

Mai Zheng

Data Storage Lab @Iowa State

Dept. of Electrical & Computer Engineering
Iowa State University, Ames, IA 50011
☎ (515) 294-6285
✉ mai@iastate.edu
🌐 www.ece.iastate.edu/~mai

RESEARCH INTERESTS

Data Storage Systems; Data Integrity & Security; Data Intensive Computing

EMPLOYMENT

- 2023–present **Associate Professor**, *Department of Electrical and Computer Engineering*, Iowa State University, Ames, IA
- 2020–2023 **Harpole-Pentair Assistant Professor**, *Department of Electrical and Computer Engineering*, Iowa State University, Ames, IA
- 2018–2020 **Assistant Professor**, *Department of Electrical and Computer Engineering*, Iowa State University, Ames, IA
- 2015–2018 **Assistant Professor**, *Department of Computer Science*, New Mexico State University, Las Cruces, NM
- 2012, 2013 **Summer Research Intern**, *HP Labs*, Palo Alto, CA.
[HP Labs Mentors: Dr. Joseph Tucek, Dr. Mark Lillibridge, Dr. Kimberly Keeton]

EDUCATION

- 2009–2015 **Ph.D. in Computer Science**, *The Ohio State University*, Columbus, OH.
[Advisor: Dr. Feng Qin]
- 2006–2009 **M.S. in Electronic Science & Technology**, *University of Science and Technology of China*, Hefei, China. [Advisor: Dr. Li Guo].
- 2002–2006 **B.S. in Electronic Science & Technology**, *Qingdao University*, Qingdao, China.

HONORS & AWARDS

- 2022 **Samsung Global Research Outreach (GRO) Award**, *Samsung Advanced Institute of Technology*
- 2022 **HotStorage Best Paper Nominee**, *ACM Hot Topics in Storage & File Systems Committee*
- 2022 **Cybersecurity Faculty Fellow**, *Iowa State University*
- 2021 **HotStorage Best Paper Nominee**, *ACM Hot Topics in Storage & File Systems Committee*
- 2020 **Harpole-Pentair Professorship**, *Iowa State University*
- 2020 **NSF CAREER Award**, *National Science Foundation*
- 2019 **IDEMA Research Award**, *Western Digital Research & International Disk Drive Equipment and Materials Association*
- 2018 **FAST Best Paper Nominee**, *USENIX File & Storage Technologies Program Committee*
- 2018 **IDEMA Research Award**, *Western Digital Research & International Disk Drive Equipment and Materials Association*
- 2016 **NSF CRII Award**, *National Science Foundation*
- 2009 **University Fellowship**, *The Ohio State University*
- 2008 **Huawei Fellowship**, *University of Science and Technology of China*

SELECTED RESEARCH PRODUCTS

- ARA **arawireless.org**, *A multi-city data infrastructure for wireless & rural America*.
Publications: **MobiCom'23-Demo, WiNTECH'22, WiNTECH'21**.
Sponsors: National Science Foundation (NSF), National Institute of Food and Agriculture (NIFA), US Ignite, & PAWR Industry Consortium.
- ConfD **github**, *A framework for analyzing configuration dependencies and detecting configuration-related issues in storage systems*.
Publications: **FAST'23, NVMW'23, HotStorage'22**.
Sponsor: National Science Foundation (NSF).
- PROV-IO+ **github**, *A provenance framework for scientific data on HPC systems*.
Publications: **TPDS'24, HUG'23, HPDC'22**.
Sponsors: National Science Foundation (NSF), Lawrence Berkeley National Lab (LBNL), Samsung Advanced Institute of Technology (SAIT).
- BugBench^k **github**, *Characterization of system bugs for improving reliability, security, reproducibility, etc.*
Publications: **TOS'23, FAST'23, NAS'22, TBench'21, SYSTOR'21**.
Sponsors: National Science Foundation (NSF), Western Digital Research (WD).
- PFault **github**, *A framework for analyzing the reliability of parallel file systems*.
Publications: **TOS'22, HotStorage'21, PDSW'20, MSST'19, ICS'18, PDSW'16**.
Sponsor: National Science Foundation (NSF).
- rFSCK **github**, *A framework for analyzing the fault resilience of file system checkers & a transactional library for enhancing reliability*.
Publications: **TOS'18, FAST'18, HotStorage'17**.
Sponsors: National Science Foundation (NSF), Western Digital Research (WD).

GOVERNMENT^G & INDUSTRY^I RESEARCH FUNDING

[SUMMARY: About \$10.2M in total as PI/Co-PI; Mai Zheng's share: \$2.85M]

- ^{G,I}2023 US Ignite: **ARA Launchpad: Open Wireless Platforms, Data and Co-Prototyping for Integrative Research and Community Building**. Total \$784K, Mai Zheng (Co-PI) 15%. 2023 - 2026.
- ^G2023 CloudBank: Commerical Cloud Research Credits for NSF CAREER Award. Mai Zheng (Sole PI): \$15K. 2023 - 2024.
- ^I2022 Samsung Semiconductor, Inc.: **Samsung Global Research Outreach Award**. Mai Zheng (Sole PI): \$150K. 2022 - 2023.
- ^G2022 CloudBank: Commerical Cloud Research Credits for NSF CAREER Award. Mai Zheng (Sole PI): \$15K. 2022 - 2023.
- ^I2022 Google LLC: Google Cloud Platform (GCP) Research Credits. Mai Zheng (Sole PI), \$5K, 2022.
- ^{G,I}2021 National Science Foundation (NSF), United States Department of Agriculture's (USDA's) National Institute of Food and Agriculture (NIFA), US Ignite & PAWR Industry Consortium: **ARA: Wireless Living Lab for Smart and Connected Rural Communities**. Total \$7.1M, Mai Zheng (Co-PI) 16.7%. 2021 - 2026.
- ^G2020 National Science Foundation (NSF): **CAREER: Towards Full-Stack Crash Consistency**. Mai Zheng (Sole PI). \$544K (include \$32K REU Supplement). 2020 - 2025.
- ^{G,I}2020 National Science Foundation (NSF) & US Ignite: **Planning: ARA: Wireless Living Lab for Smart and Connected Rural Communities**. Total \$300K, Mai Zheng (Co-PI) 16.7%. 2020.
- ^I2019 International Disk Drive Equipment and Materials Association (IDEMA) & Western Digital Research: University Research Award. Mai Zheng (Sole PI), \$25K. 2019.
- ^G2019 National Science Foundation (NSF): **Collaborative Research: A Parallel Graph-Based Paradigm for HPC Parallel File System Checkers**. Total \$600K, Mai Zheng (ISU PI) 50%. 2019 - 2023.

- ¹2018 International Disk Drive Equipment and Materials Association (IDEMA) & Western Digital Research: University Research Gift. PI: Mai Zheng. \$35K, 2018.
- ^G2017 National Science Foundation (NSF): **Collaborative Research: Uncovering Vulnerabilities in Parallel File Systems for Reliable High Performance Computing**. Total \$466K, Mai Zheng (NMSU/ISU PI) 50%. 2017 - 2021.
- ^G2016 National Science Foundation (NSF): **CRII: Towards Pinpointing the Root Causes of Failures in Flash-based Storage Systems**. Mai Zheng (Sole PI): \$174K. 2016 - 2020.

PUBLICATIONS

Peer-Reviewed Journals

- TPDS'24 [12] Mai Zheng, Vignesh T. Ravi, Feng Qin, and Gagan Agrawal, **PROV-IO+: A Cross-Platform Provenance Framework for Scientific Data on HPC Systems**. *IEEE Transactions on Parallel and Distributed Systems*, 2024.
- TOS'23 [11] Om R. Gatla, Duo Zhang, Wei Xu, and Mai Zheng, **Understanding Persistent-Memory Related Issues in the Linux Kernel**. *ACM Transactions on Storage*, 2023.
- TC'22 [10] Jiang Zhou, Yong Chen, Mai Zheng, Weiping Wang. **Data Distribution for Heterogeneous Storage Systems**. *IEEE Transactions on Computers*, 2022.
- TOS'22 [9] Runzhou Han, Om R. Gatla, Mai Zheng, Jinrui Cao, Di Zhang, Dong Dai, Yong Chen, and Jonathan Cook, **A Study of Failure Recovery and Logging of High-Performance Parallel File Systems**. *ACM Transactions on Storage*, 2022.
- CPE'21 [8] Chunxue Zuo, Fang Wang, Mai Zheng, Yuchong Hu, and Dan Feng. **Ensuring High Reliability and Performance with Low Space Overhead for Deduplicated and Delta-Compressed Storage Systems**. *Concurrency and Computation: Practice and Experience*, 2021.
- TBench'21 [7] Duo Zhang, Mai Zheng, **Benchmarking for Observability: The Case of Diagnosing Storage Failures**. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 2021.
- TOS'18 [6] Om R. Gatla, Mai Zheng, Muhammad Hameed, Viacheslav Dubeyko, Adam Manzanares, Filip Blagojevic, Cyril Guyot, and Robert Mateescu, **Towards Robust File System Checkers**. *ACM Transactions on Storage*, 2018. *Fast-tracked!*
- TOCS'16 [5] Mai Zheng, Joseph Tucek, Feng Qin, Mark Lillibridge, Bill W Zhao, and Elizabeth S Yang, **Reliability Analysis of SSDs under Power Fault**. *ACM Transactions on Computer Systems*, 2016.
- JKSYS'16 [4] Yongsheng Hao, Lina Wang, and Mai Zheng, **An Adaptive Algorithm for Scheduling Parallel Jobs in Meteorological Cloud**. *Journal of Knowledge-based Systems*, 2016.
- TPDS'14 [3] Mai Zheng, Vignesh T. Ravi, Feng Qin, and Gagan Agrawal, **GMRace: Detecting Data Races in GPU Programs via A Low-Overhead Scheme**. *IEEE Transactions on Parallel and Distributed Systems*, 2014.
- JCEA'09 [2] Jian Ji, Li Guo, Mai Zheng, and Lu Gao, **A Programmable Pixel Shader for Mobile Devices**. *Journal of Computer Engineering and Applications*, 2009.
- JCEA'09 [1] Bingqin Wang, Li Guo, and Mai Zheng, **A Sub-Pixel Image Registration Algorithm in Panoramic Image Mosaics**. *Journal of Computer Engineering and Applications*, 2008.

Peer-Reviewed Conference/Workshop Papers [3 Best Paper Award Nominations]

[X% Acceptance Rate; Y% Best Paper Nomination Rate]

- ASPLOS'23 [21%] [37] Benjamin Carver, Runzhou Han, Jingyuan Zhang, Mai Zheng, and Yue Cheng **λ FS: Elastically Scaling Distributed File System Metadata Service using Serverless Functions**. *Proceedings of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2023.

- MobiCom'23 -Demo [36] Taimoor U. Islam, Joshua O. Boateng, Guouing Zu, Mukaram Shahid, Md Nadim, Wei Xu, Tianyi Zhang, Salil Reddy, Xun Li, Ataberk Atalar, Yung-Fu Chen, Sarath Babu, Hongwei Zhang, Daji Qiao, Mai Zheng Yong Guan, Ozdal Boyraz, Anish Arora, Mohamed Selim, Myra B. Cohen. **ARA PAWR: Wireless Living Lab for Smart and Connected Rural Communities**. *Proceedings of the 29th ACM Annual International Conference on Mobile Computing and Networking*, 2023.
- IPDPS'23 [26%] [35] Di Zhang, Chris Egersdoerfer, Tabassum Mahmud, Mai Zheng, and Dong Dai, **Drill: Log-based Anomaly Detection for Large-scale Storage Systems Using Source Code Analysis**. *Proceedings of 37th IEEE International Parallel and Distributed Processing Symposium*, 2023.
- IPDPS'23 [26%] [34] Saisha Kamat, Abdullah Al Raqibul Islam, Mai Zheng, and Dong Dai **FaultyRank: A Graph-based Parallel File System Checker**. *Proceedings of 37th IEEE International Parallel and Distributed Processing Symposium*, 2023.
- NVMW'23 [33] Tabassum Mahmud, Om R. Gatla, Duo Zhang, Carson Love, Ryan Bumann, and Mai Zheng, **Analyzing Configuration Dependencies of DAX File Systems**. *The 14th Annual Non-Volatile Memories Workshop*, 2023.
- FAST'23 [23%] [32] Tabassum Mahmud, Om R. Gatla, Duo Zhang, Carson Love, Ryan Bumann, and Mai Zheng, **CONF: Analyzing Configuration Dependencies of File Systems for Fun and Profit**. *Proceedings of the 21st USENIX Conference on File and Storage Technologies*, 2023.
- WiNTECH'22 [31] Mukaram Shahid, Sarath Babu, Hongwei Zhang, Daji Qiao, Yong Guan, Joshua Ofori Boateng, Taimoor UI Islam, Guoying Zu, Ahmed Kamal, and Mai Zheng. **Wireless Guard for Trustworthy Spectrum Management**. *Proceedings of 16th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization at MobiCom*. 2022.
- NAS'22 [30] Duo Zhang, Tabassum Mahmud, Om R. Gatla, Runzhou Han, Yong Chen, and Mai Zheng, **On the Reproducibility of Bugs in File-System Aware Storage Applications**. *Proceedings of the 16th IEEE International Conference on Networking, Architecture, and Storage*. 2022.
- HotStorage'22 [6.4%] [29] Tabassum Mahmud, Duo Zhang, Om R. Gatla, and Mai Zheng, **Understanding Configuration Dependencies of File Systems**. *Proceedings of the 14th ACM Workshop on Hot Topics in Storage and File Systems*. 2022. *Best Paper Nominee!*
- HPDC'22 [19%] [28] Runzhou Han, Suren Byna, Houjun Tang, Bin Dong, and Mai Zheng, **PROV-IO: An I/O-Centric Provenance Framework for Scientific Data on HPC Systems**. *Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing*, 2022.
- SYSTOR'21 [30%] [27] Duo Zhang, Om R. Gatla, Wei Xu, and Mai Zheng, **A Study of Persistent Memory Bugs in the Linux Kernel**. *Proceedings of the 14th ACM International Systems and Storage Conference*, 2021.
- HotStorage'21 [7.5%] [26] Di Zhang, Dong Dai, Runzhou Han, and Mai Zheng, **SentiLog: Anomaly Detecting on Parallel File Systems via Log-based Sentiment Analysis**. *Proceedings of the 13th ACM Workshop on Hot Topics in Storage and File Systems*. 2021. *Best Paper Nominee!*
- WiNTECH'21 [25] Hongwei Zhang, Yong Guan, Ahmed Kamal, Daji Qiao, Mai Zheng, et.al., **ARA: A Wireless Living Lab Vision for Smart and Connected Rural Communities**. *Proceedings of 15th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & CHaracterization at ACM MobiCom*. 2021.
- PDSW'20 [24] Runzhou Han, Duo Zhang, and Mai Zheng, **Fingerprinting the Checker Policies of Parallel File Systems**. *Proceedings of the 5th ACM/IEEE International Parallel Data Systems Workshop at ACM/IEEE Supercomputing Conference*. 2020.
- ATC'19 [19.9%] [23] Erci Xu, Mai Zheng, Feng Qin, Yikang Xu, and Jiesheng Wu. **Lessons and Actions: What We Learned from 10K SSD-Related Storage System Failures**. *Proceedings of the USENIX Annual Technical Conference*. 2019.
- MSST'19 [22] Dong Dai, Om R. Gatla, and Mai Zheng. **A Performance Study of Lustre File System Checker: Bottlenecks and Potentials**. *Proceedings of the 35th International Conference on Massive Storage Systems and Technology*, 2019.

- PDSW'18 [21] Erci Xu, Mai Zheng, Feng Qin, Yikang Xu, and Jiasheng Wu, **Understanding SSD Reliability in Large-Scale Cloud Systems**. *Proceedings of the 3rd ACM/IEEE Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems at ACM/IEEE Supercomputing Conference*, 2018.
- ICS'18 [18.6%] [20] Jinrui Cao, Om R. Gatla, Mai Zheng, Dong Dai, Vidya Eswarappa, Yan Mu, and Yong Chen, **PFault: A General Framework for Analyzing the Reliability of High-Performance Parallel File Systems**. *Proceedings of the 32nd ACM International Conference on Supercomputing*, 2018.
- HPEC'18 [19] Yehia Arafa, Atanu Barai, Mai Zheng, and Abdel-Hameed Badawy, **Fault Tolerance Performance Evaluation of Large Scale Distributed Storage Systems: HDFS and Ceph Case Study**. *Proceedings of the 22nd IEEE High Performance Extreme Computing Conference*, 2018.
- FAST'18 [16.5%;3.6%] [18] Om R. Gatla, Muhammad Hameed, Mai Zheng, Viacheslav Dubeyko, Adam Manzanares, Filip Blagojevic, Cyril Guyot, and Robert Mateescu, **Towards Robust File System Checkers**. *Proceedings of the 16th USENIX Conference on File and Storage Technologies*, 2018. *Best Paper Nominee!*
- IGSC'17 [17] Li Li, Bruce Beitman, Mai Zheng, Xiaorui Wang and Feng Qin, **eDelta: Pinpointing Energy Deviations in Smartphone Apps via Comparative Trace Analysis**. *Proceedings of the 8th International Green and Sustainable Computing Conference*, 2017.
- IGSC'17 [16] Li Li, Yunhao Bai, Xiaorui Wang, Mai Zheng and Feng Qin, **Selective Checkpointing for Minimizing Recovery Energy and Efforts of Smartphone Apps**. *Proceedings of the 8th International Green and Sustainable Computing Conference*, 2017.
- HotStorage'17 [15] Om R. Gatla and Mai Zheng, **Understanding the Fault Resilience of File System Checkers**. *Proceedings of the 9th USENIX Workshop on Hot Topics in Storage and File Systems*, 2017.
- ICNC'17 [14] Yiliang Shi, Danny V. Murillo, Simeng Wang, Jinrui Cao, and Mai Zheng, **A Command-Level Study of Linux Kernel Bugs**. *Proceedings of International Conference on Computing, Networking and Communications*, 2017.
- PDSW'16 [13] Jinrui Cao, Simeng Wang, Dong Dai, Mai Zheng, and Yong Chen, **A Generic Framework for Testing Parallel File Systems**. *Proceedings of the ACM/IEEE Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems at ACM/IEEE Supercomputing Conference*, 2016.
- NAS'16 [12] Simeng Wang, Jinrui Cao, Danny V. Murillo, Yiliang Shi, and Mai Zheng, **Emulating Realistic Flash Device Errors with High Fidelity**. *Proceedings of the 11th IEEE International Conference on Networking, Architecture, and Storage*, 2016.
- CSCI'15 [11] Junwen Lu, Yongsheng Hao, Lina Wang, Mai Zheng, **Towards efficient service composition in multi-cloud environment**. *Proceedings of the 2015 International Conference on Computational Science and Computational Intelligence*, 2015.
- OSDI'14 [18.4%] [10] Mai Zheng, Joseph Tucek, Dachuan Huang, Feng Qin, Mark Lillibridge, Elizabeth S. Yang, Bill W. Zhao, and Shashank Singh, **Torturing Databases for Fun and Profit**. *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation*, 2014.
- FAST'13 [18.8%] [9] Mai Zheng, Joseph Tucek, Feng Qin, and Mark Lillibridge, **Understanding the Robustness of SSDs under Power Fault**. *Proceedings of the 11th USENIX Conference on File and Storage Technologies*, 2013.
- MASCOTS'13 [8] Dachuan Huang, Xuechen Zhang, Wei Shi, Mai Zheng, Song Jiang, and Feng Qin, **LiU: Hiding Disk Access Latency for HPC Applications with a New SSD-Enabled Data Layout**. *Proceedings of the 21st IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, 2013.
- HiPC'13 [7] Mai Zheng, Vignesh T. Ravi, Wenjing Ma, Feng Qin, and Gagan Agrawal, **GMPProf: A Low-Overhead, Fine-Grained Profiling Approach for GPU Programs**. *Proceedings of the 19th IEEE International Conference on High Performance Computing*, 2012.

- WCRE'12 [6] Dawei Qi, William Sumner, Feng Qin, Mai Zheng, Xiangyu Zhang and Abhik Roychoudhury, **Modeling Software Execution Environment**. *Proceedings of the 19th Working Conference on Reverse Engineering*, 2012.
- ASPLOS'11 [21%] [5] Qi Gao, Wenbin Zhang, Zhezhe Chen, Mai Zheng, and Feng Qin, **2ndStrike: Towards Manifesting Hidden Concurrency Typestate Bugs**. *Proceedings of the 16th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2011.
- PPoPP'11 [15.7%] [4] Mai Zheng, Vignesh T. Ravi, Feng Qin, and Gagan Agrawal, **GRace: A Low-Overhead Mechanism for Detecting Data Races in GPU Programs**. *Proceedings of the 16th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, 2011.
- ICSP'08 [3] Mai Zheng, Jian Ji, Li Guo, and Junzhu Zhu, **A Phase-Fitting Method for Sub-pixel Displacement Measurements Using Digital Speckle Images**. *Proceedings of the 9th IEEE International Conference on Signal Processing*, 2008.
- ICCSIT'08 [2] Mai Zheng, Antai Guo, Wei Zhong, and Li Guo, **Image Stitching of Scenes with Large Misregistration**. *Proceedings of the 1st IEEE International Conference on Computer Science and Information Technology*, 2008.
- ISVC'08 [1] Mai Zheng, Xiaolin Chen, and Li Guo, **Stitching Video from Webcams**. *Proceedings of the 4th International Symposium on Visual Computing*, 2008.
- Peer-Reviewed Position Papers/Work-in-Progress Reports/Posters
- HUG'23 [14] Runzhou Han, Mai Zheng, Suren Byna, Houjun Tang, Bin Dong, Dong Dai, Yong Chen, Dongkyun Kim, Joseph Hassoun, David Thorsley, Matthew Wolf, **PROV-IO+: A Cross-Platform Provenance Framework for Scientific Data on HPC Systems**. *HDF5 User Group (HUG) Meeting*, 2023.
- FAST'23 [13] Duo Zhang, Om Rameshwar Gatla, Abdullah Al Raqibul Islam, Dong Dai, and Mai Zheng. **On the Scalability of Testing the Crash Consistency of PM Systems**. *Work-in-Progress (WiP) & Poster Sessions, 21th USENIX Conference on File and Storage Technologies*, 2023.
- DOEData'22 [12] Mai Zheng, Runzhou Han, Haojun Tang, **Towards Unified FAIR Metadata Services for Scientific Data**. *Department of Energy (DOE) ASCR Workshop on the Management and Storage of Scientific Data*, 2022.
- FAST'22 [11] Tabassum Mahmud, Mai Zheng. **Understanding Configuration Issues in Storage Systems**. *Work-in-Progress (WiP) & Poster Sessions, 20th USENIX Conference on File and Storage Technologies*, 2022.
- FAST'22 [10] Runzhou Han, Suren Byna, Mai Zheng. **Towards A Practical Provenance Framework for Scientific Data on HPC Systems**. *Work-in-Progress (WiP) & Poster Sessions, The 20th USENIX Conference on File and Storage Technologies*, 2022.
- HotStorage'20 [9] Duo Zhang, Om R. Gatla, Runzhou Han, Mai Zheng, **Position: On Failure Diagnosis of the Storage Stack**. *Position Paper & Poster Sessions, The 12th USENIX Workshop on Hot Topics in Storage and File Systems*, 2020.
- FAST'20 [8] Duo Zhang, Chander B. Gupta, Mai Zheng, Adam Manzanares, Filip Blagojevic, and Cyril Guyot, **A Cross-Layer Approach for Diagnosing Storage System Failures**. *Work-in-Progress (WiP) & Poster Sessions, The 18th USENIX Conference on File and Storage Technologies*, 2020.
- FAST'19 [7] Om R. Gatla, Yealim Sung, and Mai Zheng, **Let the Device Talk**. *Work in Progress (WiP) & Poster Sessions, The 17th USENIX Conference on File and Storage Technologies*, 2019.
- FAST'19 [6] Ryan Chatier, Om Rameshwar Gatla, Mai Zheng, and Henry Duwe, **On the Recoverability of Persistent Memory Systems**. *Work in Progress (WiP) & Poster Sessions, 17th USENIX Conference on File and Storage Technologies*, 2019.
- CoDA'18 [5] Panika Valecha, Huiping Cao, Qixu Gong, Mai Zheng, Feng Yan, Xing Lin and Art Harkin, **Analysis and Prediction of Storage Error Events for High Performance Computing Systems**. *Poster Session, Department of Energy (DOE) Conference on Data Analysis*, 2018.

- FAST'17 [4] Om R. Gatla and Mai Zheng, **On Fault Resilience of File System Checkers**. *Work in Progress -wip (WiP) & Poster Sessions, The 15th USENIX Conference on File and Storage Technologies*, 2017.
- NVMW'17 [3] Simeng Wang, Jinrui Cao, Om R. Gatla, Muhammad Hameed, and Mai Zheng, **Do Not Blame -poster Devices for All Failures**. *The 8th Annual Non-Volatile Memories Workshop*, 2017.
- DOENet'16 [2] Satyajayant Misra and Mai Zheng **Rethinking Networking in a Non-volatile, Heterogeneous World**. *Department of Energy (DOE) Workshop on Network Research Problems and Challenges 2025*, 2016.
- NSFCloud'14 [1] Mai Zheng, Joseph Tucek, Feng Qin, and Mark Lillibridge, **A Reliability Analysis Framework for Cloud Storage Systems**. *National Science Foundation (NSF) Workshop on Experimental Support for Cloud Computing*, 2014.

Non-Peer-Reviewed Technical Reports

- arXiv'23 [5] Runzhou Han, Mai Zheng, Suren Byna, Houjun Tang, Bin Dong, Dong Dai, Yong Chen, Dongkyun Kim, Joseph Hassoun, David Thorsley, Matthew Wolf, **PROV-IO+: A Cross-Platform Provenance Framework for Scientific Data on HPC Systems**. *arxiv.org/abs/2308.00891*, 2023.
- arXiv'23 [4] Om Rameshwar Gatla, Duo Zhang, Wei Xu, and Mai Zheng, **Understanding Persistent-Memory Related Issues in the Linux Kernel**. *arxiv.org/abs/2307.040953*, 2023.
- Dagstuhl'22 [3] Mai Zheng, Jack Clark, and Miryung Kim, **Towards Bug-Free DBMS Ecosystems**. *Report of Dagstuhl Seminar 21442: Ensuring the Reliability and Robustness of Database Management Systems*, Volume 11, Issue 10, 2022.
- arXiv'19 [2] Viacheslav Dubeyko, Om R. Gatla, and Mai Zheng, **Nature of System Calls in CPU-centric Computing Paradigm**. *arXiv: 1903.04075v1 [cs.OS]*, 2019.
- NSF- [1] George Amvrosiadis, Ali R. Butt, Vasily Tarasov, Erez Zadok, Ming Zhao, Irfan Ahmad, Remzi Storage'18 H. Arpaci-Dusseau, Feng Chen, Yiran Chen, Yong Chen, Yue Cheng, Vijay Chidambaram, Dilma Da Silva, Angela Demke-Brown, Peter Desnoyers, Jason Flinn, Xubin He, Song Jiang, Geoff Kuenning, Min Li, Carlos Maltzahn, Ethan L. Miller, Kathryn Mohror, Raju Rangaswami, Narasimha Reddy, David Rosenthal, Ali Saman Tosun, Nisha Talagala, Peter Varman, Sudharshan Vazhkudai, Avani Waldani, Xiaodong Zhang, Yiyang Zhang, and Mai Zheng. **Data Storage Research Vision 2025**. *Technical Report on NSF Visioning Workshop on Data Storage Research*, 2018.

Talks

- ACCORD [18] **Analyzing Configuration Dependencies of File Systems for Fun and Profit**. *Host: Dr. Myra Cohen & Dr. Paul Gazzillo, The 1st Workshop on A Community for Configurability Open Research and Development (ACCORD)*, Zoom, March, 2024.
- Samsung [17] **Scrutinizing Latent Defects in HPC Storage Systems for RAS Optimizations**. *Host: Dr. Yong Chen, Samsung*, Zoom, March, 2024.
- Tsinghua [16] **A Bug's Life vs. A Byte's Life: Systems Approaches to Enhance Data Integrity and Understandability**. *Host: Dr. Youyou Lu, Tsinghua University*, January, 2024.
- Samsung [15] **ECFault: A Systematic Framework for Analyzing Erasure Coded Storage Systems**. *Host: Dr. Yong Chen, Samsung*, Zoom, July, 2023.
- Alibaba [14] **A Bug's Life vs. A Byte's Life: Systems Approaches to Enhance Data Integrity and Understandability**. *Host: Dr. Qing Li, Alibaba*, Virtual, Feb, 2023.
- ByteDance [13] **A Bug's Life vs. A Byte's Life: Systems Approaches to Enhance Data Integrity and Understandability**. *Host: Dr. Tongping Liu, ByteDance & UMass Amherst*, Virtual, Nov, 2022.
- Wisconsin [12] **A Bug's Life vs. A Byte's Life: Systems Approaches to Enhance Data Integrity and Understandability**. *Host: Dr. Remzi Arpaci-Dusseau, Department of Computer Sciences, University of Wisconsin - Madison*, Madison, WI, Oct, 2022.

- Samsung [11] **Analyzing the Blast Radius Problem of SSD Failures on HPC Storage Systems via Holistic Testbed Framework.** *Host: Dr. Yong Chen, Samsung, Zoom, July, 2022.*
- IDEMA [10] **Heterogenous Memory Management: Design Tradeoffs for Performance & Durability.** *Host: Dr. Cyril Guyot, IDEMA Advanced Storage Research Committee (ASRC) Spring Meeting, Virtual, 2020.*
- IDEMA [9] **Heterogenous Memory Management.** *Host: Dr. Cyril Guyot, IDEMA Advanced Storage Research -WD Committee (ASRC) Spring Meeting at Western Digital (WD) Research, Milpitas, CA, October 2019.*
- IDEMA- [8] **System Performance Improvements through DRAM Replacement by NVM.** *Host: Dr. Cyril Seagate Guyot, IDEMA Advanced Storage Research Committee (ASRC) Spring Meeting at Seagate, Fremont, CA, March 2019.*
- PDSW'18 [7] **Understanding SSD Reliability in Large-Scale Cloud Systems.** *The 3rd ACM/IEEE Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS) at ACM/IEEE Supercomputing (SC), Dallas, TX, 2018.*
- SDC'15 [6] **Torturing Databases for Fun and Profit.** *SNIA Storage Developer Conference (SDC), Santa Clara, CA, 2015.*
- NSF-Cloud'14 [5] **A Reliability Analysis Framework for Cloud Storage Systems.** *National Science Foundation (NSF) Workshop on Experimental Support for Cloud Computing, Arlington, VA, 2014.*
- OSDI'14 [4] **Torturing Databases for Fun and Profit.** *The 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Broomfield, CO, 2014.*
- FAST'13 [3] **Understanding the Robustness of SSDs under Power Fault.** *The 11th USENIX Conference on File and Storage Technologies (FAST), San Jose, CA, 2013.*
- HP Labs [2] **Understanding the Robustness of SSDs under Power Fault.** *Host: Dr. Joseph Tucek, HP Labs, Palo Alto, CA, 2013.*
- PPoPP'11 [1] **GRace: A Low Overhead Mechanism for Detecting Data Races in GPU Programs.** *The 16th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming (PPoPP), San Antonio, TX, 2011.*

TEACHING

- ISU CprE308 Operating Systems (FA'18, FA'19, FA'20, FA'21, SP'22, SP'23, FA'23, SP'24)
- ISU CprE563 Advanced Data Storage Systems (SP'20, SP'21, SP'22, SP'23, SP'24)
- ISU CprE588 Embedded Computer Systems (SP'19)
- NMSU CS579 Reliable Storage Systems (FA'17)
- NMSU CS579 Modern Storage Systems: Flash, Cloud, & Beyond (SP'16)
- NMSU CS474 Operating Systems I (FA'15, FA'16)
- NMSU CS574 Operating Systems II (SP'17, SP'18)
- NMSU CS473 Computer Architecture I (SP'18)
- NMSU CS573 Computer Architecture II (FA'17)
- NMSU CS491 Parallel Programming (FA'16)
- OSU CSE4251 The UNIX Programming Environment (FA'14, SP'15)
- USTC EE4303 Embedded Systems (SP'08)

STUDENT MENTORING

[SUMMARY: 7 Ph.D. (3 Graduated); 16 Masters (15 Graduated); 17 Undergraduates]

Current Students

- Ph.D. Runzhou Han (passed preliminary exam)
- Tabassum Mahmud (passed preliminary exam)

Wei Xu
 Chao Shi
 Master Varun S. Girimaji
 Jack Kelly
 Undergraduate Jacob Lyons
 John Huaracha
 Hang Thang
 Hussein Baderddeine
 Zeren Yang (University of Wisconsin-Madison)

Former Students

Ph.D. [3] Duo Zhang, Ph.D.(ISU) 2023
 [2] Om R. Gatla, Ph.D.(ISU) 2022 [→ SK Hynix]
 [1] Jinrui Cao, Ph.D.(NMSU) 2020 [→ Tenure-Track Assistant Professor@The State University of New York (SUNY) at Plattsburgh]

Master [15] Manish Manepalli, M.S.(ISU) 2023 [→ Ph.D. Program@ISU]
 [14] Jacob Betsworth, M.S.(ISU) 2023
 [13] Kajal Kattige, M.S.(ISU) 2022 [→ ASML]
 [12] Haolun Ping, M.S.(ISU) 2022 [→ Fortinet]
 [11] Nidhi Milind Dalvi, M.S.(ISU) 2022
 [10] Siyuan Lu, M.S.(ISU) 2021 [→ ASML]
 [9] Prakhar Bansal, M.S.(ISU) 2020 [→ AeroSpike]
 [8] Satya Prakash, M.S.(ISU) 2020 [→ Qualcomm]
 [7] Chander B. Gupta, M.S.(ISU) 2020 [→ Intel]
 [6] Ryan Chartier, M.S.(NMSU) 2019 [→ Allscripts]
 [5] Chase Gilbert, M.S.(NMSU) 2018 [→ Honeywell Aerospace]
 [4] Yuan Xu, M.S.(NMSU) 2018 [→ Microsoft]
 [3] Muhammad Hameed, M.S.(NMSU) 2018 [→ Self-Employed]
 [2] Om R. Gatla, M.S.(NMSU) 2017 [→ Ph.D. Program @NMSU&ISU]
 [1] Simeng Wang, M.S.(NMSU) 2017 [→ Meituan]

Undergraduate [16] Ella Cook (ISU), Fall 2023
 [15] Varun Advani (ISU), Summer 2023
 [14] Michael Less (ISU), Summer 2023
 [13] Varun S. Girimaji (PES University), Summer 2023, [→ MS Program @Iowa State Univ.]
 [12] Carson Love (ISU, NSF REU), 2022 – 2023 [→ PwC]
 [11] Ryan Bumann (ISU, NSF REU), 2022 – 2023 [→ Vermeer]
 [10] Katie Wolf (ISU), 2021
 [9] Aaron Goff (ISU), 2021 [→ ITS @ISU]
 [8] Joshua Kalyanapu (ISU), 2021 [→ Ph.D. Program @NCSU]
 [7] Jesrik Gomez (ISU), 2020 [→ Western Digital]
 [6] Zhenyu Zhao (ISU), 2019 [→ Google]
 [5] Yealim Sung (ISU), 2019 [→ TTA]
 [4] Chelsea Deane (George Washington Univ., NSF REU), 2017

- [3] Kristopher Chesney (California State Univ., NSF REU), 2017
- [2] Danny V. Murillo (Mesa Community College, NSF REU), 2016
- [1] Yiliang Shi (U. of Utah, NSF REU), 2016 [→ Ph.D. Program @Columbia Univ.]

Visitor Chunxue Zuo, Visiting PhD Student (HUST), 2019 – 2020
 Xining Yuan, M.S. Intern (Northwestern University), 2023
 Philip Ma, Undergraduate Intern (Grinnell College), 2021

PROFESSIONAL SERVICE

Book: Reviewer & Question Writer

MOS Modern Operating Systems (5th edition) by Andrew S. Tanenbaum and Herbert Bos, Publisher: Pearson (ISBN-13: 9780137618880), 2022.

Journal: Editorial Board

TBench BenchCouncil Transactions on Benchmarks, Standards and Evaluations, since 2022.

Journal: Reviewer

- TOPS [36] ACM Transactions on Privacy and Security, Mar 2024.
- TSC [35] IEEE Transactions on Services Computing, Nov 2023.
- JPDC [34] Journal of Parallel and Distributed Computing, Nov 2023.
- TR [33] IEEE Transactions on Reliability, Nov 2023.
- TOS [32] ACM Transactions on Storage, Aug 2023.
- TR [31] IEEE Transactions on Reliability, Jun 2023.
- TPDS [30] IEEE Transactions on Parallel and Distributed Systems, Jun 2023.
- TR [29] IEEE Transactions on Reliability, 2022.
- TCC [28] IEEE Transactions on Cloud Computing, 2021.
- TPDS [27] IEEE Transactions on Parallel and Distributed Systems, 2021.
- TCC [26] IEEE Transactions on Cloud Computing, 2020.
- TR [25] IEEE Transactions on Reliability, 2020.
- TON [24] ACM/IEEE Transactions on Networking, 2020.
- TMPECS [23] ACM Transactions on Modeling and Performance Eval. of Comp. Systems, 2020.
- TCC [22] IEEE Transactions on Cloud Computing, 2020.
- COMMAG [21] IEEE Communications Magazine, Aug 2020.
- COMMAG [20] IEEE Communications Magazine, May 2020.
- TR [19] IEEE Transactions on Reliability, Dec 2019.
- TPDS [18] IEEE Transactions on Parallel and Distributed Systems, Aug 2019.
- TR [17] IEEE Transactions on Reliability, May 2019.
- TPDS [16] IEEE Transactions on Parallel and Distributed Systems, Feb 2019.
- JPDC [15] Journal of Parallel and Distributed Computing, 2019.
- TR [14] IEEE Transactions on Reliability, 2018.
- TPDS [13] IEEE Transactions on Parallel and Distributed Systems, Nov 2018.
- TPDS [12] IEEE Transactions on Parallel and Distributed Systems, Sept 2018.
- TPDS [11] IEEE Transactions on Parallel and Distributed Systems, Aug 2018.
- JPDC [10] Journal of Parallel and Distributed Computing, 2018.
- TPDS [9] IEEE Transactions on Parallel and Distributed Systems, 2017.
- IET [8] IET Computers & Digital Techniques, 2017.

- TOS [7] ACM Transactions on Storage, 2016.
- TPDS [6] IEEE Transactions on Parallel and Distributed Systems, 2016.
- TECS [5] ACM Transactions on Embedded Computer Systems, 2015.
- TC [4] IEEE Transactions on Computers, 2015.
- TPDS [3] IEEE Transactions on Parallel and Distributed Systems, 2015.
- TC [2] IEEE Transactions on Computers, 2014.
- TPDS [1] IEEE Transactions on Parallel and Distributed Systems, 2013.

Conference: Organization

- MSST [7] Posters Co-Chair, 38th International Conference on Massive Storage Systems and Technology, June 2024.
- ARA [6] Session Chair, ARA Public Launch Program, Sept 2023.
- HPDC [5] Publicity Co-Chair, ACM International Symposium on High-Performance Parallel and Distributed Computing, Jun 2023.
- IPDPS [4] Session Chair, IEEE International Parallel & Distributed Processing Symposium, May 2023.
- NVMW [3] Session Chair, Annual Non-Volatile Memories Workshop, Mar 2023
- HPDC [2] Proceedings Chair, ACM International Symposium on High-Performance Parallel and Distributed Computing, 2018.
- DAAC [1] Proceedings Chair, International Workshop on Data-center Automation, Analytics, and Control, 2018.

Conference: Technical Program Committee

- SC [33] Research/ACM SRC Posters Committee, International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), Nov 2024
- SYSTOR [32] The ACM International Systems and Storage Conference, Jun 2024.
- FAST [31] Artifact Evaluation Committee, USENIX Conference on File and Storage Technologies (FAST), Feb 2024
- REX-IO [30] Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads, held in conjunction with IEEE Cluster, Oct 2023.
- ICDCS [29] IEEE International Conference on Distributed Computing Systems, Jul 2023.
- IPDPS [28] IEEE International Parallel & Distributed Processing Symposium, May 2023.
- NAS [27] IEEE International Conference on Networking, Architecture, and Storage, 2022.
- REX-IO [26] Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads, held in conjunction with IEEE Cluster, 2022.
- ICDCS [25] IEEE International Conference on Distributed Computing Systems, 2022.
- IPDPS [24] IEEE International Parallel & Distributed Processing Symposium, 2022.
- REUNS [23] National Workshop for REU Research in Networking and Systems, 2022.
- ICDCS [22] IEEE International Conference on Distributed Computing Systems, 2021.
- HiPC [21] IEEE International Conference on High Performance Computing, Data, and Analytics, 2021.
- Bench [20] BenchCouncil International Symposium on Benchmarking, Measuring and Optimizing, 2021.
- REUNS [19] National Workshop for REU Research in Networking and Systems, 2021.
- CCGrid [18] ACM/IEEE International Symposium on Cluster, Cloud and Internet Computing, 2020.
- IPDPS [17] IEEE International Parallel & Distributed Processing Symposium, 2019.
- DAAC [16] International Workshop on Data-center Automation, Analytics, and Control, held in conjunction with ACM/IEEE Supercomputing (SC), 2019.
- REUNS [15] National Workshop for REU Research in Networking and Systems, 2019.

- DAAC [14] International Workshop on Data-center Automation, Analytics, and Control, held in conjunction with ACM/IEEE Supercomputing (SC), 2018,
- HiPC [13] IEEE International Conference on High Performance Computing, Data, and Analytics, 2018.
- MLCS [12] Workshop on Machine Learning for Computing Systems, 2018.
- FiCloud [11] IEEE International Conference on Future Internet of Things and Cloud, 2018.
- FMEC [10] IEEE International Conference on Fog and Mobile Edge Computing, 2018.
- REUNS [9] National Workshop for REU Research in Networking and Systems, 2018.
- HiPC [8] IEEE International Conference on High Performance Computing, Data, and Analytics, 2017.
- IWQoS [7] ACM/IEEE International Symposium on Quality of Service, 2017.
- UCC [6] ACM/IEEE International Conference on Utility and Cloud Computing, 2017.
- REUNS [5] National Workshop for REU Research in Networking and Systems, 2017.
- CCWC [4] IEEE Annual Computing and Communication Workshop and Conference, 2017.
- IWQoS [3] ACM/IEEE International Symposium on Quality of Service, 2016.
- REUNS [2] National Workshop for REU Research in Networking and Systems, 2016.
- DIBS [1] Workshop on Data-Centric Infrastructure for Big Data Science, held in conjunction with IEEE BigData, 2015.

Grant Reviewer/Panelist

- NSF [5] National Science Foundation, 2020-b.
- NSF [4] National Science Foundation, 2020-a.
- NSF [3] National Science Foundation, 2018.
- NSF [2] National Science Foundation, 2016.
- KSEF [1] Kentucky Science & Engineering Foundation, 2016.

Others

- USENIX Campus Representative of USENIX Association at ISU, 2018 – Present.
- USENIX Campus Representative of USENIX Association at NMSU, 2016 – 2018.
Member of ACM, IEEE, USENIX.