASPLOS '14 papers.
- [4] **Honorable mention in IEEE Micro's Top Picks in Computer Architecture:** MICRO '17, ISCA '18 papers.
- [5] **Best paper award nominations:** ASPLOS '17, MICRO '15, PACT '09 papers.
- [6] **Google research award:** Efficient Virtual Memory Support for Emerging Big-Data Workloads, 2013.
- [7] **VMware research award:** Virtualization for Systems with Emerging Storage and Memory Technologies, 2013.
- [8] **National Science Foundation's CAREER award:** NSF's highest award in support of junior faculty who exemplify the role of teacher-scholars through outstanding research, excellent education, and integration of education and research, 2013.
- [9] **Wu Prize for Academic Excellence, Princeton University:** Selected as one of the top PhD students in the School of Engineering and Applied Sciences, 2009.
- [10] **British Association Medal for Great Distinction, McGill University:** Graduated at the top of class in the Department of Electrical, Computer, and Software Engineering, 2005.
- [11] **James McGill Award, McGill University:** Selected for being in the top 1-3% of students in the Honours Electrical Engineering program, 2004.
- [12] **James McGill Entrance Scholarship, McGill University:** Awarded during undergraduate program, 2001-2005.

SCIENTIFIC PUBLICATIONS AND PATENTS

Bold authors: My students and myself.

Conference Articles.

- [1] **Zi Yan**, Daniel Lustig, David Nellans, **Abhishek Bhattacharjee**. “Nimble Page Management for Tiered Memory Systems”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXIV)*, April 2019.
 - We have integrated our transparent hugepage migration techniques into the Linux kernel.
- [2] Srikant Bharadwaj•, **Guilherme Cox**•, Tushar Krishna, **Abhishek Bhattacharjee**. “Scalable Distributed Last-Level TLBs Using Low-Latency Interconnects”, *International Symposium on Microarchitecture (MICRO-51)*, October 2018.
 - Joint first authors.
- [3] **Ján Veselý**, Arkaprava Basu, **Abhishek Bhattacharjee**, Gabriel H. Loh, Mark H. Oskin, Steven K. Reinhardt. “Generic System Calls for GPUs”, *International Symposium on Computer Architecture (ISCA-45)*, June 2018.
 - Selected for honorable mention in IEEE Micro's Top Picks in Computer Architecture journal.
 - Open-sourced under the Radeon Open Compute project.

- [4] Mayank Parasar, **Abhishek Bhattacharjee**, Tushar Krishna. “SEESAW: Using Superpages to Improve VIPT Caches”, *International Symposium on Computer Architecture (ISCA-45)*, June 2018.
- [5] Seunghee Shin, **Guilherme Cox**, Mark H. Oskin, Gabriel H. Loh, Yan Solihin, **Abhishek Bhattacharjee**, Arkaprava Basu. “Scheduling Page Table Walks for Irregular GPU Applications”, *International Symposium on Computer Architecture (ISCA-45)*, June 2018.
- [6] **Guilherme Cox**, **Zi Yan**, **Abhishek Bhattacharjee**, Vinod Ganapathy. “Secure, Consistent, and High-Performance Memory Snapshotting”, *Conference on Data and Application Security and Privacy (CODASPY-8)*, March 2018.
- [7] Mohan Kumar, Steffen Maass, Sanidhya Kashyap, **Ján Veselý**, **Zi Yan**, Taesoo Kim, **Abhishek Bhattacharjee**, Tushar Krishna. “LATR: Lazy Translation Coherence”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXIII)*, March 2018.
- [8] **Abhishek Bhattacharjee**. “Using Branch Predictors to Predict Brain Activity in Brain-Machine Implants”, *International Symposium on Microarchitecture (MICRO-50)*, October 2017.
- Selected for honorable mention in IEEE Micro’s Top Picks in Computer Architecture journal.
- [9] **Zi Yan**, **Ján Veselý**, **Guilherme Cox**, **Abhishek Bhattacharjee**. “Hardware Translation Coherence for Virtualized Systems”, *International Symposium on Computer Architecture (ISCA-44)*, June 2017.
- [10] **Abhishek Bhattacharjee**. “Translation-Triggered Prefetching”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXII)*, April 2017.
- Nominated for the Best Paper Award.
 - Selected for inclusion in IEEE Micro’s Top Picks in Computer Architecture journal.
- [11] **Guilherme Cox**, **Abhishek Bhattacharjee**. “Efficient Address Translation for Architectures with Multiple Page Sizes”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXII)*, April 2017.
- [12] **Ján Veselý**, Arkaprava Basu, Mark H. Oskin, Gabriel H. Loh, **Abhishek Bhattacharjee**. “Observations and Opportunities in Architecting Shared Virtual Memory for Heterogeneous Systems”, *International Symposium on Performance Analysis of Systems and Software (ISPASS)*, September 2016.
- [13] Daniel Lustig[•], **Geet Sethi**[•], Margaret R. Martonosi, **Abhishek Bhattacharjee**. “COATCheck: Verifying Memory Ordering at the Hardware-OS Interface”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXI)*, March 2016.
- Joint first authors.
 - Selected for inclusion in IEEE Micro’s Top Picks in Computer Architecture journal.
- [14] **Binh Q. Pham**, **Ján Veselý**, Gabriel H. Loh, **Abhishek Bhattacharjee**. “Large Pages and Lightweight Memory Management in Virtualized Environments: Can You Have it Both Ways?”, *International Symposium on Microarchitecture (MICRO-48)*, December 2015.
- Nominated for the Best Paper Award.
- [15] **Bharath Pichai**, Lisa Hsu, **Abhishek Bhattacharjee**. “Architectural Support for Address Translation on GPUs”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XIX)*, March 2014.
- Selected for inclusion in IEEE Micro’s Top Picks in Computer Architecture journal.

[16] **Binh Q. Pham, Abhishek Bhattacharjee**, Yasuko Eckert, Gabriel H. Loh. “Increasing TLB Reach by Exploiting Clustering in Page Translations”, *International Symposium on High Performance Computer Architecture (HPCA-20)*, February 2014.

[17] **Abhishek Bhattacharjee**. “Large-Reach Memory Management Unit Caches”, *International Symposium on Microarchitecture (MICRO-46)*, December 2013.

[18] Cheng Li, Íñigo Goiri, **Abhishek Bhattacharjee**, Ricardo Bianchini, Thu D. Nguyen. “Quantifying and Improving I/O Predictability in Virtualized Systems”, *International Symposium on Quality of Service (IWQoS)*, June 2013.

[19] **Binh Q. Pham, Viswanathan Vaidyanathan**, Aamer Jaleel, **Abhishek Bhattacharjee**. “CoLT: Coalesced Large-Reach TLBs”, *International Symposium on Microarchitecture (MICRO-45)*, December 2012.

- AMD implements coalesced TLBs on its chips today, beginning with its Ryzen architecture launched in 2017.

[20] Qingyuan Deng, David Meisner, **Abhishek Bhattacharjee**, Thomas F. Wenisch, Ricardo Bianchini. “CoScale: Coordinating CPU and Memory System DVFS in Server Systems”, *International Symposium on Microarchitecture (MICRO-45)*, December 2012.

[21] Qingyuan Deng, David Meisner, **Abhishek Bhattacharjee**, Thomas F. Wenisch, Ricardo Bianchini. “MultiScale: Memory System DVFS with Multiple Memory Controllers”, *International Symposium on Low Power Electronics and Design (ISLPED-17)*, July 2012.

[22] **Abhishek Bhattacharjee**, Daniel Lustig, Margaret R. Martonosi. “Shared Last-Level TLBs for Chip Multiprocessors”, *International Symposium on High Performance Computer Architecture (HPCA-17)*, February 2011.

[23] **Abhishek Bhattacharjee**, Margaret R. Martonosi. “Inter-Core Cooperative TLB Prefetchers for Chip Multiprocessors”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XV)*, March 2010.

[24] **Abhishek Bhattacharjee**, Margaret R. Martonosi. “Characterizing the TLB Behavior of Emerging Parallel Workloads on Chip Multiprocessors”, *International Conference on Parallel Architectures and Compilation Techniques (PACT-18)*, September 2009.

- Nominated for the Best Paper Award.

[25] **Abhishek Bhattacharjee**, Margaret R. Martonosi. “Thread Criticality Predictors for Dynamic Performance, Power, and Resource Management in Chip Multiprocessors”, *International Symposium on Computer Architecture (ISCA-36)*, June 2009.

[26] **Abhishek Bhattacharjee**, Gilberto Contreras, Margaret R. Martonosi. “Full-System Chip Multiprocessor Power Evaluations Using FPGA-Based Emulation”, *International Symposium on Low Power Electronics and Design (ISLPED-13)*, August 2008.

Journal Articles.

[1] **Zi Yan, Ján Veselý, Guilherme Cox, Abhishek Bhattacharjee**. “Hardware Translation Coherence for Virtualized Systems”, *SIGOPS Operating Systems Review*, 2018.

[2] **Abhishek Bhattacharjee**, “Breaking the Address Translation Wall by Accelerating Memory Replays”, *IEEE Micro’s Top Picks in Computer Architecture journal (Top Picks)*, May 2018.

[3] Daniel Lustig, **Geet Sethi, Abhishek Bhattacharjee**, Margaret Martonosi. “Transistency Models: Memory Ordering at the Hardware-OS Interface”, *IEEE Micro’s Top Picks in Computer Architecture journal (Top Picks)*, May 2017.

[4] **Bharath Pichai**, Lisa Hsu, **Abhishek Bhattacharjee**. “Address Translation for Throughput Oriented Accelerators”, *IEEE Micro’s Top Picks in Computer Architecture journal (Top Picks)*, May 2015.

[5] Daniel Lustig, **Abhishek Bhattacharjee**, Margaret R. Martonosi. “TLB Improvements for Chip Multiprocessors: Inter-Core Cooperative Prefetchers and Shared Last-Level TLBs”, *ACM Transactions on Architecture and Code Optimization (TACO)*, Vol. 10, No. 1, Article 5, April 2013.

[6] **Abhishek Bhattacharjee**, Gilberto Contreras, Margaret R. Martonosi. “Parallelization Libraries: Characterizing and Reducing Overheads”, *ACM Transactions on Architecture and Code Optimization (TACO)*, Vol. 8, No. 1, Article 5, April 2011.

Textbooks.

[1] **Abhishek Bhattacharjee**. “Appendix L: Advanced Concepts on Address Translation”, Appendix contributed to Hennessy and Patterson, *Computer Architecture: A Quantitative Approach* (6th edition), 2018.

[2] **Abhishek Bhattacharjee**, Daniel Lustig. “Architectural and Operating System Support for Virtual Memory”, *Synthesis Lectures on Computer Architecture*, Morgan Claypool Publishers, 2017.

Workshops and Others.

[1] Reto Achermann, Jayneel Gandhi, **Abhishek Bhattacharjee**, Timothy Roscoe. “Poster: Mitosis: Transparently Self-Replicating Page Tables for Big Memory Machines”, *Operating System Design and Implementation (OSDI)*, October 2018.

[2] **Binh Q. Pham**, Derek Hower, **Abhishek Bhattacharjee**, Trey Cain. “TLB Shutdown Mitigation for Low-Power Many-Core Servers with L1 Virtual Caches”, *IEEE Computer Architecture Letters*, April 2017.

[3] Liu Liu, Sibren Isaacman, **Abhishek Bhattacharjee**, Ulrich Kremer. “Exploiting Approximations for Energy/Quality Tradeoffs in Service-Based Applications”, *International Conference on Parallel Architectures and Compilation Techniques (PACT)*, September 2017.

Patents.

[1] **Abhishek Bhattacharjee**, Margaret R. Martonosi. “Inter-Core Cooperative TLB Prefetchers”, patent number 9524232, Dec. 20, 2016.

TALKS

[1] “Architectural Techniques to Build Energy-Efficient Brain Implants”, *ARM Research Summit*, Cambridge, UK, September 2018.

[2] “Architectural Support for Virtual Memory”, *International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems*, Fiuggi, Italy, July 2018.

[3] “Efficient Heterogeneous Systems Across Computing Scales”, *University of California, Santa Barbara*, April 2018.

[4] “Computer Systems for the Brain Sciences”, *University of Wisconsin-Madison*, March 2018.

[5] “Efficient Heterogeneous Systems Across Computing Scales”, *Yale University*, February 2018.

[6] “Computer Systems for the Brain Sciences”, *Columbia University*, February 2018.

[7] “Attacking the Address Translation Bottleneck”, *VMware Research*, January 2018.

[8] “Computer Systems for the Brain Sciences”, *University of Wisconsin-Madison*, December 2017.

[9] “Computer Systems for the Brain Sciences”, *University of California, Santa Barbara*, December 2017.

[10] “Computer Systems for the Brain Sciences”, *University of California, Los Angeles*, December 2017.

[11] “Computer Systems for the Brain Sciences”, *Rice University*, November 2017.

[12] “Computer Systems for the Brain Sciences”, *University of Texas, Austin*, November 2017.

- [13] “Computer Systems for the Brain Sciences”, *Yale University*, November 2017.
- [14] “Computer Systems for the Brain Sciences”, *Dartmouth College*, October 2017.
- [15] “Computer Systems for the Brain Sciences”, *Rutgers University*, October 2017.
- [16] “Using Branch Predictors to Predict Brain Activity in Brain-Machine Implants”, *International Symposium on Microarchitecture (MICRO-50)*, October 2017.
- [17] “Computer Systems for Neuroscience”, *Brown University*, September 2017.
- [18] “Translation-Triggered Prefetching”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXII)*, April 2017.
- [19] “Efficient Virtual Memory for Next-Generation Systems”, *Dartmouth College*, November 2016.
- [20] “Efficient Virtual Memory for Next-Generation Systems”, *Northwestern University*, September 2016.
- [21] “Sweet Spots and Limits for Virtualization”, *Virtual Execution and Environments Invited Panelist*, March 2016.
- [22] “Efficient Virtual Memory for Next-Generation Systems”, *University of Pennsylvania*, March 2016.
- [23] “Taming Heterogeneous Big-Data Systems with Efficient Virtual Memory”, *Georgia Institute of Technology*, November 2015.
- [24] “Virtual Memory in Next-Generation Heterogeneous Manycore Systems”, *University of Pennsylvania*, March 2015.
- [25] “Virtual Memory in Next-Generation Heterogeneous Manycore Systems”, *Carnegie Mellon University*, March 2015.
- [26] “Virtual Memory in Next-Generation Heterogeneous Manycore Systems”, *Indian Institute of Technology, Delhi*, February 2015.
- [27] “Virtual Memory in Next-Generation Heterogeneous Manycore Systems”, *Intel Bangalore*, February 2015.
- [28] “Virtual Memory in Next-Generation Heterogeneous Manycore Systems”, *University of Texas, Austin*, December 2014.
- [29] “Challenges in Address Translation for Next-Generation Heterogeneous Manycore Systems”, *Microsoft Research*, October 2014.
- [30] “Challenges in Address Translation for Next-Generation Heterogeneous Manycore Systems”, *University of Washington*, October 2014.
- [31] “Challenges in Address Translation for Next-Generation Heterogeneous Manycore Systems”, *University of Wisconsin-Madison*, October 2014.
- [32] “Challenges in Address Translation for Next-Generation Heterogeneous Manycore Systems”, *Ecole Polytechnique Federale de Lausanne*, October 2014.
- [33] “Architectural Support for Virtual Memory on Next-Generation Heterogeneous Systems with Memory-Intensive Workloads”, *Brown University*, April 2014.
- [34] “Architectural Support for Address Translation on GPUs”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XIX)*, March 2014.
- [35] “Large-Reach Memory Management Unit Caches”, *International Symposium on Microarchitecture (MICRO-46)*, December 2013.
- [36] “High Performance Virtual Memory for Manycore Heterogeneous Systems”, *AMD Research*, May 2013.
- [37] “High Performance Virtual Memory for Manycore Heterogeneous Systems”, *Harvard University*, May 2013.
- [38] “High Performance Virtual Memory for Manycore Heterogeneous Systems”, *IBM TJ Watson Research Center*, March 2013.

- [39] “Address Translation for Emerging Applications and Architectures”, *Columbia University*, November 2012.
- [40] “Address Translation for Emerging Applications and Architectures“, *VSSAD, Intel Corporation*, September 2012.
- [41] “Shared Last-Level TLBs for Chip Multiprocessors”, *International Symposium on High Performance Computer Architecture (HPCA-17)*, February 2011.
- [42] “Inter-Core Cooperative TLB Prefetchers for Chip Multiprocessors”, *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XV)*, March 2010.
- [43] “Thread Criticality and TLB Enhancement Techniques for Chip Multiprocessors”, *Rutgers*, February 2010.
- [44] “Characterizing the TLB Behavior of Emerging Parallel Workloads on Chip Multiprocessors”, *International Conference on Parallel Architectures and Compilation Techniques (PACT-18)*, September 2009.
- [45] “Thread Criticality Predictors for Dynamic Performance, Power, and Resource Management in Chip Multiprocessors”, *International Symposium on Computer Architecture (ISCA-36)*, June 2009.
- [46] “Full-System Chip Multiprocessor Power Evaluations Using FPGA-Based Emulation”, *International Symposium on Low Power Electronics and Design (ISLPED-13)*, 2008.

RESEARCH FUNDING SOURCES

- [1] **National Science Foundation, CCF-SHF:** “Architectural Techniques for Energy-Efficient Brain-Machine Implants”, PI: Abhishek Bhattacharjee; 2018; **\$465,995.**
- [2] **VMware Research, Research Gift:** “Virtual Memory Abstractions for Heterogeneous Server-Scale Systems in Virtualized Environments”, PI: Abhishek Bhattacharjee; March 2018; **\$10,000.**
- [3] **Intel Hardware Accelerator Research Program:** “Accelerating Brain Modeling Frameworks with FPGAs”, PI: Abhishek Bhattacharjee; October 2017.
- [4] **National Science Foundation II-EN: Collaborative Research:** “Enhancing the Parasol Experimental Testbed for Sustainable Computing”, Rutgers PI: Thu D. Nguyen, Rutgers Co-PIs: Abhishek Bhattacharjee, Ulrich Kremer, Manish Parashar, Ivan Rodero; 2017; Rutgers share: **\$729,725.**
- [5] **Google Research Award:** “Efficient Virtual Memory Support for Emerging Big-Data Workloads”, PI: Abhishek Bhattacharjee; 2013; **\$42,000.**
- [6] **National Science Foundation, XPS:** “Enhancing the Programmability of Heterogeneous Manycore Systems”, PI: Abhishek Bhattacharjee, Co-PI: Ricardo Bianchini; 2013; **\$749,995.**
- [7] **National Science Foundation, CCF-SHF:** “Taming the Combinatorial Explosion of Power Management for Future Manycore Systems”, PI: Ricardo Bianchini, Co-PI: Abhishek Bhattacharjee; 2013; **\$450,000.**
- [8] **VMware Labs, Research Gift:** “Virtualization for Systems with Emerging Storage and Memory Technologies”, PI: Abhishek Bhattacharjee; 2013; **\$118,000.**
- [9] **National Science Foundation CAREER-SHF:** “Cross-Core Learning in Future Manycore Systems”, PI: Abhishek Bhattacharjee, 2013; **\$520,000.**
- [10] **National Science Foundation CCF-SHF:** “Heterogeneous Memory Architectures for Future Many-core Systems”, PI: Abhishek Bhattacharjee, Co-PI: Ricardo Bianchini; 2012; **\$150,000.**

[11] **Rutgers University Faculty Research Grant:** “Redesigning Architectural Support for Virtual Memory”, PI: Abhishek Bhattacharjee; 2012; **\$24,000.**

T E A C H I N G

Associate Professor, Department of Computer Science, Yale University, 2018-Present.

[1] Spring 2019: Topics on the Hardware/Software Interface (635); Course rating TBD, Instructor rating TBD, TBD graduate and undergraduate students.

Invited Instructor, International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems (ACACES), HIPEAC, 2018.

[1] Summer 2018: Architectural Support for Virtual Memory; Course rating 4.39/5.0, Average rating for all courses offered at the summer school 4.13/5.0, 69 graduate and undergraduate students as well as industry practitioners.

Associate Professor, Department of Computer Science, Rutgers University, 2016-2018.

[1] Fall 2018: Operating System Design (416); Course rating 4.58/5.0, Instructor rating 4.71/5.0, 109 undergraduates.

[2] Spring 2018: Computer Architecture (211); Course rating 4.50/5.0, Instructor rating 4.56/5.0, 214 undergraduates.

[3] Fall 2017: Computer Structures (505); Course rating 5.0/5.0, Instructor rating 4.75/5.0, 4 graduate students.

[4] Fall 2016: Computer Architecture (211); Course rating 4.44/5.0, Instructor rating 4.50/5.0, 229 undergraduates.

Assistant Professor, Department of Computer Science, Rutgers University, 2010-2016.

[1] Spring 2016: Computer Structures (505); Course rating 4.80/5.0, Instructor rating 4.80/5.0, 14 graduate students.

[2] Fall 2015: Computer Architecture (211); Course rating 4.50/5.0, Instructor rating 4.69/5.0, 101 undergraduates.

[3] Spring 2015: Computer Structures (505); Course rating 4.50/5.0, Instructor rating 4.60/5.0, 13 graduate students.

[4] Fall 2014: Computer Architecture (211); Course rating 3.45/5.0, Instructor rating 3.74/5.0, 155 undergraduates.

[5] Spring 2014: Computer Architecture (211); Course rating 4.30/5.0, Instructor rating 4.00/5.0, 158 undergraduates.

[6] Spring 2013: Adv. Computer Arch. (507); Course rating 4.88/5.0, Instructor rating 4.88/5.0, 8 graduate students.

[7] Fall 2012: Computer Structures (505); Course rating 4.53/5.0, Instructor rating 4.67/5.0, 15 graduate students.

[8] Spring 2012: Computer Structures (505); Course rating 5.0/5.0, Instructor rating 5.0/5.0, 12 graduate students.

[9] Fall 2011: Computer Architecture (211); Course rating 4.22/5.0, Instructor rating 4.10/5.0, 100 undergraduates.

[10] Spring 2011: Computer Architecture (211); Course rating 4.06/5.0, Instructor rating 4.25/5.0, 70 undergraduates.

[11] Fall 2010: Computer Structures (505); Course rating 4.69/5.0, Instructor rating 4.69/5.0, 32 graduate students.

ADVISING

Current Students.

[1] **Ján Veselý**, post-qualifiers PhD candidate, Computer Science, Rutgers University.

[2] **Karthik Sriram**, pre-qualifiers PhD student, Computer Science, Yale University.

Alumni.

PhD students

[1] **Zi Yan, PhD, Computer Science, Rutgers University, 2019.**

Thesis: Virtual Memory for Next-Generation Tiered Memory Architectures.

First employment: Software Engineer, NVIDIA Research.

[2] **Guilherme Cox, PhD, Computer Science, Rutgers University, 2018.**

Thesis: Improving and Complementing Virtual Memory Using Hardware Techniques.

First employment: Software Engineer, NVIDIA.

[3] **Binh Q. Pham, PhD, Computer Science, Rutgers University, 2015.**

Thesis: Architectural Support for Virtual Memory on Big-Memory Systems.

Awards: Best paper nomination at MICRO '15, Rizvi Family Prize.

First employment: Research Scientist, Intel Labs.

MS students

[1] **Ronil Mehta, MS, Computer Science, Rutgers University, 2016.**

Research: GPU Virtual Memory Performance Profiling.

First employment: Software Engineer, Bloomberg.

[2] **Bharath Pichai, MS, Computer Science, Rutgers University, 2013.**

Thesis: Architectural Support for Address Translation on GPUs.

First employment: Software Development Engineer, Amazon.

Awards: ASPLOS '14 paper chosen for Top Picks '15.

[3] **Viswanathan Vaidyanathan, MS, Computer Science, Rutgers University, 2012.**

Thesis: Infrastructure Development and Characterization of TLBs and Superpage Behavior on Modern Processors.

First employment: Systems Engineer, Riverbed Technologies.

Undergraduate students

[1] **Geet Sethi, BS, Computer Science & Math, Rutgers University, 2016.**

Research: Memory Transistency Models.

Next institution: Stanford PhD program.

Awards: CRA Undergraduate Research Award honorable mention, first-author publication as an undergraduate at ASPLOS '16 (also selected for inclusion in IEEE Micro's Top Picks in Computer Architecture journal).

[2] **Jonathan Risinger, BS, Computer Science, Rutgers University, 2016.**

Next institution: Rutgers MS program.

[3] **Timothy Yong, BS, Electrical and Computer Engineering, Rutgers University, 2016.**

Next institution: Rutgers MS program.

High-school students

[1] **David Liao, Governor's School of Science and Technology in NJ, 2013.**

First institution: BSE, University of Pennsylvania.

[2] **Ryan Morey, Governor's School of Science and Technology in NJ, 2013.**

First institution: BS, Rutgers University.

[3] **Alex Rucker, Governor's School of Science and Technology in NJ, 2013.**

First institution: BS, Cornell University.

Other PhD Committees.

[1] Ari Weinstein, Rutgers.

[2] Qingyuan Deng, Rutgers.

[3] Monica Babes-Vroman, Rutgers.

[4] Cheng Li, Rutgers.

[5] Md. Haque, Rutgers.

[6] Rajat Shuvro Roy, Rutgers.

[7] Luiz Ramos, Rutgers.

[8] Javier Picorel, EPFL.

[9] Ioannis Manousakis, Rutgers.

[10] Myrto Papadopoulou, Toronto.

[11] Hai Nguyen, Rutgers.

[12] Daeyoung Kim, Rutgers.

PROFESSIONAL SERVICE

Conference Program Committees.

Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2019.

Int'l Symposium on Computer Architecture (ISCA) 2018.

Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) 2018.

Int'l Symposium on Microarchitecture (MICRO) 2017.

Int'l Symposium on Computer Architecture (ISCA) 2017.

Int'l Symposium on High Performance Computer Architecture (HPCA) 2017.

Int'l Symposium on Workload Characterization (IISWC) 2016.

Int'l Symposium on Computer Architecture (ISCA) 2016.

Int'l Conference on Parallel Architectures and Compilation Techniques (PACT) 2016.

Int'l Symposium on High Performance Computer Architecture (HPCA) 2015.

Int'l Symposium on Microarchitecture (MICRO) 2014.

Int'l Symposium on Parallel and Distributed Systems (IPDPS) 2014.

Int'l Conference on High Performance and Embedded Architectures and Compilers (HiPEAC) 2014.

Int'l Conference on High Performance and Embedded Architectures and Compilers (HiPEAC) 2013.

Int'l Symposium on Performance Analysis of Systems and Software (ISPASS) 2012.

Int'l Symposium on High Performance Computer Architecture (HPCA) 2012.

Other Activities.

Lightning Session Chair, Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2019.

Industry Sponsors Chair, Int'l Symposium on Performance Analysis of Systems and Software (ISPASS), 2019.

Rutgers School of Arts and Sciences Honors Program Faculty Mentor, 2018-2019.

Students awards chair, Int'l Symposium on Computer Architecture (ISCA) 2018.

Industry liaison co-chair, Int'l Symposium on Computer Architecture (ISCA) 2013.