Yiying Zhang

Updated Dec 2017

Assistant Professor
School of Electrical and Computer Engineering
Purdue University

yiying@purdue.edu
https://engineering.purdue.edu/~yiying
(765) 494-5916

ACADEMIC EXPERIENCE

Postdoctoral Scholar — University of California, San Diego — September 2013 - 2015
Advisor: Prof. Steven Swanson

Ph.D. in Computer Science — University of Wisconsin - Madison — August 2013
Advisors: Prof. Remzi H. Arpaci-Dusseau and Prof. Andrea C. Arpaci-Dusseau
Dissertation: De-indirection for Flash-based Solid State Drives

M.S. in Computer Engineering — University of Florida — December 2006

B.S. in Computer Science — Fudan University — June 2005

APPOINTMENTS

Purdue University, Assistant Professor — 2015-
University of California, San Diego, Postdoctoral Scholar — 2013-2015
University of Wisconsin-Madison, Research Assistant — 2007-2013
University of Florida, Research Assistant — 2005-2006

INDUSTRY EXPERIENCE

Research Intern, **NetApp Advanced Technology Group** — Sunnyvale, CA, Summer 2012
Mentors: Dr. Gokul Soundararajan and Dr. Mark W. Storer

Research Intern, **Microsoft Research** — Mountain View, CA, Summer 2010
Mentor: Dr. Vijayan Prabhakaran

Software Engineer Intern, **Microsoft,** Business Intelligence Group — Shanghai, China, Summer 2008

Software Design Engineer, **Optym** — Gainesville, FL, January - August 2007

PUBLICATIONS

Yizhou Shan, Yutong Huang, Yilun Chen, **Yiying Zhang**, "Lego: a New Distributed, Decomposed OS for Datacenter Resource Disaggregation," *in preparation*

Shin-Yeh Tsai, **Yiying Zhang**, "Mitsume: a Smart Data Store on Dumb Remote Memory," *in preparation*

Ke Liu, Shin-Yeh Tsai, **Yiying Zhang**, "TAP: a New Datacenter Network Layer for Approximate Applications," *in preparation*

Xiaoyu Liu, Linhai Song, **Yiying Zhang**, "Understanding Thread Communication and Synchronization in Go," *in preparation*

Linhai Song, Ce Zhang, **Yiying Zhang**, Heqing Huang and Wu Zhou, "Understanding Real-World Malware and Anti-Virus Engines," *in preparation*

Liwei Guo, **Yiying Zhang**, Xiaozhu Lin, "Let the Cloud Watch Over Your IoT File Systems," *in*

preparation

Shin-Yeh Tsai, **Yiying Zhang**, "LITE Kernel RDMA Support for Datacenter Applications," *Proceedings of the 26th ACM Symposium on Operating Systems Principles (**SOSP '17**)* (Acceptance Rate: 17%)

Yizhou Shan, Shin-Yeh Tsai, **Yiying Zhang**, "Distributed Shared Persistent Memory," *Proceedings of the ACM Symposium on Cloud Computing 2017 (**SoCC '17**)* (Acceptance Rate: 23%)

Yiying Zhang, Yizhou Shan, Sumukh Hallymysore, "Disaggregated Operating System," *17th International Workshop on High Performance Transaction Systems (**HPTS '17**)*

Linhai Song, Heqing Huang, Wu Zhou, Wenfei Wu, **Yiying Zhang**, "Learning from Big Malwares," *Proceedings of the 6th ACM Asia-Pacific Workshop on Systems (**APSys '16**)*

Yiying Zhang, Jian Yang, Amirsaman Memaripour, Steven Swanson, "Mojim: A Reliable and Highly-Available Non-Volatile Memory System," *Proceedings of the 20th International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS '15**)* (Acceptance Rate: 17%)

Yiying Zhang, Steven Swanson, "A Study of Application Performance with Non-Volatile Main Memory," *Proceedings of the 31st IEEE Conference on Massive Data Storage (**MSST '15**)* (Acceptance Rate: 30%)

Yiying Zhang, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Removing the Costs and Retaining the Benefits of Flash-Based SSD Virtualization with FSDV," *Proceedings of the 31st IEEE Conference on Massive Data Storage (**MSST '15**)* (Acceptance Rate: 30%)

Yiying Zhang, Gokul Soundararajan, Mark W. Storer, Lakshmi N. Bairavasundaram, Sethuraman Subbiah, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Warming up Storage-Level Caches with Bonfire," *Proceedings of the 11th Conference on File and Storage Technologies (**FAST '13**)* (Acceptance Rate: 19%)

Mohit Saxena, **Yiying Zhang**, Michael M. Swift, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Getting Real: Lessons in Transitioning Research Simulations into Hardware Systems," *Proceedings of the 11th Conference on File and Storage Technologies (**FAST '13**)* (Acceptance Rate: 19%)

Yiying Zhang, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Warped Mirrors for Flash," *Proceedings of the 29th IEEE Conference on Massive Data Storage (**MSST '13**)* (Acceptance Rate: 13%)

Yiying Zhang, Leo Prasath Arulraj, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "De-indirection for Flash-based SSDs with Nameless Writes," *Proceedings of the 10th Conference on File and Storage Technologies (**FAST '12**)* (Acceptance Rate: 19%)

Mohit Saxena,Michael M. Swift, **Yiying Zhang**, "FlashTier: a Lightweight, Consistent and Reliable Storage Cache," *Proceedings of the 7th European Conference on Computer Systems (**EuroSys '12**)* (Acceptance Rate: 15%)

Hyeoncheol Kim, **Yiying Zhang**, Yong-Seok Heo, Heung-Bum Oh, Su-Shing Chen, "Specificity Rule Discovery in HIV-1 Protease Cleavage Site Analysis," ***Computational Biology and Chemistry** 32(1): 72-79 (2008)* (Impact Factor: 1.37)

Hyeoncheol Kim, Tae-Sun Yoon, **Yiying Zhang**, Anupam Dikshit, Su-Shing Chen, "Predictability of Rules in HIV-1 Protease Cleavage Site Analysis," *Proceedings of the 2006 International Conference on Computational Science (**ICCS '06**)* (Acceptance Rate: 35%)

POSTERS AND TECHNICAL REPORTS

Yizhou Shan, Sumukh Hallymysore, Yutong Huang, Yilun Chen, **Yiying Zhang** "Disaggregated Operating System" *Poster at the ACM Symposium on Cloud Computing 2017 (**SoCC '17**)*

Shin-Yeh Tsai, Linzhe Li, Yiying Zhang "Rockies: A Network System for Future Data Center Racks" *WIP and Poster at the 14th USENIX Conference on File and Storage Technologies (**FAST '16**)*

Yiying Zhang, Jian Yang, Amirsaman Memaripour, Steven Swanson, "Mojim: A Reliable and Highly-Available Non-Volatile Memory System," *Poster at the 11th USENIX Symposium on Operating Systems Design and Implementation (**OSDI '14**)*

Yiying Zhang, Vijayan Prabhakaran, "Duplication Aware Disk Array," *Microsoft Technical Report (MSR-TR-2012-127)*

Yiying Zhang, Vijayan Prabhakaran, "DADA: Duplication Aware Disk Array," *Poster at the 9th Conference on File and Storage Technologies (**FAST '11**)*

Yiying Zhang, Leo Prasath Arulraj, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, "Porting File System Structures to Nameless Writes," *Poster at the 9th USENIX Symposium on Operating Systems Design and Implementation (**OSDI '10**)*

PATENTS

"Software Platform for Building Data Center Software Applications," Yiying Zhang, Shin-Yeh Tsai, Under Filing

"Duplicate-Aware Disk Arrays," Vijayan Prabhakaran, Yiying Zhang, US Patent 20120226936 A1

"System and Method for an Efficient Cache Warm-up," Lakshmi N. Bairavasundaram, Gokul Soundararajan, Mark W. Storer, Yiying Zhang, US Patent WO2014100253 A1

RESEARCH GRANTS AND CONTRACTS RECEIVED

Principal Investigator, NSF CNS-1719215, CSR: Small: Distributed Shared Persistent Memory, $404,572.

Principal Investigator, Samsung Co. Ltd., Hardware donation of 22x PM953 SSDs and 2x PM1725s (high-end) SSDs.

Principal Investigator, Purdue FY1617 PRF Research Grant, $29,130.

RESEARCH EXPERIENCE

Datacenter Resource Disaggregation

New Operating System for Disaggregated Resources 2016-

Proposed a new OS architecture called *splitkernel* to manage disaggregated hardware resources in future datacenters. The basic idea of splitkernel is to decompose traditional OSes into independent, stateless *micro-OS services* and to deploy them on disaggregated hardware devices. These services (together with the devices they manage) can be heterogeneous and can be added, restarted, or reconfigured dynamically without affecting the rest of splitkernel.

Designed and led the development of *Lego*, a splitkernel and the *first* OS built (from scratch) for datacenter resource disaggregation. Lego is a distributed, loosely-coupled, fault-tolerant OS that appears to applications as a set of distributed non-cache-coherent servers. Lego uses coarse-grained, global resource management to allocate, schedule, and coordinate across components, and it uses a combination of process checkpointing and memory replication to handle component failure.

Disaggregated Processor and Memory Architecture 2017-

Designed a set of new processor and memory architectures for resource disaggregation. We cleanly separate processor functionalities from those of memory, in that processors have no local memory and see only virtual memory addresses, while memory components handle all memory addressing, mapping, and protection. Processors thus use virtually-indexed and virtually-tagged caches at all levels; page tables and TLBs (or other types of memory mapping) are placed at memory components. To meet datacenter applications' performance needs, we propose a software-managed stacked-DRAM-based cache together with a small victim DRAM cache as the last-level cache for processors.

Remote and Distributed Memory Systems

Smart Data Store on Dumb Remote Memory 2017-

This project explores the idea of building data stores on network-attached remote memory devices (e.g., NVMe over fabrics), which can offer great flexibility and resource utilization. We propose to treat remote memory as "dumb" devices that do not have processing power and to access them via pure one-sided RDMA operations. To manage this data store, we propose to use a separate control plane running at distributed set of client nodes that have processing power and can communicate with each other using two-sided operations.

Distributed Shared Persistent Memory 2015-2016

Proposed *Distributed Shared Persistent Memory (DSPM)*, a new framework for using persistent memories in distributed datacenter environments. DSPM provides a new abstraction that allows applications to both perform traditional memory load and store instructions and to name, share, and persist their data. Designed and led the development of *Hotpot*, the first DSPM system and one of the first distributed persistent memory systems, in the Linux kernel. Hotpot provides low-latency, transparent memory accesses, data persistence, data reliability, and high availability.

Datacenter Network Systems

Network Design for Datacenter Approximate Applications 2017-

Designed and guided the implementation of a new network layer for datacenter approximate applications. This network layer builds on lossy datacenter network environments and uses the combination of three techniques to achieve high-throughput performance and low network utilization, while guaranteeing application approximation requirements. These techniques include selective packet dropping, adaptive approximate-aware routing, and approximate flow control.

Kernel-Level RDMA Indirection Layer for Datacenter Applications 2016-2017

Identified three fundamental issues of using current RDMA architecture for datacenter applications: RDMA lacks a flexible, high-level abstraction; its performance does not scale; and it does not provide resource sharing or flexible protection. Proposed a ground-up approach that goes the opposite direction from RDMA's design philosophy of kernel bypassing: to add kernel back. We built *LITE*, a *Local Indirection TiEr* for RDMA, in the Linux kernel. LITE virtualizes native RDMA into a flexible, high-level, easy-to-use abstraction and allows applications to safely share resources. Despite the widely-held belief that kernel bypassing is essential to RDMA's low-latency performance, LITE demonstrates that using a kernel-level indirection can achieve both flexibility *and* low-latency, scalable performance at the same time.

InfiniBand Network for Future Racks 2015-2016

Designed and guided the implementation of a new InfiniBand-based network layer for future datacenter racks with disaggregated resources, including a new topology, routing algorithm, congestion control, and data placement policies. This new network layer closely cooperates with the disaggregated operating system and it achieves low latency, high bandwidth, quality of service, low monetary cost, and fault-tolerance.

Modern Datacenter Application Programming Model

Container Application Study on Thread and Process Communication 2017-

Designed and led the study of container applications (e.g., applications on Docker Hub) on how they model parallelism and how their threads and processes communicate within a machine and across machines.

Go and Other Emerging Programming Language Concurrency Study 2017-

Collaborated with another faculty on the study of how Go and other modern programming languages handle concurrency. Our study focuses on thread/process communication methods, concurrency bugs caused by new programming models, and performance comparison and optimization of different languages.

Secure IoT File Systems

Building Secure IoT File Systems with Cloud Verification 2017-

Collaborated with another faculty on the design of secure IoT file systems. Our overall idea is to extract an IoT file systems data path and protect it within an trusted execution environment on the device. We further run a trusted, metadata-only replica of the same file system in the cloud. At run time, IoT apps forward their file I/O operations to both the local and the cloud replicas for execution, and only trust the outcome if two replicas agree on the resultant data operations.

Understanding Malwares in the Wild

Big Data Study on Online Malware Repositories 2016-

Participated in the design of a big data study of malware files in open malware repositories provided by online malware analysis services like VirusTotal. Our study focus not only on the analysis of malware, but also anti-virus engines and user community behaviors. We also designed malware prediction models that rely only on "metadata" and no original files.

Non-Volatile Main Memory (NVMM)

Reliable and Highly-Available NVMM 2014

Designed and developed Mojim, a system that provides reliability and availability to NVMM, while preserving NVMM's good performance. Mojim achieves these goals with highly-optimized RDMA-based replication protocols, software, and networking stacks, plus a two-tier architecture that provides flexible modes with different reliability, availability, consistency, and monetary costs. Surprisingly, Mojim provides replicated NVMM with only 29% to 73% the average latency of the un-replicated NVMM and 0.5 to 3.5× the bandwidth. It is also 3.4 to 4× faster than existing replication schemes.

Performance Study with NVMM 2014

Conducted a study of application performance with NVMM using a hardware NVMM emulator. Designed and developed an optimization of the data persistence process using selective cache flushing, which significantly improves performance (up to 240×) for applications that require strict durability and consistency guarantees over large regions of memory.

Removing Excess Indirection

File System De-Virtualizer 2013

Designed and implemented the File System De-Virtualizer (FSDV), a system that dynamically removes storage device indirection by changing file system pointers from virtual to physical addresses.

De-indirection with Nameless Writes 2010-2012

Designed and implemented a new interface, Nameless Writes, to remove the need for indirection in flash-based SSDs by storing physical addresses returned by the device in the file system metadata. Ported the Linux ext3 file system to use nameless writes and implemented an emulated nameless-writing device. Prototyped the nameless writes interfaces and FTL on SATA and the OpenSSD hardware platform. Nameless writes reduce FTL mapping table size by 14-50 × and improve random write performance by 20 × compared to a typical traditional SSD.

Storage-Level Caches

Storage-level Cache Warmup 2012

Designed and implemented Bonfire, a system for accelerating storage-level cache warmup, after analyzing a set of data-center traces that provided insights into heuristics of cache warmup. Bonfire speeds up the cache warmup time by 14% to 100% and has 44% to 228% more server load reduction compared to traditional cache warmup.

Solid-State Cache 2011
Developed an emulator for Solid-State Cache, a flash device designed for caching.

Redundant Arrays

Duplicate-aware Disk Arrays 2010
Designed and implemented Duplicate-Aware Disk Arrays (DADA), which keep track of block duplication and use it to improve the reliability and availability of storage arrays. DADA reduces disk scrubbing and recovery time by 17% to 26%.

Correlated Failure of Flash-based Arrays 2009-2010
Designed, simulated, and evaluated Warped Mirror, a system that separates correlated device failures in flash-based arrays in a predictable way.

Machine Learning and Bioinformatics

Rule Discovery in Biological Data 2005-2006
Investigated machine learning techniques including decision trees, neural networks, and support vector machines to learn the cleavage properties of HIV-1 virus protein sequences and other properties of DNA and RNA sequences.

MENTORING AND ADVISING EXPERIENCE

Ph.D. students at Purdue: Shin-Yeh Tsai, Yizhou Shan, Yilun Chen
Visiting scholar at Purdue: Ke Liu
Master students at Purdue: Xiaoyu Liu
Undergraduate students at Purdue: Yutong Huang, Junjie Wang
Master's completed at Purdue: Sumukh Hallymysore (08/2017), Linzhe Li (12/2016), Nan Xiang (05/2016), Tim Wingender (exchange student from 08/2015 to 05/2016)
Ph.D. students at UCSD: Jian Yang, Amirsaman Memaripour, Jian Xu, Meenakshi Sundaram Bhaskaran
Visiting students and scholar at UCSD: Jiaxin Ou (Tsinghua), Matias Bjorling (IT University of Copenhagen), Kosuke Suzuki (Fujitsu)
Master student at UCSD: Pavan Kumar Pavagada Nagaraja

TEACHING

ECE695 Modern Datacenter Systems (Fall 2016, Fall 2017)
ECE469 Operating Systems Engineering (Spring 2016, Spring 2017)
ECE565 Computer Architecture (Fall 2015)

SERVICES AND OTHER ACTIVITIES

Program Committee:
2018: OSDI '18, USENIX ATC '18
2017: ASPLOS '18, SoCC 2017, HotStorage '17
2016: ASPLOS '17 (external), Micro '16 (external)
2015: FAST '16, SoCC '15, MSST '15, Micro '15 (external)
2014: HotStorage '14, IEEE Cluster '14

Journal Review:
ACM Transactions on Storage (TOS): 2017, 2014, 2013
ACM Transactions on Embedded Computing Systems (TECS): 2016
ACM Transactions on Computer Systems (TOCS):2014
IEEE Transactions on Computers (TOC): 2014, 2012
IEEE Transactions on VLSI Systems (TVLSI): 2014
IEEE Transactions on Parallel and Distributed Systems (TPDS): 2014

Diversity Activities:
Attended Diversity Workshop (2017, 2010, 2009) and lead penal discussions in 2017

the Grace Hopper Celebration of Women in Computing Conference (2014)

TALKS

"Remote and Distributed Memory in the Age of Modern Datacenters", Databricks, San Francisco, CA, Dec 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", Yale, New Haven, CT, Dec 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", Columbia, New York, NY, Dec 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", UC-Irvine, Irvine, CA, Nov 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", UC-Davis, Davis, CA, Nov 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", Amazon, East Palo Alto, CA, Oct 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", NetApp, Sunnyvale, CA, Oct 2017
"Disaggregated Operating System", HPTS '17, Pacific Grove, CA, Oct 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", VMware Research, Palo Alto, CA, July 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", UCSB, Santa Barbara, CA, July 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", USC, Los Angeles, CA, July 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", Tsinghua University, Beijing, China, May 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", Chinese Academy of Science, Beijing, China, May 2017
"Remote and Distributed Memory in the Age of Modern Datacenters", SJTU, Shanghai, China, May 2017
"Towards Future Data Center Racks", Samsung, San Jose, CA, Feb 2016
"Rethinking Storage Vertically", HP, Palo Alto, CA, June 2015
"Rethinking Storage Vertically", NetApp, Sunnyvale, CA, June 2015
"Rethinking Storage Vertically", EMC, Santa Clara, CA, June 2015
"Rethinking Storage Vertically", Samsung, Milpitas, CA, June 2015
"A Study of Application Performance with Non-Volatile Main Memory,", MSST '15, Santa Clara, CA, June 2015
"Removing the Costs and Retaining the Benefits of Flash-Based SSD Virtualization with FSDV,", MSST '15, Santa Clara, CA, June 2015
"System Design for Emerging Storage Technologies", Microsoft Research Asia, Beijing, China, Dec 2014
"System Design for Emerging Storage Technologies", Tsinghua University, Beijing, China, Dec 2014
"System Design for Emerging Storage Technologies", Chinese Academy of Sciences, Beijing, China, Dec 2014
"Replicating Non-Volatile Main Memory", Center for Magnetic Recording Research, UCSD, San Diego, CA, Oct 2014
"Replicating Non-Volatile Main Memory", Center for Networked Systems, UCSD, San Diego, CA, Oct 2014
"Lock-Free, Resilient Data Structures for Fast I/O Accesses", Center for Networked Systems, UCSD, San Diego, CA, Oct 2014
"Non-Volatile Main Memory Systems", Intel, Hillsboro, PO, May 2014
"Virtualization and Non-Volatile Main Memory", Center for Magnetic Recording Research, UCSD, San Diego, CA, May 2014
"Virtualization and Non-Volatile Main Memory", Center for Networked Systems, UCSD, San Diego, CA, May 2014
"A Reliable and Highly-Available Non-Volatile Memory System", University of Chicago, Chicago, IL, May 2014

"A Reliable and Highly-Available Non-Volatile Memory System", University of Wisconsin-Madison, Madison, WI, May 2014
"De-indirection in Computer Systems", HPTS '13, Pacific Grove, CA, Sept 2013
"De-virtualization in Storage Systems", UCSD, San Diego, CA, May 2013
"De-virtualization in Storage Systems", Microsoft Research, Redmond, WA, April 2013
"Warped Mirrors for Flash", MSST '13, Long Beach, CA, May 2013
"Warming up Storage-Level Caches with Bonfire", FAST '13, San Jose, CA, Feb. 2013
"De-virtualization for Flash-based SSDs", WISDoM Workshop, Madison, WI, Nov. 2012
"Warming up Storage-level Caches with Bonfire", NetApp, Sunnyvale, CA, Aug. 2012
"De-indirection for Flash-based SSDs with Nameless Writes", Google, Madison, WI, Mar. 2012
"De-indirection for Flash-based SSDs with Nameless Writes", Microsoft Research, Mountain View, CA, Feb. 2012
"De-indirection for Flash-based SSDs with Nameless Writes", FAST '12, San Jose, CA, Feb. 2012
"Duplicate-aware Disk Arrays", Microsoft Research, Mountain View, CA, Aug. 2010

REFERENCES

Available upon request