

Michael Ferdman

Assistant Professor, Computer Science, Stony Brook University

Co-Director, Computer Architecture at Stony Brook (COMPAS) Laboratory

Curriculum Vitae - April 2018

<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

Ph.D. in Electrical and Computer Engineering

June 2012

M.S. in Electrical and Computer Engineering

December 2002

B.S. in Electrical and Computer Engineering

December 2002

B.S. in Computer Science

May 2002

Honors and Awards

David R. Smith Young Scholar in Computer Science Award (2016-2020)

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 (Total: 35, Conference: 23, Journal: 7; Google Scholar: 2449 citations; ISI Web of Science: 505 citations)

Peer-reviewed Conference Papers

- [2018] **Panning for gold.com: Understanding the dynamics of domain dropcatching**
Najmeh Miramirkhani, Timothy Barron, Michael Ferdman, Nick Nikiforakis, *In Proceedings of the Web Conference (WWW)*, 2018.
- [2018] **Mantis: A Fast, Small, and Exact Large-Scale Sequence Search Index**
Prashant Pandey, Fatemeh Almodaresi, Michael A. Bender, Michael Ferdman, Rob Johnson, Rob Patro , *In 21st Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, 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 Proceedings of the 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 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 Proceedings of the 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

- [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**
Shenghsun 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, Fatemeh Almodaresi, Michael A. Bender, Michael Ferdman, Rob Johnson, Rob Patro , Cold Spring Harbor Laboratory, 2017.

Patents

System and Method for Fused Computation of Convolutional Neural Network Layers.

Manoj Alwani, Michael Ferdman, and Peter Milder. Filed October 11, 2016. (pending)

Teaching

Stony Brook University	Stony Brook, NY
<i>CSE 506 - Graduate Operating Systems</i>	Fall 2017
<i>CSE 502 - Graduate Computer Architecture</i>	Spring 2017
<i>CSE 356 - Cloud Computing</i>	Spring 2017
<i>CSE 506 - Graduate Operating Systems</i>	Fall 2015

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

Industry Experience

Telinta, Inc.	Springfield, NJ
<i>Chief Technology Officer</i>	2002-
Cadence Design Systems	Pittsburgh, PA
<i>Software Engineer</i>	April 2004-August 2007
Neoliner, Inc. (startup acquired by Cadence)	Pittsburgh, PA
<i>Software Engineer</i>	March 2003-April 2004
Automatika, Inc.	Pittsburgh, PA
<i>Independent Contractor</i>	September 2002-January 2003
National Robotics Engineering Consortium	Pittsburgh, PA
<i>Circuit Designer and Software Engineer</i>	February 2001-May 2002

Professional Service

Organizing committees: *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)*

Program committees: *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*

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

External reviewer: *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: 2016 (panelist, reviewer), 2014 (panelist)

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

PhD committees: *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)*

Department Service: *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'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)*

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).

Member, IEEE Computer Society, ACM SIGARCH, ACM SIGMICRO, ACM SIGOPS, HiPEAC Associate.

Funding (Total: \$2.4M grants, \$150K gifts, \$46K equipment)

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 (5 students)

Varun Agrawal, 2013-present
Shenghsun Cho, 2014-present
Yongming Shen, 2014-present
Mina Abbasi Dinani, 2016-present
Sergey Madaminov, 2016-present

MS Thesis (2 students)

Tapti Palit, 2014-2015, Benchmarking Network-Intensive Applications
Manoj Alwani, 2015-2016, Fused Convolutional Neural Network Accelerators

MS Advanced Project (50 students)

Aakshintala Amogh, Aavuty Rajesh, Abhiroop Dabral, Abhishek Chauhan, Aditi Singh, Akshay Kale, Arnabjyoti Kalita, Balaji Srinivasan, Basavaraj Kaladagi, Bhavya Agarwal, Biswaranjan Panda, Chaitanya Chakka Krishna, Chidambaram Ramanathan, Devashish Thakur, Dhruva Kumar Devineni, Gangabarani Balakrishnan, Jerrin Shaji George, Jihyu Yang, Karan Pugla, Mandar Naik, Mingchen Zhang, Neeraj Dixit, Palit Tapti, Parag Gupta,

Parikshit Bhattacharjee, Paul Mathew, Raghav Dogra, Rahul Gadi, Rajendra Kumar Raghupatruni, Ravi Prakash Pandey, Ravikiran Patil, Rohit Chandramohanan, Romeyo Dsouza, Sahil Parmar, Saptarshi Sen, Scott Harvey, Shantanu Potdar, Shreyas Prabhu Binnamangala, Shyam Sundar Chandrasekaran, Srinath Battula Yagna Reddy, Subin Mathew, Sumeeth Kyathanahalli, Sunad S Bhandary, Tamilmani Manoharan, Varsha Venkatesh, Vinay Krishnamurthy, Vinay Shetty, Vishal Nayak, Yigong Wang, Yuxuan Shui

BS Honors Project (1 student)

Mehrab Hoque

BS Research (4 students)

Edgar Samudio, Eniola Abdul, Mrunal Patel, Roman Scher