Ehsan Yousefzadeh-Asl-Miandoab

Department of Computer Science
IT University of Copenhagen,
Rued Langgaards Vej 7,
2300 **Copenhagen, Denmark**.

E-mail 1: ehyo@itu.dk
E-mail 2: ehsanyusefzadehasl@gmail.com
Homepage: My GitHub Pages Resume

EDUCATION

◇ **IT University of Copenhagen**, Copenhagen, Denmark **2021 - 2025**
— **Ph.D.** in **Computer Science**

* Thesis: ***Collocation and Memory Estimation for Efficient GPU Resource Management in Deep Learning Training***

...

◇ **Sharif University of Technology**, Tehran, Iran **2016 - 2018**
— **M.Sc.** in **Computer Engineering - Hardware** **GPA: 17.20/20**

* Thesis: ***Unifying L1 Data Cache and Shared Memory in GPUs***

...

◇ **University of Tabriz**, Tabriz, Iran **2012 - 2016**
— **B.Sc.** in **Computer Engineering - Hardware** **GPA: 18.42/20**

◇ **Shahid Foroughi High School**, Miandoab, Iran **2008 - 2012**
— High School Diploma in **Mathematics and Physics** **GPA: 19.05/20**

* * *

RESEARCH INTERESTS

◇ Computer Architecture and Systems Optimizations
◇ GPUs and Heterogeneous Systems
◇ Systems and Hardware Optimizations and Accelerators for Deep Learning
◇ High Performance Computing (HPC) and Parallel Systems

* * *

RESEARCH EXPERIENCE

◇ **IT University of Copenhagen** **November 2021 - Now**
- Data Intensive Systems and Applications (DASYA)
- Supervisor: **Prof. Pinar Tozun**
The focus of my research is on Resource-Aware Data Science, with the goal of improving the efficiency and utilization of GPUs for deep learning training workloads. My work has produced insights into GPU profiling, monitoring tools, effective collocation of training tasks to address GPU underutilization, developed a deep-learning-based GPU memory usage estimator tailored for deep learning training jobs, and resulted in a robust, collocation-aware resource manager for training environments. This is an ongoing research project. The results so far include a paper at (*EuroMLSys 2023*, a paper at (*EuroMLSys 2024*, and a paper at (*DEEM 2023*. Currently, two additional work are under submission.

◇ **Sharif University of Technology** **Sep. 16 - Aug. 19/ Sep. 20 - Feb. 21**
- High Performance Computing Architectures and Networks (HPCAN)
- Supervisor: **Prof. Hamid Sarbazi-Azad**
The focus of my research was to introduce an energy and area efficient on-chip memory design with negligible performance overhead for GPU Streaming Multiprocessors (SMs). I implemented and analyzed proposed designs using simulators such as GPGPU-Sim, and self-written C codes.
The results of those efforts were two papers, one published in *IEEE Transactions on Parallel and Distributed Systems (TPDS)* in 2022, and another in *Proceedings of the ACM on Measurement and Analysis of Computing Systems (MACS)* in 2022.

* * *

PUBLICATIONS

◇ *Under-submission.*

- **GPUMemNet: Deep Leaning-based Estimation of GPU Memory Requirement for Neural Network Training Tasks**
- **CARMA: Collocation-Aware Resource Manager with GPU Memory Estimator**
- **DP-Morph: Improving the Privacy–Utility–Performance Trade-off for Differentially Private OCT Segmentation**

◇ "**PiMPiC: An Overlap-Aware Contrastive Learning Framework for 3D Patch-Based Medical Image Segmentation**," got accepted to **DEMI workshop** @ *MICCAI 2025*

◇ Ties Robroek, **Ehsan Yousefzadeh-Asl-Miandoab**, and Pnar Tzn. "**An Analysis of Collocation on GPUs for Deep Learning Training.**" In Proceedings of the Fourth Workshop on Machine Learning and Systems, 2024 (**EuroMLSys '24**).

◇ Ties Robroek, Aaron Duane, **Ehsan Yousefzadeh-Asl-Miandoab**, and Pinar Tozun. "**Data Management and Visualization for Benchmarking Deep Learning Training Systems.**" In Proceedings of the Seventh Workshop on Data Management for End-to-End Machine Learning, 2023. (**DEEM '23**).

◇ **Ehsan Yousefzadeh-Asl-Miandoab**, Ties Robroek and Pinar Tozun. "**Profiling and Monitoring Deep Learning Training Tasks.**" Proceedings of the 3rd Workshop on Machine Learning and Systems, 2023 (**EuroMLSys '23**).

◇ Darabi, Sina, Negin Mahani, Hazhir Baxishi, **Ehsan Yousefzadeh-Asl-Miandoab**, Mohammad Sadrosadati and Hamid Sarbazi-Azad. "**NURA: A Framework for Supporting Non-Uniform Resource Accesses in GPUs.**" Proceedings of the ACM on Measurement and Analysis of Computing Systems, 2022 (**POMACS '22**).

◇ Darabi, Sina, **Ehsan Yousefzadeh-Asl-Miandoab**, Negar Akbarzadeh, Hajar Falahati, Pejman Lotfi-Kamran, Mohammad Sadrosadati, and Hamid Sarbazi-Azad. "**OSM: Off-Chip Shared Memory for GPUs.**" IEEE Transactions on Parallel and Distributed Systems, 2022 (**TPDS '22**).

* * *

TEACHING EXPERIENCE

- Department of Computer Science, IT University of Copenhagen

Course (level)	Role	Teacher	Semester	Year
Operating Systems and C (bachelor's)	Guest Lecturer	Prof. Willard Rafnsson	Fall	2024
Computer Systems Performance (bachelor's)	TA, Guest Lecturer	Prof. Pinar Tozun	Spring	2022, 2023, 2024
Large Scale Data Analysis (master's)	Guest Lecturer	Dr. Maria Sinziiana Astefanoaei	Spring	2022, 2023, 2024
Big Data Management (bachelor's)	Guest Lecturer	Dr. Maria Sinziiana Astefanoaei	Fall	2022, 2023, 2024
Advanced Data Systems (master's)	TA, Guest, Lecturer	Prof. Pinar Tozun	Fall	2022, 2023, 2024
Deep Learning: A Computational Efficiency Perspective (master's, PhD)	TA	Prof. Pinar Tozun	Spring	2022

- Department of Computer Engineering, Sharif University of Technology

Course (level)	Role	Teacher	Semester	Year
Computer Architecture and Systems Lab. (bachelor's)	Head TA	Prof. Sarbazi-Azad	Summer	2017
Microprocessor/ Microcontroller Lab. (bachelor's)	Head TA	Prof. Sarbazi-Azad	Fall	2017
Computer Architecture (bachelor's)	Head TA	Prof. Jahangir	Spring	2018
Computer Architecture and Systems Lab. (bachelor's)	Head TA	Prof. Jahangir	Spring, Summer	2018
Microprocessor/ Microcontroller Lab. (bachelor's)	Head TA	Prof. Jahangir	Spring, Summer	2018

- Department of Electrical and Computer Engineering, University of Tabriz

Course (level)	Role	Teacher	Semester	Year
Logic Circuit Design (bachelor's)	TA	Prof. Mina Zolfy	Spring	2016

◇ **Private Tutorship**

Course	Role	Teacher	Semster	Year
Teaching first-year undergraduates how to program with C	Teacher	-	Spring, Summer	2018

* * *

HONORS AND AWARDS

◇ **Ranked 1st in terms of cumulative GPA among Computer Hardware Engineering Students, University of Tabriz (2016)**

◇ **Admitted to the M.Sc. program at Sharif University of Technology as an Exceptional Talented Student (2016)**

* * *

TECHNICAL SKILLS

◇ **Programming Languages :** Python, C++, C, Java, MATLAB, VB, Bash
◇ **Hardware Description Language (HDL) :** Verilog
◇ **Simulation Tools :** Altera Quartus II, Mentor Modelsim, Proteus Design Suite, Synopsys Hspice and Pspice, GPGPU-Sim.
◇ **Document Preparation :** MS Office, LaTeX.
◇ **Machine Learning Frameworks :** TensorFlow, Keras, PyTorch
◇ **Profiling and monitoring tools:** NVIDIA compute, NSight systems, Machine learning frameworks' profilers, e.g., TensorFlow profiler with TensorBoard, Linux terminal based system monitors
◇ **Other :** Linux Terminal and Bash scripting, Git, MVC, MySQL, VB programming for Microsoft Excel

* * *

PROJECTS

◇ **Resource-aware Data Science**
- This is the wider umbrella encompassing multiple more narrow projects, which each one includes code contributions to the community, under it including:

1. Benchmarking profiling/ monitoring tools
2. Benchmarking NVIDIA A100 under collocating multiple deep learning training tasks
3. Building an deep-learning-based GPU memory estimator for deep learning training tasks

4. Design and implementing a resource manager/ scheduler system prototype being aware of the collocation for tackling energy inefficiency and low GPU utilization challenge for deep learning training tasks.

◇ **A Unified On-chip Memory for Shared and L1 Cache Accesses for GPUs in C++**
- For my master's thesis, I worked on the implementation of a unified on-chip memory to serve both the shared memory and level one data cache accesses. I logged the addresses generated by the GPGPU-Sim simulator and studied their locality, lifetime, read-after-write frequency behavior by developing C++ programs. Then, based on the observations, a unified structure with locking capability replaced level one data cache and shared memory in the GPGPU-Sim simulator.

◇ **Morris Mano's Basic Computer**
- During the Computer Architecture course in my bachelor's, I designed and implemented a basic computer in the Quartus II Schematics. Then in my master's when I was a computer architecture course's TA, I implemented this basic computer once more in Verilog HDL with more details. It can be checked on my GitHub page Here.

◇ **Cache with different Configurations**
- I implemented direct-mapped and set-associative cache for my advanced computer architecture course. Then, I did experiments on it to observe the effects of different policies on hit/ miss rate. It is accessible here on GitHub.

◇ **Python Basics; Data Structures and Algorithms, Programming with CUDA, Neural Networks Tutorials with Examples on GitHub**
- In my free time, I enjoy creating tutorials for learners who appreciate clear and accessible explanations. This practice not only helps others but also allows me to review and strengthen my own understanding, following the principle that the best way to learn is to teach. Some of these tutorials are available on my GitHub and YouTube

◇ **Web Development Projects**
- I have worked as a back-end developer on multiple web development teams, where I was responsible for writing database queries to retrieve data required by the front-end forms. In addition, I developed and maintained APIs to deliver structured data to mobile developers. I also contributed to an open-source URL-shortener project on GitHub. Furthermore, I developed a lightweight API framework in Go, aimed at helping beginners get started with writing API endpoints by providing a simple and approachable starting point.

◇ **Design and Implementation of An Efficient Archiving System in Excel and some other automation systems in MS**
- During my conscription, I designed, implemented, and organized an efficient document archiving system using Excel and Access with VBA, tailored to document types and their associated transactions. Prior to this, the workflow was largely manual and highly time-consuming. By automating key processes, I significantly improved operational efficiency, which reduced manual workload to the extent that the staff had substantial idle time.

* * *

LANGUAGE SKILLS

◇ **English**: Fluent.

Taken Test	Reading	Listening	Speaking	Writing	overall
TOEFL on July 10, 2021	26	26	24	24	100

◇ **Azerbaijani**: Mother Tongue.
◇ **Persian**: Native.
◇ **Danish**: Elementary.

PROFESSIONAL SERVICE

◇ **Shadow Reviewer** (*EuroSys 2026*). Selected to participate in a mentoring program to gain experience in reviewing systems research papers.