

Hazem Abdelhafez

+1 (778) 316 2491 - hazem@ece.ubc.ca

Vancouver, Canada

Webpage: hazemabdelhafez.github.io

PROFESSIONAL EXPERIENCE

University of British Columbia, Vancouver, Canada

Research Assistant, September 2016 - Present

- Characterize and model deep-learning workloads on heterogeneous computing platforms to identify performance bottlenecks and optimize power-consumption to promote energy-efficient computing.
- Develop a statistical analysis-based methodology to study the inter- and intra-node performance and power consumption variability amongst identical Edge computing devices.
- Study and instrument PyTorch framework Virtual Machine stack interpreter, part of the TorchScript module, to insert performance and kernel-level information gathering code that allows us to study deployed deep learning networks' runtime behavior and hot kernels on CPU/GPU-based systems.

Huawei Heterogeneous Compilers Lab, Markham, Canada

Research and Development Intern, January 2020 - June 2020

- Developed and implemented a proof-of-concept for efficient compiler-based data-dependency management algorithm for in-order instruction issue processors that led to up to 15% and average of 5% improvement in end-to-end latency across several benchmarks.
- Created and submitted an accepted patent application (currently in the filling process) for the algorithm.
- Evaluated the impact of re-ordering operands of associative instructions as a compiler optimization phase across several benchmarks, which allowed the management to make a well-informed decision before allocating resources and time for full scale implementation.
- Implemented and contributed several components from SPIR-V specification in MLIR project (part of the open-source LLVM compiler project) which allowed me to achieve the contributor status in the LLVM project.

Oak Ridge National Lab, Oak Ridge, United States

Astro Program Participant (Research Intern), May 2018 - August 2018

- Built an open-source profiling-based tool that allows developers and system designers to analyze the performance of GPU-accelerated applications in fields such as high-performance computing.
- Created and developed an analytical model for projecting intra-node connectivity impact on data transfer time in GPU-accelerated systems using data sheet information. The model projects data transfer times on next-generation nodes without actual deployment with error ranging from 19% to 23%.

Avidbeam, Cairo, Egypt

Software Engineer, September 2014 - September 2016

- Took a leadership role in the development of a large-scale video analytics platform (ATUN). This platform currently serves as a base for scalable computer vision algorithms shipped by Avidbeam.
- Contributed to joint research and development effort between Avidbeam and Intel corporation to create a real-time large-scale video analytics platform and participated in Network Function Virtualization project proposal which led to a collaborative proof-of-concept between international telecommunication service provider and Avidbeam.

Intel Labs, Cairo, Egypt

Research and Development Engineer, July 2012 – April 2014

- Developed several demos for Licensed Shared Access systems that were demonstrated at various Intel technology events such as Intel Developer Forum (IDF) and Research at Intel (R@I).
- Created algorithms for dynamically allocating radio spectrum using predictive machine learning algorithms which led to two publications and three granted patents.

University of British Columbia, Vancouver, Canada

Teaching Assistant, Spring 2017 to 2023

- Worked as a teacher assistant for three semesters for the Design of Distributed Software Applications undergraduate course (CPEN431). Provided guidance and assistance to students during lab hours and online to help them build an end-to-end distributed key-value store application.
- Built an automated grading software to reduce marking effort and time; And developed a software-based management tool for PlanetLab distributed computing resources. Managed our AWS cloud infrastructure that serves the course's UI web-based services and testing software.

EDUCATION

University of British Columbia, Vancouver, Canada

- Ph.D. in Electrical and Computer Engineering, September 2016 - Present, GPA: A+
- Awards received: Four Year Fellowship (2016 - 2020), and Graduate Support Initiative (2017).

Cairo University, Giza, Egypt

- M.Sc. in Electronics and Telecommunications Engineering, March 2014 - September 2016, GPA: A-
- B.Sc. in Electronics and Telecommunications Engineering, September 2007 - May 2012, Grade: (76.6%)

SELECTED PUBLICATIONS

- **Hazem A. Abdelhafez**, Hassan Halawa, Amr Almoallim, Amirhossein Ahmadi, Karthik Pattabiraman, and Matei Ripeanu. Characterizing Variability at the Edge: A Methodology and Case Study. Proceedings of the seventh ACM/IEEE Symposium on Edge Computing (SEC), Dec. 2022. Acceptance rate 27%.
- **Hazem A. Abdelhafez**, Hassan Halawa, Mohamed Osama Ahmed, Karthik Pattabiraman, and Matei Ripeanu. MIRAGE: Machine Learning-based Modeling of Identical Replicas of the Jetson AGX Embedded Platform. Proceedings of the sixth ACM/IEEE Symposium on Edge Computing (SEC), Dec. 2021. Acceptance rate 27%.
- **Hazem A. Abdelhafez**, Hassan Halawa, Karthik Pattabiraman, and Matei Ripeanu. Snowflakes at the Edge: A Study of Variability among NVIDIA Jetson AGX Xavier Boards. Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking (EdgeSys). ACM, New York, USA, April 2021.
- **Hazem A. Abdelhafez**, Christopher Zimmer, Sudharshan Vazhkudai, and Matei Ripeanu. AHEAD: A Tool for Projecting Next-Generation Hardware Enhancements on GPU-Accelerated Systems. In 2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), May 20th, pp. 583-592, (*Best Paper Award*).
- **Hazem A. Abdelhafez** and Matei Ripeanu. Studying the Impact of CPU and Memory Controller Frequencies on Power Consumption of the Jetson TX1. 2019 IEEE International Conference on Fog and Mobile Edge Computing (FMEC), June 10th.
- Halawa H., **Hazem A. Abdelhafez**, Boktor A., Ripeanu M., NVIDIA Jetson Platform Characterization. In EuroPar 2017. Lecture Notes in Computer Science, vol. 10417. Springer. Acceptance rate: 25%.

PROGRAMMING SKILLS

- Use on a regular basis: Python, and Java
- Use intermittently: C++, Bash
- Somewhat Familiar: CUDA
- Previous coursework: MatLab, Erlang

PATENTS

- **Hazem A. Abdelhafez**, Ning Xie, Ahmed Mohammed ElShafiey Mohammed Eltantawy. Devices, methods, and media for efficient data dependency management for in-order issue processors. US Patent Application: US20220107811A1, April 2022.
- Mohamed El-Refaey, Markus Dominik MECK, Kirk D. SMITH, **Hazem Abdelmegeed**, Hani H. El-Gebaly, Norhan M. Osman. Spectrum reclaiming in a leased spectrum system. US Patent US9326156B2, April 2016.
- Mohamed El-Refaey, **Hazem Abdelmegeed** and Hani H. El-Gebaly. "Coordination of spectrum usage rights among entities." US Patent Grant US9313665B2, April 2016.
- Mohamed El-Refaey, Norhan M. Osman, **Hazem Abdelmegeed**, Hani Elgebaly. "Requesting extra spectrum." US Patent Application Number US20160095016 A1, 31 March 2016.