

Abhinav Vishnu

Curriculum Vitae

Education

- 2002–2007 **Ph. D.**, *The Ohio State University*, Computer Science and Engineering.
- 1998–2002 **Bachelor of Technology**, *Institute of Technology, Banaras Hindu University*, Computer Science and Engineering.

Work Experience

- 2017–Present **Team Lead, Scalable Machine Learning**, PACIFIC NORTHWEST NATIONAL LABORATORY.
Extreme Scale Deep Learning Algorithms on Leadership Class Facilities (LCFs)
- 2016–Present **Research Scientist V**, PACIFIC NORTHWEST NATIONAL LABORATORY.
Scalable, Fault Tolerant and Energy Efficient Programming Models for Machine Learning and Data Mining (MLDM) Algorithms
- 2009–2015 **Research Scientist IV**, PACIFIC NORTHWEST NATIONAL LABORATORY.
Scalable, Fault Tolerant and Energy Efficient Programming Models for Machine Learning and Data Mining (MLDM) Algorithms
- 2008–2008 **Software Engineer**, IBM.
Scalable Communication Protocols for IBM PERCS Architecture
- 2006–2007 **IBM Ph. D. Fellow**, THE OHIO STATE UNIVERSITY.
High Performance MPI with Multi-Pathing on InfiniBand
- 2006–2006 **Summer Intern**, IBM.
Scalable LAPI on Modern Interconnects
- 2005–2005 **Summer Intern**, IBM.
Scalable Message Passing Interface on IBM InfiniBand Architecture

Research Grants

- 2016-2019 **xGA: Global Arrays on Exascale Systems**, *PI*, DoE Exascale Computing Project (ECP).
- 2016-2017 **Convergence of Deep Learning and Machine Learning for HPC Modeling and Simulation**, *PI*, Total Award Amount: 450,000, DoE Advanced Scientific Computing Research.
- 2016-2017 **Task Lead**, High Performance Data Analytics Program on Two Tasks.
- 2016-2017 **Learning Control on Building Systems**, *PI*, Total Award Amount: 200,000, Laboratory Directed Research and Development.
- 2015-2016 **Task Lead**, High Performance Data Analytics Program.
- 2015-2016 **Scalable Subsampling for Extreme Scale Machine Learning Algorithms**, *PI*, Total Award Amount: 400,000, Laboratory Directed Research and Development.
- 2014-2015 **Scalable Feature Extraction and Sampling Algorithms for Streaming Data**, *PI*, Total Award Amount: 300,000, Laboratory Directed Research and Development.
- 2013-2014 **Scalable Aggregate Remote Memory Copy Interface on Portals4**, *PI*, Total Award Amount: 155,000, Intel.

- 2012-2013 **Scalable Knowledge Extraction on Large Scale Systems**, *PI*, Total Award Amount: 125,000, Laboratory Directed Research and Development.
- 2009-2012 **A Scalable Fault Tolerance Infrastructure and Algorithms with Programming Models and Scientific Applications**, *PI*, Total Award Amount: 900,000, Laboratory Directed Research and Development.

Supercomputer Time Allocations

- 2017-2019 **xGA on Leadership Class Facilities**, *PI*, Total Award Amount: 128K SUs, Argonne Leadership Class Facility(ALCF).
- 2016-2017 **Scaling MaTE_x on Intel KNL**, *PI*, Total Award Amount: 40K SUs, Texas Advanced Computing Center (TACC).
- 2016-2017 **Scalable Machine Learning and Data Mining on OLCF Titan**, *PI*, Total Award Amount: 1 million core hours, Oak Ridge Director's Discretionary Award.
- 2015-2016 **Scalable Machine Learning and Data Mining on Blue Gene/Q**, *PI*, Total Award Amount: 2 million core hours, Argonne Director's Discretionary Award.
- 2015-2016 **Scalable Machine Learning and Data Mining on Large Scale Systems**, *PI*, Total Award Amount: 2 million core hours, PNNL Institutional Computing (PIC) Constance.
- 2014-2016 **Scalable Global Arrays on InfiniBand**, *co-PI*, Total Award Amount: 1.6 million core hours, PNNL Cascade Supercomputer.

Awards and Accomplishments

- 2015 **Best Student Paper Finalist**, *International Conference in High Performance Computing, Networking, Storage and Analysis (SC'15)*.
- 2015 **DOE Machine Learning Workshop**, *Committee Member and Panelist — Deep Learning on Big Data*.
- 2014 **Best Poster Finalist**, *International Conference in High Performance Computing, Networking, Storage and Analysis (SC'14)*.
- 2014 **DOE Data Council**, *Committee Member*.
- 2013 **DOE Resilience Council**, *Committee Member*.
- 2007 **Best Paper Finalist**, *International Conference on Cluster and Grid Computing (CCGrid'07)*.
- 2007 **TCPP Travel Award**, *International Parallel and Distributed Processing Symposium (IPDPS'07)*.
- 2006 **Ph. D. Fellowship**, *IBM*.
- 2006 **TCSC Travel Award**, *International Conference on High Performance Computing, Networking, Storage and Analysis (SC'06)*.

Patents

- 2013 **Flow Control For Reliable Message Passing**, *US8452888 B2*, IBM.
with Tsai Yang Jea, Hung Thai, Hanhong Xue, Chulho Kim, Uman Chan and Zen Piatek

Books Edited

- ParCo'16 **Programming Models and Systems Software. Special Issue of the Parallel Computing**, *Y. Chen, P. Balaji, and A. Vishnu*, 2016.
- ParCo'15 **Energy Efficient Supercomputing. Special Issue of the Parallel Computing**, *A. Vishnu, A. Marquez, and D. Nikolopoulos*, 2015.

- JoSC'15 **Programming Models and Systems Software. Special Issue of the Journal of Supercomputing** , *Y. Chen, P. Balaji, and A. Vishnu*, 2015.
- ParCo'15 **Programming Models and Systems Software. Special Issue of the Parallel Computing** , *P. Balaji, Y. Chen. and A. Vishnu*, 2015.
- ParCo'13 **Programming Models and Systems Software. Special Issue of the Parallel Computing** , *A. Vishnu, P. Balaji, and Y. Chen.*, 2013.
- JoSC'12 **Programming Models and Systems Software. Special Issue of the Journal of Supercomputing** , *A. Vishnu, P. Balaji, and Y. Chen.*, 2013.
- IJHPCA'11 **Programming Models and Systems Software for High End Applications. Special Issue of the Journal of Supercomputing** , *P. Balaji, and A. Vishnu*, 2011.
- IJHPCA'10 **Programming Models and Systems Software for High End Applications. Special Issue of the Journal of Supercomputing** , *P. Balaji, and A. Vishnu*, 2010.

Publications

- EuroMPI'17 **What does fault tolerant Deep Learning need from MPI?**, *V. Amaty, A. Vishnu, C. Siegel and J. Daily*, EuroMPI/USA, 2017.
- ParLearning'17 **Scaling Deep Learning Workloads: NVIDIA DGX-1/Pascal and Intel Knights Landing** , *N. Gawande, J. Landwehr, J. Daily, N. Tallent, A. Vishnu, and D. Kerbyson*, International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analytics, 2017.
- ICS'17 **Enabling Scalability-Sensitive Speculative Parallelization for FSM Computations**, *J. Qiu, Z. Zhao, B. Wu, A. Vishnu, and S. Song*, International Conference on Supercomputing, 2017.
- IPDPS'17 **Generating Performance Models for Irregular Applications**, *R. Friese, N. Tallent, A. Vishnu, D. Kerbyson and A. Hoisie*, International Parallel and Distributed Processing Symposium, 2017.
- JCC'17 **Deep Learning for Computational Chemistry**, *G. Goh, N. Hodas and, A. Vishnu*, Journal of Computational Chemistry, 2017.
- Arxiv'17 - 1 **User-transparent Distributed TensorFlow** , *A. Vishnu, J. Manzano, C. Siegel and J. Daily*, Arxiv, 2017.
- Arxiv'17 - 2 **Chemception: A Deep Neural Network with Minimal Chemistry Knowledge Matches the Performance of Expert-developed QSAR/QSPR Models.** , *G. Goh, C. Siegel, A. Vishnu, N. Hodas and N. Baker*, Arxiv, 2017.
- BigData'16 **Adaptive Neuron Apoptosis for Accelerating Deep Learning on Large Scale Systems**, *C. Siegel, A. Vishnu and J. Daily*, International Conference on Big Data, 2016.
- IPDPS'16 **Fault Modeling of Extreme Scale Applications using Machine Learning**, *A. Vishnu, N. Tallent, H. v. Dam, D. Kerbyson and A. Hoisie*, International Parallel and Distributed Processing Symposium, 2016.
- ICPADS'16 **Accelerating Deep Learning with Shrinkage and Recall**, *S. Zheng, A. Vishnu and C. Ding*, International Conference on Parallel and Distributed Systems, 2016.
- HiPC'16 **Fault Tolerant Frequent Pattern Mining**, *S. Abdulah, A. Vishnu and G. Agrawal*, International Conference on High Performance Computing, Data and Analytics, 2016.
- ICPP'16 **Fault Tolerant Support Vector Machines** , *S. Abdulah, A. Vishnu and G. Agrawal*, International Conference on Parallel Processing, 2016.
- SC'15 **A Case for Application-Oblivious Energy-Efficient MPI Runtime**, *A. Venkatesh, A. Vishnu, K. Hamidouche, N. Tallent, D. Panda, D. Kerbyson and A. Hoisie*, International Conference on High Performance Computing, Networking, Storage and Analysis, 2015.

- Cluster'15a **Fast and Accurate Support Vector Machines on Large Scale Systems**, A. Vishnu, J. Narasimhan, L. Holder, D. Kerbyson and A. Hoisie, IEEE Cluster, 2015.
- Cluster'15b **Large Scale Frequent Pattern Mining using MPI One-Sided Model**, A. Vishnu, and K. Agarwal, IEEE Cluster, 2015.
- LSPP'15 **On the Impact of Execution Models: A Case Study in Computational Chemistry**, D. Chavarria, M. Halappanavar, S. Krishnamoorthy, A. Vishnu, and A. Hoisie, International Workshop on Large Scale Parallel Processing, IPDPS, 2015.
- PPoPP'15 **Diagnosing the Causes and Severity of One-sided Message Contention**, N. Tallent, A. Vishnu, H. v. Dam, J. Daily, D. Kerbyson, A. Hoisie, Principle and Practice of Parallel Programming, 2015.
- JPDC'14 **A Work Stealing Based Approach for Enabling Scalable Optimal Sequence Homology Detection**, J. Daily, A. Kalyanaraman, S. Krishnamoorthy, and A. Vishnu, Journal of Parallel and Distributed Computing, 2014.
- HiPC'14 **On the Suitability of MPI as a PGAS Runtime**, J. Daily, A. Vishnu, B. Palmer, H. v. Dam, and D. Kerbyson, International Conference on High Performance Computing, 2014.
- FGCS'14 **A Performance Comparison of Current HPC systems: Blue Gene/Q, Cray XE6 and InfiniBand systems.**, J. Daily, D. Kerbyson, K. Barker, A. Vishnu and A. Hoisie, Future Generation Computer Science, 2013.
- JCTC'13 **A Case for Soft Error Detection and Correction in Computational Chemistry**, H. v. Dam, A. Vishnu, and W. d. Jong, Journal of Chemical Theory and Computation, 2013.
- Cluster'13 **An overview of energy efficiency techniques in cluster computing systems**, Giorgio Valentini, Walter Lassonde, Samee Ullah Khan, Nasro Min-Allah, Sajjad Ahmad Madani, Juan Li, Limin Zhang, Lizhe Wang, Nasir Ghani, Joanna Kolodziej, Hongxiang Li, Albert Y. Zomaya, Cheng-Zhong Xu, Pavan Balaji, Abhinav Vishnu, Frédéric Pinel, Johnatan E. Pecero, Dzmityr Kliazovich, Pascal Bouvry, International Journal on Cluster Computing, 2013.
- ParCo'13 **A survey on resource allocation in high performance distributed computing systems**, Hameed Hussain, Saif Ur Rehman Malik, Abdul Hameed, Samee Ullah Khan, Gage Bickler, Nasro Min-Allah, Muhammad Bilal Qureshi, Limin Zhang, Yongji Wang, Nasir Ghani, Joanna Kolodziej, Albert Y. Zomaya, Cheng-Zhong Xu, Pavan Balaji, Abhinav Vishnu, Frédéric Pinel, Johnatan E. Pecero, Dzmityr Kliazovich, Pascal Bouvry, Hongxiang Li, Lizhe Wang, Dan Chen, Ammar Rayes: A survey on resource allocation in high performance distributed computing systems. *Parallel Computing* 39(11): 709-736 (2013), Parallel Computing, 2013.
- JoSC'13 **Designing Energy Efficient Communication Runtime Systems: A View from PGAS Models**, A. Vishnu, S. Song, A. Marquez, K. Barker, D. Kerbyson, K. Cameron, P. Balaji, Journal of Supercomputing, 2013.
- CASS'13 **Building Scalable PGAS Communication Subsystem on Blue Gene/Q**, A. Vishnu, D. Kerbyson, K. Barker, and H. v. Dam, Communication Architecture for Scalable Systems, 2013.
- ICPADS'12 **A. Kerbyson, A. Vishnu, K. Barker, and A. Hoisie**, A. Vishnu, J. Daily and B. Palmer, International Conference on High Performance Computing, 2012.
- HiPC'12 **Designing scalable PGAS communication subsystems on cray gemini interconnect**, A. Vishnu, J. Daily and B. Palmer, International Conference on High Performance Computing, 2012.
- CCGrid'12 **Global Futures: A Multi-threaded Execution Model for Global Arrays-based Applications**, D. Chavarria, S. Krishnamoorthy, and A. Vishnu, International Conference on Cluster, Cloud and Grid Computing, 2012.

- Computer'11 **Codesign Challenges for Exascale Systems: Performance, Power, and Reliability.**, *D. Kerbyson, A. Vishnu, K. Barker, and A. Hoisie*, IEEE Computer, 2011.
- ISC'11 **Mapping communication layouts to network hardware characteristics on massive-scale blue gene systems**, *P. Balaji, R. Gupta, A. Vishnu and P. Beckman*, International Supercomputing Conference, 2011.
- Cluster'11 **Energy Templates: Exploiting Application Information to Save Energy**, *D. Kerbyson, A. Vishnu, and K. Barker*, IEEE Cluster, 2011.
- HotI'11 **Evaluating the Potential of Cray Gemini Interconnect for PGAS Communication Runtime Systems**, *A. Vishnu, M. Bruggencate, and R. Olson*, Symposium on High Performance Interconnects, 2011.
- IPDPS'11 **Iso-Energy-Efficiency: An Approach to Power-Constrained Parallel Computation**, *S. Song, C. Su, R. Ge, A. Vishnu, K. Cameron*, International Parallel and Distributed Processing Symposium, 2011.
- CASS'11 **Dynamic Time-Variant Connection Management for PGAS Models on InfiniBand**, *A. Vishnu, M. Krishnan and P. Balaji*, Communication Architecture for Scalable Systems, 2011.
- EuroMPI'11 **Noncollective Communicator Creation in MPI**, *J. Dinan, S. Krishnamoorthy, P. Balaji, J. Hammond, M. Krishnan, V. Tipparaju, and A. Vishnu*, EuroMPI, 2011.
- JCTC'11 **Designing a Scalable Fault Tolerance Model for High Performance Computational Chemistry: A Case Study with Coupled Cluster Perturbative Triples**, *H. v. Dam, A. Vishnu and W. d. Jong*, Journal of Chemical Theory and Computation, 2011.
- CCGrid'10 **Efficient On-Demand Connection Management Mechanisms with PGAS Models over InfiniBand**, *A. Vishnu and M. Krishnan*, International Conference on Cluster, Cloud and Grid Computing, 2010.
- GreenCom'10 **Designing Energy Efficient Communication Runtime Systems for Data Centric Programming Models**, *A. Vishnu, S. Song, A. Marquez, K. Barker, D. Kerbyson and P. Balaji*, International Conference on Green Computing and Communications, 2010.
- CAC'10 **Designing topology-aware collective communication algorithms for large scale InfiniBand clusters: Case studies with Scatter and Gather**, *K. Kandalla, H. Subramoni, A. Vishnu and D. Panda*, Communication Architecture for Clusters, 2010.
- HiPC'10 **Fault Tolerant Communication Runtime Support for Data Centric Programming Models**, *A. Vishnu, H. v. Dam, W. d. Jong, P. Balaji and S. Song*, International Conference on High Performance Computing, 2010.
- CCPE'09 **Topology Agnostic Hot-Spot Avoidance with InfiniBand**, *A. Vishnu, M. Koop, A. Moody, A. Mamidala, S. Narravula and D. Panda*, Concurrency and Computation: Practice and Experience, 2009.
- Cluster'09 **A Hardware-Software Approach to Network Fault Tolerance with InfiniBand**, *A. Vishnu, M. Krishnan and D. Panda*, International Conference on Cluster Computing, 2009.
- PPoPP'07 **On using connection-oriented vs. connection-less transport for performance and scalability of collective and one-sided operations: trade-offs and impact.**, *A. Mamidala, S. Narravula, A. Vishnu, G. Santhanaraman and D. Panda*, Principles and Practice of Parallel Programming, 2007.
- CCGrid'07 **Hot-Spot Avoidance With Multi-Pathing Over InfiniBand: An MPI Perspective**, *A. Vishnu, M. Koop, A. Moody, A. Mamidala, S. Narravula and D. Panda*, International Conference on Cluster and Grid Computing, 2007.
- CCGrid'07 **High Performance Distributed Lock Management Services using Network-based Remote Atomic Operations**, *S. Narravula, A. Mamidala, A. Vishnu, K. Vaidyanathan and D. Panda*, International Conference on Cluster and Grid Computing, 2007.

- ICPP'07 **High Performance MPI over iWARP: Early Experiences**, *S. Narravula, A. Mamidala, A. Vishnu, G. Santhanaraman and D. Panda*, International Conference on Parallel Processing, 2007.
- SC'06 **A software based approach for providing network fault tolerance in clusters with uDAPL interface: MPI level design and performance evaluation**, *A. Vishnu, P. Gupta, A. Mamidala, and D. Panda*, International Conference on High Performance Networking and Storage, 2006.
- EuroMPI'06 **Efficient Shared Memory and RDMA Based Design for MPI-Allgather over InfiniBand**, *A. Mamidala, A. Vishnu, and D. Panda*, European Conference on MPI, 2006.
- HotI'06 **Memory Scalability Evaluation of the Next-Generation Intel Bensley Platform with InfiniBand**, *M. Koop, W. Huang, A. Vishnu and D. Panda*, Symposium on High Performance Interconnects, 2006.
- Micro'05 **Evaluating InfiniBand Performance with PCI Express**, *J. Liu, A. Mamidala, A. Vishnu and D. Panda*, IEEE Micro, 2005.
- HiPC'05 **Supporting MPI-2 One Sided Communication on Multi-rail InfiniBand Clusters: Design Challenges and Performance Benefits.**, *A. Vishnu, G. Santhanaraman, W. Huang, H.-W. Jin and D. Panda*, International Conference on High Performance Computing, 2005.
- HotI'05 **Can Memory-Less Network Adapters Benefit Next-Generation InfiniBand Systems?**, *S. Sur, A. Vishnu, H.-W. Jin, W. Huang and D. Panda*, Symposium on High Performance Interconnects, 2005.
- SC'04 **Building Multirail InfiniBand Clusters: MPI-Level Design and Performance Evaluation**, *J. Liu, A. Vishnu and D. Panda*, International Conference on High Performance Networking and Storage, 2004.
- HotI'04 **Performance Evaluation of InfiniBand with PCI Express**, *J. Liu, A. Mamidala, A. Vishnu and D. Panda*, Symposium on High Performance Interconnects, 2004.

Posters

- ModSim'15 **Fault Modeling of Extreme Scale Applications using Machine Learning**, *A. Vishnu, H. v. Dam, N. Tallent, D. Kerbyson, A. Hoisie*, DOE Modeling and Simulation Workshop, 2015.
- SC'14 **Diagnosing Network Bottlenecks: One-sided Message Contention**, *N. Tallent, A. Vishnu, H. v. Dam, J. Daily, D. Kerbyson, A. Hoisie*, International Conference on High Performance Networking, Storage and Analysis, 2014.
- SC'13 **PGAS Models using MPI Runtime: Design Alternatives and Performance Evaluation**, *J. Daily, A. Vishnu, B. Palmer, and H. v. Dam*, International Conference on High Performance Networking, Storage and Analysis, 2013.

Position Papers

- ModSim'14 **Exploring Machine Learning Techniques for Dynamic Modeling on Future Exascale Systems**, *S. Song, N. Tallent and A. Vishnu*, Modeling and Simulation of Exascale Systems and Applications, 2014.
- ModSim'13 **Application-specific Modeling of Performance and Power in Concert**, *K. Barker, D. Kerbyson, A. Hoisie and A. Vishnu*, Modeling and Simulation of Exascale Systems and Applications, 2013.
- ScalPerf'10 **Exploring Power Consumption in Extreme Scale Systems**, *K. Barker, D. Kerbyson, A. Hoisie, A. Marquez, and A. Vishnu*, Scalable Approaches to High Performance and High Productivity Computing, 2010.

Professional Activities

Chairman and Editorships

- P2S2'17 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2017.
- P2S2'16 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2016.
- P2S2'15 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2015.
- ParLearning'14 **International Workshop on Parallel and Distributed Computing for Large Scale Machine Learning and Big Data Analysis**, A. Vishnu, and Y. Xia, International Parallel and Distributed Processing Symposium, 2014.
- P2S2'14 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2014.
- P2S2'13 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2013.
- P2S2'12 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2012.
- P2S2'11 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2011.
- P2S2'10 **International Workshop on Parallel Programming Models and Systems Software**, P. Balaji, A. Vishnu, and Y. Chen, International Conference on Parallel Processing, 2010.

Technical Program Committee

- IPDPS **2017,2016**, *International Parallel and Distributed Processing Symposium*.
- HiPC **2017, 2016, 2015, 2014, 2013, 2012**, *International Conference on High Performance Computing*.
- CCGrid **2017, 2016, 2014, 2012, 2011**, *International Conference on Cluster, Cloud and Grid Computing*.
- Cluster **2017, 2012, 2010**, *International Conference on Cluster Computing*.
- ICPP **2012**, *International Conference on Parallel Processing*.
- NCP **2012**, *International Conference on Network and Parallel Computing*.
- PASA **2013**, *International Workshop on Power Aware Systems and Architecture*.
- CASS **2013,2012**, *International Workshop on Communication Architecture for Scalable Systems*.
- PGAS **2012**, *International Conferece on Partitioned Global Address Space Models*.
- HPCC **2010**, *International Conferece on High Performance Computing and Communications*.
- CloudCom **2012**, *International Conferece on Cloud Computing*.
- ICCCN **2011**, *International Conferece on Computer and Communication Networks*.