

**Dhabaleswar K. (DK) Panda**  
**Vita**  
**July 2020**

**Communication Address**

Dhabaleswar K. Panda  
Dept. of Computer Science and Engineering  
Dreese Lab 785  
The Ohio State University  
Columbus, OH 43210-1277, USA  
Tel: (614)-292-5199, Fax: (614)-292-2911  
E-mail: panda@cse.ohio-state.edu  
<http://www.cse.ohio-state.edu/~panda>

**Education**

- Ph. D. in Computer Engineering  
University of Southern California, Los Angeles, December 1991.  
Thesis: *Vectorized Interprocessor Communication and Data Movement in Shared-memory Multiprocessors*
- M. S. in Electrical and Communication Engineering  
Indian Institute of Science, Bangalore, India, 1986, **Gold Medalist**.
- B. S. in Electrical Engineering  
Indian Institute of Technology, Kanpur, India, 1984.

**Employment**

2015–	University Distinguished Scholar and Professor	Dept. of Computer Science and Engineering, Ohio State University, Columbus, OH.
2001–	Professor	Dept. of Computer Science and Engineering, Ohio State University, Columbus, OH.
1997–2001	Associate Professor	Dept. of Computer Science and Engineering, Ohio State University, Columbus, OH.
1998–99	Academic Visitor	IBM T.J. Watson Research Center, NY
1991–97	Assistant Professor	Dept. of Computer and Information Science, Ohio State University, Columbus, OH.
1988–91	Research and Teaching Assistant	University of Southern California, Los Angeles, CA.
1986–88	Senior Hardware Engineer	WIPRO Information Technology (R & D), Bangalore, India.
1984–86	Research and Teaching Assistant	Indian Institute of Science, Bangalore, India.
1984	Scientist/Engineer-SB	Indian Space Research Satellite Center, Bangalore, India.

## Research Interests

Parallel computer architecture, high performance networking, InfiniBand, network-based computing, 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

- **Keynote Speaker**, International Symposium on Benchmarking, Measuring and Optimizing (Bench '19) November 2019.
- **Keynote Speaker**, HPC-AI Advisory Council UK Conference, September 2019.
- **Keynote Speaker**, HPC-AI Advisory Council Perth Conference, August 2019.
- **Keynote Speaker**, Int'l Symposium on Intelligent Computers, June 2019.
- **Keynote Speaker**, Scalable Data Analytics in Scientific Computing (SDASC) Workshop, in conjunction with ISC '19 conference, June 2019.
- **Keynote Speaker**, High Performance Machine Learning (HPML) Workshop, in conjunction with CCGrid '19 conference, May 2019.
- **Keynote Speaker**, SCAsia '19 Conference, March 2019.
- **Keynote Speaker**, HiPINEB Workshop, in conjunction with HPCA Conference, February 2019.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, October 2018.
- **Keynote Speaker**, HPC Advisory Council China Conference, October 2018.
- **Best Paper Award (Architecture)**, IEEE Cluster '18, September 2018.
- **Keynote Speaker**, KBNNet Workshop (in conjunction with SigComm Conference), August 2018.
- **Keynote Speaker**, IEEE HPSR Conference, June 2018.
- **Keynote Speaker**, Charm++ Workshop, April 2018.
- **Keynote Speaker**, HPC Advisory Council Switzerland Conference, April 2018.
- **Keynote Speaker**, NCAR SEA Symposium, April 2018.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, February 2018.
- **Keynote Speaker**, IEEE CloudTech Conference, October 2017.
- **Best Paper Award**, EuroMPI '17, September 2017.
- **Hans Meuer Best Paper Award**, ISC '17, June 2017.

- **Rockstar of HPC Award**, InsideHPC, May 2017.
- **Keynote Speaker**, Hot Data Center '17, in conjunction with HPC China Conference, November 2017.
- **Keynote Speaker**, HPC Advisory Council Spain Conference, November 2017.
- **Keynote Speaker**, HPC Advisory Council Spain Conference, September 2017.
- **Keynote Speaker**, HPC Advisory Council Switzerland Conference, April 2017.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, February 2017.
- **Keynote Speaker**, Int'l Workshop on Virtualization Solutions for High-Performance Computing (VisorHPC), January 2017.
- **Keynote Speaker**, Int'l Conference on Cloud Computing Technology and Science, December 2016.
- **Keynote Speaker**, HPC Advisory Council China Conference, October 2016.
- **Keynote Speaker**, HPC Advisory Council Switzerland Conference, March 2016.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, February 2016.
- **Keynote Speaker**, Int'l Symposium on Computational Sciences (ISCS '15), December 2015.
- **Plenary Speaker**, South African High-Performance Computing Conference (CHPC '15), Dec 2015.
- **Keynote Speaker**, HPC Advisory Council China Conference, November 2015.
- **Keynote Speaker**, HPC Advisory Council Spain Conference, September 2015.
- **Keynote Speaker**, Int'l Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-6), held in conjunction with VLDB '15, September 2015.
- **Keynote Speaker**, HPC Advisory Council Brazil Conference, August 2015.
- **University Distinguished Scholar**, The Ohio State University, April 2015.
- **Keynote Speaker**, HPC Advisory Council Switzerland Conference, Lugano, Mar 2015.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, Stanford, Feb 2015.
- **Keynote Speaker**, Int'l Workshop on Performance Analysis, of Big Data Systems (PABS), Austin, Feb 2015.
- **Plenary Speaker**, South African High-Performance Computing Conference, Kruger, Dec 2014.
- **Keynote Speaker**, China High Performance Interconnect Forum, Guangzhou, Nov 2014.

- **Keynote Speaker**, HPC China Conference, Guangzhou, Nov 2014.
- **Keynote Speaker**, Int'l Workshop on Big Data Benchmarks, Performance Optimization and Emerging Hardware (BPOE-6), Guangzhou, Nov 2014.
- **Keynote Speaker**, HPC Advisory Council China Conference, Guangzhou, Nov 2014.
- **Keynote Speaker**, Int'l Workshop on Advanced Data Management Systems (ADMS), Hangzhou, China September 2014.
- **Keynote Speaker**, HPC Advisory Council Brazil Conference, Sao Paulo, May 2014.
- **Keynote Speaker**, HPC Advisory Council Switzerland Conference, Lugano, March 2014.
- **Keynote Speaker**, Int'l Workshop on Big Data Benchmarks, Performance Optimization and Emerging Hardware (BPOE), Salt Lake City, March 2014.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference and Exascale Workshop, Stanford, Feb 2014.
- **Keynote Speaker**, CCF Big Data Conference, Beijing, China, December 2013.
- **Keynote Speaker**, Big Data Technology Conference, Beijing, China, December 2013.
- **Keynote Speaker**, Big Data Forum, HPC China, Guilin, October 2013.
- **Keynote Speaker**, HPC Advisory Council China Conference, Guilin, October 2013.
- **Keynote Speaker**, HPC Advisory Council Stanford Conference, Stanford, Feb 2013.
- **Keynote Speaker**, HP-CAST Conference, Salt Lake City, November 2012.
- **Keynote Speaker**, HPC Advisory Council China Conference, Zhangjiajie, October 2012.
- **Keynote Speaker**, IEEE Cluster Conference, Beijing, China, September 2012.
- **Keynote Speaker**, HPC Advisory Council Spain Conference, Malaga, Spain, September 2012.
- **Keynote Speaker**, Israel Supercomputing Conference, Tel Aviv, Israel, Feb 2012.
- **Keynote Speaker**, Int'l Workshop on Co-Design, Beijing, China, Oct 2011.
- **Keynote Speaker**, HPC Advisory Council Workshop, Jinan, China, Oct 2011.
- **Innovator Award**, College of Engineering, The Ohio State University, May 2011.
- **Keynote Speaker**, Int'l Conference on Parallel and Distributed Systems (ICPADS), Shanghai, Dec 2010.
- **Keynote Speaker**, HPC China Conference, Beijing, Oct 2010.
- **Keynote Speaker**, HPC Advisory Council Workshop, Beijing, Oct 2010.

- **Best Paper**, Int'l Supercomputing Conference, (ISC '10), Hamburg, Germany, May 2010.
- **Plenary Speaker**, Int'l Cluster Computing Conference, (Cluster '09), New Orleans, Sept. 2009.
- **Plenary Speaker**, Int'l Cluster Computing Conference, (Cluster '08), Tsukuba, Japan, Oct. 2008.
- **Best Paper**, Int'l Cluster Computing Conference, (Cluster '08), Tsukuba, Japan, Oct. 2008.
- **Keynote Speaker**, Int'l Workshop on System Management Techniques, Processes and Services (SMPTS), April 2008.
- **IEEE Fellow**, IEEE, 2008.
- **Best Paper**, Int'l Cluster Computing Conference, (Cluster '07), Austin, Sept. 2007.
- **Keynote Speaker**, Int'l Conference on Parallel and Distributed Systems (ICPADS), July 2006.
- **Outstanding Teaching Award**, Department of Computer Science and Engineering, The Ohio State University, May 2006.
- **Lumley Research Award**, College of Engineering, The Ohio State University, May 2006.
- **Invited visit to Capitol Hill**, Washington, DC, Demonstration of InfiniBand Research and its Benefits to High-End Computing, June 2004.
- **Best Paper in Software Track** (jointly with Dr. Jarek Nieplocha and Vinod Tipparaju from PNNL), Int'l Parallel and Distributed Processing Symposium (IPDPS), April 2003.
- **Keynote Speaker**, Int'l Workshop on High Speed Local Area Networks (HSLN), Nov. 2002.
- **Lumley Research Award**, College of Engineering, The Ohio State University, May 2001.
- **Keynote Speaker**, Scientific Computing Workshop, NASA Glenn Research Center, Apr. 2000.
- **Keynote Speaker**, Int'l Conference on Parallel and Distributed Computing and Systems (PDCS '98), Oct. 1998.
- **Ameritech Faculty Fellow Award**, Ameritech Corporation, 1998.
- **IEEE Chapter Tutorials Program Speaker**, IEEE Computer Society, 1997-2000.
- **IEEE Distinguished Visitors Program Speaker**, IEEE Computer Society, 1997-2000.
- Listed in **Who'sWho in America**, 1997-current.
- Listed in **American Men & Women of Science (AMWS)**, 1997-current.
- **Lumley Research Award**, College of Engineering, The Ohio State University, May 1997.

- **Faculty Early CAREER Development Award**, National Science Foundation, 1995–1998.
- **Outstanding Faculty Recognition Award** by Mortar Board and Sphinx Honor Society, The Ohio State University, February 1992.
- **Graduate Student Recognition Award, Best Teaching Assistant** at the University of Southern California, 1988–89.
- J. Watumull Scholarship, 1988, One of four TOP students enrolled at the University of Southern California from Eastern India.
- Prof. S. V. C. Aiya **Gold Medal, Best M. S. student**, 1984–85, Department of Electrical and Communication Engineering, Indian Institute of Science, Bangalore, India.
- Certificate of Merit, Senate of the Indian Institute of Technology, Kanpur, 1979, **12th rank in the All India Joint Entrance Examination** (taken by over 60,000 candidates) for admission to undergraduate programs at the I. I. Ts.

## Research Grants

1. PI, *MRI: RADiCAL: Reconfigurable Major Research Cyberinfrastructure for Advanced Computational Data Analytics and Machine Learning*, National Science Foundation, Jointly with R. Machiraju, S. Parthasarathy, A. Parwani, and R. Ramnath (OSU), \$770K (total \$1.1M, including OBOR Cost-share), Oct 20–Nov 23.
2. PI, *OAC Core: Small: Next-generation Communication and I/O Middleware for HPC and Deep Learning with Smart NICs*, National Science Foundation, Jointly with Hari Subramoni (OSU), \$500K, July 20–June 23.
3. PI, *Designing MVAPICH2 Support for Sandia Structural Simulation Toolkit*, Sandia National Laboratory, \$79,861, April 20–Sept 20.
4. PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Dec 19–Oct 20.
5. Co-PI, *Operations & Maintenance for the Endless Frontier*, PI - Dan Stanzione, Texas Advanced Computing Center (TACC), National Science Foundation), \$60M, OSU Share (PI) \$2.5M, Sept 19–August 24.
6. PI, *Collaborative Research: Frameworks: Designing Next-Generation MPI Libraries for Emerging Dense GPU Systems*, National Science Foundation, Jointly with Karen Tomko (OSC), Bill Barth (TACC) and Amit Majumdar (SDSC), \$2.09M (OSU+OSC Share: \$1.36M), Nov. 19–Oct. 22.
7. PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, May 19–April 20.
8. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$227,115, April 19–Mar 20.

9. PI, *Optimizing and Tuning MVAPICH2 for ARM*, ARM, \$100,000, Jan 19–Dec 20.
10. PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Nov 18–Nov 19.
11. Co-PI, *Computation for the Endless Frontier*, PI - Dan Stanzione, Texas Advanced Computing Center (TACC), National Science Foundation), \$60M, Sept 18–August 23.
12. PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, Aug 18–July 19.
13. Senior Personnel, *BD Spoke: Community-Driven Data Engineering for Opioid and Substance Abuse in the Rural Midwest*, National Science Foundation \$651,000, April 18–March 21.
14. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$220,963, April 18–Mar 19.
15. PI, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel, \$64,741, Feb. 18–Jan 19.
16. PI, *Designing and Optimizing MVAPICH2 MPI Library on FPGA Systems*, Coventry Computer, \$114,000, Jan 18–Jan 19.
17. PI, *An Infrastructure for Performance Engineering using the MPI Tools Interface*, Lawrence Livermore National Laboratory, \$120,000, Nov 17–Nov 18.
18. PI, *Stampede 2: Operations and Maintenance for the Next Generation of Petascale Computing*, University of Texas Austin (National Science Foundation Sub Award), \$600,000, Oct 17–June 20.
19. PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, Aug 17–July 18.
20. PI, *SI2-SSI: Collaborative Project: FAMILI: High performance and scalable fabric analysis, monitoring and introspection infrastructure for HPC and Big Data*, National Science Foundation, Jointly with Karen Tomko (OSC), \$800,000, July 17–June 20.
21. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$216,070, April 17–Mar 18.
22. PI, *MVAPICH on CNTK w/ Evaluations of MPI-GDR-level Designs and Optimizations*, Microsoft, \$120,000, April '17–March '18.
23. PI, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel, \$64,424, Feb. 17–Jan 18.
24. PI, *MVAPICH2 Enhancements for GPUDirect RDMA*, NVIDIA, \$177,500, Jan. 17–Dec. 17.
25. PI (OSU), *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.

26. PI, *SHF: Large: Collaborative Research: Next Generation Communication Mechanisms Exploiting Heterogeneity, Hierarchy and Concurrency for Emerging HPC systems*, National Science Foundation, Jointly with Karen Tomko (OSC), Bill Barth (TACC) and Amit Majumdar (SDSC), \$2.0M (OSU+OSC Share: \$1.17M), Aug. 16–July. 19.
27. PI, *Student Travel Support for MVAPICH User Group (MUG) Meeting*, National Science Foundation, \$10,000, Aug 16–July 17.
28. PI, *Accelerator-enabled version of the MVAPICH2 Library for use in Numerical Weather Prediction (NWP) software*, CSCS- Switzerland National Supercomputing Centre, \$75,000, May 16–April 17.
29. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$213,768, April 16–Mar 17.
30. PI, *Enhanced PSM Support in MVAPICH2 with Omni-Path*, Intel, \$64,171, Feb. 16–Jan 17.
31. PI, *MVAPICH2 Enhancements for GPUDirect RDMA*, NVIDIA, \$183,000, Jan. 16–Dec. 16.
32. PI, *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.
33. PI, *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.
34. PI, *Failure Recovery Models and Interfaces in MVAPICH*, Lawrence Livermore National Laboratory, \$71,466, Aug 15–Aug 16.
35. Co-PI, *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.
36. PI, *Support and Custom Software Development Relating to the MVAPICH2 Library, Accelerators, MPI and Contemporary Networking*, CSCS- Switzerland National Supercomputing Centre, \$77,500, May 15–April 16.
37. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$212,030, April 15–Mar 16.
38. PI, *MVAPICH2 Enhancements for GPUDirect RDMA*, NVIDIA, \$182,600, Jan. 15–Dec. 15.
39. PI (OSU), *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.
40. PI, *BIGDATA: Scalable Middleware for Managing and Processing Big Data on Next Generation HPC Systems*, in collaboration with Amitava Majumdar and Mahidhar Tatineni (SDSC), National Science Foundation, \$1,080,190, (OSU share: \$720,191), Sept. 14–Aug 17.



41. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$205,709, April 14–Mar 15.
42. PI, *PMI2 Support in MVAPICH2 and Tight Integration with SLURM*, Cray, \$277,856, January 14–April 15.
43. PI, *MVAPICH2 Enhancements for GPUDirect RDMA*, NVIDIA, \$150,000, October 13–December 14.
44. PI, *Eager-HPC Virtualization with SR-IOV*, National Science Foundation, \$98,291, October 13–September 14.
45. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$190,585, April 13–Mar 14.
46. PI, *Enabling, Enhancing and Extending Petascale Computing for Science and Engineering*, University of Texas Austin (National Science Foundation Sub Award), \$600,000, March 13–February 17.
47. Co-PI, *CC-NIE Integration II: Innovations to Transition a Campus Core Cyberinfrastructure to Serve Diverse and Emerging Researcher Needs*, in collaboration with C. Whitacre (PI), P. Calyam, D. Gaitonde and U. Catalyurek, National Science Foundation, \$987,019, October 12–September 14.
48. PI, *MVAPICH2 Enhancements for GPUDirect RDMA and OpenACC Extensions*, NVIDIA, \$115,000, October 12–September 12.
49. PI, *Enhanced MVAPICH2 Design for Stampede with Intel MIC*, Intel, \$50,000, September 12–August 13.
50. PI, *High Performance and Scalable Design of HDFS over InfiniBand*, Mellanox Technologies, Inc., \$200,000, Aug 12–Aug 13.
51. PI, *SHF: Unified Runtime for Supporting Hybrid Programming Models on Heterogeneous Architecture*, in collaboration with Karen Tomko (OSC), K. Schulz and B. Barth (TACC) and Amitava Majumdar (SDSC), National Science Foundation, \$1,799,999, (OSU+OSC share: \$1,045,822), July 12–June 15.
52. PI, *SI2-SSI: A Comprehensive Performance Tuning Framework for the MPI Stack*, in collaboration with Karen Tomko (OSC), K. Schulz and B. Barth (TACC) and Amitava Majumdar (SDSC), National Science Foundation, \$2,152,411, (OSU+OSC share: \$1,251,644), June 12–May 15.
53. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$181,264, April 12–March 13.
54. PI, *High-Performance and Scalable designs for Transactional Middleware*, IBM Open Collaborative Faculty Award, \$100,000, April 12–Mar 13.
55. PI, *HPC Energy Measurement and Optimization*, STTR Phase II, Department of Energy, \$325,000, Oct 11–Sept 13.

56. PI, *High Performance MPI Design for InfiniBand Clusters with GPUs*, NVIDIA, \$115,237, July 11–March 12.
57. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$176,835, April 11–March 12.
58. PI, *Optimizing MVAPICH/MVAPICH2 Library on TACC InfiniBand Systems*, NSF/TACC, \$172,616, Oct 10–Sept 13.
59. PI, *Application of Low-latency Networking and Application Techniques to Transactional and Applications Middlewares*, with Sayantan Sur (Co-PI), IBM Open Collaborative Research Program, \$100,000, Feb 11–Jan 12.
60. PI, *Study of WAN-Level Storage and Distributed Computing with Obsidian Longbow Routers with Encryption*, Avetec/DICE, \$55,389, Nov 10–Jun 11.
61. PI, *Performance Evaluation of Obsidian Longbow Routers with Encryption for Parallel File Systems*, Avetec/DICE, \$79,059, Aug 10–May 11.
62. PI, *Scalability and Fault-Tolerance Properties of InfiniBand Subnet Management*, with Sayantan Sur (Co-PI), Sandia National Laboratory, \$115,885, May 10–Sept 11.
63. PI, *HPC Application Energy Profiling for Energy Optimization*, STTR Phase I, Department of Energy, \$33,000, June 10–Mar 11.
64. PI, *Green Storage for HPC with Solid State Disk (SSD) Technologies*, STTR Phase I, Department of Energy, \$33,000, June 10–Mar 11.
65. PI, *Dynamic Staging Architecture for Accelerating I/O Pipelines*, National Science Foundation, \$90,000, May 10–April 13.
66. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$117,890, April 10–March 11.
67. PI, *Creating Petascale File Systems using Application-Aware Network Offloading*, STTR Phase II, Department of Energy, \$275,000, Sept 09–June 12.
68. PI, *Designing QoS-Aware MPI and File Systems Protocols for Emerging InfiniBand Clusters*, National Science Foundation, \$491,570, Sept 09–Aug 12.
69. PI, *Topology-Aware MPI Communication and Scheduling for Petascale Systems*, in collaboration with K. Schulz and B. Barth (TACC) and Amitava Majumdar (SDSC), National Science Foundation, \$1,840,000, (OSU share: \$920,000), Oct 09–Sept 12.
70. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$117,812, April 09–March 10.
71. PI, *Extending One-Sided Communication in MPI Programming Model for Next-Generation Ultra-Scale HEC*, in collaboration with K. Schulz (TACC) and Amitava Majumdar (SDSC), National Science Foundation, \$797,000, (OSU share: \$399,000), Sept 08–Sept 10.

72. PI, *Design and Optimization of MVAPICH2 over QLogic InfiniBand Adapter*, QLogic, \$52,930, Sept. 08– Aug. 09.
73. PI, *Creating Petascale File Systems using Network Offloading*, DOE STTR Phase I Grant with RNet Technologies, \$33,000, June '08–Dec '08.
74. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$112,599, April 08–March 09.
75. PI, *Open Solaris pNFS-RDMA over InfiniBand*, Sun Microsystems, \$150,000, Jan. 08–Dec. 08.
76. PI, *Designing Next Generation Communication and I/O Subsystems with Multi-Core Architecture*, with P. Balaji (Co-PI) Univ. of Chicago, National Science Foundation, \$462,512, (OSU share: \$375,000), May 07–April 10.
77. PI, *Design and Optimization of MVAPICH over QLogic InfiniBand Adapter*, QLogic, \$50,000, Feb. 07–Jan 08.
78. Co-PI, *High Performance and Scalable Data-Centers with Multi-core Architectures and Emerging Networking Technologies*, jointly with Stu Zweben (PI), Wright Center for Innovation (WCI) on Advanced Data Management (ADMA) Equipment Grant, State of Ohio, \$600,000, Jan 07–June 09.
79. PI, *Accelerator for Off-loading Services of Next Generation Data-Centers*, STTR Phase I, National Science Foundation, \$74,999, Jan. 07–Dec. 07.
80. PI, *Open Solaris pNFS-RDMA over InfiniBand*, Sun Microsystems, \$150,000, Jan. 07–Dec. 07.
81. PI, *On-site Technical Assistance for Computer Cluster*, Central State University, \$5,000, Jan. 06–May 07.
82. PI, *Center for Programming Models for Scalable Parallel Computing*, part of the *Center for Programming Models for Scalable Parallel Computing* project with four other universities and two National Labs, Department of Energy, OSU Share: \$1,500,000, OSU Co-PI: P. Sadayappan, Sept. 06 – August 12.
83. PI, *Coordinated Fault Tolerance for High Performance Computing*, part of the *Coordinated Fault Tolerance for High Performance Computing* project with three other universities and three National Labs, Department of Energy, OSU Share: \$1,000,000, Sept. 06 – August 11.
84. PI, *Center for Performance Evaluation of Cluster Networking and I/O Technologies (PEC-NIT)*, DICE/AVETEC, \$749,996, July 06 – June 09.
85. PI, *Open Solaris NFS-RDMA over InfiniBand*, Sun Microsystems and Network Appliance, \$150,000, Dec. 05–Nov. 06.
86. PI, *High Performance Implementation of Cluster OpenMP over InfiniBand*, Intel, \$53,000, Nov 05–Oct 06.

87. PI, *Scalable, Fault-Tolerant and High Performance MPI over Multi-thousand node InfiniBand Clusters*, Cisco Systems, \$96,596, Dec. 05–Nov. 06.
88. PI, *Clustered and Distributed File Systems over InfiniBand and iWARP for PetaByte Storage*, Cisco Systems, \$96,596, Dec. 05–Nov. 06.
89. PI, *Efficient and Scalable Checkpoint Restart Scheme in MVAPICH*, Linux Networx, \$50,000, Sept. 05–Aug 06.
90. PI, *Multi-Gigabit OSbypass Systems for Grid Computing*, Department of Energy (SBIR Phase II), \$339,570, July 05–June 07.
91. PI, *Designing Next Generation Data-Centers with Advanced Communication Protocols and Systems Services*, National Science Foundation, \$150,000, July 05–June 07.
92. PI, *MVAPICH on Solaris 10*, SUN Microsystems, \$129,370, April 05–Dec 05.
93. PI, *High-End Computing and Networking Research Testbed for Next Generation Data Driven, Interactive Applications*, jointly with (Co-PIs and investigators) G. Agrawal, P. Sadayappan, J. Saltz, H.-W. Shen, S. Ahalt, U. Catalyurek, H. Ferhatosmanoglu, H.-W. Jin, T. Kurc, M. Lauria, D. Lee, R. Machiraju, S. Parthasarathy, P. Sinha, D. Stredney, A. E. Stutz, and P. Wyckoff, National Science Foundation, Total: \$3,014,063 (\$1,529,997 from NSF + \$1,484,066 from Ohio Board of Regents and OSU), Sept 04–Aug 09.
94. PI, *Advanced Message Passing Algorithms for RDMA-enabled Interconnects*, Intel, \$103,000, Sept 04–Aug 06.
95. PI, *Multi-Gigabit OSbypass Systems for Grid Computing*, Department of Energy (SBIR Phase I), \$33,000, July 04–April 05.
96. PI, *Research on High Performance and Scalable MPI over InfiniBand*, Mellanox Technologies, Inc., \$416,821, April 04–March 08.
97. PI, *Designing High Performance and Scalable Communication Subsystems for Next Generation Clusters with Infiniband Architecture*, National Science Foundation, \$150,000, August 03–July 05.
98. Co-PI, *Center for Grid-Enabled Medical Image Analysis*, jointly with J. Saltz (PI), M. Knopp, J. Zweier, C. Eng, A. Friedman, J. Au, A. Stutz, D. Stredney, P. Wyckoff, P. Kuppusamy, T. Kurc, U. Catalyurek, B. Clymer, R. Lee, and R. Machiraju, National Institute of Health, \$677,298, August 03–July 04.
99. PI, *Analysis and enhancements to MVAPICH wrt IBAL, SMP, connection management, and MPICH2*, Intel, \$25,000, Aug 03–Dec 03.
100. PI, *Efficient Implementation of LA-MPI on Myrinet/GM2 and Quadrics/Elan*, Los Alamos National Laboratory, \$149,991, Aug 03–Sept. 04.
101. PI, *Designing Smart High-Speed Network Interface Card*, Department of Energy (SBIR Phase I), \$33,000, Aug 03–April 04.

102. PI, *Collective Communication and Connection Management Issues in InfiniBand-Based Clusters*, jointly with Pete Wyckoff, Los Alamos National Laboratory, \$99,989, July 03–Sept. 03.
103. Co-PI, *High Performance Computing and Bio-medical Informatics at OSU*, jointly with M. Lauria (PI), Tahsin Kurc, and Daniel Janies, HP Equipment grant, \$288,054, June 03.
104. PI, *Supporting MPI Collective Communication Operations with Application Bypass*, jointly with P. Sadayappan and Pete Wyckoff, Sandia National Laboratory, \$85,000, July 03–June 04.
105. PI, *Network Interface Support for High Performance and Scalable Communication Services in Clusters*, National Science Foundation, \$260,601, July 02–June 05.
106. Co-PI, *Job Scheduling for Petaflop Systems*, jointly with P. Sadayappan (PI) and Pete Wyckoff, Sandia National Laboratory, \$145,118, July 02–Sept. 02.
107. PI, *High Performance and Scalable MPI Implementation on InfiniBand*, jointly with Pete Wyckoff, Sandia National Laboratory, \$201,340, May 02–Sept. 03.
108. PI, *Communication and Scheduling Issues in Buffered Co-Scheduling*, jointly with P. Sadayappan, Los Alamos National Laboratory, \$48,951, May 02–Sept. 02.
109. PI, *Communication Support for Programming Models for Scalable Parallel Computing*, part of the *Center for Programming Models for Scalable Parallel Computing* with six other Universities and four National Labs, Department of Energy, OSU Share: \$750,000, Sept. 01 - August 06.
110. PI, *Efficient Strategies and Implementations for One-sided Communication in Clusters*, Pacific Northwest National Laboratories, \$21,219, April 01–Sept. 2001.
111. Co-PI, *Evaluation of the Influence of Network Topology and Mapping on the Performance of Parallel Applications on Clusters*, jointly with P. Sadayappan (PI), J. Duato and P. Wyckoff. Sandia National Laboratories, \$129,724, April 01–March 02.
112. Co-PI, *Analysis of Message Passing Environments on Large Clusters, and Future Directions*, jointly with P. Wyckoff (PI), and P. Sadayappan, Sandia National Laboratories, \$150,016, Jan. 01–Dec 01.
113. Co-PI, *Low Latency Gigabit Ethernet Message Passing*, jointly with P. Wyckoff (PI), Sandia National Laboratories, \$274,028, Sept. 00–August 02.
114. PI, *Scalable Communication Support on SMP Clusters for Network-Based Computing*, jointly with P. Farrell (KSU), Ohio Board of Regents, \$85,104, July 00–June 02.
115. PI, *Network Computing Testbed for Interactive Visualization, Multimedia, and Metacomputing*, jointly with R. Crawfis, J. Duato, W-C. Feng, P. Sadayappan, D. Stredeny, and A. Stutz, National Science Foundation, (total \$350,000 = NSF \$175,000 + State of Ohio \$87,500 + University \$87,500), May 00–April 03.

116. PI, *Efficient Communication Support for Networks of Workstations Environments*, jointly with M. Banikazemi, IBM Research (Cooperative Fellowship, IBM Mentor - Dr. Rama Govindaraju), \$30,405, Oct. 99–June 00.
117. Co-PI, *Interactive Medical Data on Demand: A High-Performance Image-Based Warehouse Across Heterogeneous Environments*, jointly with D. Stredney (PI), R. Crawfis, W.-C. Feng, J. Hou, R. Miller, P. Sadayappan, and A. Stutz, National Library of Medicine, \$95,809, Sept. 1998 – June 1999.
118. PI, *Scalable Parallel Computing Over Geographically Distributed ATM-Interconnected Workstation Clusters*, Ameritech Faculty Fellowship Award, \$29,991, Oct. 1998 – March 2000.
119. PI, *Scalable Collective Communication on Network-Based Computing Systems*, jointly with F. Annexstein (UC), K. Berman (UC), and T. Page (OSU), Ohio Board of Regents Collaborative CS Research Fund, \$81,236, June 1998 – May 2000.
120. Co-PI, *Exploiting Multiple Interconnection Technologies in Networked PC Multicomputers*, jointly with T. Page (PI), P. Sadayappan, and N. Nagarajan (Systran Corporation), NSF Small Business Innovative Research (SBIR) Grant, \$100,000, Jan 1998 – Aug. 1998.
121. PI, *Scalable Collective Communication Support for Heterogeneous Networks of Workstations*. National Science Foundation, \$225,959, July 97 – Dec. 2000.
122. PI, *Collective Communication on Parallel Systems with Indirect Interconnects*, jointly with R. Sivaram, IBM Research (Cooperative Fellowship, IBM Mentor - Dr. Craig Stunkel), \$26,140, Oct. 97–June 98.
123. Co-PI, *An Imaging and Visualization System for Super High Field MR*, jointly with A. Abduljali, D. Stredney, A. E. Stutz, and R. Yagel, OSU Interdisciplinary Seed Grant, \$100,000, June 97–Dec. 98.
124. Co-PI, *vBNS Gateway for Ohio via OCARnet and OARnet*, jointly with C. Bender (PI), A. Stutz, G. Wallis, and R. Jain, National Science Foundation, \$350,000, Jan. 97–Dec. 98.
125. Co-PI, *OCARNet: Ohio Computing and Communications ATM Research Network*, jointly with R. Jain (PI) and other Ohio Universities/Institutions (OSC, OARnet, KSU, CSU, and UD), Ohio Board of Regents Investment Fund, (Total \$1,721,730, OSU share \$691,000), March 96–Feb. 99.
126. PI, *Collective Communication on Parallel Systems with Indirect Interconnects*, jointly with R. Sivaram, IBM Research (Cooperative Fellowship, IBM Mentor - Dr. Craig Stunkel), \$25,338, Oct. 96–June 97.
127. Co-PI, *ATM Testbed for Multimedia and Distributed Computing*, jointly with R. Jain (PI), A. Arora, and T. Page, National Science Foundation, (total \$400,000 = NSF \$175,300 + State of Ohio \$100,000 + University \$124,700), Feb. 96–Jan. 99.
128. PI, *Communication and Architectural Supports for Implementing Distributed Shared Memory on Wormhole Networks*, National Science Foundation Faculty Early CAREER Development Award, \$105,998, July 1995–June 98.

129. Co-PI, *Acquisition of a Mid-range Scalable Parallel Computer*, National Science Foundation Academic Research Infrastructure Grant, Co-PI, jointly with P. Sadayappan (PI), C.-H. Huang, D. N. Jayasimha, T.-H. Lai, D. Wang, R. Yagel, and F. Zhao, \$239,500, Sept 94–Aug. 97.
130. PI, *Communication and Synchronization in  $k$ -ary  $n$ -cube cluster- $c$  Scalable Systems*, National Science Foundation Research Initiation Award (RIA), \$99,549, July 93–June 96.
131. PI, *Development of Flexible Parallel Architectures: A Unified Platform to Support Shared Memory and Message Passing Paradigms*, The Ohio State University Seed Grant, \$15,000, Jan. 92–June 93.

#### Ph. D. Dissertations Supervised

1. Ching-Hsiang Chu, *Accelerator-enabled Communication Middleware for Large-scale Heterogeneous HPC Systems with Modern Interconnects*, July 2020.
2. Jahanzeb Hashmi, *Designing High Performance Shared-Address-Space and Adaptive Communication Middlewares for Next-Generation HPC Systems*, April 2020.
3. Ammar Awan, *Co-designing Communication Middleware and Deep Learning Frameworks for High-Performance DNN Training on HPC Systems*, April 2020.
4. Dipti Shankar, *Designing Fast, Resilient and Heterogeneity-Aware Key-Value Storage for Modern HPC Clusters*, July 2019.
5. Sourav Chakraborty, *High Performance and Scalable Cooperative Communication Middleware for Next Generation Architectures*, June 2019.
6. Jie Zhang, *Designing and Building Efficient HPC Cloud with Modern Networking Technologies on Heterogeneous HPC Clusters*, July 2018.
7. Mingzhe Li, *Designing High-Performance Remote Memory Access for MPI and PGAS Models with Modern Networking Technologies on Heterogeneous Clusters*, December 2017.
8. Akshay Venkatesh, *High-Performance Heterogeneity/Energy-Aware Communication for Multi-Petaflop HPC Systems*, December 2016.
9. Nusrat Islam, *High-Performance File System and I/O Middleware Design for Big Data on HPC Clusters*, December 2016.
10. Wasiur Rahman, *Designing and Modeling High-Performance MapReduce and DAG Execution Framework on Modern HPC Systems*, December 2016.
11. Raghunath Rajachandrasekar, *Designing Scalable And Efficient I/O Middleware for Fault-Resilient High-performance Computing Clusters*, November 2014.
12. Jithin Jose, *Designing High Performance and Scalable Unified Communication Runtime (UCR) for HPC and Big Data Middleware*, August 2014.
13. Sreeram Potluri, *Enabling Efficient Use of MPI and PGAS Programming Models on Heterogeneous Clusters with High Performance Interconnects*, May 2014.

14. Miao Luo, *Designing Efficient MPI and UPC Runtime for Multicore Clusters with InfiniBand and Heterogeneous System*, July 2013.
15. Krishna Kandalla, *High Performance Non-Blocking Collective Communication for Next Generation InfiniBand Clusters*, July 2013.
16. Hari Subramoni, *Topology-Aware MPI communication and Scheduling for High Performance Computing Systems*, July 2013.
17. Xiangyong Ouyang, *Efficient Storage Middleware Design in InfiniBand Clusters for High-End Computing*, March 2012.
18. Matthew Koop, *High-Performance Multi-Transport MPI Design for Ultra-Scale InfiniBand Clusters*, June 2009.
19. Gopal Santhanaraman, *Designing Scalable and High Performance One-Sided Communication Middleware for Modern Interconnects*, June 2009.
20. Lei Chai, *High Performance and Scalable MPI Intra-Node Communication Middleware for Multi-core Clusters*, February 2009.
21. Wei Huang, *High Performance Network I/O in Virtual Machines over Modern Interconnects*, August 2008.
22. Ranjit Noronha, *Designing High Performance and Scalable Clustered Network Attached Storage with InfiniBand*, August 2008.
23. Sundeepp Narravul, *Designing High Performance and Scalable Distributed Datacenter Services over Modern Interconnects*, August 2008,
24. Amith Mamidala, *Scalable and High Performance MPI Collective Communication over Multicore InfiniBand Clusters*, May 2008.
25. Karthikeyan Vaidyanathan, *High Performance and Scalable Soft Shared State for Next-Generation Datacenters*, May 2008.
26. Abhinav Vishnu, *High Performance and Network Fault-Tolerant MPI with Multi-pathing over InfiniBand*, November 2007.
27. Sayantan Sur, *Scalable and High Performance MPI Design for Very Large InfiniBand Clusters*, August 2007.
28. Pavan Balaji, *High Performance Communication Support for Sockets-Based Applications over High-Speed Networks*, June 2006.
29. Weikuan Yu, *Enhancing MPI with Modern Networking Mechanisms in Cluster Interconnects*, June 2006.
30. Jiasheng Wu, *Communication and Memory Management in Networked Storage Systems*, Sept 2004.
31. Jiuxing Liu, *Designing High Performance and Scalable MPI over InfiniBand*, Sept 2004.
32. Darius Buntinas, *Improving Cluster Performance through the Use of Programmable Network Interfaces*, June 2003.
33. Mohammad Banikazemi, *Design and Implementation of High Performance Communication Subsystems for Clusters*, December 2000.



34. Donglai Dai, *Designing Efficient Communication Subsystems for Distributed Shared Memory (DSM) Systems*, February 1999.
35. Ram Kesavan, *Communication Mechanisms and Algorithms for Supporting Collective Communication on Parallel Systems*, September 1998.
36. Rajeev Sivaram, *Architectural Support for Efficient Communication in Scalable Parallel Systems*, Co-adviser: Dr. Craig Stunkel (IBM T.J. Watson Research center), August 1998.
37. Debashis Basak, *Designing High Performance Parallel Systems: A Processor-Cluster Based Approach*, July 1996.
38. Vibha A. Dixit-Radiya, *Mapping on Wormhole-routed Distributed-Memory Systems: A Temporal Communication Graph-based Approach*, Mar 1995.

### M. S. Thesis Supervised

1. Kamal Raj Sankarapandian, *Profiling MPI Primitives in Real-time Using OSU INAM*, April 2020.
2. Rajarshi Biswas, *Benchmarking and Accelerating TensorFlow-based Deep Learning on Modern HPC Systems*, July 2018.
3. Kunal Kulkarni, *Performance Characterization and Improvements of SQL-on-Hadoop Systems*, August 2016.
4. Albert Augustine, *Designing a Scalable Network Analysis and Monitoring Tool with MPI Support*, July 2016.
5. Adithya Bhat, *RDMA-based Plugin Design and Profiler for Apache and Enterprise Hadoop Distributed Filesystem*, August 2015.
6. Vijay Dhanraj, *Enhancement of LIMIC-Based Collectives for Multi-core Clusters*, August 2012.
7. Ashish Singh, *Optimizing All-to-all and Allgather Communications on GPGPU Clusters*, April 2012.
8. Siddesh Pai Raikar, *Network Fault-Resilient MPI for Multi-Rail InfiniBand Clusters*, December 2011.
9. Nishant Dandapanthula, *InfiniBand Network Analysis and Monitoring using OpenSM*, August 2011.
10. Vilobh Meshram, *Distributed Metadata Management for Parallel Systems*, August 2011.
11. Greg Marsh, *Evaluation of High Performance Financial Messaging on Modern Multi-core Systems*, March 2010.
12. Karthik Gopalakrishnan, *Enhancing Fault Tolerance in MPI for Modern InfiniBand Clusters*, August 2009.
13. Jaidev Sridhar, *Scalable Job Startup and Inter-Node Communication in Multi-Core InfiniBand Clusters*, June 2009.

14. Tejus Gangadharappa, *Designing Support for MPI-2 Programming Interfaces on Modern Interconnects*, June 2009.
15. Rahul Kumar, *Enhancing MPI Point-to-point and Collectives for Clusters with Onloaded/Offloaded InfiniBand Adapters*, August 2008.
16. Sitha Bhagvat, *Designing and Enhancing the Sockets Direct Protocol (SDP) over iWARP and InfiniBand*, August 2006.
17. Weihang Jiang, *High Performance MPICH2 One-Sided Communication Implementation over InfiniBand*, June 2004.
18. Savitha Krishnamoorthy, *Dynamic Reconfigurability Support to Provide Soft QoS Guarantees in Cluster-Based Multi-Tier Data-Centers over InfiniBand*, June 2004.
19. Adam Wagner, *Static and Dynamic Processing Offload on Myrinet Clusters with Programmable NIC Support*, June 2004.
20. Adam Moody, *NIC-based Reduction on Large-Scale Quadrics Cluster*, September 2003.
21. Balasubramanian Chandrasekharan, *Micro-benchmark Level Performance Evaluation and Comparison of High Speed Cluster Interconnects*, August 2003.
22. Sushmita Kini, *Efficient Collective Communication using Multicast and RDMA Operations for InfiniBand-based Clusters*, June 2003.
23. Sandhya Senapathi, *QoS-Aware Middleware to Support Interactive and Resource Adaptive Applications on Myrinet Clusters*, Sept. 2002.
24. Rinku Gupta, *Efficient Collective Communication using Remote Memory Operations on VIA-Based Clusters*, Aug 2002.
25. Piyush Shivam, *High Performance User Level Protocol on Gigabit Ethernet*, Aug 2002.
26. Amina Saify, *Optimizing Collective Communication Operations in ARMCI*, June 2002.
27. Shreyas Desai, *Mechanisms for Implementing Efficient Collective Communication in Clusters with Application Bypass*, June 2002.
28. Vinod Tipparaju, *Optimizing ARMCI Get and Put Operations on Myrinet/GM*, September 2001.
29. Abhisekh Gulati, *A Proportional Bandwidth Allocation Scheme for Myrinet Clusters*, June 2001.
30. Vijay Kota, *Designing Efficient Inter-Cluster Communication Layer for Distributed Computing*, June 2001.
31. Sencer Kutlug, *Performance Evaluation and Analysis of User Level Networking Protocols in Clusters*, June 2000.

## Software Design, Development, and Distribution

Dr. Panda and his students are involved in designing a high performance and scalable MPI (Message Passing Interface standard) for clusters with the emerging InfiniBand, Omni-Path, iWARP, EFA, and RDMA over Converged Ethernet (RoCE) networking technologies. Currently, there are two versions of this MPI: The MVAPICH *open-source* software was first demonstrated

at Supercomputing (SC '02) and after that it has been steadily gaining acceptance in the HPC, Cluster, Networking, and InfiniBand communities. More than 3,100 organizations (National Labs, Universities, and Industry) in 89 countries have downloaded this software from OSU's web site directly together with voluntary registration. As of July '20, more than 798,000 downloads have taken place from the OSU Web site. In addition, many server vendors, networking vendors (InfiniBand, Omni-Path, iWARP, EFA, 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, SuSE, Spack and OpenHPC). Several InfiniBand systems using MVAPICH2 have obtained positions in the TOP500 ranking of supercomputers in the world ([www.top500.org](http://www.top500.org)) during the last decade. The June '20 list includes the following systems: 4th ranked Sunway TaihuLight system in China with 10,649,600-cores, 8th ranked 448,448 cores (Frontera) at TACC, 12th ranked 391,680 cores (ABCI) in Japan, 18th ranked 570,020 cores (Neurion) in South Korea, 19th ranked 556,104 cores (Oakforest-PACS) in Japan, 22nd ranked 367,024 cores (Stampede2) at TACC, and 40th ranked 241,108-core (Pleiades) at NASA. 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.ohio-state.edu/>

Dr. Panda and his students are involved in designing high performance and scalable versions of Hadoop, Spark and Memcached for Big Data processing to exploit RDMA technology like InfiniBand and RoCE on modern clusters. The RDMA-enabled Apache version of Hadoop (including RDMA-based designs for multiple components of Hadoop including HDFS, MapReduce and RPC), RDMA-Spark, RDMA-HBase, and RDMA-enabled Memcached version are publicly available from <http://hibd.cse.ohio-state.edu>. As of July '20, more than 330 organizations from 36 countries are using these software libraries. More than 37,000 downloads of this software have taken place from the project website alone.

The group has also been focusing on co-designing Deep Learning Frameworks. High-performance and scalable versions of the Caffe framework and Tensorflow-gRPC framework are available from <http://hidl.cse.ohio-state.edu>. As of July '20, more than 1,200 downloads of this software have taken place and it is being used by more than 70 organizations.

## Professional Activities

- **General Chair/Program Chair/Vice Chair:**

1. Program Chair, Supercomputing Asia (SCAsia), 2020
2. General Vice Chair, Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2020
3. Vice Chair (Architecture, Networks and Infrastructure), Int'l Supercomputing Conference (ISC), 2020
4. General Co-Chair, IEEE Micro, 2019
5. Program Chair, Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2018
6. Program Co-Chair, Int'l Workshop on High-Performance Big Data Computing (HPBDC), 2019, 2018, 2017, 2016 and 2015

7. Vice Co-Chair (Architecture and Networking), Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2016
8. Program Co-Chair, Int'l Workshop on Extreme Scale Programming Models and Middleware (ESPM2), 2020, 2018, 2017, 2016 and 2015
9. Program Chair, 9th Int'l Conference on Partitioned Global Address Programming Models (PGAS), 2015
10. Program Co-Chair, Int'l Workshop on Communication Architectures at Extreme Scale (ExaComm), 2020, 2019, 2018, 2017, 2016 and 2015
11. Program Chair, Int'l Symposium on High Performance Computing (HiPC), 2012
12. Vice Chair (Architecture), Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2012
13. Vice Chair (Communications and Networks), Int'l Symposium on High Performance Computing (HiPC), 2011
14. Vice Chair, Int'l Conference on Cluster Computing (Cluster), 2010.
15. Program Co-Chair (with Dimitri Stiliadis), ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS '08), 2008.
16. Technical Papers Co-Chair (with Darren Kerbyson), Int'l Conference on Supercomputing (SC '08), 2008.
17. Working Group Co-Chair on Network Interface (with Keith Underwood), Interconnection Networks Workshop, Department of Energy (DOE), July 2008.
18. Vice Chair (Communications and Networks), Int'l Symposium on High Performance Computing (HiPC), 2007
19. Program Co-Chair (with Ron Brightwell), IEEE Symposium on Hot Interconnects (HotI15), 2007
20. Program Chair, International Parallel and Distributed Processing Symposium (IPDPS), 2007
21. Program Co-Chair (with Dr. Tal Lavian), IEEE Symposium on Hot Interconnects (HotI14), 2006
22. General Chair, Int'l Conference on Parallel Processing (ICPP '06)
23. Program Co-Chair, Int'l Workshop on High Performance Interconnect for Distributed Computing (HPI-DC '05)
24. General Co-Chair, Int'l Conference on Information Technology (CIT '05)
25. Program Co-Chair (with Prof. Jose Duato and Dr. Craig Stunkel), Workshop on Communication Architecture for Clusters (CAC '04)
26. Program Co-Chair (with Prof. Jose Duato and Dr. Craig Stunkel), Workshop on Communication Architecture for Clusters (CAC '03)
27. Program Co-Chair (with Dr. Jarek Nielplocha), DOE Workshop on Communication and Middleware for Parallel Programming Models (CMPPM '02)
28. Program Co-Chair (with Prof. Jose Duato and Dr. Craig Stunkel), Workshop on Communication Architecture for Clusters (CAC '02)

29. General Co-Chair (with Prof. Jose Duato), International Conference on Parallel Processing (ICPP '01)
30. Program Co-Chair (with Prof. Jose Duato), Workshop on Communication Architecture for Clusters (CAC '01)
31. Program Co-Chair (with Prof. Norio Shiratori), International Conference on Parallel Processing (ICPP '99)
32. Vice Chair (Architecture Track), International Conference on Parallel Processing (ICPP '98)
33. Program Co-Chair (with Dr. Craig Stunkel), Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing (CANPC '98)
34. Program Co-Chair (with Dr. Craig Stunkel), Workshop on Communication and Architectural Support for Network-based Parallel Computing (CANPC '97)

• **Program Committee Member:**

1. Int'l Conference on Cluster Computing (Cluster '21)
2. Int'l Symposium on Workload Characterization(IISWC '20)
3. Int'l Conference on Supercomputing (ICS '20)
4. Int'l Conference on Distributed Computing Symposium (ICDCS '20)
5. Int'l Conference on Supercomputing (SC '20)
6. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '20)
7. Int'l TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC '19)
8. Int'l Symposium on Computer Architecture and High Performance Computing (SBAC-PAD 2019)
9. Int'l Workshop on High Performance Interconnection Networks and Interconnects (HIP-INI '19)
10. Int'l Conference on Supercomputing (ICS '19)
11. ACM Int'l Conference on Computing Frontiers (CF '19)
12. Int'l Conference on Computer Design (ICCD '19)
13. Int'l Supercomputing Conference, Research Paper Track (ISC '19)
14. Int'l Conference on Supercomputing (SC '19)
15. Int'l Supercomputing Conference, Tutorial Track (ISC '19)
16. OpenSHMEM 2018
17. Int'l Workshop on Kernel-Bypass Networks (KBNets '18)
18. Int'l Workshop on the Convergence of Extreme Scale Computing and Big Data Analysis (CEBDA '18)
19. EuroMPI/USA 2018
20. ACM Int'l Conference on Computing Frontiers (CF '18)

21. Int'l Conference on Supercomputing, Poster Track (SC '18)
22. Int'l Conference on Supercomputing (ICS '18)
23. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '18)
24. OpenSHMEM 2017
25. Int'l Workshop on Advances in Software and Hardware for Big Data to Knowledge Discovery (ASH '17)
26. Int'l Conference on Parallel Architectures and Compilation Techniques (PACT '17)
27. IEEE Int'l Conference on Networking, Architecture, and Storage (NAS '17)
28. Int'l Conference on Supercomputing (SC '17)
29. ACM Int'l Conference on Computing Frontiers (CF '17)
30. Int'l Conference on Supercomputing (ICS '17)
31. Int'l Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB '17)
32. EuroMPI/USA 2017
33. Int'l Workshop on Communication Optimizations in High Performance Computing (COMHPC '16)
34. Int'l Workshop on Irregular Applications: Architectures and Algorithms (IA<sup>3</sup> '16)
35. OpenSHMEM 2016
36. Int'l Workshop on Advances in Software and Hardware for Big Data to Knowledge Discovery (ASH '16)
37. Int'l TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC '16)
38. ACM Int'l Conference on Computing Frontiers (CF '16)
39. Int'l Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB '16)
40. Int'l Workshop on Performance Analysis of Big Data Systems (PABS '16)
41. Int'l Conference on Distributed Computing Symposium (ICDCS '16)
42. Int'l Conference on Cluster Computing (Cluster '16)
43. Int'l Supercomputing Conference (ISC '16)
44. Int'l Workshop on Advances in Software and Hardware for Big Data to Knowledge Discovery (ASH '15)
45. Int'l TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC '15)
46. Int'l Workshop on Software-Defined Ecosystems (BigSystem '15)
47. Int'l Conference on Supercomputing, Emerging Technologies Track (SC '15)
48. Int'l Conference on Distributed Computing Symposium (ICDCS '15)
49. Int'l Workshop on Performance Analysis of Big Data Systems (PABS '15)

50. Int'l Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB '15)
51. Int'l Conference on Cluster Computing (Cluster '15)
52. Int'l TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC '14)
53. Int'l Conference on Cluster Computing (Cluster '14)
54. Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid '14)
55. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '14)
56. Int'l Symposium on Architectures for Networking and Communications Systems (ANCS '13)
57. Int'l Workshop on Benchmarks, Performance Optimization, and Emerging hardware of Big Data Systems and Applications (BPOE '13)
58. Int'l Workshop on BigData Benchmarking, China (WBDB '13)
59. Int'l Conference on Parallel and Distributed Systems (ICPADS '13)
60. IEEE CloudNet (CloudNet'13)
61. Int'l Workshop on Virtualization Technologies in Distributed Computing (VTDC-2013)
62. Int'l Conference on Supercomputing (SC '13)
63. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '13)
64. Int'l Workshop on BigData Benchmarking (WBDB '12)
65. Int'l Workshop on Network-Aware Data Management (NDM '12)
66. Int'l Conference on Partitioned Global Address Space Programming Model (PGAS '12)
67. TACC-Intel Highly Parallel Computing Symposium (TI-HPCS '12)
68. Int'l Workshop on Resiliency in High Performance Computing (Resiliency '12)
69. Int'l Workshop on Virtualization Technologies in Distributed Computing (VTDC-2012)
70. Int'l Workshop on Advanced Data Management Systems (ADMS '12)
71. Int'l Conference on e-Science (eScience '12)
72. Int'l Workshop on In-Situ Data Processing (ISDP '12)
73. Int'l Conference on Cloud Computing and Services Science (CLOSER '12)
74. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '12)
75. Int'l Workshop on Network-Aware Data Management (NDM '11)
76. Int'l Conference on Supercomputing (SC '11)
77. Int'l Workshop on Resiliency in High Performance Computing (Resilience 2011)
78. Int'l Conference on Cluster Computing (Cluster '11)
79. IEEE Infocom Workshop on Cloud Computing (CC '11)
80. Int'l Workshop on Data Center Converged And Virtual Ethernet Switching (DC CAVES '11)
81. Int'l Workshop on Virtualization Technologies in Distributed Computing (VTDC-2011)

82. Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid '11)
83. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '11)
84. Int'l Conference on Parallel and Distributed Systems (ICPADS '10)
85. Int'l Workshop on Virtualization Performance: Analysis, Characterization, and Tools (VPACT '10)
86. Int'l Conference on Computer Communication Networks (ICCCN '10)
87. Int'l Workshop on Virtualization Technologies in Distributed Computing (VTDC-2010)
88. Int'l Workshop on System-level Virtualization for High Performance Computing (HPCVirt 2010)
89. Int'l Conference on Supercomputing (SC '10)
90. Int'l Workshop on Resiliency in High Performance Computing (Resilience 2010)
91. Int'l Workshop on Data Center Converged And Virtual Ethernet Switching (DC CAVES '10)
92. Int'l Conference on Parallel Processing (ICPP '10)
93. Int'l Workshop on Storage Network Architecture and Parallel I/O (SNAPI '10)
94. Int'l Workshop on Grid Computing, Applications and Technologies (GridCAT '10)
95. Int'l Conference on Cluster Computing (Cluster '09)
96. Int'l Symposium on Architectures for Networking and Communications Systems (ANCS '09)
97. Int'l Conference on Parallel and Distributed Systems (ICPADS '09)
98. Int'l Workshop on High Performance Interconnects for Distributed Computing (HPI-DC '09)
99. Int'l Conference on Parallel Processing (ICPP '09)
100. Int'l Workshop on Data Center Converged And Virtual Ethernet Switching (DC CAVES '09)
101. Int'l Workshop on System-level Virtualization for High Performance Computing (HPCVirt 2009)
102. Int'l Conference on High Performance Computing (HiPC '09)
103. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '09)
104. Int'l Symposium on High Performance Computer Architecture (HPCA-15, 2009)
105. Int'l Symposium on Computer Architecture and High Performance Computing (SBAC-PAD '08)
106. Int'l Workshop on Storage Network Architecture and Parallel I/O (SNAPI '08)
107. Int'l Workshop on System-level Virtualization for High Performance Computing (HPCVirt 2008)
108. Int'l Conference on Electro/Information Technology (EIT 2008)
109. Int'l Conference on Cluster Computing (Cluster '08)



110. Int'l Symposium on Parallel and Distributed Processing and Applications (ISPA '08)
111. Int'l Conference on Parallel Processing (ICPP '08)
112. Int'l Conference on Cluster Computing (Cluster '07)
113. Int'l Workshop on Storage Network Architecture and Parallel I/O (SNAPI '07)
114. Int'l Symposium on Architectures for Networking and Communications Systems (ANCS '07)
115. Int'l Workshop on System-level Virtualization for High Performance Computing (HPCVirt 2007)
116. Int'l Conference on Supercomputing (SC '07)
117. Int'l Conference on Computer Communications and Networks (ICCCN '07)
118. Int'l Workshop on RDMA Applications, Implementations, and Technologies (RAIT '06)
119. Int'l Symposium on Architectures for Networking and Communications Systems (ANCS '06)
120. Int'l Conference on Supercomputing (SC '06)
121. Int'l Symposium on High Performance Computer Architecture (HPCA-12, 2006)
122. Int'l Conference on Cluster Computing (Cluster '05)
123. Int'l Workshop on RDMA Applications, Implementations, and Technologies (RAIT '05)
124. Int'l Conference on High Performance Computing (HiPC '05)
125. Int'l Conference on Supercomputing (SC '05)
126. Int'l Conference on High Performance Distributed Computing (HPDC '05)
127. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '05)
128. Int'l Symposium on Performance Analysis of Systems and Software (ISPASS '05)
129. Int'l Workshop on Storage Network Architecture and Parallel I/Os (SNAPI '04)
130. Int'l Workshop on RDMA Applications, Implementations, and Technologies (RAIT '04)
131. Int'l Conference on Cluster Computing (Cluster '04)
132. Int'l Conference on Supercomputing (SC '04)
133. Int'l Workshop on High-Speed Local Networks (HSLN '04)
134. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '04)
135. Int'l Symposium on High Performance Computer Architecture (HPCA-10, 2004)
136. Int'l Conference on Cluster Computing (Cluster '03)
137. Int'l Workshop on High-Speed Local Networks (HSLN '03)
138. Int'l Conference on Parallel Processing (ICPP '03)
139. Int'l Conference on Supercomputing (SC '03)
140. Int'l Conference on High Performance Computing (HiPC '03)
141. Int'l Conference on Parallel and Distributed Processing Symposium (IPDPS '03)
142. Int'l Conference on Information Technology (CIT '02)

143. Int'l Workshop on High-Speed Local Networks (HSLN '02)
144. Int'l Conference on High Performance Computing (HiPC '02)
145. Int'l Conference on Parallel Processing (ICPP '02)
146. Int'l Conference on Information Technology (CIT '01)
147. Int'l Conference on Parallel Processing (ICPP '01)
148. Int'l Symposium on Parallel Architectures, Algorithms and Networks (I-SPAN '00)
149. Int'l Conference on Parallel Processing (ICPP '00)
150. Int'l Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing (CANPC '00)
151. Int'l Workshop on Heterogeneous Computing (HCW '00)
152. Int'l Conference on High Performance Computing (HiPC '99)
153. Int'l Conference on Parallel and Distributed Computing Systems (PDCS '99)
154. Int'l Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing (CANPC '99)
155. Int'l Conference on Computer Communications and Networks (IC3N '98)
156. Int'l Conference on Information Technology (CIT' 98)
157. Int'l Conference on High Performance Computing (HiPC '98)
158. Int'l Workshop on Parallel Computer, Routing, and Communication (PCRCW '97)
159. Int'l Symposium on High Performance Computer Architecture (HPCA-4, 1997)
160. Int'l Conference on Computer Communications and Networks (IC3N '97)
161. Int'l Conference on Massively Parallel Processing with Optical Interconnects (MPPOI '97)
162. Int'l Conference on Parallel Processing (ICPP '96)
163. Int'l Conference on Distributed Computing Systems (ICDCS-16, 1996)
164. Int'l Parallel Processing Symposium (IPPS '96)
165. Int'l Conference on High Performance Computing (HiPC '95)
166. Int'l Parallel Processing Symposium (IPPS '95)
167. Int'l Workshop on High-Speed Network Computing (HiNet '95)

• **Panel Moderator:**

1. *Current and Future State of HPC Interconnects*, Linux Networx User Group meeting, Sept. 13, 2005.
2. *Subnet Management Features and MPI: How to Integrate them Together*, BOF, OpenIB Workshop, Feb 2005.
3. *The Emergence of Workstation Clusters: Should we Continue to Build MPPs?*, Int'l Symposium on High Performance Computer Architecture (HPCA-4), Feb 2, 1998.

4. *Designing High-Performance Communication Subsystems: Top Five Problems to Solve and Five Problems Not to Solve During the Next Five Years*, Parallel Computing, Routing, and Communication Workshop (PCRCW'97), Atlanta, June 26, 1997.

• **Panel Member:**

1. *Scaling Deep Learning on ORNL Summit*, Industry Panel at HiPC '19 Conference, December 18, 2019.
2. *HPC and Cloud Convergence. What about HPC and Edge?*, Panel at HPBDC Workshop, in conjunction with IPDPS '19, May 20, 2019.
3. *Heterogeneity in Computing: Now and in the Future*, Panel at HCW Workshop, in conjunction with IPDPS '19, May 20, 2019.
4. *New Generation Cloud Computing: Challenges and Future Directions*, Panel at CCGrid '19, May 16, 2019.
5. *The Next Wave of HPC in the Datacenter*, Panel at Supercomputing '18, November 15, 2018.
6. *High-Performance Training for Deep Learning and Computer Vision HPC*, Panel at the ECV Workshop (in conjunction with CVPR '18 conference), June 22, 2018.
7. *Which Framework is the Best for High-Performance Deep Learning: Big Data Framework or HPC Framework?*, Panel at HPBDC '18 Workshop (in conjunction with IPDPS '18 conference), May 21, 2018.
8. *Combining Software and Data Programs*, NSF SI2 PI Meeting Panel, May 1, 2018.
9. *Architectural Convergence of Big Data and Extreme-Scale Computing: Marriage of Convenience or Conviction*, Charm++ Workshop, April 12, 2018.
10. *Sunrise or Sunset: Exploring the Design Space of Big Data Software Stack*, Panel at HPBDC '17 Workshop (in conjunction with IPDPS '17 conference), June 2, 2017.
11. *Exascale Era*, Panel at HPC Advisory Council Stanford Conference, April 10, 2017.
12. *The Exascale Endeavor*, Panel at HPC Advisory Council Stanford Conference, February 7, 2017.
13. *Merge or Split: Mutual Influence between Big Data and HPC Techniques*, Panel at HPBDC '16 Workshop (in conjunction with IPDPS '16 conference), May 27, 2016.
14. *The Road to Exascale*, Panel at HPC Advisory Council Lugano Conference, March 22, 2016.
15. *The Road to Exascale*, Panel at HPC Advisory Council Stanford Conference, February 24, 2016.
16. *How Can we Dramatically Increase Network Scalability?*, Panel at HiPINEB '15 Workshop (held in conjunction with Cluster '15 conference), September 8, 2015.
17. *Intelligent Networking with SDN: Current State of the Art, Emerging Technologies and Future Visions*, Panel at DOE Intelligent Optical Networking Infrastructure (IONI) Workshop, Aug 5, 2014.

18. *Challenges for Exascale Systems*, Panel at HPC Advisory Council European Conference, June 22, 2014.
19. *The Road to Exascale*, Panel at HPC Advisory Council Stanford Conference, February 3, 2014.
20. *The Path to Exascale*, Panel at HPC Advisory Council China Workshop, October 28, 2013.
21. *Needs for Improvement in Extreme Scaling Systems with Heterogeneity*, Panel at Extreme Scaling Workshop, August 15, 2013.
22. *Scaling with PGAS Languages*, Panel at OpenFabrics International Developers Workshop, April 22, 2013.
23. *Do's and Don'ts of Managing a Software Project*, Panel at NSF SI2 PI Meeting, January 18, 2013.
24. *Challenges and Issues of Scaling both High Throughput and High Performance Computing in Clusters*, Panel at IEEE Cluster Conference (Cluster), September 26, 2012.
25. *Will Exascale Computing Really Require New Algorithms and Programming Models?*, Panel at Int'l Parallel and Distributed Processing Symposium (IPDPS), May 23, 2012.
26. *Data Management in Exascale Computing and Terabit Networking Era*, Panel at Network-Aware Data Management Workshop, held in conjunction with SC '11 Conference, November 14, 2011.
27. *Interconnects: Is Exascale End-of-the-line for Commodity Networks?* Panel at Cluster '11 Conference, September 27, 2011.
28. *In the Multi-core age, How do Larger, Faster, Cheaper and more Responsive Memory Sub-systems Affect Data Management?* Panel at Workshop on Accelerating Data Management Systems Using Modern Processors and Storage Architectures (ADMS '11), in conjunction with VLDB '11, September 2, 2011.
29. *Programming Models at Exascale: Are we ready for the Challenges?*, Panel at Future Approaches to Data Centric Programming for Exascale (DCPM) Workshop, in conjunction with IPDPS '11, May 20, 2011.
30. *Virtualization for HPC*, Disruptive Technology Panel, Supercomputing '10 Conference, Nov. 16, 2010.
31. *Cluster Applications in the Petascale Decade*, Cluster '09 Conference, Sept. 1, 2009.
32. *Applications and Tools for Petascale Systems*, Petascale Summer Workshop, June 2008.
33. *Ethernet: Convergence or Schism*, Int'l Workshop on Communication Architecture for Clusters (CAC), in conjunction with IPDPS '08.
34. *The Role of Accelerators in Cluster Communication*, Int'l Workshop on Communication Architecture for Clusters (CAC), in conjunction with IPDPS '07.
35. *Data Intensive Computing*, Supercomputing 2006, Nov. 16, 2006.
36. *Challenges for Future Interconnection Networks: Power, Reliability and Performance Scalability?*, Int'l Symposium on Hot Interconnect (HotI), August 24, 2006.

37. *How will we develop and program emerging robust, low-power, adaptive multicore computing systems?*, Int'l Conference on Parallel and Distributed Computing Systems (ICPADS), July 12, 2006.
38. *Is InfiniBand Ready for Production?*, ARL Cluster Symposium, July 27, 2005.
39. *Student Pipelining Issues and Challenges*, NSF RI PI Workshop, July 25, 2005.
40. *The Future (of) RDMA*, High Performance Interconnects for Distributed Computing (HPI-DC) Workshop, July 24, 2005.
41. *How would you architect a 100,000-processor cluster?*, Communication for Clusters (CAC) Workshop, April 4, 2005.
42. *What happened to the I/O network?*, IEEE Workshop on High Speed Local Networks (HSLN '04), Nov. 16, 2004.
43. *Trends in Designing High Performance Computing*, Intel Developers Forum, Feb. 18, 2004.
44. *Top Problems in Cluster Computing and Systems and Possible Solutions*, Cluster 2003, Dec. 3, 2003.
45. *Battle of the Network Stars*, Supercomputing 2003, Nov. 21, 2003.
46. *Emerging Hardware Issues for Building Future Clusters*, Cluster 2001, Oct. 9, 2001.
47. *Scalable Computing Infrastructure for Scientific and Commercial Applications*, Int'l Conference on Parallel Processing, ICPP '00, Aug. 23, 2000.
48. *Future Computing Requirements for NASA*, Scientific Computing Workshop, NASA Glenn Research Center, Cleveland, Apr. 13, 2000.

- **BoF Organizer:**

1. *Building Efficient Clouds for HPC, Big Data, and Deep Learning Middleware and Applications*, ISC '18, June 27, 2018.
2. *Building Efficient Clouds for HPC, Big Data, and Deep Learning Middleware and Applications*, Supercomputing '17, Nov. 15, 2017.
3. *Accelerating Big Data Processing and Machine/Deep Learning Middleware on Modern HPC Clusters*, Supercomputing '17, Nov. 14, 2017.
4. *Accelerating Big Data Processing System Software on Modern HPC Clusters*, ISC '17, June 20, 2017.

- **Editorship:**

1. Co-Editor-in-Chief, *CCF Transactions on High-Performance Computing*, March 2018 – current.
2. Associate Editor, *IEEE Transactions on Computers*, January 2010–2014.
3. Subject Area Editor, *Journal of Parallel and Distributed Computing*, March 2006–2016.
4. Guest Editor, Special issue of *Journal of Parallel and Distributed Computing* with Best Papers from IPDPS '07 conference, August 2009.

5. Guest Editor (with Prof. Jose Duato), Special issue on *Communication Architecture for Clusters*, Cluster Computing Journal, 2003.
6. Associate Editor, *IEEE Transactions on Parallel and Distributed Systems*, 1998–2001.
7. Guest Editor (with Prof. Lionel Ni), Two special issues on *Workstation Clusters and Network-based Computing*, *Journal of Parallel and Distributed Computing*, one published in January 1997 and the second one in September 1997.
8. Guest Editor, Special issue on *Interconnection Networks for High Performance Computing Systems* of IEEE Technical Committee on Computer Architecture (TCCA) Newsletter, Fall 1994 and Winter 1995.
9. Co-editor, IEEE Technical Committee on Parallel Processing (TCPP) Newsletter, 1992–95. Have been awarded a **Meritorious Service Certificate**, IEEE Computer Society, 1996, for contributions as a co-editor.

• **Organizational:**

1. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2018.
2. Session Chair, Int'l Conference on Cluster Computing and Grid (CCGrid), May 15, 2017.
3. Session Chair, Int'l Conference on Supercomputing (ICS), June 1, 2016.
4. Session Chair, Int'l Conference on Cluster Computing and Grid (CCGrid), May 17, 2016.
5. Session Chair, South African High-Performance Computing Conference (CHPC '15), Dec 4, 2015.
6. Session Chair, Int'l Conference on Cluster Computing (Cluster), September 10, 2015.
7. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2014.
8. Session Chair, Int'l Conference on Supercomputing (SC), 2013.
9. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2013.
10. Session Chair, Int'l Conference on Supercomputing (ICS), 2012.
11. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2012.
12. Session Chair, Int'l Conference on Cluster Computing (Cluster), 2009.
13. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2009.
14. Session Chair, ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS), 2008.
15. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2008.
16. Session Chair, Int'l Symposium on High Performance Computing (HiPC), 2007.
17. Session Chair, Hot Interconnect (HotI), 2007.
18. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2007.
19. Session Chair, Hot Interconnect (HotI), 2006.
20. Session Chair, Int'l Symposium on High Performance Computing (HiPC), 2005.
21. Session Chair, Euro PVM/MPI Conference, 2005.

22. Session Chair, Int'l Conference on Parallel Processing (ICPP), 2005.
23. Session Chair, Int'l Conference on Cluster Computing (Cluster), 2004.
24. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2004.
25. Session Chair, Int'l Conference on Cluster Computing (Cluster), 2003.
26. Session Chair, Int'l Conference on Supercomputing (SC), 2003.
27. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2003.
28. Session Chair, Int'l Conference on Cluster Computing (Cluster '02), 2002.
29. Session Chair, Int'l Parallel and Distributed Processing Symposium (IPDPS), 2001.
30. Session Chair, Int'l Conference on Parallel Processing (ICPP), 2000.
31. Awards Chair, Int'l Conference on Parallel Processing (ICPP), 2000.
32. Session Chair, Int'l Workshop on Heterogeneous Computing (HCW), 2000.
33. Session Chair, Int'l Conference on Parallel Processing (ICPP), 1998.
34. Session Chair, Int'l Parallel Processing Symposium (IPPS), 1998.
35. Session Chair, Int'l Symposium on High-Performance Computer Architecture (HPCA-4), 1998.
36. Session Chair, Int'l Conference on Parallel Processing (ICPP), 1997.
37. Session Chair, Parallel Computer, Communication, and Routing Workshop (PCRCW), 1997.
38. Tutorial Chair, Int'l Parallel Processing Symposium (IPPS), 1998.
39. Tutorial Chair, Int'l Conference on Parallel Processing (ICPP), 1997.
40. Publications Chair, Third Int'l Symposium on High-Performance Computer Architecture (HPCA-3), 1997.
41. Session Chair, Int'l Symposium on Parallel and Distributed Processing (SPDP), 1996.
42. Session Chair, Int'l Conference on Parallel Processing (ICPP), 1996.
43. Session Chair, Int'l Parallel Processing Symposium (IPPS), 1996.
44. Session Chair, Int'l Parallel Processing Symposium (IPPS), 1995.
45. Organizing Committee Member, Int'l Conference on Parallel and Distributed Systems (ICPADS), 1996.
46. Executive Committee Member in IEEE Technical Committee on Parallel Processing (TCPP), 1992–97.
47. Organizing Committee Member, Int'l Parallel Processing Symposium (IPPS), 1994.
48. Organizing Committee Member, Int'l Workshop on Parallel Processing (IWPP), 1994.
49. Session Chair, Int'l Parallel Processing Symposium (IPPS), 1993.

- **Member of Steering Committee/Advisory Board:**

Large-Scale IMage Processing Infrastructure Development (LIMPID) project, 2019-2020  
 2019 Workshop on the Future Directions for the Cyberinfrastructure for Sustained Scientific

Innovation (CSSI), 2019  
Int'l Workshop on Parallel Programming Models in High-Performance Cloud (ParaMo) 2019-20  
Int'l Open Benchmarking Council, 2019-current  
IEEE Cluster Conference, 2000-2018  
Int'l Workshop on Benchmarks, Performance Optimization, and Emerging hardware of Big Data Systems and Applications (BPOE), 2013-current  
Workshop on Communication Architecture for Scalable Systems (CASS), 2011-2014  
Workshop on Communication and Architecture (CAC), 2001-2010  
Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), 2008-2010

- **Evaluation Committee:**

IEEE Fellows Evaluation Committee, 2020  
IEEE Fellows Evaluation Committee, 2017  
IEEE Fellows Evaluation Committee, 2015  
IEEE Fellows Evaluation Committee, 2014  
IEEE Fellows Evaluation Committee, 2013  
IEEE Fellows Evaluation Committee, 2009  
IEEE Fellows Evaluation Committee, 2008

- **Proposal Review Panel:**

National Science Foundation, CSSI Panel, January 2020  
National Science Foundation, Cybertraining Panel, February 2019  
National Science Foundation, CSR Panel, Oct 2017  
National Science Foundation, S2I2 Panel, June 2017  
National Science Foundation, OAC Cybertraining Panel, Mar 2017  
National Science Foundation, CNS Virtual Panel, Feb 2016  
National Science Foundation, CSR Virtual Panel, March 2015  
National Science Foundation, SHF Panel, Jan 2015  
National Science Foundation, CCF Panel, Oct 2014  
INRIA, France, On-Site Proposal Review Panel, October 2012  
National Science Foundation, CNS Panel, April 2012  
National Science Foundation, MRI Panel, October 2009  
National Science Foundation, STCI Panel, July 2009  
National Science Foundation, Architecture Panel, Mar 2008  
National Science Foundation, NBD Panel, May 2007  
National Science Foundation, CRI Panel, March 2007  
National Science Foundation, DDDRS Panel, August 2005  
National Science Foundation, SBIR Panel, Sept 2002  
National Science Foundation, ITR Small-scale Panel, May 2002  
National Science Foundation, Research Infrastructure Panel, Dec 2001  
National Science Foundation, Operating Systems and Compilers Panel, Feb 2000  
National Science Foundation, Architecture CAREER Panel, Oct 1999  
National Science Foundation, Architecture Program Panel, April 1998



- **Referee:**

*Funding/Grant Agencies:*

- National Science Foundation (NSF), USA
- NSERC/CRSNG, Canada
- Cornell National Supercomputer Facility, USA
- Arkansas Science and Technology Authority, USA
- Research Grants Council, Hong Kong
- Norway Research Council, Norway
- Austrian Research Agency, Austria
- Louisiana Board of Regents

*Journals:*

- IEEE Transactions on Parallel and Distributed Systems
- IEEE Transactions on Computers
- Journal of Parallel and Distributed Computing
- Parallel Computing
- IEEE Parallel and Distributed Technology
- IEEE Computer
- ACM Computing Surveys
- Parallel Processing Letters
- Int'l Journal of Computers and Mathematics with Applications

Recognized as a member of **Senior Reviewers List** of IEEE COMPUTER by Prof. Ted Lewis, Editor-in-chief, March 1994.

Recognized as an **Outstanding Reviewer** of IEEE COMPUTER by Prof. Edward A. Parrish, Editor-in-chief, Jan 1995.

Recognized as a **Meritorious Reviewer** of IEEE Parallel and Distributed Technology by Prof. Michael J. Quinn, Summer 1995.

*Conferences:*

- International Symposium on Computer Architecture (ISCA)
- High Performance Computing Architecture (HPCA)
- International Conference on Parallel Processing (ICPP)
- International Parallel and Distributed Processing Symposium (IPDPS)
- Int'l Conference on High Performance Computing (HiPC)
- Supercomputing (SC)
- High Performance Distributed Computing (HPDC)
- International Conference on Parallel and Distributed Systems (ICPADS)
- International Conference on Supercomputing (ICS)
- Cluster Computing (Cluster)

- **Memberships:**

- IEEE, Fellow, special interest in Computer Society.
- ACM, special interest in computer architecture.
- Computer Society of India, special interest in hardware.

## Invited Keynote Talks, Tutorials, and Presentations

• **Invited Keynote/Plenary Talks**

1. *Scalable and Distributed DNN Training on Modern HPC Systems*, Int'l Conference on Machine Learning and Data Science, December 17, 2019.
2. *Benchmarks and Middleware for Designing Convergent HPC, Big Data and Deep Learning Software Stacks for Exascale Systems*, International Symposium on Benchmarking, Measuring and Optimizing (Bench '19), November 15, 2019.
3. *Designing Scalable HPC, Deep Learning, Big Data, and Cloud Middleware for Exascale Systems*, HPC-AI UK Conference, September 16, 2019.
4. *Scalable and Distributed DNN Training on Modern HPC Systems: Challenges and Solutions* HPC-AI Perth Conference, August 27, 2019.
5. *How to Design Convergent HPC, Big Data Analytics and Deep Learning Software Stacks for Exascale Systems?* Int'l Symposium on Intelligent Computers, June 27, 2019.
6. *Scalable and Distributed DNN Training on Modern HPC Systems: Challenges and Solutions*, Scalable Data Analytics in Scientific Computing (SDASC) Workshop, in conjunction with ISC '19 conference, June 20, 2019.
7. *Scalable and Distributed DNN Training on Modern HPC Systems: Challenges and Solutions*, High Performance Machine Learning (HPML) Workshop, in conjunction with CCGrid '19 conference, May 14, 2019.
8. *How to Design Convergent HPC, Deep Learning and Big Data Analytics Software Stacks for Exascale Systems?*, SCAsia '19 Conference, March 13, 2019.
9. *RDMA-Based Networking Technologies and Middleware for Next-Generation Clusters and Data Centers*, HiPINEB Workshop, in conjunction with HPCA Conference, February 17, 2019.
10. *How to Design Scalable HPC, Deep Learning and Cloud Middleware for Exascale Systems?*, HPC Advisory Council Stanford Conference, February 14, 2019.
11. *Designing Scalable HPC, Big Data, Deep Learning, and Cloud Computing Middleware for Exascale Systems*, HPC Advisory Council China Conference, October 17, 2018.
12. *RDMA-Based Networking Technologies and Middleware for Next-Generation Clusters and Data Centers*, Workshop on Kernel Bypassing Networks (KBNets), in conjunction with SigComm conference, August 20, 2018.
13. *Networking Technologies and Middleware for Next-Generation Clusters and Data Centers: Opportunities and Challenges*, IEEE High-Performance Switching and Routing (HPSR) Conference, June 18, 2018.
14. *Exploiting Computation and Communication Overlap in MVAPICH2 MPI Library*, Charm++ Workshop, April 11, 2018.
15. *Exploiting HPC Technologies for Accelerating Big Data Processing and Associated Deep Learning*, HPC Advisory Council Lugano Conference, April 10, 2018.
16. *Designing HPC, Big Data, Deep Learning, and Cloud Middleware for Exascale Systems: Challenges and Opportunities*, NCAR SEA Symposium, April 3, 2018.

17. *Designing HPC, Deep Learning and Cloud Middleware for Exascale Systems*, HPC Advisory Council Stanford Conference, February 20, 2018.
18. *HPC Meets Cloud: Opportunities and Challenges in Designing High-Performance MPI and Big Data Libraries on Virtualized InfiniBand Clusters*, Int'l Conference on CloudTech, Rabat, Morocco, October 24, 2017.
19. *Networking Technologies and Middleware for Next-Generation Data Centers: Challenges and Opportunities*, Hot DataCenter '17, in Conjunction with HPC China Conference, October 18, 2017.
20. *Designing Efficient HPC, Big Data, Deep Learning, and Cloud Computing Middleware for Exascale Systems*, HPC Advisory Council China Conference, October 18, 2017.
21. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, HPC Advisory Council Spain Conference, September 29, 2017.
22. *High-Performance and Scalable Designs of Programming Models for Exascale Systems*, HPC Advisory Council Lugano Conference, April 10, 2017.
23. *HPC Meets Big Data: Accelerating Hadoop, Spark, and Memcached*, HPC Advisory Council Stanford Conference, February 7, 2017.
24. *HPC Meets Cloud: Opportunities and Challenges in Designing High-Performance MPI and Big Data Libraries on Virtualized InfiniBand Clusters*, Int'l Workshop on Virtualization Solutions for High-Performance Computing (VisorHPC), in conjunction with HiPEAC 2017, Stockholm, January 24, 2017.
25. *HPC Meets Cloud: Opportunities and Challenges in Designing High-Performance MPI and Big Data Libraries on Virtualized InfiniBand Clusters*, 8th Int'l Conference on Cloud Computing Technology and Science, Luxembourg, December 13, 2016.
26. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, HPC Advisory Council China Conference, October 26, 2016.
27. *Designing OpenSHMEM and Hybrid MPI+OpenSHMEM Libraries for Exascale Systems: MVAPICH2-X Experience*, Invited Plenary Talk, OpenSHMEM 2016 Conference, August 3, 2016.
28. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Big Data Processing*, Keynote Talk, HPC Advisory Council Switzerland Conference, March 22, 2016.
29. *Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council Stanford Conference, February 24, 2016.
30. *Designing HPC and Big Data Middleware for Exascale Systems: Opportunities and Challenges*, Keynote Talk, Int'l Symposium on Computational Sciences (ISCS '15), December 14, 2015.
31. *Designing HPC and Big Data Middleware for Exascale Systems: Opportunities and Challenges*, Plenary Talk, South African High-Performance Computing Conference (CHPC '15), Dec 4, 2015.
32. *High Performance and Scalable Communication Libraries for HPC and Big Data*, Keynote Talk, HPC Advisory Council China Conference, November 9, 2015.

33. *Communication Frameworks for HPC and Big Data*, Keynote Talk, HPC Advisory Council Spain Conference, September 22, 2015.
34. *Accelerating and Benchmarking Big Data Processing on Modern Clusters*, Keynote Talk, Int'l Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-6), held in conjunction with VLDB '15, September 4, 2015.
35. *Challenges in Designing HPC and Big Data Middleware for Exascale Systems*, Keynote Talk, HPC Advisory Council Brazil Conference, August 26, 2015.
36. *High-Performance and Scalable Designs of Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council Switzerland Conference, March 24, 2015.
37. *High-Performance and Scalable Designs of Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council Stanford Conference, February 2, 2015.
38. *Accelerating Big Data Processing on Modern Clusters*, Keynote Talk, Int'l Workshop on Performance Analysis, of Big Data Systems (PABS), Austin, Feb 1, 2015.
39. *Designing Software Libraries and Middleware for Exascale Computing: Opportunities and Challenges*, Plenary Talk, South African High-Performance Computing Conference (CHPC '14), Dec 4, 2014.
40. *Networking Technologies for Exascale Computing Systems: Opportunities and Challenges*, Keynote Talk, China High Performance Interconnect Forum, Nov 7, 2014.
41. *Challenges in Designing Communication Libraries for Exascale Systems*, Keynote Talk, HPC China Conference, Nov 7, 2014.
42. *Accelerating Big Data Processing on Modern HPC Clusters*, Keynote Talk, Sixth Workshop on Big Data, Performance Optimization and Emerging Technologies (BPOE-6), Nov 6, 2014.
43. *High-Performance and Scalable Designs of Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council China Conference, Nov 5, 2014.
44. *Accelerating Data Management and Processing on Modern Clusters with RDMA-Enabled Interconnects*, Keynote Talk, ADMS '14 Workshop, in conjunction with VLDB '14 Conference, Sep 1, 2014.
45. *Designing Software Libraries and Middleware for Exascale Systems: Opportunities and Challenges*, Keynote Talk, HPC Advisory Council Brazil Conference, May 26, 2014.
46. *Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council Switzerland Conference, March 31, 2014.
47. *Accelerating Big Data Processing with RDMA-Enhanced Apache Hadoop*, Keynote Talk, Fourth Workshop on Big Data, Performance Optimization and Emerging Technologies (BPOE-4), March 1, 2014.
48. *Programming Models for Exascale Systems*, Keynote Talk, HPC Advisory Council Stanford Conference, February 3, 2014.
49. *Challenges in Accelerating Big Data Processing on Modern Clusters*, Keynote Talk, CCF Big Data Technology Conference, Beijing, China, December 6, 2013.
50. *Impact of Networking Technologies and Protocols on Hadoop*, Keynote Talk, Big Data Technology Conference, Beijing, China, December 5, 2013.

51. *Accelerating Big Data Processing with Hadoop-RDMA*, Keynote Talk, Big Data Technology Conference, Beijing, China, December 4, 2013.
52. *Challenges in Benchmarking and Evaluating Big Data Processing Middleware on Modern Clusters*, Keynote Talk, Fourth Workshop on Big Data, Performance Optimization and Emerging Technologies (BPOE-3), October 30, 2013.
53. *Accelerating Big Data Processing with Hadoop-RDMA*, Keynote Talk, Big Data Forum, HPC China Conference, October 29, 2013.
54. *New Developments for Message Passing, PGAS and Big Data*, Keynote Talk, HPC Advisory Council China Conference, October 28, 2013.
55. *Accelerating Big Data with Hadoop and Memcached*, Keynote Talk, HPC Advisory Council Stanford Conference, Feb 8, 2013.
56. *Programming Models and their Designs for Exascale Systems*, Keynote Talk, HPC Advisory Council Stanford Conference, Feb 7, 2013.
57. *The Future of MPI and PGAS Programming Models for Exascale Computing*, Keynote Talk, HP-CAST, November 10, 2012.
58. *Software Libraries and Middleware for Exascale Systems*, Keynote Talk, HPC Advisory Council Workshop, Zhangjiajie, China, October 28, 2012.
59. *Designing Networking Technologies and Communication Libraries for Exascale Computing Systems: Opportunities and Challenges*, Keynote Talk, IEEE Cluster Conference (Cluster 2012), September 25, 2012.
60. *The Future of MPI*, Keynote Talk, HPC Advisory Council Workshop, Malaga, Spain, September 13, 2012.
61. *The Future of Supercomputers Software Libraries*, Keynote Talk, Israel Supercomputing Conference, Tel Aviv, Israel, Feb 7, 2012.
62. *Driving Toward Exascale Computing: PGAS and Latest MPI Development*, Keynote Talk, HPC Advisory Council Workshop, Jinan, China, Oct 26, 2011.
63. *Co-Designing MPI Library and Applications for InfiniBand Clusters*, Keynote Talk, Int'l Workshop on Co-Design, Beijing, China, Oct 25, 2011.
64. *Networking Technologies for Clusters: Where do We Stand and What Lies Ahead?*, Keynote Talk, Int'l Conference on Parallel and Distributed Systems (ICPADS '10), Shanghai, China, Dec 9, 2010.
65. *Networking Technologies for Exascale Computing Systems: Opportunities and Challenges*, Keynote Talk, HPC China Conference, Beijing, Oct 28, 2010.
66. *Design of Collectives and One-Sided Operations in MPI and their Impact on Application-Level Performance and Scalability*, Keynote Talk, HPC Advisory Council Workshop, Beijing, China, Oct 27, 2010.
67. *Networking Technologies for Clusters: Where do We Stand and What Lies Ahead?*, Plenary Talk, Int'l Conference on Cluster Computing (Cluster '09), Sept. 2, 2009.
68. *Designing Next Generation Clusters with InfiniBand and 10GE/iWARP: Opportunities and Challenges*, Plenary Talk, Int'l Conference on Cluster Computing (Cluster '08), Oct. 2008.

69. *Fault-Tolerant/System-Management Issues in InfiniBand*, Int'l Workshop on System Management Techniques, Processes and Services (SMPTS), April 14, 2008.
70. *Emerging Networking Technologies and Protocols for Next Generation Clusters*, Int'l Conference on Parallel and Distributed Systems (ICPADS), July 13, 2006.
71. *Active Network Interface: Opportunities and Challenges*, Workshop on High-Speed Local Networks, held in conjunction with 27th Annual IEEE Conference on Local Computer Networks (LCN), Nov. 6, 2002.
72. *Network-Based Computing: Current Trends and Its Future*, Scientific Computing Workshop, NASA Glenn Research Center, Cleveland, Ohio, Apr. 13, 2000.
73. *Network-Based Computing: The Next Frontier*, Int'l Conference on Parallel and Distributed Computing and Systems (PDCS '98), Las Vegas, Oct. 30, 1998.

- **Invited Tutorials**

1. *High Performance Distributed Deep Learning: A Beginner's Guide*, ISCA '20. May 31, 2020.
2. *High Performance Distributed Deep Learning*, NVIDIA GTC '20. March 23, 2020.
3. *High Performance Distributed Deep Learning*, PPOPP '20. February 22, 2020.
4. *How to Boost the Performance of HPC/AI Applications using MVAPICH2 Libraries?*, HiPEAC '20. January 22, 2020.
5. *High Performance Distributed Deep Learning: A Beginner's Guide*, HiPEAC '20. January 22, 2020.
6. *InfiniBand, Omni-Path, and High-Speed Ethernet for Beginners*, SC '19, Nov 18, 2019.
7. *InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SC '19, Nov 18, 2019.
8. *High Performance Distributed Deep Learning: A Beginner's Guide*, SC '19. November 17, 2019.
9. *High-Performance Architectures for Distributed Deep Learning*, IEEE Micro '19. October 13, 2019.
10. *HPC Meets Distributed Deep Learning*, Hot Interconnect '19. August 14, 2019.
11. *High Performance Distributed Deep Learning: A Beginner's Guide*, PEARC '19. July 29, 2019.
12. *Accelerating Big Data Processing and Associated Deep Learning on Modern Datacenters*, ISCA '19. June 23, 2019.
13. *High Performance Distributed Deep Learning: A Beginner's Guide*, ISCA '19. June 22, 2019.
14. *High Performance Distributed Deep Learning: A Beginner's Guide*, ISC '19. June 16, 2019.
15. *InfiniBand and High-Speed Ethernet for Beginners*, ISC '19, June 16, 2019.
16. *Accelerating Big Data Processing on Modern HPC Clusters*, CCGrid '19. May 16, 2019.

17. *High Performance Distributed Deep Learning: A Beginner's Guide*, CCGrid '19. May 15, 2019.
18. *How to Boost the Performance of Your HPC/AI Applications with MVAPICH2 Libraries?*, NCAR SEA Symposium, April 12, 2019.
19. *High Performance Distributed Deep Learning: A Beginner's Guide*, NCAR SEA Symposium, April 12, 2019.
20. *How to Boost the Performance of HPC/AI Applications Using MVAPICH2 Library?*, NVIDIA GTC '19. March 20, 2019.
21. *High Performance Distributed Deep Learning: A Beginner's Guide*, NVIDIA GTC '19. March 18, 2019.
22. *High Performance Distributed Deep Learning*, PPOPP '19. February 17, 2019.
23. *Exploiting HPC Technologies for Accelerating Big Data Processing and Associated Deep Learning*, SC '18. November 12, 2018.
24. *InfiniBand, Omni-Path, and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SC '18, November 11, 2018.
25. *InfiniBand and High-Speed Ethernet for Beginners*, SC '18, November 11, 2018.
26. *How to Accelerate Your Big Data and Associated Deep Learning Applications with Hadoop and Spark: A Hands-On Approach*, PEARC '18, July 24, 2018.
27. *InfiniBand and High-Speed Ethernet for Beginners*, ISC '18, June 24, 2018.
28. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, ISC '18, June 24, 2018.
29. *Accelerating Big Data Processing and Associated Deep Learning on Datacenters with Modern Architectures*, ISCA '18. June 2, 2018.
30. *Performance Tuning of MVAPICH2 and MVAPICH2-GDR MPI Libraries*, HPCMP PETTT Technical Interchange Meeting, May 15, 2018.
31. *High-Performance Distributed Deep Learning: A Beginner's Guide*, HPCMP PETTT Technical Interchange Meeting, May 15, 2018.
32. *Acceleration of Big Data Stacks and Associated Deep Learning on Modern Clusters*, HPCMP PETTT Technical Interchange Meeting, May 15, 2018.
33. *How to Boost the Performance of Your MPI and PGAS Applications with MVAPICH2 Libraries*, NCAR SEA Symposium, April 5, 2018.
34. *High-Performance Distributed Deep Learning: A Beginner's Guide*, NCAR SEA Symposium, April 5, 2018.
35. *Accelerating Big Data Processing and Associated Deep Learning on Datacenters and HPC Clouds with Modern Architectures*, ASPLOS '18. March 24, 2018.
36. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, HPCA '18. February 25, 2018.
37. *High-Performance Distributed Deep Learning: A Beginner's Guide*, PPOPP '18, February 25, 2018.

38. *High-Performance Distributed Deep Learning for Dummies*, IT4 Innovations, January 24, 2018.
39. *InfiniBand and High-Speed Ethernet for Dummies*, IT4 Innovations, January 23, 2018.
40. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, IT4 Innovations, January 23, 2018.
41. *HPC Meets Cloud: Building Efficient Clouds for HPC, Big Data, and Deep Learning Middleware and Applications*, UCC '17, December 8, 2017.
42. *InfiniBand and High-Speed Ethernet for Dummies*, Supercomputing '17, November 13, 2017.
43. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, Supercomputing '17, November 13, 2017.
44. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Big Data Processing and Management*, Supercomputing '17, November 12, 2017.
45. *Exploiting High-Performance Interconnects to Accelerate Big Data Processing with Hadoop, Spark, Memcached, and gRPC/TensorFlow*, Hot Interconnect '17, August 28, 2017.
46. *High-Performance Distributed Deep Learning for Dummies*, Hot Interconnect '17, August 28, 2017.
47. *How to Accelerate Your Big Data Applications with Hadoop and Spark?*, PEARC '17, July 10, 2017.
48. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ISCA '17, June 25, 2017.
49. *InfiniBand and High-Speed Ethernet for Dummies*, ISC '17, June 18, 2017.
50. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, ISC '17, June 18, 2017.
51. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Modern Clusters*, CCGRID '17. May 15, 2017.
52. *Programming, Optimizing and Tuning your Message Passing Interface (MPI) Application with MVAPICH2 and MVAPICH2-GDR Libraries*, HPCMP PETTT Technical Interchange Meeting, May 2, 2017.
53. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ASPLOS '17. April 8, 2017.
54. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, HPCA '17. February 5, 2017.
55. *InfiniBand and High-Speed Ethernet for Dummies*, SC '16, November 13, 2016.
56. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, SC '16, November 13, 2016.
57. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Apache Hadoop, Spark and Memcached*, SC '16. November 13, 2016.
58. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Apache Hadoop, Spark and Memcached*, IEEE Cluster '16. September 13, 2016.



59. *PGAS And Hybrid MPI+PGAS Programming Models On Modern HPC Clusters With Accelerators*, IEEE Cluster '16. September 13, 2016.
60. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Modern Clusters*, Int'l Conference on Field Programmable Logic (FPL) 2016. September 2, 2016.
61. *Accelerating Big Data Processing with Hadoop, Spark and Memcached Over High-Performance Interconnects*, Hot Interconnect 2016. August 26, 2016.
62. *PGAS And Hybrid MPI+PGAS Programming Models On Modern HPC Clusters With Accelerators*, ICS '16. May 31, 2016.
63. *How to Boost the Performance of your MPI and PGAS Applications with MVAPICH2 Libraries*, XSEDE '16, July 18, 2016.
64. *InfiniBand and High-Speed Ethernet for Dummies*, ISC '16, June 19, 2016.
65. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, ISC '16, June 19, 2016.
66. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ISCA 2016. June 18, 2016.
67. *PGAS And Hybrid MPI+PGAS Programming Models On Modern HPC Clusters With Accelerators*, ICS '16. May 31, 2016.
68. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Modern Clusters*, CCGrid 2016. May 16, 2016.
69. *Programming, Optimizing and Tuning your Message Passing Interface (MPI) application with MVAPICH2 and MVAPICH2-GDR Libraries*, HPCMP PETTT KY721 Technical Interchange Meeting on High Performance Signal Processing, April 25, 2016.
70. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ASPLOS 2016. April 2, 2016.
71. *PGAS And Hybrid MPI+PGAS Programming Models On Modern HPC Clusters With Accelerators*, PPOPP '16. Mar 13, 2016.
72. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, HPCA 2016. Mar 12, 2016.
73. *InfiniBand and High-Speed Ethernet for Dummies*, South African High-Performance Computing Conference (CHPC '15), Dec 1, 2015.
74. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Clusters*, South African High-Performance Computing Conference (CHPC '15), Dec 1, 2015.
75. *InfiniBand and High-Speed Ethernet for Dummies*, SC '15, Nov 16, 2015.
76. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Clusters*, SC '15. Nov 16, 2015.
77. *PGAS and Hybrid MPI+PGAS Programming Models on Modern Clusters with Accelerators*, PGAS '15. Sept 16, 2015.
78. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Clusters*, HotI '15. Aug 28, 2015.

79. *Optimization and Tuning of MPI and PGAS Applications using MVAPICH2 Software Stacks* XSEDE'15, July 27, 2015.
80. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Clusters*, XSEDE '15. July 27, 2015.
81. *InfiniBand and High-Speed Ethernet for Dummies*, ISC '15, July 12, 2015.
82. *InfiniBand and High-Speed Ethernet: Advanced Features, Challenges in Designing HEC Systems, and Usage*, ISC '15, July 12, 2015.
83. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ISCA '15, June 13, 2015.
84. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures*, ASPLOS '15, March 15, 2015.
85. *PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters*, PPOPP '15, February 8, 2015.
86. *Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Networking and Storage Architecture*, HPCA '15, February 7, 2015.
87. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends*, South African Conference on High-Performance Computing (CHPC '14), December 1, 2014.
88. *Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters*, South African Conference on High-Performance Computing (CHPC '14), December 1, 2014.
89. *InfiniBand and High-Speed Ethernet for Dummies*, Supercomputing (SC '14), November 17, 2014.
90. *Advanced Topics in InfiniBand and High-Speed Ethernet for Designing High-End Computing Systems*, Supercomputing (SC '14), November 17, 2014.
91. *PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters*, Supercomputing (SC '14), November 16, 2014.
92. *MVAPICH2-X: Supporting PGAS and Hybrid MPI+PGAS Programming Models on Modern Clusters*, Intel Xeon Phi for PGAS Tutorial, in conjunction with PGAS '14 conference, October 8, 2014.
93. *Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters*, IEEE Cluster (Cluster '14), September 22, 2014.
94. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends*, IEEE Cluster (Cluster '14), September 22, 2014.
95. *Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters*, Hot Interconnect, (HotI), August 28, 2014.
96. *Optimization and Tuning of MPI and PGAS Applications using MVAPICH2 and MVAPICH2-X*, XSEDE '14, July 13, 2014.
97. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends*, Int'l Conference on High Performance Switching and Routing (HPSR '14), July 1, 2014.

98. *InfiniBand and High-Speed Ethernet: Overview, Latest Status and Trends*, Int'l Supercomputing Conference (ISC '14), June 22, 2014.
99. *Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters*, Int'l Symposium on Computer Architecture, (ISCA), June 14, 2014.
100. *PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters*, Int'l Conference on Supercomputing (ICS), June 11, 2014.
101. *Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters*, Int'l Symposium on Cluster, Cloud and Grid Computing, (CCGrid), May 29, 2014.
102. *Accelerating Big Data Processing with Hadoop and Memcached on Datacenters with Modern Networking and Storage Architecture*, Int'l Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), March 1, 2014.
103. *Accelerating Big Data Processing with Hadoop and Memcached on Datacenters with Modern Networking and Storage Architecture*, Int'l Symposium on High Performance Computer Architecture (HPCA), February 16, 2014.
104. *PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters*, Int'l Symposium on Principles and Practice of Parallel Programming (PPoPP), February 15, 2014.
105. *InfiniBand and High-Speed Ethernet for Dummies*, Supercomputing (SC '13), November 17, 2013.
106. *Advanced Topics in InfiniBand and High-Speed Ethernet for Designing HEC Systems*, Supercomputing (SC '13), November 17, 2013.
107. *Scalable and High-Performance Support for PGAS and Hybrid MPI+PGAS Programming Models on Modern Clusters: Opportunities and Challenges* Int'l Conference on PGAS Programming Models (PGAS), October 2, 2013.
108. *Accelerating Big Data with Hadoop and Memcached Using High Performance Interconnects: Opportunities and Challenges*, Hot Interconnect (HotI), August 23, 2013.
109. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '13), June 16, 2013.
110. *Advanced Topics in InfiniBand and High-Speed Ethernet for Designing HEC Systems*, Int'l Supercomputing Conference (ISC '13), June 16, 2013.
111. *InfiniBand and High-Speed Ethernet for Dummies*, Supercomputing (SC '12), November 12, 2012.
112. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, Supercomputing (SC '12), November 12, 2012.
113. *Designing Clusters with InfiniBand and High-Speed Ethernet: Current Status and Trends*, IEEE Cluster, September 24, 2012.
114. *Designing Scientific, Enterprise, and Cloud Computing Systems with InfiniBand and High-Speed Ethernet: Current Status and Trends*, Hot Interconnect (HOTI), August 24, 2012.
115. *InfiniBand and High-Speed Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '12), June 17, 2012.

116. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, Int'l Supercomputing Conference (ISC '12), June 17, 2012.
117. *Designing Cloud and Grid Computing Systems with InfiniBand and High-Speed Ethernet*, Int'l Conference on Cluster, Cloud and Grid Computing (CCGrid '12), May 13, 2012.
118. *InfiniBand and Ethernet Architectures for Scientific and Enterprise Computing: Opportunities and Challenges*, Int'l Symposium on High Performance Computing Architecture (HPCA '12), February, 25, 2012.
119. *InfiniBand and 10-Gigabit Ethernet for Dummies*, Supercomputing (SC '11), Nov. 13, 2011.
120. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, Supercomputing (SC '11), Nov. 13, 2011.
121. *Designing Clusters with InfiniBand and High-Speed Ethernet: Current Status and Trends*, IEEE Cluster, September 26, 2011.
122. *Designing Scientific and Enterprise Computing Systems with InfiniBand and High-Speed Ethernet: Current Status and Trends*, Hot Interconnect (HOTI), August 26, 2011.
123. *InfiniBand and 10-Gigabit Ethernet for Dummies*, Int'l Supercomputing Conference (ISC '11), jointly with S. Sur, June 19, 2011.
124. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, Int'l Supercomputing Conference (ISC '11), jointly with S. Sur, June 19, 2011.
125. *Designing High-End Computing Systems and Programming Models with InfiniBand and High-Speed Ethernet* Int'l Conference on Supercomputing (ICS '11), May 31, 2011.
126. *Designing Cloud and Grid Computing Systems with InfiniBand and High-Speed Ethernet*, Int'l Conference on Cluster and Grid Computing (CCGrid '11), jointly with S. Sur, May 23, 2011.
127. *InfiniBand and Ethernet Architectures for Scientific and Enterprise Computing: Opportunities and Challenges*, jointly with P. Balaji and S. Sur, Int'l Symposium on High Performance Computing Architecture (HPCA '11), Feb. 12, 2011.
128. *InfiniBand and 10-Gigabit Ethernet for Dummies*, jointly with P. Balaji and S. Sur, Int'l Conference on Supercomputing (SC '10), Nov. 15, 2010.
129. *Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet*, jointly with P. Balaji and S. Sur, Int'l Conference on Supercomputing (SC '10), Nov. 15, 2010.
130. *Designing High-End Computing Systems and Programming Models with InfiniBand and High-speed Ethernet*, jointly with P. Balaji and S. Sur, Int'l Conference on Partitioned Global Address Space (PGAS '10), October 15, 2010.
131. *Designing High-End Computing Systems with InfiniBand and High-speed Ethernet*, jointly with P. Balaji and S. Sur, Int'l Conference on Cluster Computing (Cluster '10), September 20, 2010.
132. *Designing High-End Computing Systems with InfiniBand and High-speed Ethernet*, jointly with P. Balaji and S. Sur, Hot Interconnect (HOTI), August 20, 2010.
133. *InfiniBand Software Networking Technologies*, Discovery 2015 Workshop, Oak Ridge National Laboratory, July 13, 2010.

134. *Designing High-End Computing Systems with InfiniBand and High-speed Ethernet*, jointly with P. Balaji and S. Sur, Int'l Conference on Supercomputing (ICS), June 1, 2010.
135. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji, Int'l Supercomputing Conference (ISC), May 30, 2010.
136. *InfiniBand and Ethernet Architectures for Scientific and Enterprise Computing: Opportunities and Challenges*, jointly with P. Balaji, Int'l Symposium on High-Performance Computer Architecture (HPCA-16), Jan 10, 2010.
137. *InfiniBand and 10-Gigabit Ethernet for Dummies*, jointly with P. Balaji and M. Koop, Int'l Conference on Supercomputing (SC '09), Nov. 15, 2009.
138. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji and M. Koop, Int'l Conference on Supercomputing (SC '09), Nov. 15, 2009.
139. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji and M. Koop, IEEE Cluster (Cluster '09), September 4, 2009.
140. *InfiniBand and 10-Gigabit Ethernet for Dummies*, jointly with P. Balaji and M. Koop, Hot Interconnect (HotI '09), August 25, 2009.
141. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji and M. Koop, Hot Interconnect (HotI '09), August 25, 2009.
142. *InfiniBand and 10-Gigabit Ethernet Architecture for Emerging HPC Clusters and Enterprise Datacenters*, jointly with P. Balaji, Int'l Conference on High Performance Computing Architecture (HPCA), Feb. 14, 2009.
143. *InfiniBand and 10-Gigabit Ethernet for Dummies*, jointly with P. Balaji and M. Koop, Int'l Conference on Supercomputing (SC '08), Nov. 16, 2008.
144. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji and M. Koop, Int'l Conference on Supercomputing (SC '08), Nov. 16, 2008.
145. *Designing HPC Clusters and Enterprise Datacenters: The InfiniBand and 10GE Way*, jointly with P. Balaji, Int'l Symposium on Hot Interconnect (HOTI), August 26, 2008.
146. *High-Speed Network Architectures for Clusters: Designs and Trends*, jointly with P. Balaji, Int'l Conference on High Performance Computing Architecture (HPCA), Feb. 17, 2008,
147. *Designing High-End Computing Systems with InfiniBand and iWARP Standards*, jointly with P. Balaji and S. Sur, Int'l Conference on Supercomputing (SC '07), Nov. 12, 2007.
148. *Designing High-End Computing Systems with InfiniBand and 10-Gigabit Ethernet*, jointly with P. Balaji, Int'l Conference on IEEE Cluster (Cluster'07), Sept. 17, 2007.
149. *Designing Clusters and Grid Computing Systems with InfiniBand and iWARP*, jointly with P. Balaji, Int'l Conference on Cluster Computing and Grid (CCGrid'07), May 14, 2007.
150. *Storage Networks, Protocols and File Systems: Latest Trends?*, jointly with W. Yu, Hot Interconnect, August 25, 2006.

151. *Designing Next Generation High-End Computing Systems with InfiniBand*, Proctor and Gamble, Cincinnati, July 19, 2006.
152. *State of InfiniBand in Designing Next Generation Clusters, File/Storage Systems and Datacenters*, Int'l Conference on High Performance Computing (HiPC '05), Dec 21, 2005.
153. *State of InfiniBand in Designing Next Generation Clusters, File/Storage Systems and Datacenters*, Int'l Conference on Supercomputing (SC '05), Nov. 13, 2005.
154. *State of InfiniBand in Designing Next Generation Clusters, File/Storage Systems and Datacenters*, Int'l Cluster Symposium (Cluster '05), September 26, 2005.
155. *State of InfiniBand in Designing Next Generation Clusters, File/Storage Systems and Datacenters*, Commodity Cluster Symposium (CCS), July 28, 2005.
156. *Designing Next Generation HPC Clusters, Storage/File Systems, and Datacenters with InfiniBand: Opportunities and Challenges*, Int'l Conference on High Performance Computing (HiPC '04), Dec. 22, 2004.
157. *State of InfiniBand in Designing HPC Clusters, Storage/File Systems, and Datacenters*, Int'l Conference on Supercomputing (SC '04), Nov. 7, 2004.
158. *State of InfiniBand in Designing Next HPC Clusters, Storage/File Systems, and Datacenters*, Int'l Conference on Cluster Computing (Cluster '04), Sept, 20, 2004.
159. *State of InfiniBand in Designing Next Generation Clusters, Cluster-based Servers and Datacenters*, Int'l Conference on Cluster Computing and Grid (CCGrid'04), April 19, 2004.
160. *Designing Next Generation Clusters with InfiniBand: Opportunities and Challenges*, Int'l Conference on Cluster Computing (Cluster '03), Hong Kong, Dec 1, 2003.
161. *InfiniBand Architecture and Its Impact on High Performance Computing*, Int'l Conference on Parallel Processing (ICPP), Kahosiung, Taiwan, Oct 6, 2003.
162. *Designing Next Generation Clusters, Cluster-based Servers and Datacenters with InfiniBand: Opportunities and Challenges*, Hot Interconnect (HotI 11), Stanford, August 22, 2003.
163. *InfiniBand Architecture: Is it Ready for Prime Time?*, DOE InfiniBand Workshop, Jan 20, 2003.
164. *InfiniBand Architecture and Where is it Headed*, Systran Corporation, Dayton, Dec 9, 2002.
165. *InfiniBand Architecture and Its Impact on High Performance Computing*, Supercomputing (SC '02), Nov. 17, 2002.
166. *InfiniBand Architecture: Where is it Headed and What will be the Impact on Cluster Computing*, Cluster Computing (Cluster '02), Sept.23, 2002.
167. *InfiniBand Architecture and Its Impact on Designing Next Generation Networked Computing Systems*, Int'l Conference on Parallel and Distributed Computing Systems (PDCS '02), Sept. 19, 2002.
168. *InfiniBand Architecture and Where it is Headed?*, Dell, Austin, August 30, 2002.
169. *InfiniBand Architecture and Where it is Headed?*, Hot Interconnect (HotI 10), Stanford, August 23, 2002.

170. *InfiniBand Architecture and Its Impact on High Performance Computing*, Int'l Conference on Supercomputing (ICS '02), New York, June 22, 2002.
171. *InfiniBand Architecture and Its Impact on High Performance Computing*, Supercomputing '01, Denver, Nov. 12, 2001.
172. *InfiniBand Architecture*, Cluster 2001, Newport Beach, October 8, 2001.
173. *InfiniBand Architecture*, Spanish Parallel Computing Conference, Valencia, Spain, September 4, 2001.
174. *InfiniBand Architecture*, Hot Interconnect (Hot 9), Stanford, August 24, 2001.
175. *Network-Based Computing: Current Trend, Challenges, and the Future*, MASCOTS 2001, Cincinnati, August 15, 2001.
176. *VIA and InfiniBand Architecture*, Int'l Symposium on Computer Architecture (ISCA), Gotberg, Sweden, June 30, 2001.
177. *VIA and InfiniBand Architecture*, Int'l Symposium on High Performance Computer Architecture (HPCA-7), Monterrey, Mexico, Jan 22, 2001.
178. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE CTP Program), IEEE Computer Society Chapter, San Jose, Costa Rica, July 19, 1999.
179. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE CTP Program), IEEE Computer Society Chapter, Halifax, Nova Scotia, Canada, Oct 3, 1998.
180. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE CTP Program), IEEE Computer Society Chapter, Salvador, Brazil, August 27, 1998.

• **Invited Presentations**

1. *Scalable Distributed Deep Learning on Modern HPC Systems*, IBM OpenPOWER Cineca Workshop, July 9, 2020.
2. *MVAPICH Touches the Cloud: New Frontiers for MPI in High-Performance Clouds*, Open Fabrics Workshop, June 12, 2020.
3. *Designing a Deep-Learning Aware MPI Library: An MVAPICH2 Approach*, Open Fabrics Workshop, June 8, 2020.
4. *High-Performance, Scalable and Distributed Training on Modern HPC Systems with MPI Library: An MVAPICH2 Approach*, IBM OpenPOWER Academic and Research Webinar, May 15, 2020.
5. *How to Achieve High-Performance, Scalable and Distributed DNN Training on Modern HPC Systems?*, HPC-AI Advisory Council Stanford Conference, April 22, 2020.
6. *Building the Software Foundation for HPC and AI*, HPC-AI Advisory Council Webinar, Singapore, March 24, 2020.
7. *Designing Next-Generation MPI Libraries for Emerging Dense GPU Systems*, NSF CSSI PI Meeting, February 13, 2020.
8. *Sustaining Software Design and Development in Academia: Experiences with the MVAPICH2, HiBD and HiDL Libraries*, NSF Scalable Image Informatics (LIMPID) Workshop, February 3, 2020.

9. *Accelerating HPC, Big Data and Deep Learning on OpenPOWER Platforms*, Jawaharlal Nehru Technological University, Hyderabad, India, December 18, 2019.
10. *Exploiting Multi-core Processors for HPC and Deep Learning: The MVAPICH2 Approach*, Intel Booth, Supercomputing '19, November 20, 2019.
11. *Achieving High Performance on Microsoft Azure HPC Cloud using MVAPICH2*, Microsoft Booth, Supercomputing '19, November 20, 2019.
12. *MVAPICH2-GDR: Pushing the Frontier of HPC and Deep Learning*, Mellanox Booth, Supercomputing '19, November 20, 2019.
13. *Achieving High Performance on AWS HPC Cloud using MVAPICH2-AWS*, AWS Booth, Supercomputing '19, November 19, 2019.
14. *Enhancing MPI Communication using Hardware Tag Matching: The MVAPICH Approach*, UCX BoF, Supercomputing '19, November 19, 2019.
15. *The MVAPICH2 Project: Latest Developments and Plans Towards Exascale Computing*, Mellanox Booth, Supercomputing '19, November 19, 2019.
16. *MVAPICH Performance on Arm at Scale*, ARM HPC User group, Supercomputing '19, November 18, 2019.
17. *Full Stack View of AI Influence on HPC Hardware Systems*, HPC Day with The Next Platform (in conjunction with SC '19), November 17, 2019.
18. *Accelerating HPC, Big Data and Deep Learning on OpenPOWER Platforms*, 4th OpenPOWER Academia Discussion Group Workshop, November 16, 2019.
19. *Designing and Deploying High-Performance and Scalable MPI Library for Amazon-AWS and Microsoft-Azure Cloud*, NSF Workshop on Next-Generation Cloud Research, November 12, 2019.
20. *Exploiting Emerging Multi-core Processors for HPC and Deep Learning using MVAPICH2 MPI Library*, Intel Extreme Performance User Group (IXPUG), September 24, 2019.
21. *Big Data, HPC, and Deep Learning to Understand Neuroscience*, Advanced Computational Neuroscience Network (ACNN) Workshop, September 19, 2019.
22. *Designing Convergent HPC, Big Data Analytics and Deep Learning Software Stacks for Exascale Systems*, Army Research Laboratory (ARL), July 17, 2019.
23. *Designing Convergent HPC, Big Data Analytics and Deep Learning Software Stacks for Exascale Systems*, Shenzhen Institutes of Advanced Technology (SIAT), Shenzhen, China, June 28, 2019.
24. *Accelerating HPC Applications on HPC Systems with Intel Omni-Path: The MVAPICH Approach*, Omni-Path User Group (OPUG) Meeting, ISC '19, June 19, 2019.
25. *Exploiting Multi-/Many-core Intel CPUs for Deep Learning using MVAPICH2 MPI Library*, Intel Booth, ISC '19, June 18, 2019.
26. *Enhancing MPI Communication using Accelerated Verbs and Tag Matching: the MVAPICH Approach*, UCX BoF, ISC '19, June 18, 2019.
27. *High-Performance and Scalable MPI, Deep Learning, Big Data and Cloud Middleware for HPC Clusters*, Microsoft-Azure, May 8, 2019.



28. *Designing Scalable HPC, Deep Learning and Cloud Middleware for Exascale Systems*, EPFL, Switzerland, April 3, 2019.
29. *Designing Convergent HPC, Deep Learning and Big Data Analytics Software Stacks for Exascale Systems*, HPC Advisory Council Lugano Conference, April 2, 2019.
30. *Scalable and Distributed DNN Training on Modern HPC Systems*, HPC Advisory Council Lugano Conference, April 1, 2019.
31. *Designing Convergent HPC, Deep Learning and Big Data Analytics Software Stacks for Exascale Systems*, IBM TJ Watson Research Center, March 25, 2019.
32. *MVAPICH2-GDR: High-Performance and Scalable CUDA-Aware MPI Library for HPC and AI*, NVIDIA GTC Conference, March 19, 2019.
33. *Scalable and Distributed DNN Training on Modern HPC Systems*, HPC-AI Competition Event, in conjunction with SCAsia '19, March 12, 2019.
34. *Accelerating HPC, Big Data and Deep Learning in Datacenters using OpenPOWER Systems*, OpenPOWER and AI Workshop, National University of Singapore, March 11, 2019.
35. *Designing Scalable HPC, Deep Learning, Big Data and Cloud Middleware for Exascale Systems*, IIT Bhubaneswar, India, March 7, 2019.
36. *Designing Convergent HPC and BigData Software Stacks: An Overview of the HiBD Project*, HPC Advisory Council Stanford Conference, February 14, 2019.
37. *Scalable and Distributed DNN Training on Modern HPC Systems*, Workshop on Deep Learning - Past, Present and Future of AI, in conjunction with HiPC '18, December 19, 2018.
38. *Designing Scalable HPC, Deep Learning, Big Data and Cloud Middleware for Exascale Systems*, Second Workshop on Software Challenges to Exascale Computing (SCEC) , December 13, 2018.
39. *Scalable and Distributed DNN Training on Modern HPC Systems*, Deep Learning Workshop, Supercomputing '18, November 16, 2018.
40. *MVAPICH2-GDR Library: Pushing the Frontier of HPC and Deep Learning*, Mellanox Theater, Supercomputing '18, November 15, 2018.
41. *MVAPICH2-GDR Library: Pushing the Frontier of HPC and Deep Learning*, NVIDIA Booth, Supercomputing '18, November 13, 2018.
42. *Unified Communication X (UCX) Community*, UCX BoF, Supercomputing '18, November 13, 2018.
43. *The MVAPICH2 Project: Latest Developments and Plans Towards Exascale Computing*, Oct 26, 2018. Mellanox Theater, Supercomputing '18, November 13, 2018.
44. *Designing Scalable HPC, Deep Learning, Big Data and Cloud Middleware for Exascale Systems*, AMD, Oct 26, 2018.
45. *Exploiting HPC Technologies for Accelerating Big Data Processing and Associated Deep Learning*, Int'l Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware, in conjunction with HPC-China Conference, October 19, 2018.

46. *Opportunities and Challenges for Middleware on Exascale Systems*, Int'l Workshop on Co-Design, in conjunction with HPC-China Conference, October 19, 2018.
47. *Designing High-Performance MPI Libraries for Multi-/Many-core Era*, Intel IXPUG Conference, Sept 25, 2018.
48. *Big Data, HPC, and Deep Learning to Understand Neuroscience*, ACNN 2018 Workshop, Sept 07, 2018.
49. *Building Efficient HPC Cloud with SR-IOV Enabled InfiniBand: The MVAPICH2 and HiBD Approach*, Cloud BoF at ISC '18, June 27, 2018.
50. *Enhancing MPI Communication using Accelerated Verbs: The MVAPICH Approach*, UCX BoF at ISC '18, June 27, 2018.
51. *Exploiting HPC Technologies to Accelerate Big Data Processing and Associated Deep Learning*, Intel Collaboration Hub at ISC '18, June 25, 2018.
52. *DLoBD: An Emerging Paradigm of Deep Learning Over Big Data Stacks*, Spark + AI Summit, San Francisco, June 5, 2018.
53. *Building Efficient HPC Clouds with MVAPICH2 and OpenStack over SR-IOV-enabled Heterogeneous Clusters*, OpenStack Summit, Vancouver, May 24, 2018.
54. *Accelerating TensorFlow with RDMA for High-Performance Deep Learning*, DataWorks Summit, Berlin, April 19, 2018.
55. *Designing Scalable HPC, Deep Learning and Cloud Middleware for Exascale Systems*, HPC Advisory Council Lugano Conference, April 10, 2018.
56. *Big Data and HPC to Understand Brain Health and Performance*, Brain Health and Performance Summit, April 6, 2018.
57. *Exploiting Computation and Communication Overlap in MVAPICH2 and MVAPICH2-GDR MPI Libraries*, Overlapping Communication with Computation Symposium, NCAR, April 4, 2018.
58. *MVAPICH2-GDR: Pushing the Frontier of MPI Libraries Enabling GPUDirect Technologies*, NVIDIA GTC '18 Conference, March 28, 2018.
59. *High-Performance MPI and Deep Learning on OpenPOWER Platform*, OpenPOWER Summit, March 19, 2018.
60. *High-Performance Hadoop and Spark on OpenPOWER Platform*, OpenPOWER Summit, March 19, 2018.
61. *Designing Converged Software Stacks for HPC, Big Data, and Deep Learning: Opportunities and Challenges*, RWBC-OIL Workshop, Tokyo, March 6, 2018.
62. *Co-Designing Interconnection Networks and HPC Middleware for Exascale Systems: Latest Status and Trends*, STEP HPC Co-Design Workshop, March 1, 2018.
63. *Big Data Meets HPC - Exploiting HPC Technologies for Accelerating Big Data Processing*, HPC Advisory Council Stanford Conference, February 21, 2018.
64. *Designing Efficient HPC, Big Data, Deep Learning, and Cloud Computing Middleware for Exascale Systems*, Simula Research Laboratory, Univ. of Oslo, December 14, 2018.

65. *MVAPICH2-GDR for HPC and Deep Learning*, Mellanox Theater, Supercomputing '17, November 16, 2017.
66. *MVAPICH2-GDR Library: Pushing the Frontier of HPC and Deep Learning*, NVIDIA Theater, Supercomputing '17, November 15, 2017.
67. *High-Performance Big Data (HiBD): Accelerating Hadoop, Spark, and Memcached on Modern Clusters*, Mellanox Theater, Supercomputing '17, November 15, 2017.
68. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Big Data Processing and Management*, SigHPC Big Data BoF Presentation, Supercomputing '17, November 15, 2017.
69. *Exploiting Latest Networking and Accelerator Technologies for MPI, Streaming, and Deep Learning: An MVAPICH2-Based Approach*, U.S. Naval Research Laboratory Booth, Supercomputing '17, November 15, 2017.
70. *Performance of PGAS Models on KNL: A Comprehensive Study with MVAPICH2-X*, Intel Nerve Center, Supercomputing '17, November 14, 2017.
71. *The MVAPICH2 Project: Latest Developments and Plans Towards Exascale Computing*, Mellanox Theater, Supercomputing '17, November 14, 2017.
72. *Scalability and Performance of MVAPICH2 on OakForest-PACS*, JCAHPC Booth, Supercomputing '17, November 14, 2017.
73. *Accelerate Big Data Processing (Hadoop, Spark, Memcached and TensorFlow) with HPC Technologies*, Intel HPC Dev Conference, in conjunction with Supercomputing '17, November 11, 2017.
74. *High-performance and Scalable MPI+X Library for Emerging HPC Clusters & Cloud Platforms*, Intel HPC Dev Conference, in conjunction with Supercomputing '17, November 11, 2017.
75. *Opportunities and Challenges for Middleware on Exascale Systems*, Int'l Workshop on Co-Design, in conjunction with HPC-China Conference, October 20, 2017.
76. *Exploiting High-Performance Computing (HPC) and Big Data to Accelerate Processing of Neuroscience Data*, Big Data Neuroscience Workshop, September 8, 2017.
77. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, US-China Workshop, in conjunction with PEARC '17 Conference, July 10, 2017.
78. *Computational Neuroscience Network (ACNN), Accelerated Processing of Big Neuroscience Data*, NSF BigData Workshop, in conjunction with IPDPS '17 Conference, June 02, 2017.
79. *MVAPICH2-GDR: Pushing the Frontier of HPC and Deep Learning*, NVIDIA GTC '17 Conference, May 11, 2017.
80. *Bringing NVIDIA GPUs to the PGAS/OpenSHMEM World: Challenges and Solutions*, NVIDIA GTC '17 Conference, May 11, 2017.
81. *High-Performance Broadcast Designs for Streaming Applications on Multi-GPU Infini-Band Clusters*, NVIDIA GTC '17 Conference, May 11, 2017.

82. *Building Efficient HPC Clouds with MVAPICH2 and OpenStack over SR-IOV enabled InfiniBand Clusters*, OpenStack Boston Summit, May 11, 2017.
83. *Accelerating OpenStack Swift with RDMA for Building Efficient HPC Clouds*, OpenStack Boston Summit, May 10, 2017.
84. *Accelerating Big Data Processing and Framework Provisioning with OpenStack Heat-based Hadoop/Spark*, OpenStack Boston Summit, May 9, 2017.
85. *High-Performance Broadcast with GPUDirect RDMA and InfiniBand Hardware Multicast for Streaming Applications*, PETTT TIM, May 3, 2017.
86. *Network-Based Computing for HPC, Big Data and Deep Learning*, Battelle and Midwest BigData Hub Meeting on Data Analytics, April 27, 2017.
87. *Big Data Meets HPC - Exploiting HPC Technologies for Accelerating Big Data Processing*, HPC Advisory Council Lugano Conference, April 11, 2017.
88. *NVM-aware RDMA-based Communication and I/O Schemes for High-Performance Big Data Analytics*, OpenFabrics Workshop, March 31, 2017.
89. *Building Efficient HPC Clouds with MVAPICH2 and RDMA-Hadoop over SR-IOV IB Clusters*, OpenFabrics Workshop, March 30, 2017.
90. *Designing MPI and PGAS Libraries for Exascale Systems: The MVAPICH2 Approach*, OpenFabrics Workshop, March 30, 2017.
91. *HPC Meets Big Data: Accelerating Hadoop, Spark, and Memcached with HPC Technologies*, OpenFabrics Workshop, March 28, 2017.
92. *Designing HPC and Deep Learning Middleware for Exascale Systems*, HPC Advisory Council Stanford Conference, February 8, 2017.
93. *MVAPICH2-GDR: Pushing the Frontier of HPC and Deep Learning*, Mellanox Theater, SC '16, November 17, 2016.
94. *Designing HPC, Big Data, and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, HPC Connection Workshop, SC '16, November 16, 2016.
95. *High Performance Big Data (HiBD): Accelerating Hadoop, Spark and Memcached on Modern Clusters*, Mellanox Theater, SC '16, November 16, 2016.
96. *MVAPICH2-GDR: Pushing the Frontier of HPC and Deep Learning*, NVIDIA Theater, SC '16, November 16, 2016.
97. *The MVAPICH2 Project: Latest Developments and Plans Towards Exascale Computing*, Mellanox Theater, SC '16, November 15, 2016.
98. *MVAPICH2 MPI Libraries to Exploit Latest Networking and Accelerator Technologies*, Naval Research Laboratory (NRL) Booth, SC '16, November 15, 2016.
99. *Exploiting HPC Technologies to Accelerate Big Data Processing (Hadoop, Spark, and Memcached)*, Intel HPC Developers Conference, November 13, 2016.
100. *High-Performance and Scalable MPI+X Library for Emerging HPC Clusters*, Intel HPC Developers Conference, November 12, 2016.
101. *MPI+X - The Right Programming Paradigm for Exascale?*, HP-CAST, November 12, 2016.

102. *Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Big Data Processing*, Int'l Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-8), October 29, 2016.
103. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Int'l Workshop on Co-Design, in conjunction with HPC-China Conference, October 27, 2016.
104. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Tsinghua University, China, October 25, 2016.
105. *High-Performance Broadcast with GPUDirect RDMA and InfiniBand Hardware Multicast for Streaming Applications*, DOD PETTT Annual TFA Review and Planning Meeting, Ann Arbor, October 18, 2016.
106. *Big Neuroscience Infrastructure*, ACNN Midwest Workshop on Big Neuroscience Data, Tools, Protocols & Services, Ann Arbor, September 20, 2016.
107. *Overview of OSU Network Based Computing*, ACNN Midwest Workshop on Big Neuroscience Data, Tools, Protocols & Services, Ann Arbor, September 20, 2016.
108. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Int'l Workshop on HPC Architecture, Software, and Application at an Extreme Scale, Guangzhou, China, September 18, 2016.
109. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Int'l Workshop on HPC Architecture, Software, and Application at an Extreme Scale, Wuxi, China, September 17, 2016.
110. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, National Tsinghua University, Taiwan, September 12, 2016.
111. *Designing HPC, Big Data and Deep Learning Middleware for Exascale Systems: Challenges and Opportunities*, Lawrence Livermore National Laboratory, August 25, 2016.
112. *Power-Performance Modeling of data Movement Operations on Next-Generation Systems with High-Bandwidth Memory*, DOE ModSim Workshop on Modeling and Simulation of Systems and Applications, August 10, 2016.
113. *Designing HPC and Big Data Middleware for Exascale Systems: Challenges and Opportunities*, NASA Langley Research Center, May 11, 2016.
114. *Building Efficient HPC Clouds with MVAPICH2 and OpenStack over SR-IOV enabled InfiniBand Clusters*, OpenStack Austin Summit, April 26, 2016.
115. *How to Tune and Extract Higher Performance with MVAPICH2 Libraries?*, XSEDE EECS Webinar, April 19, 2016.
116. *Designing HPC and Big Data Middleware for Exascale Systems: Challenges and Opportunities*, IBM Ireland, April 14, 2016.
117. *Accelerating Apache Hadoop through High-Performance Networking and I/O Technologies*, Apache Hadoop Summit Ireland, April 13, 2016.
118. *MVAPICH2-GDR: Pushing the Frontier of Designing MPI Libraries Enabling GPUDirect Technologies*, NVIDIA GTC '16 Conference, April 6, 2016.

119. *Bringing NVIDIA GPUs to the PGAS/OpenSHMEM World: Challenges and Solutions*, NVIDIA GTC '16 Conference, April 6, 2016.
120. *Coupling GPUDirect RDMA and InfiniBand Hardware Multicast Technologies for Streaming Applications*, NVIDIA GTC '16 Conference, April 6, 2016.
121. *Accelerating Big Data Processing (Hadoop, Spark and Memcached) on Modern HPC Clusters*, OpenFabrics Alliance 2016 Workshop, April 5, 2016.
122. *Designing MPI and PGAS Libraries for Exascale Systems: The MVAPICH2 Approach*, OpenFabrics Alliance 2016 Workshop, April 5, 2016.
123. *High-Performance MPI Library with SR-IOV and SLURM for Virtualized InfiniBand Clusters*, OpenFabrics Alliance 2016 Workshop, April 5, 2016.
124. *Exploiting HPC Technologies to Accelerate Big Data Processing*, Int'l Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-7), April 3, 2016.
125. *High-Performance and Scalable Designs of Programming Models for Exascale Systems*, HPC Advisory Council Switzerland Conference, March 21, 2016.
126. *Big Data Acceleration*, HPC Advisory Council Stanford Conference, February 25, 2016.
127. *Accelerating Big Data Applications using Hadoop, Spark and Memcached on Modern Clusters*, OSU TDA Forum, Jan 20, 2016.
128. *Exploiting Full Potential of GPU Clusters with InfiniBand using MVAPICH2-GDR*, Mellanox Booth Presentation (in conjunction with SC '15), Nov 19, 2015.
129. *How to Exploit MPI, PGAS and Hybrid MPI+PGAS Programming through MVAPICH2-X?*, OSC Booth Presentation (in conjunction with SC '15), Nov 19, 2015.
130. *MVAPICH2-GDR: Pushing the Frontier of Designing MPI Libraries Enabling GPUDirect Technologies*, NVIDIA Theater Presentation (in conjunction with SC '15), Nov 18, 2015.
131. *High Performance Big Data (HiBD): Accelerating Hadoop, Spark and Memcached on Modern Clusters*, Mellanox Booth Presentation (in conjunction with SC '15), Nov 18, 2015.
132. *The MVAPICH2 Project: Heading Towards New Horizons in Energy-Awareness, Virtualization and Network/Job-Level Introspection*, OSC Booth Presentation (in conjunction with SC '15), Nov 18, 2015.
133. *Update on MVAPICH2-X Project*, PGAS BOF (in conjunction with SC '15), Nov 18, 2015.
134. *The MVAPICH2 Project: Latest Developments and Plans Towards Exascale Computing*, Mellanox Booth Presentation (in conjunction with SC '15), Nov 18, 2015.
135. *HiBD - High-Performance Big Data*, SDSC Booth Presentation (in conjunction with SC '15), Nov 18, 2015.
136. *The MVAPICH2 Project: Pushing the Frontier of InfiniBand and RDMA Networking Technologies*, OSC Booth Presentation (in conjunction with SC '15), Nov 17, 2015.
137. *MVAPICH2 - High-Performance MPI and PGAS Libraries for Modern Clusters*, SDSC Booth Presentation (in conjunction with SC '15), Nov 17, 2015.

138. *An Assessment of MPI 3.x*, HP-CAST (in conjunction with SC '15), Nov 14, 2015.
139. *High Performance MPI Support for Clouds with IB and SR-IOV*, HP-CAST (in conjunction with SC '15), Nov 14, 2015.
140. *Challenges in Designing HPC and Big Data Middleware for Exascale Systems*, Co-Design Workshop, HPC China Conference, Nov 10, 2015.
141. *Designing Hybrid MPI+PGAS Library for Exascale Systems: MVAPICH2-X Experience*, LENS workshop, Japan, Oct 30, 2015.
142. *High-Performance Broadcast with GPUDirect RDMA and InfiniBand Hardware Multicast for Streaming Applications*, SAR-TIM, Wright Patterson Airforce Base, May 8, 2015.
143. *Accelerating Big Data Processing on Modern HPC Clusters*, IBM TJ Watson Research Center, April 29, 2015.
144. *Designing High Performance and Scalable MPI Library with InfiniBand-RDMA and Accelerators/Co-processors*, IBM TJ Watson Research Center, April 29, 2015.
145. *Accelerating Big Data Processing with Hadoop, Spark and Memcached*, HPC Advisory Council Switzerland Conference, March 25, 2015.
146. *Latest Advances in MVAPICH2 MPI Library for NVIDIA GPU Clusters with InfiniBand*, NVIDIA GTC Conference, March 19, 2015.
147. *High-Performance Broadcast with GPUDirect RDMA and InfiniBand Hardware Multicast for Streaming Applications*, NVIDIA GTC Conference, March 19, 2015.
148. *Enabling Efficient Use of UPC and OpenSHMEM PGAS Models on GPU Clusters*, NVIDIA GTC Conference, March 19, 2015.
149. *Accelerating Big Data Processing on Modern Clusters* Samsung, March 18, 2015.
150. *Challenges in Designing Communication Libraries and Middleware for Exascale Systems*, Bilkent University, Turkey March 13, 2015.
151. *Accelerating Data Management and Processing on Modern Clusters with RDMA-Enabled Interconnects* Bilkent University, Turkey March 13, 2015.
152. *Challenges in Designing Communication Libraries and Middleware for Exascale Systems*, Univ. of Illinois at Urbana-Champaign, February 12, 2015.
153. *Accelerating Big Data Processing with Hadoop, Spark and Memcached*, HPC Advisory Council Stanford Conference, February 3, 2015.
154. *Advances in MPI and PGAS Programming Models*, HPC Advisory Council South African Workshop, December 2, 2014.
155. *Latest Advances in MVAPICH2 MPI Library for NVIDIA GPU Clusters with InfiniBand*, Presentation at NVIDIA Theater, in conjunction with Supercomputing (SC '14), Nov 20, 2014.
156. *Accelerating Big Data Middleware (Hadoop, Spark and Memcached) with RDMA*, Presentation at SDSC Theater, in conjunction with Supercomputing (SC '14), Nov 19, 2014.

157. *High Performance MPI, PGAS and MPI + PGAS for InfiniBand Clusters*, Presentation at OSC Theater, in conjunction with Supercomputing (SC '14), Nov 19, 2014.
158. *MVAPICH2 and MVAPICH2-MIC: Designs and Experience on Stampede*, Presentation at TACC Theater, in conjunction with Supercomputing (SC '14), Nov 19, 2014.
159. *High Performance MPI, PGAS and Big Data (Hadoop and Spark) over InfiniBand*, Presentation at Mellanox Theater, in conjunction with Supercomputing (SC '14), Nov 18, 2014.
160. *High Performance MPI over InfiniBand*, Presentation at OSC Theater, in conjunction with Supercomputing (SC '14), Nov 18, 2014.
161. *High Performance MPI Support for Cloud with IB and SR-IOV*, HP-CAST, November 15, 2014.
162. *The Future of MPI - An Assessment*, HP-CAST, November 15, 2014.
163. *Designing Software Libraries and Middleware for Exascale Systems: Opportunities and Challenges*, Brookhaven National Laboratory, October 17, 2014.
164. *Designing Software Libraries and Middleware for Exascale Systems: Opportunities and Challenges*, Stony Brook University, October 16, 2014.
165. *Performance Optimization of Hadoop using InfiniBand RDMA*, ISC Big Data Conference, October 1, 2014.
166. *Overview and Latest Status of MVAPICH2 and High-performance Big Data (HiBD) Projects*, Lawrence Livermore National Laboratory, August 28, 2014.
167. *Power-Performance Modeling of Collectives for Exascale Systems with Heterogeneous Architectures*, DOE Workshop on Modeling and Simulation of Systems and Applications, August 13, 2014.
168. *Supporting PGAS Models (UPC and OpenSHMEM) on MIC Clusters*, Intel Phi User's Group Meeting (IXPUG), July 17, 2014.
169. *MVAPICH2 and MVAPICH2-MIC: Latest Status*, Intel Phi User's Group Meeting (IXPUG), July 16, 2014.
170. *MVAPICH2 and MVAPICH2-X: Latest Developments and Future Plans*, HPC Advisory Council European Workshop, June 22, 2014.
171. *Latest Advances in MVAPICH2 MPI Library for NVIDIA GPU Clusters with InfiniBand*, NVIDIA-Mellanox-OSU Webinar, June 5, 2014.
172. *MPI for Many-Core Processor*, Virtual School of Computational Science and Engineering (VSCSE) Webinar, With Dr. Karen Tomko, May 8, 2014.
173. *MPI/PGAS Hybrid Programming*, Virtual School of Computational Science and Engineering (VSCSE) Webinar, With Dr. Karen Tomko, May 7, 2014.
174. *Advanced MPI Capabilities - MPI-3 Capabilities*, Virtual School of Computational Science and Engineering (VSCSE) Webinar, With Dr. Karen Tomko, May 6, 2014.
175. *Big Data: Hadoop and Memcached*, HPC Advisory Council Switzerland Conference, April 1, 2014.



176. *Accelerating HPL on Heterogeneous GPU Clusters*, NVIDIA GTC Conference, March 25, 2014.
177. *Latest Advances in MVAPICH2 MPI Library for NVIDIA GPU Clusters with Infini-Band*, NVIDIA GTC Conference, March 25, 2014.
178. *Enabling Efficient Use of UPC and OpenSHMEM PGAS Models on GPU Clusters*, NVIDIA GTC Conference, March 25, 2014.
179. *Designing Software Libraries and Middleware for Exascale Systems: Opportunities and Challenges*, University of Oslo, Norway, Feb 10, 2014.
180. *Accelerating Big Data Processing on Modern HPC Clusters: Opportunities and Challenges*, TACTiCS Symposium, Tata Consultancy Service, Mumbai, India, Dec 16, 2013.
181. *Accelerating Big Data Processing on Modern HPC Clusters: Opportunities and Challenges*, Huawei, Beijing, China, Dec 5, 2013.
182. *MVAPICH2-GDR: Optimized MPI Communication using GPUDirect RDMA*, Presentation at NVIDIA Booth, in conjunction with Supercomputing (SC '13), Nov 21, 2013.
183. *RDMA for Apache Hadoop: High Performance and Scalable Hadoop using RDMA for Big Data Analytics*, Presentation at Mellanox Booth, in conjunction with Supercomputing (SC '13), Nov 21, 2013.
184. *Overview of MVAPICH2 and RDMA for Apache Hadoop Projects*, Presentation at SDSC Booth, in conjunction with Supercomputing (SC '13), Nov 20, 2013.
185. *MVAPICH2-X: Hybrid MPI+PGAS (UPC and OpenSHMEM) Support in MVAPICH2*, Presentation at Mellanox Booth, in conjunction with Supercomputing (SC '13), Nov 20, 2013.
186. *High Performance Runtime for PGAS and Hybrid MPI+PGAS*, PGAS Birds of Feather, in conjunction with Supercomputing (SC '13), Nov 20, 2013.
187. *MVAPICH2 and MVAPICH2-MIC Designs and Experience on Stampede*, Presentation at TACC Booth, in conjunction with Supercomputing (SC '13), Nov 20, 2013.
188. *High Performance, Scalable, and Fault-Resilient MPI and GPUDirect RDMA*, Presentation at Mellanox Booth, in conjunction with Supercomputing (SC '13), Nov 19, 2013.
189. *MVAPICH2 MPI Library for HPC and A New Hadoop-RDMA Project for Big Data Processing*, Lawrence Livermore National Laboratory, August 23, 2013.
190. *MVAPICH2-MIC: A High-performance MPI Library for Xeon Phi Clusters with Infini-Band*, Extreme Scaling Workshop, August 15, 2013.
191. *MVAPICH2-GDR: MVAPICH2 with GPUDirect RDMA*, NVIDIA, Mellanox and OSU Webinar, August 6, 2013.
192. *MVAPICH2-MIC: A High-performance MPI Library for Xeon Phi Clusters*, jointly with Sreeram Potluri, MIC User Experience Workshop, Texas Advanced Computing Center, July 17, 2013.
193. *MVAPICH2 and GPUDirect RDMA*, jointly with Hari Subramoni and Sreeram Potluri, HPC Advisory Council European Conference, June 16, 2013.
194. *SSD-Assisted Designs for MPI Fault-Tolerance and Accelerating Web 2.0*, OpenFabrics International Developers Workshop, April 23, 2013.

195. *High Performance RDMA based Design for Big Data and Web 2.0 Memcached*, OpenFabrics International Developers Workshop, April 23, 2013.
196. *MVAPICH2 for Intel MIC*, OpenFabrics International Developers Workshop, April 22, 2013.
197. *MVAPICH2: A High Performance MPI Library for NVIDIA GPU Clusters with InfiniBand*, NVIDIA GPU Technology Conference (GTC), Mar 20, 2013.
198. *Accelerating Big Data with Hadoop (HDFS, MapReduce and HBase) and Memcached*, HPC Advisory Council Switzerland Conference, March 14, 2013.
199. *Programming Models and their Designs for Exascale Systems (Part II): Accelerators, Coprocessors, QoS and Fault-Tolerance*, HPC Advisory Council Switzerland Conference, March 14, 2013.
200. *Programming Models and their Designs for Exascale Systems (Part I): MPI, PGAS and Hybrid*, HPC Advisory Council Switzerland Conference, March 13, 2013.
201. *Designing Software Libraries and Middleware for Exascale Systems: Opportunities and Challenges*, Univ. of Cincinnati, March 4, 2013.
202. *MPI Communication on Stampede with MIC using MVAPICH2: Early Experience*, SIAM CSE 2013, Feb 28, 2013.
203. *InfiniBand for Next Generation HPC Exascale Needs and Challenges to be Solved*, Tata Consultancy Service (TCS), Pune, India, Dec 21, 2012.
204. *High-Performance and Scalable Processing of Big Data using Hadoop on Modern Clusters*, Big Data Workshop, Ohio State University, Dec 5, 2012.
205. *High Performance MPI, Hybrid MPI+PGAS (UPC and OpenSHMEM) and BigData-Hadoop over InfiniBand* Presentation at Mellanox Booth, in conjunction with Supercomputing (SC '12), Nov 14, 2012.
206. *HPC Cloud: Can Infrastructure Clouds Provide a Viable Platform for HPC?* Presentation at BOF on HPC Cloud, in conjunction with Supercomputing (SC '12), Nov 13, 2012.
207. *High Performance Computing Programming Techniques For Big Data Hadoop*, Presentation at BOF on Big Data, in conjunction with Supercomputing (SC '12), Nov 13, 2012.
208. *Supporting High-Performance and Scalable MPI and PGAS Programming Models for Exascale Computing: An MVAPICH2-X Approach*, HPC in China Workshop, held in conjunction with Supercomputing '12, Nov. 13, 2012.
209. *Designing High Performance, Scalable and Fault-Resilient MPI Library for Modern Clusters*, Lawrence Livermore National Laboratory, August 23, 2012.
210. *Can We Take Advantage of InfiniBand and RDMA for Hadoop?*, IBM Almaden Research Center, August 22, 2012.
211. *Challenges in Designing Communication Libraries and Middleware for Exascale Systems*, Distinguished Talk, Georgia Institute of Technology, April 25, 2012.
212. *RDMA for Heterogeneous Parallel Computing*, OpenFabrics Monterey Workshop, March 28, 2012.

213. *Acceleration for Big Data, Hadoop and Memcached*, OpenFabrics Monterey Workshop, March 27, 2012.
214. *Experience with SSDs*, OpenFabrics Monterey Workshop, OpenFabrics Monterey Workshop, March 26, 2012.
215. *Unified MPI and PGAS (UPC and OpenShmem) Design with RDMA to Support Hybrid Programming Environment for Exascale Systems*, OpenFabrics Monterey Workshop, March 26, 2012.
216. *MVAPICH/MVAPICH2: Latest Status and Future Plans*, OpenFabrics Monterey Workshop, March 26, 2012.
217. *Accelerations for Big Data, Hadoop and Memcached*, HPC Advisory Council Lugano Workshop, March 15, 2012.
218. *MPI: Advanced Topics and Future Trends*, HPC Advisory Council Lugano Workshop, March 14, 2012.
219. *MPI: Overview, Performance Optimization and Tuning*, HPC Advisory Council Lugano Workshop, March 13, 2012.
220. *Challenges in Designing High Performance, Scalable and Fault-Tolerant MPI*, HPC Advisory Council Stanford Workshop, Dec. 8, 2011.
221. *High Performance MPI Support in MVAPICH2 for InfiniBand Clusters*, Presentation at NVIDIA Booth, Supercomputing Conference (SC '11), Nov. 15, 2011.
222. *MVAPICH2 Project: Latest Status and Future Plans*, Presentation at BOF on MPICH, in conjunction with Supercomputing (SC '11), Nov 15, 2011.
223. *MVAPICH/MVAPICH2 Project: Latest Status and Future Plans*, Presentation at Mellanox Booth, Supercomputing Conference (SC '11), Nov. 15, 2011.
224. *High Performance MPI Support for NVIDIA GPU Accelerator: MVAPICH2 Approach*, Presentation at BOF on Accelerators, Supercomputing Conference (SC '11), Nov. 15, 2011.
225. *Designing High Performance, Scalable and Fault-Resilient MPI Library for Modern Clusters*, National University of Defense Technology (NUDT), Changsha, China, October 24, 2011.
226. *High Performance and Scalable Checkpoint-Restart and Migration on InfiniBand Clusters with Staging*, ADIOS Meeting, OakRidge National Laboratory, September 12, 2011.
227. *Can High-Performance Interconnects Benefit High-Performance Computing*, Oracle, August 26, 2011.
228. *MVAPICH and Futures of MPI*, A series of three talks, jointly with S. Sur, HPC Advisory European Council Workshop, Hamburg, June 19, 2011.
229. *MVAPICH2-GPU: Optimized GPU to GPU Communication for Infiniband Clusters*, jointly with S. Sur, Int'l Workshop on Future Approaches to Data Centric Programming for Exascale (DCPM), in conjunction with IPDPS '11, May 20, 2011.
230. *Unified Communication Runtime for MPI and PGAS (UPC): Experience with MVAPICH2*, Int'l Workshop on Future Approaches to Data Centric Programming for Exascale (DCPM), in conjunction with IPDPS '11, May 20, 2011.

231. *Can HPC Interconnects Benefit MemCached and Hadoop*, jointly with S. Sur, Int'l Open-Fabrics Workshop, Monterey, April 5, 2011.
232. *Unified Runtime for PGAS and MPI over OFED*, jointly with S. Sur, Int'l OpenFabrics Workshop, Monterey, April 4, 2011.
233. *Futures of MPI*, jointly with S. Sur, HPC Advisory Council Workshop, Switzerland, March 23, 2011.
234. *MPI Performance and Optimizations*, jointly with S. Sur, HPC Advisory Council Workshop, Switzerland, March 22, 2011.
235. *Introduction to MPI*, jointly with S. Sur, HPC Advisory Council Workshop, Switzerland, March 21, 2011.
236. *Designing Scalable and High Performance MPI and PGAS (UPC) Libraries for Multi-Petaflop and Exaflop Systems*, jointly with S. Sur, AMD, Mar 11, 2011.
237. *Virtualization for HPC*, Presentation at Disruptive Technology Booth, Supercomputing (SC '10), Nov 18, 2010.
238. *MVAPICH2 Project: Latest Status and Future Plans*, BOF on MPICH2, in conjunction with Supercomputing (SC '10), Nov 17, 2010.
239. *MVAPICH/MVAPICH2 Project: Latest Status and Future Plans*, Presentation at Mellanox Booth, Supercomputing Conference (SC '10), Nov. 17, 2010.
240. *Developing FTB-Enabled InfiniBand Monitoring Framework and MVAPICH2*, BOF on CIFTS, in conjunction with Supercomputing (SC '10), Nov 16, 2010.
241. *Designing High Performance, Scalable and Fault-Resilient MPI Library for Modern Cluster*, jointly with S. Sur, Presentation at Ohio Supercomputer Center Booth, Supercomputing (SC '10), Nov 15, 2010.
242. *Designing High Performance, Scalable and Fault-Resilient MPI Library for Modern Clusters*, Institute of Computing Technology (ICT), Chinese Academy of Sciences, October 26, 2010.
243. *Networking Technologies for Clusters: Where do We Stand and What Lies Ahead?*, Indiana University - Purdue University Indianapolis, Oct 01, 2010.
244. *Co-Designing MPI Library and Applications for InfiniBand Clusters*, Workshop on Application/Architecture Co-Design for Extreme-Scale Computing (AAEC), in conjunction with Cluster '10 conference, Sept 24, 2010.
245. *MVAPICH/MVAPICH2 Update, Future Plans and Path Towards Exascale*, Open Fabrics Sonoma Workshop, March 16, 2010.
246. *Center for Performance Evaluation of Cluster Networking and I/O Technologies (PEC-NIT): Overview and Performance Results*, Avetec, March 5, 2010.
247. *Designing Fault Resilient and Fault Tolerant Systems with InfiniBand*, SIAM Conference on Parallel Computing, February 24, 2010.
248. *Designing High Performance, Scalable and Fault-Resilient MPI Library for Modern Clusters*, Pacific Northwest National Library (PNNL), February 23, 2010.
249. *New Research Areas in HPC Technologies Development*, Center for Development of Advanced Computing (C-DAC), Pune, India, January 8, 2010.

250. *Evolution of HPC Interconnects in next 5 years and their Role in Peta/Exascale Systems*, Center for Development of Advanced Computing (C-DAC), Pune, India, January 8, 2010.
251. *MVAPICH2 Project: Latest Status and Future Plans*, BOF on MPICH2, in conjunction with Supercomputing (SC '09), Nov 19, 2009.
252. *MVAPICH/MVAPICH2 Project: Latest Status and Future Plans*, Presentation at Mellanox Booth, Supercomputing Conference (SC '09), Nov. 17, 2009.
253. *FTB Framework for InfiniBand*, BOF on CIFTs, in conjunction with Supercomputing (SC '09), Nov 17, 2009.
254. *Designing Fault Resilient and Fault Tolerant Systems with InfiniBand*, HPC Resiliency Workshop, October 14, 2009.
255. *Designing Fault Resilient and Fault Tolerant Systems with InfiniBand*, jointly with Abhinav Vishnu (PNNL) and K. Gopalakrishnan (OSU), Invited Poster Presentation at 2009 National Workshop on Resiliency, August 12, 2009.
256. *Staging Capabilities with High Performance Networking (InfiniBand)*, Oak Ridge ADIOS Meeting, Oak Ridge, July 8, 2009.
257. *Supporting Fault-Tolerance in Modern High-End Computing Systems with InfiniBand*, Dagstuhl Seminar on Fault-Tolerance in High Performance Computing, Dagstuhl, Germany, May 21, 2009.
258. *MVAPICH/MVAPICH2 Update*, Open Fabrics Sonoma Workshop, March 24, 2009.
259. *Designing Next Generation Clusters with InfiniBand and 10GE/iWARP*, Intel, Bangalore, India, Dec. 19, 2008.
260. *Research on Network-Based Computing*, TATA CRL Laboratory, India, Dec. 18, 2008.
261. *Overview of MVAPICH/MVAPICH2 Project*, TATA CRL Laboratory, India, Dec. 18, 2008.
262. *MVAPICH2 Project: Latest Status and Future Plans*, BOF on MPICH2, in conjunction with Supercomputing (SC '08), Nov 20, 2008.
263. *Extending One-Sided Communication in MPI Programming Model for Next Generation HEC*, BOF on HECURA, in conjunction with Supercomputing (SC '08), Nov 18, 2008.
264. *FTB Framework for InfiniBand*, BOF on CIFTs, in conjunction with Supercomputing (SC '08), Nov 18, 2008.
265. *MVAPICH/MVAPICH2 Project: Latest Status and Future Plans*, Presentation at Mellanox Booth, Supercomputing Conference (SC '08), Nov. 18, 2008.
266. *High Performance, Scalable and Fault-Tolerant MPI over InfiniBand: An Overview of MVAPICH/MVAPICH2 Project*, Tsukuba University, Japan, Oct. 2, 2008.
267. *Center for Performance Evaluation of Cluster Networking and I/O Technologies (PEC-NIT): Overview and Performance Results*, DICE Alliance Workshop, May 7, 2008.
268. *MVAPICH/MVAPICH2 Update*, Open Fabrics Sonoma Workshop, April 08, 2008.
269. *Trends in MPI Libraries and Potential Benefits from the SIF Architecture*, jointly with W. Huang and G. Santhanaraman, SUN SIF Workshop, March 10, 2008.

270. *Emerging Networking Technologies and Protocols for Next Generation Clusters*, Florida State University, Nov. 26, 2007.
271. *MVAPICH with XRC: Early Experiences*, Open Fabrics Developers Summit, Nov 16, 2007.
272. *MVAPICH/MVAPICH2 Update*, Open Fabrics Developers Summit, Nov 15, 2007.
273. *MVAPICH2 Project: Latest Status and Future Plans*, BOF on MPICH2, in conjunction with Supercomputing (SC '07), Nov 15, 2007.
274. *FTB Framework for InfiniBand*, BOF on CIFTs, in conjunction with Supercomputing (SC '07), Nov 13, 2007.
275. *Handling Reliability at the MPI layer*, BOF on Reliability of High-Speed Networks, in conjunction with Supercomputing (SC '07), Nov 13, 2007.
276. *Center for Performance Evaluation of Cluster Networking and I/O Technologies (PEC-NIT): Overview and Performance Results*, DICE/SC '07 Event, Nov 13, 2007.
277. *Challenges in Scaling to Large Fabrics for Making Efficient Communication Libraries*, SUN Microsystems, Norway, Nov 1, 2007.
278. *Performance Evaluation of Intel Connects Cable for InfiniBand*, Intel HPC Roundtable Conference, Washington, DC, Oct 9, 2007.
279. *Center for Performance Evaluation of Cluster Networking and I/O Technologies (PEC-NIT): Overview and Performance Results*, DICE Alliance Workshop, May 8, 2007.
280. *Designing Next Generation Clusters and Data-Centers with InfiniBand and iWARP*, HP, Palo Alto, May 2, 2007.
281. *HPC with Virtual Machines: Experiences with Xen InfiniBand and MPI*, OpenFabrics Sonoma Workshop, May 1, 2007.
282. *MVAPICH/MVAPICH2 Update*, OpenFabrics Sonoma Workshop, April 30, 2007.
283. *Emerging Networking Technologies and Protocols for Next Generation Clusters*, Kent State University, March 7, 2007.
284. *Designing Multi-core Aware Middleware for HPC and Data-Centers*, Multi-cores/Many-cores Workshop, March 2, 2007.
285. *OSU MPI (MVAPICH and MVAPICH2): Latest Status, Performance Numbers and Future Plans*, OpenFabrics Developers Summit, November 17, 2006.
286. *Benchmarking for Clusters with High Performance Networks*, Wright Patterson Airforce Base, October 15, 2006.
287. *Memory Architecture for High Performance Computing Optimization*, Intel Developers Forum, September 26, 2006.
288. *HPC with Virtual Machines: Experiences with Xen and InfiniBand*, Virtualization Workshop, organized by Oak Ridge National Laboratory, September 20, 2006.
289. *An InfiniBand Message Passing Model for Multicore Clusters*, Interactive Cinema Summit, held in conjunction with SigGraph, August 2, 2006.
290. *Optimizing InfiniBand on Bensley Platforms*, Intel Developers Forum, March 08, 2006.

291. *OSU MPI (MVAPICH and MVAPICH2): Latest Status, Performance Numbers and Future Plans*, OpenIB Workshop, February 07, 2006.
292. *Designing High Performance DSM Support over InfiniBand*, Intel Tele-Seminar, Dec 08, 2005.
293. *InfiniBand: Where is it Headed?*, Ohio Supercomputing Center Statewide Users Group, Nov. 19, 2005.
294. *OSU MPI over InfiniBand (MVAPICH): Latest Status, Performance Numbers and Future Plans*, Datacenter Fabric Workshop, August 22, 2005.
295. *Network-Centric Designs for High-End I/O and File Systems*, IHEC-IWG File Systems and I/O R&D Workshop, August 16, 2005.
296. *High Performance MPI for InfiniBand Cluster*, Central State University, Aug. 4, 2005.
297. *Designing Next Generation High Performance Clusters and Datacenters with InfiniBand*, Microsoft, July 18, 2005.
298. *Designing High Performance MPI-1 and MPI-2 on RDMA-Enabled Interconnects*, Intel Tele-Seminar, May 17, 2005.
299. *InfiniBand: An Emerging Networking Technology for Designing Next Generation High Performance Clusters, File/Storage Systems and Datacenters*, Central Ohio ACM Chapter, April 20, 2005.
300. *Designing Next Generation High Performance Clusters and Datacenters with InfiniBand*, Oak Ridge National Laboratory, March 18, 2005.
301. *Emerging Clusters and Data-Centers with InfiniBand Networking for Supporting Next Generation Data-Intensive Applications*, Imaging Workshop, Columbus, Feb 16, 2005.
302. *OSU MPI (MVAPICH) over InfiniBand: Latest status, Performance Numbers, and Results*, OpenIB Workshop, Feb 7, 2005.
303. *InfiniBand: A Modern Networking Technology*, WIPRO Information Technology, Dec 20, 2004.
304. *Designing Next Generation High Performance Clusters and Datacenters with InfiniBand*, Livermore National Laboratory, Dec. 6, 2004.
305. *Designing High Performance Distributed Data-Centers for Security Applications*, National Security and Defense-Related Research Meeting, Columbus, Oct 7, 2004.
306. *Designing Next Generation High Performance Clusters and Datacenters with InfiniBand*, IBM TJ Watson Research Center, Sept. 9, 2004.
307. *High Performance MPI for InfiniBand-based Clusters*, Central State University, Aug. 11, 2004.
308. *Designing Next Generation High Performance Clusters and Datacenters with InfiniBand*, ARL Cluster Symposium, July 28, 2004.
309. *MPI, File Systems, Storage, and Datacenter Research with InfiniBand: Latest Research Results and Plans*, Second annual DOE InfiniBand Workshop, January 20, 2004.
310. *Designing Next Generation High Performance Clusters with InfiniBand*, ARL Cluster Symposium, July 22, 2003.

311. *InfiniBand: Is it Ready for Prime-Time High Performance Computing?*, Univ. of Cincinnati, May 7, 2003.
312. *InfiniBand: Is it Ready for Prime-Time High Performance Computing?*, Los Alamos National Laboratory, Apr 8, 2003.
313. *Efficient Support for Programming Models on Modern Interconnects*, Intel KAI Lab, Urbana, Oct 16, 2002.
314. *Active Network Interface Support for Efficient Communication, Synchronization, and QoS in Clusters*, University of Texas, Austin, Aug 29, 2002.
315. *Designing High Performance Communication Subsystems for Clusters*, HP Labs, California, Aug 21, 2002.
316. *Cluster Computing: Where are We Headed?*, Fermi National Laboratory, June 4, 2002.
317. *Active Network Interface Support for Efficient Communication and Synchronization*, DOE Workshop on Communication and Middleware for Parallel Programming Models, April 14, 2002.
318. *Network-Based Computing and InfiniBand*, Paceline Systems, Cambridge, MA, Feb 5, 2002.
319. *Designing High Performance Communication Subsystems for Clusters with VIA and InfiniBand*, IBM Almaden, August 24, 2001.
320. *Designing High Performance Communication Subsystems for Clusters*, Fujitsu Research Laboratory, San Jose, August 23, 2001.
321. *End-to-end QoS Support for Collaborative Networked Computing*, Windows on the Future Conference, Internet 2 Technology Evaluation Center (ITEC), Columbus, March 30, 2001.
322. *Designing High Performance Communication Subsystems for Clusters*, Argonne National Laboratories, Feb 8, 2001.
323. *Network-Based Computing Research at OSU*, Sandia National Laboratories, Dec 12, 2000.
324. *Network-Based Computing: Issues, Trends, and Challenges* ACM Chapter, Columbus, Nov. 21, 2000.
325. *Supporting Efficient Communication in Network-Based Computing Systems*, Mississippi State University, Oct 2, 2000.
326. *Experiences on Implementing VIA and Higher Layers on Modern Interconnects*, Intel Corporation, Portland, July 28, 2000.
327. *Cluster Computing: Where are We Headed?*, Pacific Northwest National Laboratory, Pasco, July 27, 2000.
328. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), IEEE Calumet Section, Chicago, March 21, 2000.
329. *Experiences with Network-Based Computing Research and Beowulf Cluster at OSU-CIS*, Beowulf Focus Group Meeting, Ohio Supercomputer Center, Columbus, Sept. 28, 1999.



330. *Trends in Parallel Architecture and Communication*, Computational Chemistry Workshop IV, Columbus, June 18, 1999.
331. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), University of Nevada, Reno, April 1, 1999.
332. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), Drexel University, March 18, 1999.
333. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), IEEE Computer Society Chapter, Dayton, July 16, 1998.
334. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), IEEE Computer Society Chapter, Dallas, May 28, 1998.
335. *Network-Based Computing with OCARNET*, OCARNET State-level Meeting, Columbus, Jan 30, 1998.
336. *Challenges in Designing Communication Subsystems for High-Performance Computing Systems*, Univ. of Nebraska, Omaha, Jan 22, 1998.
337. *Network-Based Computing: Issues, Trends, and Challenges* (Under the IEEE DVP Program), Univ. of Nebraska, Lincoln, Jan 21, 1998.
338. *Networks of Workstations (NOW): An Emerging Trend for Cost-Effective Parallel Computing*, Center for Materials Research (CMR) Faculty Seminar Series, Ohio State University, May 29, 1997.
339. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Univ. of Minnesota, Minneapolis, August 17, 1996.
340. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Georgia Institute of Technology, Atlanta, May 17, 1996.
341. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Univ. of Texas, San Antonio, April 2, 1996.
342. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Texas A&M University, April 1, 1996.
343. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Univ. of Southern California, Los Angeles, November 30, 1995.
344. *Fast and Scalable Collective Communication in Wormhole Systems: Unicast or Multidestination Message-Passing?*, Univ. of Illinois, Urbana-Champaign, November 14, 1995.
345. *Can We do Better than Unicast Message-Passing in Wormhole-Routed Parallel Systems?*, Pennsylvania State University, November 2, 1995.
346. *Can We do Better than Unicast Message-Passing in Wormhole-Routed Parallel Systems?*, Univ. of Wisconsin, Madison, August 17, 1995.
347. *Designing Scalable Interconnection for Parallel Systems with Multidestination Wormhole Message Passing*, University of California, Santa Cruz, Apr. 24, 1995.
348. *Designing Scalable Interconnection for Parallel Systems with Multidestination Wormhole Message Passing*, Wayne State University, Detroit, Mar. 22, 1995.

349. *Designing Scalable Interconnection for Parallel Systems with Multidestination Wormhole Message Passing*, Univ. of North Carolina, Raleigh, Jan. 25, 1995.
350. *Designing Scalable Interconnection for Parallel Systems with Multidestination Wormhole Message Passing*, New Jersey Institute of Technology, New Jersey, Sept. 13, 1994.
351. *Designing Scalable Interconnection for Parallel Systems with Multidestination Wormhole Message Passing*, IBM T. J. Watson Research Center, New York, Sept. 12, 1994.
352. *Trends in Parallel Architectures*, WIPRO Information Technology Ltd., Bangalore, India, Dec. 26, 1994.
353. *Medium-Grain Heterogeneous Parallel Processing Over Gigabit ATM Interconnection*, ARPA/NSF Gigabit Workshop, Washington, DC, July 20, 1994.
354. *Communication and Synchronization in Wormhole-routed Direct Networks*, University of Southern California, April 16, 1993.

## Publications

### A. Journal Articles

1. J. Hashmi, C. Chu, S. Chakraborty, M. Bayatpour, H. Subramoni, and D. K. Panda, FALCON-X: Zero-copy MPI Derived Datatype Processing on Modern CPU and GPU Architectures, *Journal of Parallel and Distributed Computing (JPDC)*, May 2020.
2. 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.
3. 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.
4. 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.
5. C. Chu, X. Lu, A. Awan, H. Subramoni, B. Elton, and *D. K. Panda*, Exploiting Hardware Multicast and GPUDirect RDMA for Efficient Broadcast, *IEEE Transactions on Parallel and Distributed Systems*, March 2019.
6. *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.
7. 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.
8. 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.

9. X. Lu, H. Shi, R. Biswas, H. Javed, and D. K. Panda, DLoBD: A Comprehensive Study of Deep Learning over Big Data Stacks on HPC Clusters, *IEEE Transactions on Multi-Scale Computing Systems*, Jun 2018.
10. M. W. Rahman, N. Islam, X. Lu, and D. K. Panda, A Comprehensive Study of MapReduce over Lustre for Intermediate Data Placement and Shuffle Strategies on HPC Clusters, *IEEE Transactions on Parallel and Distributed Systems*, Volume 28, Issue 3, March 2017.
11. D. Shankar, X. Lu and D. K. Panda, Characterizing and Benchmarking Stand-alone Hadoop MapReduce on Modern HPC Clusters, *Journal of Supercomputing*, December 2016, Volume 72, Issue 12, pp. 4573-4600.
12. 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, *Parallel Computing Journal*, Volume 58, October 2016, pp. 27-36.
13. H. Wang, S. Potluri, D. Bureddy, C. Rosales and D. K. Panda, GPU-Aware MPI on RDMA-Enabled Cluster: Design, Implementation and Evaluation, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 25, No. 10, pp. 2595-2605, Oct 2014.
14. N. Islam, X. Lu, M.-W. Rahman, J. Jose and D. K. Panda, A Micro-Benchmark Suite for Evaluating HDFS Operations on Modern Clusters, *Special Issue of LNCS on papers from WBDB '12 Workshop*, Vol. 8163, pp. 129-147, January 2014.
15. S. Sur, S. Potluri, K. Kandalla, H. Subramoni, D. K. Panda and Karen Tomko, Codesign for InfiniBand Clusters, *IEEE Computer*, Vol. 44, No. 11, pp. 31-36, Nov. 2011.
16. P. Lai, P. Balaji, R. Thakur, and D. K. Panda, ProOnE: A General-Purpose Protocol Onload Engine for Multi- and Many-Core Architectures, *Computer Science: Research and Development*, Special Issue of Scientific Papers from ISC '09, 2009.
17. A. Vishnu, M. Koop, A. Moody, A. Mamidala, S. Narravula, and D. K. Panda, Topology Agnostic Hot-Spot Avoidance with InfiniBand, *Concurrency and Computation: Practice and Experience*, Special Issue of Best Papers from CCGrid '07, Vol. 21, Issue 3, March 2009.
18. P. Balaji, W. Feng and D. K. Panda, The Convergence of Ethernet and Ethernet: A 10-Gigabit Ethernet Perspective, *IEEE Micro Special Issue on High Performance Interconnect*, May/June 2006, Vol. 26, No. 3., pp. 24-40.
19. F. Petrini, A. Moody, J. Fernandez, E. Frachtenberg, and D. K. Panda, NIC-based Reduction Algorithms for Large-scale Clusters, *Int'l Journal on High Performance Computing and Networking (IJHPCN)*, special issue on Design and Performance Evaluation of Group Communication, 2007.
20. H.-W. Jin, P. Balaji, C. Yoo, J. Y. Choi, and D. K. Panda, Exploiting NIC Architectural Support for Enhancing IP based Protocols on High Performance Networks, *Journal of Parallel and Distributed Computing (JPDC)*, Special issue on Design and Performance of Networks for Super-, Cluster-, and Grid-Computing Part II, Vol. 65, No. 11, pp. 1348-1365, Nov. 2005.

21. R. Kettimuthu, V. Subramani, S. Srinivasan, T. Gopalsamy, D. K. Panda, and P. Sadayappan, Selective Preemption Strategies for Parallel Job Scheduling, *Int'l Journal on High Performance Computing and Networking* (IJHPCN), Vol. 3, Nos. 2/3, 2005, pp. 122-152.
22. G. Santhanaraman, J. Wu, W. Huang, and D. K. Panda, Zero-Copy MPI Derived Datatype Communication over InfiniBand: Alternative Approaches and Performance Evaluation, Special issue of the International Journal of High Performance Computing Applications (IJHPCA) with Best Papers of EuroPVM/MPI 2004, Volume 19, Number 2, Summer 2005, pp. 129-142.
23. W. Yu, D. K. Panda, R. T. Aulwes, and R. Graham, High Performance Broadcast Support in LA-MPI over Quadrics, *Journal of High Performance Computing Applications* (JHPCA), special issue of selected papers from LACSI '03 conference, Volume 19, Winter 2005, pp. 453-463.
24. J. Liu, A. Mamidala, A. Vishnu, and D. K. Panda, Performance Evaluation of InfiniBand with PCI Express, Special issue of *IEEE Micro*, Best Papers of Hot Interconnect '04 conference, January/February 2005, Vol. 25, No. 1, pp. 20-29.
25. J. Nieplocha, V. Tipparaju, M. Krishnan, and D. K. Panda, High Performance Remote Memory Access Communication: The ARMCi Approach, *Journal of High Performance Computing Applications* (JHPCA), special issue, Fall 2004.
26. A. Wagner, D. Buntinas, R. Brightwell, and D. K. Panda, Application-Bypass Reduction for Large-Scale Clusters, *Int'l Journal on High Performance Computing and Networking* (IJHPCN), special issue of selected papers from Cluster '03 conference, Vol. 2, Issue 2-4, 2004, pp. 99-109.
27. V. Tipparaju, M. Krishnan, J. Nieplocha, G. Santhanaraman, and D. K. Panda, Optimization and Performance Evaluation of Mechanisms for Latency Tolerance in Remote Memory Access Communication on Clusters, *Int'l Journal on High Performance Computing and Networking* (IJHPCN), special issue of selected papers from Cluster '03 conference, Vol. 2, Issue 2-4, 2004, pp. 198-204.
28. J. Liu, J. Wu, and D. K. Panda, High Performance RDMA-Based MPI Implementation over InfiniBand, *International Journal of Parallel Programming*, special issue with selected papers from ICS '03 conference, Feb 2004.
29. J. Liu, B. Chandrasekharan, W. Yu, J. Wu, D. Buntinas, S. Kini, P. Wyckoff, and D. K. Panda, Microbenchmark Performance Comparison of High-Speed Cluster Interconnects, *IEEE Micro*, Vol. 24, No. 1, Jan-Feb 2004, pp. 42-51.
30. M. Islam, P. Balaji, P. Sadayappan and D. K. Panda, QoPS: A QoS based scheme for Parallel Job Scheduling, *IEEE Springer LNCS Journal Series*, 2003.
31. R. Sivaram, C. Stunkel, and D. K. Panda, HIPIQS: A High-Performance Switch Architecture using Input Queuing, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 13, No. 3, March 2002, pp. 275-289.

32. N. Sundar, J. N. Jayasimha, D. K. Panda, and P. Sadayappan, Hybrid Algorithms for Complete Exchange in 2D Meshes, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 12, No. 12, December 2001, pp. 1201-1218.
33. M. Banikazemi, B. Abali, L. Herger, and D. K. Panda, Design Alternatives for Virtual Interface Architecture (VIA) and an Implementation on IBM Netfinity NT Cluster, *Journal of Parallel and Distributed Computing*, special issue on cluster computing, Vol. 61, No. 11, November 2001, pp. 1512-1545.
34. M. Banikazemi, R. Govindaraju, R. Blackmore, and D. K. Panda, Implementing Efficient MPI on LAPI for the IBM SP System, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 12, No. 10, October 2001, pp. 1081-1093.
35. B. Abali, C. B. Stunkel, J. Herring, M. Banikazemi, D. K. Panda, C. Ayakanat, and Y. Ayadogan, Adaptive Routing on the New Switch Chip for IBM SP Systems, *Journal of Parallel and Distributed Computing*, special issue on routing in Computer and Communication Networks, Volume 61, Number 9, September 2001, pp. 1148-1179.
36. R. Kesavan and D. K. Panda, Efficient Multicast on Irregular Switch-based Cut-Through Networks with UP-Down Routing, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 12, No. 8, August 2001, pp. 808-828.
37. R. Sivaram, R. Kesavan, C. Stunkel, and D. K. Panda, Architectural Support for Efficient Multicasting in Irregular Networks, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 12, No. 5, May 2001, pp. 489-513.
38. R. Sivaram, C. B. Stunkel, and D. K. Panda, Implementing Multidestination Worms in Switch-Based Parallel Systems: Architectural Alternatives and their Impact, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 11, No. 8, August 2000, pp. 794-812.
39. R. Kesavan and D. K. Panda, Multiple Multicast with Minimized Node Contention on Wormhole k-ary n-cube Networks, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 10, No. 4, April 1999, pp. 371-393.
40. D. Dai and D. K. Panda, Exploiting the Benefits of Multiple-Path Network in DSM Systems: Architectural Alternatives and Performance Evaluation, *IEEE Transactions on Computers, Special Issue on Cache Memory*, Vol. 48, No. 2, Feb. 1999, pp. 236-244.
41. D. K. Panda, S. Singal, and R. Kesavan, Multidestination Message Passing in Wormhole k-ary n-cube Networks with Base Routing Conformed Paths, *IEEE Transactions on Parallel and Distributed Systems*, Volume 10, No. 1, Jan. 1999, pp. 76-96.
42. R. Prakash and D. K. Panda, Designing Communication Strategies for Heterogeneous Parallel Systems, *Parallel Computing*, Elsevier Science Publishers, Volume 24, pp. 2035-2052, 1998.
43. R. Sivaram, D. K. Panda, and C. B. Stunkel, Efficient Broadcast and Multicast on Multi-stage Interconnection Networks using Multiport Encoding, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 9, No. 10, Oct 1998, pp. 1004-1028.

44. D. Basak and D. K. Panda, Alleviating Consumption Channel Bottleneck in k-ary n-cube Wormhole Routed Systems, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 9, No. 5, May 1998, pp. 481–496.
45. Y.-C. Tseng, T.-H. Lin, S. K. S. Gupta, and D. K. Panda, Bandwidth-Optimal Complete Exchange on Wormhole-Routed 2D/3D Torus Networks: A Diagonal-Propagation Approach, *IEEE Transactions on Parallel and Distributed Systems*, Vol.8, No. 4, Apr. 1997, pp. 380–396.
46. D. Basak and D. K. Panda, Designing Clustered Multiprocessor Systems under Packaging and Technological Advancements, *IEEE Transactions on Parallel and Distributed Systems*, Vol 7, No. 9, Sept. 1996 , pp. 962–978.
47. Y.-C. Tseng, D. K. Panda, and T.-H. Lai, A Trip-based Multicasting Model in Wormhole-routed Networks with Virtual Channels, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 7, No. 2, Feb. 1996, pp. 138–150.
48. D. K. Panda, Fast Barrier Synchronization in Wormhole k-ary n-cube Networks with Multi-destination Worms, *Future Generation Computer Systems (FGCS)*, Vol. 11, Nov. 1995, pp. 585–602.  
Special Issue on High-Performance Computer Architecture, Collection of **Best Eight Papers** from *Int'l Symposium on High Performance Computer Architecture (HPCA 95)*.
49. K. Hwang and D. K. Panda, High-Radix Symbolic Substitution and Superposition Techniques for Optical Matrix Algebraic Computations, *Optical Engineering*, Nov. 1992, Vol. 31, pp. 2422-2433.
50. K. Hwang and D. K. Panda, Architectural Design of Orthogonal Multiprocessor for Multidimensional Information Processing, *Journal of Information Science and Engineering, Special Issue on Parallel processing*, Dec. 1991, pp. 459-485.
51. D. K. Panda and K. Hwang, Fast Data Movement and Index Manipulation in Multiprocessors Using Parallel Pipelined Memories, *Journal of Parallel and Distributed Computing, Special Issue on Shared-memory Multiprocessors*, June 1991, pp. 130-145.
52. M. Eshaghian, D. K. Panda and V. K. Prasanna Kumar, On the Resource Requirements for Digital Computations on Electro-Optical Systems, *Applied Optics*, Mar. 1991, pp. 928-935.
53. D. K. Panda and T. Vishwanathan, A Parallel-Serial Binary Arbitration Scheme for Collision-free Multi-access Techniques, *Computer Networks and ISDN Systems*, Aug. 1988, pp. 217-223.

## B. Invited Papers (Book Chapters/Conference/Workshop)

1. X. Lu and D. K. Panda, Contribution on Multiple Chapters related to OpenStack, Virtualized HPC, HPC Network Fabric, and HPC Workload Management , Book "The Crossroads of Cloud and HPC: OpenStack for Scientific Research; Exploring OpenStack Cloud Computing for Scientific Workloads", Edited by Stig Telfer - OpenStack Foundation Publishing (Invited Book Chapter), Nov. 2016.

2. X. Lu, M. W. Rahman, N. Islam, D. Shankar, and D. K. Panda, Accelerating Big Data Processing on Modern HPC Clusters, *Conquering Big Data with High Performance Computing* - Springer International Publishing (Invited Book Chapter), Jul. 2016.
3. D. K. Panda and S. Sur, InfiniBand, *Encyclopedia of Parallel Computing*, 2011.
4. D. K. Panda, S. Sur, H. Subramoni and K. Kandalla, Network Support for Collective Communication, *Encyclopedia of Parallel Computing*, 2011.
5. S. Sur, H. Wang, J. Huang, X. Ouyang and D. K. Panda, Can High-Performance Interconnects Benefit Hadoop Distributed File System? Workshop on Micro Architectural Support for Virtualization, Data Center Computing, and Clouds, held in conjunction with Micro 2010.
6. D. K. Panda, P. Balaji, S. Sur and M. Koop, Commodity High Performance Interconnects, Book Chapter in *Attaining High Performance Communication: A Vertical Approach*, Edited by Ada Gavrilovska, CRC Press, Sept. 2009.
7. K. Vaidyanathan, S. Narravula, P. Balaji and D. K. Panda, Designing Efficient Systems Services and Primitives for Next-Generation Data-Centers, Workshop on NSF Next Generation Software (NGS) Program; held in conjunction with IPDPS, Long Beach, 2007.
8. P. Balaji, K. Vaidyanathan, S. Narravula, H.-W. Jin and D. K. Panda, Designing Next-Generation Data-Centers with Advanced Communication Protocols and Systems Services, Workshop on NSF Next Generation Software (NGS) Program; held in conjunction with IPDPS, Greece, 2006.
9. D. K. Panda, Network-Based Computing: Current Trend, Challenges, and the Future, Int'l Conference on Information Technology, Dec. 2000, pp. 17-18.
10. D. K. Panda, Distributed Memory Parallel Systems, *Wiley Encyclopedia of Electrical and Electronics Engineering*, John Wiley and Sons, Vol. 5, Feb. 1999, pp. 714-720.
11. D. K. Panda, D. Basak, D. Dai, R. Kesavan, R. Sivaram, M. Banikazemi and V. Moorthy, Simulation of Modern Parallel Systems: A CSIM-Based Approach 1997 Winter Simulation Conference (WSC '97), Dec. 1997, pp. 1013-1020.
12. D. K. Panda, Issues in Designing Efficient and Practical Algorithms for Collective Communication in Wormhole-Routed Systems, *Proc. of the 1995 ICCP Workshop on Challenges for Parallel Processing*, Aug. 1995, pp. 8-15.
13. K. Hwang and D. K. Panda, The USC Orthogonal Multiprocessor for Image Understanding, *Parallel Architectures and Algorithms for Image Understanding*, edited by V. K. Prasanna Kumar, Academic Press, Jan 1991, pp. 59-94.

### C. Refereed Conference/Workshop Papers

1. 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

2. 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
3. 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
4. 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
5. 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
6. 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 Int'l Supercomputing Conference (ISC '20), June 2020
7. 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, Int'l Supercomputing Conference (ISC '20), June 2020
8. 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 Int'l Parallel and Distributed Processing Symposium (IPDPS '20), May 2020
9. 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
10. 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 Int'l Parallel and Distributed Processing Symposium (IPDPS '20), May 2020
11. 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 Int'l Parallel and Distributed Processing Symposium (IPDPS '20), May 2020
12. 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



13. 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
14. D. Shankar, X. Lu, and D. K. Panda, SIMD-KV: Accelerating End-to-End Performance in Key-Value Stores with SIMD and RDMA over Emerging CPU Architectures, 26th IEEE International Conference on High Performance Computing, Data, Analytics and Data Science (HiPC '19), Dec 2019
15. 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
16. 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
17. K. V. Manian, C. Chu, A. Awan, K. S. Khorassani, H. Subramoni, and D. K. Panda, OMB-UM: 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
18. D. Shankar, X. Lu, and D. K. Panda, SimdHT-Bench: Characterizing SIMD-Aware Hash Table Designs on Emerging CPU Architectures, 2019 IEEE International Symposium on Workload Characterization(IISWC), Nov 2019 **Best Paper Finalist**
19. 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
20. S. Chakraborty, S. Xu, H. Subramoni, and D. K. Panda, Designing Scalable and High-performance MPI Libraries on Amazon Elastic Fabric Adapter, HOT Interconnects 26, Aug 2019
21. 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
22. 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 ISC'19, Jun 2019
23. J. Zhang, X. Lu, C. Chu, and D. K. Panda, C-GDR: High-Performance Container-aware GPUDirect MPI Communication Schemes on RDMA Networks, 33rd IEEE International Parallel & Distributed Processing Symposium (IPDPS '19), May 2019

24. J. Hashmi, S. Chakraborty, M. Bayatpour, H. Subramoni, and D. K. 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 Nominee**
25. 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
26. 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
27. K. Vadambacheri Manian, A. Awan, A. Ruhela, C. Chu, 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
28. S. Gugnani, X. Lu, and D. K. Panda, Analyzing, Modeling, and Provisioning QoS for NVMe SSDs, 11th IEEE/ACM International Conference on Utility and Cloud Computing, Dec 2018
29. A. Awan, C. Chu, H. Subramoni, X. Lu, and D. K. 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
30. R. Biswas, X. Lu, and D. K. Panda, Accelerating TensorFlow with Adaptive RDMA-based gRPC, 25th IEEE International Conference on High Performance Computing, Data, and Analytics, Dec 2018
31. X. Lu, D. Shankar, H. Shi, and D. K. Panda, Spark-uDAPL: Cost-Saving Big Data Analytics on Microsoft Azure Cloud with RDMA Networks 2018 IEEE International Conference on Big Data, Dec 2018
32. H. Shi, X. Lu, and D. K. Panda, EC-Bench: Benchmarking Onload and Offload Erasure Coders on Modern Hardware Architectures, 2018 International Symposium on Benchmarking, Measuring and Optimizing, Dec 2018 **Best Paper Award**.
33. S. Chakraborty, M. Bayatpour, J. Hashmi, H. Subramoni, and D. K. Panda, Cooperative Rendezvous Protocols for Improved Performance and Overlap, 2018 The International Conference for High Performance Computing, Networking, Storage, and Analysis, Nov 2018 **Best Student Paper Finalist**
34. H. Shi, X. Lu, D. Shankar, and D. K. Panda, High-Performance Multi-Rail Erasure Coding Library over Modern Data Center Architectures: Early Experiences, ACM Symposium on Cloud Computing (SoCC) 2018, Oct 2018, Poster Paper.

35. A. Awan, C. Chu, H. Subramoni, and D. K. Panda, Optimized Broadcast for Deep Learning Workloads on Dense-GPU InfiniBand Clusters: MPI or NCCL? EuroMPI 2018 Conference, Sep 2018
36. A. Ruhela, H. Subramoni, S. Chakraborty, M. Bayatpour, P. Kousha, and D. K. Panda, Efficient Asynchronous Communication Progress for MPI without Dedicated Resources, EuroMPI 2018 Conference, Sep 2018
37. M. Li, X. Lu, H. Subramoni, and D. K. Panda, Multi-Threading and Lock-Free MPI RMA Based Graph Processing on KNL and POWER Architectures, EuroMPI 2018 Conference, Sep 2018
38. M. Bayatpour, J. Hashmi, S. Chakraborty, H. Subramoni, P. Kousha, and D. K. Panda, SALaR: Scalable and Adaptive Designs for Large Message Reduction Collectives, IEEE Cluster 2018, Sep 2018 **Best Paper Award (Architecture)**
39. H. Javed, X. Lu, and D. K. Panda, Cutting the Tail: Designing High Performance Message Brokers to Reduce Tail Latencies in Stream Processing, IEEE Cluster 2018, Sep 2018
40. 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
41. R. Biswas, X. Lu, and D. K. Panda, Designing a Micro-Benchmark Suite to Evaluate gRPC for TensorFlow: Early Experiences, The Ninth Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE), in conjunction with ASPLOS '18, Mar 2018
42. J. Hashmi, K. Hamidouche, H. Subramoni, and D. K. Panda, Kernel-assisted Communication Engine for MPI on Emerging Manycore Processors, 24th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'17), Dec 2017
43. 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, 24th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'17), Dec 2017
44. S. Gugnani, X. Lu, F. Pestilli, C.F. Caiafa, and D. K. Panda, MPI-LiFE: Designing High-Performance Linear Fascicle Evaluation of Brain Connectome with MPI, 24th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC'17), Dec 2017
45. S. Gugnani, X. Lu, H. Qi, L. Zha, and D. K. Panda, Characterizing and Accelerating Indexing Techniques on Distributed Ordered Tables, 2017 IEEE International Conference on Big Data (IEEE Big Data 2017), Dec 2017
46. X. Lu, H. Shi, D. Shankar, and D. K. Panda, Performance Characterization and Acceleration of Big Data Workloads on OpenPOWER System, 2017 IEEE International Conference on Big Data (IEEE Big Data 2017), Dec 2017

47. M. W. Rahman, N. Islam, X. Lu, and D. K. Panda, NVMD: Non-Volatile Memory Assisted Design for Accelerating MapReduce and DAG Execution Frameworks on HPC Systems, 2017 IEEE International Conference on Big Data (IEEE Big Data 2017), Dec 2017 [Short Paper]
48. J. Zhang, X. Lu, and D. K. Panda, Is Singularity-based Container Technology Ready for Running MPI Applications on HPC Clouds? 10th IEEE/ACM International Conference on Utility and Cloud Computing, Dec 2017 **Best Student Paper Award**
49. H. Javed, X. Lu, and D. K. Panda, Characterization of Big Data Stream Processing Pipeline: A Case Study using Flink and Kafka, 4th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies, Dec 2017
50. A. Awan, H. Subramoni, and D. K. 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
51. 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
52. 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, ExaMPI2017 - Workshop on Exascale MPI 2017, in conjunction with Supercomputing 2017, Nov 2017
53. J. Hashmi, M. Li, H. Subramoni, and D. K. Panda, Performance of PGAS Models on KNL: A Comprehensive Study with MVAPICH2-X, Intel Xeon Phi User's Group (IXPUG) 2017, Sep 2017
54. S. Chakraborty, M. Bayatpour, H. Subramoni, and D. K. Panda, Advancing MPI Libraries to the Many-core Era: Designs and Evaluations with MVAPICH2 Intel Xeon Phi User's Group (IXPUG) 2017, Sep 2017
55. 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, The 24th European MPI Users Group Meeting, Sep 2017 **Best Paper Award**
56. S. Chakraborty, H. Subramoni, and D. K. Panda, Contention Aware Kernel-Assisted MPI Collectives for Multi/Many-core Systems 2017 IEEE International Conference on Cluster Computing, Sep 2017 **Best Paper Finalist**
57. X. Lu, H. Shi, H. Javed, R. Biswas, and D. K. Panda, Characterizing Deep Learning over Big Data (DLoBD) Stacks on RDMA-capable Networks, The 25th Annual Symposium on High-Performance Interconnects (HotI), Aug 2017
58. A. Venkatesh, C. Chu, K. Hamidouche, S. Potluri, D. Rossetti, and D. K. Panda, MPI-GDS: High Performance MPI Designs with GPUDirect-aSync for CPU-GPU Control Flow Decoupling, International Conference on Parallel Processing, ICPP '17, Aug 2017

59. 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 '17, Aug 2017
60. J. Hashmi, M. Li, H. Subramoni, and D. K. Panda, Exploiting and Evaluating OpenSHMEM on KNL Architecture Fourth Workshop on OpenSHMEM and Related Technologies, Aug 2017
61. H. Subramoni, S. Chakraborty, and D. K. Panda, Designing Dynamic and Adaptive MPI Point-to-point Communication Protocols for Efficient Overlap of Computation and Communication, International Supercomputing Conference (ISC 17), Jun 2017. **Hans Meuer Best Paper Award.**
62. D. Shankar, X. Lu, and D. K. Panda, High-Performance and Resilient Key-Value Store with Online Erasure Coding for Big Data Workloads, 37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017), Jun 2017
63. J. Zhang, X. Lu, and D. K. Panda, High-Performance Virtual Machine Migration Framework for MPI Applications on SR-IOV enabled InfiniBand Clusters, 31st IEEE International Parallel & Distributed Processing Symposium (IPDPS '17), May 2017
64. S. Gugnani, X. Lu, and D. K. Panda, Swift-X: Accelerating OpenStack Swift with RDMA for Building an Efficient HPC Cloud, 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid '17), May 2017
65. N. Islam, M. W. Rahman, X. Lu, and D. K. Panda, Benchmarking Kudu Distributed Storage Engine on High-Performance Interconnects and Storage Devices, The 8th Workshop on Big Data Benchmarks, Performance, Optimization, and Emerging Hardware (BPOE-8), Apr 2017
66. J. Zhang, X. Lu, and D. K. Panda, Designing Locality and NUMA Aware MPI Runtime for Nested Virtualization based HPC Cloud with SR-IOV Enabled InfiniBand, 13th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE '17), Apr 2017
67. X. Lu, N. Islam, M. W. Rahman, and D. K. Panda, NRCIO: NVM-aware RDMA-based Communication and I/O Schemes for Big Data Analytics, The 8th Annual Non-Volatile Memories Workshop (NVMW '17), Mar 2017
68. A. Awan, K. Hamidouche, J. Hashmi, and D. K. Panda, S-Caffe: Co-designing MPI Runtimes and Caffe for Scalable Deep Learning on Modern GPU Clusters, 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Feb 2017
69. M. Li, X. Lu, K. Hamidouche, J. Zhang, and D. K. Panda, Mizan-RMA: Accelerating Mizan Graph Processing Framework with MPI RMA, 23rd IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC), Dec 2016
70. K. Hamidouche, A. Awan, A. Venkatesh, and D. K. Panda, CUDA M3: Designing Efficient CUDA Managed Memory-aware MPI by Exploiting GDR and IPC, 23rd IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC), Dec 2016

71. J. M. Hashmi, K. Hamidouche, and D. K. Panda, Enabling Performance Efficient Runtime Support for Hybrid MPI+UPC++ Programming Models, 18th IEEE Int'l Conference on High Performance Computing and Communications, Dec 2016.
72. D. Banerjee, K. Hamidouche, and D. K. Panda, Re-designing CNTK Deep Learning Framework on Modern GPU Enabled Clusters, 8th IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom '16), Dec 2016
73. S. Gugnani, X. Lu and D. K. Panda, Designing Virtualization-aware and Automatic Topology Detection Schemes for Accelerating Hadoop on SR-IOV-enabled Clouds, 8th IEEE International Conference on Cloud Computing Technology and Science (IEEE CloudCom '16), Dec 2016
74. 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
75. S. Gugnani, X. Lu and D. K. Panda, Performance Characterization of Hadoop Workloads on SR-IOV-enabled Virtualized InfiniBand Clusters, 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT'16), Dec 2016
76. N. Islam, M. W. Rahman, X. Lu, and D. K. Panda, Efficient Data Access Strategies for Hadoop and Spark on HPC Cluster with Heterogeneous Storage, 2016 IEEE International Conference on Big Data, Dec 2016
77. X. Lu, D. Shankar, S. Gugnani, and D. K. Panda, High-Performance Design of Apache Spark with RDMA and Its Benefits on Various Workloads, 2016 IEEE International Conference on Big Data, Dec 2016
78. D. Shankar, X. Lu, and D. K. Panda, Boldio: A Hybrid and Resilient Burst-Buffer Over Lustre for Accelerating Big Data I/O, 2016 IEEE International Conference on Big Data, Dec 2016
79. 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 Int'l Workshop on Optimization of Communication in HPC runtime systems (COMHPC, SC Workshop), Nov 2016
80. K. Hamidouche, J. Zhang, K. Tomko, and D. K. Panda, OpenSHMEM NonBlocking Data Movement Operations with MVAPICH2-X: Early Experiences, PGAS Applications Workshop (in conjunction with SC '16), Nov 2016
81. M. W. Rahman, N. Islam, X. Lu, and D. K. Panda, Can Non-Volatile Memory Benefit MapReduce Applications on HPC Clusters? First Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS, SC Workshop), Nov 2016
82. 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

83. 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
84. M. W. Rahman, N. Islam, X. Lu, D. Shankar, and D. K. Panda, MR-Advisor: A Comprehensive Tuning Tool for Advising HPC Users to Accelerate MapReduce Applications on Supercomputers, 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD'16), Oct 2016
85. A. Awan, K. Hamidouche, A. Venkatesh, and D. K. Panda, Efficient Large Message Broadcast using NCCL and CUDA-Aware MPI for Deep Learning, The 23rd European MPI Users' Group Meeting (EuroMPI 16), Sep 2016, **Best Paper Runner-Up**.
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. J. Zhang, X. Lu, S. Chakraborty, and D. K. Panda, SLURM-V: Extending SLURM for Building Efficient HPC Cloud with SR-IOV and IVShmem, Int'l European Conference on Parallel and Distributed Computing (Euro-Par '16), Aug 2016.
88. J. Zhang, X. Lu, and D. K. Panda, High Performance MPI Library for Container-based HPC Cloud on InfiniBand Clusters, The 45th International Conference on Parallel Processing (ICPP '16), Aug 2016
89. M. Tatineni, X. Lu, D. J. Choi, A. Majumdar, and D. K. Panda, Experiences and Benefits of Running RDMA Hadoop and Spark on SDSC Comet The 5th Annual Conference on Extreme Science and Engineering Discovery Environment (XSEDE), Jul 2016
90. H. Subramoni, A. Augustine, M. Arnold, J. Perkins, X. Lu, K. Hamidouche, and D. K. Panda, INAM<sup>2</sup>: InfiniBand Network Analysis & Monitoring with MPI, Int'l Supercomputing Conference (ISC), Jun 2016.
91. N. Islam, M. W. Rahman, X. Lu, and D. K. Panda, High Performance Design for HDFS with Byte-Addressability of NVM and RDMA, Int'l Conference on Supercomputing (ICS '16), Jun 2016.
92. J. Zhang, X. Lu, and D. K. Panda, Performance Characterization of Hypervisor- and Container-based Virtualization for HPC on SR-IOV Enabled InfiniBand Clusters, Int'l Workshop on Parallel and Distributed Runtime Systems and Middleware (IPDRM '16), in conjunction with IPDPS '16, May 2016.
93. D. Shankar, X. Lu, N. Islam, M. W. Rahman, and D. K. Panda, High-Performance Hybrid Key-Value Store on Modern Clusters with RDMA Interconnects and SSDs: Non-blocking Extensions, Designs, and Benefits, Int'l Parallel and Distributed Processing Symposium (IPDPS '16), May 2016.

94. 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, Int'l Parallel and Distributed Processing Symposium (IPDPS '16), May 2016.
95. C. Chu, K. Hamidouche, A. Venkatesh, A. Awan, and D. K. Panda, CUDA Kernel based Collective Reduction Operations on Large-scale GPU Clusters, Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid'16), May 2016.
96. S. Chakraborty, H. Subramoni, J. Perkins, and D. K. Panda, SHMEMPMI - Shared Memory based PMI for Improved Performance and Scalability, Int'l Symposium on Cluster, Cloud and Grid Computing (CCGrid'16), May 2016.
97. K. Kulkarni, X. Lu, and D. K. Panda, Characterizing Cloudera Impala Workloads with Big-DataBench on InfiniBand Clusters, Int'l Workshop on Big Data Benchmarks, Performance, Optimization, and Emerging Hardware (BPOE-7), in conjunction with ASPLOS '16, Apr 2016.
98. A. Venkatesh, K. Hamidouche, H. Subramoni, and D. K. Panda, Offloaded GPU Collectives using CORE-Direct and CUDA Capabilities on IB Clusters, Int'l Conference on High Performance Computing (HiPC '15), Dec 2015.
99. M. Li, K. Hamidouche, X. Lu, J. Zhang, J. Lin, and D. K. Panda, High Performance OpenSHMEM Strided Communication Support with InfiniBand UMR, Int'l Conference on High Performance Computing (HiPC '15), Dec 2015.
100. 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**
101. N. Islam, M. W. Rahman, X. Lu, D. Shankar, and D. K. Panda, Performance Characterization and Acceleration of In-Memory File Systems for Hadoop and Spark Applications on HPC Clusters, Int'l Conference on Big Data, Oct 2015.
102. D. Shankar, X. Lu, M. W. Rahman, N. Islam, and D. K. Panda, Benchmarking Key-Value Stores on High-Performance Storage and Interconnects for Web-Scale Workloads, Int'l Conference on Big Data, Oct 2015.
103. 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 2015, Sep 2015.
104. M. Li, H. Subramoni, K. Hamidouche, and D. K. Panda, High Performance MPI Datatype Support with User-mode Memory Registration: Challenges, Designs and Benefits, IEEE Cluster 2015, Sep 2015.
105. 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.



106. N. Islam, D. Shankar, X. Lu, M. W. Rahman, and D. K. Panda, Accelerating I/O Performance of Big Data Analytics on HPC Clusters through RDMA-based Key-Value Store Int'l Conference on Parallel Processing (ICPP '15), Sep 2015.
107. A. Bhat, N. Islam, X. Lu, M. W. Rahman, D. Shankar, and D. K. Panda, A Plugin-based Approach to Exploit RDMA Benefits for Apache and Enterprise HDFS Int'l workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-6), in conjunction with VLDB '15, Aug 2015.
108. H. Subramoni, A. Venkatesh, K. Hamidouche, K. Tomko, and D. K. Panda, Impact of InfiniBand DC Transport Protocol on Energy Consumption of All-to-all Collective Algorithms, Int'l Symposium on High Performance Interconnects (HotI '15), Aug 2015.
109. M. Li, K. Hamidouche, X. Lu, J. Lin, and D. K. Panda, High Performance and Scalable Design of MPI-3 RMA on Xeon Phi Clusters, Euro-Par Conference 2015, Aug 2015.
110. A. Awan, K. Hamidouche, C. Chu, and D. K. Panda, A Case for Non-Blocking Collectives in OpenSHMEM: Design, Implementation, and Performance Evaluation using MVAPICH2-X, OpenSHMEM 2015 for PGAS Programming in the Exascale Era, Aug 2015.
111. J. Lin, K. Hamidouche, J. Zhang, X. Lu, A. Vishnu, and D. K. Panda, Accelerating k-NN Algorithm with Hybrid MPI and OpenSHMEM, OpenSHMEM 2015 for PGAS Programming in the Exascale Era, Aug 2015.
112. 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 Conference on Supercomputing (ISC '15), Jul 2015.
113. W. Rahman, X. Lu, N. Islam, R. Rajachandrasekar, and D. K. Panda, High-Performance Design of YARN MapReduce on Modern HPC Clusters with Lustre and RDMA, Int'l Parallel and Distributed Processing Symposium (IPDPS '15), May 2015.
114. J. Lin, K. Hamidouche, X. Lu, M. Li, and D. K. Panda, High-Performance Coarray Fortran Support with MVAPICH2-X: Initial Experience and Evaluation, Int'l Workshop on High-Level Parallel Programming Models and Supporting Environments (HIPS '15), held in conjunction with IPDPS '15, May 2015.
115. S. Chakraborty, H. Subramoni, J. Perkins, A. Awan, and D. K. Panda, On-demand Connection Management for OpenSHMEM and OpenSHMEM+MPI, Int'l Workshop on High-Level Parallel Programming Models and Supporting Environments (HIPS '15), held in conjunction with IPDPS '15, May 2015.
116. S. Chakraborty, H. Subramoni, A. Moody, A. Venkatesh, J. Perkins, and D. K. Panda, Non-blocking PMI Extensions for Fast MPI Startup, Int'l Conference on Cluster and Grid Computing (CCGrid '15), May 2015.
117. J. Zhang, X. Lu, M. Arnold, and D. K. Panda, MVAPICH2 over OpenStack with SR-IOV: An Efficient Approach to Build HPC Clouds, Int'l Conference on Cluster and Grid Computing (CCGrid '15), May 2015.

118. R. Rajachandrasekar, A. Venkatesh, K. Hamidouche, and D. K. Panda, Power-Check: An Energy-Efficient Checkpointing Framework for HPC Clusters, Int'l Conference on Cluster and Grid Computing (CCGrid '15), May 2015.
119. N. Islam, X. Lu, W. Rahman, D. Shankar, and D. K. Panda, Triple-H: A Hybrid Approach to Accelerate HDFS on HPC Clusters with Heterogeneous Storage Architecture, Int'l Conference on Cluster and Grid Computing (CCGrid '15), May 2015.
120. D. Shankar, X. Lu, J. Jose, W. Rahman, N. Islam, and D. K. Panda, Can RDMA Benefit On-Line Data Processing Workloads with Memcached and MySQL, Int'l Conference on Cluster and Grid Computing (CCGrid '15), May 2015.
121. R. Shi, S. Potluri, K. Hamidouche, J. Perkins, M. Li, D. Rossetti and D. K. Panda, Designing Efficient Small Message Transfer Mechanism for Inter-node MPI Communication on InfiniBand GPU Clusters, Int'l Conference on High Performance Computing (HiPC14), December 2014
122. 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, Int'l Conference on High Performance Computing (HiPC 14), December 2014
123. J. Zhang, X. Lu, J. Jose, M. Li, R. Shi and D. K. Panda, High Performance MPI Library over SR-IOV Enabled InfiniBand Clusters, Int'l Conference on High Performance Computing (HiPC '14), December 2014
124. N. Islam, X. Lu, Md. Rahman, R. Rajachandrasekar, and D. K. Panda, In-Memory I/O and Replication for HDFS with Memcached: Early Experiences, Short Paper, IEEE BigData '14, October 2014
125. J. Jose, S. Potluri, H. Subramoni, X. Lu, K. Hamidouche, K. Schulz, H. Sundar and D. K. Panda, Designing Scalable Out-of-core Sorting with Hybrid MPI+PGAS Programming Models, Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '14), October 2014
126. M. Li, J. Lin, X. Lu, K. Hamidouche, K. Tomko and D. K. Panda, Scalable MiniMD Design with Hybrid MPI and OpenSHMEM, OpenSHMEM User Group Conference, in conjunction with Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '14), October 2014
127. J. Jose, K. Hamidouche, X. Lu, S. Potluri, J. Zhang, K. Tomko and D. K. Panda, High Performance OpenSHMEM for MIC Clusters: Extensions, Runtime Designs and Application Co-design, IEEE Cluster (Cluster '14), September 2014
128. M. Li, X. Lu, S. Potluri, K. Hamidouche, J. Jose, K. Tomko and D. K. Panda, Scalable Graph500 Design with MPI-3 RMA, IEEE Cluster (Cluster 14), September 2014
129. R. Shi, X. Lu, S. Potluri, K. Hamidouche, J. Zhang and D. K. Panda, HAND: A Hybrid Approach to Accelerate Non-contiguous Data Movement using MPI Datatypes on GPU Clusters, International Conference on Parallel Processing (ICPP14), September 2014

130. H. Subramoni, K. Kandalla, J. Jose, K. Tomko, K. Schulz, D. Pekurovsky, and D. K. Panda, Designing Topology-Aware Communication Schedules for Alltoall Operations in Large InfiniBand Clusters, International Conference on Parallel Processing (ICPP14), September 2014
131. W. Rahman, X. Lu, N. Islam, and D. K. Panda, Performance Modeling for RDMA-Enhanced Hadoop MapReduce, International Conference on Parallel Processing (ICPP), September 2014
132. R. Rajachandrasekar, J. Perkins, K. Hamidouche, M. Arnold and D. K. Panda, Understanding the Memory-Utilization of MPI Libraries: Challenges and Designs in Implementing the MPI-T Interface. EuroMPI/ASIA14. September 2014.
133. S. Chakraborty, H. Subramoni, J. Perkins, A. Moody, M. Arnold, D. K. Panda, PMI Extensions for Scalable MPI Startup, EuroMPI/ASIA14. September 2014.
134. D. Shankar, X. Lu, W. Rahman, N. Islam, and D. K. Panda, A Micro-benchmark Suite for Evaluating Hadoop MapReduce on High-Performance Networks, The 5th Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE-5), September 2014
135. X. Lu, W. Rahman, N. Islam, D. Shankar, and D. K. Panda, Accelerating Spark with RDMA for Big Data Processing: Early Experiences, Int'l Symposium on High Performance Interconnects (HotI'14), August 2014
136. J. Zhang, X. Lu, J. Jose, R. Shi and D. K. Panda, Can Inter-VM Shmem Benefit MPI Applications on SR-IOV based Virtualized InfiniBand Clusters?, 20th International European Conference on Parallel Processing (Euro-Par), August, 2014
137. M. W. Rahman, X. Lu, N. S. Islam, R. Rajachandrasekar, and D. K. Panda, MapReduce over Lustre: Can RDMA-based Approach Benefit?, 20th International European Conference on Parallel Processing (Euro-Par), August, 2014
138. 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.
139. R. Rajachandrasekar, S. Potluri, A. Venkatesh, K. Hamidouche, M. Wasi-ur-Rahman and D. K. Panda, MIC-Check: A Distributed Checkpointing Framework for the Intel Many Integrated Cores Architecture, Int'l Symposium on High Performance and Distributed Computing (HPDC), June 2014, short paper.
140. N. Islam, X. Lu, M. Wasi-ur-Rahman and D. K. Panda, SOR-HDFS: A SEDA-based Approach to Exploit Maximum Overlapping in RDMA-Enhanced HDFS over InfiniBand, Int'l Symposium on High Performance and Distributed Computing (HPDC), June 2014, short paper.
141. M. Wasi-ur-Rahman, X. Lu, N. Islam and D. K. Panda, HOMR: A Hybrid Approach to Exploit Maximum Overlapping in MapReduce over High Performance Interconnects, International Conference on Supercomputing, (ICS14), June 2014.

142. 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 (IPDPS14), May 2014
143. J. Jose, K. Hamidouche, J. Zhang, A. Venkatesh, and D. K. Panda, Optimizing Collective Communication in UPC, Int'l Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS '14), held in conjunction with International Parallel and Distributed Processing Symposium (IPDPS14), May 2014
144. J. Jose, J. Zhang, A. Venkatesh, S. Potluri and D. K. Panda, A Comprehensive Performance Evaluation of OpenSHMEM Libraries on InfiniBand Clusters, OpenSHMEM Workshop, March 2014.
145. M. Luo, X. Lu, K. Hamidouche, K. Kandalla and D. K. Panda, Initial Study of Multi-Endpoint Runtime for MPI+OpenMP Hybrid Programming Model on Multi-Core Systems, International Symposium on Principles and Practice of Parallel Programming (PPoPP '14), February 2014. (Poster Paper)
146. D. K. Panda, K. Tomko, K. Schulz and A. Majumdar, The MVAPICH Project: Evolution and Sustainability of an Open Source Production Quality MPI Library for HPC, Int'l Workshop on Sustainable Software for Science: Practice and Experiences, held in conjunction with Int'l Conference on Supercomputing (SC '13), November 2013.
147. 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.
148. 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.
149. J. Jose, K. Kandalla, S. Potluri, J. Zhang and D. K. Panda, Optimizing Collective Communication in OpenSHMEM, Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '13), October 2013.
150. K. C. Kandalla, H. Subramoni, K. Tomko, D. Pekurovsky and D. K. Panda. A Novel Functional Partitioning Approach to Design High-Performance MPI-3 Non-Blocking Alltoallv Collective on Multi-core Systems, Int'l Conference on Parallel Processing (ICPP '13), October 2013.
151. X. Lu, N. Islam, M. Wasi-ur-Rahman, J. Jose, H. Subramoni, H. Wang and D. K. Panda. High-Performance Design of Hadoop RPC with RDMA over InfiniBand, Int'l Conference on Parallel Processing (ICPP '13), October 2013.
152. S. Potluri, K. Hamidouche, A. Venkatesh, D. Bureddy and D. K. 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.

153. 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 (Cluster '13), September 2013.
154. R. Shi, S. Potluri, K. Hamidouche, X. Lu, K. Tomko and D. K. Panda, A Scalable and Portable Approach to Accelerate Hybrid HPL on Heterogeneous CPU-GPU Clusters, IEEE Cluster (Cluster '13), September 2013. **Best Student Paper Award**
155. M. Li, S. Potluri, K. Hamidouche, J. Jose and D. K. Panda, Efficient and Truly Passive MPI-3 RMA Using InfiniBand Atomics, EuroMPI 2013, September 2013.
156. 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.
157. N. Islam, X. Lu, M. Rahman and D. K. Panda, Can Parallel Replication Benefit HDFS for High-Performance Interconnects? Int'l Symposium on High-Performance Interconnects (HotI '13), August 2013.
158. S. Potluri, K. Hamidouche, D. Bureddy and D. K. Panda, MVAPICH2-MIC: A High-Performance MPI Library for Xeon Phi Clusters with InfiniBand, Extreme Scaling Workshop, August 2013.
159. A. Venkatesh, K. Kandalla and D. K. Panda, Optimized MPI Gather collective for Many Integrated Core (MIC) InfiniBand Clusters, Extreme Scaling Workshop, August 2013.
160. X. Lu, M.-W. Rahman, N. Islam and D. K. Panda, A Micro-Benchmark Suite for Evaluating Hadoop RPC on High-Performance Networks, Int'l Workshop on Big Data Benchmarking (WBDB '13), July 2013.
161. R. Rajachandrasekar, A. Moody, K. Mohror and D. K. Panda, A 1PB/s File System to Checkpoint Three Million MPI Tasks, Int'l Conference on High Performance Distributed Computing (HPDC '13), June 2013.
162. J. Jose, S. Potluri, K. Tomko and D. K. Panda, Designing Scalable Graph500 Benchmark with Hybrid MPI+OpenSHMEM Programming Models, Int'l Supercomputing Conference (ISC '13), June 2013.
163. 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.
164. 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.
165. M.-W. Rahman, N. S. Islam, X. Lu, J. Jose, H. Subramoni, H. Wang and D. K. Panda, High-Performance RDMA-based Design of Hadoop MapReduce over InfiniBand, Int'l Workshop on High Performance Data Intensive Computing (HPDIC), held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '13), May 2013.

166. A. Venkatesh, K. Kandalla and D. K. Panda, Evaluation of Energy Characteristics of MPI Communication Primitives with RAPL, Int'l Workshop on High Performance, Power-Aware Computing (HPPAC), held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '13), May 2013.
167. S. Potluri, A. Venkatesh, D. Bureddy, K. Kandalla and D. K. Panda, Efficient Intra-node Communication on Intel-MIC Clusters, Int'l Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2013), May 2013.
168. J. Jose, M. Li, X. Lu, K. Kandalla, M. Arnold and D. K. Panda, SR-IOV Support for Virtualization on InfiniBand Clusters: Early Experience, Int'l Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2013), May 2013.
169. N. S. Islam, X. Lu, M. W. Rahman, J. Jose, H. Wang and D. K. Panda, A Micro-benchmark Suite for Evaluating HDFS Operations on Modern Clusters, Int'l Workshop on Big Data Benchmarking (WBDB '12), December 2012.
170. 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), November 2012. Slides
171. 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), November 2012. **Best Paper and Best Student Paper Finalist.**
172. M. Luo, H. Wang and D. K. Panda, Multi-Threaded UPC Runtime for GPU to GPU communication over InfiniBand, Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '12), October 2012.
173. S. Potluri, K. Kandalla, D. Bureddy, M. Li and D. K. Panda, Multi-Threaded UPC Runtime for GPU to GPU communication over InfiniBand, Int'l Conference on Partitioned Global Address Space Programming Models (PGAS '12), October 2012.
174. D. Bureddy, H. Wang, A. Venkatesh, S. Potluri and D. K. Panda, OMB-GPU: A Micro-benchmark suite for Evaluating MPI Libraries on GPU Clusters, EuroMPI 2012, September 2012.
175. R. Rajachandrasekar, J. Jaswani, H. Subramoni and D. K. Panda, Minimizing Network Contention in InfiniBand Clusters with a QoS-Aware Data-Staging Framework, IEEE Cluster (Cluster '12), September 2012.
176. K. Kandalla, A. Buluc H. Subramoni, K. Tomko, J. Vienne, L. Oliker and D. K. Panda, Can Network-Offload based Non-Blocking Neighborhood MPI Collectives Improve Communication Overheads of Irregular Graph Algorithms? Int'l Workshop on Parallel Algorithm and Parallel Software (IWPAPS12), held in conjunction with IEEE Cluster (Cluster '12), September 2012.
177. J. Jose, K. Kandalla, M. Luo and D. K. Panda, Supporting Hybrid MPI and OpenSHMEM over InfiniBand: Design and Performance Evaluation, Int'l Conference on Parallel Processing (ICPP '12), September 2012.

178. X. Ouyang, N. Islam, R. Rajachandrasekhar, J. Jose, M. Luo, H. Wang and D. K. Panda, SSD-Assisted Hybrid Memory to Accelerate Memcached over High Performance Networks, Int'l Conference on Parallel Processing (ICPP '12), September 2012.
179. 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), held in conjunction with EuroPar, August 2012.
180. J. Vienne, J. Chen, M. W. Rahman, N. Islam, H. Subramoni and D. K. Panda, Performance Analysis and Evaluation of InfiniBand FDR and 40GigE RoCE on HPC and Cloud Computing System, Int'l Symposium on High-Performance Interconnects (HotI 2012), August 2012.
181. M. Luo, D. K. Panda, C. Iancu and K. Z. Ibrahim, Congestion Avoidance on Manycore High Performance Computing Systems, Int'l Conference on Supercomputing (ICS '12), June 2012.
182. M. Luo, H. Wang, J. Vienne and D. K. Panda, Redesigning MPI Shared Memory Communication for Large Multi-Core Architecture Int'l Supercomputing Conference (ISC '12), June 2012.
183. J. Huang, X. Ouyang, J. Jose, M. Wasi-ur-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.
184. 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.
185. 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.
186. R. Rajachandrasekhar, X. Besseron and D. K. Panda, Monitoring and Predicting Hardware Failures in HPC Clusters with FTB-IPMI, 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.
187. S. Potluri, H. Wang, D. Bureddy, A. K. Singh, C. Rosales and D. K. Panda, Optimizing MPI Communication on Multi-GPU Systems using CUDA Inter-Process Communication, Int'l Workshop on Accelerators and Hybrid Exascale Systems (AsHES), in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '12), May 2012.
188. J. Jose, H. Subramoni, K. Kandalla, M. Wasi-ur-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.
189. S. Potluri, K. Tomko, D. Bureddy and D. K. Panda, Intra-MIC MPI Communication using MVAPICH2: Early Experience, TACC-Intel Highly-Parallel Computing Symposium, April 2012. **Best Student Paper Award**

190. M. Wasi-ur-Rahman, J. Huang, J. Jose, X. Ouyang, H. Wang, N. S. Islam, H. Subramoni, C. Murthy and D. K. Panda, Understanding the Communication Characteristics in HBase: What are the Fundamental Bottlenecks? Int'l Symposium on Performance Analysis of Systems and Software (ISPASS '12), April 2012. Poster Paper
191. M. Luo, J. Jose, S. Sur and D. K. Panda, Multi-threaded UPC Runtime with Network Endpoints: Design Alternatives and Evaluation on Multi-core Architectures, Int'l Conference on High Performance Computing (HiPC '11), Dec. 2011.
192. J. Jose, S. Potluri, M. Luo, S. Sur and D. K. Panda, UPC Queues for Scalable Graph Traversals: Design and Evaluation on InfiniBand Clusters, Fifth Conference on Partitioned Global Address Space Programming Model (PGAS '11), Oct. 2011.
193. V. Meshram, X. Besseron, X. Ouyang, R. Rajachandrasekhar and D. K. Panda, Can a Decentralized Metadata Service Layer benefit Parallel Filesystems? Workshop on Interfaces and Architectures for Scientific Data Storage (IASDS '11), held in conjunction with Cluster '11, Sept. 2011.
194. A. Singh, S. Potluri, H. Wang, K. Kandalla, S. Sur and D. K. Panda, MPI Alltoall Personalized Exchange on GPGPU Clusters: Design Alternatives and Benefits, Int'l Workshop on Parallel Programming on Accelerator Clusters (PPAC '11), held in conjunction with Cluster '11, Sept. 2011.
195. 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, IEEE Cluster '11, Sept. 2011.
196. H. Wang, S. Potluri, M. Luo, A. Singh, X. Ouyang, S. Sur and D. K. Panda, Optimized Non-contiguous MPI Datatype Communication for GPU Clusters: Design, Implementation and Evaluation with MVAPICH2, IEEE Cluster '11, Sept. 2011.
197. S. Potluri, H. Wang, V. Dhanraj, S. Sur and D. K. Panda, Optimizing MPI One Sided Communication on Multi-core InfiniBand Clusters using Shared Memory Backed Windows, EuroMPI '11, Sept. 2011.
198. S. Potluri, S. Sur, D. Bureddy and D. K. Panda, Design and Implementation of Key Proposed MPI-3 One-Sided Communication Semantics on InfiniBand, Poster/Short Paper, EuroMPI '11, Sept. 2011.
199. J. Jose, H. Subramoni, M. Luo, M. Zhang, J. Huang, M. W. Rahman, N. S. Islam, X. Ouyang, S. Sur and D. K. Panda, Memcached Design on High Performance RDMA Capable Interconnects, Int'l Conference on Parallel Processing (ICPP '11), Sept. 2011.
200. X. Ouyang, R. Rajachandrasekhar, X. Besseron, H. Wang, J. Huang and D. K. Panda, CRFS: A Lightweight User-Level Filesystem for Generic Checkpoint/Restart, Int'l Conference on Parallel Processing (ICPP '11), Sept. 2011.
201. R. Rajachandrasekar, X. Ouyang, X. Besseron, V. Meshram and D. K. Panda, Can Checkpoint/Restart Mechanisms Benefit from Hierarchical Data Staging? Workshop on Resiliency



in High Performance Computing in Clusters, Clouds, and Grids (Resilience '11), held in conjunction with EuroPar, Aug. 2011.

202. 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.
203. 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.
204. 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.
205. H. Wang, S. Potluri, M. Luo, A. Singh, S. Sur and D. K. Panda, MVAPICH2-GPU: Optimized GPU to GPU Communication for InfiniBand Clusters, Int'l Supercomputing Conference (ISC), June 2011.
206. X. Ouyang, R. Rajachandrasekar, X. Besseron and D. K. Panda, High Performance Pipelined Process Migration with RDMA, Int'l Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2011), May 2011.
207. X. Ouyang, D. Nellans, R. Wipfel, D. Flynn and D. K. Panda, Beyond Block I/O: Rethinking Traditional Storage Primitives, 17th IEEE International Symposium on High Performance Computer Architecture (HPCA-17), Commercial Track, February 2011.
208. Y. Cui, K. B. Olsen, T. H. Jordan, K. Lee, J. Zhou, P. Small, D. Roten, G. Ely, D. K. Panda, A. Chourasia, J. Levesque, S. M. Day, P. Maechling, Scalable Earthquake Simulation on Petascale Supercomputers, SuperComputing (SC), November 2010. **Gordon Bell Prize Finalist.**
209. J. Jose, M. Luo, S. Sur and D. K. Panda, Unifying UPC and MPI Runtimes: Experience with MVAPICH, Int'l Workshop on Partitioned Global Address Space (PGAS '10), Oct. 2010.
210. X. Ouyang, S. Marcarelli, R. Rajachandrasekar and D. K. Panda, RDMA-Based Job Migration Framework for MPI over InfiniBand Int'l Conference on Cluster Computing (Cluster '10), Sept. 2010.
211. 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), Sept. 2010.
212. K. Kandalla, E. Mancini, S. Sur and D. K. Panda, Designing Power-Aware Collective Communication Algorithms for InfiniBand Clusters, Int'l Conference on Parallel Processing (ICPP '10), Sept. 2010.

213. 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 Semantics in MVAPICH2, Int'l Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2 '10), in Conjunction with ICPP '10, Sept. 2010.
214. 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.
215. S. Potluri, P. Lai, K. Tomko, S. Sur, Y. Cui, M. Tatineni, K. Schulz, W. Barth, A. Majumdar and D. K. Panda, Quantifying Performance Benefits of Overlap using MPI-2 in a Seismic Modeling Application, 24th International Conference on Supercomputing (ICS), June 2010.
216. P. Lai, S. Sur and D. K. Panda, Designing Truly One-Sided MPI-2 RMA Intra-node Communication on Multi-core Systems, International Supercomputing Conference (ISC'10), June 2010. **Best Paper Award.**
217. 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), May 2010.
218. E. P. Mancini, G. Marsh and D. K. Panda, An MPI-Stream Hybrid Programming Model for Computational Clusters, Int'l Symposium on Cluster Computing and the Grid (CCGrid), May 2010.
219. X. Ouyang, S. Marcarelli and D. K. Panda, Enhancing Checkpoint Performance with Staging IO and SSD, IEEE International Workshop on Storage Network Architecture and Parallel I/Os (SNAPI), May 2010.
220. 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, Int'l Workshop on Communication Architecture for Clusters (CAC 10), Apr. 2010.
221. M. Koop, P. Shamis, I. Rabinovitz and D. K. Panda, Designing High-Performance and Resilient Message Passing on InfiniBand, Int'l Workshop on Communication Architecture for Clusters (CAC 10), Apr. 2010.
222. X. Ouyang, K. Gopalakrishnan, T. Gangadharappa D. K. Panda, Fast Checkpointing by Write Aggregation with Dynamic Buffer and Interleaving on Multicore Architecture, Int'l Conference on High Performance Computing (HiPC '09), Dec. 2009.
223. 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), Sept. 2009.
224. X. Ouyang, K. Gopalakrishnan and D. K. Panda, Accelerating Checkpoint Operation by Node-Level Write Aggregation on Multicore Systems, Int'l Conference on Parallel Processing (ICPP '09), Sept. 2009.

225. R. Gupta, P. Beckman, H. Park, E. Lusk, P. Hargove, A. Geist, D. K. Panda, A. Lumsdaine and J. Dongarra, CIFTS: A Coordinated Infrastructure for Fault-Tolerant Systems, Int'l Conference on Parallel Processing (ICPP '09), Sept. 2009.
226. T. Gangadharappa, M. Koop and D. K. Panda, Designing and Evaluating MPI-2 Dynamic Process Management Support for InfiniBand, Int'l Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2 '09), in Conjunction with ICPP '09, Sept. 2009.
227. J. Sridhar and D. K. Panda, Impact of Node Level Caching in MPI Job Launch Mechanisms, EuroPVM/MPI '09, Sept. 2009.
228. A. Vishnu, M. Krishnan and D. K. Panda, An Efficient Hardware-Software Approach to Network Fault Tolerance with InfiniBand, Int'l Conference on Cluster Computing (Cluster '09), Sept. 2009.
229. M. Koop, M. Luo and D. K. Panda, Reducing Network Contention with Mixed Workloads on Modern Multicore Clusters, Int'l Conference on Cluster Computing (Cluster '09), Sept. 2009.
230. G. Santhanaraman, T. Gangadharappa, S. Narravula, A. Mamidala and D. K. Panda, Design Alternatives for Implementing Fence Synchronization in MPI-2 One-sided Communication on InfiniBand Clusters, Int'l Conference on Cluster Computing (Cluster '09), Sept. 2009.
231. H. Subramoni, P. Lai, M. Luo and D. K. Panda, RDMA over Ethernet - A Preliminary Study, Int'l Workshop on High Performance Distributed Computing (HPI-DC '09), Aug. '09.
232. H. Subramoni, M. Koop and D. K. Panda, Designing Next Generation Clusters: Evaluation of InfiniBand DDR/QDR on Intel Computing Platforms, Int'l Symposium on Hot Interconnects (HotI), Aug. 2009.
233. P. Lai, P. Balaji, R. Thakur and D. K. Panda, ProOnE: A General Purpose Protocol Onload Engine for Multi- and Many-Core Architectures, Int'l Supercomputing Conference (ISC), June 2009.
234. G. Santhanaraman, P. Balaji, K. Gopalakrishnan, R. Thakur, W. Gropp and D. K. Panda, Natively Supporting True One-sided Communication in MPI on Multi-core Systems with InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid), May 2009.
235. K. Kandalla, H. Subramoni, G. Santhanaraman, M. Koop and D. K. Panda, Designing Multi-Leader-based AllGather Algorithms for Multicore Clusters, International Workshop on Communication Architecture for Clusters (CAC), held in conjunction with IPDPS '09, May 2009.
236. M. Koop, J. Sridhar and D. K. Panda, TupleQ: Fully-Asynchronous and Zero-Copy MPI over InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS), May 2009.
237. J. Sridhar, M. Koop, J. Perkins and D. K. Panda, ScELA: Scalable and Extensible Launching Architecture for Clusters, Int'l Symposium on High Performance Computing (HiPC), Dec. 2008.

238. R. Noronha, X. Ouyang and D. K. Panda, Designing High Performance pNFS With RDMA on InfiniBand, Int'l Symposium on High Performance Computing (HiPC), Dec. 2008.
239. P. Balaji, S. Bhagvat, R. Thakur and D. K. Panda, Sockets Direct Protocol for Hybrid Network Stacks: A Case Study with iWARP over 10G Ethernet, Int'l Symposium on High Performance Computing (HiPC), Dec. 2008.
240. 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.
241. M. Koop, J. Sridhar and D. K. Panda, Scalable MPI Design over InfiniBand using eXtended Reliable Connection, IEEE Cluster 2008, Sept. 2008.
242. W. Huang, M. Koop and D. K. Panda, Efficient One-Copy MPI Shared Memory Communication in Virtual Machines, IEEE Cluster 2008, Sept. 2008. **Best Paper Award.**
243. R. Noronha and D. K. Panda, IMCa: A High Performance Caching Frontend for GlusterFS on InfiniBand, Int'l Conference on Parallel Processing (ICPP '08), Sept. 2008.
244. S. Narravul, 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), Sept. 2008.
245. L. Chai, P. Lai, H.-Y. Jin and D. K. Panda, Designing An Efficient Kernel-level and User-level Hybrid Approach for MPI Intra-node Communication on Multi-core Systems, Int'l Conference on Parallel Processing (ICPP '08), Sept. 2008.
246. R. Kumar, A. Mamidala, M. Koop, G. Santhanaraman D. K. Panda, Lock-free Asynchronous Rendezvous Design for MPI Point-to-point Communication, EuroPVM/MPI '08, Sept. 2008.
247. M. Koop, W. Huang, K. Gopalakrishnan and D. K. Panda, Performance Analysis and Evaluation of PCIe 2.0 and Quad-Data Rate InfiniBand, Int'l Symposium on Hot Interconnects (HotI), Aug. 2008.
248. M. Koop, R. Kumar and D. K. Panda, Can Software Reliability Outperform Hardware Reliability on High Performance Interconnects? A Case Study with MPI over InfiniBand, 22nd ACM International Conference on Supercomputing (ICS '08), Greece, June 2008
249. P. Lai, S. Narravula, K. Vaidyanathan and D. K. Panda, Advanced RDMA-based Admission Control for Modern Data-Centers, Int'l Symposium on Cluster Computing and the Grid (CCGrid), May 2008
250. K. Vaidyanathan, S. Narravula, P.Lai and D. K. Panda, Optimized Distributed Data Sharing Substrate in Multi-Core Commodity Clusters: A Comprehensive Study with Applications, Int'l Symposium on Cluster Computing and the Grid (CCGrid), May 2008
251. A. Mamidala, R. Kumar, D. De and D. K. Panda, MPI Collectives on modern Multicore clusters: Performance Optimizations and Communication Characteristics. Int'l Symposium on Cluster Computing and the Grid (CCGrid), May 2008

252. R. Kumar, A. Mamidala and D. K. Panda, Scaling Alltoall Collective on Multi-core Systems, International Workshop on Communication Architecture for Clusters, held in conjunction with IPDPS '08, Apr. 2008.
253. M. Koop, T. Jones and D. K. Panda, MVAPICH-Aptus: Scalable High-Performance Multi-Transport MPI over InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS), Apr. 2008.
254. G. Santhanaraman, S. Narravul and D. K. Panda, Designing Passive Synchronization for MPI-2 One-Sided Communication to Maximize Overlap, Int'l Parallel and Distributed Processing Symposium (IPDPS), Apr. 2008.
255. L. Chai, X. Ouyang, R. Noronha and D. K. Panda, pNFS/PVFS2 over InfiniBand: Early Experiences, Petascale Data Storage Workshop, held in conjunction with Supercomputing (SC), Nov. 2007.
256. W. Huang, M. Koop, Q. Gao and D. K. Panda, Virtual Machine Aware Communication Libraries for High Performance Computing, Supercomputing (SC), Nov. 2007. **Best Student Paper Finalist.**
257. Q. Gao, F. Qin and D. K. Panda, Finding Bugs in Large-Scale Parallel Programs by Detecting Anomaly in Data Movements, Supercomputing (SC), Nov. 2007. **Best Student Paper and Best Paper Finalist.**
258. P. Balaji, W. Feng, S. Bhagvat, D. K. Panda, R. Thakur and W. Gropp, Analyzing the Impact of Supporting Out-of-Order Communication on In-order Performance with iWARP, Supercomputing (SC), Nov. 2007.
259. R. Noronha, L. Chai, S. Shepler and D. K. Panda, Enhancing the Performance of NFSv4 with RDMA, Int'l Workshop on Storage Network Architecture and Parallel I/Os (SNAPI'07), Sept. 2007.
260. G. Santhanaraman, S. Narravula, A. Mamidala and D. K. Panda, MPI-2 One Sided Usage and Implementation for Read Modify Write operations: A case study with HPCC, EuroPVM/MPI 2007, Sept. 2007.
261. M. Koop, S. Sur and D. K. Panda, Zero-Copy Protocol for MPI using InfiniBand Unreliable Datagram, IEEE Cluster 2007, Austin, TX, Sept. 2007.
262. W. Huang, Q. Gao, J. Liu and D. K. Panda, High Performance Virtual Machine Migration with RDMA over Modern Interconnects, IEEE Cluster 2007, Austin, TX, Sept. 2007. **Selected as a BEST Paper.**
263. K. Vaidyanathan, L. Chai, W. Huang and D. K. Panda, Efficient Asynchronous Memory Copy Operations on Multi-Core Systems and I/OAT, IEEE Cluster 2007, Austin, TX, Sept. 2007.
264. Q. Gao, W. Huang, M. Koop, and D. K. Panda, Group-based Coordinated Checkpointing for MPI: A Case Study on InfiniBand, Int'l Conference on Parallel Processing (ICPP'07), Xian, China, Sept. 2007

265. S. Narravul, A. R. Mamidala, A. Vishnu and G. Santhanaraman, and D. K. Panda, High Performance MPI over iWARP: Early Experiences, Int'l Conference on Parallel Processing, Xian, China, Sept. 2007.
266. R. Noronha, L. Chai, T. Talpey and D. K. Panda, Designing NFS With RDMA For Security, Performance and Scalability, Int'l Conference on Parallel Processing, Xian, China, Sept. 2007.
267. S. Sur, M. Koop, L. Chai and D. K. Panda, Performance Analysis and Evaluation of Mellanox ConnectX InfiniBand Architecture with Multi-Core Platforms, Int'l Symposium on Hot Interconnects (HotI), Aug. 2007.
268. M. Koop, S. Sur, Q. Gao and D. K. Panda, High Performance MPI Design using Unreliable Datagram for Ultra-Scale InfiniBand Clusters, 21st Int'l ACM Conference on Supercomputing (ICS '07), June 2007.
269. W. Huang, J. Liu, M. Koop, B. Abali and D. K. Panda, Nomad: Migrating OS-bypass Networks in Virtual Machines, Third Int'l SIGPLAN/SIGOPS Conference on Virtual Execution Environments (VEE), June 2007.
270. M. Koop, T. Jones and D. K. Panda, Reducing Connection Memory Requirements of MPI for InfiniBand Clusters: A Message Coalescing Approach, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Rio de Janeiro - Brazil, May 2007
271. S. Narravul, A. Mamidala, A. Vishnu, K. Vaidyanathan and D. K. Panda, High Performance Distributed Lock Management Services using Network-based Remote Atomic Operations, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Rio de Janeiro - Brazil, May 2007
272. L. Chai, Q. Gao and D. K. Panda, Understanding the Impact of Multi-Core Architecture in Cluster Computing: A Case Study with Intel Dual-Core System, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Rio de Janeiro - Brazil, May 2007
273. A. Vishnu, M. Koop, A. Moody, A. Mamidala, S. Narravul and D. K. Panda, Hot-Spot Avoidance With Multi-Pathing Over InfiniBand: An MPI Perspective, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Rio de Janeiro - Brazil, May 2007
274. K. Vaidyanathan and D. K. Panda, Benefits of I/O Acceleration Technology (I/OAT) in Clusters, International Symposium on Performance Analysis of Systems and Software (ISPASS), San Jose, April 2007.
275. R. Noronha and D. K. Panda, Improving Scalability of OpenMP Applications on MultiCore Systems Using Large Page Support, International Workshop on Multithreaded Architectures and Applications (MTAAP), held in conjunction with IPDPS '07, March 2007.
276. A. Vishnu, B. Benton and D. K. Panda, High Performance MPI on IBM 12x InfiniBand Architecture, International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS), held in conjunction with IPDPS '07, March 2007.

277. A. Vishnu, A. Mamidala, S. Narravul and D. K. Panda, Automatic Path Migration over InfiniBand: Early Experience, Third International Workshop on System Management Techniques, Processes, and Services (SMTPS), held in conjunction with IPDPS '07, March 2007.
278. K. Vaidyanathan, W. Huang, L. Chai and D. K. Panda, Designing Efficient Asynchronous Memory Operations Using Hardware Copy Engine: A Case Study with I/OAT, International Workshop on Communication Architecture for Clusters (CAC), held in conjunction with IPDPS '07, March 2007.
279. A. R. Mamidala, S. Narravula, A. Vishnu, G. Santhanaraman, and D. K. Panda, Using Connection-Oriented and Connection-Less Transport on Performance and Scalability of Collective and One-sided operations: Trade-offs and Impact, International Symposium on Principles and Practice of Parallel Programming (PPoPP 2007), March 2007, San Jose, California.
280. K. Vaidyanathan, S. Narravul and D. K. Panda, DDSS: A Low-Overhead Distributed Data Sharing Substrate for Cluster-Based Data-Centers over Modern Interconnects, Int'l Conference on High Performance Computing (HiPC), December 2006.
281. S. Sur, M. Koop and D. K. Panda, High-Performance and Scalable MPI over InfiniBand with Reduced Memory Usage: An In-Depth Performance Analysis, Supercomputing (SC), November 2006.
282. A. Vishnu, P. Gupta, A. Mamidala and D. K. Panda, A Software Based Approach for Providing Network Fault Tolerance in Clusters Using the uDAPL Interface: MPI Level Design and Performance Evaluation, Supercomputing (SC), November 2006.
283. H.-W. Jin, S. Narravul, K. Vaidyanathan and D. K. Panda, NemC: A Network Emulator for Cluster-of-Clusters, Int'l Conf. on Computer Commn. and Networks, October 2006.
284. K. Vaidyanathan, H.-W. Jin and D. K. Panda, Exploiting RDMA operations for Providing Efficient Fine-Grained Resource Monitoring in Cluster-based Servers, Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies (RAIT 2006), in conjunction with the IEEE Cluster 2006, Sept. 2006.
285. L. Chai, A. Hartono and D. K. Panda, Designing Efficient MPI Intra-node Communication Support for Modern Computer Architectures, Int'l Conference on Cluster Computing, September 2006.
286. A. Mamidala, A. Vishnu and D. K. Panda, Efficient Shared Memory and RDMA based design for MPI\_Allgather over InfiniBand, EuroPVM/MPI, September 2006.
287. M. Koop, W. Huang, A. Vishnu and D. K. Panda, Memory Scalability Evaluation of the Next-Generation Intel Bensley Platform with InfiniBand, Int'l Symposium on Hot Interconnect (HotI), August 2006.
288. Q. Gao, W. Yu, W. Huang and D. K. Panda, Application-Transparent Checkpoint/Restart for MPI Programs over InfiniBand, Int'l Conference on Parallel Processing (ICPP), August 2006.

289. S. Liang, W. Yu and D. K. Panda, High Performance Block I/O for Global File System (GFS) with InfiniBand RDMA, Int'l Conference on Parallel Processing (ICPP), August 2006.
290. W. Huang, J. Liu, B. Abali and D. K. Panda, A Case for High Performance Computing with Virtual Machines, Int'l Conference on Supercomputing (ICS), June 2006.
291. J. Liu, W. Huang, B. Abali and D. K. Panda, High Performance VMM-Bypass I/O in Virtual Machines, USENIX Annual Technical Conference, June 2006.
292. W. Huang, G. Santhanaraman, H.-W. Jin, Q. Gao and D. K. Panda, Design and Implementation of High Performance MVAPICH2: MPI2 over InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
293. L. Chai, R. Noronha and D. K. Panda, MPI over uDAPL: Can High Performance and Portability Exist Across Architectures? Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
294. S. Narravul, H.-W. Jin, K. Vaidyanathan and D. K. Panda, Designing Efficient Cooperative Caching Schemes for Multi-Tier Data-Centers over RDMA-enabled Networks, Int'l Symposium on Cluster Computing and the Grid (CCGrid), Singapore, May 2006
295. S. Sur, L. Chai, H.-W. Jin and D. K. Panda, Shared Receive Queue based Scalable MPI Design for InfiniBand Clusters, Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006, Rhode Island, Greece.
296. W. Yu, Qi Gao and D. K. Panda, Adaptive Connection Management for Scalable MPI over InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006, Rhode Island, Greece.
297. A. Mamidala, L. Chai, H.-W. Jin and D. K. Panda, Efficient SMP-Aware MPI-Level Broadcast over InfiniBand's Hardware Multicast, Communication Architecture for Clusters (CAC) Workshop, to be held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006, Rhode Island, Greece.
298. P. Balaji, S. Bhagvat, H.-W. Jin and D. K. Panda, Asynchronous Zero-Copy Communication for Synchronous Sockets Direct Protocol (SDP) over InfiniBand, Communication Architecture for Clusters (CAC) Workshop, to be held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006, Rhode Island, Greece.
299. W. Yu, R. Noronha, S. Liang, and D. K. Panda, Benefits of High Speed Interconnects to Cluster File Systems: A Case Study with Lustre, Communication Architecture for Clusters (CAC) Workshop, to be held in conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS '06), April 2006.
300. S. Sur, L. Chai, H.-W. Jin and D. K. Panda, RDMA Read Based Rendezvous Protocol for MPI over InfiniBand: Design Alternatives and Benefits, International Symposium on Principles and Practice of Parallel Programming (PPoPP 2006), March 29-31, 2006, New York, New York.



301. V. Vishwanathz, P. Balaji, W. Feng, J. Leigh and D. K. Panda, A Case for UDP Offload Engines in LambdaGrids, International Workshop on Protocols for Fast Long-Distance Networks (PFLDnet 2006), February 2006, Nara, Japan.
302. S. Sur, U. Bondhugula, A. Mamidala, H.-W. Jin, and D. K. Panda, High Performance RDMA Based All-to-all Broadcast for InfiniBand Clusters, International Conference on High Performance Computing (HiPC 2005), December 18-21, 2005, Goa, India. Accepted for publication.
303. A. Vishnu, G. Santhanaraman, W. Huang, H. -W. Jin and D. K. Panda, Supporting MPI-2 One Sided Communication on Multi-Rail InfiniBand Clusters: Design Challenges and Performance Benefits, International Conference on High Performance Computing (HiPC 2005), December 18-21, 2005, Goa, India. Accepted for publication.
304. P. Balaji, H.-W. Jin, K. Vaidyanathan and D. K. Panda, Supporting iWARP Compatibility and Features for Regular Network Adapters, Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies (RAIT 2005), Sept. 2005, in conjunction with the IEEE Cluster 2005.
305. P. Balaji, W. Feng, Q. Gao, R. Noronha, W. Yu and D. K. Panda, Head-to-TOE Evaluation of High-Performance Sockets over Protocol Offload Engines, IEEE Cluster Computing 2005, Sept. 2005.
306. S. Liang, R. Noronha and D. K. Panda, Swapping to Remote Memory over InfiniBand: An Approach using a High Performance Network Block Device, IEEE Cluster Computing 2005, Sept. 2005.
307. L. Chai, R. Noronha, P. Gupta, G. Brown, and D. K. Panda, Designing a Portable MPI-2 over Modern Interconnects using uDAPL Interface, Euro PVM/MPI Conference, Sept. 2005.
308. A. Mamidala, H.-W. Jin, and D. K. Panda, Efficient Hardware Multicast Group Management for Multiple MPI Communicators over InfiniBand, Euro PVM/MPI Conference, Sept. 2005.
309. W. Huang, G. Santhanaraman, H.-W. Jin, and D. K. Panda, Design Alternatives and Performance Trade-offs for Implementing MPI-2 over InfiniBand, Euro PVM/MPI Conference, Sept. 2005.
310. W. Yu and D. K. Panda, Benefits of Quadrics Scatter/Gather to PVFS2 Non-contiguous I/O, International Workshop on Storage Network Architecture and Parallel I/Os (SNAPI) 2005. Sept. 2005.
311. S. Sur, A. Vishnu, H.-Y. Jin, W. Huang, and D. K. Panda, Can Memory-Less Network Adapters Benefit Next-Generation InfiniBand Systems?, Hot Interconnect 13 (HOTI 05), August, 2005.
312. W. Feng, P. Balaji, C. Baron, L. N. Bhuyan, and D. K. Panda, Performance Characterization of a 10-Gigabit Ethernet TOE, Hot Interconnect 13 (HOTI 05), August, 2005.
313. R. Noronha and D. K. Panda, Performance Evaluation of MM5 on Clusters With Modern Interconnects: Scalability and Impact, Euro-Par, August 2005.

314. H.-Y. Jin, S. Narravul, K. Vaidyanathan, P. Balaji, and D. K. Panda, Performance Evaluation of RDMA over IP: A Case Study with Ammasso Gigabit Ethernet NIC, HPI-DC Workshop, in conjunction with HPDC Conference, July, 2005.
315. W. Yu, S. Liang, and D. K. Panda, High Performance Support of Parallel Virtual File System (PVFS2) over Quadrics, Int'l Conference on Supercomputing (ICS '05), June 2005.
316. H. -W. Jin, S. Sur, L. Chai, and D. K. Panda, LiMIC: Support for High-Performance MPI Intra-Node Communication on Linux Cluster, International Conference on Parallel Processing (ICPP-05), June 2005.
317. S. Narravula, P. Balaji, K. Vaidyanathan, H. -W. Jin, D. K. Panda, Architecture for Caching Responses with Multiple Dynamic Dependencies in Multi-Tier Data-Centers over InfiniBand, IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 05), May 2005.
318. R. Noronha and D. K. Panda, Can High Performance Software DSM Systems Designed With InfiniBand Features Benefit from PCI-Express?, DSM Workshop, in conjunction with IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 05), May 2005.
319. D. K. Panda, Designing Multi-Level, Multi-Tier Data Center Architecture for Securing Distributed Infrastructure and Assets, *DHS Homeland Security Conference*, April 2005.
320. L. Chai, S. Sur, H. -W. Jin, and D. K. Panda, Analysis of Design Considerations for Optimizing Multi-Channel MPI over InfiniBand, Workshop on Communication Architecture on Clusters (CAC '05), in conjunction with International Parallel and Distributed Processing Symposium (IPDPS 2005). April 2005.
321. W. Huang, G. Santhanaraman, H. -W. Jin, and D. K. Panda, Scheduling of MPI-2 One Sided Operations over InfiniBand, Workshop on Communication Architecture on Clusters (CAC '05), in conjunction with International Parallel and Distributed Processing Symposium (IPDPS 2005). April 2005.
322. A. Vishnu, A. R Mamidala, H.- W, Jin, and D. K. Panda, Performance Modeling of Subnet Management on Fat Tree InfiniBand Networks using OpenSM, Workshop on System Management Tools on Large Scale Parallel Systems, held in Conjunction with IPDPS 2005. April 2005.
323. W. Yu, T. S. Woodall, R. L. Graham, and D. K. Panda, Design and Implementation of Open MPI over Quadrics/Elan4, International Parallel and Distributed Processing Symposium (IPDPS 2005). April 2005.
324. P. Balaji, S. Narravula, K. Vaidyanathan, H. -W. Jin and D. K. Panda, On the Provision of Prioritization and Soft QoS in Dynamically Reconfigurable Shared Data-Centers over InfiniBand, IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 05), March, 2005.

325. K. Vaidyanathan, P. Balaji, H.-W. Jin and D. K. Panda, Workload-driven Analysis of File Systems in Shared Multi-Tier Data-Centers over InfiniBand Computer Architecture Evaluation using Commercial Workloads (CAECW-8 ), (in conjunction with HPCA), Feb 2005.
326. W. Yu, J. Wu and D. K. Panda, Scalable Startup of Parallel Programs over InfiniBand, Int'l Conference on High Performance Computing (HiPC '04), December, 2004.
327. J. Liu, A. Vishnu, and D. K. Panda, Building Multirail InfiniBand Clusters: MPI-Level Design and Performance Evaluation, SuperComputing 2004 Conference (SC 04), November, 2004
328. W. Yu, D. Buntinas, and D. K. Panda, Scalable and High Performance NIC-Based Allgather over Myrinet/GM, Int'l Conference on Cluster Computing 2004, September, 2004.
329. A. Mamidala, J. Liu, and D. K. Panda, Efficient Barrier and Allreduce on IBA Clusters using Hardware Multicast and Adaptive Algorithms, Int'l Conference on Cluster Computing 2004, September, 2004.
330. A. Wagner, H.-W. Jin, R. Riesen and D. K. Panda, NIC-Based Offload of Dynamic User-Defined Modules for Myrinet Clusters. Int'l Conference on Cluster Computing 2004, September, 2004.
331. R. Noronha and D. K. Panda, Reducing Diff Overhead in Software DSM Systems using RDMA Operations in InfiniBand, Int'l Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies (RAIT 2004), held in conjunction with Cluster 2004 conference, September, 2004.
332. P. Balaji, K. Vaidyanathan, S. Narravula, K. Savitha, H.-W. Jin and D. K. Panda, Exploiting Remote Memory Operations to Design Efficient Reconfiguration for Shared Data-Centers over InfiniBand, Int'l Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies (RAIT 2004), held in conjunction with Cluster 2004 conference, September, 2004.
333. P. Balaji, H. V. Shah and D. K. Panda, Sockets vs RDMA Interface over 10-Gigabit Networks: An In-depth analysis of the Memory Traffic Bottleneck, Int'l Workshop on Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies (RAIT 2004), held in conjunction with Cluster 2004 conference, September, 2004.
334. G. Santhanaraman, J. Wu and D. K. Panda, Zero-Copy MPI Derived Datatype Communication over InfiniBand, EuroPVM/MPI 2004, September, 2004.
335. W. Jiang, J. Liu, H.-W. Jin, D. K. Panda, D. Buntinas, R. Thakur, and W. Gropp, Efficient Implementation of MPI-2 Passive One-Sided Communication on InfiniBand Clusters, EuroPVM/MPI 2004, September, 2004.
336. J. Liu, A. Mamidala, A. Vishnu, and D. K. Panda, Performance Evaluation of InfiniBand with PCI Express, Hot Interconnect 12 (HOTI 04), August, 2004.

337. S. Sur, H.-W. Jin, and D. K. Panda, Efficient and Scalable All-to-All Personalized Exchange for InfiniBand-based Clusters, Int'l Conference on Parallel Processing (ICPP '04), August, 2004.
338. J. Liu, W. Jiang, P. Wyckoff, D. K. Panda, D. Ashton, D. Buntinas, W. Gropp, and B. Toonen, Design and Implementation of MPICH2 over InfiniBand with RDMA Support, Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
339. J. Liu, A. Mamidala and D. K. Panda, Fast and Scalable MPI-Level Broadcast using InfiniBand's Hardware Multicast Support, Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
340. J. Wu, P. Wyckoff, and D. K. Panda, High Performance Implementation of MPI Datatype Communication over InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
341. V. Tipparaju, G. Santhanaraman, J. Nieplocha, and D. K. Panda, Host-Assisted Zero-Copy Remote Memory Access Communication on InfiniBand, Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
342. J. Liu and D. K. Panda, Implementing Efficient and Scalable Flow Control Schemes in MPI over InfiniBand, Int'l Workshop on Communication Architecture for Clusters (CAC 04), Held in Conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
343. W. Yu and D. K. Panda, Efficient and Scalable Barrier over Quadrics and Myrinet with a New NIC-Based Collective Message Passing Protocol, Int'l Workshop on Communication Architecture for Clusters (CAC 04), Held in Conjunction with Int'l Parallel and Distributed Processing Symposium (IPDPS 04), April, 2004.
344. W. Jiang, J. Liu, H.-W. Jin, D. K. Panda, W. Gropp, and R. Thakur, High Performance MPI-2 One-Sided Communication over InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid 04), April, 2004.
345. J. Wu, P. Wyckoff, D. K. Panda, and R. Ross, Unifier: Unifying Cache Management and Communication Buffer Management for PVFS over InfiniBand, Int'l Symposium on Cluster Computing and the Grid (CCGrid 04), April, 2004.
346. R. Noronha and D. K. Panda, Designing High Performance DSM Systems using InfiniBand Features, Int'l Workshop on Distributed Shared Memory Systems, held in conjunction with CCGrid '04, April, 2004.
347. P. Balaji, S. Narravul, K. Vaidyanathan, S. Krishnamoorthy, J. Wu, and D. K. Panda, Sockets Direct Protocol over InfiniBand in Clusters: Is it Beneficial? Int'l Symposium on Performance Analysis of Systems and Software (ISPASS 04). March, 2004
348. S. Narravul, P. Balaji, K. Vaidyanathan, S. Krishnamoorthy, J. Wu, and D. K. Panda, Supporting Strong Coherency for Active Caches in Multi-Tier Data-Centers over InfiniBand, In SAN-03 Workshop (in conjunction with HPCA), Feb. 2004.

349. J. Liu, D. K. Panda, and M. Banikazemi, Evaluating the Impact of RDMA on Storage I/O over InfiniBand, In SAN-03 Workshop (in conjunction with HPCA), Feb. 2004.
350. V. Tipparaju, M. Krishnan, J. Nieplocha, G. Santhanaraman, and D. K. Panda, Exploiting Non-blocking Remote Memory Access Communication in Scientific Benchmarks, High-Performance Computing Conference (HiPC), December 17-20, 2003.
351. A. Wagner, D. Buntinas, R. Brightwell, and D. K. Panda, Application-Bypass Reduction for Large-Scale Clusters, Cluster 2003 Conference, December 1-4, 2003.
352. J. Wu, P. Wyckoff, and D. K. Panda, Supporting Efficient Non-contiguous Access in PVFS over InfiniBand, Cluster 2003 Conference, December 1-4, 2003.
353. V. Tipparaju, M. Krishnan, J. Nieplocha, G. Santhanaraman, and D. K. Panda, Optimizing Mechanisms for Latency Tolerance in Remote Memory Access Communication, Cluster 2003 Conference, December 1-4, 2003.
354. J. Liu, B. Chandrasekaran, J. Wu, W. Jiang, S. Kini, W. Yu, D. Buntinas, P. Wyckoff, and D. K. Panda, Performance Comparison of MPI Implementations over InfiniBand, Myrinet and Quadrics, SuperComputing (SC) Conference, November, 2003.
355. A. Moody, J. Fernandez, F. Petrini, and D. K. Panda, Scalable NIC-based Reduction on Large-scale Clusters, SuperComputing (SC) Conference, November, 2003.
356. W. Yu, S. Sur, D. K. Panda, R. T. Aulwes, and R. Graham, High Performance Broadcast Support in LA-MPI over Quadrics, Los Alamos Computer Science Institute (LACSI) Symposium, Oct. 27-29, 2003.
357. W. Yu, D. Buntinas, and D. K. Panda, High Performance and Reliable NIC-Based Multicast over Myrinet/GM-2, International Conference on Parallel Processing (ICPP 03). Oct. 6-9, 2003.
358. J. Wu, P. Wyckoff, and D. K. Panda, PVFS over InfiniBand: Design and Performance Evaluation, International Conference on Parallel Processing (ICPP 03). Oct. 6-9, 2003.
359. J. Wu, P. Wyckoff, and D. K. Panda, Demotion-Based Exclusive Caching through Demote Buffering: Design and Evaluations over Different Networks, Workshop on Storage Network Architecture and Parallel I/O (SNAPI), In conjunction with PACT '03 conference, Sept. 28, 2003.
360. S. P. Kini, J. Liu, J. Wu, P. Wyckoff, and D. K. Panda, Fast and Scalable Barrier using RDMA and Multicast Mechanisms for InfiniBand-Based Clusters, Euro PVM/MPI Conference, September 29-Oct 2, 2003.
361. B. Chandrasekaran, P. Wyckoff, and D. K. Panda, MIBA: A Micro-benchmark Suite for Evaluating InfiniBand Architecture Implementations, Performance TOOLS 2003, in conjunction with International Multiconference on Measurement, Modeling, and Evaluation of Computer-Communication Systems, Sept. 2003.

362. J. Liu, B. Chandrasekaran, W. Yu, J. Wu, D. Buntinas, S. P. Kinis, P. Wyckoff, and D. K. Panda, Micro-Benchmark Level Performance Comparison of High-Speed Cluster Interconnects, Hot Interconnects 10, August 2003.
363. J. Liu, J. Wu, S. P. Kini, P. Wyckoff, and D. K. Panda, High Performance RDMA-Based MPI Implementation over InfiniBand, Int'l Conference on Supercomputing (ICS '03), June 2003.
364. S. Senapathi, B. Chandrasekharan, D. Stredney, H.-W. Shen, and D. K. Panda, QoS-aware Middleware for Cluster-based Servers to Support Interactive and Resource-Adaptive Applications, High Performance Distributed Computing (HPDC-12), June 2003.
365. P. Balaji, J. Wu, T. Kurc, U. Catalyurek, D. K. Panda, and J. Saltz, Impact of High Performance Sockets on Data Intensive Applications, High Performance Distributed Computing (HPDC-12), June 2003.
366. D. Buntinas, D. K. Panda, and R. Brightwell, Application-Bypass Broadcast in MPICH over GM, Cluster Computing and Grid (CCGrid '03), May 2003.
367. D. Buntinas, A. Saify, D. K. Panda, and Jarek Nieplocha, Optimizing Barrier and Lock Operations in ARMCI, Int'l Workshop on Communication Architecture for Clusters (CAC '03), held in conjunction with IPDPS '03, April 2003.
368. R. Noronha and D. K. Panda, Implementing TreadMarks over GM on Myrinet: Challenges, Design Experience, and Performance Evaluation, Int'l Workshop on Communication Architecture for Clusters (CAC '03), held in conjunction with IPDPS '03, April 2003.
369. R. Gupta, P. Balaji, D. K. Panda, and J. Nieplocha, Efficient Collective Operations using Remote Memory Operations on VIA-Based Clusters, Int'l Parallel and Distributed Processing Symposium (IPDPS '03), April 2003.
370. V. Tipparaju, J. Nieplocha, D. K. Panda, Fast Collective Operations Using Shared and Remote Memory Access Protocols on Clusters, Int'l Parallel and Distributed Processing Symposium (IPDPS '03), April 2003. BEST paper in the software track.
371. D. Buntinas and D. K. Panda, NIC-Based Reduction in Myrinet Clusters: Is It Beneficial? SAN-02 Workshop (in conjunction with HPCA), Feb. 2003.
372. J. Liu, M. Banikazemi, B. Abali, and D. K. Panda, A Portable Client/Server Communication Middleware over SANs: Design and Performance Evaluation with InfiniBand SAN-02 Workshop (in conjunction with HPCA), Feb. 2003.
373. J. Wu, J. Liu, P. Wyckoff, and D. K. Panda, Impact of On-Demand Connection Management in MPI over VIA, Cluster '02, Sept. 2002.
374. R. Gupta, V. Tipparaju, J. Nieplocha, and D. K. Panda, Efficient Barrier using Remote Memory Operations on VIA-Based Clusters, Cluster '02, Sept. 2002.
375. P. Balaji, P. Shivam, P. Wyckoff and D. K. Panda, High Performance User-Level Sockets over Gigabit Ethernet, Cluster '02, Sept. 2002.

376. S. Senapathi, D. K. Panda, D. Stredney, and H.-W. Shen, A QoS Framework for Clusters to support Applications with Resource Adaptivity and Predictable Performance, Int'l Workshop on Quality of Service (IWQoS), May 2002, pp. 180-190.
377. P. Shivam, P. Wyckoff and D. K. Panda, Can User Level Protocols Take Advantage of Multi-CPU NICs?, Int'l Parallel and Distributed Processing Symposium (IPDPS '02), April 2002.
378. J. Wu and D. K. Panda, MPI/IO on DAFS Over VIA: Implementation and Performance Evaluation, Communication Architecture for Clusters (CAC'02) Workshop, held in conjunction with IPDPS '02, April 2002.
379. J. Nielplocha, V. Tipparaju, A. Saify, and D. K. Panda, Protocols and Strategies for Optimizing Remote Memory Operations on Clusters (CAC'02) Workshop, held in conjunction with IPDPS '02, April 2002.
380. D. Buntinas, D. K. Panda, and W. Gropp, NIC-Based Atomic Operations on Myrinet/GM, SAN-1 Workshop, held in conjunction with High Performance Computer Architecture (HPCA) Conference, Feb 2002.
381. P. Shivam, P. Wyckoff, and D. K. Panda, EMP: Zero-copy OS-bypass NIC-driven Gigabit Ethernet Message Passing, Supercomputing '01.
382. M. Banikazemi, J. Liu, D. K. Panda, and P. Sadayappan, Implementing TreadMarks over VIA on Myrinet and Gigabit Ethernet: Challenges, Design Experience, and Performance Evaluation, Int'l Conference on Parallel Processing, September 2001, pp. 167-174.
383. A. Gulati, D. K. Panda, P. Sadayappan, and P. Wyckoff, NIC-based Rate Control for Proportional Bandwidth Allocation in Myrinet Clusters, Int'l Conference on Parallel Processing, September 2001, pp. 305-312.
384. M. Banikazemi, J. Liu, S. Kutlug, A. Ramakrishna, P. Sadayappan, H. Sah, and D. K. Panda, VIBe: A Micro-benchmark Suite for Evaluating Virtual Interface Architecture (VIA) Implementations, Int'l Parallel and Distributed Processing Symposium (IPDPS), April 2001.
385. D. Buntinas, D. K. Panda and P. Sadayappan, Fast NIC-Based Barrier over Myrinet/GM, Int'l Parallel and Distributed Processing Symposium (IPDPS), April 2001.
386. A. Singhal, M. Banikazemi, P. Sadayappan, and D. K. Panda, Efficient Multicast Algorithms for Heterogeneous Switch-based Irregular Networks of Workstations, Int'l Parallel and Distributed Processing Symposium (IPDPS), April 2001.
387. D. Buntinas, D. K. Panda, and P. Sadayappan, Performance Benefits of NIC-Based Barrier on Myrinet/GM, Workshop on Communication Architecture for Clusters (CAC '01), held in conjunction with IPDPS 2001.
388. M. Banikazemi and D. K. Panda, Can Scatter Communication Take Advantage of Multidestination Message Passing? Int'l Symposium on High Performance Computing (HiPC '00), December 2000.

389. Praveen Holenarsipur, V. Yarmolenko, J. Duato, D. K. Panda, and P. Sadayappan, Characterization and Enhancement of Static Mapping Heuristics for Heterogeneous Systems, Int'l Symposium on High Performance Computing (HiPC '00), December 2000.
390. V. Yarmolenko, J. Duato, D. K. Panda, and P. Sadayappan, Dynamic Mapping Heuristics in Heterogeneous Systems, Workshop on Network-Based Computing, held in conjunction with ICPP, August 2000, pp. 437-444.
391. A. Paul, W.-C. Feng, D. K. Panda, and P. Sadayappan, Balancing Web Server Load for Adaptive Video Distribution, Workshop on Multimedia Computing, held in conjunction with ICPP, August 2000, pp. 469-475.
392. M. Banikazemi, D. K. Panda, and P. Sadayappan, Implementing TreadMarks on Virtual Interface Architecture (VIA): Design Issues and Alternatives, Ninth Workshop on Scalable Shared Memory Multiprocessors, held in conjunction with ISCA, June 2000.
393. M. Banikazemi, V. Moorthy, L. Herger, D. K. Panda, and B. Abali, Efficient Virtual Interface Architecture Support for the IBM SP Switch-Connected NT Clusters, Int'l Parallel and Distributed Processing Symposium (IPDPS), May 2000, pp. 33-42.
394. M. Banikazemi, C. B. Stunkel, D. K. Panda, and B. Abali, Adaptive Routing in RS/6000 SP-like Bidirectional Multistage Interconnection Networks, Int'l Parallel and Distributed Processing Symposium (IPDPS), May 2000, pp. 43-52.
395. M. Banikazemi, B. Abali, and D. K. Panda, Comparison and Evaluation of Design Choices for Implementing the Virtual Interface Architecture (VIA), Fourth Int'l Workshop on Communication, Architecture, and Applications for Network-Based Parallel Computing (CANPC '00), Jan. 2000, pp. 145-161.
396. D. Buntinas, D. K. Panda, J. Duato, and P. Sadayappan, Broadcast/Multicast over Myrinet Using NIC-Assisted Multidestination Messages, Fourth Int'l Workshop on Communication and Architectural Support for Network-Based Parallel Computing (CANPC'00), Jan. 2000, pp. 115-129.
397. V. Moorthy, D. K. Panda, and P. Sadayappan, Fast Collective Communication Algorithms for Reflective Memory Network Clusters, Fourth Int'l Workshop on Communication and Architectural Support for Network-Based Parallel Computing (CANPC'00), Jan. 2000, pp. 100-114.
398. M. Banikazemi, R. Govindaraju, R. Blackmore, and D. K. Panda, Implementing Efficient MPI on LAPI for the IBM-SP: Experiences and Performance Evaluation, *International Parallel Processing Symposium (IPPS'99)*, April 99, pp. 183-190.
399. V. Moorthy, M. Jacunski, M. Pillai, P. Ware, D. K. Panda, T. Page, P. Sadayappan, V. Nagarajan, and J. Daniel, Low Latency Message Passing on Workstation Clusters using SCRAM-Net, *International Parallel Processing Symposium (IPPS'99)*, April 99, pp. 148-152.
400. M. Jacunski, P. Sadayappan, and D. K. Panda, All-to-All Broadcast on Switch-Based Clusters of Workstations *International Parallel Processing Symposium (IPPS'99)*, April 99, pp. 325-329.



401. M. Banikazemi, S. Prabhu, J. Sampathkumar, D. K. Panda, and P. Sadayappan, Communication Modeling of Heterogeneous Networks of Workstations for Performance Characterization of Collective Operations, *International Workshop on Heterogeneous Computing (HCW'99)*, in conjunction with IPPS'99, April 99, pp. 125-131.
402. M. Jacunski, V. Moorthy, P. Ware, M. Pillai, D. K. Panda, and P. Sadayappan, Low Latency Message-Passing for Reflective Memory Networks, *International Workshop on Communication, Architecture, and Applications for Network-Based Parallel Computing (CANPC '99)*, Jan 1999, pp. 211-224.
403. M. Banikazemi, V. Moorthy, and D. K. Panda, Efficient Collective Communication on Heterogeneous Networks of Workstations, *International Conference on Parallel Processing*, Aug. 1998, pp. 460-467.
404. R. Sivaram, R. Kesavan, D. K. Panda, and Craig B. Stunkel, Where to Provide Support for Efficient Multicasting in Irregular Networks: Network Interface or Switch? *International Conference on Parallel Processing*, Aug. 1998, pp. 452-459.
405. F. Silla, M. P. Malumbres, J. Duato, D. Dai, and D. K. Panda, Impact of Adaptivity on the Behavior of Networks of Workstations under Bursty Traffic, *International Conference on Parallel Processing*, Aug. 1998, pp. 88-95.
406. A. Bala, D. Shah, W.-C. Feng, and D. K. Panda, Experiences with Software MPEG-2 Video Decompression on an SMP PC, *ICPP Workshop*, Aug. 1998, pp. 29-35.
407. R. Sivaram, C. Stunkel, and D. K. Panda, HIPIQS: A High-Performance Switch Architecture using Input Queuing, *International Parallel Processing Symposium (IPPS '98)*, pp. 134-143.
408. A-H. Smai, D. K. Panda, and L-E. Thorelli, Prioritized Demand Multiplexing (PDM): A Low-Latency Virtual Channel Flow Control Framework for Prioritized Traffic, *International Conference on High Performance Computing, (HiPC'97)*, Dec. 1997, pp. 449-454.
409. D. Dai and D. K. Panda, How Much Does Network Contention Affect Distributed Shared Memory Performance? *International Conference on Parallel Processing (ICPP'97)*, pp. 454-461.
410. R. Kesavan and D. K. Panda, Optimal Multicast with Packetization and Network Interface Support, *International Conference on Parallel Processing (ICPP'97)*, pp. 370-377.
411. R. Kesavan and D. K. Panda, Multicasting on Switch-based Irregular Networks using Multi-drop Path-based Multidestination Worms, *Parallel Computing, Routing, and Communication Workshop, (PCRCW'97)*, pp. 179-192.
412. R. Sivaram, D. K. Panda, and C. B. Stunkel, Multicasting in Irregular Networks with Cut-Through Switches using Tree-Based Multidestination Worms, *Parallel Computing, Routing, and Communication Workshop, (PCRCW'97)*, pp.35-48.
413. D. Dai and D. K. Panda, How Can We Design Better Networks for DSM Systems? *Parallel Computing, Routing, and Communication Workshop, (PCRCW'97)*, pp. 133-146.

414. C. B. Stunkel, R. Sivaram, and D. K. Panda, Implementing Multidestination Worms in Switch-Based Parallel Systems: Architectural Alternatives and their Impact, *International Symposium on Computer Architecture (ISCA'97)*, June 1997, pp. 50-61.
415. R. Sivaram, C. B. Stunkel, and D. K. Panda, A Reliable Hardware Barrier Synchronization Scheme, *International Parallel Processing Symposium (IPPS'97)*, Apr. 1997, pp. 274-280.
416. R. Kesavan, K. Bondalapati, and D. K. Panda, Multicast on Irregular Switch-based Networks with Wormhole Routing, *Proceedings of the Third International Symposium on High Performance Computer Architecture (HPCA-3)*, Feb. 1997, pp. 48-57.
417. R. Sivaram, D. K. Panda, and C. B. Stunkel, Efficient Broadcast and Multicast on Multistage Interconnection Networks using Multiport Encoding, *Proceedings of the Eighth IEEE Symposium on Parallel and Distributed Computing*, Oct. 1996, pp. 36-45.
418. D. Basak and D. K. Panda, Designing Processor-cluster Based Systems: Interplay Between Cluster Organizations and Collective Communication Algorithms, *International Conference on Parallel Processing*, Aug. 1996, pp. 271-274.
419. D. Dai and D. K. Panda, Reducing Cache Invalidation Overheads in Wormhole DSMs using Multidestination Message Passing, *International Conference on Parallel Processing*, Aug. 1996, pp. 138-145.
420. R. Kesavan and D. K. Panda, Minimizing Node Contention in Multiple Multicast on Wormhole k-ary n-cube Networks, *International Conference on Parallel Processing*, Aug. 1996, 188-195.
421. N. S. Sundar, D. N. Jayasimha, D. K. Panda, and P. Sadayappan, Hybrid Algorithms for Complete Exchange in 2D Meshes, *Proceedings of the International Conference on Supercomputing*, May 1996, pp. 181-188.
422. D. Basak, D. K. Panda, M. Banikazemi, Benefits of Processor Clustering in Designing Large Parallel Systems: When and How?, *Proc. of the International Parallel Processing Symposium*, Apr. 1996, pp. 286-290.
423. M. Mandal, D. N. Jayasimha, D. K. Panda, and P. Sadayappan, Reusable Single-Assignment Variables in a Distributed Shared Memory System, *Proc. of the International Conference on High Performance Computing*, Dec. 1995, pp. 36-41.
424. D. K. Panda, Global Reduction in Wormhole k-ary n-cube Networks with Multidestination Exchange Worms, *Proc. of the International Parallel Processing Symposium*, Apr. 1995, pp. 652-659.
425. Y.-C. Tseng, S. K. S. Gupta, and D. K. Panda, An Efficient Scheme for Complete Exchange in 2D Tori, *Proc. of the International Parallel Processing Symposium*, Apr. 1995, pp. 532-536.
426. D. K. Panda, Fast Barrier Synchronization in Wormhole k-ary n-cube Networks with Multidestination Worms, *Proc. of the International Symposium on High Performance Computer Architecture*, Jan. 1995, pp. 200-209,

427. D. K. Panda and D. Basak, Issues in Designing Scalable Systems with k-ary n-cube cluster-c organization, *Proc. of the International Workshop on Parallel Processing*, Dec. 1994, pp. 5-10.
428. R. Prakash and D. K. Panda, Architectural Issues in Designing Heterogeneous Parallel Systems with Passive Star-Coupled Optical Interconnection, *Proc. of the International Symposium on Parallel Architectures, Algorithms and Networks*, Dec. 1994, pp. 246-253. Received an **Outstanding Paper Award**.
429. D. Basak and D. K. Panda, Designing Large Hierarchical Multiprocessor Systems under Processor, Interconnection, and Packaging Advancements, *Proc. of the International Conference on Parallel Processing*, Aug. 1994, pp. I:63-66.
430. D. K. Panda and V. A. Dixit-Radiya, Message-Ordering for Wormhole-Routed Multiport Systems with Link Contention and Routing Adaptivity, *Proc. of the Scalable High Performance Computing Conference*, May. 1994, pp.191-198.
431. N. S. Sundar, D. N. Jayasimha, D. K. Panda and P. Sadayappan, Complete Exchange in 2D Meshes, *Proc. of the Scalable High Performance Computing Conference*, May. 1994, pp. 406-413.
432. D. K. Panda, S. Singal and P. Prabhakaran, Multidestination Message Passing Mechanism Conforming to Base Wormhole Routing Scheme, *Proc. of the Parallel Routing and Communication Workshop*, May 1994, Lecture Notes in Computer Science, 853, Springer-Verlag, pp. 131-145.
433. V. A. Dixit-Radiya and D. K. Panda, Clustering and Intra-Processor Scheduling for Explicitly-Parallel Programs on Distributed-Memory Systems, *Proc. of the International Parallel Processing Symposium*, Apr. 1994, pp. 609-616.
434. D. Basak and D. K. Panda, Scalable Architecture with k-ary n-cube cluster-c Organizations, *Proc. of the Symposium on Parallel and Distributed Processing*, Dec. 1993, pp. 780-787.
435. V. A. Dixit-Radiya and D. K. Panda, Task Assignment in Distributed-Memory Systems with Adaptive Wormhole Routing, *Proc. of the Symposium on Parallel and Distributed Processing*, Dec. 1993, pp. 674-681.
436. D. K. Panda, Optimal Phase Barrier Synchronization in k-ary n-cube Wormhole-routed Systems using Multirendezvous Primitives, *Workshop on Fine-Grain Massively Parallel Coordination*, International Symposium on Computer Architecture, San Diego, May 1993, pp. 24-26.
437. T. Mzaik, S. Chandra, J. M. Jagadeesh, and D. K. Panda, Analysis of Routing in Pyramid Architectures, *Proc. of the IEEE National Aerospace and Electronics Conference (NAECON)*, May 24-28, 1993, pp. 831-837.
438. S. Balakrishnan and D. K. Panda, Impact of Multiple Consumption Channels on Wormhole Routed k-ary n-cube Networks, *Proc. of the International Parallel Processing Symposium*, April 1993, pp. 163-167.

439. S. K. S. Gupta and D. K. Panda, Barrier Synchronization in Distributed-Memory Multiprocessors using Rendezvous Primitives, *Proc. of the International Parallel Processing Symposium*, April 1993, pp. 501-505.
440. Y. C. Tseng and D. K. Panda, A Trip-based Multicasting Model for Wormhole-routed Networks with Virtual Channels, *Proc. of the International Parallel Processing Symposium*, April 1993, pp. 276-283.
441. D. K. Panda and K. Hwang, Message Vectorization for Converting Multicomputer Programs to Shared-Memory Multiprocessors, *Proc. of the 1991 International Conference on Parallel processing*, Aug 1991, pp. I: 204-211.
442. D. K. Panda and K. Hwang, Reconfigurable Vector Register Windows for Fast Matrix Computation On The Orthogonal Multiprocessor, *Proc. of the International Conference on Application Specific Array Processor*, Princeton, New Jersey, Sept 1990, pp. 202-213.
443. S. Mehrotra, C. M. Cheng, Kai Hwang, M. Dubois, and D. K. Panda, Algorithm-Driven Simulation and Performance Projection of a RISC-based Orthogonal Multiprocessor, *Proc. of the International Conference on Parallel Processing*, St. Charles, IL., Aug 1990, pp. III: 244-253.
444. K. Hwang, M. Dubois, D. K. Panda, S. Rao, S. Shang, A. Uresin, W. Mao, H. Nair, M. Lytwyn, F. Hsieh, J. Liu, S. Mehrotra and C. M. Cheng, OMP: A RISC-based Multiprocessor using Orthogonal-Access Memories and Multiple Spanning Buses, *Proc. of the ACM International Conference on Supercomputing*, Amsterdam, The Netherlands, June 1990, pp. 7-22.
445. K. Hwang, D. K. Panda, and N. Haddadi, The USC Orthogonal Multiprocessor for Image Processing with Neural Networks, *Proc. of the 1990 SPIE/SPSE Symposium on Electronic Imaging*, Santa Clara, Feb 1990, pp. 70-84.
446. D. K. Panda and K. Hwang, A Multiple-Bus Network for Implementing Very-Large Neural Networks with Back-Propagation Learning, *Proc. of the International Joint Conference on Neural Networks*, Washington DC, Jan 1990, pp. II:175-178.
447. K. Hwang and D. K. Panda, High-Radix Symbolic Substitution Rules for Optical Arithmetic using Redundant Code, *Proc. of the 9th Symposium on Computer Arithmetic*, Santa Monica, Sept 1989, pp. 226-232.
448. M. Eshaghian, D. K. Panda and V. K. Prasanna Kumar, Resource Requirements for 2-D Image Convolution in Electro-Optical Systems, *Proc. of the 1989 Topical Meetings on Optical Computing*, Salt Lake city, Utah, Mar 1989, pp. 330-333.

#### **D. Edited Books/Proceedings**

1. D. K. Panda and N. Shiratori (Editors), *Proceedings of the 1999 International Conference on Parallel Processing (ICPP '99)*, IEEE Computer Society Press, 1999.

2. D. K. Panda and M. Takizawa (Editors), *Proceedings of the 1999 ICPP Workshops*, IEEE Computer Society Press, 1999.
3. D. K. Panda and C. B. Stunkel (Editors), *Proceedings of the Second International Workshop on Communication, Architecture, and Applications for Network-based Parallel Computing (CANPC'98)*, Lecture Notes in Computer Science, Volume 1362, Springer-Verlag, 1998.
4. D. K. Panda and C. B. Stunkel (Editors), *Proceedings of the First International Workshop on Communication and Architectural Support for Network-based Parallel Computing (CANPC'97)*, Lecture Notes in Computer Science, Volume 1199, Springer-Verlag, 1997.

## E. Other Articles

1. D. K. Panda and L. M. Ni, Guest Editor's Introduction, Special Issue (Volume II) on Workstation Clusters and Network-based Computing, *Journal of Parallel and Distributed Computing*, Sept. 97, pp. 1-2.
2. D. K. Panda and L. M. Ni, Guest Editor's Introduction, Special Issue (Volume I) on Workstation Clusters and Network-based Computing, *Journal of Parallel and Distributed Computing*, Jan. 97, pp. 1-4.
3. V. K. Prasanna, S. Jelinek, and D. K. Panda, TC on Parallel Processing, *IEEE Computer*, Aug 95, pp. 86-87.
4. L. Ni and D. K. Panda, A Report of the ICPP '94 Panel on *Sea of Interconnection Networks: What's Your Choice?*, IEEE Technical Committee on Computer Architecture (TCCA) Newsletter, Winter 95, pp.31-44.
5. D. K. Panda, Guest Editor's Introduction, Special issue on Interconnection Networks for High Performance Computing Systems, *IEEE Technical Committee on Computer Architecture (TCCA) Newsletter*, Winter 95, pp. 9.
6. D. K. Panda, Guest Editor's Introduction, Special issue on Interconnection Networks for High Performance Computing Systems, *IEEE Technical Committee on Computer Architecture (TCCA) Newsletter*, Fall 94, pp. 12-13.