Feng Li

755 Lockefield ST APT C, Indianapolis, IN 46202

https://fengggli.github.io/ | 317-527-0028 | fengggli@yahoo.com

Education

Purdue University — Ph.D. IN, US Aug 2015 — Aug 2023

Major: Computer Science Research Areas: *high-performance computing, scientific workflows*

Dissertation title: *Efficient in-situ workflows for time-critical applications on heterogeneous ecosystems*

Huazhong University of Science and Technology — B.Eng. Wuhan, China Aug 2011 — June 2015

Major: Computer Science and Technologies GPA: 3.8/4.0

Skills & Abilities

- Programming languages (C/C++/Python/Java/Scala)
- HPC toolchains (MPI, OpenMP, CUDA, Slurm)
- Performance analysis tools (Intel VTune, Linux Perf, TAU)
- Cloud/container solutions (Google GCP, Amazon AWS, OpenStack, Docker, Kubernetes)
- Big Data/Deep learning frameworks (Tensorflow/Pytorch/Apache Spark)
- Storage related: Redis, spdk, dpdk, pmdk, fuse, fio
- CI/Build tools(cmake, Apache Maven, Travis CI, Google Gtest, Spack)

Work Experience

Research Intern — IBM Research Almaden, CA May 2019 — Aug 2019

- Designed and implemented a unified file system interface (KVFS) for multiple key-value store backends, so that file operations are translated into key-value store put/get operations.
- Used FUSE (Filesystem in Userspace) to implement KVFS, and designed mechanisms to handle the mappings between file abstractions and data objects.
- Code base in IBM Comanche: https://github.com/IBM/comanche/tree/unstable/src/fuse.

Research Intern — IBM Research Almaden, CA May 2018 — Aug 2018

- Designed and implemented NVMeStore, which is a key-value store that uses NVMe SSDs as data storage while keeping critical metadata in persistent memory.
- Added continuous integration for the Comanche project using Linux containers and Travis-CI.

Research Intern — IBM Research Almaden, CA May 2017 – Aug 2017

- Designed and implemented an NVMe-backed light-weight memory service — CO-PAGER (Collaborative Paging). A paper [3] is published in HP3C'19.
- CO-PAGER captures virtual memory page faults and performs paging operations on NVMe SSDs using fast userspace I/O.

Selected Research Projects

(More projects are described on my homepage at https://fenggli.github.io/)

Cross-environment in-situ workflow management Jan 2020 — Now

- Design, and prototype a framework that can be used to launch in-situ workflows across HPC and Cloud systems. Work published in **PASC'21** [2].
- Formalize the resource planning problem for in-situ workflows, design and evaluate heuristic-based algorithms to improve workflow metrics such as throughput and resource efficiency. Under review.

Performance Analysis of In-situ Methods in HPC April 2017 — Nov 2017

- Design and implement the first-of-its-kind in-situ workflow benchmark that supports flexible configurations of HPC workflows, transport libraries, and execution environments.
- Artifacts are available at: https://github.com/IUPU-HPC/workflow-bench.
- Work published in **HPDC'18** [6].

Services & Awards

- Reviewer: BIGCOM'16, BDCAT'18, PEARC'20, PEARC'21, Springer JCST (2022).
- Best Student Paper Award: PEARC'17.
- Student volunteer, SC'21.

Publications

1. **Feng Li**, and Fengguang Song. 2023. INSTANT: A Runtime Framework to Orchestrate In-Situ Workflows. To appear in Proceedings of the 27th International European Conference on Parallel and Distributed Computing (EuroPar'23), Springer.
2. **Feng Li**, and Fengguang Song. 2023. Efficient In-situ Workflow Planning for Geographically Distributed Heterogeneous Environments. To appear in Future Generation Computer Systems, Elsevier.
3. Ranran Chen, **Feng Li**, Drew Bieger, Fengguang Song, Yao Liang, Daniel Luna, Ryan Young, Xu Liang, and Sudhakar Pamidighantam. 2022. CyberWater: An Open Framework for Data and Model Integration in Water Science and Engineering, in Proceedings of the 31st ACM International Conference on Information & Knowledge Management` (CIKM '22), Atlanta, Georgia, USA.
4. **Feng Li**, Ranran Chen, Yuankun Fu, Fengguang Song, Yao Liang, Isuru Ranawaka, Sudhakar Pamidighantam, Daniel Luna, and Xu Liang. 2021. Accelerating complex modeling workflows in CyberWater using on-demand HPC/Cloud resources. In Proceedings of 2021 IEEE eScience21.
5. **Feng Li**, Dali Wang, Feng Yan, and Fengguang Song. 2021. X-composer: enabling cross-environments in-situ workflows between HPC and cloud. In Proceedings of ACM PASC '21.
6. **Feng Li**, Daniel G. Waddington, and Fengguang Song. 2019. Userland CO-PAGER: boosting data-intensive applications with non-volatile memory, userspace paging. In Proceedings of ACM HP3C'19.

7. **Feng Li** and Fengguang Song. 2019. Building a scientific workflow framework to enable real-time machine learning and visualization. Concurrency Computat Pract Exper. 2019; 31:e4703.
8. Yuankun Fu, **Feng Li**, Fengguang Song, and Luoding Zhu. 2018. Designing a Parallel Memory-Aware Lattice Boltzmann Algorithm on Manycore Systems. In Proceedings of SBAC-PAD'18.
9. Yuankun Fu, **Feng Li**, Fengguang Song, and Zizhong Chen. 2018. Performance analysis and optimization of in-situ integration of simulation with data analysis: zipping applications up. In Proceedings of ACM HPDC'18.
10. **Feng Li** and Fengguang Song. 2017. A Real-Time Machine Learning and Visualization Framework for Scientific Workflows. In Proceedings of ACM PEARC'17.
11. Xiao Bian, **Feng Li**, and Xia Ning. 2016. Kernelized Sparse Self-Representation for Clustering and Recommendation. In Proceedings of SIAM SDM'16.