

Dong Dai, Ph.D

Assistant Professor

Department of Computer Science
University of North Carolina at Charlotte
9201 University City Blvd, Charlotte, NC

Phone: 704-687-1978

Email: dong.dai@uncc.edu

Homepage: <http://webpages.uncc.edu/ddai/>

Research Interests

- I am interested in a wide range of topics related to build and optimize high-performance data-intensive systems, such as job and I/O scheduler, parallel file systems, metadata management, graph storage, and machine learning infrastructure.

Education

- Ph.D. Computer Science, University of Science and Technology of China, 2013.
Thesis: *Research and Implementation on Cloud Software Infrastructure*. Advisor: Prof. Xuehai Zhou.
- B.S. Computer Science, University of Science and Technology of China, 2006.

Professional Experience

- **Assistant Professor** 2018 – current
Computer Science Department, University of North Carolina at Charlotte
- **Research Assistant Professor** 2016 - 2018
Computer Science Department, Texas Tech University
- **Post-doctoral Researcher** 2013 - 2016
Texas Tech University and Argonne National Lab.

Research Projects

- **Active Projects**
 - PI**, National Science Foundation, SHF 2019 – 2022
A Parallel Graph-Based Paradigm for HPC Parallel File System Checkers \$307,682
PI: Mai Zheng (Iowa State University)
 - Single PI**, National Science Foundation, CRII 2018 – 2021
Partitioning Large Graphs in Deep Storage Architecture \$168,201
 - Co-PI**, National Science Foundation, CNS 2018 – 2021
Tuning Extreme-scale Storage Stack through Deep Reinforcement Learning \$240,026
PI: Forrest Sheng Bao (Iowa State University), Yong Chen (Texas Tech University)

Co-PI , National Science Foundation, CCF <i>Uncovering Vulnerabilities in Parallel File Systems for Reliable HPC</i> PI: Yong Chen (Texas Tech University)	2017 – 2021 \$233,000
Co-PI , National Science Foundation, OAC <i>Empowering Data-driven Discovery with a Provenance Collection, Management, and Analysis Software Infrastructure</i> PI: Yong Chen, Co-PI: William Hase, Brian Ancell (Texas Tech University)	2018 – 2021 \$599,982

Selected Publications

Names with (*) are the Ph.D students I mentored; with (†) are the undergraduate students I mentored.

- [1] Di Zhang*, **Dong Dai**, Runzhou Han, Mai Zheng. SentiLog: Anomaly Detecting on Parallel File Systems via Log-based Sentiment Analysis. *Accepted to appear in 13th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage'21), 2021*
- [2] Jiang Zhou, Yong Chen, **Dong Dai**, Yu Zhuang, Weiping Wang. I/O characteristic discovery for storage system optimizations. *Journal of Parallel and Distributed Computing (JPDC'21), Vol 148, Pages 1-13, 2021*
- [3] Di Zhang*, **Dong Dai**, Youbiao He, Forrest Sheng Bao, and Bing Xie. RLScheduler: An Automated HPC Batch Job Scheduler Using Reinforcement Learning. *in the proceeding of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20), 2020.* (acceptance rate: 22.3%).
- [4] Abdullah Al Raqibul Islam*, Anirudh Narayanan, Christopher York†, and **Dong Dai**. A Performance Study of Optane Persistent Memory: From Indexing Data Structures' Perspective. *Accepted to appear in the 36th International Conference on Massive Storage Systems and Technology (MSST'20), 2020.*
- [5] Abdullah Al Raqibul Islam*, and **Dong Dai**. Understand the Overheads of Storage Data Structures on Persistent Memory (Poster). *In 25th ACM SIGPLAN Symposium on Principle and Practice of Parallel Programming (PPoPP'20), 2020.*
- [6] Jiang Zhou, Yong Chen, Wei Xie, **Dong Dai**, Shuibing He, and Weiping Wang. PRS: A Pattern-Directed Replication Scheme for Heterogeneous Object-Based Storage. *IEEE Transactions on Computers (TC'19), 2019.*
- [7] **Dong Dai**, Om Rameshwar Gatla, and Mai Zheng. A Performance Study of Lustre File System Checker: Bottlenecks and Potentials. *Accepted to appear in the proceedings of the 35th International Conference on Massive Storage Systems and Technology (MSST'19), 2019.*
- [8] **Dong Dai**, Yong Chen, Philip Carns, John Jenkins, Wei Zhang, and Robert Ross. Managing Rich Metadata in High-Performance Computing Systems Using a Graph Model. *Accepted to appear in the IEEE Transactions on Parallel and Distributed Systems (TPDS'18), 2018.*
- [9] Jinrui Cao, Om Rameshwar 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. *Accepted to appear in the proceedings of the 32nd ACM/SIGARCH International Conference on Supercomputing (ICS'18), 2018.* (acceptance rate: 18.7%).
- [10] **Dong Dai**, Yong Chen, Philip Carns, John Jenkins, and Robert Ross. Lightweight Provenance Service for High Performance Computing. *In Proceedings of the 26th International Conference on Parallel Architectures and Compilation Techniques (PACT'17), 2017.* (acceptance rate: 23%).

- [11] **Dong Dai**, Wei Zhang, and Yong Chen. IOGP: An Incremental Online Graph Partitioning Algorithm for Distributed Graph Databases. In *Proceedings of the 26th ACM International Symposium on High Performance Parallel and Distributed Computing (HPDC'17)*, 2017. (acceptance rate: 19%).
- [12] **Dong Dai**, Yong Chen, Dries Kimpe, and Robert Ross. Two-Choice Randomized Dynamic I/O Scheduler for Object Storage Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'14)*, 2014. (acceptance rate: 82/394=20.8%).
- [13] **Dong Dai**, Yong Chen, Dries Kimpe, and Robert Ross. Trigger-based Incremental Data Processing with Unified Sync and Async Model. *Accepted to appear in Transaction on Cloud Computing (TCC'18)*
- [14] **Dong Dai**, Forrest Sheng Bao, Jiang Zhou, Xuanhua Shi, and Yong Chen. Vectorizing Disk Blocks for Efficient Storage Systems via Deep Learning. *International Journal of Parallel Computing (ParCo'18)*.
- [15] **Dong Dai**, Phil Carns, Robert Ross, John Jenkins, Nicholas Muirhead, and Yong Chen. An Asynchronous Traversal Engine for Graph-Based Rich Metadata Management. *International Journal of Parallel Computing (ParCo'16)*.

Other Conference/ Journal Publications

- [16] Jiang Zhou, **Dong Dai**, Yu Mao, Xin Chen, Yu Zhuang, and Yong Chen. I/O Characteristics Discovery in Cloud Storage Systems. In *Proceedings of the 11th International Conference on Cloud Computing (CLOUD'18)*, 2018.
- [17] Wei Zhang, **Dong Dai**, and Yong Chen. AKIN: A Streaming Graph Partitioning Algorithm for Distributed Graph Storage Systems. In *Proceedings of the 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'18)*, 2018. (acceptance rate: 20.8%).
- [18] **Dong Dai**, Wei Zhang, and Yong Chen. IOGP: An Incremental Online Graph Partitioning for Large-Scale Distributed Graph Databases. In *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'17)*, 2017. (Short Paper).
- [19] Jiang Zhou, Wei Xie, **Dong Dai**, and Yong Chen. Pattern-Directed Replication Scheme for Heterogeneous Object-based Storage. In *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'17)*, 2017.
- [20] Chao Wang, **Dong Dai**, Xi Li, Aili Wang, and Xuehai Zhou. SuperMIC: Analyzing Large Biological Datasets in Bioinformatics with Maximal Information Coefficient. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB'17)*, 2017.
- [21] **Dong Dai**, Yong Chen, Phil Carns, John Jenkins, Wei Zhang, and Robert Ross. GraphMeta: A Graph-based Engine for Managing Large-Scale HPC Rich Metadata. In *Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER'16)*, 2016. (acceptance rate: 39/162=24%).
- [22] Jinrui Cao, Simeng Wang, **Dong Dai**, Mai Zheng, and Yong Chen. A Generic Framework for Testing Parallel File Systems. In *Proceedings of the Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems held in conjunction with SC'16 (PDSW-DISCS'16)*, 2016.
- [23] **Dong Dai**, Forrest Sheng Bao, Jiang Zhou, and Yong Chen. Block2Vec: A Deep Learning Strategy on Mining Block Correlations in Storage Systems. In *Proceedings of the 9th International Workshop on Parallel Programming Models and Systems Software for High-End Computing held in conjunction with ICPP'16 (P2S2'16)*, 2016.
- [24] **Dong Dai**, Phil Carns, Robert Ross, John Jenkins, Kyle Blauer[†], and Yong Chen. GraphTrek: Asynchronous Graph Traversal for Property Graph Based Metadata Management. In *Proceedings of the IEEE International Conference on Cluster Computing (CLUSTER'15)*, 2015. (acceptance rate: 38/157=24.2%).

- [25] **Dong Dai**, Yong Chen, Dries Kimpe, and Robert Ross. Provenance-Based Object Storage Prediction Scheme for Scientific Big Data Applications. In *Proceedings of the 2014 IEEE International Conference on Big Data (BigData'14)*, 2014. (acceptance rate: 49/264=18.6%).
- [26] **Dong Dai**, Robert Ross, Philip Carns, Dries Kimpe, and Yong Chen. Using Property Graphs for Rich Metadata Management in HPC Systems. In *Proceedings of the 9th Parallel Data Storage Workshop held in conjunction with SC'14 (PDSW'14)*, 2014.
- [27] **Dong Dai**, Xuehai Zhou, Dries Kimpe, Robert Ross, and Yong Chen. Domino: An Incremental Computing Framework in Cloud with Eventual Synchronization. In *Proceedings of the 23rd ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC'14)*, 2014. (Short Paper).
- [28] **Dong Dai**, Xi Li, Chao Wang, Junneng Zhang, and Xuehai Zhou. Detecting Associations in Large Dataset on MapReduce. In *Proceedings of the 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom'13)*, 2013.
- [29] **Dong Dai**, Xi Li, Chao Wang, Mingming Sun, and Xuehai Zhou. Sedna: A Memory Based Key-Value Storage System for Realtime Processing in Cloud. In *Proceedings of the 2012 IEEE International Conference on Cluster Computing Workshops (CLUSTER WORKSHOPS'12)*, 2012.
- [30] **Dong Dai**, Xuehai Zhou, Feng Yang, and Chao Wang. An Auto-configuration Tool for Heterogeneous Hadoop Cluster. *Journal of the Graduate School of Chinese Academy of Sciences*, 2012. (Chinese)

Conference Posters

- [31] Neda Tavakoli, **Dong Dai**, John Jenkins, Philip Carns, Robert Ross, and Yong Chen. A Software-Defined Approach for QoS Control in High-Performance Computing Storage Systems. *SC'16, Salt Lake City, UT*, 2016.
- [32] **Dong Dai**, Robert Ross, Dounia Khaldi, Yonghong Yan, Matthieu Dorier, Neda Tavakoli, and Yong Chen. Exploiting Locality in Scientific Workflow System: A Cross-Layer Solution. *SC'16, Salt Lake City, UT*, 2016.
- [33] **Dong Dai**, Yong Chen, Dries Kimpe, and Robert Ross. Provenance-Based Prediction Scheme for Object Storage System in HPC. *CCGrid'14, Chicago, IL*, 2014.
- [34] Kun Lu, **Dong Dai**, and Mingming Sun. HDFS+: Concurrent Writes Improvements for HDFS. *CC-Grid'13, Delft, Netherlands*, 2013.
- [35] **Dong Dai**, Xi Li, Chao Wang, and Xuehai Zhou. Cloud Based Short Read Mapping Service. *CLUSTER'12, Beijing, China*, 2012.

Teaching

University of North Carolina at Charlotte, Charlotte, NC

- ITSC 3181 Introduction to Computer Architecture 2020 Fall, 2021 Spring
- ITCS 5145 Parallel Computing 2019-21 Spring & Fall
- ITCS 6144/8144 Operating Systems Design 2018-19 Fall

Texas Tech University

- CS4352: Operating Systems 2016
- CS5352: Advanced Operating Systems Design 2017, 2015

- CS5331: Big Data Infrastructure and Data Management 2014

University of Science and Technology of China

- Parallel Algorithm 2012
- Practical Optimization Algorithm Design 2011
- Principles of Computer Organization 2011, 2010

Mentoring Experience

- **REU Undergraduate Students:**
 - At UNCC, mentor UG student Christopher York in a REU supplement award in 2019. Outcomes include a conference publication in MSST'20.
 - At TTU, participate REU'15 and mentor one undergraduate student (Nicholas Muirheada) on HPC metadata management system. Outcomes include a poster, a technical report, and a journal publication in ParCo'16.
- **Graduate Students:**
 - At UNCC, mentor four Ph.D students: Abdullah Al Raqibul Islam, Di Zhang, Md. Hasanur Rashid, and Saisha Kamat. Among them, Saisha Kamat is from under-representative group.
 - At TTU, Mentor two Ph.D. students (Neda Tavakoli and Wei Zhang) on HPC job scheduling and distributed graph storage systems. Outcomes include a poster (SC'16), a workshop paper (P2S2'16), conference papers (HPDC'17, CLUSTER'16), and a journal publication (ParCo'18).

Professional Service

Grant Panel Service

- **Panelist:** National Science Foundation, Computer Systems Research (CSR) 2017, 2018, 2019, 2021

Conference Service

- Program Chair/Co-Chair
 - The 4th Industry/University Joint International Workshop on Data Center Automation, Analytics, and Control, held in conjunction with UCC 2021.
 - The 3rd International Industry/University Workshop on Data-center Automation, Analytic, and Control, held in conjunction with SC'19 (DAAC'19)
 - The 2nd International Industry/University Workshop on Data-center Automation, Analytic, and Control, held in conjunction with SC'18 (DAAC'18)
 - The 5th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT'18) Big Data and HPC Track
 - The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17) Poster Program
 - The 1st International Industry/University Workshop on Data-center Automation, Analytic, and Control, held in conjunction with UCC'17 (DAAC'17)
- Program Committee Member

- The 2021 IEEE International Conference on Big Data (IEEE BigData 2021)
- The 19th IEEE International Symposium on Parallel and Distributed Processing with Applications (IEEE ISPA 2021)
- The 49th International Conference on Parallel Processing (IEEE ICPP'20)
- The 35th IEEE International Parallel & Distributed Processing Symposium (IEEE IPDPS'20)
- The 21th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (IEEE/ACM CCGrid'20)
- International Conference for High Performance Computing, Networking, Storage and Analysis (IEEE/ACM SC'20 Poster)
- The 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom'18)
- The National Cyber Security Summit (NCS'18)
- IEEE International Conference on Big Data (BigData Congress'18)
- Workshop on Advances in Software and Hardware for Big Data to Knowledge Discovery held in conjunction with BigData'14 (ASH'14)
- Session Chair
 - The 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS'20)
 - The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17)
 - The 4th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT'17)
 - International Workshop on Parallel Programming Models and Systems Software for High-End Computing held in conjunction with ICPP'16 (P2S2'16)
 - International Workshop on Data Intensive Scalable Computing Systems held in conjunction with SC'14 (DISCS'14)

Journal Editorial Service

- Editor, *Cluster Computing Special Issue*
- Guest Editor, *Applied Soft Computing (BigData Special Issue)*

Journal Reviewing Service

- Reviewer for IEEE Transactions on Storage 2020, 2021
- Reviewer for IEEE Transactions on Parallel and Distributed Systems 2015-2021
- Reviewer for IEEE Transactions on Cloud Computing 2017, 2015
- Reviewer for IEEE Transactions on Industrial Informatics 2017, 2016
- Reviewer for Applied Soft Computing 2017, 2015
- Reviewer for International Journal of Parallel Programming 2015
- Reviewer for International Journal of High Performance Systems Architecture 2015

Department Service

- Faculty Search Committee, Computer Science Department, UNC-Charlotte, 2020-21.

Professional Memberships

- Institute for Electrical and Electronics Engineers (IEEE), Member, 2012–present
- Association for Computing Machinery (ACM), Member, 2014–present

Last updated: June 13, 2021