

Dayi Fan

PERSONAL INFO

Email: fan.1090@osu.edu
Homepage: <https://davis-fan.github.io/>
GitHub: <https://github.com/Davis-Fan>
Address: Columbus, Ohio, U.S.

RESEARCH INTEREST

My current work focuses on parallel and high-performance computing, algorithm engineering, and data management for large-scale systems. I have developed state-of-the-art parallel algorithms for three fundamental problems: tree edit distance, graph maximum matching, and bipartite weighted matching. By combining novel algorithmic designs with optimized implementations on multicore CPUs and GPUs, my research aims to accelerate irregular data analytics workloads through scalable parallelization.

My research agenda also extends to system-level challenges in AI applications. I am particularly interested in the use of LLM-based agents for parallel programming, where multiple agents collaborate on algorithm design, code generation, correctness testing, and benchmark-driven performance tuning. More broadly, I hope to explore how AI systems can support algorithm engineering for irregular workloads, including graph analytics, dynamic programming, and combinatorial optimization, by automatically exploring data structures, parallel strategies, and hardware-specific optimizations. I also have hands-on experience with PyTorch, Transformers, and Hugging Face libraries.

EDUCATION

- | | |
|-------------------------------|--|
| 2022.01-2026.12
(expected) | Ph.D. Candidate in Computer Science and Engineering,
The Ohio State University, U.S.
Advisor: Prof. Xiaodong Zhang
Topic: Parallel Algorithm Design and Hardware Acceleration for
Irregular Data-Centric Applications |
| 2021.01-2021.05 | Undergraduate Exchange Student in Computer Science,
National University of Singapore, Singapore
Project: LSTM Implementation with Spiking Neural Networks |

2017.09-2021.07 B.Eng. in Microelectronics Science and Engineering,
Southern University of Science and Technology, China
Project: ReRAM-based Neural Network Optimization

CONFERENCE PAPER

- 2026 **X-Wim: Massive Parallelization of Weighted Matching in Bipartite Graphs**
Dayi Fan, Simon Zhang, Rubao Lee, Hanqi Guo, Xiaodong Zhang
International Conference on Very Large Data Bases (VLDB 2026) [[PDF](#)][[Code](#)]
- Developed the first large-scale parallel framework for bipartite weighted matching; proposed phase-decoupled sequential and parallel Hungarian algorithms with optimized multicore implementations
 - Enables efficient large-scale assignment and matching for cloud resource allocation, multi-robot coordination, and marketplace platforms
- 2025 **X-Blossom: Massive Parallelization of Graph Maximum Matching**
Dayi Fan, Rubao Lee, Xiaodong Zhang
International Conference on Very Large Data Bases (VLDB 2025) [[PDF](#)][[Code](#)]
- Developed the first large-scale parallel framework for maximum matching; proposed recursion-free sequential and parallel Blossom algorithms that eliminate dynamic graph contraction, implemented on multicore CPUs
 - Supports high-performance graph applications in financial analysis, social platforms, and recommendation systems
- 2024 **X-TED: Massive Parallelization of Tree Edit Distance**
Dayi Fan, Rubao Lee, Xiaodong Zhang
International Conference on Very Large Data Bases (VLDB 2024) [[PDF](#)][[Code](#)]
- Developed the first large-scale parallel framework for tree edit distance; proposed a novel algorithm that breaks complex dynamic programming dependencies, with optimized multicore CPU and GPU implementations
 - Facilitates efficient tree-structured data comparison in LLM evaluation workflows, code analysis, and hierarchical data processing

JOURNAL ARTICLE

2024 RR-Compound: RDMA-Fused gRPC for Low Latency, High Throughput, and Easy Interface

Liang Geng, Hao Wang, Jingsong Meng, **Dayi Fan**, Sami Ben-Romdhane, Hari Kadayam Pichumani, Vinay Phegade, Xiaodong Zhang

IEEE Transactions on Parallel and Distributed Systems (TPDS 2024) [[PDF](#)]

- Integrated gRPC with RDMA to reduce latency and improve throughput; Validated the system on key-value storage and TensorFlow workloads

PREPRINTS

2026 Performance Insights and Implications of Traversal-based Graph Analytics on both Multi-core CPUs and GPUs

Jiaxin Liu, **Dayi Fan**, Rubao Lee, Cathy Xia, Xiaodong Zhang

Revision invited, SIGMETRICS 2027

ACADEMIC SERVICE

Shadow Program Committee Reviewer, VLDB 2027

Reviewer, The Journal of Supercomputing, 2026

EXPERIENCE

2022.01-CURRENT Graduate Research Associate
The Ohio State University

2022.01-2024.05 Graduate Teaching Associate
The Ohio State University
CSE 3421 Computer Architecture
OSU C4 Oral English Proficiency Certification

SKILLS

Programming Languages: C++, CUDA, Python, SQL

Machine Learning: PyTorch, Transformers, Hugging Face, OpenEvolve

Systems/HPC: Multithreading, GPU Programming, System Profiling, gRPC, RDMA

ACADEMIC PERFORMANCE

Postgraduate GPA: 3.96/4.00

Undergraduate GPA: 3.90/4.00

2021.06 Excellent Graduation Award
Southern University of Science and Technology

2017-2020 First Class Student Scholarship (Top 5%, three consecutive years)
Southern University of Science and Technology