
INTERESTS	distributed systems, database systems, storage systems
EMPLOYMENT	<p>Google</p> <ul style="list-style-type: none"> Software Engineer, F1 Query 2019 to present
EDUCATION	<p>University of Wisconsin-Madison, Madison, WI</p> <ul style="list-style-type: none"> Ph.D., Computer Science, 2013 to 2019 <ul style="list-style-type: none"> Advisors: Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau Dissertation: Beyond Storage Interfaces: Finding and Exploiting Unwritten Contracts In Storage Devices M.S., Computer Science, 2013 to 2015 <p>Illinois Institute of Technology, Chicago, IL</p> <ul style="list-style-type: none"> Attended, Department of Computer Science, 2010 to 2013 <p>Hunan University, Hunan, China</p> <ul style="list-style-type: none"> M.E., Computer Science, School of Computer and Communication 2008-2010 B.E., Information Security, School of Computer and Communication 2004-2008
SELECTED PUBLICATIONS	<p>Jun He, Kan Wu, Sudarsun Kannan, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, “Read As Needed: Building WiSER, a Flash-Optimized Search Engine”. in Proceedings of the 18th USENIX Conference on File and Storage Technologies (FAST ’20), Santa Clara, CA, Feb 2020. FAST</p> <p>Jun He, Sudarsun Kannan, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, “The Unwritten Contract of Solid State Drives”. in Proceedings of the 2017 EuroSys Conference (EuroSys ’17), Belgrade, Serbia, Apr 2017. EuroSys</p> <p>Jun He, Duy Nguyen, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, “Reducing File System Tail Latencies with Chopper”. in Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST ’15), Santa Clara, CA, Feb 2015. FAST</p> <p>Jun He, John Bent, Aaron Torres, Gary Grider, Garth Gibson, Carlos Maltzahn, and Xian-He Sun, “I/O Acceleration with Pattern Detection”. in Proceeding of the 22nd International ACM Symposium on High Performance Distributed Computing, New York City, NY, June 2013. HPDC</p>
CONFERENCE PUBLICATIONS	<p>Jun He, Sudarsun Kannan, Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau, “The Unwritten Contract of Solid State Drives” (extended abstract). the 9th Annual Non-Volatile Memories Workshop (NVMW ’18), San Diego, CA, Mar 2018. NVMW</p> <p>Yanlong Yin, Jibing Li, Jun He, Xian-He Sun, and Rajeev Thakur, “Pattern-Direct and Layout-Aware Replication Scheme for Parallel I/O Systems”. in Proceeding of the 27th IEEE International Parallel & Distributed Processing Symposium, May 2013. IPDPS</p> <p>Jun He, John Bent, Aaron Torres, Gary Grider, Garth Gibson, Carlos Maltzahn, and Xian-He Sun, “Discovering Structure in Unstructured I/O”. In Proceeding of 7th Parallel Data Storage Workshop, held in conjunction with SuperComputing’12, Salt Lake City, UT, Nov. 2012. PDSW</p> <p>Jun He, Xian-He Sun and Rajeev Thakur, “KNOWAC: I/O Prefetch via Accumulated Knowledge”. in the Proceeding of IEEE International Conference on Cluster Computing, Beijing, China, 2012. CLUSTER</p>

Huaiming Song, Hui Jin, Jun He, Xian-He Sun and Rajeev Thakur, “A Server-Level Adaptive Data Layout Strategy for Parallel File Systems”. In the Proc. of 2012 International Workshop on High Performance Data Intensive Computing, in conjunction with IEEE IPDPS 2012. **HPDIC**

Jun He, Huaiming Song, Xian-He Sun, Yanlong Yin, and Rajeev Thakur. “Pattern-aware File Reorganization in MPI-IO”. In the 6th Parallel Data Storage Workshop, in conjunction with ACM/IEEE SuperComputing 2011, Seattle, WA, Nov. 2011. **PDSW**

Jun He, Jim Kowalkowski, Marc Paterno, Don Holmgren, James Simone, and Xian-He Sun. “Layout-aware Scientific Computing - A Case Study using MILC”. In the Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, in conjunction with ACM/IEEE SuperComputing 2011, Seattle, WA, Nov. 2011. **Scala**

JOURNAL
PUBLICATIONS

Jun He, Jim Kowalkowski, Marc Paterno, Don Holmgren, James Simone, and Xian-He Sun. “Layout-aware Scientific Computing: A Case Study using the MILC Code”. In Journal of Computational Science, 2013.

PATENTS

United States Patent 10558618 (Metadata compression)

ACTIVITIES

Program Committee of Massive Storage Systems and Technology (MSST’2019, ’2020)

Reviewer of IEEE Transactions on Parallel and Distributed Systems

Reviewer of Software: Practice and Experience

Reviewer of ACM Transactions on Storage

Reviewer of IEEE Transactions on Computers

External Reviewer of OSDI’16

External Reviewer of FAST’16, ’18

External Reviewer of the 8th IEEE International Conference on Networking, Architecture, and Storage (NAS 2013), Xi’an, ShanXi, China, July 2013.

External Reviewer of the 22nd International ACM Symposium on High Performance Distributed Computing (HPDC’13), New York City, NY, June 2013.

External Reviewer of *The International Workshop on Data-Intensive Scalable Computing Systems (DISCS)*, in conjunction with the 2012 ACM/IEEE Supercomputing Conference (SC’12), Salt Lake City, UT, Nov 2012.

External Reviewer of the 21st International ACM Symposium on High Performance Distributed Computing (HPDC’12), 2012.

COLLABORATIONS/
INTERNSHIPS

Microsoft Jim Gray Systems Lab, Madison, Wisconsin September 2016 to 2019

New Mexico Consortium

Los Alamos National Laboratory, Los Alamos, New Mexico

Parallel Log-structured File System(PLFS) May 2012 to May 2013

- Mentors: John Bent (EMC), Aaron Torres, Gary Grider, Michael Lang (LANL)

Fermi National Accelerator Laboratory, Batavia, Illinois

Parallel Scientific Application Optimization May 2011 to August 2011

- Mentors: Jim Kowalkowski, Marc Paterno

SELECTED
AWARDS

National Natural Science Foundation of China

- Champion, The Future Challenge: Intelligent Vehicles and Beyond (National Self-Driving Vehicle Contest), 2009
- Development of the laser radar guiding system, and Information & Control Units for the self-driving vehicle.

Intel, Ministry of Education (China), and Ministry of Industry and Information Technology (China)

- First-class Prize (top 13/160), National Undergraduate Electronic Design Contest - Embedded System Design Invitational Contest (Intel Cup), 2008

TEACHING
EXPERIENCE

Hunan University, Hunan, China

Teaching Assistant

- Computer System & Architecture, 2009

TALKS

Read As Needed: Building WiSER, a Flash-Optimized Search Engine - the 18th USENIX Conference on File and Storage Technologies (FAST '20), Santa Clara, CA, Feb 2020.

The Unwritten Contract of Solid State Drives - the 9th Annual Non-Volatile Memories Workshop (NVMW '18), San Diego, CA, March 2018.

The Unwritten Contract of Solid State Drives, Vault - Linux Storage and Filesystems Conference, Cambridge, MA, March 2017.

The Unwritten Contract of Solid State Drives, EuroSys'2017.

Reducing File System Tail Latencies with Chopper, Vault - Linux Storage and Filesystems Conference, Boston, MA, March 2015.

Reducing File System Tail Latencies with Chopper, FAST'15.

I/O Acceleration with Pattern Detection - HPDC'2013.

Layout-aware Scientific Computing - A Case Study using MILC, EMC, Beijing, Sept 2012.

Metadata Reduction in an Exascale File System, Los Alamos National Lab, Los Alamos, Aug 2012.

Pattern-aware File Reorganization in MPI-IO, Los Alamos National Lab, Los Alamos, May 2012.

Layout-aware Scientific Computing - A Case Study using MILC, Hunan University, Changsha, China, Sept 2012.