Jahanzeb Maqbool Hashmi

https://jahanzeb-hashmi.github.io

Summary

With over 15 years of experience in High Performance Computing (HPC), I drive technical strategy for next-generation GPU and datacenter systems, collaborating with cross-functional teams spanning Hardware, Software, and Product to define architecture direction and evaluate design trade-offs from chip-level to rack-scale NVLink platforms. My work has shaped multiple NVIDIA GPU generations—from Hopper through Rubin and beyond—through architectural feature exploration, hardware/software co-design, workload modeling, and competitive analysis, with recommendations presented to executive leadership, including the CEO.

Education

Ph.D. in High Performance Computing, The Ohio State University, Columbus, Ohio, USA 2015 – 2020
M.S. in Computer Engineering, Ajou University, Suwon, South Korea 2012 – 2014
B.S. Information Technology, National University of Science and Technology, Islamabad, Pakistan 2007 – 2011

Professional Experience

- **Senior High Performance Compute Architect,** NVIDIA Corporation, USA. March 2021 – Present
  - Leading cross-org path-finding efforts to analyze design trade-offs of building scale-up and scale-out systems. My work provides insights into NVIDIA's product positioning and is directly reported to E-staff and the CEO of the company.
  - Leading the architectural feature exploration and value proposition of next-generation NVIDIA GPU architecture for HPC, Industrial, and AI applications. I have contributed to several design PORs and provided quantitative data to get them approved across several generations of NVIDIA products starting from Hopper to Rubin and several upcoming products.
  - Identify application scaling challenges and co-design applications across HW/SW design space to achieve maximum performance, energy efficiency, and TCO on NVIDIA platforms.
  - Model business critical workloads to establish performance targets for NVIDIA's upcoming architectures to serve as north star for system bringup teams.
  - Lead the HPC competitive analysis working group, projecting application and system performance — from kernel/chip-level analysis to full-scale datacenter architecture — on competitive GPU/accelerator products
- **Senior Research Associate,** The Ohio State University, USA. June 2020 – March 2021
  - Worked on the design and development of high-performance MPI library for next-generation HPC and Cloud systems with multi-core CPUs (AMD Rome, Intel Xeon, IBM POWER9, ARM A64FX) and many-core GPUs (NVIDIA, AMD)
  - Led the design and development of a generalized hierarchical MPI collective communications framework for modern CPU and GPU systems
  - Led the design and development of MVAPICH2-GDR, a high-performance GPU-aware MPI library, for NVIDIA and AMD based multi-GPU systems
  - Mentored Ph.D. and Masters students' thesis on diverse research areas such as communication over multi-core CPUs and many-core GPUs, novel designs for message-passing communication over high-speed interconnects.
- **Graduate Research Associate,** The Ohio State University, USA. August 2016 – May 2020
  - Designed an adaptive and topology aware algorithm for mapping of MPI processes to hardware cores by capturing the communication-patterns of AI and HPC applications
  - Worked collaboratively on efficient parallelization of large-scale distributed DNN training (data and model parallel) on CPU and GPU systems
  - Designed and developed a truly zero-copy XPMEM-based inter-process (IPC) communication with shared address space principle targeting manycore architectures
  - Designed a novel algorithm to cache data layouts to help mitigate the performance costs of layout translation of MPI derived datatypes. This work was nominated for the Best Paper Award at IPDPS'19
  - Worked on PGAS libraries e.g., OpenSHMEM, UPC++ and task-based programming models e.g., Kokkos with MPI backend
- **Department Fellow,** The Ohio State University, USA. Aug 2015 - July 2016

- While working as a Department of CSE Fellow, I made research contributions to *PGAS over MPI* project at Network Based Computing Laboratory where I designed, developed, and analyzed an MPI based communication conduit for UPC++ asynchronous PGAS programming model.

Technical Skills

- **GPU & Accelerator Systems:** CUDA/HIP, NCCL collective communication, GPU-aware MPI/SHMEM, NVIDIA Nsight Suite, multi-GPU/multi-node performance profiling
- **Performance Engineering:** Roofline modeling, FLOPs/bandwidth/interconnect throughput analysis, performance engineering
- **System Design:** Hardware/software co-design for HPC and AI, end-to-end system design and optimization for perf, power, and TCO
- **Programming Languages:** C, C++, Python, Java

Select Publications

For complete list of publications, please refer to my Google Scholar.

1. B. Ramesh, **J. Hashmi**, S. Xu, A. Shafi, M. Ghazimirsaeed, M. Bayatpour, H. Subramoni, and D. K. Panda. "Towards Architecture-aware Hierarchical Communication Trees on Modern HPC Systems", in proceeding of *28th IEEE International Conference on High Performance Computing, Data, Analytics and Data Science*, Dec. 2021. [Best Paper Finalist]

2. **J. Hashmi**, C. Chu, S. Chakraborty, M. Bayatpour, H. Subramoni, and DK Panda. "FALCON-X: Zero-copy MPI Derived Datatype Processing on Modern CPU and GPU Architectures", Journal of Parallel and Distributed Computing (JPDC), Volume 144, October 2020, Pages 1-13, doi.org/10.1016/j.jpdc.2020.05.008

3. **J. Hashmi**, S. Xu, B. Ramesh, M. Bayatpour, H. Subramoni, and D. K. Panda. "Machine-agnostic and Communication-aware Designs for MPI on Emerging Architectures", in proceeding of *34th IEEE International Parallel and Distributed Processing Symposium* (IPDPS '20) , May 2020

4. **J. Hashmi**, S. Chakraborty, M. Bayatpour, H. Subramoni, D. K. Panda. "FALCON: Efficient Designs for Zero-copy MPI Datatype Processing on Emerging Architectures", in proceeding of *33rd IEEE International Parallel and Distributed Processing Symposium* (IPDPS '19) , May 2019 [Best Paper Finalist]

5. **J. Hashmi**, S. Chakraborty, M. Bayatpour, H. Subramoni, D. K. Panda. "Designing Efficient Shared Address Space Reduction Collectives for Multi-/Many-cores", in proceeding of *32nd IEEE International Parallel and Distributed Processing Symposium* (IPDPS '18) , May 2018

6. **J. Hashmi**, S. Chakraborty, M. Bayatpour, H. Subramoni, D. K. Panda. "Design and Characterization of Shared Address Space MPI Collectives on Modern Architectures", in proceeding of *The 19th Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing* (CCGRID '19) , May 2019

7. S. Chakraborty, M. Bayatpour, **J. Hashmi**, H. Subramoni, D. K. Panda. "Cooperative Rendezvous Protocols for Improved Performance and Overlap", in proceeding of *IEEE/ACM International Conference for High Performance Computing, Networking, Storage, and Analysis* (SC '18) , Nov 2018 [Best Paper Finalist]

8. M. Bayatpour, **J. Hashmi**, S. Chakraborty, H. Subramoni, P. Kousha, D. K. Panda. "SALaR: Scalable and Adaptive Designs for Large Message Reduction Collectives", in proceeding of *IEEE International Conference on Cluster Computing* (CLUSTER 2018) , Sep 2018 [Best Paper Award]

9. A. Awan, K. Hamidouche, **J. Hashmi**, and D. K. Panda. "S-Caffe: Co-designing MPI Runtimes and Caffe for Scalable Deep Learning on Modern GPU Clusters", in proceeding of *22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (PPoPP '17) , February 2017