50 Central Campus Dr #3163
Salt Lake City, UT 84112
+1-951-222-9581
ashfaq@cs.utah.edu
aagontuk.github.io

Md Ashfaqur Rahaman

Research Interests

Kernel bypass, disaggregated memory, hardware offloading and systems for machine learning

Education

Aug. 2021 **Ph.D. in Computer Science**, *University of Utah*, Salt Lake City, Utah, USA
Advisor: Ryan Stutsman

2012-2019 **B.Sc. in Naval Architecture and Marine Engineering**, *Bangladesh University of Engineering and Technology (BUET)*, Dhaka, Bangladesh

Experience

Research

Summer 2024 **Research Intern**, *Hewlett Packard Labs*, Milpitas, California
Optimizing the communication infrastructure to support low-latency high-throughput LLM inference

2021-Present **Graduate Research Assistant**, *Utah Scalable Computer Systems Lab*, University of Utah, Utah
- A new efficient, secure, and scalable framework for remote memory access and function offloading exploiting the programmability and offloading capability of smartNICs
- A new software architecture for building services that centers around coherent accelerators and rack-scale shared memory

2019-2021 **Research Assistant**, *Prof. Baris Kasikci's Lab*, University of Michigan, Ann Arbor
Mentor: Tanvir Ahmed Khan
Load-time code layout optimization of large application binaries in warehouse scale computers

2018-2019 **Research Assistant**, *Climate Modeling and Simulation Lab*, IWFM, BUET
Advisor: A.K.M. Saiful Islam
I worked as a system developer in real-time Flash Flood Early Warning System (FFEWS) project

Professional

2018-2019 **Software Engineer**, *NextGen DigiTech*, Dhaka
I worked on NextGen Tower, a desktop application for designing wind turbines. I contributed in the core software architecture and developed the GUI.

2017-2018 **Firmware Engineer**, *2RA Technology Limited*, Dhaka
I worked on various embedded systems projects based on Raspberry Pi and AVR Microcontrollers.

Selected Research Projects

2021-Present **NIC Accelerated Active Messaging**
RDMA is gaining popularity in datacenters for high-throughput and low-latency network communication for building dis-aggregated systems. However, there are many issues that are holding RDMA back from being widely deployed. Current RDMA verbs are limited for diverse workloads, they are difficult to program, and multiple round trips are required to do complicated memory operations e.g. walking a hash table. We are working on creating a new network abstraction to make remote memory access more efficient, secure, and scalable exploiting the programmability and offloading feature of smartNICs. In our system, NIC offloads can be written in a high level language ensuring easy programmability, then this will be converted into verifiable bytecode tailored to a specific workload and run on the NIC or host based on the dynamic load. Program transformation, dynamic decisions, all these will be transparent to the developer.

2023-Present **Software Architectures for Large-Scale Coherent Shared Memory**

Emerging standards for cache coherent accelerators (e.g. CXL) will soon transform how memory-intensive large-scale systems are developed. Cache coherent accelerators are programmable (via FPGAs), and they can interpose on CPU memory accesses at cache line granularitiy. Low-overhead, granular access tracking with these coherent accelerators enable efficient memory disaggregation. But disaggregation alone will not fundamentally change application architecture. In this work we are designing a new software architecture for building large scale services that centers around coherent accelerators and rack-scale shared memory.

Selected Courses

Fall 2024 Advanced Compilers, University of Utah
Spring 2024 Advanced Computer Architecture, University of Utah
Fall 2023 Advanced Networking, University of Utah
Spring 2022 Software Security, University of Utah
Fall 2021 Advanced Operating Systems, University Of Utah

Teaching Assistantship

Fall 2022 Distributed Systems, University Of Utah
Spring 2022 Operating Systems, University Of Utah

Services

2022 Artifact Evaluation Committee Member, OSDI'22
2021 Artifact Evaluation Committee Member, SOSP'21

Bachelor Thesis

Title *Power Efficient Remotely Operated Underwater Vehicle Using Buoyancy Chambers*
Supervisor Dr. Md. Mashud Karim

Skills

Languages C, C++, Rust, CUDA, Python, Go, Assembly(ARM, X86)
Tools CXL, eBPF, RDMA, DPDK, Pytorch, NVIDIA DOCA, LLVM, Linux perf, BOLT
Platforms NVIDIA BlueField 2, AVR Microcontrollers, Raspberry Pi, Arduino
Text Editing Vim, LaTeX