

Hari Subramoni

Assistant Professor, Dept. of Computer Science and Engineering
The Ohio State University, Columbus, OH 43210-1277
Tel: (614)-961-2383, Fax: (614)-292-2911, E-mail: subramoni.1@osu.edu
<http://www.cse.ohio-state.edu/~subramon>

Education

B. Tech.	June 2004	University College of Engineering, Trivandrum	Computer Science and Engineering
Ph.D.	August 2013	The Ohio State University, Columbus, Ohio	Computer Science and Engineering

Employment

Aug'22	–	Current	Assistant Professor	The Ohio State University
Aug'15	–	Aug'22	Research Scientist	The Ohio State University
Aug'13	–	Jul'15	Senior Research Associate	The Ohio State University
Apr'08	–	Jul'13	Graduate Research Assistant	The Ohio State University
Jun'11	–	Aug'11	Research Intern	Lawrence Livermore National Laboratory
Oct'09	–	Jan'10	Research Intern	IBM TJ Watson Research Center
Dec'06	–	Aug'07	Member Technical Staff	Force10 Networks, India
Aug'04	–	Nov'06	Software Engineer	ISoftTech-Sasken Solutions, India

Research Interests

Parallel computer architecture, high-performance networking, conversational AI-based systems, network-based computing, network topology aware computing, QoS, power-aware LAN-WAN communication, network analysis and monitoring, exascale computing, programming models, GPUs and accelerators, Big Data, Deep Learning, high-performance file systems and storage, virtualization and cloud computing.

Awards, Honors, and Recognitions

Research Staff Award - Department of Computer Science and Engineering, The Ohio State University, April 2022.

Lumley Research Award - College of Engineering, The Ohio State University, April 2019.

Invited Keynote Talk - HPC Advisory Council Australian Conference, Aug 2018.

Best Student Paper Finalist - 28th IEEE International Conference on High Performance Computing, Data, Analytics, and Data Science (HiPC'21), Dec 2021.

Best Student Paper Finalist - IEEE International Parallel & Distributed Processing Symposium (IPDPS '21), 2021.

Best Student Poster - Graduate Student Research Poster Competition, The International Supercomputing Conference (ISC'20) 2020.

Best Student Paper Finalist - IEEE International Parallel & Distributed Processing Symposium (IPDPS '19), 2019.

Best Student Poster - Graduate Student Research Poster Competition, The International Supercomputing Conference (ISC'19) 2019.

Best Student Paper Finalist - ACM Graduate Student Research Poster Competition The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2018).

Best Student Paper Finalist - IEEE Cluster 2018.

Best Student Poster - Graduate Student Research Poster Competition, The International Supercomputing Conference (ISC'18), 2018.

Best Student Paper Finalist - IEEE Cluster 2017.

Best Student Paper - European MPI Users' Group Meeting (EuroMPI), 2017.

First Author of Hans Meuer Best Paper Award - The International Supercomputing Conference, (ISC 2017).

Best Student Paper Finalist - IEEE Cluster (Cluster) 2016.

Best Student Poster - ACM Graduate Student Research Poster Competition The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2016).

Outstanding Graduate Research Award - Department of Computer Science and Engineering, The Ohio State University, 2013.

First Author of Best Paper/Best Student Paper Finalist - The International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2012).

Best Capstone Project - Department of Computer Science and Engineering, The Ohio State University, Capstone Project Competition, Autumn 2010.

Honorable Mention - 7th Annual Graduate Student Poster Exhibition, Department of Computer Science and Engineering, The Ohio State University, 2013.

Research Grants

Principal Investigator / Co-Principal Investigator:

1. Co-PI, *Development of a welding CCT diagram and AI techniques for weld defect and weld signature detection and classification*, American Manufacturing and Engineering Co., \$ 240,000, Sep 23 - Aug 25.
2. Co-PI, *CSR: Small: CONCERT: Designing Scalable Communication Runtimes with On-the-fly Compression for HPC and AI Applications on Heterogeneous Architectures*, National Science Foundation, \$600,000, Aug 23 - Jul 26.
3. Co-PI, *Collaborative Research: Frameworks: Performance Engineering Scientific Applications with MVA-PICH and TAU using Emerging Communication Primitives*, Jointly with S. Shende (UO), S. Tomov (UTK), and Y. Cui (SDSC), National Science Foundation, \$3.2M (OSU share - \$900K), Nov 23 - Oct 26.
4. Co-PI, *Optimization of MVAPICH2 MPI Libraries for Broadcom Thor Adapters*, Broadcom, \$335,000, May 23 - April 24.
5. Co-PI, *Enhancement of MVAPICH2 Library with AMD GPUs*, AMD, \$25,000, Jan 23 - Dec 23.
6. Co-PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$125,000, Nov 22 - Oct 23.
7. Co-PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, Aug 23 - July 24.
8. Co-PI, *Student Travel Support for HDF5 User Group (HUG) Meeting*, National Science Foundation, \$10,000, Aug 23 - July 24.
9. Co-PI, *Optimization of MVAPICH2 MPI Libraries for Broadcom Thor Adapters*, Broadcom, \$350,000, May 22 - April 23.
10. PI, *HPC for Buckeyes: Preparing Next-Generation Buckeyes for a Career in HPC*, \$50,000/-, Intel Corporation, Aug 22 - Aug 23.
11. Co-PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$248,264, May 22 - Apr 23.
12. Co-PI, *Buckeye OpenSHMEM*, Los Alamos National Laboratory, \$181,999, Feb 22 - Feb 25.
13. Co-PI, *Enhancement of MVAPICH2 Library with AMD GPUs*, AMD, \$25,000, Jan 22 - Dec 22.
14. Co-PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$125,000, Nov 21 - Oct 22.
15. Co-PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, \$246,281, May 21 - Apr 22.

16. Co-PI, *Enhancement of MVAPICH2 Library with AMD GPUs*, AMD, \$80,000, Jan 21 - Dec 21.
17. Co-PI, *Design and Optimization of MVAPICH2 MPI Library on Rockport Networks*, Rockport Networks, \$135,000, Jan 21 - Dec 21.
18. Co-PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Nov 20 - Oct 21.
19. Co-PI, *PPoSS: Planning: A Cross-Layer Observable Approach to Extreme Scale Machine Learning and Analytics*, National Science Foundation, \$204,484, Oct 20 - Oct 21.
20. Co-PI, *High Performance and Scalable MPI Library for HPC and Deep Learning on Oracle HPC Cloud*, Oracle, \$90,000, Sept 20 - Aug 21.
21. Co-PI, *OAC Core: Small: Next-generation Communication and I/O Middleware for HPC and Deep Learning with Smart NICs*, National Science Foundation, Jointly with D. K. Panda (OSU), \$500K, July 20 - June 23.
22. Co-PI, *Enhancement of MVAPICH2 Library with Broadcom Network Adapter*, Broadcom, \$50,000, June 20 - May 21.
23. Co-PI, *Enhancement of MVAPICH2 Library with AMD GPUs*, AMD, \$80,000, May 20 - Jan 21.
24. Co-PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$232,066, April 20 - Mar 21.
25. Co-PI, *Designing MVAPICH2 Support for Sandia Structural Simulation Toolkit*, Sandia National Laboratory, \$130,000, April 20 - Sept 21.
26. Co-PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Dec 19 - Oct 20.
27. Co-PI (OSU), *Operations & Maintenance for the Endless Frontier, Overall PI - Dan Stanzione, OSU PI - D. K. Panda, Texas Advanced Computing Center (TACC), National Science Foundation, \$60M, OSU Share \$2.5M*, Sept 19 - August 24.
28. Co-PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, May 19 - April 20.
29. Co-PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$227,115, April 19-Mar 20.
30. Co-PI, *Optimizing and Tuning MVAPICH2 for ARM*, ARM, \$100,000, Jan 19 - Dec 20.
31. Co-PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Nov 18 - Nov 19.
32. Co-PI, *Collaborative Research: Frameworks: Designing Next-Generation MPI Libraries for Emerging Dense GPU Systems*, National Science Foundation, Jointly with D. K. Panda, Karen Tomko (OSC), Bill Barth (TACC) and Amit Majumdar (SDSC), \$2.0M (OSU+OSC Share: \$1.4M), Nov'19 - Oct'22
33. Co-PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, Aug 18 - July 19.
34. Co-PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$220,963, April 18 - Mar 19.
35. Co-PI, *SI2-SSI: FAMILI: High Performance and Scalable Fabric Analysis, Monitoring and Introspection Infrastructure for HPC and Big Data*, National Science Foundation, Jointly with D. K. Panda, Karen Tomko (OSC), \$800,000, July'17 - June'20
36. Co-PI, *BD Spokes: SPOKE: MIDWEST: Collaborative: Advanced Computational Neuroscience Network (ACNN)*, National Science Foundation, Jointly with UMICH, CWRU, IU and Northwestern, (OSU Share: \$166,454), Sept'16 - Aug'19
37. Co-PI, *SHF: Large: Collaborative Research: Next Generation Communication Mechanisms exploiting Heterogeneity, Hierarchy and Concurrency for Emerging HPC Systems*, National Science Foundation, Jointly with D. K. Panda, Karen Tomko (OSC), Bill Barth (TACC) and Amit Majumdar (SDSC), \$2.0M (OSU+OSC Share: \$1.17M), Aug'16 - July'19

Senior Personnel:

1. Senior Personnel, *AI Institute for Intelligent CyberInfrastructure with Computational Learning in the Environment (ICICLE)*, National Science Foundation, \$19.99M, Nov'21 - Oct'26
2. Senior Personnel, *Cybertraining Pilot: An Artificial Intelligence Bootcamp for Cyberinfrastructure Professionals*, National Science Foundation, \$299,858, Sept'21 - Aug'23
3. Senior Personnel, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel, \$64,741, Feb'18 - Jan'19.
4. Senior Personnel, *Designing and Optimizing MVAPICH2 MPI Library on FPGA Systems*, Coventry Computer, \$114,000, Jan'18 - Jan'19.
5. Senior Personnel, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Nov'17 - Nov'18.
6. Senior Personnel, *Stampede2: Operations and Maintenance for the Next Generation of Petascale Computing*, University of Texas Austin (National Science Foundation Sub Award), \$600,000, Oct'17 - Jun'20.
7. Senior Personnel, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$216,070, Apr'17 - Mar'18.
8. Senior Personnel, *MVAPICH on CNTK w/ Evaluations of MPI-GDR-level Designs and Optimizations*, Microsoft, \$120,000, Apr'17 - Mar'18.
9. Senior Personnel, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel \$64,424, Feb'16 - Jan'17
10. Senior Personnel, *MVAPICH2 Enhancements for GPUDirect RDMA*, NVIDIA, \$177,500, Jan'17 - Dec'17
11. Senior Personnel, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$213,768, April'16 - Mar'17
12. Senior Personnel, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel \$64,171, Feb'16 - Jan'17.
13. Senior Personnel, *SI2-SSI: Collaborative Research: A Software Infrastructure for MPI Performance Engineering: Integrating MVAPICH and TAU via the MPI Tools Interface*, National Science Foundation, Jointly with Univ. of Oregon, \$2.4M (OSU Share: \$1.2M), Sept'15 - Aug'19.
14. Senior Personnel, *Coupling InfiniBand Hardware Features and Network-to-Accelerator Remote Data Memory Access (RDMA) in the Message Passing Interface (MPI)*, Engility Corporation (US DOD), \$400,000, Sept'15 - Aug'17.
15. Senior Personnel, *Failure Recovery Models and Interfaces in MVAPICH*, Lawrence Livermore National Laboratory, \$71,466, Aug'15 - Aug'16.
16. Senior Personnel, *II-New: Research Infrastructure for Energy-aware High Performance Computing (HPC) and Data Analytics on Heterogeneous Systems*, National Science Foundation, \$898,685, July'15 - June'18.
17. Senior Personnel, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies Inc., \$212,030, April'15 - Mar'16.
18. Senior Personnel, *Chameleon: A Large-Scale, Reconfigurable Experimental Environment for Cloud Research*, in collaboration with Kate Keahy (PI, ANL), TACC, Northwestern and UT-San Antonio, National Science Foundation, \$10M, (OSU share: \$600,000), Oct'14 - Sept'17.
19. Senior Personnel, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies Inc., \$205,709, April'14 - Mar'15.
20. Senior Personnel, *PMI2 Support in MVAPICH2 and Tight Integration with SLURM*, Cray, \$277,856, Jan'14 - April'15.

Ph.D. Thesis Co-Supervised

1. K. S. Khorassani, *High-Performance, Adaptive, and Scalable GPU-aware MPI Libraries for Next-Generation Heterogeneous Systems*, Aug 2023.
2. A. Jain, *Novel Parallelization Strategies for High-Performance DNN Training on HPC Systems*, Dec 2022.

3. M. Bayatpour, *Designing High Performance Hardware-assisted Communication Middlewares for Next-Generation HPC Systems*, May 2021.
4. Ching-Hsiang Chu, *Accelerator-enabled Communication Middleware for Large-scale Heterogeneous HPC Systems with Modern Interconnects*, July 2020.
5. Jahanzeb Maqbool Hashmi, *Designing High Performance Shared-Address-Space and Adaptive Communication Middlewares for Next-Generation HPC Systems*, April 2020.
6. Ammar Ahmad Awan, *Co-designing Communication Middleware and Deep Learning Frameworks for High-Performance DNN Training on HPC Systems*, April 2020.
7. Sourav Chakraborty, *High Performance and Scalable Cooperative Communication Middleware for Next Generation Architectures*, July 2019.

M. S. Thesis Co-Supervised

1. A. P. Gupta, *Enhancing OSU Micro-Benchmarks to be an All-In-One Solution for MPI Benchmarking*, Jul 2023.
2. K. Al-Attar, *Optimizing Apache Spark using the MVAPICH2 MPI Library for High Performance Computing*, Jul 2023.
3. S. Srivastava, *MVAPICH2-AutoTune: An Automatic Collective Tuning Framework for the MVAPICH2 MPI Library*, May 2021.
4. N. Senthil Kumar, *Designing Optimized MPI+NCCL Hybrid Collective Communication Routines for Dense Many-GPU Clusters*, May 2021.
5. Mansa Kedia, *Profile, Monitor, and Introspect Spark Jobs Using OSU INAM*, December 2020.
6. Kamal Raj, *Profiling MPI Primitives in Real-time Using OSU INAM*, April 2020.
7. Albert Augustine, *Designing a Scalable Network Analysis and Monitoring Tool with MPI Support*, July 2016.

M. S. Thesis Co-Supervised

1. N. Iten, *The Ohio State University Micro-Benchmark Suite Orchestrator*, Apr 2023.
2. Z. Salem, *The Ohio State University Micro-Benchmark Suite Orchestrator*, Apr 2023.
3. M. Talyor, *The Ohio State University Micro-Benchmark Suite Orchestrator*, Apr 2023.

Undergraduate Honors Research Co-Supervised

1. Nick Sarkauskas, *Framework for End-to-End Tuning and Regression for a High Performance MPI Library on Modern Supercomputers*, April 2020.

Mentoring Experience

Have significant experience in mentoring under-graduate and graduate (master and doctorate level) students. Worked closely with two masters and seven doctorate level students in the past. Currently working closely with two masters level, 15 doctorate level students, and three staff members. Have participated in student mentoring programs that take place with major conferences like SuperComputing and PEARC.

Curriculum Development and Teaching

As the state-of-the-art in knowledge is advanced through cutting edge research, it is equally important to ensure that this knowledge gets translated and transferred to students through appropriate changes in pedagogy. With this in mind, we have been continually developing new course materials based on the cutting edge research and development being done in our group.

I teach CSE 5449: Introduction Distributed Deep Learning. The objectives of this course are to understand the principles and the practice of this emerging trend, the open set of challenges, how modern HPC technologies can be

used to accelerate DL/ML training and inferencing and apply these benefits to the real-world applications. Recent advancements in Artificial Intelligence (AI) have been fueled by the resurgence of Deep Neural Networks (DNNs) and various Deep Learning (DL) frameworks. DNNs have found widespread applications in several classical AI areas like Image Recognition, Speech Processing, Textual Analysis, as well as areas emerging areas like Cancer Detection, Medical Imaging, Physics, Materials Science, and even Autonomous Vehicle systems. However, scaling distributed training with scale-up and scale-out approaches are still challenging. This is leading to the emergence of a new field called “High-Performance Deep Learning”. The course will involve laboratory exercises for students to experiment with Deep/Machine Learning Frameworks. These exercises will be carried out on OSC (Ohio Supercomputing Center) clusters using CPUs and GPUs. This will provide hands-on knowledge to the students in the area of high-performance deep/machine learning.

I teach CSE 5441: Introduction to Parallel Computing. The course aims to enable students to be competent with certain fundamental factors affecting the performance of sequential programs, understand architectural characteristics of common high-performance computers which directly affect both serial and parallel performance, as well as the implications of computing architecture design decisions on high-performance algorithmic implementations and be familiar with a) prevalent parallel programming models understand the alternative computing architectures for implementing multi-processing systems; b) the design, implementation, and performance analysis of parallel programs, and c) current parallel programming APIs and tools for instrumentation. I have completely revamped the course with a greater focus on hands-on experience for several of the modern parallel programming models including OpenMP, CUDA, and MPI.

I teach CSE 2431: Introduction to Operating Systems. The course aims to introduce students to operating system concepts such as process, CPU scheduling, memory management, file system and storage, and multi-threaded programming. The course goals/objectives are as follows: a) be competent with process concepts and CPU scheduling; b) be competent with memory hierarchy and memory management; c) be familiar with process control blocks, system calls, context switching, interrupts, and exception control flows; d) be familiar with process synchronization, inter-process communication, and threads; e) be familiar with multi-threaded programming; f) be familiar with file systems and disk scheduling algorithms. I have completely revamped the course with a greater focus on hands-on experience for several the basic OS concepts such as threads, locks, synchronization etc.

I have also been teaching various sections of CSE 6422: Advanced Computer Architecture (<https://web.cse.ohio-state.edu/panda.2/6422/>) and CSE 6429: Network-Based Computing for HPC, Cloud, Big Data and Deep Learning (<https://web.cse.ohio-state.edu/panda.2/6429/>) for the past several years.

Software Design, Development, and Distribution

1. Leads the design, development, testing and distribution (involved in various major/minor releases - 1.2p1 through 3.0a) of the MVAPICH2 software stacks - an open-source implementation of the MPI-3.1 specification over modern high-speed networks such as InfiniBand, Omni-Path, 10/40/100/200 GigE, iWARP and RDMA over Converged Ethernet (RoCE). This software is being used by more than 3,325 organizations (National Labs, Universities, and Industry) in 90 countries have downloaded this software from OSUs web site directly together with voluntary registration. As of September’23, more than 1,720,000 downloads have taken place from the OSU Web site. In addition, many server vendors, networking vendors (InfiniBand, Omni-Path, iWARP, and RoCE) and system integrators are incorporating MVAPICH2 into their software stacks and distributing it. MVAPICH2 is also available in many Linux distributions (such as RedHat and SuSE). Several InfiniBand systems using MVAPICH2 have obtained positions in the TOP500 ranking of supercomputers in the world (www.top500.org) during the last decade. The Jun’23 list includes the following systems: 7th, 10,649,600-core (Sunway TaihuLight) at National Supercomputing Center in Wuxi, China 21st, 448, 448 cores (Frontera) at TACC 36th, 288,288 cores (Lassen) at LLNL 49th, 570,020 cores (Nurion) in South Korea. More details on the MVAPICH2 project, publications, users, and its impact in the community can be obtained by visiting the following URL: <http://mvapich.cse.ohiostate.edu/>.

The MVAPICH2-X software package provides support for hybrid MPI+PGAS (UPC, UPC++, OpenSHMEM, and CAF) programming models with unified communication runtime for emerging exascale systems.

2. Leads the design, development, testing and distribution of the OSU InfiniBand Network Analysis and Monitor-

ing (OSU INAM). OSU INAM monitors InfiniBand clusters in real time by querying various subnet management entities in the network. It is also capable of interacting with the MVAPICH2-X software stack to gain insights into the communication pattern of the application and classify the data transferred into Point-to-Point, Collective and Remote Memory Access (RMA). OSU INAM can also remotely monitor several parameters of MPI processes in conjunction with MVAPICH2-X. As of March'22, the OSU INAM package, which is derived from this work, has been downloaded more than 4,000 times from the OSU site directly.

3. Leads the design, development and release of the MVAPICH2-GDR package based on the standard MVAPICH2 software stack, incorporates designs that take advantage of the new GPUDirect RDMA technology for inter-node data movement on NVIDIA GPU clusters with Mellanox InfiniBand interconnect. It provides significantly improved performance for small and medium messages while achieving close to peak network bandwidth for large messages.
4. Co-Leads the HiDL project focusing on co-designing Deep Learning Frameworks and MPI Libraries. A high-performance and scalable version of the Caffe framework is available from <http://hidl.cse.ohio-state.edu>. As of Jan 19, more than 1,400 downloads of this software have taken place and it is being used by more than 78 organizations.
5. Leads the design, development and release of the MVAPICH2-EA package based on the standard MVAPICH2 software stack. MVAPICH2-EA incorporates designs that has energy-efficient support for point-to-point and collective communication protocols for IB, RoCE and iWARP networks. It provides significantly improved savings in energy while delivering performance within bounds specified by the user.
6. Leads the integration of MVAPICH and TAU via the MPI Tools Interface. This is a research being jointly done with the University of Oregon. The proposed research, undertaken by a team of computer scientists from OSU and UO representing the MVAPICH and TAU projects, aims to create an integrated software infrastructure built on the MPLT interface which defines the API for interaction and information interchange to enable fine grained performance optimizations for HPC applications.
7. Leads the design, development, testing and distribution of the project to have support for PMI2 in MVAPICH2 and to have a tight integration with the SLURM job launcher. The work done as part of this project has lead to several-fold increase in the performance of job-startup in MVAPICH2.
8. Leads the design, development and release of MVAPICH2-Virt (MVAPICH2 with Virtualization Support). MVAPICH2-Virt incorporates designs that take advantage of the new features and mechanisms of high-performance networking technologies with SR-IOV as well as other virtualization technologies such as Inter-VM Shared Memory (IVSHMEM). For an InfiniBand SR-IOV-based virtualized environment, MVAPICH2-Virt has very little overhead compared with MVAPICH2 running over InfiniBand in native mode.

Professional Activities

Chair/Co-Chair/Vice-Chair:

1. CCGrid-2022: 22nd Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing 2022 (Distributed middleware and network architectures track)
2. Seventh International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2022
3. Seventh International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm), 2022
4. Sixth International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2021
5. Sixth International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm), 2021
6. CCGrid-2021: 21st Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing 2021 (Posters Track)
7. CCGrid-2020: 20th Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing 2020 (Posters Track)
8. Fifth International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2020

9. Fifth International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm), 2019
10. Fourth International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2018
11. Fourth International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm), 2018
12. Third International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2017
13. Third International Workshop on Communication Architectures for HPC, Big Data, Deep Learning and Clouds at Extreme Scale (ExaComm), 2017
14. Second International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2016
15. Second International Workshop on Communication Architectures at Extreme Scale (ExaComm), 2016
16. First International Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2015
17. First International Workshop on Communication Architectures at Extreme Scale (ExaComm), 2015

Program Committee Member:

1. SC-2022: The International Conference for High Performance Computing, Networking, Storage, and Analysis 2022
2. OFAWS-2022: Open Fabrics Alliance Workshop, 2022
3. IPDPS-2022, International Parallel & Distributed Processing Symposium, 2022
4. ExaMPI-2021: Workshop on Exascale MPI (Held in conjunction with SC21), 2021
5. CCGrid-2021: 21st Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing 2021
6. EuroMPI-2021: European MPI Users's Group Meeting, 2021
7. OFAWS-2021: Open Fabrics Alliance Workshop, 2021
8. IPDPS-2020, International Parallel & Distributed Processing Symposium, 2020
9. HotI-2020: IEEE Hot Interconnects Symposium, 2020
10. ExaMPI-2020: Workshop on Exascale MPI (Held in conjunction with SC20), 2020
11. HPCS-2020: International Conference on High Performance Computing and Simulation, 2020
12. EuroMPI-2020: European MPI Users's Group Meeting, 2020
13. OFAWS-2020: Open Fabrics Alliance Workshop, 2020
14. SC-2019: The International Conference for High Performance Computing, Networking, Storage, and Analysis 2019
15. CCGrid-2019: 19th Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing 2019
16. HPCS-2019: International Conference on High Performance Computing and Simulation, 2019
17. SNACS-2019: Workshop on Scalable Networks for Advanced Computing Systems, in conjunction with IPDPS'19, 2019
18. EuroMPI-2019: European MPI Users's Group Meeting, 2019
19. ICA3PP-2019: 19th International Conference on Algorithms and Architectures for Parallel Processing, 2019
20. ExaMPI-2019: Workshop on Exascale MPI (Held in conjunction with SC19), 2019
21. HiPC-2018: IEEE International Conference on High Performance Computing, Data, and Analytics 2018
22. HPCS-2018: International Conference on High Performance Computing and Simulation, 2018
23. PFFT18: Parallel Fast Fourier Transforms 2018
24. ExaMPI-2018: Workshop on Exascale MPI (Held in conjunction with SC18), 2018

25. SCEC-2018: Second Workshop on Software Challenges to Exascale Computing, 2018
26. ICA3PP-2018: 18th International Conference on Algorithms and Architectures for Parallel Processing, 2018
27. CCGrid-2018: 18th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing 2018
28. HiPC-2017: IEEE International Conference on High Performance Computing, Data, and Analytics 2017
29. ICPP' 17: International Conference on Parallel Processing 2017
30. IEEE ANTS: IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS) 2017
31. ExaMPI-2016: Workshop on Exascale MPI (Held in conjunction with SC16), 2016
32. ICPP' 15: International Conference on Parallel Processing, 2015

Panel Member:

1. *Quieting Noisy Neighbors: Addressing Network Congestion in Multi-Workload Environments, HPCwire, November 2021.*
2. *When is enough, enough? With so many Parallel Programming Technologies, is it Time to Focus on Consolidating them?, Birds of Feather, International Supercomputing Conference (ISC), June 2019.*
3. *FFT in the Exascale: Opportunities and Challenges, Birds of Feather, International Supercomputing Conference (ISC), June 2018.*

Organizational:

1. Organizing Committee Member, Eleventh MVAPICH User Group Meeting (MUG'23)
2. Organizing Committee Member, 2023 HDF5 User Group Meeting (HUG'23)
3. Organizing Committee Member, Tenth MVAPICH User Group Meeting (MUG'23)
4. Organizing Committee Member, Ninth MVAPICH User Group Meeting (MUG'21)
5. Organizing Committee Member, Eighth MVAPICH User Group Meeting (MUG'20)
6. Organizing Committee Member, Seventh MVAPICH User Group Meeting (MUG'19)
7. Organizing Committee Member, Sixth MVAPICH User Group Meeting (MUG'18)
8. Organizing Committee Member, Fifth MVAPICH User Group Meeting (MUG'17)
9. Organizing Committee Member, Fourth MVAPICH User Group Meeting (MUG'16)
10. Organizing Committee Member, Third MVAPICH User Group Meeting (MUG'15)
11. Organizing Committee Member, Second MVAPICH User Group Meeting (MUG'14)
12. Session Chair - Int'l Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2014
13. Committee Member - Outstanding Graduate Research Award, Department of Computer Science and Engineering, 2014
14. Organizing Committee Member, First MVAPICH User Group Meeting (MUG'13)

Referee:

1. Sub-Reviewer, EuroPar'22
2. Reviewer, SuperComputing (SC), 2013, 2019, 2022, 2023
3. Reviewer, ACM Transactions on Parallel Computing (TOPC), 2014
4. Reviewer, Transactions on Parallel and Distributed Systems (TPDS), 2013, 2015, 2020, 2021, 2023
5. Reviewer, Journal of Parallel and Distributed Computing (JPDC), 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2023
6. Reviewer, Parallel Computing - Systems & Applications (ParCo), 2014, 2015, 2016, 2017, 2018

7. Reviewer, International Conference on Parallel Processing (ICPP), 2009, 2010, 2011, 2015
8. Reviewer, International Parallel & Distributed Processing Symposium (IPDPS), 2010, 2011, 2013, 2019, 2020, 2022, 2023
9. Reviewer, IEEE Cluster, 2010, 2011
10. Reviewer, Open Fabrics Workshop, 2020, 2021

Judge:

1. APAC HPC AI competition, The HPC-AI Advisory Council, Oct 2020
2. Student Cluster Competition, held in conjunction with International Supercomputing Conference (ISC), 2019
3. The Seventh Student RDMA Programming Competition, The HPC-AI Advisory Council, May 2019
4. APAC HPC AI competition, The HPC-AI Advisory Council, Oct 2018

Service:

- Member, The Process Management Interface - Exascale (PMIx) Forum, 2020 - Present
- Member, The Message Passing Interface (MPI) Forum, 2008 - Present
- Member, OpenFabrics Alliance Technical Advisory Committee, 2021 - Present
- Incoming Students Panel, Department of Computer Science and Engineering, 2017.
- Graduate Student Awards Committee, Department of Computer Science and Engineering, 2014.

Memberships:

- IEEE Member
- ACM Member

Invited Keynote Talk

1. *Designing HPC, Deep Learning and Cloud Middleware for Exascale Systems, HPC Advisory Council Australian Conference, Aug 2018*

Invited Talks

1. *Offload Framework for SMART NICs, SMART NIC Birds of Feather, in conjunction with Supercomputing, Nov 2021.*
2. *The MVAPICH2 Project: Latest Status and Future Plans, MPICH Birds of Feather, in conjunction with Supercomputing, Nov 2021.*
3. *Designing High-Performance Scalable Middleware for HPC, AI and Data Science in Exascale Systems and Clouds, Open Source Summit, Sept 2021.*
4. *Designing a High-performance MPI library using In-Network Computing, Workshop on Scalable Networks for Advanced Computing Systems (SNACS), May 2021.*
5. *Designing a High-performance MPI library using In-Network Computing, Open Fabrics Alliance Workshop, March, 2021.*
6. *Designing a ROCm-aware MPI Library for AMD GPUs over High-speed Networks, Open Fabrics Alliance Workshop, March, 2021.*
7. *Extracting Peak Performance for your Applications on Frontera with MVAPICH2 Libraries, Frontera User Meeting, Jan 2021.*
8. *Visualize and Analyze Your Network Activities Using OSU INAM, Open Fabrics Alliance Workshop, March, 2020.*
9. *The MVAPICH2 Project: Latest Status and Future Plans, MPICH Birds of Feather, in conjunction with Supercomputing, Nov 2019.*

10. *Designing High Performance MPI Collectives in MVAPICH2 for HPC and Deep Learning*, Open Fabrics Alliance Workshop, March, 2019.
11. *Visualize and Analyze Your Network Monitoring Enhancements*, Open Fabrics Alliance Workshop, March, 2019.
12. *Experiences in Designing, Developing, Packaging, and Deploying the MVAPICH2 Libraries*, The First Extreme-scale Scientific Software Stack Forum, Sep 2019.
13. *The MVAPICH2 Project: Latest Status and Future Plans*, MPICH Birds of Feather, in conjunction with Supercomputing, Nov 2018.
14. *Network Computing in HPC and AI*, APAC HPC-AI Competition Award Ceremony, August 2018.
15. *FFT in the Exascale: Performance Bottlenecks and Solutions*, FFT in the Exascale: Opportunities and Challenges Birds of Feather, International Supercomputing Conference, June 2018.
16. *Designing Scalable, High-Performance Communication Runtimes for HPC and Deep Learning: The MVAPICH2 Approach*, Open Fabrics Alliance Workshop, April, 2018.
17. *MVAPICH2-GDR: High-Performance and Scalable CUDA-Aware MPI Library for HPC and AI*, NVIDIA GTC Conference, March 19, 2019.
18. *MVAPICH2-GDR: Pushing the Frontier of MPI Libraries Enabling GPUDirect Technologies*, NVIDIA GTC Conference, March 28, 2018.
19. *Designing HPC, Big Data, and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Indian Institute of Technology, Mumbai, 2017.
20. *Designing HPC, Big Data, and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Indian Institute of Space Technology, Trivandrum, 2017.
21. *Designing HPC, Big Data, and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Indian Institute of Technology, Chennai, 2017.
22. *Optimizing and Tuning Techniques for Running MVAPICH2 over IB*, InfiniBand User Group Meeting (IBUG), April 2014.
23. *MVAPICH2 Project Update and GPUDirect RDMA*, HPC Advisory Council European Conference, June 2013.
24. *Experiencing HPC for Undergraduates - Graduate Student Perspective*, SuperComputing (SC '12), Nov 2012.
25. *MVAPICH2 Project Update and Big Data Acceleration*, HPC Advisory Council European Conference, June 2012.

Invited Tutorials

1. *Introduction to Networking Technologies for High-Performance Computing*, IEEE HPSR 23, June 5, 2023.
2. *Introduction to Networking Technologies for High-Performance Computing*, ISC 23, May 21, 2022.
3. *Principles and Practice of High Performance Deep/Machine Learning*, PPOPP 23, February 25, 2023.
4. *Exploiting High-performance Computing for Deep Learning: Why and How? High Performance Machine Learning, Deep Learning, and Data Science*, DeepLearn 23, January 19-20, 2023.
5. *Introduction to Networking Technologies for High-Performance Computing*, SC 22, November 13, 2022.
6. *High Performance Machine Learning, Deep Learning, and Data Science*, Hot Interconnect 22, August 19, 2022.
7. *Introduction to Networking Technologies for High-Performance Computing*, Hot Interconnect 22, August 19, 2022.
8. *High Performance Machine Learning, Deep Learning, and Data Science*, ISCA 22, June 18, 2022.
9. *High Performance Machine Learning, Deep Learning, and Data Science*, ISC 22, April 29, 2022.
10. *High Performance Machine Learning, Deep Learning, and Data Science, Tutorial and Training, As a part of the NSF Cybertraining Grant*, April 19-27, 2022.

11. *High Performance Machine Learning, Deep Learning, and Data Science, Open Fabrics Workshop, April 28, 2022.*
12. *Introduction to Networking Technologies for High-Performance Computing, Open Fabrics Workshop, April 27, 2022.*
13. *High Performance Machine Learning, Deep Learning, and Data Science, PPOPP 22, April 3, 2022.*
14. *Designing HPC Systems with High-Performance Networks: Advanced Features, Challenges, and Usage, SC'21, November 2021.*
15. *Introduction to Networking Technologies for High-Performance Computing, SC'21, November 2021.*
16. *Introduction to Networking Technologies for High-Performance Computing, HotI 21, Aug' 20, 2021.*
17. *Visualize, Analyze, and Correlate Networking Activities for Your Parallel Programs on InfiniBand HPC Clusters using the OSU INAM Tool, PEARC 21, July 18, 2021.*
18. *InfiniBand, High-speed Ethernet, RoCE, Omni-Path, EFA and Slingshot for Beginners, ISC 21, June 2021.*
19. *High Performance Distributed Deep Learning: A Beginners Guide, ISC '21, June 2021.*
20. *High Performance Distributed Deep Learning: A Beginners Guide, ISCA21, June 2021.*
21. *High Performance Machine Learning, Deep Learning, and Data Science, ICS '21, June 2021.*
22. *High Performance Distributed Deep Learning: A Beginners Guide, ASPLOS 21.*
23. *High Performance Distributed Deep Learning: A Beginners Guide, PPOPP 21. Feb. 28, 2021.*
24. *InfiniBand, High-speed Ethernet, RoCE, Omni-Path, EFA and Slingshot for Beginners, Supercomputing 20. November 10, 2020.*
25. *InfiniBand, High-speed Ethernet, RoCE, Omni-Path, EFA and Slingshot for Beginners, HotI 20. Aug. 21, 2020.*
26. *Exploiting HPC for distributed Deep Learning, HotI 20. Aug. 21, 2020.*
27. *High Performance Distributed Deep Learning: A Beginners Guide, ISCA 20. May 31, 2020.*
28. *High Performance Distributed Deep Learning, NVIDIA GTC 20. March 23, 2020.*
29. *High Performance Distributed Deep Learning, PPOPP 20. February 22, 2020.*
30. *How to Boost the Performance of HPC/AI Applications using MVAPICH2 Libraries?, HiPEAC 20. January 22, 2020.*
31. *High Performance Distributed Deep Learning: A Beginners Guide, HiPEAC 20. January 22, 2020.*
32. *InfiniBand, Omni-Path, and High-Speed Ethernet for Beginners, SuperComputing (SC '19), Nov 2019.*
33. *InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage, SuperComputing (SC '19), Nov 2019.*
34. *High-Performance Distributed Deep Learning: A Beginner's Guide, MICRO '19, Oct, 2019.*
35. *Extracting Peak Performance for your Applications on Frontera with MVAPICH2 Libraries?, Texas Advanced Computing Center, Aug 2019.*
36. *High-Performance Distributed Deep Learning: A Beginner's Guide, PEARC '19, July, 2019.*
37. *High-Performance Distributed Deep Learning: A Beginner's Guide, Int'l Supercomputing Conference (ISC '19), June, 2019.*
38. *InfiniBand and High-Speed Ethernet for Dummies, Int'l Supercomputing Conference (ISC '19), June, 2019.*
39. *High-Performance Distributed Deep Learning: A Beginner's Guide, ISCA '19, June 23, 2019.*
40. *High-Performance Distributed Deep Learning: A Beginner's Guide, CCGrid '19, May 14, 2019.*
41. *How to Boost the Performance of Your HPC/AI Applications with MVAPICH2 Libraries?, NCAR SEA Symposium, April 12, 2019.*
42. *High Performance Distributed Deep Learning: A Beginner's Guide, NCAR SEA Symposium, April 12, 2019.*

43. *How to Boost the Performance of HPC/AI Applications Using MVAPICH2 Library?*, NVIDIA GTC'19. March 20, 2019.
44. *High Performance Distributed Deep Learning: A Beginner's Guide*, NVIDIA GTC'19. March 18, 2019
45. *High-Performance Distributed Deep Learning: A Beginner's Guide*, PPOPP '19, February 17, 2019.
46. *InfiniBand, Omni-Path, and High-Speed Ethernet for Dummies*, SuperComputing (SC '18), Nov 2018.
47. *InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SuperComputing (SC '18), Nov 2018.
48. *Engineering your Application for Peak Performance with TAU and MVAPICH2* PEARC'18, July, 2018.
49. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '18), June, 2018.
50. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, Int'l Supercomputing Conference (ISC '18), June, 2018.
51. *Performance Tuning of MVAPICH2 and MVAPICH2-GDR MPI Libraries*, HPCMP PETTT Technical Interchange Meeting, May 15, 2018.
52. *High-Performance Distributed Deep Learning: A Beginners Guide*, HPCMP PETTT Technical Interchange Meeting, May 15, 2018.
53. *High-Performance Distributed Deep Learning: A Beginners Guide*, NCAR SEA Symposium, April 5, 2018.
54. *How to Boost the Performance of Your MPI and PGAS Applications with MVAPICH2 Libraries*, NCAR SEA Symposium, April 5, 2018.
55. *High-Performance Distributed Deep Learning: A Beginners Guide*, PPOPP 18, February 25, 2018.
56. *High-Performance Distributed Deep Learning for Dummies*, IT4 Innovations, January 24, 2018.
57. *InfiniBand and High-Speed Ethernet for Dummies*, IT4 Innovations, January 23, 2018.
58. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, IT4 Innovations, January 23, 2018.
59. *InfiniBand, Omni-Path, and High-Speed Ethernet for Dummies*, SuperComputing (SC '17), Nov 2017.
60. *InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SuperComputing (SC '17), Nov 2017.
61. *High Performance Distributed Deep Learning for Dummies*, Hot Interconnects (HotI'17), Aug, 2017.
62. *Engineering your Application for Peak Performance with TAU and MVAPICH2* PEARC'17, July, 2017.
63. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '17), June, 2017.
64. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, Int'l Supercomputing Conference (ISC '17), June, 2017.
65. *Co-designing MPI Runtimes and Deep Learning Frameworks or Modern Multi-GPU Clusters*, HPCMP PETTT Technical Interchange Meeting on High Performance Signal Processing, Dayton, OH, May, 2017.
66. *Programming, Optimizing and Tuning your Message Passing Interface (MPI) application with MVAPICH2 and MVAPICH2-GDR Libraries*, HPCMP PETTT Technical Interchange Meeting on High Performance Signal Processing, Dayton, OH, May, 2017.
67. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '16), Nov 2016.
68. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SuperComputing (SC '16), Nov 2016.
69. *How to Boost the Performance of your MPI and PGAS Applications with MVAPICH2 Libraries?* XSEDE'16, July, 2016.
70. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '16), June, 2016.
71. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, Int'l Supercomputing Conference (ISC '16), July, 2016.

72. *Programming, Optimizing and Tuning your Message Passing Interface (MPI) application with MVAPICH2 and MVAPICH2-GDR Libraries*, HPCMP PETTT Technical Interchange Meeting on High Performance Signal Processing, Dayton, OH, May, 2016.
73. *InfiniBand and High-Speed Ethernet for Dummies*, CHPC National Meeting (CHPC '15), Dec 2015.
74. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '15), Nov 2015.
75. *Optimization and Tuning of MPI and PGAS Applications using MVAPICH2 Software Stacks* XSEDE'15, July, 2015.
76. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '15), Jul 2015.
77. *Designing and Using High-End Computing Systems with InfiniBand and High-Speed Ethernet* Int'l Supercomputing Conference (ISC '15), Jul 2015.
78. *InfiniBand and High-Speed Ethernet for Dummies*, CHPC National Meeting (CHPC '14), Dec 2014.
79. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '14), Nov 2014.
80. *Designing and Using High-End Computing Systems with InfiniBand and High-Speed Ethernet* SuperComputing (SC '14), Nov 2014.
81. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends* Cluster'14, Sep, 2014.
82. *Optimization and Tuning of MPI and PGAS Applications using MVAPICH2 and MVAPICH2-X* XSEDE'14, July, 2014.
83. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends* Int'l Conference on High Performance Switching and Routing (HPSR'14), July, 2014.
84. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends* Int'l Supercomputing Conference (ISC '14), June, 2014.
85. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '13), Nov 2013.
86. *Advanced Topics in InfiniBand and High-speed Ethernet for Designing HEC Systems* SuperComputing (SC '13), Nov 2013.
87. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '13), June, 2013.
88. *Advanced Topics in InfiniBand and High-speed Ethernet for Designing HEC Systems* Int'l Supercomputing Conference (ISC '13), June, 2013.
89. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '12), Nov 2012.
90. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, SuperComputing (SC '12), Nov 2012.
91. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '12), June, 2012.
92. *Designing High-End Computing Systems and Programming Models with InfiniBand and High-Speed Ethernet*, Int'l Supercomputing Conference (ISC '12), June, 2012.
93. *InfiniBand and High-Speed Ethernet for Dummies*, SuperComputing (SC '11), Nov 2011.
94. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, SuperComputing (SC '11), Nov 2011.

Publications

The following publications have cumulatively been cited more than 2,783 times in various top international conferences and journals with an “h-index” of 27 and an “i10-index” of 70 as listed by Google Scholar ¹. A subset of these publications that were published with ACM has been downloaded over 7,300 times and cited over 460 times as indicated by ACM ². Another subset of these publications with IEEE have been cited over 993 times ³.

¹<https://scholar.google.com/citations?user=ykYaLmAAAAAJ&hl=en>

²<https://dl.acm.org/profile/81440602821>

³<https://ieeexplore.ieee.org/author/37393051700>

Journal Articles

1. T. Tran, B. Ramesh, B. Michalowicz, M. Abduljabbar, H. Subramoni, A. Shafi, and DK Panda, Accelerating Communication with Multi-HCA Aware Collectives in MPI, *Concurrency and Computation: Practice and Experience (CCPE)*, July 2023.
2. K. Khorassani, C. Chen, B. Ramesh, A. Shafi, H. Subramoni, and DK Panda, High Performance MPI over the Slingshot Interconnect, *Special Issue of Journal of Computer Science and Technology (JCST)*, Feb 2023.
3. K. Suresh, K. Khorassani, C. Chen, B. Ramesh, M. Abduljabbar, A. Shafi, H. Subramoni, and DK Panda, Network Assisted Non-Contiguous Transfers for GPU-Aware MPI Libraries, *IEEE Micro*, Jan 2023.
4. A. Jain, N. Alnaasan, A. Shafi, **H. Subramoni**, and DK Panda, Optimizing Distributed DNN Training using CPUs and BlueField-2 DPUs , *IEEE Micro*, doi: 10.1109/MM.2021.3139027
5. D. K. Panda, **H. Subramoni**, C. Chu, and M. Bayatpour, The MVAPICH project: Transforming Research into High-Performance MPI library for HPC Community, *Journal of Computational Science (JOCS)*, Special Issue on Translational Computer Science, Oct 2020.
6. J. Hashmi, C. Chu, S. Chakraborty, M. Bayatpour, **H. Subramoni**, and D. K. Panda, FALCONX: Zero-copy MPI Derived Datatype Processing on Modern CPU and GPU Architectures, *Journal of Parallel and Distributed Computing (JPDC)*, May 2020.
7. A. Awan, A. Jain, C. Chu, **H. Subramoni**, and D. K. Panda, Communication Profiling and Characterization of Deep Learning Workloads on Clusters with High-Performance Interconnects , *IEEE Micro*, vol. 40, no. 1, pp. 35-43, 1 Jan.-Feb. 2020
8. A. Awan, K. Vadambacheri Manian, C. Chu, **H. Subramoni**, and D. K. Panda, Optimized Large-Message Broadcast for Deep Learning Workloads: MPI, MPI+NCCL, or NCCL2?, *Journal of Parallel Computing (PARCO)*, Volume 85, July 2019, Pages 141-152.
9. A. Ruhela, **H. Subramoni**, S. Chakraborty, M. Bayatpour, P. Kousha, and D. K. Panda, Efficient Design for MPI Asynchronous Progress without Dedicated Resources, *Parallel Computing - Systems and Applications*, Volume 85, July 2019, Pages 13-26.
10. C. Chu, X. Lu, A. Awan, **H. Subramoni**, B. Elton, and underline D. K. Panda, Exploiting Hardware Multicast and GPUDirect RDMA for Efficient Broadcast, *IEEE Transactions on Parallel and Distributed Systems*, March 2019.
11. D. K. Panda, X. Lu, and **H. Subramoni**, Networking and Communication Challenges or Post-Exascale Systems, *Frontiers of Information technology and Electronic Engineering (FITEE)*, Volume 19, October 2018, pp. 1-6.
12. S. Ramesh, A. Maho, S. Shende, A. Malony, **H. Subramoni**, A. Ruhela, and D. K. Panda, MPI performance engineering with the MPI tool interface: The integration of MVAPICH and TAU, September 2018.
13. S. Chakraborty, I. Laguna, M. Emani, K. Mohror, D. K. Panda, M. Schulz, and **H. Subramoni**, EReinit: Scalable and Efficient Fault Tolerance for Bulk-Synchronous MPI Applications, *Concurrency and Computation: Practice and Experience*, August 2018.
14. K. Hamidouche, A. Venkatesh, **H. Subramoni**, and D. K. Panda, Designing High Performance MPI and PGAS Communication Library for GPU Clusters, Book chapter in *Advances in GPU Research and Practice*.
15. K. Hamidouche, A. Venkatesh, A. A. Awan, **H. Subramoni**, C-H. Chu and D. K. Panda, CUDA-Aware OpenSHMEM: Extensions and Designs for High Performance OpenSHMEM on GPU Clusters, Special issue on Parallel Programming Models and Systems Software, associated with the Elsevier *Parallel Computing (PARCO)* journal.
16. S. Sur, S. Potluri, K. Kandalla, **H. Subramoni**, K. Tomko, D. K. Panda, Co-Designing MPI Library and Applications for InfiniBand Clusters, *Computer*, 06 Sep. 2011. IEEE Computer Society Digital Library.
17. D. K. Panda, S. Sur, **H. Subramoni** and K. Kandalla, Network Support for Collective Communication, *Encyclopedia of Parallel Computing*, Sep. 2011.
18. **H. Subramoni**, F. Petrini, V. Agarwal and D. Pasetto, Intra-Socket and Inter-Socket Communication in Multi-core Systems, *Computer Architecture Letters* 9(1): 13-16, 2010.

Invited Papers (Book Chapters/Conference/Workshop)

1. D. K. Panda, S. Sur, **H. Subramoni** and K. Kandalla, Network Support for Collective Communication, Encyclopedia of Parallel Computing, 2011.

Refereed Conference/Workshop Publications:

1. C. Chen, K. Khorassani, P. Kousha, Q. Zhou, J. Yao, **H. Subramoni**, and DK Panda, MPI-xCCL: A Portable MPI Library over Collective Communication Libraries for Various Accelerators Sixth Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, Nov 2023.
2. P. Kousha, V. Sathu, M. Lieber, **H. Subramoni**, and DK Panda, Democratizing HPC Access and Use with Knowledge Graphs, D-HPC 2023: The First International Workshop on Democratizing High-Performance Computing, Nov 2023.
3. B. Michalowicz, K. Suresh, **H. Subramoni**, DK Panda, and S. Poole, Battle of the BlueFields: An In-Depth Comparison of the BlueField-2 and BlueField-3 SmartNICs, Hot Interconnects 2023, Aug 2023.
4. B. Michalowicz, K. Suresh, **H. Subramoni**, DK Panda, and S. Poole, DPU-Bench: A Micro-Benchmark Suite to Measure Offload Efficiency Of SmartNICs, Practice and Experience in Advanced Research Computing 23, Jul 2023.
5. S. Khuvis, K. Tomko, S. Brozell, C. Chen, **H. Subramoni**, and DK Panda, Optimizing Amber for Device-to-Device GPU Communication, Practice and Experience in Advanced Research Computing 23, Jul 2023.
6. N. Contini, B. Ramesh, K. Suresh, T. Tran, B. Michalowicz, M. Abduljabbar, **H. Subramoni**, and DK Panda, Enabling Reconfigurable HPC through MPI-based Inter-FPGA Communication, International Conference on Supercomputing 2023, Jun 2023.
7. P. Kousha, A. Jain, A. Kolli, M. Lieber, M. Han, N. Contini, H. Subramoni, and DK Panda, SAI: AI-Enabled Speech Assistant Interface for Science Gateways in HPC, ISC 2023, May 2023.
8. Q. Zhou, Q. Anthony, L. Xu, A. Shafi, M. Abduljabbar, **H. Subramoni**, and DK Panda, Accelerating Distributed Deep Learning Training with Compression Assisted Allgather and Reduce-Scatter Communication, 37th IEEE International Parallel & Distributed Processing Symposium (IPDPS '23), May 2023.
9. Q. Anthony, Ammar Awan, J. Rasley, Y. He, A. Shafi, M. Abduljabbar, **H. Subramoni**, and DK Panda, MCR-DL: Mix-and-Match Communication Runtime for Deep Learning, 37th IEEE International Parallel & Distributed Processing Symposium (IPDPS '23), May 2023.
10. K. Suresh, B. Michalowicz, B. Ramesh, N. Contini, J. Yao, S. Xu, A. Shafi, **H. Subramoni**, and DK Panda, A Novel Framework for Efficient Offloading of Communication Operations to Bluefield SmartNICs, 37th IEEE International Parallel & Distributed Processing Symposium (IPDPS '23), May 2023.
11. K. Khorassani, C. Chen, **H. Subramoni**, and DK Panda, Designing and Optimizing GPU-aware Nonblocking MPI Neighborhood Collective Communication for PETSc, 37th IEEE International Parallel & Distributed Processing Symposium (IPDPS '23), May 2023.
12. B. Michalowicz, K. Suresh, B. Ramesh, A. Shafi, **H. Subramoni**, M. Abduljabbar, and DK Panda, In-Depth Evaluation of a Lower-Level Direct-Verbs API on InfiniBand-based Clusters: Early Experiences, 25th Workshop on Advances in Parallel and Distributed Computational Models, May 2023.
13. C. Chen, K. Khorassani, G. Kuncham, R. Vaidya, M. Abduljabbar, A. Shafi, **H. Subramoni**, and DK Panda, Designing and Optimizing a GPU-aware MPI Library for Intel GPUs: Early Experiences, The 23RD IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, May 2023.
14. Q. Anthony, L. Xu, A. Shafi, **H. Subramoni**, and DK Panda, ScaMP: Scalable Meta-Parallelism for Deep Learning Search, The 23RD IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, May 2023.
15. H. Ahn, T. Chen, N. Alnaasan, A. Shafi, M. Abduljabbar, **H. Subramoni**, and DK Panda, Performance Characterization of using Quantization for DNN Inference on Edge Devices, 7TH IEEE International Conference on Fog and Edge Computing, May 2023.

16. N. Alnaasan, A. Jain, A. Shafi, **H. Subramoni**, and DK Panda, AccDP: Accelerated Data-Parallel Distributed DNN Training for Modern GPU-Based HPC Clusters, 29th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2022.
17. B. Ramesh, Q. Zhou, A. Shafi, M. Abduljabbar, **H. Subramoni**, and DK Panda, Designing Efficient Pipelined Communication Schemes using Compression in MPI Libraries, 29th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2022.
18. Q. Zhou, Q. Anthony, A. Shafi, **H. Subramoni**, and DK Panda, Accelerating Broadcast Communication with GPU Compression for Deep Learning Workloads, 29th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2022.
19. K. Suresh, A. Paniraja Guptha, B. Michalowicz, B. Ramesh, M. Abduljabbar, A. Shafi, **H. Subramoni**, and DK Panda, Efficient Personalized and Non-Personalized Alltoall Communication for Modern Multi-HCA GPU-Based Clusters, 29th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2022.
20. K. Al Attar, A. Shafi, M. Abduljabbar, **H. Subramoni**, and DK Panda, Spark Meets MPI: Towards High-Performance Communication Framework for Spark using MPI, IEEE Cluster '22, Sep 2022.
21. T. Tran, B. Michalowicz, B. Ramesh, **H. Subramoni**, A. Shafi, and DK Panda, Designing Hierarchical Multi-HCA Aware Allgather in MPI, Fifteenth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2022, Aug 2022.
22. K. Khorassani, C. Chen, B. Ramesh, A. Shafi, **H. Subramoni**, and DK Panda, High Performance MPI over the Slingshot Interconnect: Early Experiences, Practice and Experience in Advanced Research Computing, Jul 2022 **Best Student Paper Award**
23. S. Xu, A. Shafi, **H. Subramoni**, and DK Panda, Arm meets Cloud: A Case Study of MPI Library Performance on AWS Arm-based HPC Cloud with Elastic Fabric Adapter 24th Workshop on Advances in Parallel and Distributed Computational Models, May 2022.
24. K. Al Attar, A. Shafi, **H. Subramoni**, and DK Panda, Towards Java-based HPC using the MVAPICH2 Library: Early Experiences, HIPS '22 (IPDPSW), May 2022
25. C. Chen, K. Khorassani, Q. Anthony, A. Shafi, **H. Subramoni**, and DK Panda, Highly Efficient Alltoall and Alltoallv Communication Algorithms for GPU Systems, Heterogeneity in Computing Workshop (HCW 2022), May 2022
26. Q. Zhou, P. Kousha, Q. Anthony, K. Khorassani, A. Shafi, **H. Subramoni**, and DK Panda, Accelerating MPI All-to-All Communication with Online Compression on Modern GPU Clusters, Intl Supercomputing Conference 2022 2022, May 2022
27. A. Jain, A. Shafi, Q. Anthony, P. Kousha, **H. Subramoni**, and DK Panda, Hy-Fi: Hybrid Five-Dimensional Parallel DNN Training on High-Performance GPU Clusters, Intl Supercomputing Conference 2022, May 2022
28. P. Kousha, A. Kolli, S. Miriyala, S. Sainath, A. Jain, **H. Subramoni**, A. Shafi, and DK Panda, "Hey CAI" - Enhancing User Productivity through a Conversational AI Enabled User Interface for HPC Tools, Intl Supercomputing Conference 2022, May 2022
29. N. Sarkauskas, M. Bayatpour, A. Tran, B. Ramesh, **H. Subramoni**, and DK Panda, Large-Message Nonblocking MPI_Allgather and MPI_Ibcast Offload via BlueField-2 DPU, 28th IEEE International Conference on High Performance Computing, Data, Analytics, and Data Science (HiPC'21), Dec 2021.
30. K. Suresh, B. Ramesh, C. Chen, M. Ghazimirsaeed, M. Bayatpour, A. Shafi, **H. Subramoni**, and DK Panda, Layout aware Hardware assisted Designs for Derived Data Types in MPI, 28th IEEE International Conference on High Performance Computing, Data, Analytics, and Data Science (HiPC'21), Dec 2021.
31. B. Ramesh, J. Hashmi, S. Xu, A. Shafi, M. Ghazimirsaeed, M. Bayatpour, **H. Subramoni**, and DK Panda, Towards Architecture-aware Hierarchical Communication Trees on Modern HPC Systems, 28th IEEE International Conference on High Performance Computing, Data, Analytics, and Data Science (HiPC'21), Dec 2021.

32. P. Kousha, Kamal Raj Sankarapandian Dayala Ganesh Ram, M. Kedia, **H. Subramoni**, A. Jain, A. Shafi, DK Panda, Trey Dockendorf, Heechang Na, and K. Tomko, INAM: Cross-stack Profiling and Analysis of Communication in MPI-based Applications, Practice and Experience in Advanced Research Computing 2021 (PEARC'21), Jul 2021.
33. M. Bayatpour, N. Sarkauskas, **H. Subramoni**, J. Hashmi, and DK Panda, BluesMPI: Efficient MPI Non-blocking Alltoall Offloading Designs on Modern BlueField Smart NICs, Intl Supercomputing Conference 2021 (ISC'21), Jun 2021.
34. K. Khorassani, J. Hashmi, C. Chu, C. Chen, **H. Subramoni**, and D. K. Panda, Designing a ROCm-aware MPI Library for AMD GPUs: Early Experiences, International Supercomputing Conference 2021, Jun 2021.
35. M. Bayatpour, N. Sarkauskas, **H. Subramoni**, J. Hashmi, and D. K. Panda, BluesMPI: Efficient MPI Non-blocking Alltoall Offloading Designs on Modern BlueField Smart NICs, International Supercomputing Conference 2021, Jun 2021.
36. Q. Anthony, S. Xu, **H. Subramoni**, and D. K. Panda, Scaling Single-Image Super-Resolution Training on Modern HPC Clusters: Early Experiences, Scalable Deep Learning over Parallel And Distributed Infrastructures (ScaDL), May 2021.
37. A. Jain, T. Moon, T. Benson, **H. Subramoni**, S. Jacobs, D. K. Panda, and B. Essen, SUPER: SUB-Graph Parallelism for TransformerS, 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS), May 2021.
38. Q. Zhou, C. Chu, N. Senthil Kumar, P. Kousha, M. Ghazimirsaeed, H. Subramoni, and D. K. Panda, Designing High-Performance MPI Libraries with On-the-fly Compression for Modern GPU Clusters, 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS), May 2021.
39. K. Khorassani, C. Chu, Q. Anthony, **H. Subramoni**, and D. K. Panda, Adaptive and Hierarchical Large Message All-to-all Communication Algorithms for Large-scale Dense GPU Systems, The 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, (CCGrid), May 2021.
40. A. Shafi, J. Hashmi, **H. Subramoni**, and D. K. Panda, Efficient MPI-based Communication for GPU-Accelerated Dask Applications, The 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, (CCGrid), May 2021.
41. A. Shafi, J. Hashmi, **H. Subramoni**, and D. K. Panda, Blink: Towards Efficient RDMAbased Communication Coroutines for Parallel Python Applications, 27th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC), Dec 2020.
42. M. Ghazimirsaeed, Q. Zhou, A. Ruhela, M. Bayatpour, **H. Subramoni**, and D. K. Panda, A Hierarchical and Load-Aware Design for Large Message Neighborhood Collectives, SC 2020, Nov 2020.
43. A. Jain, A. Awan, A. Aljuhani, J. Hashmi, Q. Anthony, **H. Subramoni**, D. K. Panda, R. Machiraju, and A. Parwani, GEMS: GPU Enabled Memory Aware Model Parallelism System for Distributed DNN Training SC 2020, Nov 2020.
44. P. Kousha, S. D. Kamal Raj, **H. Subramoni**, D. K. Panda, H. Na, T. Dockendorf, and K. Tomko, Accelerated Real-time Network Monitoring and Profiling at Scale using OSU INAM, Practice and Experience in Advanced Research Computing (PEARC) 2020, Jul 2020.
45. M. Bayatpour, M. Ghazimirsaeed, S. Xu, **H. Subramoni**, and D. K. Panda, Design and Characterization of Infiniband Hardware Tag Matching in MPI, The 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, Nov 2020.
46. C. Chu, P. Kousha, A. Awan, K. S. Khorassani, **H. Subramoni**, and D. K. Panda, NV-Group: Link-Efficient Reductions for Distributed Deep Learning on Modern Dense GPU Systems The 34th ACM International Conference on Supercomputing (ICS 20), June 2020.
47. M. Bayatpour, J. Hashmi, S. Chakraborty, K. K. Suresh, M. Ghazimirsaeed, B. Ramesh, **H. Subramoni**, and D. K. Panda, Communication-Aware Hardware-Assisted MPI Overlap Engine, Intl Supercomputing Conference (ISC 20), June 2020.

48. A. Awan, A. Jain, Q. Anthony, **H. Subramoni**, and D. K. Panda, HyPar-Flow: Exploiting MPI and Keras for Scalable Hybrid-Parallel DNN Training with TensorFlow, Intl Supercomputing Conference (ISC 20), June 2020.
49. Q. Anthony, A. Awan, A. Jain, **H. Subramoni**, and D. K. Panda, Efficient Training of Semantic Image Segmentation on Summit using Horovod and MVAPICH2-GDR, Scalable Deep Learning over Parallel and Distributed Infrastructures Workshop (ScaDL), held in conjunction with Intl Parallel and Distributed Processing Symposium (IPDPS 20), May 2020.
50. 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, 34th IEEE International Parallel & Distributed Processing Symposium (IPDPS 20), May 2020.
51. K. Kandadai Suresh, B. Ramesh, M. Ghazimirsaeed, M. Bayatpour, J. Hashmi, **H. Subramoni**, and D. K. Panda, Performance Characterization of Network Mechanisms for Non-Contiguous Data Transfers in MPI, Workshop on Scalable Networks for Advanced Computing Systems (SNACS), held in conjunction with Intl Parallel and Distributed Processing Symposium (IPDPS 20), May 2020.
52. A. Ruhela, S. Xu, K. Vadambacheri Manian, **H. Subramoni**, and D. K. Panda, Analyzing and Understanding the Impact of Interconnect Performance on HPC, Big Data, and Deep Learning Applications: A Case Study with InfiniBand EDR and HDR, Workshop on Scalable Networks for Advanced Computing Systems (SNACS), held in conjunction with Intl Parallel and Distributed Processing Symposium (IPDPS 20), May 2020.
53. P. Kousha, B. Ramesh, K. Suresh, C. Chu, A. Jain, N. Sarkauskas, and D. K. Panda, Designing a Profiling and Visualization Tool for Scalable and In-Depth Analysis of High-Performance GPU Clusters, 26th IEEE International Conference on High Performance Computing, Data, Analytics and Data Science (HiPC 19), Dec 2019
54. C. Chu, J. Hashmi, K. Shafie Khorassani, **H. Subramoni**, and D. K. Panda, High-Performance Adaptive MPI Derived Datatype Communication for Modern Multi-GPU Systems, 26th IEEE International Conference on High Performance Computing, Data, Analytics and Data Science (HiPC 19), Dec 2019.
55. S. Xu, J. Hashmi, S. Chakraborty, **H. Subramoni**, and D. K. Panda, Design and Evaluation of Shared Memory Communication Benchmarks on Emerging Architectures using MVAPICH2, Third Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, held in conjunction with SC 19, Nov 2019.
56. A. Ruhela, B. Ramesh, S. Chakraborty, **H. Subramoni**, J. Hashmi, and D. K. Panda, Leveraging Network-level parallelism with Multiple Process-Endpoints for MPI Broadcast, Third Annual Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware, held in conjunction with SC 19, Nov 2019
57. K. V. Manian, C. Chu, A. Awan, K. S. Khorassani, **H. Subramoni**, and D. K. Panda, OMBUM: Design, Implementation, and Evaluation of CUDA Unified Memory Aware MPI Benchmarks, 10th International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, held in conjunction with SC 19, Nov 2019.
58. A. Jain, A. Awan, Q. Anthony, **H. Subramoni**, and D. K. Panda, Performance Characterization of DNN Training using TensorFlow and PyTorch on Modern Clusters, 21st IEEE International Conference on Cluster Computing, Sep 2019.
59. S. Chakraborty, S. Xu, **H. Subramoni**, and D. K. Panda, Designing Scalable and Highperformance MPI Libraries on Amazon Elastic Fabric Adapter, HOT Interconnects 26, Aug 2019.
60. A. Awan, A. Jain, C. Chu, **H. Subramoni**, and D. K. Panda, Communication Profiling and Characterization of Deep Learning Workloads on Clusters with High-Performance Interconnects, HOT Interconnects 26, Aug 2019.
61. K. Shafie Khorassani, C. Chu, **H. Subramoni**, and D. K. Panda, Performance Evaluation of MPI Libraries on GPU-enabled OpenPOWER Architectures: Early Experiences, International Workshop on OpenPOWER for HPC, held in conjunction with ISC19, Jun 2019.
62. J. Hashmi, S. Chakraborty, M. Bayatpour, **H. Subramoni**, and DK Panda, FALCON: Efficient Designs for Zero-copy MPI Datatype Processing on Emerging Architectures, 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS '19), May 2019 **Best Paper Finalist**

63. J. Hashmi, S. Chakraborty, M. Bayatpour, **H. Subramoni**, and D. K. Panda, Design and Characterization of Shared Address Space MPI Collectives on Modern Architectures, The 19th Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CCGRID 2019), May 2019.
64. A. Awan, J. Bedorf, C. Chu, **H. Subramoni**, and D. K. Panda, Scalable Distributed DNN Training using TensorFlow and CUDA-Aware MPI: Characterization, Designs, and Performance Evaluation, The 19th Annual IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CCGRID 2019), May 2019.
65. K. Vadambacheri Manian, A. Awan, A. Ruhela, C. Chu, **H. Subramoni** and D. K. Panda, Characterizing CUDA Unified Memory (UM)-Aware MPI Designs on Modern GPU Architectures, 12th Workshop on General Purpose Processing Using GPU (GPGPU 2019), held in conjunction with ASPLOS 2019, Apr 2019
66. A. Awan, C. Chu, **H. Subramoni**, X. Lu, and DK Panda, OC-DNN: Exploiting Advanced Unified Memory Capabilities in CUDA 9 and Volta GPUs for Out-of-Core DNN Training, 25th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2018
67. S. Chakraborty, M. Bayatpour, J. Hashmi, **H. Subramoni**, and DK Panda, Cooperative Rendezvous Protocols for Improved Performance and Overlap, 2018 The International Conference for High Performance Computing, Networking, Storage, and Analysis, Nov 2018, **Best Paper Finalist**
68. M. Bayatpour, J. Hashmi, S. Chakraborty, P. Kousha, **H. Subramoni**, and D. K. Panda, SALaR: Scalable and Adaptive Designs for Large Message Reduction Collectives, 2018 IEEE International Conference on Cluster Computing, Sep 2018
69. A. A. Awan, C-H. Chu, **H. Subramoni**, D. K. Panda, Optimized Broadcast for Deep Learning Workloads on Dense-GPU InfiniBand Clusters: MPI or NCCL? The 25th European MPI Users Group Meeting, Sep 2018
70. M. Li, X. Lu, **H. Subramoni**, D. K. Panda, Multi-Threading and Lock-Free MPI RMA Based Graph Processing on KNL and POWER Architectures, The 25th European MPI Users Group Meeting, Sep 2018
71. A. Ruhela, **H. Subramoni**, S. Chakraborty, M. Bayatpour, P. Kousha, D. K. Panda, Efficient Asynchronous Communication Progress for MPI without Dedicated Resources, The 25th European MPI Users Group Meeting, Sep 2018
72. J. Hashmi, S. Chakraborty, M. Bayatpour, **H. Subramoni**, and D. K. Panda, Designing Efficient Shared Address Space Reduction Collectives for Multi-/Many-cores, 32nd IEEE International Parallel & Distributed Processing Symposium (IPDPS '18), May 2018
73. M. Li, X. Lu, **H. Subramoni**, and D. K. Panda, Designing Registration Caching Free High-Performance MPI Library with Implicit On-Demand Paging (ODP) of InfiniBand, IEEE International Conference on High Performance Computing (HiPC'17), Dec 2017
74. A. Awan, **H. Subramoni**, and DK Panda, An In-depth Performance Characterization of CPU- and GPU-based DNN Training on Modern Architectures, 3rd Workshop on Machine Learning in High Performance Computing Environments, held in conjunction with SC17, Nov 2017
75. J. Hashmi, K. Hamidouche, **H. Subramoni**, and D. K. Panda, Kernel-assisted Communication Engine for MPI on Emerging Manycore Processors, IEEE International Conference on High Performance Computing (HiPC'17), Dec 2017
76. M. Bayatpour, S. Chakraborty, **H. Subramoni**, X. Lu, and D. K. Panda, Scalable Reduction Collectives with Data Partitioning-based Multi-Leader Design, SuperComputing 2017, Nov 2017
77. S. Ramesh, A. Maheo, S. Shende, A. Malony, **H. Subramoni**, and D. K. Panda, MPI Performance Engineering with the MPI Tool Interface: the Integration of MVAPICH and TAU EuroMPI 2017, Sep 2017 **Best Paper**
78. Contention Aware Kernel-Assisted MPI Collectives for Multi/Many-core Systems, IEEE International Conference on Cluster Computing (Cluster), Sep 2017 **Best Paper Finalist**
79. C. Chu, X. Lu, A. Awan, **H. Subramoni**, J. Hashmi, Bracy Elton, and D. K. Panda, Efficient and Scalable Multi-Source Streaming Broadcast on GPU Clusters for Deep Learning, International Conference on Parallel Processing (ICPP), Aug 2017

80. J. Hashmi, M. Li, **H. Subramoni**, and D. K. Panda, Exploiting and Evaluating OpenSHMEM on KNL Architecture, Fourth Workshop on OpenSHMEM and Related Technologies (OpenSHMEM), Aug 2017
81. **H. Subramoni**, S. Chakraborty, and D. K. Panda, Dynamic and Adaptive MPI Point-to-point Communication Protocols for Efficient Overlap of Computation and Communication, IEEE International Supercomputing Conference (ISC 17), Jun 2017 **Best Paper**
82. X. Lu, D. Shankar, S. Gugnani, **H. Subramoni**, and D. K. Panda, Impact of HPC Cloud Networking Technologies on Accelerating Hadoop RPC and HBase, 8th IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom '16), Dec 2016
83. C. Chu, K. Hamidouche, **H. Subramoni**, A. Venkatesh, B. Elton, and D. K. Panda, Efficient Reliability Support for Hardware Multicast-based Broadcast in GPU-enabled Streaming Applications, First Workshop on Optimization of Communication in HPC runtime systems (COMHPC, SC Workshop), Nov 2016
84. M. Li, K. Hamidouche, X. Lu, **H. Subramoni**, J. Zhang, and D. K. Panda, Designing MPI Library with On-Demand Paging (ODP) of InfiniBand: Challenges and Benefits, SuperComputing 2016, Nov 2016
85. C. Chu, K. Hamidouche, **H. Subramoni**, A. Venkatesh, B. Elton, and D. K. Panda, Designing High Performance Heterogeneous Broadcast for Streaming Applications on GPU Clusters, 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD'16), Oct 2016
86. M. Bayatpour, **H. Subramoni**, S. Chakraborty, and D. K. Panda, Adaptive and Dynamic Design for MPI Tag Matching, IEEE Cluster 2016, Sep 2016, **Best Paper Nominee**
87. **H. Subramoni**, A. Augustine, M. Arnold, J. Perkins, X. Lu, K. Hamidouche, and D. K. Panda, INAM²: InfiniBand Network Analysis and Monitoring with MPI, IEEE International Supercomputing Conference (ISC'16). June 2016
88. C-H. Chu, K. Hamidouche, A. Venkatesh, D. S. Banerjee, **H. Subramoni** and D. K. Panda, Offloaded GPU Collectives using CORE-Direct and CUDA Capabilities on IB Clusters, IEEE International Parallel & Distributed Processing Symposium (IPDPS'16) May 2016
89. A. Venkatesh, K. Hamidouche, **H. Subramoni**, and D. K. Panda, Offloaded GPU Collectives using CORE-Direct and CUDA Capabilities on IB Clusters, IEEE International Conference on High Performance Computing (HiPC'15), Dec 2015
90. **H. Subramoni**, A. Venkatesh, K. Hamidouche, K. Tomko and D. Panda, Impact of InfiniBand DC Transport Protocol on Energy Consumption of All-to-all Collective Algorithms, Int'l Symposium on Hot Interconnects (HotI), Aug 2015
91. A. Awan, K. Hamidouche, A. Venkatesh, J. Perkins, **H. Subramoni**, and D. K. Panda, GPU-Aware Design, Implementation, and Evaluation of Non-blocking Collective Benchmarks, EuroMPI, Sep 2015
92. M. Li, **H. Subramoni**, K. Hamidouche, X. Lu, and D. K. Panda, High Performance MPI Datatype Support with User-mode Memory Registration: Challenges, Designs and Benefits, IEEE Cluster'15, Sep 2015
93. Exploiting GPUDirect RDMA in Designing High Performance OpenSHMEM for NVIDIA GPU Clusters, K. Hamidouche, A. Venkatesh, A. Awan, **H. Subramoni**, and D. K. Panda, IEEE Cluster'15, Sep 2015
94. **H. Subramoni**, A. A. Awan, K. Hamidouche, D. Pekurovsky, A. Venkatesh, S. Chakraborty, K. Tomko, and D. K. Panda, Designing Non-Blocking Personalized Collectives with Near Perfect Overlap for RDMA-Enabled Clusters, IEEE International Supercomputing Conference (ISC'15). July 2015
95. S. Chakraborty, **H. Subramoni**, J. Perkins and D. K. Panda, On-demand Connection Management for OpenSHMEM and OpenSHMEM+MPI, Workshop on High-level Parallel Programming Models and Supportive Environments (HIPS'15), May 2015
96. S. Chakraborty, **H. Subramoni**, A. Moody, A. Venkatesh, J. Perkins and D. K. Panda, Non-blocking PMI Extensions for Fast MPI Startup, Int'l Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2015), May 2015.

97. A. Venkatesh, **H. Subramoni**, K. Hamidouche and D. K. Panda, A High Performance Broadcast Design with Hardware Multicast and GPUDirect RDMA for Streaming Applications on InfiniBand Clusters, IEEE International Conference on High Performance Computing (HiPC'14). Dec. 2014.
98. **H. Subramoni**, K. Kandalla, J. Jose, K. Tomko, K. Schulz, and D. Pekurovsky, Designing Topology Aware Communication Schedules for Alltoall Operations in Large InfiniBand Clusters, Int'l Conference on Parallel Processing (ICPP'14). Sep. 2014.
99. S. Chakraborty, **H. Subramoni**, J. Perkins, A. Moody, M. Arnold, D. K. Panda, PMI Extensions for Scalable MPI Startup, EuroMPI'14. Sep. 2014.
100. **H. Subramoni**, K. Hamidouche, A. Venkatesh, S. Chakraborty and D. K. Panda, Designing MPI Library with Dynamic Connected Transport (DCT) of InfiniBand : Early Experiences. IEEE International Supercomputing Conference (ISC'14). June 2014.
101. **H. Subramoni**, D. Bureddy, K. Kandalla, K. Schulz, B. Barth, J. Perkins, M. Arnold, and D. K. Panda, Design of Network Topology Aware Scheduling Services for Large InfiniBand Clusters, IEEE Cluster '13, Sep. 2013.
102. S. Potluri, D. Bureddy, K. Hamidouche, A. Venkatesh, K. Kandalla, **H. Subramoni**, and D. K. Panda, MVAPICH-PRISM: A Proxy-based Communication Framework using InfiniBand and SCIF for Intel MIC Clusters, Int'l Conference on Supercomputing (SC '13), Nov. 2013.
103. K. Hamidouche, S. Potluri, **H. Subramoni**, K. Kandalla and D. K. Panda, MIC-RO: Enabling Efficient Remote Offload on Heterogeneous Many Integrated Core (MIC) Clusters with InfiniBand, Int'l Conference on Supercomputing (ICS '13), June 2013.
104. S. Potluri, D. Bureddy, H. Wang, **H. Subramoni** and D. K. Panda, Extending OpenSHMEM for GPU Computing, Int'l Parallel and Distributed Processing Symposium (IPDPS '13), May 2013.
105. **H. Subramoni**, S. Potluri, K. Kandalla, B. Barth, J. Vienne, J. Keasler, K. Tomko, K. Schulz, A. Moody and D. K. Panda, Design of a Scalable InfiniBand Topology Service to Enable Network-Topology-Aware Placement of Processes, Int'l Conference on Supercomputing (SC '12), Nov. 2012. (**Best Paper and Best Student Paper Finalist**)
106. **H. Subramoni**, J. Vienne and D. K. Panda, A Scalable InfiniBand Network-Topology-Aware Performance Analysis Tool for MPI, Int'l Workshop on Productivity and Performance (Proper '12), in conjunction with EuroPar, Aug. 2012.
107. N. S. Islam, M. W. Rahman, J. Jose, R. Rajachandrasekar, H. Wang, **H. Subramoni**, C. Murthy and D. K. Panda, High Performance RDMA-Based Design of HDFS over InfiniBand, Int'l Conference on Supercomputing (SC '12), Nov. 2012.
108. R. Rajachandrasekar, J. Jaswani, **H. Subramoni** and D. K. Panda, Minimizing Network Contention in InfiniBand Clusters with a QoS-Aware Data-Staging Framework, Cluster (Cluster '12), Sep. 2012.
109. K. Kandalla, A. Buluç, **H. Subramoni**, K. Tomko, J. Vienne, L. Olikier, and D. K. Panda, Can Network-Offload based Non-Blocking Neighborhood MPI Collectives Improve Communication Overheads of Irregular Graph Algorithms?. International Workshop on Parallel Algorithms and Parallel Software (IWPAPS 2012), in conjunction with Cluster, Sep. 2012.
110. J. Huang, X. Ouyang, J. Jose, M. W. Rahman, H. Wang, M. Luo, **H. Subramoni**, C. Murthy and D. K. Panda, High-Performance Design of HBase with RDMA over InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS '12), May 2012.
111. K. Kandalla, U. Yang, J. Keasler, T. Kolev, A. Moody, **H. Subramoni**, K. Tomko, J. Vienne and D. K. Panda, Designing Non-blocking Allreduce with Collective Offload on InfiniBand Clusters: A Case Study with Conjugate Gradient Solvers Int'l Parallel and Distributed Processing Symposium (IPDPS '12), May 2012.
112. S. P. Raikar, **H. Subramoni**, K. Kandalla, J. Vienne and D. K. Panda, Designing Network Failover and Recovery in MPI for Multi-Rail InfiniBand Clusters, Int'l Workshop on System Management Techniques, Processes, and Services (SMTPS), in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '12), May 2012.

113. J. Jose, **H. Subramoni**, K. Kandalla, M. W. Rahman, H. Wang, S. Narravula and D. K. Panda, Scalable Memcached design for InfiniBand Clusters using Hybrid Transports, Int'l Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2012), May 2012.
114. **H. Subramoni**, K. Kandalla, J. Vienne, S. Sur, B. Barth, K. Tomko, R. McLay, K. Schulz and D. K. Panda, Design and Evaluation of Network Topology-/Speed-Aware Broadcast Algorithms for InfiniBand Clusters, Cluster '11, Sep. 2011.
115. J. Jose, **H. Subramoni**, M. Luo, M. Zhang, J. Huang, M. W. Rahman, N. S. Islam, X. Ouyang, H. Wang, S. Sur and D. K. Panda, Memcached Design on High Performance RDMA Capable Interconnects, Int'l Conference on Parallel Processing (ICPP '11), Sep. 2011.
116. N. Dandapanthula, **H. Subramoni**, J. Vienne, K. Kandalla, S. Sur, D. K. Panda, and R. Brightwell, INAM - A Scalable InfiniBand Network Analysis and Monitoring Tool, 4th Int'l Workshop on Productivity and Performance (PROPER 2011), in conjunction with EuroPar, Aug. 2011.
117. K. Kandalla, **H. Subramoni**, J. Vienne, K. Tomko, S. Sur and D. K. Panda, Designing Non-blocking Broadcast with Collective Offload on InfiniBand Clusters: A Case Study with HPL, Hot Interconnect '11, Aug. 2011.
118. K. Kandalla, **H. Subramoni**, K. Tomko, D. Pekurovsky, S. Sur and D. K. Panda, High-Performance and Scalable Non-Blocking All-to-All with Collective Offload on InfiniBand Clusters: A Study with Parallel 3D FFT, Int'l Supercomputing Conference (ISC), June 2011.
119. **H. Subramoni**, P. Lai, S. Sur and D. K. Panda, Improving Application Performance and Predictability using Multiple Virtual Lanes in Modern Multi-Core InfiniBand Clusters, Int'l Conference on Parallel Processing (ICPP'10), Sep. 2010.
120. M. Luo, S. Potluri, P. Lai, E. P. Mancini, **H. Subramoni**, K. Kandalla, S. Sur and D. K. Panda, High Performance Design and Implementation of Nemesis Communication Layer for Two-sided and One-Sided MPI, The Workshop on Parallel Programming Models and System Software (P2S2 2012), in conjunction with Int'l Conference on Parallel Processing (ICPP'10), Sep. 2010.
121. **H. Subramoni**, K. Kandalla, S. Sur and D. K. Panda, Design and Evaluation of Generalized Collective Communication Primitives with Overlap using ConnectX-2 Offload Engine, Int'l Symposium on Hot Interconnects (HotI), Aug. 2010.
122. **H. Subramoni**, P. Lai, R. Kettimuthu and D. K. Panda, High Performance Data Transfer in Grid Environment Using GridFTP over InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid '10), May 2010.
123. **H. Subramoni**, F. Petrini, V. Agarwal and D. Pasetto, Streaming, low-latency communication in on-line trading systems, Int'l Workshop on System Management Techniques, Processes, and Services (SMTPS), in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '10), April 2010.
124. K. Kandalla, **H. Subramoni**, A. Vishnu and D. K. Panda, Designing Topology-Aware Collective Communication Algorithms for Large Scale InfiniBand Clusters: Case Studies with Scatter and Gather, The 10th Workshop on Communication Architecture for Clusters (CAC '10), in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '10), April 2010.
125. P. Lai, **H. Subramoni**, S. Narravula, A. Mamidala and D. K. Panda, Designing Efficient FTP Mechanisms for High Performance Data-Transfer over InfiniBand, Int'l Conference on Parallel Processing (ICPP '09), Sep. 2009.
126. **H. Subramoni**, P. Lai, M. Luo, Dhabaleswar K. Panda, RDMA over Ethernet - A Preliminary Study, Workshop on High Performance Interconnects for Distributed Computing (HPIDC '09), in conjunction with Cluster, Sep. 2009.
127. **H. Subramoni**, M. Koop, and D. K. Panda, Designing Next Generation Clusters: Evaluation of InfiniBand DDR/QDR on Intel Computing Platforms, 17th Annual Symposium on High-Performance Interconnects (HotI '09), Aug. 2009.
128. K. Kandalla, **H. Subramoni**, G. Santhanaraman, M. Koop and D. K. Panda, The 9th Workshop on Communication Architecture for Clusters (CAC '09), in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '09), May 2009.

129. **H. Subramoni**, G. Marsh, S. Narravula, P. Lai and D.K. Panda, Design and Evaluation of Benchmarks for Financial Applications using Advanced Message Queuing Protocol (AMQP) over InfiniBand, Workshop on High Performance Computational Finance (in conjunction with SC '08), Austin, TX, November 2008.
130. S. Narravula, **H. Subramoni**, P. Lai, R. Noronha and D. K. Panda Performance of HPC Middleware over InfiniBand WAN, Int'l Conference on Parallel Processing (ICPP '08), Portland, Oregon, Sep. 2008.