

# KAI SHEN

## Employment:

<i>Senior Staff Engineer</i> , Google, New York, NY.	2016 – present
<i>Visiting Researcher</i> , Google, New York, NY.	2015–2016
<i>Associate Professor</i> in Computer Science, University of Rochester, Rochester, NY.	2008–2016
<i>Assistant Professor</i> in Computer Science, University of Rochester, Rochester, NY.	2002–2008
<i>Software Engineer</i> , Teoma Technologies, Piscataway, NJ.	Summer 2000/2001

## Education:

<i>Ph.D.</i> in Computer Science, University of California at Santa Barbara. Thesis: “Clustering, Resource Management, and Replication Support for Scalable Network Services”. Advisor: Tao Yang.	2002
<i>B.Eng.</i> in Computer Science and Engineering, Shanghai Jiaotong University, China.	1996

## Publications:

- [1] Meng Zhu and Kai Shen, “Energy Discounted Computing on Multicore Smartphones”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’16)*, Denver, CO, June 2016.
- [2] Sharanyan Srikanthan, Sandhya Dwarkadas, and Kai Shen, “Coherence Stalls or Latency Tolerance: Informed CPU Scheduling for Socket and Core Sharing”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’16)*, Denver, CO, June 2016.
- [3] Zhuan Chen and Kai Shen, “OrderMergeDedup: Efficient, Failure-Consistent Deduplication on Flash”. In *Proceedings of the Fourteenth USENIX Conference on File and Storage Technologies (FAST’16)*, Santa Clara, CA, February 2016.
- [4] Sharanyan Srikanthan, Sandhya Dwarkadas, and Kai Shen, “Data Sharing or Resource Contention: Toward Performance Transparency on Multicore Systems”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’15)*, Pages 529–540, Santa Clara, CA, July 2015.
- [5] Qi Zhu, Meng Zhu, Bo Wu, Xipeng Shen, Kai Shen, and Zhiying Wang, “Software Engagement with Sleeping CPUs”. In *Proceedings of the Fifteenth Workshop on Hot Topics in Operating Systems (HotOS XV)*, Kartause Ittingen, Switzerland, May 2015.
- [6] Konstantinos Menychtas, Kai Shen, and Michael L. Scott, “Disengaged Scheduling for Fair, Protected Access to Fast Computational Accelerators”. In *Proceedings of the Nineteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’14)*, Pages 301–316, Salt Lake City, UT, March 2014.
- [7] Kai Shen, Stan Park, and Meng Zhu, “Journaling of Journal Is (Almost) Free”. In *Proceedings of the Twelfth USENIX Conference on File and Storage Technologies (FAST’14)*, Pages 287–293, Santa Clara, CA, February 2014.
- [8] Kai Shen and Stan Park, “FlashFQ: A Fair Queueing I/O Scheduler for Flash-Based SSDs”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’13)*, Pages 67–78, San Jose, CA, June 2013.

- [9] Konstantinos Menychtas, Kai Shen, and Michael L. Scott, “Enabling OS Research by Inferring Interactions in the Black-Box GPU Stack”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’13)*, Pages 291–296, San Jose, CA, June 2013.
- [10] Stan Park, Terence Kelly, and Kai Shen, “Failure-Atomic msync(): A Simple and Efficient Mechanism for Preserving the Integrity of Durable Data”. In *Proceedings of the European Conference on Computer Systems (EuroSys’13)*, Pages 225–238, Prague, Czech Republic, April 2013.
- [11] Kai Shen, Arrvindh Shriraman, Sandhya Dwarkadas, Xiao Zhang, and Zhuan Chen, “Power Containers: An OS Facility for Fine-Grained Power and Energy Management on Multicore Servers”. In *Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’13)*, Pages 65–76, Houston, TX, March 2013.
- [12] Xiao Zhang, Rongrong Zhong, Sandhya Dwarkadas, and Kai Shen, “A Flexible Framework for Throttling-Enabled Multicore Management (TEMM)”. In *Proceedings of the Forty-First International Conference on Parallel Processing (ICPP’12)*, Pages 389–398, Pittsburgh, PA, September 2012.
- [13] Amal Fahad, Tolga Soyata, Tai Wang, Gaurav Sharma, Wendi Heinzelman, and Kai Shen, “SOLAR-CAP: Super Capacitor Buffering of Solar Energy for Self-Sustainable Field Systems”. In *Proceedings of the Twenty-Fifth IEEE International System-on-Chip Conference*, Pages 236–241, Niagara Falls, NY, September 2012.
- [14] Amal Fahad, Zhuan Chen, Kai Shen, and Jeffrey Bigham, “An Evaluation of Web Acceleration Techniques for the Developing World”. In *Proceedings of the Sixth USENIX/ACM Workshop on Networked Systems for Developing Regions*, Boston, MA, June 2012.
- [15] Xiaoya Xiang, Bin Bao, Chen Ding, and Kai Shen, “Cache Conscious Task Regrouping on Multicore Processors”. In *Proceedings of the Twelfth IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, Pages 603–611, Ottawa, Canada, May 2012.
- [16] Stan Park and Kai Shen, “FIOS: A Fair, Efficient Flash I/O Scheduler”. In *Proceedings of the Tenth USENIX Conference on File and Storage Technologies (FAST’12)*, San Jose, CA, February 2012.
- [17] Christopher Stewart, Kai Shen, Arun Iyengar, and Jian Yin, “EntomoModel: Understanding and Avoiding Performance Anomaly Manifestations”. In *Proceedings of the Eighteenth IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS’10)*, Best Paper Award, Pages 3–13, Miami Beach, FL, August 2010.
- [18] Xin Li, Michael C. Huang, Kai Shen, and Lingkun Chu, “A Realistic Evaluation of Memory Hardware Errors and Software System Susceptibility”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’10)*, Boston, MA, June 2010.
- [19] Xiao Zhang, Kai Shen, Sandhya Dwarkadas, and Rongrong Zhong, “An Evaluation of Per-Chip Nonuniform Frequency Scaling on Multicores”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’10)*, Boston, MA, June 2010.
- [20] Kai Shen, “Request Behavior Variations”. In *Proceedings of the Fifteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’10)*, Pages 103–116, Pittsburgh, PA, March 2010.
- [21] Christopher Stewart and Kai Shen, “Some Joules Are More Precious Than Others: Managing Renewable Energy in the Datacenter”. In *Workshop on Power Aware Computing and Systems (HotPower’09)*, Big Sky, MT, October 2009.
- [22] Stan Park and Kai Shen, “A Performance Evaluation of Scientific I/O Workloads on Flash-Based SSDs”. In *Workshop on Interfaces and Architectures for Scientific Data Storage*, New Orleans, LA, September 2009.

- [23] Xiao Zhang, Sandhya Dwarkadas, and Kai Shen, “Hardware Execution Throttling for Multi-core Resource Management”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’09)*, San Diego, CA, June 2009.
- [24] Kai Shen, Christopher Stewart, Chuanpeng Li, and Xin Li, “Reference-Driven Performance Anomaly Identification”. In *Proceedings of the ACM SIGMETRICS*, Pages 85–96, Seattle, WA, June 2009.
- [25] Xiao Zhang, Sandhya Dwarkadas, and Kai Shen, “Towards Practical Page Coloring-based Multi-core Cache Management”. In *Proceedings of the Fourth European Conference on Computer Systems (EuroSys’09)*, Pages 89–102, Nuremberg, Germany, April 2009.
- [26] Terence Kelly, Kai Shen, Alex Zhang, and Christopher Stewart, “Operational Analysis of Parallel Servers”. In *Proceedings of the Sixteenth IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS’08)*, Baltimore, MD, September 2008.
- [27] Christopher Stewart, Kai Shen, Arun Iyengar, and Jian Yin, “Configuration-Space Performance Anomaly Depiction”. In *the Second Workshop on Large-Scale Distributed Systems and Middleware*, IBM Watson Research Center, New York, September 2008.
- [28] Christopher Stewart, Matthew Leventi, and Kai Shen, “Empirical Examination of A Collaborative Web Application”. In *Proceedings of the IEEE International Symposium on Workload Characterization - Special Session on Benchmark Innovation*, Seattle, WA, September 2008.
- [29] Ming Zhong, Pin Lu, Kai Shen, and Joel Seiferas, “Optimizing Data Popularity Conscious Bloom Filters”. In *Proceedings of the Twenty-seventh ACM Symposium on Principles of Distributed Computing (PODC’08)*, Pages 355–364, Toronto, Canada, August 2008.
- [30] Christopher Stewart, Terence Kelly, Alex Zhang, and Kai Shen, “A Dollar from 15 Cents: Cross-Platform Management for Internet Services”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’08)*, Pages 199–212, Boston, MA, June 2008.
- [31] Ming Zhong, Kai Shen, and Joel Seiferas, “Correlation-Aware Object Placement for Multi-Object Operations”. In *Proceedings of the Twenty-Eighth International Conference on Distributed Computing Systems (ICDCS’08)*, Pages 512–521, Beijing, China, June 2008.
- [32] Ming Zhong, Kai Shen, and Joel Seiferas, “The Convergence-Guaranteed Random Walk and Its Applications in Peer-to-Peer Networks”. In *IEEE Transactions on Computers*, Volume 57, Number 5, Pages 619–633, May 2008.
- [33] Ming Zhong, Kai Shen, and Joel Seiferas, “Object Replication Degree Customization for High Availability”. In *Proceedings of the Third European Conference on Computer Systems (EuroSys’08)*, Pages 55–68, Glasgow, Scotland, April 2008.
- [34] Kai Shen, Ming Zhong, Sandhya Dwarkadas, Chuanpeng Li, Christopher Stewart, and Xiao Zhang, “Hardware Counter Driven On-the-Fly Request Signatures”. In *Proceedings of the Thirteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS’08)*, Pages 189–200, Seattle, WA, March 2008.
- [35] Pin Lu and Kai Shen, “Multi-Layer Event Trace Analysis for Parallel I/O Performance Tuning”. In *Proceedings of the Thirty-Sixth International Conference on Parallel Processing (ICPP’07)*, XiAn, China, September 2007.
- [36] Pin Lu and Kai Shen, “Virtual Machine Memory Access Tracing With Hypervisor Exclusive Cache”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’07)*, Pages 29–43, Santa Clara, CA, June 2007.
- [37] Xin Li, Kai Shen, Michael C. Huang, and Lingkun Chu, “A Memory Soft Error Measurement on Production Systems”. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC’07)*, Pages 275–280, Santa Clara, CA, June 2007.

- [38] Xin Li, Michael C. Huang, Kai Shen, and Lingkun Chu, “An Empirical Study of Memory Hardware Errors in A Server Farm”. In *the Third Workshop on Hot Topics in System Dependability (HotDep’07)*, Edinburgh, UK, June 2007.
- [39] Chuangpeng Li, Chen Ding, and Kai Shen, “Quantifying The Cost of Context Switch”. In *the Workshop on Experimental Computer Science* at the Federated Computing Research Conference, San Diego, CA, June 2007.
- [40] Xiao Zhang, Sandhya Dwarkadas, Girts Folkmanis, and Kai Shen, “Processor Hardware Counter Statistics As A First-Class System Resource”. In *the Eleventh Workshop on Hot Topics in Operating Systems (HotOS XI)*, San Diego, CA, May 2007.
- [41] Chuanpeng Li, Kai Shen, and Athanasios E. Papathanasiou, “Competitive Prefetching for Concurrent Sequential I/O”. In *Proceedings of the Second European Conference on Computer Systems (EuroSys’07)*, Pages 189–202, Lisbon, Portugal, March 2007.
- [42] Kai Shen, “Parallel Sparse LU Factorization on Different Message Passing Platforms”. In *Journal of Parallel and Distributed Computing (JPDC)*, Volume 66, Issue 11, Pages 1387–1403, November 2006.
- [43] Christopher Stewart, Ming Zhong, Kai Shen, and Thomas O’Neill, “Comprehensive Depiction of Configuration-dependent Performance Anomalies in Distributed Server Systems”. In *the Second Workshop on Hot Topics in System Dependability (HotDep’06)*, Seattle, WA, November 2006.
- [44] Ming Zhong and Kai Shen, “Random Walk Based Node Sampling in Self-Organizing Networks”. In *ACM SIGOPS Operating Systems Review*, Volume 40, Number 3, Pages 49–55, July 2006.
- [45] Peter DeRosa, Kai Shen, Christopher Stewart, and Jonathan Pearson, “Realism and Simplicity: Disk Simulation for Instructional OS Performance Evaluation”. In *Proceedings of the Thirty-Seventh ACM Symposium on Computer Science Education (SIGCSE’06)*, Pages 308–312, Houston, TX, March 2006.
- [46] Ming Zhong and Kai Shen, “Popularity-Biased Random Walks for Peer-to-Peer Search under the Square-Root Principle”. In *Proceedings of the Fifth International Workshop on Peer-to-Peer Systems*, Santa Barbara, CA, February 2006.
- [47] Kai Shen, Ming Zhong, and Chuanpeng Li, “I/O System Performance Debugging Using Model-driven Anomaly Characterization”. In *Proceedings of the Fourth USENIX Conference on File and Storage Technologies (FAST’05)*, Pages 309–322, San Francisco, CA, December 2005.
- [48] Chuanpeng Li and Kai Shen, “Managing Prefetch Memory for Data-Intensive Online Servers”. In *Proceedings of the Fourth USENIX Conference on File and Storage Technologies (FAST’05)*, Pages 253–266, San Francisco, CA, December 2005.
- [49] Kai Shen, “Parallel Sparse LU Factorization on Second-class Message Passing Platforms”. In *Proceedings of the Nineteenth ACM International Conference on Supercomputing (ICS’05)*, Pages 351–360, Cambridge, MA, June 2005.
- [50] Ming Zhong, Justin Moore, Kai Shen, and Amy Murphy, “An Evaluation and Comparison of Current Peer-to-Peer Full-Text Keyword Search Techniques”. In *the Eighth International Workshop on the Web and Databases*, Baltimore, MD, June 2005.
- [51] Christopher Stewart and Kai Shen, “Performance Modeling and System Management for Multi-component Online Services”. In *Proceedings of the Second USENIX Symposium on Networked Systems Design and Implementation (NSDI’05)*, Pages 71–84, Boston, MA, May 2005.
- [52] Ming Zhong, Kai Shen, and Joel Seiferas, “Non-uniform Random Membership Management in Peer-to-Peer Networks”. In *Proceedings of the Twenty-Fourth IEEE INFOCOM*, Pages 1151–1161, Miami, FL, March 2005.

- [53] Lingkun Chu, Kai Shen, Hong Tang, Tao Yang, and Jingyu Zhou, “Dependency Isolation for Thread-based Multi-tier Internet Services”. In *Proceedings of the Twenty-Fourth IEEE INFOCOM*, Pages 796–806, Miami, FL, March 2005.
- [54] Kai Shen, Lingkun Chu, and Tao Yang, “Supporting Cluster-based Network Services on Functionally Symmetric Software Architecture”. In *Proceedings of the High Performance Computing, Networking and Storage Conference (SuperComputing’04)*, Pittsburgh, PA, November 2004.
- [55] Christopher Stewart, Kai Shen, Sandhya Dwarkadas, Michael L. Scott, and Jian Yin, “Profile-driven Component Placement for Cluster-based Online Services”. In the (Refereed) Work-in-Progress Track of the *Fifth ACM/IFIP/USENIX International Middleware Conference*, Toronto, Canada, October 2004. Also published in *IEEE Distributed Systems Online*, Volume 5, Number 10, October 2004.
- [56] Kai Shen and Yuan Sun, “Distributed Hashtable on Pre-Structured Overlay Networks”. In *Proceedings of the Ninth International Workshop on Web Caching and Content Distribution*, Pages 63–81, Beijing, China, October 2004.
- [57] Chuanpeng Li, Athanasios E. Papathanasiou, and Kai Shen, “Competitive Prefetching for Data-Intensive Online Servers”. In *the First Workshop on Operating System and Architectural Support for the on demand IT InfraStructure*, Boston, MA, October 2004.
- [58] Kai Shen, “Structure Management for Scalable Overlay Service Construction”. In *Proceedings of the First USENIX Symposium on Networked Systems Design and Implementation (NSDI’04)*, Pages 281–294, San Francisco, CA, March 2004.
- [59] Kai Shen, Tao Yang, and Lingkun Chu, “Clustering Support and Replication Management for Scalable Network Services”. In *IEEE Transactions on Parallel and Distributed Systems – Special Issue on Middleware*, Volume 14, Number 11, Pages 1168–1179, November 2003.
- [60] Lingkun Chu, Hong Tang, Tao Yang, and Kai Shen, “Optimizing Data Aggregation for Cluster-based Internet Services”. In *Proceedings of the Ninth ACM Symposium on Principles and Practice of Parallel Programming (PPoPP’03)*, Pages 119–130, San Diego, CA, June 2003.
- [61] Kai Shen, Hong Tang, Tao Yang, and Lingkun Chu, “Integrated Resource Management for Cluster-based Internet Services”. In *Proceedings of the Fifth USENIX Symposium on Operating Systems Design and Implementation (OSDI’02)*, Pages 225–238, Boston, MA, December 2002.
- [62] Kai Shen, Tao Yang, and Lingkun Chu, “Cluster Load Balancing for Fine-grain Network Services”. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS’02)*, Fort Lauderdale, FL, April 2002.
- [63] Kai Shen, Tao Yang, Lingkun Chu, JoAnne L. Holliday, Douglas A. Kuschner, and Huican Zhu, “Neptune: Scalable Replica Management and Programming Support for Cluster-based Network Services”. In *Proceedings of the Third USENIX Symposium on Internet Technologies and Systems (USITS’01)*, Pages 197–208, San Francisco, CA, March 2001.
- [64] Hong Tang, Kai Shen, and Tao Yang, “Program Transformation and Runtime Support for Threaded MPI Execution on Shared Memory Machines”. In *ACM Transactions on Programming Languages and Systems (TOPLAS)*, Volume 22, Number 4, Pages 673–700, November 2000.
- [65] Kai Shen, Tao Yang, and Xiangmin Jiao, “S<sup>+</sup>: Efficient 2D Sparse LU Factorization on Parallel Machines”. In *SIAM Journal on Matrix Analysis and Applications*, Volume 22, Number 1, Pages 282–305, 2000.
- [66] Kai Shen, Hong Tang, and Tao Yang, “Adaptive Two-level Thread Management for Fast MPI Execution on Shared Memory Machines”. In *Proceedings of the IEEE/ACM High Performance Networking and Computing Conference (SuperComputing’99)*, Portland, OR, November 1999.

- [67] Hong Tang, Kai Shen, and Tao Yang, “Compile/Run-time Support for Threaded MPI Execution on Multiprogrammed Shared Memory Machines”. In *Proceedings of the Seventh ACM Symposium on Principles and Practice of Parallel Programming (PPoPP’99)*, Pages 107–118, Atlanta, GA, May 1999.
- [68] Bin Jiang, Steven Richman, Kai Shen, and Tao Yang, “Fast Sparse LU Factorization with Lazy Space Allocation”. In *Proceedings of the Ninth SIAM Parallel Processing Conference on Scientific Computing*, San Antonio, TX, March 1999.
- [69] Anurag Acharya, Huican Zhu, and Kai Shen, “Adaptive Algorithm for Cache-efficient Trie Search”. In *the ACM/SIAM Workshop on Algorithm Engineering and Experimentation*, Baltimore, MD, January 1999.
- [70] Kai Shen, Xiangmin Jiao, and Tao Yang, “Elimination Forest Guided 2D Sparse LU Factorization”. In *Proceedings of the Tenth ACM Symposium on Parallel Architectures and Algorithms (SPAA’98)*, Pages 5–15, Puerto Vallarta, Mexico, June 1998.

## Research Funding:

### As the Principal Investigator:

- National Science Foundation CNS-1239423, “CPS: Synergy: Self-Sustainable Data-Driven Systems In the Field”. \$836,000, October 2012 – September 2016. PI: Kai Shen. Co-PI: Tolga Soyata, Gaurav Sharma, and Wendi B. Heinzelman.
- National Science Foundation CNS-1217372, “CSR: Small: System Support for SSD-Backed Recoverable Network Applications”. \$400,000, September 2012 – August 2016. PI: Kai Shen.
- National Science Foundation CCF-1255729 / Semiconductor Research Corporation Contract No. 2013-HJ-2405, “Software Susceptibility-Driven Non-Uniform Memory Error Protection”. \$345,000, April 2013 – March 2016. PI: Kai Shen. Co-PI: Michael C. Huang.
- Google Research Award. “Online Request Adaptation on Multicore Servers”. \$65,000, August 2010 – July 2013. PI: Kai Shen.
- National Science Foundation CCF-0937571, “Collaborative Research: Automatic Extraction of Parallel I/O Benchmarks from HEC Applications”. \$245,974, September 2009 – August 2013. PI: Kai Shen.  
In collaboration with grant CCF-0937908 awarded to the North Carolina State University (lead institution, PI: Xiaosong Ma, co-PI: Frank Mueller) and grant CCF-0938064 awarded to the University of Illinois (PI: Marianne Winslett).
- IBM Faculty Award, “Performance Anomaly Depiction for Distributed Systems”. \$30,000, August 2007 – July 2012. PI: Kai Shen.
- National Science Foundation CNS-0615045, “CSR—PDOS: Model-Driven Comprehensive Performance Anomaly Characterization for System Software”. \$256,998, September 2006 – August 2010. PI: Kai Shen.
- National Science Foundation CCF-0621472, “HECURA: Concurrent I/O Management for Cluster-based Parallel Storages”. \$295,282, August 2006 – July 2009. PI: Kai Shen.
- IBM Faculty Award, “Comprehensive Characterization of Configuration-dependent Performance Anomalies in Distributed Server Systems”. \$30,000, August 2006 – July 2007. PI: Kai Shen.
- Ask.com, A gift in distributed networking research. \$14,375, March 2005 – February 2006. PI: Kai Shen.
- National Science Foundation CCF-0448413, “CAREER: System Support for Data-Intensive Online Applications”. \$401,459, January 2005 – December 2011. PI: Kai Shen.

National Science Foundation CCR-0306473, “Improving the Robustness of Cluster-based Network Services Using Peer-to-Peer Design Principles”. \$249,941, August 2003 – July 2007. PI: Kai Shen.

National Science Foundation ITR/IIS-0312925, “Keyword Searching in Loosely-Coupled Distributed Systems”. \$369,877, July 2003 – June 2007. PI: Kai Shen. Co-PI: Amy L. Murphy.  
PI/co-PI switch after the original PI Murphy left the University of Rochester.

### **As a co-PI:**

National Science Foundation CNS-1319353, “CSR: Small: Managing Multicore Energy for Emerging Applications and Devices”. \$433,673, October 2013 – September 2016. PI: Sandhya Dwarkadas. Co-PI: Kai Shen.

National Science Foundation CNS-1319417, “CSR: Small: First-Class Operating System Management of Computational Accelerators”. \$455,255, October 2013 – September 2016. PI: Michael L. Scott. Co-PI: Kai Shen.

National Science Foundation CCF-1016902, “SHF:Small: Auxiliary Hardware/Software Mechanisms for Flexible Memory Access Control”. \$459,980, September 2010 – August 2014. PI: Sandhya Dwarkadas. Co-PI: Kai Shen.

National Science Foundation CNS-0834451, “CSR–PSCE, SM: Operating System-Level Resource Management in the Multi-Core Era”. \$400,000, September 2008 – August 2012. PI: Sandhya Dwarkadas. Co-PI: Kai Shen.

National Science Foundation CNS-0509270, “CSR–AES: Program Phase Detection and Exploitation”. \$519,834, July 2005 – June 2009. PI: Chen Ding. Co-PIs: John B. Carter, Sandhya Dwarkadas, Michael C. Huang, and Kai Shen. Senior Personnel: Michael L. Scott.

National Science Foundation ACIR-0234346, “Cluster-based Runtime Support for Data-Intensive Online Applications”. \$332,827, March 2003 – February 2008. PI: Tao Yang (University of California at Santa Barbara). Co-PI: Kai Shen.

### **Awards and Honors:**

Google Research Award, 2010.

Best Paper Award at the Eighteenth IEEE MASCOTS Symposium, 2010.

IBM Faculty Award, 2007 and 2006.

National Science Foundation CAREER Award, 2004.

First Prize (Gold Medal) in the Thirty-Third International Mathematics Olympics (for high school students), 1992.

### **Professional Activities:**

Program Committee Member for:

2014/2013/2011 USENIX Annual Technical Conference;

2014 ACM Symposium on Cloud Computing;

2010/2009/2008 IEEE MASCOTS Symposium;

2009/2008 International Conference on Distributed Computing Systems;

2007/2005/2004 International Conference on Parallel Processing;

2007 USENIX Conference on File and Storage Technologies;

2007 IEEE International Symposium on Performance Analysis of Systems and Software.

External Review Committee Member for the 2016/2015 Conference on Architectural Support for Programming Languages and Operating Systems.

Work-in-Progress Chair for the 2007 USENIX Conference on File and Storage Technologies.

Registration Chair for the 2003 ACM Symposium on Operating Systems Principles.

Principal developer of *Neptune*, a middleware system for data center networked services. It was deployed as the platform software backbone for web sites and services at Ask.com, where it managed thousands of servers.

Invited participation at the 2007/2006 Google Faculty Summits.

## Presentations and Invited Talks:

(Since September 2002)

“Resource Isolation and Performance In the Era of Fast Devices,”  
*HP Labs*, Palo Alto, CA, February 2016.

“Software Engagement with Sleeping CPUs,”  
*Fifteenth Workshop on Hot Topics in Operating Systems*, Kartause Ittingen, Switzerland, May 2015.

“Harnessing the Hardware Execution Statistics for Dependable Multicore Servers,”  
*Rutgers University Department of Computer Science*, April 2014;  
*University at Buffalo Computer Science and Engineering Department*, December 2013;  
*Google*, Mountain View, CA, June 2013.

“Journaling of Journal Is (Almost) Free,”  
*Twelfth USENIX FAST Conference*, Santa Clara, CA, February 2014.

“FlashFQ: A Fair Queueing I/O Scheduler for Flash-Based SSDs,”  
*USENIX Annual Technical Conference*, San Jose, CA, June 2013.

“Power Containers: An OS Facility for Fine-Grained Power and Energy Management on Multicore Servers,”  
*Eighteenth ASPLOS Conference*, Houston, TX, March 2013.

“Reference-Driven Performance Anomaly Identification,”  
*Oak Ridge National Laboratory*, Oak Ridge, TN, July 2010;  
*ACM SIGMETRICS Conference*, Seattle, WA, June 2009.

“Request Behavior Variations,”  
*Fifteenth ASPLOS Conference*, Pittsburgh, PA, March 2010.

“Hardware Counter Driven On-the-Fly Request Signatures,”  
*Ask Jeeves Inc.*, Edison, NJ, December 2008;  
*Thirteenth ASPLOS Conference*, Seattle, WA, March 2008.

“Configuration-Space Performance Anomaly Depiction,”  
*Second LADIS Workshop*, IBM Watson Research Center, New York, September 2008.

“Optimizing Data Popularity Conscious Bloom Filters,”  
*Twenty-Seventh ACM PODC Symposium*, Toronto, Canada, August 2008.

“Correlation-Aware Object Placement for Multi-Object Operations,”  
*Twenty-Eighth ICDCS Conference*, Beijing, China, June 2008.

- “Performance Anomaly Characterization of Complex Software Systems,”  
*Cornell University Computer Science Department*, March 2008;  
*University of California at Santa Barbara Computer Science Department*, November 2007;  
*Xerox Research Center*, Webster, NY, September 2006.
- “Multi-Layer Event Trace Analysis for Parallel I/O Performance Tuning,”  
*Thirty-Sixth ICPP Conference*, XiAn, China, September 2007.
- “Concurrent I/O Management for Cluster-based I/O Storage,”  
*High End Computing Revitalization Task Force Inter Agency Working Group (HECIWG) File Systems and I/O Research Workshop*, Arlington, VA, August 2007.
- “Performance Anomaly Identification for Cluster-based Parallel I/O Systems,”  
*High End Computing Revitalization Task Force Inter Agency Working Group (HECIWG) File Systems and I/O Research Workshop*, Washington, DC, August 2006.
- “I/O System Performance Debugging Using Model-driven Anomaly Characterization,”  
*Fourth USENIX FAST Conference*, San Francisco, CA, December 2005;  
*Ask Jeeves Inc.*, Piscataway, NJ, December 2005;  
*IBM Watson Research Center*, Hawthorne, NY, November 2005.
- “Parallel Sparse LU Factorization on Second-class Message Passing Platforms,”  
*Nineteenth ACM International Conference on Supercomputing*, Cambridge, MA, June 2005.
- “Challenge: Platform-dependent Adaptation of Irregular Applications,”  
 Invited presentation, *Workshop in Programming Models for HPCS Ultra-scale Applications*, Cambridge, MA, June 2005.
- “Supporting Cluster-based Network Services on Functionally Symmetric Software Architecture,”  
*SC2004 Conference*, Pittsburgh, PA, November 2004.
- “Structure Management for Scalable Overlay Service Construction,”  
*First USENIX NSDI Symposium*, San Francisco, CA, March 2004.
- “Competitive Prefetching for Disk I/O-Intensive Online Servers,”  
*IBM Almaden Research Center*, San Jose, CA, March 2004.
- “Integrated Resource Management for Cluster-based Internet Services,”  
*Fifth USENIX OSDI Symposium*, Boston, MA, December 2002.

## Graduated Ph.D. Advisees:

- Zhuan Chen, thesis title: “System Support for Data-Intensive Sensing in the Field,” completed in 2016.  
 Current placement: *Microsoft*, Redmond, WA.
- Konstantinos Menychtas, thesis title: “Fair, Protected OS-level Scheduling for Fast Computational Accelerators,” completed in 2015. Supervised jointly with Prof. Michael L. Scott. Current placement: *Google*, Mountain View, CA.
- Amal Fahad, thesis title: “System- and Application- Level Techniques for Limited Resources Environments,” completed in 2014. Current placement: *Microsoft*, Redmond, WA.
- Stan Park, thesis title: “Managing and Exploiting Flash-based Storage for Data-intensive Systems,” completed in 2014. Current placement: *HP Labs*, Palo Alto, CA.
- Xin Li, thesis title: “Toward a Better Understanding of the Properties and Impact of Memory Hardware Errors,” completed in 2012. Supervised jointly with Prof. Michael C. Huang. Current placement: *Google*, Mountain View, CA.

Christopher Stewart, thesis title: “Performance Modeling and System Management for Internet Services,” completed in 2008. Current placement: tenured faculty appointment at the *Ohio State University*.

Chuanpeng Li, thesis title: “Operating System Enhancements for Data-Intensive Server Systems,” completed in 2008. Current placement: *Google*, New York, NY.

Ming Zhong, thesis title: “Exploiting Object Usage Nonuniformity in Distributed Data Management,” completed in 2007. Supervised jointly with Prof. Joel Seiferas. Current placement: *Google*, Mountain View, CA.

## Courses Taught:

Computer Organization (Spring 2015, Spring 2012).

Computer Networks (Fall 2014, Fall 2013, Spring 2013, Fall 2011, Fall 2010, Fall 2009, Fall 2007, Fall 2006, Fall 2005, Fall 2004, Fall 2003, Fall 2002).

Parallel and Distributed Systems (Spring 2014, Spring 2011, Spring 2010).

Big Data Computer Systems (Fall 2013).

Seminar on Advanced Operating Systems (Spring 2008).

Operating Systems (Spring 2007, Spring 2006, Spring 2005, Spring 2004).

Seminar on Scalable Internet Services & Systems (Spring 2003).

The student overall course rating, averaged over all my courses up to Spring 2015, is 4.34 (out of the maximum of 5.0). Among the 14 courses for which the average ratings for comparable courses in the school are available, my overall course rating has exceeded the school average for all these courses.