Liang Geng

Curriculum Vitae

geng.161@osu.edu — Google Scholar — Linkedin — GitHub — Apache Sedona Committer

RESEARCH INTERESTS

High-performance spatial databases, GPU-accelerated computing (NVIDIA RT Cores/CUDA), and large-scale graph analytics.

EDUCATION

Ph.D. Candidate
Jan, 2022 - Dec, 2026 (Expected)
The Ohio State University, USA
Computer Science and Engineering

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

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

EMPLOYMENT

Graduate Research Associate
Jan 2022 - Dec 2025
The Ohio State University, Computer Science and Engineering Department

Software Engineer Intern
Jun 2025 - Aug 2025
Wherobots, Spatial Database Group

- Developed a GPU-accelerated spatial join library based on my previous research projects [4, 3] that exploit NVIDIA RT Cores to accelerate the filtering stage of spatial queries.
- Designed and implemented a GPU-accelerated RelateEngine to support diverse geometries and combinatorial predicates.
- Invited to become an official Apache Sedona Committer in recognition of core open-source database contributions.

Senior Development Engineer (Full Time)
Jun 2020 - Dec 2021
Alibaba Group, DAMO Academy

- Researched and developed state-of-the-art graph computing systems.
- Maintained and optimized major open-source projects, including libgrape-lite and GraphScope.
- Directed the implementation and performance evaluation of key internal research initiatives.

Software Engineer (Full Time)
Jan 2019 - May 2020
Chehaoduo Group, Data Platform Division

- Engineered and optimized toolchains for enterprise data warehouses, including robust SQL parsers and ETL job management systems.

Visiting Research Associate
Dec 2017 - Sep 2018
The Ohio State University, Computer Science and Engineering Department

- Architected and developed SEP-Graph [12], a highly efficient hybrid graph computing framework.

REPRESENTATIVE PROJECTS

RayBooster: RT-Accelerated Spatial Query Engine [1]
Jun 2025 - Aug 2025
Wherobots, Inc

- Architected and integrated the **first Ray Tracing (RT) core acceleration layer** into Apache SedonaDB, optimizing spatial join performance for production-grade geospatial databases.
- Engineered a GPU-optimized Structure of Arrays (SoA) storage layout to replace Well-Known Binary (WKB) formats, enabling high-performance random access and memory throughput.
- Developed a **Monolithic Bounding Volume Hierarchy (BVH)** indexing strategy that encodes geometry IDs into the Z-axis, eliminating the scalability overhead of managing millions of micro-indexes.
- Designed **RelateEngine** that exploits RT Cores to compute the *DE-9IM model*, supporting complex combinatorial spatial predicates across diverse geometries.
- Implemented a memory-aware execution scheduler to handle massive datasets, providing robust resource management and preventing Out-of-Memory failures, and reducing memory allocation overhead with RAPIDS RMM.
- Delivered a **5.8**× speedup on the *SpatialBench* suite and reduced operational cloud costs by **58.1%** by repurposing idle GPU RT units for spatial analytics.
- **This work has been part of Apache SedonaDB**.

X-HD: RT-Accelerated Hausdorff Distance Computation [2]
Jan 2025 - Dec 2025
The Ohio State University

- Developed a general-purpose Hausdorff Distance (HD) algorithm leveraging NVIDIA Ray Tracing (RT) cores to accelerate spatial data analysis for large-scale datasets.
- Engineered a custom grid-partitioned Bounding Volume Hierarchy (BVH) traversal strategy, reducing tree traversal intensity and optimizing hardware resource utilization.
- Implemented HD Estimators to prune non-contributing data points, significantly reducing redundant computations during the k-NN search phase.
- Optimized hardware throughput by developing a hybrid scheduling model, selectively offloading distance computations from RT shaders to CUDA kernels to resolve load imbalances.
- Achieved a **5.3**× speedup over the industry-standard *ITK* library and up to **6.4**× faster performance compared to existing GPU-optimized solutions.

LibRTS: A High-Performance RT-Accelerated Spatial Index Library [3]
May 2023 - Aug 2024
The Ohio State University

- Designed and implemented a **general-purpose spatial index library** that repurposes **NVIDIA Ray Tracing (RT) Cores** to accelerate non-rendering workloads like point and range queries.
- Developed a **mutable indexing framework** capable of adapting to dynamic datasets, bridging the gap between specialized RT hardware and general spatial data processing.
- Formulated complex spatial queries into **hardware-native BVH traversal problems**, abstracting the specialized RT programming model into a developer-friendly API.
- Engineered a **load-balancing mechanism** to optimize GPU utilization, ensuring high throughput across diverse query types (point-in-polygon, contains, and intersects).
- Outperformed industry-standard CPU and GPU spatial libraries by up to **93.2× for range queries** and **11.7× for point queries**.
- **The research outcomes of this project have been integrated into Apache SedonaDB.**

RayJoin: Real-Time RT-Accelerated Spatial Join Framework [4]
Aug 2022 - Jan 2024
The Ohio State University

- Developed **RayJoin**, a high-performance framework utilizing **NVIDIA RT Cores** to achieve unprecedented real-time speeds for Line Segment Intersection (LSI) and Point-in-Polygon (PIP) queries.
- Engineered a **high-precision spatial join framework** that overcomes hardware-level floating-point limitations, ensuring rigorous analytical accuracy despite underlying hardware constraints.
- Optimized BVH (Bounding Volume Hierarchy) construction pipelines, significantly reducing index buildup latency while maintaining peak query execution performance.
- Resolved critical bottlenecks in conventional plane-sweeping and tree-based algorithms, enabling the processing of **millions of polygons in under 460ms**.
- Achieved performance speedups ranging from **3.0× to 28.3×** over state-of-the-art GPU-optimized methods, establishing a new benchmark for real-time spatial analytics.
- **The research outcomes of this project have been integrated into Apache SedonaDB.**

RR-Compound: RDMA-Enabled gRPC Framework [5]
Jan 2022 - May 2023
The Ohio State University

- Engineered **RR-Compound**, a high-performance RPC framework that integrates **Remote Direct Memory Access (RDMA)** as an internal transport for **gRPC**, achieving sub-millisecond latencies for data center workloads.
- Designed **BPEV (Busy-Pooling + Event)**, a novel connection management method that decouples RDMA resource handling from gRPC internals, enhancing scalability across massive connection counts.
- Optimized the gRPC asynchronous framework to operate within constrained RDMA-registered memory, balancing high-throughput buffer management with strict memory-space utilization.
- Designed the system as a **seamless replacement** for standard gRPC, requiring zero code changes for existing production applications (e.g., TensorFlow, KV-Stores).
- Outperformed the state-of-the-art *mRPC* in micro-benchmarks, delivering a **14.77% increase in throughput** and a **42.55% reduction in latency**.
- Validated performance in real-world deployments, achieving a **2.35×** throughput boost for *TensorFlow* and *KV-Store* applications compared to standard gRPC over IPoIB.

GPU support for libgrape-lite
Nov 2020 - Dec 2021
Alibaba Group

- Engineered a **GPU-acceleration layer** for *libgrape-lite*, Alibaba's high-performance C++ library for distributed parallel graph processing.

- Designed and implemented **in-device-memory graph representations**, enabling large-scale graph structures to be maintained and processed directly in GPU VRAM for reduced latency.
- Developed **state-of-the-art load balancing strategies** to optimize GPU throughput for irregular graph topologies, mitigating thread divergence and hardware underutilization.
- Integrated **multi-GPU communication** protocols utilizing **NVIDIA NCCL**, enabling efficient data synchronization and scaling across multi-node GPU clusters.

SEP-Graph: Hybrid Graph Computing Framework on the GPU [12]
Jan 2018 - Sep 2018
The Ohio State University

- Developed **SEP-Graph**, a high-performance framework that optimizes graph processing by automatically switching between execution modes (**Sync/Async**), communication mechanisms (**Push/Pull**), and traversal schemes (**Data-driven/Topology-driven**).
- Engineered a **shortest-path execution model** that dynamically reconfigures algorithm parameters per iteration to minimize execution time based on real-time graph characteristics.
- Implemented hardware-specific optimizations for NVIDIA Pascal and Volta architectures, addressing thread divergence and memory coalescing issues inherent in irregular graph workloads.
- Achieved significant performance gains over state-of-the-art frameworks, reducing execution time by up to **45.8× compared to Groute** and **39.4× compared to Gunrock**.
- Validated system scalability and efficiency across multiple GPU generations (GTX 1080, P100, and V100) using diverse graph datasets.

PROGRAMMING SKILLS

Languages: C++, CUDA, Java, Python, SQL, Shell
Frameworks & Libraries: OptiX, NCCL, MPI, gRPC
Systems & Technologies: GPU, RDMA, Parallel Computing

PUBLICATIONS

[1] **Liang Geng**, Rubao Lee, Dewey Dunnington, Feng Zhang, Jia Yu, Xiaodong Zhang. "RayBooster: A Ray Tracing Engine to Accelerate SedonaDB". In: *Proceedings of the VLDB Endowment (VLDB Industry Track) (to appear)*. 2026.

[2] **Liang Geng**, Zhehu Yuan, Rubao Lee, Fusheng Wang, Xiaodong Zhang. "X-HD: Fast Hausdorff Distance Computation with Ray Tracing". In: *Proceedings of the ACM International Conference on Supercomputing 2026 (to appear)*. 2026.

[3] **Liang Geng**, Rubao Lee, Xiaodong Zhang. "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.

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

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

[6] Shufeng Gong, Chao Tian, Qiang Yin, Zhengdong Wang, Song Yu, Yanfeng Zhang, Wenyuan Yu, **Liang Geng**, Chong Fu, Ge Yu, Jingren Zhou. "Ingress: An Automated Incremental Graph Processing System". In: *The VLDB Journal* (2024).

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

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

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

[10] Shufeng Gong, Chao Tian, Qiang Yin, Wenyuan Yu, Yanfeng Zhang, **Liang Geng**, Song Yu, Ge Yu, Jingren Zhou. "Automating Incremental Graph Processing with Flexible Memoization". In: *Proceedings of the VLDB Endowment* (2021).

[11] Qiange Wang, Yanfeng Zhang, Hao Wang, **Liang Geng**, Rubao Lee, Xiaodong Zhang, Ge Yu. "Automating Incremental and Asynchronous Evaluation for Recursive Aggregate Data Processing". In: *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 2020.

[12] Hao Wang*, **Liang Geng***, Rubao Lee, Kaixi Hou, Yanfeng Zhang, Xiaodong Zhang. "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*. (* Co-first authors). 2019.

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

[14] Simon Zhang, Mengbai Xiao, Chengxin Guo, **Liang Geng**, Hao Wang, Xiaodong Zhang. "Hypha: A Framework based on Separation of Parallelisms to Accelerate Persistent Homology Matrix Reduction". In: *Proceedings of the ACM International Conference on Supercomputing*. 2019.

GRANTS & AWARDS

Graduate Research Award
Honored for his mastery of hardware acceleration and parallel systems, translating academic prototypes into industry-leading tools
Apr, 2026
The Ohio State University

Presidential Fellowship
One-year support for Ph.D. Dissertation
Nov, 2025
The Ohio State University

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 (Designed and implemented a general suffix tree on Apache Spark for string processing)
Apr, 2017
Nanjing, China