Duo Zhang
San Jose, CA 95134

Email : duozhang@ucmerced.edu
Mobile : 315-949-8633
Web: https://duozhangisu.github.io/

Summary

- Sufficient knowledge and experience in data storage systems, the Linux kernel, and HPC
- Extensive hands-on experience designing, executing, and analyzing benchmarks and AI workloads
- Strong end-to-end performance troubleshooting skills across compute, network, and storage stacks
- Capability of multi-tasking and collaboration with academic and industrial experts
- Intensive coding experience in system-level programming with C/C++, Python, and Bash

Education

- **Iowa State University** — Ames, IA
  Doctor of Philosophy in Computer Engineering — *Sep. 2018 - Dec. 2023*
- **Syracuse University** — Syracuse, NY
  Master of Science in Electrical Engineering — *Aug. 2016 - May. 2018*

Certificate

AWS Certified Cloud Practitioner
Internet and Computing Core Certification(IC3)

Work Experience

- **SK Hynix**
  System Engineer (Staff) — *Dec. 2025 - Present*
  - Benchmarked and evaluated new memory and storage devices using fio- and MLPerf-style workloads across diverse AI scenarios.
  - Tracked emerging technical trends and delivered software stack design guidance for software–hardware co-designed systems and demos.
- **University of California, Merced**
  Postdoc Researcher — *July 2024 - Dec. 2025*
  - Contributed to drafting competitive research proposals, including outlining technical objectives, major methodologies, and broader impact statements for funding opportunities.
  - Mentored and guided Ph.D. students on system design, debugging, and experimental workflows, providing both technical insight and project management support.
  - Designed project frameworks and established operational routines, including task prioritization, code debugging, experiment scheduling, and milestone tracking to ensure research progress and team alignment.
- **Iowa State University**
  Graduate Research/Teaching Assistant — *Aug 2018 - Feb. 2024*
  - Contributed to research proposals by developing methodological sections and surveying existing literature to support the proposed approach and demonstrate novelty.
  - Provided technical support in system failure cases reproducing and diagnosis results to Westen Digital.
  - Delivered lectures and led lab sessions to help students understand tutorials, providing hands-on support with debugging and troubleshooting during exercises.
  - Provided reliable experiment platforms for our group including lab servers and git repos management and maintenance, and new devices deployment (e.g., Intel Optane Persistent Memory, Samsung SmartSSD).
  - Published conference and journal papers and presented research findings at conferences.
- **Sunrise Technology & Development Inc**
  Assistant/Associate/Principal System Engineer — *Aug 2014 - Jun. 2016*
  - Provided technical support and resolved system issues including partial damages in hardware and customer operation errors.
  - Managed and maintained lab instruments system with servers, and deployed git service to our clusters.

Projects

- **CXL Memory–Driven Storage Design for AI Workloads**
  - Characterized CXL memory features using existing benchmarking tools including MLPerf and Intel MLC.
  - Evaluated existing memory pool solutions, including FAMFS, DPDK memory pools, and EXT4-dax.
  - Analyzed and evaluated the NVIDIA Dynamo AI serving framework workflow, classified 3 major data access patterns, and identified KV-cache bottlenecks in the prefill–decode disaggregation scenario.
  - Designed and implemented an object-based file system with CXL co-design to improve overall performance.
- **Scalable and Resilient Modeling for Federated-Learning(FL)-Based Complex Workflows**
  - Designed and implemented a discrete-event simulator by Simgrid for federated learning, supporting three aggregation strategies: FedAvg, FedAsync, and FedCompass.
  - Extended APPFL's multi-client simulation capability using Apache Ray for scalable FL experiments.
  - Enhanced the FLOWER training capability by integrating distributed training with TRL and gRPC.
- **Characterizing Large Reasoning Model Training Pipeline**
  - Reproduced training/inference pipelines of large-scale reasoning models(e.g., DeepSeek-R1) using Megatron-DeepSpeed, vLLM and VeRL frameworks on supercomputer such as Perlmutter(Top19) and Frontier(Top2).
  - Benchmarked training and distillation pipelines, leveraged IO500 to evaluate storage performance, and analyzed Lustre performance variability under HPC workloads.
  - Identified network and system-level performance bottlenecks, including model sharding inefficiencies and checkpoint I/O stalls, using NVIDIA Nsight Systems and PyTorch FSDP tracing tools, and achieved 6× faster training by optimizing compute–communication balance with Libfabric (OFI).
- **Lossy Compression with NVMe-SSD**
  - Designed Z-SPDK, a dual-pipeline I/O framework combining lossy compression with NVMe SSDs to optimize scientific data storage and retrieval.
  - Performed microbenchmarking on various scientific datasets using compressors like SZx and ZFP to assess compression speed, I/O throughput, and decompression accuracy, revealing inefficiencies in NVMe performance with lossy compression.
  - Utilized fio to model data access patterns under lossy compression scenarios and built a benchmark-driven database to support dynamic parameter tuning.
  - Integrated compressor modules with SPDK and NVMe I/O threads to enable pipeline-level parallelism and near-zero-copy processing, achieving up to 12× higher performance for high-throughput I/O workloads.
- **Uncovering Vulnerabilities in File Systems**
  - Deployed PFSes (Lustre, BeeGFS, Ceph) on VMware and AWS, analyzed I/O patterns under ML workloads, and developed a fault-injection taxonomy to investigate 28 unexpected system and checker results.
  - Analyzed parameter dependencies of the local file system (e.g., EXT4 and XFS) by LLVM, Python and C++, and identified 78 issues by a configuration-based testing method built by Bash script and Python.
- **Virtual machine-based System Failures Detection and Diagnosis**
  - Collected SSD-related bug cases via Bash scripts, reproduced 8 kernel file system and memory corruption bugs with customized workloads, and identified faulty execution paths through flow analysis.
  - Measured and evaluated the debugging observability of FTrace and PANDA(virtualization-based tool) via reproduced system failures, and identified 4 limitations via reproduced bug cases, Bash scripts and Python.
  - Modified NVMe emulation in QEMU and Linux kernel SCSI driver to capture device commands and instructions, and align captured data by Python to identify critical buggy paths, which effectively narrow down the search space for root cause (0.06% - 6.2% of the original Ftrace method.
- **Uncovering Vulnerabilities for Persistent Memory(PM) System**

- Collected 1553 PM patches in the Linux kernel source tree by Bash scripts, and analyzed patches in detail including bug patterns, triggering conditions, and fix strategies.
- Implemented inter/intra-processes analysis framework for PM kernel drivers, and identified 29 potential issues through designed security bug detectors, static analysis (LLVM) technique and C++.
- Measured and evaluated I/O patterns of Redis through YCSB under PM system, and designed representative workloads based on the measurement.
- Built a full-stack crash consistency testing framework to emulate power outage scenario by QEMU and Python, and identified 10 unexpected system behaviors under-extracted workloads.

Selected Publications

"On Fault Tolerance of Data Storage Systems: A Holistic Perspective". Mai Zheng, Duo Zhang, and Ahmed Dajani. In *Fault Tolerance in Modern Engineering Systems*, edited by Mircea Ruba and Gabriel Chindris. IntechOpen, 2025.

"HPC-R1: Characterizing R1-like Large Reasoning Models on HPC". Adam Weingram*, Duo Zhang*, Zhonghao Chen, Hao Qi, and Xiaoyi Lu. In Proceedings of the 38th International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2025.

"Can Long-Haul RDMA Benefit Federated Learning?". Zhonghao Chen, Yuke Li, Duo Zhang, and Xiaoyi Lu. In Proceedings of the 38th International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Poster Paper, 2025.

"Can Lossy Compression Benefit NVMe-based IO?". Darren Ng, Duo Zhang, Sheng Di, Zhaorui Zhang, Guanpeng Li, and Xiaoyi Lu. In Proceedings of the 38th International Conference for High Performance Computing, Networking, Storage and Analysis (SC), Poster Paper, 2025.

"Analyzing Configuration Dependencies of File Systems". Tabassum Mahmud, Om Rameshwar Gatla, Duo Zhang, Carson Love, Ryan Bumann, Varun S. Girimaji, and Mai Zheng. ACM Transactions on Computer Systems (TOCS), 2025.

"Understanding Persistent-Memory Related Issues in the Linux Kernel". Om R. Gatla, Duo Zhang, Wei Xu, and Mai Zheng. ACM Transactions on Storage (TOS), 2023.

"Analyzing Configuration Dependencies of DAX File Systems". Tabassum Mahmud, Om R. Gatla, Duo Zhang, Carson Love, Ryan Bumann, and Mai Zheng. The 14th Annual Non-Volatile Memories Workshop (NVMW), 2023.

"ConfD: Analyzing Configuration Dependencies of File Systems for Fun and Profit". Tabassum Mahmud, Om R. Gatla, Duo Zhang, Carson Love, Ryan Bumann, and Mai Zheng. Proceedings of the 21st USENIX Conference on File and Storage Technologies (FAST), 2023.

"On the Scalability of Testing the Crash Consistency of PM Systems". Duo Zhang, Om Rameshwar Gatla, Abdullah Al Raqibul Islam, Dong Dai, and Mai Zheng. The 21th USENIX Conference on File and Storage Technologies (FAST), Work-in-Progress (WiP) & Poster Sessions, 2023.

"On the Reproducibility of Bugs in File-System Aware Storage Applications". Duo Zhang, Tabassum Mahmud, Om Rameshwar Gatla, Runzhou Han, Yong Chen and Mai Zheng. 16th International Conference on Networking, Architecture, and Storage (NAS), 2022.

"Understanding configuration dependencies of file systems". Tabassum Mahmud, Duo Zhang, Om Rameshwar Gatla, Mai Zheng. 14th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage), 2022 (Best Paper Nominee!).

"Benchmarking for observability: The case of diagnosing storage failures". Duo Zhang, Mai Zheng. BenchCouncil Transactions on Benchmarks (Bench), 2021.

"A study of persistent memory bugs in the Linux kernel". Duo Zhang, Om Rameshwar Gatla, Wei Xu, Mai Zheng. Proceedings of the 14th ACM International Conference on Systems and Storage (SYSTOR), 2021.

"Position: On Failure Diagnosis of the Storage Stack". Duo Zhang, Om R. Gatla, Runzhou Han, Mai Zheng. The 12th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage), Position Paper & Poster Sessions, 2020.

"Fingerprinting the Checker Policies of Parallel File Systems". Runzhou Han, Duo Zhang, Mai Zheng. IEEE/ACM Fifth International Parallel Data Systems Workshop (PDSW), 2020.