

JIE LI

Email: jie[dot]li[at]ttu[dot]edu ◊ Tel:(+1) 806 787 8258

Data-Intensive Scalable Computing Laboratory (DISCL) ◊ Texas Tech University

Homepage ◊ Google Scholar

EDUCATION

Texas Tech University PhD candidate	Sept. 2019 - Expected Dec. 2023 <i>Lubbock, TX</i>
Texas Tech University Master of Science in Computer Science	Sept. 2017 - Aug. 2019 <i>Lubbock, TX</i>
Huaqiao University Bachelor of Architecture	Sept. 2006 - Jun. 2012 <i>Xiamen, CHN</i>

EXPERIENCE

Data-Intensive Scalable Computing Laboratory (DISCL) <i>Research Assistant</i>	Sept. 2019 - Present <i>Lubbock, TX</i>
<ul style="list-style-type: none">• Conduct research in the areas of High-Performance Computing, Computer Architecture, and Parallel and Distributed Computing. This includes designing and implementing experiments, analyzing and interpreting data, and writing research papers for publication in academic conferences and journals.• Attend conferences, workshops, and seminars to stay up-to-date with the latest research developments and technologies, including presenting research papers and posters, and networking with other researchers in the field.• Participate in the development and maintenance of research software and tools. This involves writing and testing code, documenting software features and functionality, and contributing to open-source software projects.• Mentor graduate and undergraduate students on their independent studies. This involves providing research topics, guidance on research projects, assisting with data collection and analysis, and helping students to develop their technical writing and presentation skills.• Administer two high-end servers (Hugo and Alita) hosted in the High-Performance Computing Center at Texas Tech University. This includes configuring and maintaining software, troubleshooting issues, and ensuring the availability and reliability of the servers.	

National Energy Research Scientific Computing Center (NERSC) <i>Graduate Student Intern</i>	Jun.2022-Aug.2022 <i>Berkeley, CA (remote)</i>
<ul style="list-style-type: none">• Simplified the code structure and create a more efficient and streamlined codebase by refactoring and combining the data collection codes used to access the system monitoring data collected from NERSC's Perlmutter.• Analyzed the system monitoring data on a large scale to evaluate resource utilization by examining metrics such as CPU and GPU utilization, host DRAM utilization, and GPU HBM2 utilization. Identified trends and patterns in the data to gain insights into system performance.• Summarized the analysis and write a system resource analysis paper for ISC 2023.	

National Energy Research Scientific Computing Center (NERSC) <i>Graduate Student Intern</i>	Jun.2021-Aug.2021 <i>Berkeley, CA (remote)</i>
<ul style="list-style-type: none">• Integrated data from multiple sources to analyze system-wide architectural efficiency and workload patterns.• Conducted statistical analysis of job-level monitoring data and applied various machine learning models (e.g., Random Forests, Support Vector Classification, LinearSVC) to classify jobs based on extracted time-series features.• Developed a novel approach to encoding time-series monitoring data as images and trained a Convolutional Neural Network (CNN) to classify job signatures with high accuracy.	

RESEARCH EXPERIENCE

Collecting and Storing Telemetry Metrics of HPC clusters Sept. 2019-Sept. 2021

<https://github.com/nsfcac/MonSter>

- Investigated the Integrated Dell Remote Access Controller (iDRAC) and explored pull model and push model (subscription) for obtaining remote monitoring data.
- Developed a set of tools for configuring iDRAC9 telemetry reports and initializing tables in TimescaleDB.
- Implemented the code to asynchronously collect telemetry metrics from 240 CPU nodes and 20 GPU nodes via Server-Sent Event (SSE), collected monitoring metrics from 467 CPU nodes via pull model, and utilized the resource manager (e.g. Slurm, UGE) interface to collect job accounting information.

Improving the Monitoring Data Querying Performance Jan. 2020-Apr. 2020

<https://github.com/nsfcac/MetricsBuilder>

- Optimized the database schemas in InfluxDB, reducing **71.98%** of data volume and gaining up to **1.76×** performance boost.
- Developed a middleware between the consumers (i.e. visualization and automation components) and the producers (i.e. databases) called Metrics Builder.
- Metrics Builder hides the details of querying databases, speed up querying and transmitting, and provide a unified API (based on OpenAPI) and dataframe for data analysis consumers.
- Metrics Builder improves query performance by up to **25×** and reduces the volume of data transmitted by **95%**.

Processing-in-Memory Accelerator for Stencil Computations Mar. 2019-July 2019

- Investigated the cache performance of stencil computation; implemented a simulator generating the memory traces of different types of stencils.
- Designed a Processing-In-Memory (PIM) accelerator based on the Hybrid Memory Cube (HMC) architecture.
- Evaluated the proposed accelerator using HMC-Sim library. The results showed that the proposed accelerator reduces **48.25%** of data movement on average and obtains up to **65.55%** of bank conflict reduction.

PUBLICATIONS

- [1] Tommy Dang, Ngan VT Nguyen, **Jie Li**, Alan Sill, Jon Hass, Yong Chen. "JobViewer: Graph-based Visualization for Monitoring High-Performance Computing System" Paper In *2022 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT)*, pp. 110-119. IEEE. 2022
- [2] **Jie Li**, George Michelogiannakis, Brandon Cook and Yong Chen. "A Holistic View of Memory Utilization on Perlmutter" Poster In *The International Conference for High Performance Computing, Networking, Storage, and Analysis*, 2022.
- [3] **Jie Li**, Brandon Cook and Yong Chen. "HPC Application Recognition Through Image-encoded Monitoring Data" Poster In *NSF Cloud and Autonomic Computing Industry Advisory Board Conference*, 2022, **Best Poster Award**.
- [4] **Jie Li**, Brandon Cook and Yong Chen. "Detecting and Identifying Applications by Job Signatures in HPC Systems." Poster In *The International Conference for High Performance Computing, Networking, Storage, and Analysis*, 2021.
- [5] Xi Wang, Antonino Tumeo, John D. Leidel, **Jie Li**, and Yong Chen. "HAM: Hotspot-Aware Manager for Improving Communications With 3D-Stacked Memory" in *IEEE Transactions on Computers* 70, no. 6 (2021): 833-848.

- [6] **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)*, pp. 119-129. IEEE, 2020.
- [7] Ngan Nguyen, Jon Hass, Yong Chen, **Jie Li**, Alan Sill, and Tommy Dang. "RadarViewer: Visualizing the dynamics of multivariate data" In *PEARC '20: Practice and Experience in Advanced Research Computing*, pp. 555-556. ACM, 2020.
- [8] 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*, pp. 3267-3276. IEEE, 2019.
- [9] **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*, pp. 41-52. ACM, 2019.
- [10] 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*, p. 2. ACM, 2019.