

Jakub Kurzak

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/7035671/publications.pdf>

Version: 2024-02-01

19
papers

708
citations

932766

10
h-index

996533

15
g-index

20
all docs

20
docs citations

20
times ranked

527
citing authors

#	ARTICLE	IF	CITATIONS
1	Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects. Journal of Physics: Conference Series, 2009, 180, 012037.	0.3	245
2	Parallel tiled QR factorization for multicore architectures. Concurrency Computation Practice and Experience, 2008, 20, 1573-1590.	1.4	92
3	Autotuning GEMM Kernels for the Fermi GPU. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 2045-2057.	4.0	85
4	Scheduling dense linear algebra operations on multicore processors. Concurrency Computation Practice and Experience, 2010, 22, 15-44.	1.4	53
5	Implementation and Tuning of Batched Cholesky Factorization and Solve for NVIDIA GPUs. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 2036-2048.	4.0	29
6	Accelerating collaborative filtering using concepts from high performance computing. , 2015, , .		25
7	Porting the PLASMA Numerical Library to the OpenMP Standard. International Journal of Parallel Programming, 2017, 45, 612-633.	1.1	25
8	QR Factorization for the Cell Broadband Engine. Scientific Programming, 2009, 17, 31-42.	0.5	23
9	Parallel Two-Sided Matrix Reduction to Band Bidiagonal Form on Multicore Architectures. IEEE Transactions on Parallel and Distributed Systems, 2010, 21, 417-423.	4.0	18
10	Implementing Linear Algebra Routines on Multi-core Processors with Pipelining and a Look Ahead. , 2006, , 147-156.		18
11	Experiences in autotuning matrix multiplication for energy minimization on GPUs. Concurrency Computation Practice and Experience, 2015, 27, 5096-5113.	1.4	16
12	Search Space Generation and Pruning System for Autotuners. , 2016, , .		9
13	Autotuning Numerical Dense Linear Algebra for Batched Computation With GPU Hardware Accelerators. Proceedings of the IEEE, 2018, 106, 2040-2055.	16.4	9
14	Performance-Portable Autotuning of OpenCL Kernels for Convolutional Layers of Deep Neural Networks. , 2016, , .		8
15	Scheduling Two-Sided Transformations Using Tile Algorithms on Multicore Architectures. Scientific Programming, 2010, 18, 35-50.	0.5	7
16	Symmetric Indefinite Linear Solver Using OpenMP Task on Multicore Architectures. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 1879-1892.	4.0	6
17	Linear Systems Solvers for Distributed-Memory Machines with GPU Accelerators. Lecture Notes in Computer Science, 2019, , 495-506.	1.0	5
18	Automatic Generation of FFT for Translations of Multipole Expansions in Spherical Harmonics. International Journal of High Performance Computing Applications, 2008, 22, 219-230.	2.4	3

#	ARTICLE	IF	CITATIONS
19	Autotuning batch Cholesky factorization in CUDA with interleaved layout of matrices. , 2017, , .		3