

Jiesong Liu (4th-year Undergraduate, Renmin University of China)

CONTACT INFORMATION

No. 59 Zhongguancun Street, Haidian District Beijing, 100872,
P.R. China.
jliu93@ncsu.edu; liujiesong@ruc.edu.cn
+86 17711737968; +1 9844640548
Home Page: <https://fred1031.github.io/>
GitHub: <https://github.com/Fred1031>

RESEARCH INTERESTS

- My research interests lie in the broad fields of computer architecture design, compiler optimization, programming systems, and hardware acceleration.
- I am interested in putting efforts cutting across multiple programming technologies, ranging from high-level algorithmic optimization and autotuning, to kernel library implementations (especially for GPUs), advanced compilation constructions, and computer architecture designs.
- My current research focuses on building high-performance, energy-efficient, and high-fidelity programming frameworks for the emerging machine learning.

EDUCATION

Renmin University of China, Beijing, China

B.S., School of Information Sept 2019 to Jul 2023(Expected)

University of British Columbia, Vancouver, Canada

Exchange Program Jan. 2022 to Apr 2022

- *Major: Computer Science*
- *GPA: 3.97/4.00*
- *Core Courses: Data Structures and Algorithms (Honors Course) 96, Introduction to Computer System 96, Introduction to Data Science 97, Operating System (Honors Course) 94, Advanced Programming 92, Higher Algebra 96, Introduction to Stochastic Process 100*
- *Language Proficiency: TOEFL 111 (R:30 S:25 W:28)*
- *Standard Test: GRE 327 (Q: 170)*
- *Awards: 2020 National Scholarships of China (highest scholarship for Chinese undergraduate)*

ACADEMIC AND RESEARCH EXPERIENCE

Efficient DNN Inference on Microcontrollers via Transient Redundancy Elimination-based Convolution Oct 2021 to Oct 2022

Research Assistant, North Carolina State University, Supervisor: Xipeng Shen

- Proposed the use of TREC (Transient Redundancy Elimination-based Convolution) as a new way to reduce computations in DNNs running on microcontrollers.
- Introduced a set of optimizations to mitigate the space overhead incurred by TREC.
- Empirically evaluated the effectiveness of the new solution on two models of microcontrollers, confirming the substantial benefits of the new solution in enabling efficient DNNs on microcontrollers.

Enabling Efficient Learned Index on GPU Feb 2021 to Sept 2021

Research Assistant, Key Laboratory of Data Engineering and Knowledge Engineering, Renmin University of China

Supervisor: Feng Zhang, and Xiaoyong Du

- Developed the first dynamic learned index on GPUs with an intelligent design that is optimized for querying and updating

- Devised efficient indexing strategy which, arguably, achieves coalesced memory access and reduces branch divergence.
- Achieves significant speedups compared to the state-of-the-art learned indexes.

Approximating Probabilistic Group Steiner Trees in Graphs

Sept 2021 to January 2022

Research Assistant, Renmin University of China, Supervisor: Yahui Sun

- Defined the problem of probabilistic group steiner tree (PGST) while the existing work focuses on certain cases.
- Developed new approximation algorithms to produce considerably better solutions for the PGST problem than state-of-the-art algorithms.
- Devised and implemented the parallel version of pruned landmark labeling algorithm and achieved significant speedups.

Exploring Query Processing on CPU-GPU Integrated Edge Device

Sept 2021 to Nov 2021

Research Assistant, Renmin University of China, Supervisor: Feng Zhang

- Developed the first framework enabling fine-grained query processing on the CPU-GPU integrated edge device.
- Unveiled the challenges for fine-grained query processing at the edge and provide a series of solutions.

PUBLICATION

1. [ASPLOS'23 (Revision)] "Space Efficient TREC for Enabling Deep Learning on Microcontrollers";
Jiesong Liu, Feng Zhang, Jiawei Guan, Hsing-Hsuan Sung, Xiaoyong Du, Xipeng Shen; 28th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2023. [Paper]
2. [WWW'23 (Submitted)] "G-Learned Index: Enabling Efficient Learned index on GPU";
Jiesong Liu, Feng Zhang, Lv Lu, Xiaoyong Du, Guoliang Li, Dong Deng. [Paper]
3. [VLDB'23] "Approximating Probabilistic Group Steiner Trees in Graphs";
Shuang Yang, Yahui Sun, **Jiesong Liu**, Xiaokui Xiao, Ronghua Li, Zhewei Wei; 49th International Conference on Very Large Data Bases (VLDB), 2023. [Paper]
4. [TPDS'22] "Exploring Query Processing on CPU-GPU Integrated Edge Device";
Jiesong Liu, Feng Zhang, Hourun Li, Dalin Wang, Weitao Wan, Xiaokun Fang, Jidong Zhai, Xiaoyong Du; IEEE Transactions on Parallel and Distributed Systems (TPDS), 2022. [Paper]
5. [NeurIPS'22] "TREC: Transient Redundancy Elimination-based Convolution";
Jiawei Guan, Feng Zhang, **Jiesong Liu**, Hsing-Hsuan Sung, Ruofan Wu, Xiaoyong Du, Xipeng Shen; 36th Conference on Neural Information Processing Systems (NeurIPS), 2022. [Paper]

PERSONAL PROFILE

- Specialties: Machine Learning, Parallel Computing, CPU-GPU Integrated Architectures
- Programming Language: C, C++, CUDA, OpenCL, OpenMP, Python, SQL, Verilog
- Language: English and Chinese