DONGHYUN SOHN

224-204-4582 ⋄ Chicago, IL, United States
donghyun.sohn@u.northwestern.edu ⋄ linkedin.com/in/bulelion37 ⋄ https://donghyun-sohn.github.io

EDUCATION

Northwestern University — Evanston, IL, United States
Ph.D. in Computer Science — Sep. 2022 - Present

- Advisor: Jennie Rogers

Sogang University — Seoul, Republic of Korea
B.S. in Computer Science and Engineering — Mar. 2016 - Aug. 2021

- Summa Cum Laude (2 years of military service included)

EXPERIENCE

Graduate Research Assistant, Northwestern University — Evanston, IL, United States
Sep. 2022 - Present

- Introduced Alchemy, a protocol-agnostic optimization framework for privacy-preserving query processing, combining traditional optimization techniques with circuit-aware cost modeling. Achieved up to 100× speedups in secure query execution across multiple cryptographic protocols.
- Built an OpenFHE-based columnar engine with SIMD and a hardware-calibrated cost model for optimal per-operator thread assignment; currently extending with GPU acceleration.
- Built a Differential Privacy-compliant data synthesis framework for OLAP performance regression, ensuring high fidelity simulation of statistical distributions and physical storage layouts in DuckDB/Parquet.
- Investigating LLM-based NL2SQL translation with a focus on schema incompleteness as a structural bottleneck; evaluating retrieval-augmented approaches on the BIRD benchmark.
- Designed a cost-aware optimizer for secure SQL workloads with ML-based runtime prediction, incorporating AWS hardware characteristics for performance- and cost-efficient instance selection.

Undergraduate Research Intern, Sogang University — Seoul, Republic of Korea
Aug. 2020 - Dec. 2021

- Optimized the BIRCH clustering algorithm by designing a disk-based Clustering Feature tree with selective buffering, reducing disk I/O and improving clustering quality by 7% for large-scale datasets.

PUBLICATIONS

1. **Donghyun Sohn**, Kelly Jiang, Nicolas Hammer, Jennie Rogers, "Alchemy: A Query Optimization Framework for Oblivious SQL," *Proceedings of the VLDB Endowment*, 2025. [Paper] [Code]

2. **Donghyun Sohn**, Xiling Li, Jennie Rogers, "Everything You Always Wanted to Know About Secure and Private Database Systems (but were Afraid to Ask)," *Data Engineering Bulletin*, 2023. [Paper]

3. **Donghyun Sohn**, Sungwon Jung, "Disk-based BACF Tree for Clustering Massive Datasets," *Korea Software Congress*, 2021. [Paper] [Code]

SERVICES

Program Committee Member
SIGMOD Availability & Reproducibility Initiative (ARI) — 2024, 2025

Teaching Assistant, Northwestern University
COMP_SCI 339: Intro to Database Systems — Fall 2023, 2024, 2025

SKILLS

Programming: C, C++, Java, Python, SQL
Systems/Tools: AWS, Docker, Ubuntu, EMP-Toolkit, OpenFHE