JIANSHU LIU

5500 Companile Drive, San Diego, CA 92182
Phone: (+1)225-315-4383 Email: jliu11@sdsu.edu Homepage: https://jianshuliu1721.github.io

EMPLOYMENT

Assistant Professor, Department of Computer Science *08/2025 - present*
San Diego State University, CA, United States

Assistant Professor, Department of Computer Science *07/2024 - 05/2025*
Boise State University, ID, United States

Summer Research Intern, Globus Lab *05/2023 - 08/2023*
University of Chicago, IL, United States

EDUCATION

Louisiana State University, Baton Rouge, LA, USA *08/2018 - 05/2024*
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, AI/ML systems in Science.

PUBLICATIONS

- Zhiqi Li, Ruiqi Yu and **Jianshu Liu**. "IrishBench: An Open-Source Benchmark Suite for Video Processing Systems in Cloud" accepted to the *In Companion of the 16th ACM/SPEC International Conference on Performance Engineering* (**HotCloudPerf'25**), Toronto, Canada, May 2025.
- Zhiqi Li, Ruiqi Yu and **Jianshu Liu**. "Poster: Benchmarking Video Processing Systems in the Cloud" presented at the *18th USENIX Symposium on Operating System Design and Implementation* (**OSDI'24**), San Jose, CA, USA, July 2024.
- **Jianshu Liu**, Shungeng Zhang, and Qingyang Wang. "μConAdapter: Reinforcement Learning-based Fast Concurrency Adaptation for Microservices in the Cloud" in Proceedings of *14th Symposium on Cloud Computing* (**SoCC'23**), Santa Cruz, CA, October 2023.
- **Jianshu Liu**, Qingyang Wang, Shungeng Zhang, Liting Hu, and Dilma Da Silva. "Sora: A Latency Sensitive Approach for Microservices Soft Resource Adaptation" in Proceedings of the *24th ACM/IFIP International Middleware Conference* (**Middleware'23**), Bologna, Italy, December 2023. [**Best Paper Award**]
- **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 Proceedings of the *34th IEEE International Parallel & Distributed Processing Symposium* (**IPDPS'20**), New Orleans, LA, May 2020.

- Xuhang Gu, Qingyang Wang, **Jianshu Liu**, and Jingpeng Wei. "Grunt Attack: Exploiting Execution Dependencies in Microservices" in the *54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks ASIA Conference on Computer and Communications Security* (**DSN'24**), Brisbane, Australia, June 2024.
- Xuhang Gu, Qingyang Wang, Qiben Yan, **Jianshu Liu**, and Calton Pu. "Sync-Millibottleneck Attack on Microservices Cloud Architecture" in the *19th ACM ASIA Conference on Computer and Communications Security* (**AsiaCCS'24**), Singapore, July 2024.
- Xuhang Gu, **Jianshu Liu**, and Qingyang Wang. "A BlackBox Approach to Profile Runtime Execution Dependencies in Microservices" in the *9th IEEE International Conference on Collaboration and Internet Computing* (**CIC'23**), Atalanta, GA, November 2023.
- 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 Proceedings of the *23rd ACM/IFIP International Middleware Conference* (**Middleware'22**), Quebec City, Quebec, Canada, November 2022.
- Shungeng Zhang, Qingyang Wang, Yasuhiko Kanemasa, **Jianshu Liu**, and Calton Pu. "DoubleFaceAD:A New Datastore Driver Architecture to Optimize Fanout Query Performance" in 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 Proceedings of the *39th International Conference on Distributed Computing Systems* (**ICDCS'19**), Dallas, TX, July 2019.

RESEARCH PROJECTS

SensEat: An Everyday Dietary Monitoring Framework Design *10/2024 - present*
Collaborate within Boise State CS department.

- **Edge-Cloud Continuum in Healthcare:** Designing an edge-cloud framework for dietary monitoring by leveraging ultrasound acoustic sensing and federated learning.

DT-RUTH: Digital Twin-based Deterministic Traffic Simulator Design *12/2024 - present*
Collaborate with IT4Innovations.

- **Digital Twin:** Designing a deterministic traffic simulator adopting digital twin technology and utilizing serverless functions to deploy in clouds and HPC clusters.

ML-Guided Molecular Simulations on Stream Processing Systems *06/2023 - 09/2023*
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 - 05/2024*
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×.
- **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.

TEACHING

Instructor at San Diego State University

- CS576 Computer Networks and Distributed Systems *Fall 2025*

Instructor at Boise State University

- CS452/552 Operating Systems *Fall 2024*

Teaching Assistant at Louisiana State University

- CSC7510 Cloud and Enterprise Systems *Fall 2021 - 2023*
- CSC4610 Cloud Systems and Virtualization *Fall 2018 - 2023*
- CSC2610 Cloud Fundamentals and Web Programming *Spring 2019 - 2023*
- CSC1350 Java Programming *Fall 2018, 2021 - 2022*
- CSC4890 Introduction of Theory of Computation *Fall 2020*
- CSC2730 Data Science and Analytics *Fall 2020*

PROFESSIONAL SERVICES AND AWARDS

- **TPC member in AI/ML for Cloud and IoT Track** in IEEE CloudSummit'25 *02/2025*
- **PC member in System Software** in IPDPS'25 *12/2024*
- **Reviewer** in Journals, TPDS, TIFS, TMC, TCC, TOIT, and TASE *08/2024 - 12/2024*
- **Operation Committee Member** in Boise State University CS *08/2024 - 12/2024*
- **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*
- **Shadow Program Committee** in EuroSys'23 *11/2022 - 01/2023*