

Akshay Venkatesh
180 Alicante Drive, Apt 319,
San Jose, CA, USA - 95134

E-Mail: akvenkatesh@nvidia.com; akshay.v.3.14@gmail.com
Web: <http://akshay-venkatesh.github.io/home-page>
Phone: (585) 733 7992

Objective:

Exploiting architectural advances (network/accelerators) to make applications/communication-stacks efficient; Solving problems in parallel and distributed computing using the plethora of programming models available; Accelerating machine learning applications using high performance computing techniques

Research Interests:

Heterogeneous Architectures, Power-Aware Computing, Parallel Programming Models, Network Topologies, Parallel Algorithms, Optimizing Parallel Applications, Optimizing Machine Learning Applications

Education:

The Ohio State University, Columbus, OH.

PhD Candidate, Computer Science and Engineering

Aug'11 - Dec'16

(Major: High Performance Computing; Minors: Machine Learning, Theory)

National Institute of Technology Karnataka (NITK), Surathkal, India.

Bachelor of Technology, Information Technology

Jun'11

Experience:

Senior Engineer, System Software

NVIDIA, Santa Clara, USA.

Jan'17 - Present

Currently involved in the design and development of the following projects:

MPI-GDS: Developed MPI operations to be offloaded to GPUs using GDR-aSync - a proprietary NVIDIA technology that allows GPUs to progress communication operations

NVSHMEM: Working on different components of the PGAS model variation for CPU/GPU PGAS primitives to/from CPU/GPU buffers

UCX: Working on the common communication stack for MPI, PGAS, and task-based programming models

NCCL: Worked on developing benchmarks to compare and evaluate NCCL (NVIDIA's collective library) operations with competing stacks

Graduate Research Associate

The Ohio State University, USA.

Aug'12 - Jan'17

Involved in the design and development of MVAPICH2 - an open source, high performance and scalable MPI implementation for (Message Passing Interface) for clusters with the modern InfiniBand, 10GigE/iWARP, RoCE networking technologies and Accelerators/Coprocessors. MVAPICH2 is used by over 2,450 organizations world-wide in 76 countries.

Contributions to the following projects involving MVAPICH2:

MVAPICH2-GPU: High-performance MPI for Heterogeneous Clusters with NVIDIA GPUs.

MVAPICH2-MIC: High-performance MPI for Heterogeneous Clusters with Intel MICs.

MVAPICH2-X: High-performance MPI for Clusters with a Unified PGAS and MPI communication runtime

MVAPICH2-EA: High-performance Energy-Aware MPI for Infiniband clusters.

High-performance and Scalable Designs for blocking MPI Collectives on InfiniBand Clusters.

High-performance and Scalable Designs for non-blocking MPI Collectives on InfiniBand Clusters with Mellanox Core-Direct offloading Technology.

High-performance and Scalable Designs for non-blocking MPI Collectives on InfiniBand Clusters with GPUDirect-RDMA and Mellanox Core-Direct offloading Technology.

Framework to Measure Energy/Power of MPI Primitives using Intel RAPL.

Research Intern

Pacific Northwest National Lab, USA.

May '14 - Aug '14

Involved in the design and development of an energy-aware MPI Runtime (EAM) - Now available publicly as MVAPICH2-EA InfiniBand clusters.

Conference and Workshop Publications:

1. **A. Venkatesh**, Ching-Hsiang Chu, Khaled Hamidouche, Sreeram Potluri, Davide Rossetti, and Dhableswar Panda - *MPI-GDS: High Performance MPI Designs with GPUDirect-aSync for CPU-GPU Control Flow Decoupling* (ICPP 17), August 2017
2. A. Awan, K. Hamidouche, **A. Venkatesh**, and D. K. Panda - *Efficient Large Message Broadcast using NCCL and CUDA-Aware MPI for Deep Learning* (Best Paper Runner-Up)- The 23rd European MPI Users' Group Meeting (EuroMPI 16), Sep 2016
3. C. Chu, K. Hamidouche, **A. Venkatesh**, D. Banerjee, H. Subramoni, and D. K. Panda - *Exploiting Maximal Overlap for Non-Contiguous Data Movement Processing on Modern GPU-enabled System* - The 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS '16), May 2016
4. C. Chu, K. Hamidouche, **A. Venkatesh**, A. Awan, and D. K. Panda - *CUDA Kernel based Collective Reduction Operations on Large-scale GPU Clusters* - 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'16), May 2016
5. **A. Venkatesh**, K. Hamidouche, H. Subramoni and DK Panda - *Offloaded GPU Collectives using CORE-Direct and CUDA Capabilities on IB Clusters* - High Performance Computing (HiPC '15), December 2015 (Accepted)
6. **A. Venkatesh**, A. Vishnu, K. Hamidouche, N. Tallent, D. K. Panda, D. Kerbyson, and A. Hoise - *A Case for Application-Oblivious Energy-Efficient MPI Runtime* - Supercomputing (SC '15), Nov 2015 (Best Student Paper Finalist)
7. K. Hamidouche, **A. Venkatesh**, A. Awan, H. Subramoni, and D. K. Panda - *Exploiting GPUDirect RDMA in Designing High Performance OpenSHMEM for NVIDIA GPU Clusters* - IEEE Cluster 2015, Sep 2015
8. H. Subramoni, 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* - Int'l Super Computing Conference (ISC 15), Jul 2015
9. S. Chakraborty, H. Subramoni, A. Moody, **A. Venkatesh**, J. Perkins, and D. K. Panda - *Non-blocking PMI Extensions for Fast MPI Startup* - CCGrid '15, May 2015
10. **A. Venkatesh**, H. Subramoni, K. Hamidouche, DK Panda - *A High Performance Broadcast Design with Hardware Multicast and GPUDirect RDMA for Streaming Applications on Infiniband Clusters* - High Performance Computing '14, Dec 2014
11. H. Subramoni, K. Hamidouche, **A. Venkatesh**, S. Chakraborty, DK Panda - *Designing MPI Library with Dynamic Connected Transport (DCT) of InfiniBand: Early Experiences* - Int'l Super Computing Conference (ISC 14), May 2014
12. **A. Venkatesh**, S. Potluri, R. Rajachandrasekar, M. Luo, K. Hamidouche and D. K. Panda - *High Performance Alltoall and Allgather designs for InfiniBand MIC Clusters* - International Parallel and Distributed Processing Symposium (IPDPS '14), May 2014
13. J. Jose, J. Zhang, **A. Venkatesh**, S. Potluri and D. K. Panda - *A Comprehensive Performance Evaluation of OpenSHMEM Libraries on InfiniBand Clusters* - First OpenSHMEM Workshop: Experiences, Implementations and Tools (OpenSHMEM '13), October 2013
14. M. Luo, M. Li, **A. Venkatesh**, X. Lu and D. K. Panda - *UPC on MIC: Early Experiences with Native and Symmetric Modes* - Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '13), October 2013
15. 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), November 2013
16. S. Potluri, K. Hamidouche, **A. Venkatesh**, D. Bureddy and D. Panda - *Efficient Inter-node MPI Communication using GPUDirect RDMA for InfiniBand Clusters with NVIDIA GPUs* - Int'l Conference on Parallel Processing (ICPP '13), October 2013
17. **A. Venkatesh**, K. Kandalla and D. K. Panda - *Optimized MPI Gather collective for Many Integrated Core (MIC) InfiniBand Clusters* - Extreme Scaling Workshop, August 2013
18. K. Kandalla, **A. Venkatesh**, K. Hamidouche, S. Potluri and D. K. Panda - *Designing Optimized MPI Broadcast and Allreduce for Many Integrated Core (MIC) InfiniBand Clusters* - Int'l Symposium on High-Performance Interconnects (HotI '13), August 2013
19. S. Potluri, **A. Venkatesh**, D. Bureddy, K. Kandalla, and D. K. Panda - *Efficient Intra-node Communication on Intel-MIC Clusters* - International Symposium on Cluster, Cloud and Grid Computing (CCGrid), May 2013
20. **A. Venkatesh**, K. Kandalla, D. K. Panda - *Evaluation of Energy Characteristics of MPI Communication Primitives with RAPL* - High-Performance, Power-Aware Computing (HPPAC), 2013

21. Devendar Bureddy, Hao Wang, **A. Venkatesh**, Sreeram Potluri, Dhabaleswar K. Panda - *OMB-GPU: A Micro-Benchmark Suite for Evaluating MPI Libraries on GPU Clusters* - EUROMPI 2012

Journal Publications:

1. K. Hamidouche, **A. Venkatesh**, A. Awan, H. Subramoni, and D. K. Panda - *CUDA-Aware OpenSHMEM: Extensions and Designs for High Performance OpenSHMEM on GPU Clusters*, ParCo: Elsevier Parallel Computing Journal
2. S.N. Omkar, **Akshay Venkatesh**, Mrunmaya Mudigere - *MPI-based Parallel Synchronous Vector Evaluated Particle Swarm Optimization for Multi-Objective Design Optimization of Composite Structures* - Engineering Applications of Artificial Intelligence, Volume 25, Issue 8, December 2012

Skills:

Programming Languages: C, Java, SQL, Python and Shell scripting.

Operating Systems: Linux and Windows.

Software: Eclipse IDE, GBD, mpiP, HPC-Toolkit and Valgrind.

Parallel Programming Models, Frameworks and Libraries: MPI, CUDA, UPC, OpenSHMEM, OpenMP, MPICH2, MVAPICH2, Infiniband Verbs and OFED.

Awards:

Student Travel Grant, IPDPS 2013

Student Volunteer, SC 2013.

Student Travel Grant, HiPC 2015

Runner-up, OSU-CSE Department Annual Poster Presentation 2015

Professional Service:

IEEE Student Member.

Student Volunteer, SC 2013.