Dr. Dingwen Tao

Associate Professor
School of Informatics, Computing, and Engineering, Indiana University Bloomington
Phone: (951) 236-2342 • Email: ditao@iu.edu • Web: https://www.dingwentao.com/

Education

University of California, Riverside 09/2013–06/2018
Ph.D. in Computer Science Riverside, CA, USA

University of Science and Technology of China 09/2009–06/2013
B.S. in Information and Computing Sciences Hefei, Anhui, China

Employment Experience

Indiana University, Luddy School of Informatics, Computing, and Engineering 08/2022–Present
Associate Professor Bloomington, IN, USA

Washington State University, School of EECS 08/2020–08/2022
Assistant Professor Pullman, WA, USA

University of Washington, Center of Responsibility in AI Systems & Experiences 04/2021–Present
Affiliated Member Seattle, WA, USA

University of Alabama, Department of Computer Science 08/2018–08/2020
Assistant Professor Tuscaloosa, AL, USA

Brookhaven National Laboratory, Computational Science Initiative 01/2018–06/2018
Research Assistant Upton, NY, USA

Argonne National Laboratory, Mathematics and Computer Science 06/2016–12/2017
Research Assistant Lemont, IL, USA

Pacific Northwest National Laboratory, HPC Group 06/2015–09/2015
Graduate Research Intern Richland, WA, USA

Research Interests

- High-performance computing (HPC), parallel & distributed systems
- Scientific data management, analytics & visualization
- Large-scale machine learning & deep learning
- Heterogeneous computing & reconfigurable computing
- Fault tolerance, resilience & reliability
- Scientific computing & simulation, numerical algorithm & software

Honors and Awards

- *Meta Research Award*, 2022.
- *R&D100 Award Winner*. R&D Word, 2021.
- *Third Place Award*. ACM Student Research Competition (SRC) in Supercomputing (SC), 2021.

- *Early Career Researchers Award for Excellence in High Performance Computing*. IEEE CS TCHPC, 2020.
- *CISE Research Initiation Initiative (CRII) Award*. National Science Foundation, 2020.
- *Best Paper Award*. IEEE Cluster Conference, 2018.
- *Dissertation Year Program (DYP) Award*. University of California, Riverside, 2017.
- *Dean's Distinguished Fellowship*. University of California, Riverside, 2013.
- *Outstanding Student Scholarship*. University of Science and Technology of China, 2012, 2010, and 2009.
- *First Prize in National High School Mathematics League*. Anhui Province, China, 2008.
- *First Prize in National Olympiad in Informatics*. Anhui Province, China, 2005.
- *Zhao Wei Fellowship*. Anhui Province, China, 2006.

Grants and Contracts

$12.5** million in total, my share is **$2.15M

Summary: https://eecs.wsu.edu/~dtao/funding_summary.png

Current Funded Projects

- **Site PI**, *Collaborative Research: OAC Core: CEAPA: A Systematic Approach to Minimize Compression Error Propagation in HPC Applications,* $250,000, 06/01/2022–05/31/2025, NSF (OAC-2211539).
  Lead PI: Guanpeng Li @ U. Iowa; Collaborators: Tianbao Yang @ U. Iowa, Samuel Li @ NCAR, Zarija Lukić @ LBNL, Sheng Di @ ANL, Ang Li @ PNNL.
- **Site PI**, *Collaborative Research: Elements: ROCCI: Integrated Cyberinfrastructure for In Situ Lossy Compression Optimization Based on Post Hoc Analysis Requirements,* $280,000, 09/15/2021–08/31/2024, NSF (OAC-2104024).
  Lead PI: Sheng Di @ U. Chicago; Co-PI: Franck Cappello @ U. Chicago; Collaborators: Junjing Deng @ ANL, Zarija Lukic, Suren Byna @ LBNL.
- **Lead PI**, *CDS&E: Collaborative Research: HyLoC: Objective-driven Adaptive Hybrid Lossy Compression Framework for Extreme-Scale Scientific Applications*, $527,563, 08/01/2020–07/31/2023, NSF (OAC-2042084/OAC-2003709).
  Site PI: Sheng Di @ U. Chicago; Collaborators: Ian Foster, Fred Chong, Kyle Chard @ U. Chicago, Junjing Deng @ ANL, Zarija Lukic @ LBNL.
- **Sole PI**, *CRII: OAC: An Efficient Lossy Compression Framework for Reducing Memory Footprint for Extreme-Scale Deep Learning on GPU-Based HPC Systems*, $189,593, 05/01/2020–04/30/2023, NSF (OAC-2034169).
- **Site PI**, *FAIR Benchmarks Supporting AI and Simulation Research*, $147,212, 08/16/2022–8/15/2023, Argonne National Laboratory, DOE (No. 2F-60243).
- **Site PI**, *FAIR Surrogate Benchmarks Supporting AI and Simulation Research*, $60,912, 08/15/2021–12/31/2022, Argonne National Laboratory, DOE (No. 1F-60533).
- **Site PI**, *Exploring Multiresolution Based Compression for Extreme-Scale Scientific Applications*, $149,897, 01/01/2021 –12/31/2022, Argonne National Laboratory, DOE (No. 1F-60329).
- **Site PI**, *Improving GPU Version of SZ for Scientific Applications at Extreme Scale*, $188,587, 06/01/2020–12/31/2022, Argonne National Laboratory, DOE (No. 0F-60172).
- **Site PI**, *SHF: Small: Tools for Productive High-performance Computing with GPUs*, $140,394, 01/01/2020–06/30/2022, University of Utah (10054032-WSU), NSF (SHF-2018016).
  Lead PI: P. (Saday) Sadayappan @ U. Utah.
- **Co-PI**, *CC* Compute: Accelerating Advances in Science and Engineering at The University of Alabama Through HPC Infrastructure*, $399,995 (my share: $100,000), 07/01/2020–06/30/2022, NSF (OAC-2018846).

PI: Jeff Carver @ U. Alabama; Co-PIs: David Dixon, Xiaoyan Hong @ U. Alabama.

Completed Funded Projects

- **Site PI**, *Improving Lossy Compression for Scientific Applications at Extreme Scale*, $21,588, 08/15/2019–12/31/2019, Argonne National Laboratory, DOE (No. 9F-60232).
- **Co-PI**, *Center for Remote Sensing of Snow and Soil Moisture*, $10,000,000 (my share: $500,000), 06/01/2019–05/31/2020, NOAA UCAR (No. SUBAWD000837).
  PI: S. Prasad Gogineni @ U. Alabama.

Industry Gifts

- **Lead PI,** *Accelerating Communication in DLRM via Frequency-aware Lossy Compression*, $50,000, 09/2022–08/2023, Meta Research. Co-PI: Tong Geng @ U. Rochester (award notice).
- **Sole PI**, *Accelerating Parallel I/O via Hardware-Algorithm Co-Design of Efficient and Adaptive Lossy Compression*, Xilinx Alveo U50 Data Center Accelerator Card (estimated value: $3,000), 05/2021, Xilinx.
- **Sole PI**, *CPU Inference of Sparse DNN Models on AMD EPYC*, $10,000, 02/2021, AMD.
- **Sole PI**, *FPGA-Enhanced Scientific Data Management*, Xilinx Virtex UltraScale+ FPGA VCU118 Evaluation Kit (estimated value: $7,000), 7/2020, Xilinx.
- **Sole PI**, *FPGA-Enhanced Lossy Compression for Scientific Data*, 1x Xilinx Zynq-7000 SoC ZC706 Evaluation Kit and 2x ZedBoard Xilinx Zynq-7000 (estimated value: $4,000), 12/2020, Xilinx.

Funded Computing Resources

- **Sole PI,** *Computing Resources for Engineering Cloud Computing Course at Indiana University*, 100,000 ACCESS Credits, 09/06/2022 – 08/31/2023, ACCESS CI, CIS220090.
- **Sole PI,** *CEAPA: A Systematic Approach to Minimize Compression Error Propagation in HPC Applications,* 20K SUs, 07/28/2022 – 01/28/2023, Chameleon Cloud, CHI-221007.
- **Sole PI**, *Improving Performance of Lossy Compression for Extreme-scale Scientific Applications on A64FX*, 1,000 node hours + 500 GB storage at NSF Ookami system, 05/01/2021 – 04/30/2023, Stony Brook University (abstract).
- **Sole PI**, *An Efficient Lossy Compression Framework for Reducing Memory Footprint for Extreme-Scale Deep Learning on GPU-Based HPC Systems*, 7,500 GPU Hours + 30,000 CPU Hours + 2,000 GBs at PSC Bridges-2 system (estimated value: $4,482.5), 05/01/2021 – 04/30/2023, XSEDE, ASC200032 (abstract).

Research Expenditures

- 2021 CY: $ 205,709
- 2020 CY: $ 58,173
- 2019 CY: $ 21,588

Total expenditures: $285,470 (CY2019~2021)

Travel Grants

- 12/2018 Faculty Travel Grant $500, LBL SRP Program
- 06/2018 Student Travel Grant $1,400, ACM HPDC'18
- 11/2017 Student Travel Grant $1,500, ACM/IEEE SC'17
- 06/2017 Student Travel Grant $250, IEEE IPDPS'17

- 06/2016 Student Travel Grant $1,000, ACM HPDC'16

Publications

Google citations: https://scholar.google.com/citations?user=Ppjzn_EAAAAJ&hl

*underlined: my students, *: corresponding author, AR: acceptance rate*

Referred Conference Publications

1. **[SC'22]** Sian Jin, **Dingwen Tao***, Houjun Tang, Sheng Di, Suren Byna, Zarija Lukic, Franck Cappello. "Accelerating Parallel Write via Deeply Integrating Predictive Lossy Compression with HDF5." To appear in the *International Conference for High Performance Computing, Networking, Storage and Analysis*, Dallas, TX, USA, November 13–18, 2022. [**AR:** 81/330 = **24.5%**]
2. **[PACT'22]** Xinyu Chen, Marco Minutoli, Jiannan Tian, Mahantesh Halappanavar, Ananth Kalyanaraman, **Dingwen Tao***. "HBMax: Optimizing Memory Efficiency for Parallel Influence Maximization on Multicore Architectures." To appear in the *31st International Conference on Parallel Architectures and Compilation Techniques*, Chicago, IL, October 10–12, 2022.
3. **[MASCOTS'22]** Griffin Dube, Jiannan Tian, Sheng Di, **Dingwen Tao**, Jon Calhoun, Franck Cappello. "Efficient Error-Bounded Lossy Compression for CPU Architectures." To appear in the *30th International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, Nice, France, October 18–20, 2022.
4. **[FPL'22]** Chengming Zhang, Tong Geng, Anqi Guo, Jiannan Tian, Martin Herbordt, Ang Li, **Dingwen Tao***. "H-GCN: A Graph Convolutional Network Accelerator on Xilinx Versal AI Engines." *The 32nd International Conference on Field-Programmable Logic and Applications*, Belfast, United Kingdom, August 29–September 2, 2022. DOI: 10.1109/FPL57034.2022.00040. [**AR:** 33/129 = **25.6%**]
5. **[ICS'22]** Chengming Zhang, Sian Jin, Tong Geng, Jiannan Tian, Ang Li, **Dingwen Tao***. "CEAZ: Accelerating Parallel I/O via Hardware-Algorithm Co-Designed Adaptive Lossy Compression." *The 36th ACM International Conference on Supercomputing*, Virtual, USA, June 27–30, 2022. DOI: 10.1145/3524059.3532362. [**AR:** 39/165 = **23.6%**]
6. **[ICS'22]** Heng Zhang, Lingda Li, Hang Liu, Donglin Zhuang, Rui Liu, Chengying Huan, Shuang Song, **Dingwen Tao**, Yongchao Liu, Charles He, Yanjun Wu, Shuaiwen Leon Song. "Bring Orders into Uncertainty: Enabling Efficient Uncertain Graph Processing via Novel Path Sampling on Multi-Accelerator System." *The 36th ACM International Conference on Supercomputing*, Virtual, USA, June 27–30, 2022. DOI: 10.1145/3524059.3532379. [**AR:** 39/165 = **23.6%**]
7. **[HPDC'22]** Daoce Wang, Jesus Pulido, Pascal Grosset, Sian Jin, Jiannan Tian, James Ahrens, **Dingwen Tao***. "TAC: Optimizing Error-Bounded Lossy Compression for Three-Dimensional Adaptive Mesh Refinement Simulations." *The 31st ACM International Symposium on High-Performance Parallel and Distributed Computing*, Minneapolis, MN, June 27–July 1, 2022. DOI: 10.1145/3502181.3531458. [**AR:** 21/108 = **19.4%**]
8. **[HPDC'22]** Xiaodong Yu, Sheng Di, Kai Zhao, Jiannan Tian, **Dingwen Tao**, Xin Liang, Franck Cappello. "Ultra-fast Error-bounded Lossy Compression for Scientific Dataset." *The 31st ACM International Symposium on High-Performance Parallel and Distributed Computing*, Minneapolis, MN, June 27–July 1, 2022. DOI: 10.1145/3502181.3531473. [**AR:** 21/108 = **19.4%**]
9. **[ICDE'22]** Sian Jin, Sheng Di, Jiannan Tian, Suren Byna, **Dingwen Tao***, Franck Cappello. "Improving Prediction-Based Lossy Compression Dramatically via Ratio-Quality Modeling." *The 38th IEEE International Conference on Data Engineering*, Virtual Event, May 9–12, 2022. DOI: 10.1109/ICDE53745.2022.00232. [**AR:** 211/780 = **27.1%**]
10. **[IPDPS'22]** Cody Rivera, Sheng Di, Xiaodong Yu, Jiannan Tian, **Dingwen Tao***, Franck Cappello. "Optimizing Huffman Decoding for Error-Bounded Lossy Compression on GPUs." *The 36th IEEE International Parallel and Distributed Symposium*, Lyon, France, May 30–June 3, 2022. DOI: 10.1109/IPDPS53621.2022.00075. [**AR:** 123/474 = **25.9%**]
11. **[VLDB'22]** Sian Jin, Chengming Zhang, Jiannan Tian, Yunhe Feng, Hui Guan, Guanpeng Li, Shuaiwen Leon Song, **Dingwen**

Tao*. "COMET: A Novel Memory-Efficient Deep Learning Training Framework by Using Error-Bounded Lossy Compression." *The 48th International Conference on Very Large Data Bases*, Australia, September 5–9, 2022. DOI: 10.14778/3503585.3503597.

12. **[HiPC'21]** Yuanjian Liu, Sheng Di, Kai Zhao, Sian Jin, Chen Wang, Kyle Chard, **Dingwen Tao**, Ian Foster, Franck Cappello. "Optimizing Multi-Range based Error-Bounded Lossy Compression for Scientific Datasets." *2021 IEEE 28th International Conference on High Performance Computing, Data, and Analytics*, Virtual Event, December 18–21, 2021. DOI: 10.1109/HiPC53243.2021.00036. [**AR:** 36/145 = **24.8%**]
13. **[Cluster'21]** Jiannan Tian, Sheng Di, Xiaodong Yu, Cody Rivera, Kai Zhao, Sian Jin, Yunhe Feng, Xin Liang, **Dingwen Tao***, Franck Cappello. "Optimizing Error-Bounded Lossy Compression for Scientific Data on GPUs." *2021 IEEE International Conference on Cluster Computing*, Virtual Event, Sept 7–10, 2021. DOI: 10.1109/Cluster48925.2021.00047. [**AR:** 48/163 = **29.4%**]
14. **[Cluster'21]** Bo Fang^, Daoce Wang^, Sian Jin, Quincey Koziol, Zhao Zhang, Qiang Guan, Suren Byna, Sriram Krishnamoorthy, **Dingwen Tao***. "Characterizing Impacts of Storage Faults on HPC Applications: A Methodology and Insights." *2021 IEEE International Conference on Cluster Computing*, Virtual Event, September 7–10, 2021. DOI: 10.1109/Cluster48925.2021.00048. ^equal contribution. [**AR:** 48/163 = **29.4%**]
15. **[Cluster'21]** Jinyang Liu, Sheng Di, Kai Zhao, Sian Jin, **Dingwen Tao**, Xin Liang, Zizhong Chen, Franck Cappello. "Exploring Autoencoder-Based Error-Bounded Compression for Scientific Data." *2021 IEEE International Conference on Cluster Computing*, Virtual Event, September 7–10, 2021. DOI: 10.1109/Cluster48925.2021.00034. [**AR:** 48/163 = **29.4%**]
16. **[Cluster'21]** Xiaodong Yu, Sheng Di, Ali Murat Gok, **Dingwen Tao**, Franck Cappello. "cuZ-Checker: A GPU-Based Ultra-Fast Assessment System for Lossy Compressions." *2021 IEEE International Conference on Cluster Computing*, Virtual Event, September 7–10, 2021. DOI: 10.1109/Cluster48925.2021.00065. [**AR:** 48/163 = **29.4%**]
17. **[ICS'21]** Chengming Zhang, Geng Yuan, Wei Niu, Jiannan Tian, Sian Jin, Donglin Zhuang, Zhe Jiang, Yanzhi Wang, Bin Ren, Shuaiwen Song, **Dingwen Tao***. "ClickTrain: Efficient and Accurate End-to-End Deep Learning Training via Fine-Grained Architecture-Preserving Pruning." *The 35th ACM International Conference on Supercomputing*, Virtual, USA, June 14–17, 2021. DOI: 10.1145/3447818.3459988. [**AR:** 38/157 = **24.2%**]
18. **[HPDC'21]** Sian Jin, Jesus Pulido, Pascal Grosset, Jiannan Tian, **Dingwen Tao***, James Ahrens. "Adaptive Configuration of In Situ Lossy Compression for Cosmology Simulations via Fine-Grained Rate-Quality Modeling." *The 30th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Virtual Event, Sweden, June 21–25, 2021. DOI: 10.1145/3431379.3460653. [**AR:** 19/92 = **20.7%**]
19. **[IPDPS'21]** Jiannan Tian, Cody Rivera, Sheng Di, Jieyang Chen, Xin Liang, **Dingwen Tao***, Franck Cappello. "Revisiting Huffman Coding: Toward Extreme Performance on Modern GPU Architectures." *The 35th IEEE International Parallel and Distributed Processing Symposium*, Portland, May 17–21, 2021. DOI: 10.1109/IPDPS49936.2021.00097. [**AR:** 105/462 = **22.7%**]
20. **[PPoPP'21]** Sian Jin, Guanpeng Li, Shuaiwen Leon Song, **Dingwen Tao***. "A Novel Memory-Efficient Deep Learning Training Framework via Error-Bounded Lossy Compression." *The 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Virtual, February 27–March 3, 2021. DOI: 10.1145/3437801.3441597.
21. **[PPoPP'21]** Heng Zhang, Lingda Li, Donglin Zhuang, Rui Liu, Shuang Song, **Dingwen Tao**, Yanjun Wu, Shuaiwen Leon Song. "An Efficient Uncertain Graph Processing Framework for Heterogeneous Architectures." *The 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Virtual, February 27–March 3, 2021. DOI: 10.1145/3437801.3441584.
22. **[PACT'20]** Jiannan Tian, Sheng Di, Kai Zhao, Sian Jin, Megan Hickman, Robert Underwood, Xin Liang, Jon Calhoun, **Dingwen Tao***, Franck Cappello. "cuSZ: An Efficient GPU Based Error-Bounded Lossy Compression Framework for Scientific Data." *The 29th International Conference on Parallel Architectures and Compilation Techniques*, Atlanta, GA, USA, Oct 3–7, 2020. DOI:

10.1145/3410463.3414624. [**AR:** 35/137 = **25.5%**]

23. **[ICPP'20]** Zhenbo Hu, Xiangyu Zou, Wen Xia, Sian Jin, **Dingwen Tao**, Yang Liu, Weizhe Zhang, Zheng Zhang. "Delta-DNN: Efficiently Compressing Deep Neural Networks via Exploiting Floats Similarity." *The 49th International Conference on Parallel Processing*, Edmonton, AB, Canada, August 17–20 2020. DOI: 10.1145/3404397.3404408. [**AR:** 78/269 = **28.9%**]
24. **[DAC'20]** Peiyan Dong, Siyue Wang, Wei Niu, Chengming Zhang, Sheng Lin, Zhengang Li, Yifan Gong, Bin Ren, Xue Lin, **Dingwen Tao***. "RTMobile: Beyond Real-Time Mobile Acceleration of RNNs for Speech Recognition." *The 57th Annual Design Automation Conference*, San Francisco, CA, USA, July 19–23, 2020. DOI: 10.5555/3437539.3437579. [**AR:** 228/984 = **23.2%**]
25. **[HPDC'20]** Kai Zhao, Sheng Di, Xin Liang, Sihuan Li, **Dingwen Tao**, Zizhong Chen, Franck Cappello. "Significantly Improving Lossy Compression for HPC Datasets with Second-Order Prediction and Parameter Optimization." *The 29th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Stockholm, Sweden, June 23–26, 2020. DOI: 10.1145/3369583.3392688. [**AR:** 16/71 = **22.5%**]
26. **[IPDPS'20]** Sian Jin, Pascal Grosset, Christopher M. Biwer, Jesus Pulido, Jiannan Tian, **Dingwen Tao***, James Ahrens. "Understanding GPU-Based Lossy Compression for Extreme-Scale Cosmological Simulations." *34th IEEE International Parallel and Distributed Symposium*, New Orleans, LA, May 18–22, 2020. DOI: 10.1109/IPDPS47924.2020.00021. [**AR:** 110/446 = **24.7%**]
27. **[PPoPP'20]** Jiannan Tian, Sheng Di, Chengming Zhang, Xin Liang, Sian Jin, Dazhao Cheng, **Dingwen Tao***, Franck Cappello. "waveSZ: A Hardware-Algorithm Co-Design of Efficient Lossy Compression for Scientific Data." *The 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, San Diego, CA, USA, Feb 22–26, 2020. DOI: 10.1145/3332466.3374525. [**AR:** 28/121 = **23.1%**]
28. **[HPDC'19]** Sian Jin, Sheng Di, Xin Liang, Jiannan Tian, **Dingwen Tao***, Franck Cappello. "DeepSZ: A Novel Framework to Compress Deep Neural Networks by Using Error-Bounded Lossy Compression." *The 28th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Phoenix, AZ, USA, June 24–28, 2019. DOI: 10.1145/3307681.3326608. [**AR:** 22/106 = **20.8%**]
29. **[SC'19]** Xin Liang, Sheng Di, Sihuan Li, **Dingwen Tao**, Bogdan Nicolae, Zizhong Chen, Franck Cappello. "Significantly Improving Lossy Compression Quality Based on An Optimized Hybrid Prediction Model." *The International Conference for High Performance Computing, Networking, Storage and Analysis*, Denver, CO, USA, Nov 17–22, 2019. DOI: 10.1145/3295500.3356193. [**AR:** 78/344 = **22.7%**]
30. **[BigData'19]** Yuqi Fu, Shaolun Zhang, Jose Terrero, Ying Mao, Guangya Liu, Sheng Li, **Dingwen Tao**. "Progress-based Container Scheduling for Short-lived Applications in a Kubernetes Cluster." *IEEE International Conference on Big Data*, Los Angeles, CA, USA, December 9–12, 2019. DOI: 10.1109/BigData47090.2019.9006427. [**AR:** 106/550 = **19.3%**]
31. **[BigData'19]** Donglin Yang, Wei Rang, Dazhao Cheng, Yu Wang, Jiannan Tian, **Dingwen Tao**. "Elastic Executor Provisioning for Iterative Workloads on Apache Spark." *IEEE International Conference on Big Data*, Los Angeles, CA, USA, December 9–12, 2019. DOI: 10.1109/BigData47090.2019.9006021. [**AR:** 106/550 = **19.3%**]
32. **[Cluster'19]** Xin Liang, Sheng Di, **Dingwen Tao**, Sihuan Li, Bogdan Nicolae, Zizhong Chen, Franck Cappello. "Improving Performance of Data Dumping with Lossy Compression for Scientific Simulation." *IEEE International Conference on Cluster Computing, Albuquerque*, NM, USA, September 23–26, 2019. DOI: 10.1109/CLUSTER.2019.8891037. [**AR:** 39/141 = **27.7%**]
33. **[HPCC'19]** Xiangyu Zou, Tao Lu, Sheng Di, **Dingwen Tao**, Wen Xia, Xuan Wang, Weizhe Zhang, Qing Liao. "Accelerating Lossy Compression on HPC datasets via Partitioning Computation for Parallel Processing." *IEEE International Conference on High Performance Computing and Communications*, Zhangjiajie, Aug 10–12, 2019. DOI: 10.1109/HPCC/SmartCity/DSS.2019.00246.
34. **[ICS'19]** Jieyang Chen, Nan Xiong, Xin Liang, **Dingwen Tao**, Sihuan Li, Kaiming Ouyang, Kai Zhao, Nathan DeBardeleben,

Qiang Guan, Zizhong Chen. "TSM2: Optimizing Tall-and-Skinny Matrix-Matrix Multiplication on GPUs." *The 33rd ACM International Conference on Supercomputing*, AZ, USA, June 26–28, 2019. DOI: 10.1145/3330345.3330355. [**AR:** 45/193 = **23.3%**]

35. **[MSST'19]** Xiangyu Zou, Tao Lu, Wen Xia, Xuan Wang, Weizhe Zhang, Sheng Di, **Dingwen Tao**, Franck Cappello. "Accelerating Relative-error Bounded Lossy Compression for HPC datasets with Precomputation-Based Mechanisms." *IEEE Symposium on Mass Storage Systems and Technologies*, Santa Clara, CA, USA, May 20–24, 2019. DOI: 10.1109/MSST.2019.00-15. [**AR: 29%**]
36. **[NYSDS'18]** Line Pouchard, Kevin Huck, Gyorgy Matyasfalvi, **Dingwen Tao**, Li Tang, Huub Van Dam, Shinjae Yoo. "Prescriptive Provenance for Streaming Analysis of Workflows at Scale." *2018 New York Scientific Data Summit*, New York, NY, USA, August 6–8, 2018. DOI: 10.1109/NYSDS.2018.8538951.
37. **[BigData'18]** Xin Liang, Sheng Di, **Dingwen Tao**, Sihuan Li, Shaomeng Li, Hanqi Guo, Zizhong Chen, Franck Cappello. "Error-Controlled Lossy Compression Optimized for High Compression Ratios of Scientific Datasets." *IEEE International Conference on Big Data,* Seattle, WA, USA, December 10–13. DOI: 10.1109/BigData.2018.8622520. [**AR:** 98/518 **= 18.9%**]
38. **[SC'18]** Jieyang Chen, Hongbo Li, Sihuan Li, Xin Liang, Panruo Wu, **Dingwen Tao**, Kaiming Ouyang, Yuanlai Liu, Kai Zhao, Qiang Guan, Zizhong Chen. "FT-MAGMA: Fault Tolerance Dense Matrix Decomposition on Heterogeneous Systems with GPUs." *The International Conference for High Performance Computing, Networking, Storage and Analysis*, Dallas, TX, USA, Nov 11–16, 2018. DOI: 10.1109/SC.2018.00071. [**AR:** 68/288 = **23.6%**]
39. **[Cluster'18] Dingwen Tao**, Sheng Di, Xin Liang, Zizhong Chen, Franck Cappello. "Fixed-PSNR Lossy Compression for HPC Scientific Data." *IEEE International Conference on Cluster Computing*, Belfast, UK, September 10–13, 2018. DOI: 10.1109/CLUSTER.2018.00048. [**AR:** 44/154 = **28.6%**]
40. **[Cluster'18]** Xin Liang, Sheng Di, **Dingwen Tao**, Zizhong Chen, Franck Cappello. "An Efficient Transformation Scheme for Lossy Data Compression with Point-wise Relative Error Bound." *IEEE International Conference on Cluster Computing*, Belfast, UK, September 10–13, 2018. DOI: 10.1109/CLUSTER.2018.00036. (**Best Track Paper Award**) (4/154).
41. **[Cluster'18]** Ali Murat Gok, Sheng Di, Yury Alexeev, **Dingwen Tao,** Vladimir Mironov, Franck Cappello. "PaSTRI: Error-Bounded Lossy Compression for Two-Electron Integrals in Quantum Chemistry." *IEEE International Conference on Cluster Computing*, Belfast, UK, September 10–13, 2018. DOI: 10.1109/CLUSTER.2018.00013. (**Best Overall Paper Award**) (1/154)
42. **[HPDC'18] Dingwen Tao**, Sheng Di, Xin Liang, Zizhong Chen, Franck Cappello. "Improving Performance of Iterative Methods by Lossy Checkponting." *The 27th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Tempe, AZ, USA, June 11– 15, 2018. DOI: 10.1145/3208040.3208050. [**AR:** 22/112 = **19.6%**]
43. **[BigData'17] Dingwen Tao**, Sheng Di, Zizhong Chen, Franck Cappello. "In-Depth Exploration of Single-Snapshot Lossy Compression Techniques for N-Body Simulations." *IEEE International Conference on Big Data,* Boston, MA, USA, December 11–14, 2017. DOI: 10.1109/BigData.2017.8257962. [**AR:** 87/437 = **19.9%**]
44. **[SC'17]** Xin Liang, Jieyang Chen, **Dingwen Tao**, Sihuan Li, Panruo Wu, Hongbo Li, Kaiming Ouyang, Yuanlai Liu, Fengguang Song, Zizhong Chen. "Correcting Soft Errors Online in Fast Fourier Transform." *The International Conference for High Performance Computing, Networking, Storage and Analysis*, Denver, CO, USA, Nov 12–17, 2017. DOI: 10.1145/3126908.3126915. [**AR:** 61/327 = **18.6%**]
45. **[EuroPar'17]** Ian Foster, Mark Ainsworth, Bryce Allen, Julie Bessac, Franck Cappello, Jong Youl Choi, Emil Constantinescu, Philip E Davis, Sheng Di, Wendy Di, Hanqi Guo, Scott Klasky, Kerstin Kleese Van Dam, Tahsin Kurc, Qing Liu, Abid Malik, Kshitij Mehta, Klaus Mueller, Todd Munson, George Ostouchov, Manish Parashar, Tom Peterka, Line Pouchard, **Dingwen Tao**, Ozan Tugluk, Stefan Wild, Matthew Wolf, Justin M Wozniak, Wei Xu, Shinjae Yoo. "Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales." *International European Conference on Parallel and Distributed Computing*, Santiago

de Compostela, Spain, Aug 28–Sept 1, 2017. DOI: 10.1007/978-3-319-64203-1_1.

46. **[IPDPS'17] Dingwen Tao**, Sheng Di, Zizhong Chen, Franck Cappello. "Significantly Improving Lossy Compression for Scientific Data Sets Based on Multidimensional Prediction and Error-Controlled Quantization." *The 31st IEEE International Parallel and Distributed Processing Symposium*, Orlando, FL, USA, May 29–June 2, 2017. DOI: 10.1109/IPDPS.2017.115. [**AR:** 116/508 = **22.8%**]
47. **[PPoPP'17]** Panruo Wu, Nathan Debardeleben, Qiang Guan, Sean Blanchard, Jieyang Chen, **Dingwen Tao**, Xin Liang, Kaiming Ouyang, Sihuan Li, Zizhong Chen. "Silent Data Corruption Resilient Two-sided Matrix Factorizations." *The 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Austin, TX, USA, February 4–8, 2017. DOI: 10.1145/3155284.3018750. [**AR:** 29/132 = **22.0%**]
48. **[SC'16]** Jieyang Chen, Li Tan, Panruo Wu, **Dingwen Tao**, Hongbo Li, Xin Liang, Sihuan Li, Rong Ge, Laxmi Bhuyan, Zizhong Chen. "GreenLA: Green Linear Algebra Software for GPU-Accelerated Heterogeneous Computing." *The International Conference for High Performance Computing, Networking, Storage and Analysis*, Salt Lake City, UT, USA. DOI: 10.1109/SC.2016.56. [**AR:** 81/442 = **18.3%**]
49. **[HPDC'16] Dingwen Tao**, Shuaiwen Leon Song, Sriram Krishnamoorthy, Panruo Wu, Xin Liang, Eddy Z. Zhang, Darren Kerbyson, Zizhong Chen. "New-Sum: A Novel Online ABFT Scheme for General Iterative Methods." *The 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Kyoto, JAPAN, May 31– June 4, 2016. DOI: 10.1145/2907294.2907306. [**AR:** 20/129 = **15.5%**]
50. **[HPDC'16]** Panruo Wu, Qiang Guan, Nathan DeBardeleben, Sean Blanchard, **Dingwen Tao**, Xin Liang, Jieyang Chen, Zizhong Chen. "Towards Practical Algorithm Based Fault Tolerance in Dense Linear Algebra." *The 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*, Kyoto, JAPAN, May 31–June 4, 2016. DOI: 10.1145/2907294.2907315. [**AR:** 20/129 = **15.5%**]
51. **[ICPADS'14]** Longxiang Chen, **Dingwen Tao**, Panruo Wu, Zizhong Chen. "Extending Checksum-Based ABFT to Tolerate Soft Errors Online in Iterative Methods." *IEEE International Conference on Parallel and Distributed Systems*, Hsinchu, Taiwan, December 16–19, 2014. DOI: 10.1109/PADSW.2014.7097827.

Referred Journal Publications

1. **[TPDS]** Yuanjian Liu, Sheng Di, Kai Zhao, Sian Jin, Cheng Wang, Kyle Chard, **Dingwen Tao**, Ian Foster, Franck Cappello. "Optimizing Error-Bounded Lossy Compression for Scientific Data with Diverse Constraints." *IEEE Transactions on Parallel and Distributed Systems*. 10.1109/TPDS.2022.3194695.
2. **[TBD]** Xin Liang, Kai Zhao, Sheng Di, Sihuan Li, Robert Underwood, Ali M. Gok, Jiannan Tian, Junjing Deng, Jon C. Calhoun, **Dingwen Tao**, Zizhong Chen, and Franck Cappello. "SZ3: A Modular Framework for Composing Prediction-Based Error-Bounded Lossy Compressors." *IEEE Transactions on Big Data*.
3. **[TC]** Xin Liang, Ben Whitney, Jieyang Chen, Lipeng Wan, Qing Liu, **Dingwen Tao**, James Kress, Dave Pugmire, Matthew Wolf, Nobert Podhorszki, Scott Klasky. "MGARD+: Optimizing Multilevel Methods for Error-bounded Scientific Data Reduction." *IEEE Transactions on Computers* 71(7): 1522-1536 (2022). DOI: 10.1109/TC.2021.3092201.
4. [**ISJ**] Ying Mao, Yuqi Fu, Wenjia Zheng, Long Cheng, Qingzhi Liu, **Dingwen Tao**. "Speculative Container Scheduling for Deep Learning Applications in a Kubernetes Cluster." *IEEE Systems Journal* (2021). DOI: 10.1109/JSYST.2021.3129974.
5. **[JPDC]** Cody Rivera, Jieyang Chen, Nan Xiong, Jing Zhang, Shuaiwen Leon Song, **Dingwen Tao***. "TSM2X: High-Performance Tall-and-Skinny Matrix-Matrix Multiplication on GPUs." *Journal of Parallel and Distributed Computing* 151 (2021): 70-85. DOI: 10.1016/j.jpdc.2021.02.013.

6. **[TPDS]** Xiangyu Zou, Tao Lu, Wen Xia, Xuan Wang, Weizhe Zhang, Haijun Zhang, Sheng Di, **Dingwen Tao**, Franck Cappello. "Performance Optimization for Relative-Error-Bounded Lossy Compression on Scientific Data." *IEEE Transactions on Parallel and Distributed Systems* 31(7): 1665-1680 (2020). DOI: 10.1109/TPDS.2020.2972548.
7. **[IJHPCA]** Franck Cappello, Sheng Di, Sihuan Li, Xin Liang, Ali Murat Gok, **Dingwen Tao**, Chun Hong Yoon, Xin-Chuan Wu, Yuri Alexeev, Frederic T. Chong. "Use Cases of Lossy Compression for Floating-Point Data in Scientific Datasets." *The International Journal of High Performance Computing Applications* 33(6) (2019). DOI: 10.1177/1094342019853336.
8. **[TPDS] Dingwen Tao**, Sheng Di, Xin Liang, Zizhong Chen, Franck Cappello. "Optimizing Lossy Compression Rate-Distortion from Automatic Online Selection Between SZ and ZFP." *IEEE Transactions on Parallel and Distributed Systems* 30(8): 1857-1871 (2019). DOI: 10.1109/TPDS.2019.2894404.
9. **[IJHPCA] Dingwen Tao**, Sheng Di, Hanqi Guo, Zizhong Chen, Franck Cappello. "Z-checker: A Framework for Assessing Lossy Compression of Scientific Data." *The International Journal of High Performance Computing Applications* 33(2) (2019). DOI: 10.1177/1094342017737147.
10. **[TPDS]** Sheng Di, **Dingwen Tao**, Xin Liang, Franck Cappello. "Efficient Lossy Compression for Scientific Data based on Pointwise Relative Error Bound." *IEEE Transactions on Parallel and Distributed Systems* 30(2): 331-345 (2019). DOI: 10.1109/TPDS.2018.2859932.

Referred Workshop Publications

1. **[IWBDR'21]** Jinyang Liu, Sihuan Li, Sheng Di, Xin Liang, Kai Zhao, **Dingwen Tao**, Zizhong Chen, and Franck Cappello. "Improving Lossy Compression for SZ by Exploring the Best-Fit Lossless Compression Techniques." *The 2nd International Workshop on Big Data Reduction*, co-located with 2021 IEEE International Conference on Big Data, December 17, 2021. DOI: 10.1109/BigData52589.2021.9671954.
2. **[DRBSD-7]** Yuanjian Liu, Sheng Di, Kai Zhao, Sian Jin, Cheng Wang, Kyle Chard, **Dingwen Tao**, Ian Foster, Franck Cappello. "Understanding Effectiveness of Multi-Error-Bounded Lossy Compression for Preserving Ranges of Interest in Scientific Analysis." *The 7th International Workshop on Data Analysis and Reduction for Big Scientific Data*, co-located with SC'21, St. Louis, MO, USA, Nov 14-19, 2021. DOI: 10.1109/DRBSD754563.2021.00010.
3. [**IWBDR'20**] Baodi Shan, Aabid Shamji, Jiannan Tian, Guanpeng Li, **Dingwen Tao***. "LCFI: A Fault Injection Tool for Studying Lossy Compression Error Propagation in HPC Programs." *The 1st International Workshop on Big Data Reduction*, co-located with 2020 IEEE International Conference on Big Data, December 10, 2020. DOI: 10.1109/BigData50022.2020.9378104.
4. [**IWBDR'20**] Kai Zhao, Sheng Di, Xin Liang, Sihuan Li, **Dingwen Tao**, Julie Bessac, Zizhong Chen, Franck Cappello. "SDRBench: Scientific Data Reduction Benchmark for Lossy Compressors." *The 1st International Workshop on Big Data Reduction*, co-located with 2020 IEEE International Conference on Big Data, December 10, 2020. DOI: 10.1109/BigData50022.2020.9378449.
5. **[DRBSD-4]** Xin Liang, Sheng Di, **Dingwen Tao**, Sihuan Li, Zizhong Chen, Franck Cappello. "Exploring Best Lossy Compression Strategy by Combining SZ with Spatiotemporal Decimation." *The 4th International Workshop on Data Reduction for Big Scientific Data*, co-located with SC'18, Dallas, TX, USA, Nov 12–15, 2018. https://sc18.supercomputing.org/proceedings/workshops/workshop_files/ws_drbsd108s1-file1.pdf.
6. **[DRBSD-2]** Sheng Di, **Dingwen Tao**, Franck Cappello. "An Efficient Approach to Lossy Compression with Pointwise Relative Error Bound." *The 2nd International Workshop on Data Reduction for Big Scientific Data*, co-located with SC'17, Denver, CO, USA, Nov 12–17, 2017. https://sc17.supercomputing.org/SC17%20Archive/workshops/workshop_pages/wkpr139.html.
7. **[DRBSD-1] Dingwen Tao**, Sheng Di, Zizhong Chen, Franck Cappello. "Exploration of Pattern-Matching Techniques for Lossy

Compression on Cosmology Simulation Data Sets." *The 1st International Workshop on Data Reduction for Big Scientific Data,* co-located with ISC'17, Frankfurt, Germany, June 22, 2017. DOI: 10.1007/978-3-319-67630-2_4.

Teaching Experience

Undergraduate Course

- CptS 233: *Advanced Data Structures in Java*, Spring 2021, Fall 2021 (WSU)
- CptS 455: *Introduction to Computer Networks*, Fall 2020, Fall 2021 (WSU)
- CS 470: *Computer Algorithms*, Fall 2019 (UA)
- CS 481: *High Performance Computing*, Spring 2019, Spring 2020 (UA)

Graduate Course

- ENGR 516: *Engineering Cloud Computing*, Fall 2022 (IUB)
- CS 570: *Computer Algorithms*, Fall 2019 (UA)
- CS 581: *High Performance Computing*, Spring 2019, Spring 2020 (UA)

Student Mentorship

Doctoral Dissertation Chair

Sian Jin (Fall 2018 – Current): passed qualifying exam in Fall 2019 at UA

Jiannan Tian (Spring 2019 – Current): passed qualifying exam in Spring 2020 at UA

Chengming Zhang (Fall 19 – Current): passed qualifying exam in Spring 2021 at WSU

Lizhi Xiang (Spring 2020 – Current): passed qualifying exam in Spring 2021 at WSU

Daoce Wang (Fall 2020 – Current): passed qualifying exam in Winter 2021 at WSU

Baixi Sun (Spring 2021 – Current): passed qualifying exam in Spring 2022 at WSU

Xinyu Chen (Fall 2021 – Current), Boyuan Zhang (Fall 2021 – Current), Pei-Yau Weng (Fall 2021 – Current)

Doctoral Dissertation Committee Member

Xiaoqin Fu (Spring 2021, WSU), Devjeet Roy (Fall 2020, WSU), Peyman Abbaszadeh (Spring 2019 – Spring 2020, UA), Nasir U. Eisty (Fall 2019 – Spring 2020, UA), Tasnuva Mahjabin (Fall 2019 – Spring 2020, UA)

Undergraduate Research Mentor

Eric Song (Summer 2021 – Current, WSU), Dung Hoang Le (Spring 2021 – Current, WSU), Xintong Jiang (Spring 2021 – Current, WSU), Cody Rivera (Spring 2019 – Current, UA), Philip Speegle (Fall 2019 – Spring 2020, UA), Jack O'Donohue (Summer 2019, UA), Aashman Gupta (Summer 2019, UA)

Professional Service

University Service

- *Curriculum Committee* for Computer Science at Washington State University (2021 – Present)
- *Search Committee* for Electrical Engineering at Washington State University (2021 – 2022)
- *External Search Committee* for Civil & Environmental Engineering (2022)
- *Search Committee* for Computer Science at the University of Alabama (2019 – 2020)

Journal Editor

- *Review Board* of IEEE Transactions on Parallel and Distributed Systems (TPDS)
- *Lead Guest Editor* of the Special Issue on Scalable Computing Systems and Networks in Supercomputing in the Journal of Supercomputing
- *Guest Editor* of the Special Issue on New Trends for High-Performance Computing in the Electronics journal

Conference/Workshop Organization

- *Program Chairs' Team* of 2023 IEEE International Parallel & Distributed Processing Symposium (IPDPS)
- *Program Co-chair* of 2021 IEEE International Conference on Scalable Computing & Communications (ScalCom)
- *Security Team Lead* of 2021 IEEE/ACM Supercomputing Conference (SC)
- *Program Co-chair* of the 1st & 2nd International Workshop on Big Data Reduction (IWBDR) held with 2020 & 2021 IEEE International Conference on Big Data (Big Data)
- *Program Co-chair* of the 1st & 2nd International Workshop on Big Data Analytics of Cyber-Physical Systems (CPSBigData) held with 2019 & 2020 IEEE Green and Sustainable Computing Conference (IGSC)

Conference/Workshop Program Committee

- 2023 AAAI Conference on Artificial Intelligence (AAAI)
- 2022, 2020 IEEE/ACM Supercomputing Conference (SC)
- 2022 ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC)
- 2022 IEEE/ACM International Symposium on Cluster, Cloud, and Internet Computing (CCGrid)
- 2022, 2021, 2020 IEEE International Conference on Big Data (Big Data)
- 2022, 2021 IEEE International Conference on Machine Learning and Applications (ICMLA)
- 2022, 2021, 2019, 2018 IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)
- 2023, 2021 IEEE International Parallel & Distributed Processing Symposium (IPDPS)
- 2021, 2020 IEEE International Conference on Cluster Computing (Cluster)
- 2020, 2019, 2018 IFIP International Conference on Network and Parallel Computing (NPC)
- 2020 International Conference on Parallel Processing (ICPP)
- 2022, 2020, 2019 IEEE International Conference on Services Computing (SCC)
- 2019, 2018 IEEE International Congress on Big Data (BigData Congress)
- 2018 IEEE International Conference on e-Science (eScience)
- 2021 Redefining Scalability for Diversely Heterogeneous Architectures (RSDHA) Workshop
- 2018, 2019, 2020 High-Performance Machine Learning (HPML) Workshops
- 2018 Fault Tolerant Systems (FTS) Workshop
- 2018 International Workshop on Large-Scale Deep Learning on Modern Heterogeneous Supercomputers (DLMHS)

Grant Reviewer

- 2022 U.S. NSF Cyberinfrastructure for Sustained Scientific Innovation (CSSI) Program
- 2022 National Research Foundation of Ukraine (NRFU) in Mathematical, Natural and Technical Sciences
- 2021 U.S. NSF Cyberinfrastructure for Sustained Scientific Innovation (CSSI) Program
- 2021 U.S. NASA EPSCoR Cooperative Agreement Notice (CAN) Program

Journal Reviewer

Publons verified reviews: https://publons.com/researcher/3034803/dingwen-tao/peer-review/

- IEEE Transactions on Parallel and Distributed Systems (JCR Impact Factor: 2.687)
- IEEE Transactions on Computers (JCR Impact Factor: 2.711)
- IEEE Transactions on Dependable and Secure Computing (JCR Impact Factor: 6.864)
- IEEE Transactions on Visualization and Computer Graphics (JCR Impact Factor: 4.558)
- IEEE Transactions on Knowledge and Data Engineering (JCR Impact Factor: 4.935)
- IEEE Transactions on Cloud Computing (JCR Impact Factor: 4.714)
- IEEE Transactions on Big Data (JCR Impact Factor: 3.344)
- IEEE Transactions on Smart Grid (JCR Impact Factor: 8.267)
- IEEE Transactions on Emerging Topics in Computing (JCR Impact Factor: 6.043)
- IEEE Transactions on Sustainable Computing (JCR Impact Factor: 2.456)
- IEEE Industrial Electronics Magazine (JCR Impact Factor: 13.593)
- IEEE Embedded Systems Letters (JCR Impact Factor: 1.689)
- IEEE Access (JCR Impact Factor: 5.213)
- ACM Transactions on Storage (JCR Impact Factor: 1.480)
- Journal of Parallel and Distributed Computing (JCR Impact Factor: 2.296)
- Journal of Supercomputing (JCR Impact Factor: 2.469)
- SIAM Journal on Scientific Computing (JCR Impact Factor: 1.976)
- Parallel Computing (JCR Impact Factor: 1.119)
- Journal of Systems Architecture (JCR Impact Factor: 2.552)
- Integration the VLSI Journal (JCR Impact Factor: 2.037)
- ETRI Journal (JCR Impact Factor: 1.094)

External Reviewer

- 2022 International Conference on Parallel Processing (ICPP)
- 2021 IEEE Transactions on Parallel and Distributed Systems (TPDS) Special Section on Parallel and Distributed Computing Techniques for AI, ML, and DL
- 2018, 2017, 2016, 2015 IEEE/ACM Supercomputing Conference (SC)
- 2018, 2017, 2016 IEEE International Parallel & Distributed Processing Symposium (IPDPS)
- 2017 ACM International Conference on Supercomputing (ICS)
- 2017 ACM International Symposium on High-Performance Parallel & Distributed Computing (HPDC)
- 2016 IEEE International Conference on Parallel & Distributed Systems (ICPADS)
- 2016 ACM International Conference on Supercomputing (ICS)
- 2015 International Conference on Parallel Architectures & Compilation Techniques (PACT)
- 2014 IEEE/ACM International Symposium on Cluster, Cloud, & Grid Computing (CCGrid)
- The 12th, 11th, 10th Workshop on Resiliency in High Performance Computing (Resilience) in Clusters, Clouds, and Grids held with Euro-Par 2019, 2018, 2017

Invited Talks

- "*Advancing HPC and ML Systems and Applications via Efficient Data Management*", Distinguished Lecture Series,

Department of Electrical and Computer Engineering, University of Rochester, October 2022.

- "*Advancing HPC and ML Systems and Applications via Efficient Data Management*", Technical Seminar Series, Department of Electrical Engineering and Computer Science, University of California, Merced, September 2022.
- "*Advancing HPC I/O and Storage via Efficient Data Compression*", Watson Software Paper Bar, Huawei, July 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Intelligent Systems Engineering, Indiana University, March 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Computer Science and Engineering, University of Notre Dame, March 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Computer and Data Sciences, Case Western Reserve University (CWRU), March 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Computer Science, George Mason University, March 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", School of Computing, University of Utah, March 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", University of Washington, February 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Computer Science, Norwegian University of Science and Technology (NUNT), February 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Electrical and Computer Engineering, Rutgers University, January 2022.
- "*Advancing HPC and AI Systems via Efficient Data Management*", School of Computer Science and Engineering, National Technological University (NTU), December 2021.
- "*Advancing HPC and AI Systems via Efficient Data Management*", School of Computer Science and Engineering, University of New South Wales (UNSW Sydney), November. 2021.
- "*Advancing HPC and AI Systems via Efficient Data Management*", Department of Computer Science, University of Warwick, November. 2021.
- "*Improving HPC and AI Systems via Error-Bounded Lossy Compression*", Microelectronics Thrust, Hong Kong University of Science and Technology (Guangzhou), October. 2021.
- "*Adaptive Configuration of In Situ Lossy Compression for Cosmology Simulations via Fine-Grained Rate-Quality Modeling*", DOE Computer Graphics Forum (CGF), Virtual, April 2021.
- "*Improving HPC and AI Systems and Applications via Error-Bounded Lossy Compression*", School of Computer Science and Technology, University of Science and Technology of China (USTC), October 2020.
- "*Scientific Data Reduction Challenges in the Era of Exascale Computing*", Computer Science and Mathematics Division, Oak Ridge National Laboratory, December 2019.
- "*Keeping-up with Scientific Data Explosion in the Era of Exascale Computing*", Department of Electrical and Computer Engineering, Boston University, November 2019
- "*Keeping-up with Scientific Data Explosion in the Era of Exascale Computing*", Department of Electrical and Computer Engineering, Northeastern University, November 2019
- "*Keeping-up with Scientific Data Explosion*", Data Science at Scale, Los Alamos National Laboratory, July 2019.
- "*Keeping-up with Flood of Scientific Data*", College of Information Science and Engineering, China University of Petroleum Beijing (CUPB), June 2019.
- "*High-Performance Computing at Extreme Scale: Data Reduction, Resilience, Scalability*", Institute of Computing

Technology (ICT), Chinese Academy of Sciences (CAS), October 2018.

- "*High-Performance Computing at Extreme Scale: Data Reduction, Resilience, Scalability*", College of Software Engineering, Tsinghua University, October 2018.
- "*GreenLA: Energy Efficient Linear Algebra Software for GPU-Accelerated Heterogeneous Computing*", Energy-Efficient Computing Workshop, HPC China 2018, Qingdao, China, September 2018.

Selected Software

My Lab's GitHub: https://github.com/hipdac-lab

- **SZ**: Fast, Effective, Parallel Error-bounded Exascale Scientific Data Lossy Compressor (**2021 R&D100 Award**)
- **cuSZ**: A GPU Accelerated Error-Bounded Lossy Compressor for Scientific Data
- **DeepSZ**: Lossy Compression Framework for Deep Neural Networks
- **TSM2X**: High-Performance Tall-and-Skinny Matrix-Matrix Multiplication on GPUs
- **HBMax**: Optimizing Memory Efficiency for Parallel Influence Maximization on Multicore Architectures
- **COMET**: Memory-Efficient Deep Learning Training Framework by Using Error-Bounded Lossy Compression
- **OptHuffDec**: Optimized Huffman Decoder for Error-Bounded Lossy Compression on GPUs
- **TAC**: An Error-Bounded Lossy compressor for Three-Dimensional Adaptive Mesh Refinement (AMR) Data
- **Foresight**: A Compression Benchmark Suite for Visualization and Analysis of Simulation Data
- **Z-checker**: Exascale Scientific Data Analysis and Lossy Compression Assessment Library

Professional Membership

- Member, Association for Computing Machinery (2016–Present)
- Member, Institute of Electrical and Electronics Engineers (2017–Present)
- Member, Society for Industrial and Applied Mathematics (2021–Present)