Chen Wang

Fernbach Postdoctoral Fellow · Lawrence Livermore National Lab

wang116@llnl.gov | Home Page | Google Scholar | Github

Research Interests

High-Performance Computing, Parallel I/O, Storage System, HPC for AI, Distributed System.

Professional Experience

Sept. 2022 - present	**Fernbach Postdoc Fellow**	Lawrence Livermore National Laboratory, US
May 2021 - Aug. 2021	**Research Intern**	Argonne National Laboratory, US
May 2020 - Aug. 2020	**Software Engineer Intern**	The HDF Group, US
May 2019 - Aug. 2019	**Research Intern**	Lawrence Livermore National Laboratory, US
Jan. 2017 - Aug. 2017	**Senior Software Engineer**	Nextan Pte Ltd, Singapore

Education

University of Illinois Urbana-Champaign — *Champaign, US*
Ph.D., Computer Science — *Aug. 2017 - Jul. 2022*
- Advisor: Dr. Marc Snir
- Dissertation: Parallel File System with Tunable Consistency

Tianjin University — *Tianjin, China*
M.S., Computer Science — *Sept. 2014 - Jan. 2017*
- Advisor: Dr. Jizhou Sun and Dr. Ce Yu
- Thesis: A CUDA-based Dynamic Programming Framework

Hainan University — *Haikou, China*
B.S., Computer Science — *Sept. 2010 - June 2014*
- Advisor: Dr. Qi Chen

Honors

2024	**R&D100 Award**	UnifyFS: Specialized File System for Supercomputers
2024	**MVAPICH User Group Meeting Travel Grant**	National Science Foundation, US
2022	**Sidney Fernbach Fellowship**	Lawrence Livermore National Laboratory, US
2021	**Conference Presentation Award**	University of Illinois Urbana-Champaign, US
2018	**Best Paper Award**	IEEE International Conference on Cluster Computing
2016	**Bronze Medal**	National Parallel Application Challenge, China
	China National Scholarship	China
	Second-Class Scholarship	Tianjin University, China
2015	**Huawei Scholarship**	Huawei, China
	First-Class Scholarship	Tianjin University, China
2014	**First-Class Scholarship**	Tianjin University, China
2013	**Second-Class Scholarship**	Hainan University, China
2011	**First-Class Scholarship**	Hainan University, China

Publications

Peer-Reviewed Papers

1. Hariharan Devarajan, Ian Lumsden, **Chen Wang**, Konstantia Georgouli, Jae-Seung Yeom, and Michela Taufer. "DLDM: Locality-aware Data Management for accelerating Deep Learning Training." IEEE/SBC 36th International Symposium on Computer Architecture and High Performance Computing. (**SBAC-PAD, 2024**). [Conference]
2. **Chen Wang**, Houjun Tang, Jean Luca Bez and Suren Byna. "Object-Centric Data Management in HPC Workflows - A Case Study". 4th Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads. (**REX-IO, 2024**). [Workshop]
3. Olga Kogiou, Hariharan Devarajan, **Chen Wang**, Weikuan Yu and Kathryn Mohror. "Understanding Adaptable Storage for Diverse Workloads". 4th Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads. (**REX-IO, 2024**). [Workshop]
4. **Chen Wang**, Kathryn Mohror, and Marc Snir. "Formal Definitions and Performance Comparison of Consistency Models for Parallel File Systems." IEEE Transactions on Parallel and Distributed Systems. (**TPDS, 2024, CORE Rank A***) [Journal]
5. **Chen Wang**, Yanfei Guo, Pavan Balaji, and Marc Snir. "Near-Lossless MPI Tracing and Proxy Application Autogeneration." IEEE Transactions on Parallel and Distributed Systems. (**TPDS, 2022, CORE Rank A***). [Journal]
6. **Chen Wang**, and Marc Snir. "Revisiting Storage Programming Models." ASCR Workshop on the Management and Storage of Scientific Data. 2022. [Workshop]
7. **Chen Wang**, Pavan Balaji, and Marc Snir. "Pilgrim: Scalable and (near) Lossless MPI Tracing." The International Conference for High Performance Computing, Networking, Storage and Analysis. (**SC, 2021, CORE Rank A**). [Conference]
8. Sushma Yellapragada, **Chen Wang**, and Marc Snir. "Verifying IO Synchronization from MPI Traces." IEEE/ACM 6th International Parallel Data Systems Workshop. (**PDSW, 2021**). [Workshop]
9. **Chen Wang**, Kathryn Mohror, and Marc Snir. "File System Semantics Requirements of HPC Applications." 30th International Symposium on High-Performance Parallel and Distributed Computing. (**HPDC, 2021, CORE Rank A**). [Conference]
10. Jinghan Sun, **Chen Wang**, Jian Huang, and Marc Snir. "Understanding and Finding Crash-Consistency Bugs in Parallel File Systems." 12th USENIX Workshop on Hot Topics in Storage and File Systems. (**HotStorage, 2020**). [Workshop]
11. **Chen Wang**, Jinghan Sun, Marc Snir, Kathryn Mohror, and Elsa Gonsiorowski. "Recorder 2.0: Efficient Parallel I/O Tracing and Analysis.", The IEEE International Workshop on High-Performance Storage. (**HPS, 2020**). [Workshop]
12. Na Bai, Shanjiang Tang, Ce Yu, Hao Fu, **Chen Wang**, and Xi Chen. "GMSA: A Data Sharing System for Multiple Sequence Alignment across Multiple Users.", Current Bioinformatics. 2019. (**CORE Rank C**). [Journal]
13. **Chen Wang**, Nikoli Dryden, Franck Cappello, and Marc Snir. "Neural Network based Silent Error Detector." IEEE International Conference on Cluster Computing. (**CLUSTER, 2018, CORE Rank B, Best Paper Award**). [Conference]
14. Xi Chen, **Chen Wang**, Shanjiang Tang, Ce Yu, and Quan Zou. "CMSA: A Heterogeneous CPU/GPU Computing System for Multiple Similar RNA/DNA Sequence Alignment", BMC Bioinformatics. 2017 (**CORE Rank A**). [Journal]
15. **Chen Wang**, Ce Yu, Shanjiang Tang, Jian Xiao, Jizhou Sun, and Xiangfei Meng. "A General and Fast Distributed System for Large-Scale Dynamic Programming Applications." Parallel Computing. 2016. (**CORE Rank B**). [Journal]
16. **Chen Wang**, Ce Yu, Jizhou Sun, and Xiangfei Meng. "DPX10: An Efficient X10 Framework for Dynamic Programming Applications", 44th International Conference on Parallel Processing. (**ICPP, 2015, CORE Rank B**). [Conference]

Seminar Report and Book Chapter

1. Fahim Chowdhury, Hariharan Devarajan, Ann Gentile, Jay Lofstead, Kathryn Mohror, Devesh Tiwari, and **Chen Wang**. "Tools: Cross-Cutting Issues." Report for Dagstuhl Seminar 21332: Understanding I/O Behavior in Scientific and Data-Intensive Computing. Section 5. Aug. 2021. [Report]
2. **Chen Wang**, Shanjiang Tang, and Ce Yu. "Parallel Dynamic Programming for Large-Scale Data Applications." Horizons in Computer Science Research, Volume 14. Chapter 4. 2017. [Chapter]

Dataset

1. **Chen Wang**, Marc Snir, and Kathryn Mohror. "High Performance Computing Application I/O Traces." In Lawrence Livermore National Laboratory (LLNL) Open Data Initiative. UC San Diego Library Digital Collections. 2020. https://doi.org/10.6075/J0Z899X4.

Posters

1. Hammad Ather, Jean Luca Bez, **Chen Wang**, Hariharan Devarajan, Suren Byna, Kathryn Mohror, Hank Childs, Allen D. Malony. "Enhancing I/O performance: Leveraging Runtime and Offline Optimization Frameworks." The International Conference for High Performance Computing, Networking, Storage and Analysis. 2024. (**SC**) [Poster]
2. Olga Kogiou, Hariharan Devarajan, **Chen Wang**, Weikuan Yu, Kathryn Mohror. "I/O Characterization of Heterogeneous Workflows." The International Conference for High Performance Computing, Networking, Storage and Analysis. 2024. (**SC**) [Poster]
3. Hariharan Devarajan, Ian Lumsden, **Chen Wang**, Konstantia Georgouli, Tom Scogland, Jae-Seung Yeom, and Michela Taufer. "DYAD: Accelerating Deep Learning Training via Inter-Node Access Coordination." DOE Data Day Workshop. 2024. [Poster]
4. Johntan Garcia and **Chen Wang**, "A Visual Representation for Data and Workflow of HPC Applications." ISC High Performance, 2024. (**ISC**) [Poster]
5. Olga Kogiou, Hariharan Devarajan, **Chen Wang**, Weikuan Yu, and Kathryn Mohror. "I/O Characterization and Performance Evaluation of Large-Scale Storage Architectures for Heterogeneous Workloads." 3th Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads. (**REX-IO, 2023**). [WIP]
6. **Chen Wang**, Kathryn Mohror, and Marc Snir. "Consistency Models: The Parallel File System's Perspective." Lawrence Livermore National Laboratory Postdoc Poster Seminar. 2023. [Poster]
7. **Chen Wang**, Kathryn Mohror, and Marc Snir. "I/O Characteristics of Scientific Applications." 30th International Symposium on High-Performance Parallel and Distributed Computing, 2021. (**HPDC**) [Poster]

In Review

1. **Chen Wang**, Zhaobin Zhu, Kathryn Mohror, Sarah Neuwirth, and Marc Snir. "VerifyIO: Verifying Adherence to Parallel I/O Consistency Semantics". Submitted to IPDPS'25. (Core Rank A Conference).
2. Jean Luca Bez, Houjun Tang, **Chen Wang**, and Suren Byna. "Data Without Borders: Seamless Cross-Facility Data Movement using Object Data Management Runtime." Submitted to IPDPS'25. (Core Rank A Conference).
3. **Chen Wang**, Izzet Yildirim, Hariharan Devarajan, Kathryn Mohror, and Marc Snir. "Recorder: Comprehensive Parallel IO Tracing and Analysis." Submitted to Journal of Supercomputing. (Core Rank B Journal).

Invited Talks

Aug. 2024	**"Revisiting MPI-IO Consistency: Potential Revisions and Their Impact on HDF5"** The HDF5 User Group Meeting.[slides]	*Chicago, US*
Nov. 2022	**"Detecting Data Races on Storage Systems using Recorder"** Analyzing Parallel I/O, Birds-of-a-Feather at SC'22.[slides]	*Dallas, US*
May. 2022	**"Near-Lossless MPI Tracing and Proxy Application Autogeneration"** ECP Community Birds-of-a-Feather. [slides]	*Virtual*
Jan. 2021	**"Pilgrim - A Lossless MPI Tracing Tool"** Argonne National Laboratory Mathematics and Computer Science Seminar. [slides]	*Virtual*
Oct. 2020	**"An I/O Study of ECP Applications"** The HDF Group Meeting. [slides]	*Virtual*
Dec. 2019	**"Neural Network based Silent Error Detector"** Tianjin University. [slides]	*Tianjin, China*

Grant Proposal Writing Experience

1. LLNL Fernbach Fellowship. “Revisiting Consistency Models for HPC I/O: Can we finally ditch POSIX?”. 2022. $157K per year for 2 years. DOE’s Office of Science and LLNL.
2. LLNL Lab Directed Research and Development (LDRD) Program Feasibility Study. “Autonomous I/O Knowledge Engine: Freeing Users from Manual I/O Optimizations.” 2024. Rejected.
3. Interlab (DOE labs) LDRD. “Accountable and Traceable AI Explainability for Science, Energy, and Security through Data-Driven Motifs”. 2024. Rejected.

Teaching Experience

Fall 2021	**Scientific Visualization**	Teaching Assistant, University of Illinois Urbana-Champaign
Fall 2016	**Parallel Computing**	Teaching Assistant, Tianjin University

Mentoring Experience

June. 2024 - present	**Hammad Ather**	Ph.D Student at University of Oregon, US Hammad is working on runtime I/O analysis and optimization. I served as Hammad’s mentor during his 2024 summer internship at LLNL.
May. 2024 - present	**Zhaobin Zhu**	Ph.D Student at Johannes Gutenberg University of Mainz, Germany Zhaobin is working on performance modeling of HPC applications. I served as Zhaobin’s mentor during his 2024 summer internship at LLNL.
Sept. 2022 - present	**Olga Kogiou**	Ph.D Student at Florida State University, US Olga is working on storage systems for complex workflows and HPC applications. I served as Olga’s co-mentor during her 2024 summer internship at LLNL.
May 2023 - May 2024	**Johnatan Garcia**	M.S. Student at University of Texas at El Paso, US I mentored Johnatan through the NNSA Minority Serving Internship Program. He was working on data flow analysis of HPC workflows.

Community Service

Program Chair	**ESSA 2025**	Workshop on Extreme-Scale Storage and Analysis.
Publicity Chair	**SSDBM 2025**	The International Conference on Scalable Scientific Data Management.
	ESSA 2024	Workshop on Extreme-Scale Storage and Analysis.
Workshop Review Committee	**ISC 2025**	ISC High Performance.
Program Committee	**IPDPS 2025**	IEEE International Parallel Distributed Processing Symposium.
	ICS 2025	International Conference on Supercomputing
	SC 2024	The International Conference for High Performance Computing, Networking, Storage, and Analysis.
	CLUSTER 2025, 2024	IEEE International Conference on Cluster Computing.
	PERMAVOST 2024	Workshop on Performance Engineering, Modelling, Analysis, and Visualization Strategy.
	ESSA 2023	Workshop on Extreme-Scale Storage and Analysis.
	PDSW 2023	International Parallel Data Systems Workshop.
	CHEOPS 2023	Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems.
	REX-IO 2023	Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads.
Reviewer Board	**2023 - present**	IEEE Transactions on Parallel and Distributed Systems (TPDS).
Reproducibility Reviewer Board	**2023 - present**	IEEE Transactions on Parallel and Distributed Systems (TPDS).
External Reviewer	**2024-present**	Transactions on Cloud Computing.
	2022 - present	The Journal of Supercomputing.
	2022 - present	ACM Transactions on Embedded Computing Systems (TECS).
Student Volunteer	**SC 2018, 2019**	The International Conference for High Performance Computing, Networking, Storage, and Analysis.