

# Hari Subramoni

## List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/2969621/publications.pdf>

Version: 2024-02-01

48  
papers

862  
citations

2258059

3  
h-index

2272923

4  
g-index

48  
all docs

48  
docs citations

48  
times ranked

486  
citing authors

#	ARTICLE	IF	CITATIONS
1	Optimizing Distributed DNN Training Using CPUs and BlueField-2 DPUs. IEEE Micro, 2022, 42, 53-60.	1.8	3
2	High Performance MPI over the Slingshot Interconnect: Early Experiences. , 2022, , .		6
3	Adaptive and Hierarchical Large Message All-to-all Communication Algorithms for Large-scale Dense GPU Systems. , 2021, , .		6
4	Design and Characterization of InfiniBand Hardware Tag Matching in MPI. , 2020, , .		2
5	Efficient Training of Semantic Image Segmentation on Summit using Horovod and MVAPICH2-GDR. , 2020, , .		5
6	Dynamic Kernel Fusion for Bulk Non-contiguous Data Transfer on GPU Clusters. , 2020, , .		3
7	Scalable Distributed DNN Training using TensorFlow and CUDA-Aware MPI: Characterization, Designs, and Performance Evaluation. , 2019, , .		26
8	High-Performance Adaptive MPI Derived Datatype Communication for Modern Multi-GPU Systems. , 2019, , .		5
9	Designing a Profiling and Visualization Tool for Scalable and In-depth Analysis of High-Performance GPU Clusters. , 2019, , .		9
10	Leveraging Network-level parallelism with Multiple Process-Endpoints for MPI Broadcast. , 2019, , .		3
11	Designing Scalable and High-Performance MPI Libraries on Amazon Elastic Fabric Adapter. , 2019, , .		3
12	Design and Evaluation of Shared Memory Communication Benchmarks on Emerging Architectures using MVAPICH2. , 2019, , .		0
13	Performance Characterization of DNN Training using TensorFlow and PyTorch on Modern Clusters. , 2019, , .		19
14	Exploiting Hardware Multicast and GPUDirect RDMA for Efficient Broadcast. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 575-588.	5.6	12
15	Networking and communication challenges for post-exascale systems. Frontiers of Information Technology and Electronic Engineering, 2018, 19, 1230-1235.	2.6	2
16	SALaR: Scalable and Adaptive Designs for Large Message Reduction Collectives. , 2018, , .		19
17	Optimized Broadcast for Deep Learning Workloads on Dense-GPU InfiniBand Clusters. , 2018, , .		16
18	An In-depth Performance Characterization of CPU- and GPU-based DNN Training on Modern Architectures. , 2017, , .		45

#	ARTICLE	IF	CITATIONS
19	A Scalable Network-Based Performance Analysis Tool for MPI on Large-Scale HPC Systems. , 2017, , .		5
20	Contention-Aware Kernel-Assisted MPI Collectives for Multi-/Many-Core Systems. , 2017, , .		15
21	Kernel-Assisted Communication Engine for MPI on Emerging Manycore Processors. , 2017, , .		0
22	Designing Registration Caching Free High-Performance MPI Library with Implicit On-Demand Paging (ODP) of InfiniBand. , 2017, , .		2
23	Efficient and Scalable Multi-Source Streaming Broadcast on GPU Clusters for Deep Learning. , 2017, , .		9
24	Designing MPI Library with On-Demand Paging (ODP) of InfiniBand: Challenges and Benefits. , 2016, , .		7
25	Impact of HPC Cloud Networking Technologies on Accelerating Hadoop RPC and HBase. , 2016, , .		7
26	SHMEMPMI – Shared Memory Based PMI for Improved Performance and Scalability. , 2016, , .		6
27	On-demand Connection Management for OpenSHMEM and OpenSHMEM+MPI. , 2015, , .		1
28	High Performance MPI Datatype Support with User-Mode Memory Registration: Challenges, Designs, and Benefits. , 2015, , .		12
29	Exploiting GPUDirect RDMA in Designing High Performance OpenSHMEM for NVIDIA GPU Clusters. , 2015, , .		14
30	Non-Blocking PMI Extensions for Fast MPI Startup. , 2015, , .		8
31	Designing Topology-Aware Communication Schedules for Alltoall Operations in Large InfiniBand Clusters. , 2014, , .		3
32	High-Performance RDMA-based Design of Hadoop MapReduce over InfiniBand. , 2013, , .		42
33	High-Performance Design of Hadoop RPC with RDMA over InfiniBand. , 2013, , .		86
34	Performance Analysis and Evaluation of InfiniBand FDR and 40GigE RoCE on HPC and Cloud Computing Systems. , 2012, , .		35
35	Understanding the communication characteristics in HBase: What are the fundamental bottlenecks?. , 2012, , .		7
36	High-Performance Design of HBase with RDMA over InfiniBand. , 2012, , .		58

#	ARTICLE	IF	CITATIONS
37	Scalable Memcached Design for InfiniBand Clusters Using Hybrid Transports. , 2012, , .		33
38	Memcached Design on High Performance RDMA Capable Interconnects. , 2011, , .		126
39	Codesign for InfiniBand Clusters. Computer, 2011, 44, 31-36.	1.1	4
40	Designing topology-aware collective communication algorithms for large scale InfiniBand clusters: Case studies with Scatter and Gather. , 2010, , .		56
41	High Performance Data Transfer in Grid Environment Using GridFTP over InfiniBand. , 2010, , .		17
42	Improving Application Performance and Predictability Using Multiple Virtual Lanes in Modern Multi-core InfiniBand Clusters. , 2010, , .		9
43	Design and Evaluation of Generalized Collective Communication Primitives with Overlap Using ConnectX-2 Offload Engine. , 2010, , .		18
44	Designing Efficient FTP Mechanisms for High Performance Data-Transfer over InfiniBand. , 2009, , .		12
45	Designing multi-leader-based Allgather algorithms for multi-core clusters. , 2009, , .		28
46	RDMA over Ethernet &#x2014; A preliminary study. , 2009, , .		28
47	Designing Next Generation Clusters: Evaluation of InfiniBand DDR/QDR on Intel Computing Platforms. , 2009, , .		10
48	Design and evaluation of benchmarks for financial applications using Advanced Message Queuing Protocol (AMQP) over InfiniBand. , 2008, , .		20