

JIE LI

Post-doctoral Researcher, Department of Computer Science, Texas Tech University

Email: jie[dot]li[at]ttu[dot]edu | Homepage: <https://lijie.me>

RESEARCH INTERESTS

My research focuses on high-performance computing (HPC), parallel and distributed systems, and computer architecture. I develop intelligent, scalable, and energy-efficient computing systems to support data-intensive science, AI applications, and next-generation HPC workloads. My work integrates system monitoring, adaptive scheduling, and architectural co-design to improve the performance, resilience, and sustainability of large-scale computing platforms.

EDUCATION

- **Ph.D.**, Computer Science, Texas Tech University, Lubbock, Texas 2024
Dissertation: *Optimizing High-Performance Computing Systems: Insights from System Monitoring, Workload Management, and Scheduling Strategies*
- **M.S.**, Computer Science, Texas Tech University, Lubbock, Texas 2019
Thesis: *PIMS: A Lightweight Processing-in-Memory Accelerator for Stencil Computations*
- **B.A.**, Architecture, Huaqiao University, Xiamen, China 2012
Introduced to programming through computational design and parametric modeling coursework.

PROFESSIONAL EXPERIENCE

- **Post-doctoral Researcher**, Texas Tech University 2024–present
 - Lead research in high-performance computing; continue publishing in top-tier venues.
 - Co-author competitive research proposals, including an NSF grant as Co-PI.
 - Support operations of the Data-Intensive Scalable Computing Laboratory (DISCL) and the NSF Cloud and Autonomic Computing Center (CAC IUCRC), including organizing Semi-Annual Industry Advisory Board Meetings and coordinating internal activities.
 - Supervise student researchers and lead collaborative research software development.
- **Research Assistant**, Texas Tech University 2019–2024
 - Conducted original research in HPC, parallel and distributed systems, and computer architecture.
 - Published in peer-reviewed venues, including ICPP, MemSys, BigData, CLUSTER, ISC, and CLOUD.
 - Mentored graduate and undergraduate students in research and project development, including supervision of a master's thesis published at CLOUD'23.
 - Developed research software and contributed to open-source projects, including the HPC system monitoring tool MonSter, disaggregated memory scheduler, and xBGAS simulation tools. MonSter code was partially adopted by the Dell Omnia project.
 - Administered and maintained HPC servers, ensuring system reliability, configuration, and software availability.
- **Graduate Student Intern**, Lawrence Berkeley National Laboratory Summers 2021, 2022, 2023
 - Designed and implemented data pipelines to integrate and analyze HPC telemetry (including LDMS, DCGM, and Slurm) from NERSC systems.
 - Applied machine learning and deep learning techniques (e.g., convolutional neural networks) to classify and predict HPC job behavior from time-series data.

- Authored and co-authored peer-reviewed publications in top venues such as ISC and IEEE CLUSTER, showcasing results from independent research conducted during internships.
- Led the development of DisaggregationAwareScheduler, an open-source simulation framework for job scheduling in disaggregated memory systems; used in multiple publications and ongoing research collaborations.

GRANTS AND FUNDING

- **CICI:UCSS:SHIELD – Strengthening High-Performance Infrastructure with Enhanced Layered Defense Role:** *Co-Principal Investigator (with PI Dr. Yong Chen)*. Contributed to the project vision and technical approach for enhancing cybersecurity in HPC environments.

Submitted: 2025 **Status:** Pending **Sponsor:** National Science Foundation **Total Requested:** \$600,000

- **Category II: REPACSS – Empowering Scientific Discovery through Renewable Energy Powered Advanced Computing Systems and Services**

Role: *Contributor to Proposal Development (with PI Dr. Yong Chen and Co-PI Dr. Alan Sill)*. Assisted in proposal writing, including sections on data center remote control, monitoring, visualization, and analytics.

Submitted: 2023 **Status:** Funded **Sponsor:** National Science Foundation **Total Award:** \$12,250,000

- **Frameworks: DAVinci – An Integrated Data Collection, Automation, and Visualization Framework for HPC Systems**

Role: *Lead Contributor (with PI Dr. Yong Chen and Co-PI Dr. Alan Sill)*. Led the development of the proposal's core framework design and research methodology.

Submitted: 2020 **Status:** Unfunded **Sponsor:** National Science Foundation **Total Requested:** \$1,000,000

PUBLICATIONS

- [15] **Jie Li**, George Michelogiannakis, Samuel Maloney, Brandon Cook, Estela Suarez, John Shalf, and Yong Chen. Job scheduling in high performance computing systems with disaggregated memory resources. In *2024 IEEE International Conference on Cluster Computing (CLUSTER'24)*, pages 297–309. IEEE, 2024c. doi: 10.1109/CLUSTER59578.2024.00033
- [14] **Jie Li**, George Michelogiannakis, Brandon Cook, John Shalf, and Yong Chen. Scheduling and allocation of disaggregated memory resources in hpc systems. In *2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 1202–1203. IEEE, 2024b. doi: 10.1109/IPDPSW63119.2024.00206
- [13] **Jie Li**, John D Leidel, Brian Page, and Yong Chen. Towards cycle-accurate simulation of xbgas. In *2024 International Conference on Computing, Networking and Communications (ICNC'24)*, pages 468–472. IEEE, 2024a. doi: 10.1109/ICNC59896.2024.10556078
- [12] Tommy Dang, Ngan VT Nguyen, **Jie Li**, Alan Sill, and Yong Chen. Spiro: Order-preserving visualization in high performance computing monitoring. In *International Symposium on Visual Computing (ISVC'24)*, pages 109–120. Springer, 2023
- [11] **Jie Li**, Rui Wang, Ghazanfar Ali, Tommy Dang, Alan Sill, and Yong Chen. Workload failure prediction for data centers. In *2023 IEEE 16th International Conference on Cloud Computing (CLOUD'23)*, pages 479–485, 2023b. doi: 10.1109/CLOUD60044.2023.00064
- [10] Cristiano E. Caon, **Jie Li**, and Yong Chen. Effective management of time series data. In *2023 IEEE 16th International Conference on Cloud Computing (CLOUD'23)*, pages 408–414, 2023. doi: 10.1109/CLOUD60044.2023.00055

- [9] **Jie Li**, George Michelogiannakis, Brandon Cook, Dulanya Cooray, and Yong Chen. Analyzing resource utilization in an hpc system: A case study of nersc’s perlmuter. In *International Conference on High Performance Computing (ISC’23)*, pages 297–316. Springer, 2023a. doi: 10.1007/978-3-031-32041-5_16
- [8] Tommy Dang, Ngan VT Nguyen, **Jie Li**, Alan Sill, Jon Hass, and Yong Chen. Jobviewer: Graph-based visualization for monitoring high-performance computing system. In *2022 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT’22)*, pages 110–119. IEEE, 2022. doi: 10.1109/BDCAT56447.2022.00021
- [7] Tommy Dang, Ngan Nguyen, Jon Hass, **Jie Li**, Yong Chen, and Alan Sill. The gap between visualization research and visualization software in high-performance computing center. *The Gap between Visualization Research and Visualization Software (VisGap’21)*, 2021. doi: 10.2312/visgap.20211089
- [6] Xi Wang, Antonino Tumeo, John D Leidel, **Jie Li**, and Yong Chen. Ham: Hotspot-aware manager for improving communications with 3d-stacked memory. *IEEE Transactions on Computers (IEEE Trans Comput)*, 70(6): 833–848, 2021. doi: 10.1109/TC.2021.3066982
- [5] **Jie Li**, Ghazanfar Ali, Ngan Nguyen, Jon Hass, Alan Sill, Tommy Dang, and Yong Chen. Monster: an out-of-the-box monitoring tool for high performance computing systems. In *2020 IEEE International Conference on Cluster Computing (CLUSTER’20)*, pages 119–129. IEEE, 2020. doi: 10.1109/CLUSTER49012.2020.00022
- [4] Ngan Nguyen, Jon Hass, Yong Chen, **Jie Li**, Alan Sill, and Tommy Dang. Radarviewer: visualizing the dynamics of multivariate data. In *Practice and Experience in Advanced Research Computing (PEARC’20)*, pages 555–556. 2020. doi: 10.1145/3311790.3404538
- [3] Vung Pham, Ngan Nguyen, **Jie Li**, Jon Hass, Yong Chen, and Tommy Dang. Mtsad: Multivariate time series abnormality detection and visualization. In *2019 IEEE International Conference on Big Data (BigData’19)*, pages 3267–3276. IEEE, 2019. doi: 10.1109/BigData47090.2019.9006559
- [2] **Jie Li**, Xi Wang, Antonino Tumeo, Brody Williams, John D Leidel, and Yong Chen. Pims: A lightweight processing-in-memory accelerator for stencil computations. In *Proceedings of the International Symposium on Memory Systems (MemSys’19)*, pages 41–52, 2019. doi: 10.1145/3357526.3357550
- [1] Xi Wang, Antonino Tumeo, John D Leidel, **Jie Li**, and Yong Chen. Mac: Memory access coalescer for 3d-stacked memory. In *Proceedings of the 48th International Conference on Parallel Processing (ICPP’19)*, pages 1–10, 2019. doi: 10.1145/3337821.3337867

OPEN-SOURCE SOFTWARE

- **MonSter**: <https://github.com/nsfcac/MonSter>. An “out-of-the-box” monitoring tool for high-performance computing platforms. Presented in CLUSTER’20 and partially adopted by the Dell Omnia project.
- **DisaggregationAwareScheduler**: <https://github.com/artlands/DisaggregationAwareScheduler>. Simulator for exploring job scheduling strategies in HPC systems with disaggregated memory. Used in CLUSTER’24 and IPDPSW’24 publications.
- **xBGAS REV-CPU**: <https://github.com/artlands/rev-xbgas>. Cycle-accurate simulation extension for xBGAS architecture, based on the REV-CPU framework. Developed in collaboration with Tactical Computing Laboratories (TCL); presented in ICNC’24.
- **xBGAS Runtime**: <https://github.com/artlands/rev-xbgas-runtime>. Lightweight runtime system supporting global address space extensions for RISC-V-based xBGAS simulation. Developed in collaboration with TCL.

PRESENTATIONS

Unless noted otherwise, all presentations were delivered by Jie Li.

- [9] “Integrated Data Collection and Visualization Framework for Data Centers based on Telemetry Model,” by J. Li, P. Roemsri, T. Dang, Y. Chen, and A. Sill. Fall 2024 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, December 16, 2024
- [8] “Towards Cycle-Accurate Simulation of xBGAS,” by J. Li and Y. Chen. Latch-Up 2024, Cambridge, MA, USA, April 19, 2024
- [7] “Workload Failure Prediction for Data Centers,” by J. Li, R. Wang, G. Ali, T. Dang, A. Sill, and Y. Chen. The 16th IEEE International Conference on Cloud Computing (CLOUD’23), Chicago, IL, July 6, 2023
- [6] “Integrated Data Collection and Visualization Framework for Data Centers based on Telemetry Model,” by J. Li, N. Ngan VT, G. Ali, T. Dang, Y. Chen, and A. Sill. Spring 2023 NSF CAC Industry Advisory Board Conference, Tucson, AZ, USA, May 25, 2023
- [5] “Integrated Data Collection and Visualization Framework for Data Centers based on Telemetry Model,” by J. Li, N. Ngan VT, G. Ali, C. Caon, T. Dang, Y. Chen, and A. Sill. Fall 2022 NSF CAC Industry Advisory Board Conference, Denton, TX, USA, November 11, 2022
- [4] “Integrated Data Collection and Visualization Framework for Data Centers based on Telemetry Model,” by J. Li, N. Ngan VT, G. Ali, T. Dang, Y. Chen, and A. Sill. Spring 2021 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, May 7, 2021.
- [3] “Advanced Visualization and Data Analysis of HPC Cluster and User Application Behavior,” by J. Li, N. Ngan VT, G. Ali, T. Dang, Y. Chen, and A. Sill. The 33rd ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC’21), November 2021
- [2] “MonSTer: An Out-of-the-Box Monitoring Tool for High Performance Computing Systems,” by J. Li, G. Ali, N. Ngan VT, J. Hass, A. Sill, T. Dang, and Y. Chen. The 22nd IEEE International Conference on Cluster Computing (CLUSTER’20), 2020.
- [1] “PIMS: A Lightweight Processing-In-Memory Accelerator for Stencil Computations,” by J. Li, X. Wang, A. Tumeo, B. Williams, J. D. Leidel, and Y. Chen. The 5th International Symposium on Memory Systems (MemSys’19), October 2019

REFEREED POSTERS/POSTER PAPERS

- [3] **J. Li**, B. Cook, G. Michelogiannakis, Y. Chen. A Holistic View of Memory Utilization on Perlmutter. In the 34th International Conference for High Performance Computing, Networking, Storage and Analysis (SC’22), 2022.
- [2] **J. Li**, B. Cook, Y. Chen. Detecting and Identifying Applications by Job Signatures. In the 33rd International Conference for High Performance Computing, Networking, Storage and Analysis (SC’21), 2021
- [1] X. Wang, **J. Li**, A. Tumeo, J. D. Leidel, Y. Chen. Memory Hotspot Optimizations for 3D-Stacked Memory. The 28th International Conference on Parallel Architectures and Compilation Techniques (PACT’19), 2019.

NON-REFEREED POSTERS

- [8] M. Side, B. Williams, **J. Li**, J. Leidel, and Y. Chen. A Secure Global Address Space Extension for HPC. In the Fall 2024 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, November 16th, 2024.

- [7] Y. Zhao, **J. Li**, B. Williams, J. Leidel, and Y. Chen. xBGAS subset communication and benchmarking. In the Fall 2024 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, November 16th, 2024.
- [6] M. Side, B. Williams, **J. Li**, J. Leidel, and Y. Chen. A Secure Global Address Space Extension for HPC. In the Fall 2024 Departmental Research Showcase held in conjunction with the TTU WCOE Computer Science Board Meeting, Lubbock, TX, USA, October 17th, 2024.
- [5] **J. Li**, G. Michelogiannakis, B. Cook, and Y. Chen. Job Scheduling in HPC Systems with Disaggregated Memory. In the Fall 2023 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, November 20th, 2023.
- [4] **J. Li**, G. Michelogiannakis, B. Cook, and Y. Chen. A Holistic View of Memory Utilization on Perlmutter. In the Fall 2022 NSF CAC Industry Advisory Board Conference, Denton, TX, USA, November 11th, 2022.
- [3] **J. Li**, B. Cook, and Y. Chen. HPC Application Recognition Through Image-encoded Monitoring Data. In the Spring 2022 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, April 21st, 2022.
- [2] **J. Li**, R. Wang, G. Ali, and Y. Chen. Predicting Abnormal Workloads in HPC Systems. In the Spring 2021 NSF CAC Industry Advisory Board Conference, Lubbock, TX, USA, May 7th, 2021.
- [1] **J. Li**, H. Nguyen, A. Sill, T. Dang and Y. Chen. Behavior Based Job Classification and Characterization on HPC Systems. In the Fall 2021 NSF CAC Industry Advisory Board Conference, Tucson, AZ, USA, November 11th, 2021.

MENTORING & TEACHING EXPERIENCE

Undergraduate Students (including REU participants)

- Mentoring *three* students on Independent Study (CS4000). Spring 2021, Spring2022

Graduate Students

- Mentoring *one* student on Independent Study (CS7000) and Master's Thesis. Fall 2021-Fall 2022
- Mentoring *four* students on course projects of Advanced Operating System (CS5379). Spring 2020, Spring 2021

Teaching Philosophy

Committed to student-centered teaching that bridges computer science fundamentals with practical high-performance computing experience. I aim to foster curiosity, critical thinking, and hands-on skill development through mentorship, collaboration, and real-world system engagement.

AWARDS AND HONORS

- Best Poster Award, NSF Cloud and Autonomic Computing Industry Advisory Board Conference 2022
- Summer Thesis/Dissertation Research Award (\$2300), Lubbock, Texas 2019

PROFESSIONAL SERVICE

Paper Reviewer

- The Journal of Supercomputing
- IEEE International Parallel and Distributed Processing Symposium (IPDPS'23)
- IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGrid'22)
- IEEE International Conference on Distributed Computing Systems (ICDCS'22)
- The International Conference for High-Performance Computing, Networking, Storage, and Analysis (SC'22)
- International Parallel Data Systems Workshop (PDSW'22)
- IEEE International Conference on Big Data (20-22)

- IEEE International Conference on Smart Data Services (SDMS'20)

Conference Volunteer

- Student volunteer of SC'21, St. Louis, Missouri 2021
- Student volunteer of SC'19, Denver, Colorado 2019

SYNERGISTIC ACTIVITIES

- **Contributions to Open-Source High-Performance Computing Software and Tools:** Developed and contributed to several open-source HPC monitoring and visualization tools, including *MonSter* (CLUSTER'20), *JobViewer* (BDCAT'22), and *Spiro* (ISVC'24). These tools have been adopted by Dell EMC in products such as Dell Omnia.
- **Active Dissemination of Research Results:** Regularly presented research findings at major conferences such as IEEE CLUSTER, ISC High Performance, and IEEE CLOUD.
- **Undergraduate and Graduate Research Mentorship and Outreach:** Mentored undergraduate and graduate students in research projects, independent studies, and thesis supervision— including underrepresented minority students—resulting in publications at venues like CLOUD'23.
- **Professional Service and Community Engagement:** Promoted knowledge exchange and community development through active reviewing for top-tier venues and volunteer service at major HPC events like SC (2019, 2021).
- **Industry Engagement:** Supported NSF CAC IUCRC operations and organized Semi-Annual Industry Advisory Board Meetings. Facilitated collaboration with Dell EMC, Lancium, and the NSA through industry-sponsored R&D.