Lei Zhang

Website: https://geraldleizhang.com
Tel: +1 (404) 662-8268
Email: lei.zhang@emory.edu

EDUCATION

Emory University, Atlanta — ***Aug 2018 – Jul 2021 (Expected)***

- Ph.D. candidate in Computer Science
  - **Thesis**: Measurement and Diagnosis Methods for Distributed System Performance
  - Advisor: Prof. Ymir Vigfusson
  - Committee: Prof. Robbert van Renesse, Prof. Nosayba El-Sayed, Prof. Avani Wildani
  - *Transferred as a post-qual Ph.D. student*

Georgia Institute of Technology, Atlanta — ***Aug 2015 – Aug 2018***

- Ph.D. Student in Computer Science
  - Advisor: Prof. Karsten Schwan(deceased)
- Master of Science (M.S.) in Computer Science — ***May 2017***

Tsinghua University, Beijing — ***Aug 2011 – Jul 2015***

- Bachelor of Engineering (B.E.) in Computer Science and Technology
  - **Thesis:** Panku: a matrix-centered distributed computation framework

RESEARCH INTERESTS

Distributed Systems, Cloud Systems, Cache & Memory Management, Distributed Tracing, Performance Analysis

PUBLICATIONS

"Distributed Tracing Through Retrospective Logging.". Under Review.

Lei Zhang, Juncheng Yang, Anna Blasiak, Mike McCall, Ymir Vigfusson, "When is the Cache Warm? Manufacturing a Rule of Thumb". In *USENIX HotCloud 2020*

Lei Zhang, Reza Karimi, Irfan Ahmad, Ymir Vigfusson, "Optimal Data Placement for Heterogeneous Cache, Memory, and Storage Systems". In *ACM SIGMETRICS 2020*. Kenneth C. Sevcik Outstanding Student Paper Award

Lei Zhang, Douglas M. Blough, "Deceptive Secret Sharing". In *IEEE International Concerence on Dependable Systems and Networks (DSN) 2018*

Maomeng Su, **Lei Zhang**, Yongwei Wu, Kang Chen, Keqin Li, "Systematic Data Placement Optimizationin in Multi-Cloud Storage for Complex Requirements". In *IEEE Transactions on Computers (TC) 2016*, Vol. 65, Issue 6, pp.1964-1977

MAJOR RESEARCH EXPERIENCE

Distributed Tracing Through Retrospective Logging — ***Oct 2019 – Dec 2020, Emory University***

- We propose a novel retroactive tracing abstraction where full telemetry information about a distribute request can be retrieved in runtime soon after a performance problem is detected.
- We design and implement Hindsight, a distributed tracing systems that enables always-on retroactive tracing, to overcome the current tracing systems' bottlenecks to handle large amount of data.
- We applied Hindsight onto a variety of systems and benchmarks, and show Hindsight can rapidly persist full, detailed retroactive traces with competitive overhead than current sampling based methods.

Optimal Placement Policy for Memory Hierarchy — ***Sep 2018 – Aug 2019, Emory University***

- We frame the challenges of data placement in modern memory hierarchies in a generalized paging model that makes way with anachronistic assumptions.
- We design an offline placement algorithm for providing optimal placement decisions as the upper bound of performance gain for any data placement algorithm, allowing us to chart the realm of possible future improvements. We show how spatial sampling enables this algorithm to support a variety of large traces efficiently. We also show analytically that spatial sampling gives a high-fidelity approximation.
- We present trace-driven simulation results for the algorithm that show an opportunity to improve latency performance ranges between 8.2% and 44.8% on average relative to the clairvoyant Belady's MIN optimal cache replacement algorithm. We also show how our revisited assumptions of cache bypass and performance asymmetry contributes to such improvements.
- We evaluate the performance of applying spatial sampling on the algorithm, showing an error of only 0.2% difference on average latency at 1% sampling rate on PARSEC benchmarks.

Understanding Distributed Cache Warm Up Process *Jan 2018 – Aug 2018, Emory University*

- We want to investigate how long a cache takes to become useful after being suppressed or lost cache items
- I give a definition of a cache server getting warmed-up by cache hit rate, and with the definition provide a formula to estimate the cache warm-up time.
- I measure the accuracy of provided formula through real world CDN traces from Akamai, non-CDN traces from CloudPhysics, and file system cache traces from MSR. Experimental results show that our provided formula can give an accurate estimation of cache warm-up time for most of the traces.

OTHER RESEARCH EXPERIENCE

Understanding Causality in System Logs for Cybersecurity *May 2017 – Dec 2017, Georgia Tech*

- We are trying to find causality between system behaviors and network behaviors, for achieving malware attribution and for further purposes like actor, tool, and intent identification
- I parse raw system and network logs for a real malware execution dataset, which includes clearing up useless traces, and identifying unique system patterns for further defining malware's unique features
- I design a new LSM-Tree liked data structure for effectively storing large scale malware execution traces and efficiently clustering malware families

Distributed Storage for Sensitive Data *Sep 2016 – Dec 2017, Georgia Tech*

- This work aims at achieving distributed, secret shared, and obfuscated data storage, which can provide trade-off between security of sensitive data and storage resources.
- I design and implement two novel schemes for deception within both XOR secret sharing and polynomial-based threshold secret sharing that jointly encode multiple secrets, together with quantitative security analyses for these schemes
- I build a distributed prototype implementing proposed schemes, achieving notable enhanced security with limited and reasonable infection on system performance including storage space, latency, and throughput

Message Transaction in HPC Environment *Aug 2015 – May 2016, Georgia Tech*

- EVpath is designed to be an event transport middleware layer, which is maintained by CERCS Gatech
- I implement a publish/subscribe library to expand transaction module
- I experiment on expanding EVpath's scalability over InfiniBand, working on fabric communication frameworks like Libfabric and testing message passing and memory management component on EVpath

Sparse Matrix Partition in Graph Processing *Jan 2015 – Jul 2015, Tsinghua University*

- Panku is a matrix-centered distributed computation framework, aiming at optimization of distributed storage for real-world large scale computing tasks like graph processing and machine learning, where data are always expressed in sparse matrix format
- I design a sparse matrix partition algorithm which takes both local and global storage optimization into consideration
- I implement the proposed computation framework, which achieves trade-off among computational skew, single node memory consumption, and network traffic in multi-core computing environment, providing up to 2 times improvement on overall computing performance

Formalized Data Placement Optimization in Multi-cloud Storage *Sep 2013 – Dec 2014, Tsinghua University*

- Triones is an erasure-code based systematic model, focusing on formalized data placement optimization in multi-cloud storage where current methods can only fulfill ad-hoc cloud storage requirements
- We propose a non-linear programming model for data placement configuration with complex requirements
- We design and implement a set of APIs for automatically generating optimized multi-cloud storage configurations

WORK EXPERIENCE

Teaching Assistant, Emory CS 377: Database Systems *Aug 2020 – Dec 2020*

Research Assistant, Emory University *Aug 2018 – present*

Software Engineer Intern, Facebook Inc. *May 2018 – Aug 2018*

- For video caching infrastructure, I built a utilization monitor for maintaining outside distributed storage usage towards partitioning between different customers and making file eviction decisions.
- I designed such monitor achieving aggregation of data stream through buffer caches, handling both customer and file level utilization logs, and tolerating lags of relied tracing services through approximation.
- I built a service for processing such utilization, which includes designing a data structure for tracking utilization history, and a priority table for file storage and eviction decisions.

Teaching Assistant, Gatech CS 3210: Design Operating Systems *Aug 2016 – Dec 2016*

Research Assistant, Georgia Tech *Aug 2015 – Aug 2018*

PROFESSIONAL SERVICE

External reviewer: **SRDS** 2018, **SOCC** 2018, **EuroSys** 2019, **HotCloud** 2020, **HotStorage** 2020, **ATC** 2021

OTHERS

Kenneth C. Sevcik Outstanding Student Paper Award, SIGMETRICS'20

Travel Grants, NSDI'20, FAST'20

Bronze medal, 24th, 25th China Mathematical Olympiad

Experienced in C++, C, C#, GoLang, Python

Familiar with Linux Kernel, Parallel Programming, Assembly Programming, LLVM, Apache Kafka, MySQL, MongoDB, OpenTracing