JIANSHU LIU

2303 Patrick F. Taylor Hall LSU, Baton Rouge, LA, 70803

Phone: (+1)225-315-4383 Email: jliu96@lsu.edu Homepage: https://jianshuliu1721.github.io

EDUCATION

Louisiana State University, Baton Rouge, LA, USA   *08/2018 - Present*
Ph.D. in Computer Science (CS)

Beijing University of Posts and Telecommunications, Beijing, China   *09/2014 – 06/2018*
B.S. in Electronics Engineering (EE)

RESEARCH INTERESTS

Distributed systems, cloud computing, microservices, IoT stream processing systems, performance and quality of service, and machine learning for systems.

PUBLICATIONS

- **Jianshu Liu**, Shungeng Zhang, and Qingyang Wang. "μConAdapter: Reinforcement Learning-based Fast Concurrency Adaptation for Microservices in the Cloud," in the Proceedings of *14th Symposium on Cloud Computing,* (**SoCC'23**), Santa Cruz, California, October 2023.

- **Jianshu Liu**, Qingyang Wang, Shungeng Zhang, Liting Hu, and Dilma Da Silva. "Sora: A Latency Sensitive Approach for Microservices Soft Resource Adaption," in the Proceedings of the *24th ACM/IFIP International Middleware Conference* (**Middleware'23**), Bologna, Italy, December 2023.

- **Jianshu Liu**, Shungeng Zhang, Qingyang Wang, and Jinpeng Wei. "Coordinating Fast Concurrency Adapting with Autoscaling for SLO-Oriented Web Applications," in *IEEE Transactions on Parallel and Distributed Systems* (**TPDS**), February 2022

- **Jianshu Liu**, Shungeng Zhang, Qingyang Wang, and Jinpeng Wei. "Mitigating Large Response Time Fluctuations through Fast Concurrency Adapting in the Cloud," in the Proceedings of the *34th IEEE International Parallel & Distributed Processing Symposium* (**IPDPS'20**), New Orleans, LA, USA, May 2020.

- Xuhang Gu, Qingyang Wang, Qiben Yan, **Jianshu Liu**, and Calton Pu. "Sync-Millibottleneck Attack on Microservices Cloud Architecture" to appear in the *19th ACM ASIA Conference on Computer and Communications Security* (**AsiaCCS'24**), Singapore, July 2024

- Shungeng Zhang, Qingyang Wang, Yasuhiko Kanemasa, Julius Michaelis, **Jianshu Liu**, and Calton Pu. "ShadowSync: Latency Long Tail caused by Hidden Synchronization in Real-time Stream Processing Systems," in the Proceedings of the *23rd ACM/IFIP International Middleware Conference* (**Middleware'22**), Quebec City, Quebec, Canada, November 2022.

- Calton Pu, Qingyang Wang, Yasuhiko Kanemasa, Rodrigo Alves Lima, Joshua Kimball, Shungeng Zhang, **Jianshu Liu**, and Xuhang Gu. "A Functional Model and Analysis of Next Generation Malware Attacks and Defenses," in the Proceedings of the *2021 Third IEEE International Conference on Trust, Privacy, and Security in Intelligent Systems and Applications* (**TPS-ISA'21**), Atlanta, Georgia, December 2021.

- Shungeng Zhang, Qingyang Wang, Yasuhiko Kanemasa, **Jianshu Liu**, and Calton Pu. "DoubleFaceAD:A New Datastore Driver Architecture to Optimize Fanout Query Performance," in the Proceedings of the *21st ACM/IFIP International Middleware Conference* (**Middleware'20**), Delft, Netherlands, December 2020.

- Shungeng Zhang, Huasong Shan, Qingyang Wang, **Jianshu Liu**, Qiben Yan, and Jinpeng Wei. "Tail Amplification in n-Tier Systems: A Study of Transient Cross-Resource Contention Attacks," in the Proceedings of the *39th International Conference on Distributed Computing Systems* (**ICDCS'19**), Dallas, Texas, July 2019.

RESEARCH PROJECTS

ML-Guided Molecular Simulations on Stream Processing Systems *06/2023 - Present*
Collaborate with UChicago and Argonne National Laboratory

- **AI for Science:** Designed a flexible and scalable framework for supporting ML-guided molecular simulation ensembles on a real-time stateful stream processing engine (e.g., Apache Flink), achieving comparable scientific effectiveness to Colmena.

Dracena: Real-time Platforms for Stream Processing IoT Applications *12/2021 - Present*
Collaborate with Georgia Tech. and Fujitsu Laboratories Ltd.

- **Performance Debugging:** Designed a framework to detect and diagnose the long-tail latency degradation caused by stateful object evolution/growth in persistent IoT applications that utilize digital-twins technology (e.g., HealthCare, Manufacturing, and Smart Cities).

- **QoS Support for Latency:** Participated in diagnosing and mitigating the *ShadowSync* problem, a long-tail latency issue caused by very short but intense resource demands outside critical paths, such as the overlap of flushing/compaction operations in RocksDB state backend.

Intelligent Autoscaling Frameworks with Soft Resource Adaptation *03/2019 - Present*
Collaborate with UNC Charlotte, UCSC, and Texas A&M.

- **Autoscaling:** Developed adaptive statistical models (e.g., Non-linear Regression) to quickly identify the optimal soft resource (e.g., threads) allocation for web services in monolithic and microservices systems. Mitigating SLO violations by $3\times$.

- **ML for System:** Designed an RL-based (e.g., DQN) framework to support intelligent software and hardware resource scaling for microservices cloud applications, resulting in improved system performance and cost-effectiveness.

AWARDS AND ACTIVITIES

- **Best Paper Award** in the 24th International Middleware Conference (Middleware'23) *12/2023*
- **Oral Presenter** in the 14th Symposium on Cloud Computing (SoCC'23) *10/2023*
- **Student Scholarships** in the 14th Symposium on Cloud Computing (SoCC'23) *10/2023*
- **Oral Presenter** in the 5th ParslFest Community Meeting *10/2023*
- **Summer Research Intern** in Globus Lab, UChicago *05/2023 - 08/2023*
- **Oral Presenter** in the 1st LSU Graduate Research Conference *04/2023*
- **Shadow Program Committee** in EuroSys'23 *11/2022 - 01/2023*
- **Poster Presenter** in LSU's 4th EECS Graduate Student Research *04/2022*
- **Intern** in China Telecom Beijing Research Institute *03/2018-05/2018*
- **Honorable Mention** in Mathematical Contest in Modeling *02/2017*