Liang Geng

Curriculum Vitae

geng.161@osu.edu — Google Scholar — Linkedin

EDUCATION

Ph.D. Candidate
2022 - Current
The Ohio State University, USA
Computer Science and Engineering

Master of Engineering
2016 - 2019
Northeastern University, China
Computer Science and Engineering

Bachelor of Engineering
2012 - 2016
Liaoning Technical University, China
Software Engineering

EMPLOYMENT

Graduate Research Associate
Jan 2022 - Current
The Ohio State University, Computer Science and Engineering Department

Software Engineer Intern
Jun 2025 - Aug 2025
Wherobots, Spatial Database Group
I was responsible for developing a GPU-accelerated spatial join library. The library is based on my research project, LibRTS, which exploits NVIDIA RT cores to accelerate the filtering stage of a spatial join query. A GPU-accelerated RelateEngine is also designed and implemented to support various geometries and predicates.

Senior Development Engineer (Full Time)
Jun 2020 - Dec 2021
Alibaba Group, DAMO Academy
I was responsible for the R&D of state-of-the-art graph computing systems, maintenance of open-source projects libgrape-lite and GraphScope; participation in the implementation and evaluation of research projects.

Software Engineer (Full Time)
Jan 2019 - May 2020
Chehaoduo Group, Data Platform Division
I was responsible for the R&D of the toolchains for Datawarehouse, including SQL parsers, ETL job management systems.

Visiting Research Associate
Dec 2017 - Sep 2018
The Ohio State University, Computer Science and Engineering Department
I was responsible for designing and developing a hybrid graph computing framework SEP-Graph.

REPRESENTATIVE PROJECTS

LibGPUSpatial
Jun 2025 - Aug 2025
Wherobots, Inc

`LibGPUSpatial` is a GPU-accelerated spatial join library. The library is based on my research project, LibRTS, which exploits NVIDIA RT cores to accelerate the filtering stage of a spatial join query. To support various geometries and spatial predicates, I designed a RelateEngine that implements the DE-9IM with CUDA. To improve performance with complex polygons, I designed a mechanism that adaptively applies warp-level parallelism for better load balancing and memory access patterns.

LibRTS
May 2023 - Aug 2024
The Ohio State University

`LibRTS` is a fast spatial index library that leverages RT cores to accelerate spatial queries. `LibRTS` supports both point and range queries and remains mutable to accommodate changing data. Instead of relying on a case-by-case approach, `LibRTS` provides a general, high-performance spatial indexing framework for spatial data processing. By formulating spatial queries as RT-suitable problems and overcoming load-balancing challenges, `LibRTS` delivers superior query performance through RT cores without requiring developers to master complex programming on this specialized hardware.

RayJoin
Aug 2022 - Jan 2024
The Ohio State University

`RayJoin` is a high-performance spatial join framework that utilizes NVIDIA RT Cores as accelerators to facilitate real-time spatial join queries. Specifically, `RayJoin` accelerates two critical spatial join queries: line segment intersection and point-in-polygon test, offering high performance and precision. Additionally, `RayJoin` introduces solutions to address two significant technical challenges: (1) achieving high precision computation with the limited precision support of the RT Cores, and (2) reducing the substantial setup costs of Bounding Volume Hierarchy while maintaining optimal query performance.

RR-Compound
Jan 2022 - May 2023
The Ohio State University

`RR-Compound` is an RDMA-fused gRPC variant designed for low latency and high throughput. `RR-Compound` is fully compatible with gRPC and can serve as a seamless replacement without requiring modifications to existing applications. The primary challenge involves connection management. There is a significant discrepancy in connection management between RDMA, which relies on busy polling to detect new messages, and gRPC, which uses epoll-based APIs. To address this, we developed a new method called BPEV that utilizes dedicated polling threads to monitor incoming buffers and generates I/O events through the kernel-supported event-fd mechanism.

GPU support for libgrape-lite

Nov 2020 - Dec 2021
Alibaba Group

`libgrape-lite` is a C++ library from Alibaba designed for parallel graph processing in distributed environments. During my tenure at Alibaba, I developed GPU support for `libgrape-lite`. This enhancement consists of three major components: (1) in-device-memory graph representations, which allow the graph to be maintained on the GPU; (2) state-of-the-art load balancing strategies, which improve performance for irregular graphs; and (3) multi-GPU communications utilizing NCCL.

SEP-Graph
Jan 2018 - Sep 2018
The Ohio State University

`SEP-Graph` is a graph processing framework that utilizes hybrid execution modes. Traditional graph-processing frameworks typically employ a single strategy throughout the execution of an application, often leading to variable and suboptimal performance. In contrast, `SEP-Graph` features a hybrid execution mode that automatically switches among synchronous/asynchronous execution modes, push/pull communication mechanisms, and data-driven/topology-driven traversal schemes to minimize execution time in each iteration.

PUBLICATIONS

[1] **Geng, Liang** and Lee, Rubao and Zhang, Xiaodong. "LibRTS: A Spatial Indexing Library by Ray Tracing". In: *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*. 2025.

[2] **Geng, Liang** and Lee, Rubao and Zhang, Xiaodong. "RayJoin: Fast and Precise Spatial Join". In: *Proceedings of the ACM International Conference on Supercomputing*. 2024.

[3] **Geng, Liang** and Wang, Hao and Meng, Jingsong and Fan, Dayi and Ben-Romdhane, Sami and Pichumani, Hari Kadayam and Phegade, Vinay and Zhang, Xiaodong. "RR-Compound: RDMA-Fused gRPC for Low Latency and High Throughput With an Easy Interface". In: *IEEE Transactions on Parallel and Distributed Systems* (2024).

[4] Gong, Shufeng and Tian, Chao and Yin, Qiang and Wang, Zhengdong and Yu, Song and Zhang, Yanfeng and Yu, Wenyuan and **Geng, Liang** and Fu, Chong and Yu, Ge and Zhou, Jingren. "Ingress: an automated incremental graph processing system". In: *The VLDB Journal* (2024).

[5] Meng, Ke and **Geng, Liang** and Li, Xue and Tao, Qian and Yu, Wenyuan and Zhou, Jingren. "Efficient Multi-GPU Graph Processing with Remote Work Stealing". In: *2023 IEEE 39th International Conference on Data Engineering*. 2023.

[6] Xiao, Mengbai and Wang, Hao and **Geng, Liang** and Lee, Rubao and Zhang, Xiaodong. "An RDMA-enabled In-memory Computing Platform for R-tree on Clusters". In: *ACM Transactions on Spatial Algorithms and Systems* (2022).

[7] Fan, Wenfei and **Geng, Liang** and Jin, Ruochun and Lu, Ping and Tugay, Resul and Yu, Wenyuan. "Linking Entities across Relations and Graphs". In: *2022 IEEE 38th International Conference on Data Engineering*. 2022.

[8] Gong, Shufeng and Tian, Chao and Yin, Qiang and Yu, Wenyuan and Zhang, Yanfeng and **Geng, Liang** and Yu, Song and Yu, Ge and Zhou, Jingren. "Automating incremental graph processing with flexible memoization". In: *Proceedings of the VLDB Endowment* (2021).

[9] Wang, Qiange and Zhang, Yanfeng and Wang, Hao and **Geng, Liang** and Lee, Rubao and Zhang, Xiaodong and Yu, Ge. "Automating incremental and asynchronous evaluation for recursive aggregate data processing". In: *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 2020.

[10] Xiao, Mengbai and Wang, Hao and **Geng, Liang** and Lee, Rubao and Zhang, Xiaodong. "Catfish: Adaptive RDMA-enabled r-tree for low latency and high throughput". In: *2019 IEEE 39th International Conference on Distributed Computing Systems*. 2019.

[11] Zhang, Simon and Xiao, Mengbai and Guo, Chengxin and **Geng, Liang** and Wang, Hao and Zhang, Xiaodong. "Hypha: a framework based on separation of parallelisms to accelerate persistent homology matrix reduction". In: *Proceedings of the ACM International Conference on Supercomputing*. 2019.

[12] Wang, Hao and **Geng, Liang** and Lee, Rubao and Hou, Kaixi and Zhang, Yanfeng and Zhang, Xiaodong. "Sep-graph: finding shortest execution paths for graph processing under a hybrid framework on GPU". In: *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming*. 2019.

Grants & AWARDS

Proceedings of the 38th ACM International Conference on Supercomputing
Travel Grant
Jun, 2024

CSE Graduate Student Research Poster Exhibition
Best Poster Award
Apr, 2024
The Ohio State University

The 3rd National University Cloud Computing Application Innovation Competition
First Place Award
Apr, 2017
Nanjing, China