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: ddai@uncc.edu
Homepage: http://webpages.uncc.edu/ddai/

Research Interests

I am interested in developing intelligent infrastructure for high-performance and robust data-intensive computing, including I/O scheduling, parallel file systems, metadata management, graph storage, and machine learning infrastructure.

Education

- Ph.D. Computer Science, University of Science and Technology of China, 2013.
- 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.

Selected Publications

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

[1] Abdullah Al Raqibul Islam*, Dong Dai. DGAP: Efficient Dynamic Graph Analysis on Persistent Memory. *Accepted to appear in the International Conference for High Performance Computing, Networking, Storage and Analysis (**SC'23**), 2023.* (acceptance rate: 23%, Conference CORE ***Ranking A***).

[2] Di Zhang*, Chris Egersdoerfer†, Tabassum Mahmud, Mai Zheng, Dong Dai. Drill: Log-based Anomaly Detection for Large-scale Storage Systems Using Source Code Analysis. *Accepted to appear in 37th IEEE International Parallel & Distributed Processing Symposium (**IPDPS'23**), 2023* (acceptance rate: 23%).

[3] Saisha Kamat*, Abdullah Al Raqibul Islam*, Mai Zheng, Dong Dai. FaultyRank: A Graph-based Parallel File System Checker. *Accepted to appear in 37th IEEE International Parallel & Distributed Processing Symposium (**IPDPS'23**), 2023* (acceptance rate: 23%).

[4] Abdullah Al Raqibul Islam*, Dong Dai, A Framework for Large Dynamic Graph Analysis on Persistent Memory, *The 21st USENIX Conference on File and Storage Technologies Work-in-Progress Session (**FAST'23 WiP**), 2023*

[5] Duo Zhang, Om Rameshwar Gatla, Abdullah Al Raqibul Islam*, Dong Dai, Mai Zheng, On the Scalability of Testing the Crash Consistency of PM Systems, *The 21st USENIX Conference on File and Storage Technologies Work-in-Progress Session (**FAST'23 WiP**), 2023*

[6] Chris Egersdoerfer†, Di Zhang*, Dong Dai. ClusterLog: Clustering Logs for Effective Log-based Anomaly Detection. *in the proceeding of 2022 IEEE/ACM 12th Workshop on Fault Tolerance for HPC at eXtreme Scale (**FTXS'22 @ SC'22**), 2022*

[7] Abdullah Al Raqibul Islam*, Christopher York†, Dong Dai. A performance study of Optane persistent memory: from storage data structures' perspective. *CCF Transactions on High Performance Computing (**THPC'22**), 2022*

[8] Di Zhang*, Dong Dai, Bing Xie. SchedInspector: A Batch Job Scheduling Inspector Using Reinforcement Learning. *in proceeding of the 31st International ACM Symposium on High-Performance Parallel and Distributed Computing (**HPDC'22**), 2022* (acceptance rate: 19%).

[9] Abdullah Al Raqibul Islam*, Dong Dai, Dazhao Cheng. VCSR: Mutable CSR Graph Format Using Vertex-Centric Packed Memory Array. *in the proceeding of the 22nd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (**CCGrid'22**), 2022* (acceptance rate: 28%).

[10] Runzhou Han, Om Rameshwar Gatla, Mai Zheng, Jinrui Cao, Di Zhang*, Dong Dai, Yong Chen, Jonathan Cook. A Study of Failure Recovery and Logging of High-Performance Parallel File Systems. *ACM Transactions on Storage 18, no. 2 (2022): 1-44. (**TOS'22**), 2022.*

[11] Dazhao Cheng, Yu Wang, Dong Dai. Dynamic Resource Provisioning for Iterative Workloads on Apache Spark. *IEEE Transactions on Cloud Computing (**TCC'21**), 2021.*

[12] Di Zhang*, Dong Dai, Runzhou Han, Mai Zheng. SentiLog: Anomaly Detecting on Parallel File Systems via Log-based Sentiment Analysis. *in the proceeding of the 13th ACM Workshop on Hot Topics in Storage and File Systems (**HotStorage'21**), 2021* **Best Paper Nominee!**.

[13] 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.*

[14] 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%).

[15] Abdullah Al Raqibul Islam*, Anirudh Narayanan, Christopher York†, and Dong Dai. A Performance Study of Optane Persistent Memory: From Indexing Data Structures' Perspective. *in proceeding of the 36th International Conference on Massive Storage Systems and Technology (**MSST'20**), 2020.*

[16] Abdullah Al Raqibul Islam*, and Dong Dai. Understand the Overheads of Storage Data Structures on Persistent Memory. *in proceeding of the 25th ACM SIGPLAN Symposium on Principle and Practice of Parallel Programming (**PPoPP'20 Poster**), 2020.*

[17] 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.*

[18] Dong Dai, Om Rameshwar Gatla, and Mai Zheng. A Performance Study of Lustre File System Checker: Bottlenecks and Potentials. *in proceedings of the 35th International Conference on Massive Storage Systems and Technology (**MSST'19**), 2019.*

[19] 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. *IEEE Transactions on Parallel and Distributed Systems (**TPDS'18**), 2018.*

[20] 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. *in proceedings of the 32nd ACM/SIGARCH International Conference on Supercomputing (**ICS'18**), 2018.* (acceptance rate: 18.7%).

[21] 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%).

[22] 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*.

[23] 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%, Conference CORE ***Ranking B***).

[24] 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%).

[25] 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.* (Conference CORE ***Ranking A***)

[26] 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: 24%).

[27] 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: 24%).

[28] 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: 20.8%).

[29] 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: 18.6%, ***Ranking B***).

Professional Service

Grant Panel Service

- **Panelist**: National Science Foundation, Computer Systems Research (CSR) 2017-2019, 2021-2023

- **Panelist**: National Science Foundation, Office of Advanced Cyberinfrastructure (OAC) 2022
- **Panelist**: National Science Foundation, Small Business Innovation Research (SBIR) 2021, 2022

Conference Service

- Program Chair/Co-Chair
  - *Digital Conference Organization Chairs*: The 14th ACM International Conference on Future Energy Systems, ACM Eenergy 2023.
  - *Finance Chair*: The 33rd IEEE International Symposium on Software Reliability Engineering. ISSRE 2022.
  - *Program Co-Chair*: The 3rd International Industry/University Workshop on Data-center Automation, Analytic, and Control (DAAC'19 @ SC'19)
  - *Program Co-Chair*: The 2nd International Industry/University Workshop on Data-center Automation, Analytic, and Control (DAAC'18 @ SC'18)
  - The 5th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT'18) Big Data and HPC Track
  - *Poster Program Chair*: The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17) Poster Program
  - *Program Co-Chair*: The 1st International Industry/University Workshop on Data-center Automation, Analytic, and Control, held in conjunction with UCC'17 (DAAC'17)
- Program Committee Member
  - International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20 Poster Committee, SC'22 Workshop Committee)
  - IEEE International Parallel & Distributed Processing Symposium (IPDPS 2020, 2022, 2023)
  - International Conference on Parallel Processing (ICPP 2020, 2022, 2023)
  - IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid 2020, 2022)
  - IEEE International Conference on Big Data (BigData 2021, 2022, 2023)
  - IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA 2021, 2022)

Journal Editorial Service

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

Journal Reviewing Service

- Reviewer for IEEE Transactions on Computers 2022
- Reviewer for IEEE Transactions on Storage 2020, 2021, 2022
- Reviewer for IEEE Transactions on Parallel and Distributed Systems 2015-2023
- Reviewer for IEEE Transactions on Cloud Computing 2015,2017
- 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

Last updated: October 14, 2023