

Sudarsun Kannan

Postdoctoral Research Associate,
Department of Computer Sciences, #7366
University of Wisconsin-Madison
1210 W Dayton St, Madison WI 53706

Email: sudarsun@cs.wisc.edu
<http://cs.wisc.edu/~sudarsun>
Ph: +1404-394-7370

Research Interests

Operating Systems, Computer Architecture, Storage, Distributed Systems, High Performance Computing, Software and Hardware Co-design

Research Summary

My research goal is to build scalable and efficient systems for emerging heterogeneous memory and storage technologies. Given the multi-faceted differences among these technologies, a critical step toward maximizing their impact is to shield applications and upper-level software stacks from the consequent complexity. My research approach is based on empowering operating systems and runtimes with new abstractions and techniques that can tap into the transformative benefits offered by new hardware technologies, while also effectively bridging the performance, capacity and energy gaps among them. I favor a principled approach to unearthing the systems software opportunities and challenges to build practical OS and runtime techniques. My research also studies the hardware-software interfaces and co-design opportunities, particularly in the context of data-intensive applications.

Academic Experience

University of Wisconsin-Madison

Postdoctoral Research Associate, Dept. of Computer Sciences

Sep. 2016 - Current

Topic: File Systems For Modern Storage

Advisers - Prof. Andrea C. Arpaci-Dusseau and Prof. Remzi H. Arpaci-Dusseau

Georgia Institute of Technology, Atlanta

Ph.D. in Computer Science

May 2010 - Aug 2016

Thesis: Operating System Support for Heterogeneous Memory

Advisers - The late Prof. Karsten Schwan and Prof. Ada Gavrilovska

Georgia Institute of Technology, Atlanta

M.S in Computer Science

Aug 2008 - May 2010

Masters Research Project: Virtual Storage for Mobile Devices

Advisers - The late Prof. Karsten Schwan and Prof. Ada Gavrilovska

Anna University, Chennai

B.E in Computer Science (with distinction)

July 2001- May 2005

Research project: Embedded Web Server Management

Reviewed Publications

- Sudarsun Kannan**, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau
Yuangang Wang, Jun Xu, Gopinath Palani
Designing a True Direct-Access File System with DevFS
USENIX Conference on File and Storage Technologies, 2018 (Accepted) **FAST '18**
- Sudarsun Kannan**, Ada Gavrilovska, Vishal Gupta, Karsten Schwan
HeteroOS - OS Design for Heterogeneous Memory Management in Datacenter
44th International Symposium on Computer Architecture, 2017. **ISCA '17**
- Jun He, **Sudarsun Kannan**, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau
The Unwritten Contract of Solid State Drives
European Conference on Computer Systems, 2017. **EuroSys '17**
- Sudarsun Kannan**, Moinuddin Qureshi, Ada Gavrilovska, Karsten Schwan
Reducing Energy Overheads of Memory-based Persistence in NVMs
25th International Conference on Parallel Architectures and Compilation Techniques, 2016. **PACT '16**
- Pradeep Fernando, **Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan
Phoenix: Memory Speed HPC I/O with NVM
23rd annual IEEE International Conference on High Performance Computing, Data, and Analytics, 2016. **HiPC '16**
- Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan
pVM – Persistent Virtual Memory for Efficient Capacity Scaling and Object Storage
European Conference on Computer Systems, 2016. **EuroSys '16**
- Sudarsun Kannan**, Moinuddin Qureshi, Ada Gavrilovska, Karsten Schwan
Energy Aware Persistence: Reducing the Energy Overheads of Persistent Memory
Computer Architecture Letters, 2015 (also presented at NVMW 2016). **CAL '15**
- Sudarsun Kannan**, Naila Farooqui, Ada Gavrilovska, Karsten Schwan
HeteroCheckpoint: Efficient Checkpointing for Accelerator-based Systems
4th Workshop on Fault-Tolerance for HPC at Extreme Scale (in DSN), 2014. **FTXS '14**
- Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan
Reducing the Cost of Persistence for NonvolatileHeaps in End User Devices
20th International Symposium on High Performance Computer Architecture, 2014. **HPCA '14**
- Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan
NVM Heaps for Accelerating Browser-based Applications
USENIX Interactions of NVM/Flash with Operating Systems and Workloads, 2013. **INFLOW '13**
- Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan, Dejan Milojicic
Optimizing Checkpoints Using NVM as Virtual Memory
27th IEEE International Parallel & Distributed Processing Symposium, 2013. **IPDPS '13**
- Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan, Dejan Milojicic, Vanish Talwar
Using Active NVRAM for I/O Staging
Petascale Data Analytics: Challenges and Opportunities, SC workshop, 2011. **PDAC '11**
- Sudarsun Kannan**, Ada Gavrilovska, and Karsten Schwan
Cloud4Home – Enhancing Data Services with @Home Clouds
31st International Conference on Distributed Computing Systems, 2011. **ICDCS '11**

Sudarsun Kannan, Karishma Babu, Ada Gavrilovska, Karsten Schwan
VStore++: Virtual Storage Services for Mobile Devices
International Workshop on Mobile Computing and Clouds, 2010.

MobiCloud '10

Under Review

Sudarsun Kannan, Nitish Bhat, Ada Gavrilovska, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau
Redesigning LSMs for Nonvolatile Memory

Short Papers and Posters

Sudarsun Kannan, Ada Gavrilovska, Vishal Gupta, Karsten Schwan
HeteroOS - OS Design for Heterogeneous Memory Management in Datacenter
Non-volatile Memories Workshop (NVMW), 2018, University of California San Diego.

Jun He, **Sudarsun Kannan**, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau
The Unwritten Contract of Solid State Drives
Non-volatile Memories Workshop (NVMW), 2018, University of California San Diego.

Sudarsun Kannan, Moinuddin Qureshi, Ada Gavrilovska, and Karsten Schwan
Energy Aware Persistence for Nonvolatile Memory
Non-volatile Memories Workshop (NVMW), 2016, University of California San Diego.

Pradeep Fernando, **Sudarsun Kannan**, Ada Gavrilovska, Karsten Schwan
Fast Restarts/Recovery with NVM memory for HPC systems
Non-volatile Memories Workshop (NVMW), 2015, University of California San Diego.

Sudarsun Kannan, Ada Gavrilovska, and Karsten Schwan
Reducing I/O Sandboxing Overheads for Browser-based Applications
Non-volatile Memories Workshop (NVMW), 2013, University of California San Diego.

Sudarsun Kannan, Ada Gavrilovska, and Karsten Schwan
Rich Client Services using Persistent Memory
Non-volatile Memories Workshop (NVMW), 2012, University of California San Diego.

Hrishi Amur, Alex Merritt, **Sudarsun Kannan**, Ada Gavrilovska, and Karsten Schwan
MESSY Library for Memory Consistency on 48 core Intel SCC
Intel Marc Symposium 2011, Hilsboro.

Industrial Patents

Sudarsun Kannan, Palo Faraboschi, Murray McLaren, Dejan Milojicic
Checkpoint Generation
U.S. Patent, Publication number US20140214770 A1.

Sudarsun Kannan, Dejan Milojicic, Vanish Talwar
Active Non Volatile Memory Post Processing
U.S. Patent, Publication number US9619430 B2.

James E. Donahue, Ricky Ho, **Sudarsun Kannan**, Pradnyesh S. Gore
System and Method for Installation and Management of Cloud-Independent Multi-Tenant Applications
U.S. Patent, Publication number US8812627 B2.

Talks

OS design for Heterogeneous Memory Management in Datacenter, ISCA June 2017, Toronto.

Reducing Energy Overheads of Memory-based Persistence in NVMs, PACT, Sep. 2016, Haifa.

Persistent Virtual Memory for Capacity Scaling and Object Storage, Eurosys, Mar. 2016, London.

Heterogeneous Memory Management, Univ. of Wisconsin-Madison, May 2016, Madison.

Energy Aware Persistence, NVMW, March 2016, Univ. of Sand Diego, Sand Diego.

Reducing the Cost of Persistence for Nonvolatile Heaps, HPCA, Feb. 2014, Orlando.

NVM for Rich Client Services, Intel Labs, Aug 2013, Hillsboro, Oregon.

Optimizing Checkpoints Using NVM as Virtual Memory, IPDPS, June 2013, Boston.

NVM Heaps for Accelerating Browser-based Applications, INFLOW (SOSP Workshop), Nov. 2013, Pennsylvania.

Using Active NVRAM for I/O Staging, PDAC, Nov. 2011 (SC Workshop), Seattle.

Active NVRAM for in-memory processing, HP Labs, Aug. 2011, Palo Alto.

Cloud4Home – Enhancing Data Services with @Home Clouds, ICDCS, June 2011, Minneapolis.

Teaching Experience

Guest Lecturer, UW-Madison

CS 736 - Advanced Operating Systems, Fault tolerance and recovery Spring 2017

Guest Lecturer, UW-Madison

CS 739 - Distributed Systems, Logical clocks, SSD storage Fall 2017

Head Teaching Assistant, Georgia Tech

CS 3210 - Undergraduate OS Introduction, Lectures, project design, grading Spring 2016

Guest Lecturer, Georgia Tech

CS 3210, Undergraduate OS Introduction, OS memory management, synchronization Spring 2015

Guest Lecturer, Georgia Tech

CS 6210 - Advanced OS, OS memory management, synchronization, virtualization Fall 2014

Guest Lecturer, Georgia Tech

CS 6210 - Advanced OS, Trust and Protection in the OS Fall 2013

Head Teaching Assistant, Georgia Tech

CS 4210 - Graduate OS Introduction, Lectures, project design and grading Fall 2011

Conference and Journal Services

2017 - Program committee - 46th International Conference on Parallel Processing (ICPP)

2017 - Reviewer - IEEE Transactions on Computers (TC)

2017 - Shadow PC committee - Eurosys

2017 - Reviewer - Science of Computer Programming

2016 - Shadow PC committee - Eurosys

2016 - Reviewer - IEEE Computer Architecture Letters (CAL)

2016 - Reviewer - IEEE Transactions on Computers (TC)

2014 - Reviewer - IEEE Transactions on Cloud Computing (TCC)

2012 - Reviewer - IEEE Transactions on Parallel and Distributed Systems (TPDS)

Industrial Experience

Intel Labs, System Architecture Lab, Hillsbro

May 2013 - Aug 2013

Research Intern with Dr. Sanjay Kumar

I studied the performance impact of NVM on end-user applications by developing a benchmark suite for end-user applications. Intel researchers and our research group used the benchmark for papers published at INFLOW 2013, HPCA 2014, CAL 2015, EuroSys 2016.

HP Labs, Intelligent Infrastructure Lab, Palo Alto

May 2012 - Aug 2012

Research Intern with Dr. Dejan Milojicic

I designed and developed a NVM checkpointing mechanism for HPC application. The work was published in IPDPS 2013 paper. The techniques were also patented, and I was the lead author of the patent.

HP Labs, Intelligent Infrastructure Lab, Palo Alto

May 2011 - Aug 2011

Research Intern with Dr. Dejan Milojicic, Dr. Partha Ranganathan, Dr. Vanish Talwar

I explored and designed a technique for using NVM for in-memory processing. The work was published as a workshop paper at PDAC 2011. Due to its novelty, the work was also patented, and I was the lead author of the patent.

Adobe Labs, Advanced technology Labs, San Jose

May 2010 - Aug 2010

Research Intern with Dr. Jim Donnahue

I worked with Dr. Donnahue to develop a novel application deployment framework in public clouds. The work was patented and is currently used in Adobe cloud services.

Research in Motion, OS development, Fort Lauderdale, Florida

May 2009 - Aug 2009

OS intern

I had an opportunity to learn about mobile OS. I also developed OS-level live-debugging tools for the Blackberry OS.

I|Nautix Technologies, Windows Development, Chennai

May 2005 - July 2008

Senior Application Developer

I was one of the five architects involved in designing next-generation trading platform (www2.netx360.com) for Bank of New York used by more than 100K users. Our goal was to develop a latency-critical framework, which also supports dynamic code compilation techniques for just-in-time deployment.

Midas Communication at Indian Institute of Technology, Chennai

May 2004 - May 2005

Intern

I developed a webserver for micro-controllers that manage wireless communication towers. The project was ranked third out of 60 projects in the undergraduate research competition.

Research Experience

Hardware-accelerated filesystems

Dec 2016 - Present

I worked on designing a new hardware-accelerated filesystem that provides applications with near-hardware latency and high bandwidth without compromising critical properties such as integrity, crash consistency, and security (FAST '18). I collaborated to optimize software-level file systems (EuroSys '17).

Fast key-value stores with NVM

Sept 2016 - Present

Existing disk-friendly log-structured persistent key-value stores suffer from write amplification (extra writes more than the actual data size) and lack parallelism even when using fast storage technologies such as NVM. To address this, I designed a new persistent key-value store that exploits NVM byte-addressability and lower latency. The work is under submission.

Designing OS for heterogeneous memory

Oct 2015 - August 2016

There is a growing consensus to support heterogeneous memory technologies such as stacked 3D-DRAM and NVM in the system stack for addressing the DRAM capacity scaling and performance limitations. I designed an application-transparent heterogeneous memory management mechanism for virtualized data center systems. To support memory heterogeneity, I redesigned the Linux OS's and the hypervisor's virtual memory management (ISCA '17).

Energy aware persistence

Oct 2015 - August 2016

I collaborated with Professor Moinuddin Qureshi (Georgia Tech) and my advisers to dissect the energy overheads of persistence in NVMs. I analyzed the energy cost of four main properties of persistence – atomicity, consistency, isolation, and durability (ACID). We concluded that data and metadata durability as the most energy expensive property. I also developed techniques to reduce the NVM energy usage without affecting the correctness of an application (PACT '16).

pVM - Persistent virtual memory

Aug 2013 - Oct 2015

To exploit the large capacity of NVMs, I designed pVM (persistent virtual memory) that can seamlessly expand main memory to NVM for higher capacity and also provide fast persistent object storage. pVM extends the OS virtual memory instead of the file system for persistent storage because of their cache and TLB efficiency. I presented this research at EuroSys '16.

Cache overheads on NVM-based persistence

May 2012 - Aug 2013

I studied the implications of using byte-addressable NVMs for their dual use of additional memory capacity and storage. I observed that applications using NVM hide NVM write latency by buffering writes to the cache; however, frequent cache flushes impact all applications that share the cache. To overcome this problem, I proposed an OS and library-level design that efficiently partitions the cache between applications using the NVM for persistence versus applications using NVM or DRAM for additional heap capacity. (HPCA '14).

Memory checkpointing for HPC applications

May 2012 - Aug 2013

High-performance computing applications are long-running and use thousands of cores. To recover and restart from a system failure, these applications checkpoint their intermediate state which can waste significant compute cycles. Just replacing disks with byte-addressable NVMs will not solve this problem due to substantially large checkpoint sizes coupled with limited NVM write bandwidth. To address this problem, I collaborated with HP Labs to design a high-performance checkpoint/restart mechanism that exploits the memory-ness of NVM to overcome the limited NVM bandwidth problem (IPDPS '13).

Cloud4Home

Aug 2010 - Dec 201

I studied the benefits of combining the cloud infrastructure and the hardware capabilities of home devices for time-sensitive applications (e.g., intruder detection). I designed a framework to virtualize CPU, memory, and storage across home devices and seamlessly use them across multiple home devices. The research was published at ICDCS '11.

Student Mentoring

Hakan Memisoglu, Ph.D. Student at UW-Madison	Fall 2017-current
Optimizing file system for fast storage	

Pradeep Fernando, Ph.D. Student at Georgia Tech	Fall 2015-current
Resilience in HPC applications and other topics	

Thaleia-Dimitra Doudali, Ph.D. Student at Georgia Tech	Spring 2016
Resource allocation in Cloud	

Andrea Hu, BS Student at Georgia Tech	Spring 2016
Graph analytics and storage performance	

Albert, BS Student at Georgia Tech	Spring 2016
Persistent memory durability overheads	

Nitish Bhatt, MS Student at Georgia Tech (now at VMWare)	2016
NoSQL database on new memory technologies	

Amaro Emmanuel, MS Student at Georgia Tech, (Ph.D. Student, U.C. Berkeley)	Spring 2015
Reducing virtualization system call cost	

Preethi Sreenivasan, MS Student at Georgia Tech, (now at Amazon)	Summer 2015
Architectural analysis of thin clients	

Ravi Mangal, Ph.D. Student at Georgia Tech	Spring 2011
Virtual storage for @home cloud	

Sasi Siddharth, MS Student @ Georgia Tech, (now at HP)	Spring 2011
Virtual storage for @home cloud	

Awards and Achievements

Nominated for Georgia Tech-wide Outstanding Teaching Assistant Award, **One nominee** from School of Computer Science, Fall 2012.

Memorable Paper Award Invitation, NVMW '18, "The Unwritten Contract of Solid State Drives" (results yet to be announced).

Research on NVM funded by Intel Labs Persistent Memory ISRA for multiple years. I actively collaborated and presented the research progress to Intel every quarter.

Associate of Month Award, I|Nautix Technologies, Windows development, March 2007.

Travel Grants for ICDCS '11, NVMW '11, NVMW '12, IPDPS '13, SOSP '13, NVMW '13, HPCA '14, NVMW '14, EuroSys '16 Shadow PC, EuroSys '16, NVMW '16, OSDI '16, ISCA '17.

Tools and Software Contributions

NVM memory and object storage manager. <https://github.com/SudarsunKannan/nvmalloc>

WiscSee - SSD analyzer. <https://github.com/junhe/wiscsee>

Memory bandwidth and latency throttler. <https://github.com/SudarsunKannan/Thermalthrottling>

Memory and cache performance analyzer. <https://github.com/SudarsunKannan/memlatency>

Leadership Activities

Student Representative, Faculty Recruitment Committee, Georgia Tech Jan 2015 - Aug 2016

Lab Reading Group Organizer, CERCS, Georgia Tech Jan 2014 - Aug 2016

Research Award Committee, Undergraduate research competition, Georgia Tech 2014 - 2015

Technical Interviewer, I-Nautix Technologies, Chennai 2007 - 2008

Student Dept. Secretary, Computer Society, SaiRam Eng. College, Anna University 2004 - 2005

References

Prof. Ada Gavrilovska
College of Computing
Georgia Institute of Technology
ada@cc.gatech.edu
Ph: +1 (404) 894-0387

Prof. Andrea C. Arpaci-Dusseau
Department of Computer Sciences
University of Wisconsin-Madison
dusseau@cs.wisc.edu
Ph: +1 (608) 263-7764

Prof. Remzi H. Arpaci-Dusseau
Department of Computer Sciences
University of Wisconsin-Madison
remzi@cs.wisc.edu
Ph: +1 (608) 265-6013

Prof. Moinuddin Qureshi
School of Electrical and Computer Engineering
Georgia Institute of Technology
moin@ece.gatech.edu
Ph: +1 (512) 565-9463