Manuel F Dolz

List of Publications by Year
in descending order

68 papers	435 citations	1040056 9 h-index	1125743 13 g-index
72 all docs	72 docs citations	72 times ranked	336 citing authors

#	Article	IF	Citations
1	Using machine learning to model the training scalability of convolutional neural networks on clusters of GPUs. Computing (Vienna/New York), 2023, 105, 915-934.	4.8	1
2	Analyzing the impact of the MPI allreduce in distributed training of convolutional neural networks. Computing (Vienna/New York), 2023, 105, 1101-1119.	4.8	2
3	THE JIGSAW PUZZLE: A USE CASE TO COOPERATIVELY LEARN DATA CENTER NETWORK TOPOLOGIES. INTED Proceedings, 2022, , .	0.0	0
4	High performance and energy efficient inference for deep learning on multicore ARM processors using general optimization techniques and BLIS. Journal of Systems Architecture, 2022, 125, 102459.	4.3	3
5	Towards Portable Realizations of Winograd-based Convolution with Vector Intrinsics and OpenMP. , 2022, , .		2
6	Efficient and portable GEMM-based convolution operators for deep neural network training on multicore processors. Journal of Parallel and Distributed Computing, 2022, 167, 240-254.	4.1	7
7	BestOf: an online implementation selector for the training and inference of deep neural networks. Journal of Supercomputing, 2022, 78, 17543-17558.	3.6	1
8	Convolutional neural nets for estimating the run time and energy consumption of the sparse matrix-vector product. International Journal of High Performance Computing Applications, 2021, 35, 268-281.	3.7	2
9	PyDTNN: A user-friendly and extensible framework for distributed deep learning. Journal of Supercomputing, 2021, 77, 9971-9987.	3.6	15
10	Performance Modeling for Distributed Training of Convolutional Neural Networks. , 2021, , .		3
11	Evaluation of MPI Allreduce for Distributed Training of Convolutional Neural Networks. , 2021, , .		2
12	A Flexible Research-Oriented Framework for Distributed Training of Deep Neural Networks. , 2021, , .		6
13	Detecting semantic violations of lock-free data structures through C++ contracts. Journal of Supercomputing, 2020, 76, 5057-5078.	3.6	0
14	A pipeline for the QR update in digital signal processing. Computational and Mathematical Methods, 2020, 2, e1022.	0.8	0
15	Performance modeling of the sparse matrixâ€“vector product via convolutional neural networks. Journal of Supercomputing, 2020, 76, 8883-8900.	3.6	7
16	A similarity study of I/O traces via string kernels. Journal of Supercomputing, 2019, 75, 7814-7826.	3.6	0
17	Theoretical Scalability Analysis of Distributed Deep Convolutional Neural Networks. , 2019, , .		10
18	Analysis of model parallelism for distributed neural networks. , 2019, , .		8

#	Article	IF	Citations
19	Exploring stream parallel patterns in distributed MPI environments. Parallel Computing, 2019, 84, 24-36.	2.1	7
20	A pipeline structure for the block QR update in digital signal processing. Journal of Supercomputing, 2019, 75, 1470-1482.	3.6	1
21	Hybrid staticâ€"dynamic selection of implementation alternatives in heterogeneous environments. Journal of Supercomputing, 2019, 75, 4098-4113.	3.6	2
22	An adaptive offline implementation selector for heterogeneous parallel platforms. International Journal of High Performance Computing Applications, 2018, 32, 854-863.	3.7	3
23	Finding parallel patterns through static analysis in C++ applications. International Journal of High Performance Computing Applications, 2018, 32, 779-788.	3.7	10
24	Energy monitoring as an essential building block towards sustainable ultrascale systems. Sustainable Computing: Informatics and Systems, 2018, 17, 27-42.	2.2	8
25	Understanding hardware and software metrics with respect to power consumption. Sustainable Computing: Informatics and Systems, 2018, 17, 43-54.	2.2	3
26	Supporting MPI-distributed stream parallel patterns in GrPPI. , 2018, , .		3
27	Comparison of Clang Abstract Syntax Trees using String Kernels. , 2018, , .		4
28	Paving the way towards high-level parallel pattern interfaces for data stream processing. Future Generation Computer Systems, 2018, 87, 228-241.	7.5	8
29	Parallelizing and Optimizing LHCb-Kalman for Intel Xeon Phi KNL Processors. , 2018, , .		1
30	Towards Automatic Parallelization of Stream Processing Applications. IEEE Access, 2018, 6, 39944-39961.	4.2	3
31	Supporting Advanced Patterns in GrPPI, aÂGeneric Parallel Pattern Interface. Lecture Notes in Computer Science, 2018, , 55-67.	1.3	0
32	Adapting concurrency throttling and voltageâ€"frequency scaling for dense eigensolvers. Journal of Supercomputing, 2017, 73, 29-43.	3.6	2
33	Enabling semantics to improve detection of data races and misuses of lockâ€free data structures. Concurrency Computation Practice and Experience, 2017, 29, e4114.	2.2	3
34	A generic parallel pattern interface for stream and data processing. Concurrency Computation Practice and Experience, 2017, 29, e4175.	2.2	39
35	Energy-Aware High Performance Computing. , 2017, , .		2
36	A Novel String Representation and Kernel Function for the Comparison of I/O Access Patterns. Lecture Notes in Computer Science, 2017, , 500-512.	1.3	1

#	ARTICLE	IF	CITATIONS
37	Porting Matlab Applications to High-Performance C++ Codes: CPU/GPU-Accelerated Spherical Deconvolution of Diffusion MRI Data. Lecture Notes in Computer Science, 2016, , 630-643.	1.3	7
38	A C++ Generic Parallel Pattern Interface for Stream Processing. Lecture Notes in Computer Science, 2016, , 74-87.	1.3	7
39	Analyzing the energy consumption of the storage data path. Journal of Supercomputing, 2016, 72, 4089-4106.	3.6	4
40	An analytical methodology to derive power models based on hardware and software metrics. Computer Science - Research and Development, 2016, 31, 165-174.	2.7	4
41	Discovering Pipeline Parallel Patterns in Sequential Legacy C++ Codes. , 2016, , .		6
42	Embedding Semantics of the Single-Producer/Single-Consumer Lock-Free Queue into a Race Detection Tool. , 2016, , .		2
43	CID: A Compile-Time Implementation Decider for Heterogeneous Platforms Based on C++ Attributes. , 2016, , .		2
44	Evaluating the performance and energy efficiency of the COSMO-ART model system. Computer Science - Research and Development, 2015, 30, 177-186.	2.7	5
45	ARDUPOWER: A low-cost wattmeter to improve energy efficiency of HPC applications. , 2015, , .		13
46	Are our dense linear algebra libraries energy-friendly?. Computer Science - Research and Development, 2015, 30, 187-196.	2.7	4
47	Balancing task- and data-level parallelism to improve performance and energy consumption of matrix computations on the Intel Xeon Phi. Computers and Electrical Engineering, 2015, 46, 95-111.	4.8	11
48	Tools and Methods for Measuring and Tuning the Energy Efficiency of HPC Systems. Scientific Programming, 2014, 22, 273-283.	0.7	13
49	Enhancing performance and energy consumption of runtime schedulers for dense linear algebra. Concurrency Computation Practice and Experience, 2014, 26, 2591-2611.	2.2	2
50	Modeling power and energy consumption of dense matrix factorizations on multicore processors. Concurrency Computation Practice and Experience, 2014, 26, 2743-2757.	2.2	3
51	Assessing the impact of the CPU power-saving modes on the task-parallel solution of sparse linear systems. Cluster Computing, 2014, 17, 1335-1348.	5.0	9
52	Automatic detection of power bottlenecks in parallel scientific applications. Computer Science - Research and Development, 2014, 29, 221-229.	2.7	5
53	Modeling power and energy of the task-parallel Cholesky factorization on multicore processors. Computer Science - Research and Development, 2014, 29, 105-112.	2.7	12
54	Block pivoting implementation of a symmetric Toeplitz solver. Journal of Parallel and Distributed Computing, 2014, 74, 2392-2399.	4.1	1

#	Article	IF	Citations
55	Assessing Power Monitoring Approaches for Energy and Power Analysis of Computers. Sustainable Computing: Informatics and Systems, 2014, 4, 68-82.	2.2	15
56	Evaluating Lustre's Metadata Server on a Multi-Socket Platform. , 2014, , .		1
57	Energy-efficient execution of dense linear algebra algorithms on multi-core processors. Cluster Computing, 2013, 16, 497-509.	5.0	6
58	Solving Some Mysteries in Power Monitoring of Servers: Take Care of Your Wattmeters!. Lecture Notes in Computer Science, 2013, , 3-18.	1.3	12
59	Runtime Scheduling of the LU Factorization: Performance and Energy. Lecture Notes in Computer Science, 2013, , 153-167.	1.3	1
60	Tools for Power-Energy Modelling and Analysis of Parallel Scientific Applications. , 2012, , .		46
61	Saving Energy in the LU Factorization with Partial Pivoting on Multi-core Processors. , 2012, , .		7
62	A simulator to assess energy saving strategies and policies in HPC workloads. Operating Systems Review (ACM), 2012, 46, 2-9.	1.9	3
63	DVFS-control techniques for dense linear algebra operations on multi-core processors. Computer Science - Research and Development, 2012, 27, 289-298.	2.7	12
64	Reducing Energy Consumption of Dense Linear Algebra Operations on Hybrid CPU-GPU Platforms. , 2012, , .		12
65	Binding Performance and Power of Dense Linear Algebra Operations. , 2012, , .		3
66	Improving power efficiency of dense linear algebra algorithms on multi-core processors via slack control. , 2011, , .		12
67	Evaluation of the Energy Performance of Dense Linear Algebra Kernels on Multi-core and Many-Core Processors. , 2011, , .		4
68	A flexible simulator to evaluate a power saving system for HPC clusters. , 2011, , .		0