

AmirHossein Sojoodi

PHD CANDIDATE · HPC RESEARCHER · GPU SOFTWARE ENGINEER

Kingston, ON, Canada

Available upon request | Email: amir.sojoodi@gmail.com | Website: amirsojoodi.github.io | GitHub: [amirsojoodi](https://github.com/amirsojoodi) | LinkedIn: [amirsojoodi](https://www.linkedin.com/in/amirsojoodi)

Summary

As a researcher, involved in various areas mainly focused on performance optimizations in MPI and UCX. More specifically, improving GPU-aware communications in CUDA-enabled OpenMPI and UCX libraries, for both point-to-point and collective communications including message matching engine. As a Software Developer, work with parallel-processing platforms and programming models (such as MPI) and have extensive knowledge in various programming languages and frameworks like C/C++, CUDA, OpenMP, WebGPU, etc.

Education

Queen's University

PH.D. IN ELECTRICAL AND COMPUTER ENGINEERING

Kingston, Canada

Jan. 2020 - Jan. 2025

- **Research Area:** High-performance GPU-aware communication in hybrid clusters, supervised by Dr. A. Afsahi

Shiraz University

M.S. IN SOFTWARE ENGINEERING

Shiraz, Iran

Sep. 2012 - Aug. 2015

Shiraz University

B.S. IN SOFTWARE ENGINEERING

Shiraz, Iran

Sep. 2007 - Aug. 2012

Professional Experience

Distributive Co.

GPU SOFTWARE ENGINEER, PART-TIME

Kingston, Ontario

Sep. 2022 - Sep. 2024

- Explore and develop CUDA, WebGPU, and WebAssembly solutions.
- Develop a WebGPU-based math algorithm, called LeGendre Pairs Length 117, achieving +700x speedup compared to the CPU implementation.
- Design and Develop a WebGPU microbenchmark for Distributed Compute Protocol (DCP) systems.

RE Grant Consulting Company

CUDA DEVELOPER, PART-TIME

Kingston, Ontario

Oct. 2023 - Mar. 2024

- Design and implement a CUDA microbenchmark suite for Rockport Networks (CERIO) systems.
- The suite is a C++ CMake project including Stress and Speed tests for various transfer types, sizes, and scenarios.

Microsoft Research (MSR)

RESEARCH INTERN

Seattle, Washington (Remote)

Jun. 2022 - Aug. 2022

- Enable GPU in Distributed Deep Learning (DistDL) framework utilizing CuPy.

Engineering Teaching and Learning Team (ETLT), Queen's University

COURSE DESIGN AND DEVELOPMENT SPECIALIST, PART-TIME

Kingston, Ontario

Sep. 2020 - Apr. 2022

- Maintain and update the courses' learning objectives (CLOs).
- Edit videos and create animations for the Faculty of Engineering and Applied Science (FEAS) courses.
- Automate the courses evaluation process using Python and Queen's survey platform.

HPC Group, CSE Dept, Shiraz University

XENSERVER AND LINUX ADMINISTRATOR

Shiraz, Iran

Sep. 2018 - Nov. 2019

- Setup and maintain Apache Hadoop, Spark, Tez, and Ignite clusters.

Aria Hamrah Samaneh

SYSTEM ADMINISTRATOR, DATA VISUALIZATION AND JAVA BACKEND DEVELOPER

Shiraz, Iran

Apr. 2016 - Feb. 2017

- Develop data visualization dashboards using Tableau.

Information and Communication Technology Center (ICTC), Shiraz University

BUSINESS INTELLIGENCE (BI) DEVELOPER

Shiraz, Iran

Oct. 2015 - Apr. 2016

- Develop BI solutions based on Microsoft Power BI.

HPC Group, CSE Dept, Shiraz University

XENSERVER AND LINUX ADMINISTRATOR

Shiraz, Iran

Oct. 2013 - Oct. 2015

- Setup and maintain various services such as Nexus repository manager, apt cacher, squid, SVN, project manager, etc.

Skills

Programming	C, C++, Python, JavaScript, Rust, TXL, Matlab, Java, Assembly, and Shell
Platforms/APIs/Libs	CUDA, OpenMP, MPI, UCX, WebGPU, Pthreads, MapReduce, PyTorch, NumPy/CuPy, Apache Ignite
Misc. Tools	Git, Perf, Valgrind, NVIDIA Nsight Tools, Arm DDT, Auto Tools, CMake, Nexus, \LaTeX
Operating Systems	Ubuntu, Debian, CentOS
Video Editing	Camtasia, Corel Video Studio, Proshow Producer, Adobe AfterEffects
Languages	Farsi (maternal), English (fluent), French (basic)

Publications

- [1] Amirhossein Sojoodi, Yiltan Hassan Temucin, Ahmad Afsahi, “Enhancing Intra-Node GPU-to-GPU Performance in MPI + UCX through Multi-Path Communication”*Proceedings of the International Workshop on Extreme Heterogeneity Solutions (ExHET)*, pp. 1–6, doi: 10.1145/3642961.3643800 2024
- [2] Pedram Alizadeh, Amirhossein Sojoodi, Yiltan Hassan Temucin, Ahmad Afsahi, “Efficient Process Arrival Pattern Aware Collective Communication for Deep Learning”*Proceedings of the European MPI Users’ Group Meeting (EuroMPI)*, pp. 68–78, doi: 10.1145/3555819.3555857 2022
- [3] Philipp A. Witte, Russell J. Hewett, Kumar Saurabh, AmirHossein Sojoodi, Ranveer Chandra, “SciAI4Industry – Solving PDEs for industry-scale problems with deep learning”, *arXiv* (2022), pp. 1–11, doi: 10.48550/arXiv.2211.12709 2022
- [4] Yiltan Hassan Temucin, Amirhossein Sojoodi, Pedram Alizadeh, Ahmad Afsahi, “Efficient Multi-Path NVLink / PCIe-Aware UCX based Collective Communication for Deep Learning”*Proceedings of the IEEE Symposium on High-Performance Interconnects (HOTI)*, pp. 1–10, doi: 10.1109/HOTI52880.2021.00018 2021
- [5] Yiltan Hassan Temucin, Amirhossein Sojoodi, Pedram Alizadeh, Benjamin W Kitor, Ahmad Afsahi, “Accelerating Deep Learning using Interconnect-Aware UCX Communication for MPI Collectives”, *IEEE Micro* (2021), pp. 1–9, doi: 10.1109/MM.2022.3148670 2021
- [6] Majid Salimi Beni, Amir Hossein Sojoodi, Farshad Khunjush, “A GPU-Enabled Extension for Apache Ignite to Facilitate Running Genetic Algorithms”*Proceedings of the International Symposium on Computer Architecture and Digital Systems (CADS)*, pp. 1–8, doi: 10.1109/CADS50570.2020.9211857 2020
- [7] Amir Hossein Sojoodi, Majid Salimi Beni, Farshad Khunjush, “Ignite-GPU: a GPU-enabled in-memory computing architecture on clusters”, *Journal of Supercomputing* (2020), pp. 1–28, doi: 10.1007/s11227-020-03390-z 2020

Honors & Awards

2024-03	Best Paper Award , International Workshop on Extreme Heterogeneity Solutions (ExHET)	Edinburgh, UK
2020-09	Parya Scholarship , Parya Trillium Foundation	Canada
2019-01	Best T.A. , CSE Department students’ poll	Shiraz University
2016-08	9th Place , National IoT Hackathon	IUST, Tehran
2015-02	Silver Medal (with B. Ahmadi and M. R. Katebzadeh) , 7th National JavaChallenge	Sharif University
2012-07	Gold Medal in Team Section , Chess Games, South of Iran Universities	Shiraz University
2010-07	5th Place (with M. Asadi) , Students Competitions (Ms Pacman Intelligent Controller)	IEEE CIG (Online)
2010-07	5th Place (with S. Kazemi and M. Saeedi) , Kashan 2nd International Programming contest (ACM)	Kashan University

Selected Certificates

2024-08	Modern C++ Programming , Compute Ontario Summer School
2022-08	Fundamentals of Accelerated Computing with CUDA Python , NVIDIA Deep Learning Institute (DLI)
2022-08	Fundamentals of Deep Learning , NVIDIA Deep Learning Institute (DLI)
2022-08	Accelerating CUDA C++ Applications with Concurrent Streams , NVIDIA Deep Learning Institute (DLI)
2021-08	Debugging and Performance Tuning , SCINET Summer Workshop
2021-07	Summer Workshops , PUMPS+AI at Barcelona Supercomputing Center
2021-06	Modern C++ and Parallel Programming , SHARCNET HPC Summer Workshop
2019-08	Software Testing , Udacity
2019-07	Software Development Process , Georgia Tech & Udacity
2019-07	Design Patterns in Java , Addison-Wesley, Livelessons
2014-11	Intro to Parallel Programming with GPUs , Udacity and NVIDIA

Teaching Experience

2024 Fall **ELEC 278 - Fundamentals of Information Structure**, T.A. of Dr. Tom Dean
 2024 Win **ELEC 374 - Digital Systems Engineering**, T.A. of Dr. Ahmad Afsahi
 2023 Fall **ELEC 379 - Introduction to Algorithms**, T.A. of Dr. Naraig Manjikian
 2023 Win **ELEC 374 - Digital Systems Engineering**, T.A. of Dr. Ahmad Afsahi
 2022 Fall **ELEC 278 - Fundamentals of Information Structure**, T.A. of Dr. Jianbing Ni
 2022 Win **ELEC 374 - Digital Systems Engineering**, T.A. of Dr. Ryan Grant
 2021 Fall **ELEC 278 - Fundamentals of Information Structure**, T.A. of Dr. A. Elwakeel
 2021 Win **ELEC 374 - Digital Systems Engineering**, T.A. of Dr. Sean Whitehall
 2020 Fall **ELEC 278 - Fundamentals of Information Structure**, T.A. of Dr. D. Athersych
 2020 Win **ELEC 374 - Digital Systems Engineering**, T.A. of Dr. Ahmad Afsahi
 2018 Fall **GPU Programming**, T.A. of Dr. Farshad Khunjush
 2016 Win **Introduction to Object Oriented Programming with Java**, Lecturer
 2015 Fall **GPU Programming**, T.A. of Dr. Farshad Khunjush
 2014 Win **Multicore Programming**, T.A. of Dr. Farshad Khunjush
 2013 Fall **GPU Programming**, T.A. of Dr. Farshad Khunjush
 2013 Sum **Software Engineering Lab**, Lecturer
 2012 Fall **Operating Systems**, T.A. of Dr. Mohammadreza Moosavi
 2012 Fall **System Programming (Assembly)**, T.A. of Dr. Gholamhossein Dastghaibifard
 2012 Sum **Software Engineering Lab**, Lecturer
 2012 Win **Software Engineering Lab**, Lecturer
 2012 Win **Microprocessors**, T.A. of Dr. Farshad Tajeripour
 2011 Win **Operating Systems**, T.A. of Dr. Sattar Hashemi
 2010 Win **Principles of Programming (C)**, T.A. of Dr. Ali Hamzeh
 2010 Fall **Operating Systems**, T.A. of Dr. Sattar Hashemi
 2010 Win **System Programming (Assembly)**, T.A. of Dr. Sattar Hashemi
 2009 Fall **System Programming (Assembly)**, T.A. of Dr. Sattar Hashemi

Queen's University
 Queen's University
 Queen's University
 Queen's University
 Queen's University
 Queen's University
 Queen's University
 Queen's University
 Queen's University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University
 Shiraz University

Selected Presentations

Distributive Book Club Meetings

Kingston, Ontario

PRESENTER FOR LeGENDRE PAIRS OPTIMIZATIONS

Aug. 2023

- Introduce the LeGendre Pairs algorithm and its implementation on GPUs.
- Discuss the challenges and optimizations for the algorithm.
- Shared memory utilization, coalesced memory access, warp divergence, and other optimizations.
- Slides can be obtained from [here](#) and [there](#).

Distributive Book Club Meetings

Kingston, Ontario

PRESENTER FOR INTRODUCTION TO GPUS

Oct. 2023

- Introduce the GPUs and their architecture.
- Introduce the CUDA programming model, the GPU memory hierarchy, and the memory access patterns.
- Profiling and performance tuning of GPU applications.
- Slides can be obtained from [here](#) and [there](#).

Selected Extracurricular Activities

Graduate ECE Student Council

Queen's University

PHD REPRESENTATIVE

2023-2024

- Graduate Electrical and Computer Engineering (GECE) student council is the governing body representing the graduate students of Electrical and Computer Engineering at Queen's University.
- I was the PhD representative for the 2023-2024 academic year.

Students' Scientific Group, CSE Department

Shiraz University

EXECUTIVE MEMBER

2011-2014

- In this group, we worked on different seminars and workshops for students.
- Also, each semester, courses and their exams scheduling was done by this group.
- I was a member of the group in 2009, 2010, 2012, and 2014, and its chair during 2011.