

LEI ZHANG

website: <https://geraldleizhang.com>
email: geraldleizhang@gmail.com | lei-zhang@cs.princeton.edu

RESEARCH INTEREST

I'm interested in improving performance, observability, and reliability for large-scale, distributed systems (e.g. cloud systems, datacenter applications, in-memory/storage systems, and video streaming systems).

EMPLOYMENT

Princeton University Postdoc Research Associate	Jan. 2022 – current
Emory University Graduate Research Assistant	Aug. 2018 – Dec. 2021
Facebook Inc. Ph.D. Software Engineer Intern	May 2018 – Aug. 2018
Georgia Institute of Technology Graduate Research Assistant	Aug. 2015 – Aug. 2018

EDUCATION

Emory University Ph.D. Computer Science • Advisor: Prof. Ymir Vigfusson • Transferred from Georgia Tech as a post-qualified Ph.D. candidate	Aug. 2018 – Dec 2021
Georgia Institute of Technology M.S. Computer Science • Advisor: Prof. Karsten Schwan (deceased)	Aug. 2015 – Aug. 2018
Tsinghua University B.E. Computer Science	Aug. 2011 – July 2015

PUBLICATIONS

Towards Bandwidth-adaptive Live Volumetric Video Streaming	In preparation
Automatic Instrumentation for Fine-grained Observability in Distributed Systems	In preparation
The Benefit of Hindsight: Tracing Edge-Cases in Distributed Systems <i>Lei Zhang, Zhiqiang Xie, Vaastav Anand, Ymir Vigfusson, Jonathan Mace</i>	NSDI'23
Measurement and Analysis Methods of Performance Problems in Distributed Systems <i>Lei Zhang</i>	Doctoral Thesis
When is the Cache Warm? Manufacturing a Rule of Thumb <i>Lei Zhang, Juncheng Yang, Anna Blasiak, Mike McCall, Ymir Vigfusson</i>	HotCloud'20
Optimal Data Placement for Heterogeneous Cache, Memory, and Storage Systems <i>Lei Zhang, Reza Karimi, Irfan Ahmad, Ymir Vigfusson</i>	SIGMETRICS'20 <i>Best Student Paper</i>
Deceptive Secret Sharing <i>Lei Zhang, Douglas M. Blough</i>	DSN'18
Systematic Data Placement Optimization in Multi-Cloud Storage for Complex Requirements <i>Maomeng Su, Lei Zhang, Yongwei Wu, Kang Chen, Keqin Li</i>	IEEE ToC'16

RESEARCH PROJECTS

Observability and reliability:

Tracing Edge-cases in Distributed Systems

- Developed Hindsight, a lightweight always-on distributed tracing system to track edge-case problems (e.g. tail latency) in large-scale systems
- Proposed retroactive sampling to overcome today's tracing system limitation on low sampling rate
- Evaluated Hindsight can trace 100% requests with $< 2.5\%$ overhead, provide high scalability, and practically capture low-level trace data for real-world system problems

Automatic Instrumentation for Distributed Tracing

- Proposed to leverage static analysis methods on distributed tracing to provide fine-grained information at runtime
- Developed tools to automatically instrument distributed systems for fine-grained trace data, and locate responsible code pieces for performance issues
- Developed tools to analyze code dependencies to guide heuristics on potential provenances

Distributed caching and data placement

Optimal Data Placement for Heterogeneous Memory Systems

- Developed CHOPT, a data-driven offline optimal placement framework to provide performance upper bound and examine heuristics towards online large-scale caching system design.
- Proposed design principles to revisit today's new memory hardware (e.g. NVMe) and hierarchies (e.g. remote memory) to fully leverage them on large-scale cache algorithm design
- Evaluated CHOPT with a wide range of real-world traces and exposed that state-of-the-art distributed cache can be further improved by up to 44.8%
- Concluded practical heuristics on data placement strategy for distributed cache or memory systems

Understanding Distributed Cache Warmup

- Conducted the importance of understanding cache warmup process in large-scale systems
- Proposed a practical quantification, and derived a rule-of-thumb estimation formula for cache warmup time
- Evaluated with a wide range of real-world workloads and proved high accuracy on cache warmup time estimation

Fine-grained Storage Utilization Management for Facebook's Video Cache

- Developed a utilization monitor of downstream storage usage for different customers and video contents and handled lags of tracing services through approximation
- Developed a service to maintain priorities of fine-grained utilization and automatically make eviction decisions, to ensure efficient resource utilization and fairness

Formalized Data Placement Optimization in Multi-cloud Storage

- Developed Triones, a systematic model that is based on erasure coding and formalized data placement optimization for multi-cloud storage configuration
- Proved Triones can efficiently optimize cloud application's performance and service quality, e.g. fault-tolerance, latency, and costs

Others:

Live Volumetric Video Streaming

- Proposed a new design of volumetric video streaming systems, which directly transmits adaptive selected 2D stream data to overcome the network bandwidth limitation
- Developed a bandwidth-adaptive live volumetric video streaming system
- Evaluated the system can achieve high throughput, low end-to-end latency, and eliminate dependencies on heavy computing resources.

Data-driven Malware Detection

- Proposed a systematic data-driven approach to attribute malware behaviors to specific actors, tools, and intents.
- Developed a system to parse and clean raw system and network logs from a real malware execution dataset, and identify unique system patterns to define malware's distinct features.

- Designed a LSM-Tree-like data structure to effectively manage large-scale malware execution traces effectively, and a clustering algorithm for identifying related malware families.

Deceptive Secret Sharing

- Proposed deceptive secret sharing to combine deception into secret sharing schemes and prevent attacks such as insider attacks.
- Developed a distributed storage prototype for notably enhanced privacy with limited and reasonable storage and computational overhead.

AWARDS

Kenneth C. Sevcik Outstanding Student Paper Award

SIGMETRICS'20

Bronze Medal

24th, 25th China Mathematical Olympiad

PROFESSIONAL SERVICES

Invited Program Committee

ACM SIGMETRICS'23

Program Committee

ACM SOCC'22

External Reviewer

USENIX ATC'21, HotCloud'20, HotStorage'20

External Reviewer

ACM Eurosys'19, SOCC'18, SRDS'18

MENTORSHIP

Rajrup Ghosh (Ph.D., USC), Jingyuan Chen (Ph.D., Princeton), Yazhuo Zhang (Ph.D., Emory), Tao Zhou (Undergrad, Emory)