Jayneel Gandhi

Phone: +1-919-607-9438 | Email: gandhij@vmware.com | Webpage: http://pages.cs.wisc.edu/~jayneel

Short Bio

I am a Sr. Researcher at VMware Research, and my research interests are in specialized hardware accelerators, emerging memory/storage technologies, hardware-software co-design, operating systems, virtualization, and distributed systems. My research has focused on virtualizing and easing the adoption of heterogeneous accelerators and emerging memory technologies by introducing useful operating system abstractions and efficient memory system architectures. Consequently, I have 19 publications in top-tier conferences and 9 patents issued in the US. Additionally, three of my research publications have been selected as "Top Picks in Computer Architecture". Subsequently, my research has significantly influenced OS and hypervisor development at VMware. Moreover, my research led to many open-source projects which are widely used by academia and industry.

Education

Ph.D., Computer Sciences, University of Wisconsin – Madison — Aug 2010 – Aug 2016
Advisors: Mark D. Hill and Michael M. Swift
Dissertation Title: Efficient Memory Virtualization

M.S., Computer Sciences, University of Wisconsin – Madison — Aug 2010 – May 2013
Advisors: Prof. Mark D. Hill and Prof. Michael M. Swift

M.S., Computer Engineering, North Carolina State University — Aug 2008 – Jun 2010
Advisor: Prof. Eric Rotenberg
Thesis Title: FabFetch: A Synthesizable RTL Model of a Pipelined Instruction Fetch Unit for Superscalar Processors

B.Tech., Information and Communication Technology, DA-IICT — Aug 2004 – May 2008
Advisor: Prof. Chetan Parikh

Professional Experience

Sr. Researcher, VMware Research — Jun 2021 – Present

Research Scientist, VMware Research — Sep 2016 – May 2021

Mitosis – Transparently Replicating Page-Tables on NUMA machines

- Identified a significant performance issue in large NUMA systems caused by the suboptimal placement of page-tables
- Designed mechanism and policies to mitigate the issue by replicating and migrating page-tables across sockets
- Improved performance of big-data applications by reducing the cost of page-table walks on TLB misses
- Prototyped in the Linux kernel and KVM. Additionally, the design is being implemented in VMware hypervisor

Project PBerry: FPGA Acceleration in a Virtualized Data Center

- Devised a new approach that leverages cache coherence to track application memory accesses instead of virtual memory
- Proposed using cache-coherent FPGAs to track dirty data at cache line granularity for various hypervisor services
- Established a new team in VMware with expertise in designing FPGAs and the concept is being prototyped in FPGAs

Remote Memory

- Proposed abstraction for a process to export its memory to remote hosts, and to access the memory exported by others
- Showed that our abstraction is easy to use and is performant across various remote memory technologies such as RDMA
- Prototyped the abstraction in the Linux kernel and is being productized in the VMware hypervisor

Multi-Application Concurrent GPU Memory Hierarchy

- Showed that applications have vastly different resource demands on the GPU leading to underutilization of its resources
- Redesigned the GPU virtual memory subsystems to enable multiple applications to share a single GPU
- Enabled GPU applications to efficiently use multiple page sizes which improved performance by 55.5% on average

Decoupling Address Translation and Protection

- Illustrated that coarse metadata bits such as dirty bits for larger pages limit OS's visibility to perform many system services
- Decoupled address translation from metadata bits on CPUs and GPUs in a new efficient hardware-software co-design

Composing Synchronization and Crash-Consistency in NVMs using Transactional Memory

- Proposed a single abstraction that provides guarantees on both crash-consistency and correct synchronization
- Composed various transactional memory systems with crash-consistent logging to show the efficacy of the abstraction

Graduate Intern, Intel Corporation — May 2011 – Jan 2012

- Developed an automated framework for performance correlation and debugging in Intel graphics architectures
- Identified severe performance pathologies across cycle-accurate simulator, compiler, hardware design, and test chips

Analog Design Intern, ST Microelectronics — May 2007 – Aug 2007 & Jan 2008 – May 2008

- Designed custom layout for a linear voltage regulator with improved load regulation for their embedded DRAM chips

Academic Experience

Research Assistant, University of Wisconsin – Madison — Jan 2012 – Aug 2016

Agile Paging

- Observed that more than 90% of the page-tables remain static and 10% of the page-table updates are closely clustered
- Combined two state-of-the-art techniques: nested and shadow paging, by switching between them within a page-walk
- Reduced the cost of page-table walks to achieve more than 12% better performance than the best of the two techniques

Redundant Memory Mapping

- Introduced range translations that arbitrarily sized virtual memory ranges to contiguous physical memory pages
- Mapped the entire virtual address space with standard paging and redundantly with multiple range translations
- RMM performs better than standard paging and large pages and reduces the cost of address translation to less than 1%

Direct Segments

- Introduced a form of segmentation to improve address translation while preserving properties of paging when necessary
- Mapped a single contiguous virtual address range to a contiguous physical address range with base, limit, offset registers
- Improved performance of applications by 1.3x and 2.1x while running in the native and virtualized system respectively

Research Assistant, North Carolina State University — Jan 2009 – Aug 2010

- Co-designed *FabScalar*: an automated framework for generating hardware designs of diverse superscalar processors
- Proposed hardware technique to effectively pipeline any branch predictor to arbitrary depths to improve performance

Teaching Experience

- **Teaching Assistant,** Parallel Computer Architecture, *CS/ECE 757*, University of Wisconsin – Madison, Spring'12.
- **Teaching Assistant,** Electrical and Electronic Circuits, ECE 376, University of Wisconsin – Madison, Spring'11.
- **Teaching Assistant,** Electrical and Mechanical Power Systems, ECE 377, University of Wisconsin – Madison, Spring'11.
- **Teaching Assistant,** Computer Systems Programming, *ECE 202*, North Carolina State University, Spring'09.
- **Teaching Assistant and Lab Assistant,** Basics of Electronic Circuits, *EL 102*, Dhirubhai Ambani Institute (DA-IICT), Fall'07.

Honors and Awards

- **IEEE Micro Top Picks 2016 –** Presented for "*Agile Paging for Efficient Memory Virtualization*" from ISCA'16.
- **IEEE Micro Top Picks 2015 –** Presented for "*Range Translations for Fast Virtual Memory*" from ISCA'15.
- **IEEE Micro Top Picks, Honorable Mention 2014 –** Presented for "*Efficient Memory Virtualization*" from MICRO'14.
- **IEEE Micro Top Picks 2011 –** Presented for "FabScalar: Automating Superscalar Core Design" from ISCA'11.

Reviewed Publications

- A. Panwar, R. Achermann, A. Basu, A. Bhattacharjee, G. Kanchi, **J. Gandhi**, "Fast Local Page-Tables for Virtualized NUMA Servers with vMitosis", International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)**, Apr 2021.
- R. Ausavarungnirun, T. Merrifield, **J. Gandhi**, C.J. Rossbach, *"PRISM: Architectural Support for Variable-granularity Memory Metadata"*, International Conference on Parallel Architectures and Compilation Techniques **(PACT)**, Oct 2020.
- M. Agbarya, I. Yaniv, **J. Gandhi**, D. Tsafrir, "*Predicting Execution Times with Partial Simulations in Virtual Memory Research: Why and How*", International Symposium on Microarchitecture **(MICRO)**, Oct 2020.
- C. Alverti, S. Psomadakis, V. Karakostas, **J. Gandhi,** K. Nikas, G. Goumas, N. Koziris, "*Enhancing and Exploiting Contiguity for Fast Memory Virtualization*", International Symposium on Computer Architecture **(ISCA)**, June 2020.
- R. Achermann, A. Panwar, A. Bhattacharjee, T. Roscoe, **J. Gandhi**, *"Mitosis: Transparently Self-Replicating Page-Tables for Large-Memory Machines",* International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)**, Mar 2020.
- I. Calciu, I. Puddu, A. Kolli, A. Nowatzyk, **J. Gandhi**, O. Mutlu, P. Subrahmanyam, "*Project PBerry: FPGA Acceleration for Remote Memory*", Workshop on Hot Topics in Operating Systems **(HotOS)**, May 2019.
- M.K. Aguilera, N. Amit, I. Calciu, X. Deguillard, **J. Gandhi**, S. Novakovic, A. Ramanathan, P. Subrahmanyam, L. Suresh, K. Tati, R. Venkatasubramanian, M. Wei, "*Remote regions: a simple abstraction for remote memory*", USENIX Annual Technical Conference **(USENIX ATC)**, July 2018.
- R. Ausavarungnirun, C.J. Rossbach, V. Miller, J. Landgraf, S. Ghose, **J. Gandhi**, A. Jog, O. Mutlu, "*MASK: Redesigning the GPU Memory Hierarchy to Support Multi-Application Concurrency*", International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)**, Mar 2018.
- R. Ausavarungnirun, C.J. Rossbach, J. Landgraf, V. Miller, S. Ghose, **J. Gandhi**, O. Mutlu, *"MOSAIC: A Transparent Hardware-Software Cooperative Memory Management in GPU"*, International Symposium on Microarchitecture **(MICRO)**, Oct 2017.
- M.K. Aguilera, N. Amit, I. Calciu, X. Deguillard, **J. Gandhi**, P. Subrahmanyam, L. Suresh, K. Tati, R. Venkatasubramanian, M. Wei, "*Remote memory in the age of fast networks*", Symposium on Cloud Computing **(SoCC)**, Sept 2017.
- **J. Gandhi**, M.D. Hill, M.M. Swift, "*Agile Paging for Efficient Memory Virtualization*", IEEE **Micro's Top Picks** from the Computer Architecture Conferences, May 2017.
- **J. Gandhi**, M.D. Hill, M.M. Swift, "*Exceeding the Best of Nested and Shadow Paging*", International Symposium on Computer Architecture **(ISCA)**, Jun 2016. **Also, accepted for Micro Top Picks 2016.**
- **J. Gandhi**, V. Karakostas, F. Ayar, A. Cristal, M.D. Hill, K.S. McKinley, M. Nemirovsky, M.M. Swift, O. Ünsal, "*Range Translations for Fast Virtual Memory*", IEEE **Micro's Top Picks** from the Computer Architecture Conferences, May 2016.
- V. Karakostas, **J. Gandhi**, A. Cristal, M.D. Hill, K.S. McKinley, M. Nemirovsky, M.M. Swift, O. Ünsal, "*Energy-Efficient Address Translation*", Symposium on High Performance Computer Architecture **(HPCA)**, Mar 2016.
- **J. Gandhi**, V. Karakostas, F. Ayar, A. Cristal, M.D. Hill, K.S. McKinley, M. Nemirovsky, M.M. Swift, O. Ünsal, "*Redundant Memory Mappings for Fast Access to Large Memories*", International Symposium on Computer Architecture **(ISCA)**, Jun 2015. **Also, accepted for IEEE Micro Top Picks 2015.**
- **J. Gandhi**, A. Basu, M.D. Hill, M.M. Swift, "*Efficient Memory Virtualization: Reducing Dimensionality of Nested Page Walks*", International Symposium on Microarchitecture **(MICRO)**, Dec 2014. **Also, received honorable mention in Micro Top Picks 2014.**
- A. Basu, **J. Gandhi,** J. Chang, M.D. Hill, M.M. Swift, "*Efficient Virtual Memory for Big Memory Servers*", International Symposium on Computer Architecture **(ISCA)**, June 2013.
- N.K. Choudhary, S.V. Wadhavkar, T.A. Shah, H. Mayukh, **J. Gandhi**, B.H. Dwiel, S. Navada, H.H. Najaf-abadi, E. Rotenberg, "*FabScalar: Automating Superscalar Core Design*", IEEE **Micro's Top Picks** from the Computer Architecture Conferences, May 2012.

- N.K. Choudhary, S.V. Wadhavkar, T.A. Shah, H. Mayukh, **J. Gandhi**, B.H. Dwiel, S. Navada, H.H. Najaf-abadi, E. Rotenberg, "*FabScalar: Composing Synthesizable RTL Designs of Arbitrary Cores within a Canonical Superscalar Template*", International Symposium on Computer Architecture **(ISCA)**, Jun 2011. **Also, accepted for Micro Top Picks 2011.**

Journal Papers, Posters, and Other Publications

- **Paper –** P. Fernando, I. Calciu, **J. Gandhi,** A. Kolli, A. Gavrilovska, *"Characterizing non-volatile memory transactional systems"*, Non-Volatile Memories Workshop **(NVMW)**, Mar 2021.
- **Paper –** P. Fernando, I. Calciu, **J. Gandhi,** A. Kolli, A. Gavrilovska, "*Persistence and Synchronization: Friends or Foes?",* Computing Research Repository **(arXiv)**, Dec 2020.
- **Paper –** J. Lowe-Power, Abdul M. Ahmad, A. Akram, M. Alian, R. Amslinger, M. Andreozzi, A. Armejach, N. Asmussen, S. Bharadwaj, G. Black, G. Bloom, B.R. Bruce, D.R. Carvalho, J. Castrillon, L. Chen, N. Derumigny, S. Diestelhorst, W. Elsasser, M. Fariborz, A. Farmahini-Farahani, P. Fotouhi, R. Gambord, **J. Gandhi**, D. Gope, T. Grass, B. Hanindhito, A. Hansson, S. Haria, A. Harris, T. Hayes, A. Herrera, M. Horsnell, S. Raza Jafri, R. Jagtap, H. Jang, R. Jeyapaul, T.M. Jones, M. Jung, S. Kannoth, H. Khaleghzadeh, Y. Kodama, T. Krishna, T. Marinelli, C. Menard, A. Mondelli, T. Mück, O. Naji, K. Nathella, H. Nguyen, N. Nikoleris, L.E. Olson, M. Orr, B. Pham, P. Prieto, T. Reddy, A. Roelke, M. Samani, A. Sandberg, J. Setoain, B. Shingarov, M.D. Sinclair, T. Ta, R. Thakur, G. Travaglini, M. Upton, N. Vaish, I. Vougioukas, Z. Wang, N. Wehn, C. Weis, D.A. Wood, H. Yoon, E.F. Zulian, *"The gem5 Simulator: Version 20.0+",* Computing Research Repository **(arXiv)**, Jul 2020.
- **Paper –** R. Achermann, A. Panwar, A. Bhattacharjee, T. Roscoe, **J. Gandhi**, *"Mitosis: Transparently Self-Replicating Page-Tables for Large-Memory Machines",* Computing Research Repository **(arXiv)**, Oct 2019.
- **Poster –** R. Achermann, **J. Gandhi**, A. Bhattacharjee, T. Roscoe, "*Mitosis: Transparently Replicating Page-Tables for Big-Memory Machines*", Symposium on Operating Systems Design and Implementation **(OSDI)**, Oct 2018.
- **Paper –** R. Ausavarungnirun, J. Landgraf, V. Miller, S. Ghose, **J. Gandhi**, C.J. Rossbach, O. Mutlu, *"MOSAIC: A Transparent Hardware-Software Cooperative Memory Management in GPU",* ACM SIGOPS Operating Systems Review - Special Topics **(OSR)**, Aug 2018.
- **Paper –** R. Ausavarungnirun, J. Landgraf, V. Miller, S. Ghose, **J. Gandhi**, C.J. Rossbach, O. Mutlu, *"MOSAIC: A Transparent Hardware-Software Cooperative Memory Management in GPU",* Computing Research Repository **(arXiv)**, Apr 2018.
- **Talk –** A. Kolli, **J. Gandhi**, I. Calciu, S. Novakovic, "*Remote Memory Persistency*", Workshop on Wearhouse-scale Memory Systems **(WAMS),** co-located with ASPLOS, Mar 2018.
- **Talk –** I. Calciu, A. Kolli, **J. Gandhi**, S. Novakovic, M.K. Aguilera, R. Venkatasubramanian, P. Subrahmanyam, "*Resource disaggregation for the 99%*", Workshop on Wearhouse-scale Memory Systems **(WAMS),** co-located with ASPLOS, Mar 2018.
- **Poster –** P. Fernando, I. Calciu, **J. Gandhi**, A. Kolli, A. Gavrilovska, "*Persistence and Synchronization: Friends or Foes*", Symposium on Operating Systems Principles **(SOSP)**, Oct 2017.
- **Poster –** M.K. Aguilera, N. Amit, I. Calciu, X. Deguillard, **J. Gandhi**, S. Novakovic, A. Ramanathan, P. Subrahmanyam, L. Suresh, K. Tati, R. Venkatasubramanian, M. Wei, "*Remote regions: a simple abstraction for remote memory*", Symposium on Operating Systems Principles **(SOSP)**, Oct 2017.
- **Paper –** R. Ausavarungnirun, C.J. Rossbach, V. Miller, J. Landgraf, S. Ghose, **J. Gandhi**, A. Jog, O. Mutlu, "*Improving Multi-Application Concurrency Support Within the GPU Memory System"*, Computing Research Repository **(arXiv)**, Aug 2017
- **Paper – J. Gandhi**, A. Basu, M.D. Hill, M.M. Swift, "*BadgerTrap: A Tool to Instrument x86-64 TLB Misses*", Computer Architecture News **(CAN)**, May 2014.

Patents

- **Issued –** G. Zellweger, L. Suresh, **J. Gandhi,** A. Tai, *"Efficiently Managing the Interruption of User-level Critical Sections",* US Patent 10,922,128, issued on Feb 16th, 2021, filled on Dec 19th, 2019.
- **Issued – J. Gandhi,** P. Subrahmanyam, I. Calciu, A. Kolli, *"Accelerating Replication of Page Tables for Multi-Socket Machines",* US Patent 10,929,259, issued on Feb 23rd, 2021, filled on Jan 23rd, 2019.

- **Issued –** A. Kolli, I. Calciu, **J. Gandhi,** P. Subrahmanyam, "*Failure-atomic persistent memory logging using binary translation*", US Patent 10,817,389, issued on Oct 27th, 2020, filled on Jan 24th, 2019.
- **Issued –** I. Calciu, **J. Gandhi,** A. Kolli, P. Subrahmanyam, "*Using cache-coherent FPGAs to accelerate remote access*", US Patent 10,761,984, issued on Sep 1st, 2020, filed on Jul 27th, 2018.
- **Issued – J. Gandhi**, T. Merrifield, C.J. Rossbach, "*Decoupling Memory Metadata Granularity from Page Size*", US Patent 10,430,186, issued on Apr 28th, 2020, filed on Mar 8th, 2018.
- **Issued –** M. Wei, M.K. Aguilera, I. Calciu, S. Novakovic, L. Suresh, **J. Gandhi**, N. Amit, P. Subrahmanyam, X. Deguillard, K. Tati, R. Venkatasubramanian, "*File system interface for remote direct memory access*", US Patent 10,706,005, issued on 7th July 2020, filed on Dec 8th, 2017.
- **Issued –** I. Calciu, **J. Gandhi,** P. Fernando, A. Kolli, "*Speeding up transactions in non-volatile memory using hardware transactional memory*", US Patent 10,430,186, issued on Oct 1st, 2019, filed on Oct 27th, 2017.
- **Issued – J. Gandhi,** M.D. Hill, M.M. Swift, *"Efficient Memory Management System for Computers Supporting Virtual Machines"*, US Patent 9,619,401, issued on Apr 11th, 2017, filed on Feb 20th, 2015.
- **Pending –** A. Kolli, I. Calciu, **J. Gandhi,** P. Subrahmanyam, "*Programming interfaces for accurate dirty data tracking*", US Patent App. 16/256,562, filled on Jan 24th, 2019.
- **Pending –** A. Kolli, I. Calciu, **J. Gandhi,** P. Subrahmanyam, "*Failure-atomic logging for persistent memory systems with cache-coherent FPGAs*", US Patent App. 16/256,571, filled on Jan 24th, 2019.
- **Pending – J. Gandhi,** R. Achermann, "*Transparent self-replicating page tables in computing systems*", US Patent App. 16/256,676, filed on Oct 12th, 2018.
- **Pending –** I. Calciu, **J. Gandhi,** A. Kolli, P. Subrahmanyam, "*Using cache-coherent FPGAs to accelerate live migration of virtual machines*", US Patent App. 16/048,182, filed on Jul 27th, 2018.
- **Pending –** I. Calciu, **J. Gandhi,** A. Kolli, P. Subrahmanyam, "*Using cache-coherent FPGAs to accelerate remote memory write-back*", US Patent App. 16/048,178, filed on Jul 27th, 2018.
- **Pending –** I. Calciu, **J. Gandhi,** A. Kolli, P. Subrahmanyam, "*Using cache-coherent FPGAs to accelerate post-copy migration*", US Patent App. 16/048,183, filed on Jul 27th, 2018.
- **Pending –** I. Calciu, **J. Gandhi,** A. Kolli, P. Subrahmanyam, "*Using cache-coherent FPGAs to track dirty cache lines*" US Patent App. 16/048,180, filed on Jul 27th, 2018.

Press

- Mitosis: An Efficient Way to Boost Application Performance on Large Machines (link) VMware Blog, Oct 2019

Professional Activities

Program Committee Member

- ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2022
- ACM/IEEE International Symposium on Microarchitecture (MICRO), 2021
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2021
- Young Architect Workshop (YArch), 2021
- ACM/IEEE International Symposium on Microarchitecture (MICRO), 2020
- ACM Symposium on Cloud Computing (SoCC), 2020
- Young Architect Workshop (YArch), 2020
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2020
- ACM Symposium on Cloud Computing (SoCC), 2019
- Young Architect Workshop (YArch), 2019
- IEEE International Symposium on High-Performance Computer Architecture (HPCA), 2019
- ACM Symposium on Cloud Computing (SoCC), 2018
- IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), 2018

External Review Committee Member

- IEEE/ACM International Symposium on Microarchitecture (MIRCO), 2019
- ACM International Conference on Supercomputing (ICS), 2019
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2019
- ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2018
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2017
- ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2017

Journal Reviewer

- ACM Transactions on Architecture and Code Optimization (TACO)
- ACM Journal on Emerging Technologies in Computing Systems (JETC)
- ACM Transactions on Storage (TOS)
- IEEE Access
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
- IEEE Computer Architecture Letters (CAL)
- IEEE Transactions on Computers (TC)

Professional Membership

- Association for Computing Machinery (ACM)
- Institute of Electrical and Electronics Engineers (IEEE)
- Advanced Computing Systems Association (USENIX)
- ACM Special Interest Group on Computer Architecture (SIGARCH)
- ACM Special Interest Group on Microarchitecture (SIGMICRO)
- ACM Special Interest Group on Operating Systems (SIGOPS)
- IEEE Technical Committee on Computer Architecture (TCCA)

Open-Source Development

- Mitosis: Mechanisms and Policies for explicit page-table management on large NUMA machines (link)
- MosAlloc: Mosaic memory allocator to analyze large page performance (link)
- CAPaging: Contiguity-aware physical memory allocator for Linux (link)
- BadgerTrap: A tool to instrument TLB misses (link)
- gem5: A modular platform for computer-system architecture research (link)
- FabScalar: Automating superscalar core design (link)

Panels

- *Memory system design*, Panel member, International Symposium on Computer Architecture (ISCA), worldwide event, June 2020.

Professional organization officer

- 2018 – Current: Video Chair for ACM SIGARCH (Special Interest Group on Computer Architecture)