

Georg Hager

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/6569787/publications.pdf>

Version: 2024-02-01

112
papers

2,852
citations

331670

21
h-index

315739

38
g-index

118
all docs

118
docs citations

118
times ranked

1904
citing authors

#	ARTICLE	IF	CITATIONS
1	Executionâ€Cacheâ€Memory modeling and performance tuning of sparse matrixâ€vector multiplication and Lattice quantum chromodynamics on A64FX. Concurrency Computation Practice and Experience, 2022, 34, e6512.	2.2	11
2	Analytic performance model for parallel overlapping memoryâ€bound kernels. Concurrency Computation Practice and Experience, 2022, 34, .	2.2	4
3	Addressing White-box Modeling and Simulation Challenges in Parallel Computing. , 2022, , .		1
4	A domain-specific language and matrix-free stencil code for investigating electronic properties of Dirac and topological materials. International Journal of High Performance Computing Applications, 2021, 35, 60-77.	3.7	2
5	Performance engineering for real and complex tall & skinny matrix multiplication kernels on GPUs. International Journal of High Performance Computing Applications, 2021, 35, 5-19.	3.7	5
6	YaskSite: Stencil Optimization Techniques Applied to Explicit ODE Methods on Modern Architectures. , 2021, , .		1
7	Analytic Modeling of Idle Waves in Parallel Programs: Communication, Cluster Topology, and Noise Impact. Lecture Notes in Computer Science, 2021, , 351-371.	1.3	7
8	Analytic performance modeling and analysis of detailed neuron simulations. International Journal of High Performance Computing Applications, 2020, 34, 428-449.	3.7	9
9	Performance Engineering for a Tall & Skinny Matrix Multiplication Kernels on GPUs. Lecture Notes in Computer Science, 2020, , 505-515.	1.3	3
10	Desynchronization and Wave Pattern Formation in MPI-Parallel and Hybrid Memory-Bound Programs. Lecture Notes in Computer Science, 2020, , 391-411.	1.3	8
11	Understanding HPC Benchmark Performance on Intel Broadwell and Cascade Lake Processors. Lecture Notes in Computer Science, 2020, , 412-433.	1.3	9
12	Performance Modeling of Streaming Kernels and Sparse Matrix-Vector Multiplication on A64FX. , 2020, , .		9
13	A Recursive Algebraic Coloring Technique for Hardware-efficient Symmetric Sparse Matrix-vector Multiplication. ACM Transactions on Parallel Computing, 2020, 7, 1-37.	1.4	65
14	PHIST. ACM Transactions on Mathematical Software, 2020, 46, 1-26.	2.9	5
15	ESSEX: Equipping Sparse Solvers For Exascale. Lecture Notes in Computational Science and Engineering, 2020, , 143-187.	0.3	1
16	Benefits from using mixed precision computations in the ELPA-AEO and ESSEX-II eigensolver projects. Japan Journal of Industrial and Applied Mathematics, 2019, 36, 699-717.	0.9	10
17	Automatic Throughput and Critical Path Analysis of x86 and ARM Assembly Kernels. , 2019, , .		12
18	Propagation and Decay of Injected One-Off Delays on Clusters: A Case Study. , 2019, , .		11

#	ARTICLE	IF	CITATIONS
19	CRAFT: A Library for Easier Application-Level Checkpoint/Restart and Automatic Fault Tolerance. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 501-514.	5.6	30
20	Building and utilizing fault tolerance support tools for the GASPI applications. International Journal of High Performance Computing Applications, 2018, 32, 613-626.	3.7	2
21	Multicore Performance Engineering of Sparse Triangular Solves Using a Modified Roofline Model. , 2018, , .		6
22	Automated Instruction Stream Throughput Prediction for Intel and AMD Microarchitectures. , 2018, , .		29
23	Chebyshev Filter Diagonalization on Modern Manycore Processors and GPGPUs. Lecture Notes in Computer Science, 2018, , 329-349.	1.3	10
24	On the Accuracy and Usefulness of Analytic Energy Models for Contemporary Multicore Processors. Lecture Notes in Computer Science, 2018, , 22-43.	1.3	13
25	Efficient optical simulation of nano structures in thin-film solar cells. , 2018, , .		1
26	GHOST: Building Blocks for High Performance Sparse Linear Algebra on Heterogeneous Systems. International Journal of Parallel Programming, 2017, 45, 1046-1072.	1.5	22
27	Performance analysis of the Kahanâ€enhanced scalar product on current multiâ€core and manyâ€core processors. Concurrency Computation Practice and Experience, 2017, 29, e3921.	2.2	5
28	LIKWID Monitoring Stack: A Flexible Framework Enabling Job Specific Performance monitoring for the masses. , 2017, , .		15
29	Multidimensional Intratile Parallelization for Memory-Starved Stencil Computations. ACM Transactions on Parallel Computing, 2017, 4, 1-32.	1.4	20
30	Kerncraft: A Tool for Analytic Performance Modeling of Loop Kernels. , 2017, , 1-22.		18
31	An Analysis of Core- and Chip-Level Architectural Features in Four Generations of Intel Server Processors. Lecture Notes in Computer Science, 2017, , 294-314.	1.3	16
32	Improved Coefficients for Polynomial Filtering in ESSEX. Lecture Notes in Computational Science and Engineering, 2017, , 63-79.	0.3	3
33	Chipâ€level and multiâ€node analysis of energyâ€optimized lattice Boltzmann CFD simulations. Concurrency Computation Practice and Experience, 2016, 28, 2295-2315.	2.2	21
34	Optimization of an Electromagnetics Code with Multicore Wavefront Diamond Blocking and Multi-dimensional Intra-Tile Parallelization. , 2016, , .		6
35	High-performance implementation of Chebyshev filter diagonalization for interior eigenvalue computations. Journal of Computational Physics, 2016, 325, 226-243.	3.8	28
36	Towards an Exascale Enabled Sparse Solver Repository. Lecture Notes in Computational Science and Engineering, 2016, , 295-316.	0.3	3

#	ARTICLE	IF	CITATIONS
37	Performance Engineering and Energy Efficiency of Building Blocks for Large, Sparse Eigenvalue Computations on Heterogeneous Supercomputers. Lecture Notes in Computational Science and Engineering, 2016, , 317-338.	0.3	5
38	Exploring performance and power properties of modern multi-core chips via simple machine models. Concurrency Computation Practice and Experience, 2016, 28, 189-210.	2.2	56
39	Performance and power for highly parallel systems. Concurrency Computation Practice and Experience, 2016, 28, 187-188.	2.2	0
40	Analysis of Intel's Haswell Microarchitecture Using the ECM Model and Microbenchmarks. Lecture Notes in Computer Science, 2016, , 210-222.	1.3	9
41	Performance Analysis of the Kahan-Enhanced Scalar Product on Current Multicore Processors. Lecture Notes in Computer Science, 2016, , 63-73.	1.3	4
42	Validation of Hardware Events for Successful Performance Pattern Identification in High Performance Computing. , 2016, , 17-28.		1
43	Automatic loop kernel analysis and performance modeling with Kerncraft. , 2015, , .		13
44	Multicore-Optimized Wavefront Diamond Blocking for Optimizing Stencil Updates. SIAM Journal of Scientific Computing, 2015, 37, C439-C464.	2.8	57
45	Electron confinement in graphene with gate-defined quantum dots. Physica Status Solidi (B): Basic Research, 2015, 252, 1868-1871.	1.5	18
46	Building a Fault Tolerant Application Using the GASPI Communication Layer. , 2015, , .		5
47	Increasing the Performance of the Jacobi-Davidson Method by Blocking. SIAM Journal of Scientific Computing, 2015, 37, C697-C722.	2.8	20
48	Performance Engineering of the Kernel Polynomial Method on Large-Scale CPU-GPU Systems. , 2015, , .		11
49	Quantifying Performance Bottlenecks of Stencil Computations Using the Execution-Cache-Memory Model. , 2015, , .		65
50	Domain-Specific Optimization of Two Jacobi Smoother Kernels and Their Evaluation in the ECM Performance Model. Parallel Processing Letters, 2014, 24, 1441004.	0.6	2
51	Comparing the performance of different x86 SIMD instruction sets for a medical imaging application on modern multi- and manycore chips. , 2014, , .		27
52	Overhead Analysis of Performance Counter Measurements. , 2014, , .		11
53	A Unified Sparse Matrix Data Format for Efficient General Sparse Matrix-Vector Multiplication on Modern Processors with Wide SIMD Units. SIAM Journal of Scientific Computing, 2014, 36, C401-C423.	2.8	160
54	ESSEX: Equipping Sparse Solvers for Exascale. Lecture Notes in Computer Science, 2014, , 577-588.	1.3	6

#	ARTICLE	IF	CITATIONS
55	Domain decomposition and locality optimization for large-scale lattice Boltzmann simulations. Computers and Fluids, 2013, 80, 283-289.	2.5	6
56	An Evaluation of Different I/O Techniques for Checkpoint/Restart. , 2013, , .		11
57	Comparison of different propagation steps for lattice Boltzmann methods. Computers and Mathematics With Applications, 2013, 65, 924-935.	2.7	72
58	Model-guided performance analysis of the sparse matrix-matrix multiplication. , 2013, , .		1
59	Performance engineering for the lattice Boltzmann method on GPGPUs: Architectural requirements and performance results. Computers and Fluids, 2013, 80, 276-282.	2.5	49
60	A SURVEY OF CHECKPOINT/RESTART TECHNIQUES ON DISTRIBUTED MEMORY SYSTEMS. Parallel Processing Letters, 2013, 23, 1340011.	0.6	11
61	Pushing the limits for medical image reconstruction on recent standard multicore processors. International Journal of High Performance Computing Applications, 2013, 27, 162-177.	3.7	20
62	Performance Patterns and Hardware Metrics on Modern Multicore Processors: Best Practices for Performance Engineering. Lecture Notes in Computer Science, 2013, , 451-460.	1.3	10
63	Performance Engineering: From Numbers to Insight. Lecture Notes in Computer Science, 2013, , 393-394.	1.3	0
64	Sparse Matrix-vector Multiplication on GPGPU Clusters: A New Storage Format and a Scalable Implementation. , 2012, , .		32
65	Expression Templates Revisited: A Performance Analysis of Current Methodologies. SIAM Journal of Scientific Computing, 2012, 34, C42-C69.	2.8	45
66	High performance smart expression template math libraries. , 2012, , .		23
67	likwid-bench: An Extensible Microbenchmarking Platform for x86 Multicore Compute Nodes. , 2012, , 27-36.		9
68	A flexible Patch-based lattice Boltzmann parallelization approach for heterogeneous GPUâ€CPU clusters. Parallel Computing, 2011, 37, 536-549.	2.1	49
69	Efficient multicore-aware parallelization strategies for iterative stencil computations. Journal of Computational Science, 2011, 2, 130-137.	2.9	35
70	Performance analysis and optimization strategies for a D3Q19 lattice Boltzmann kernel on nVIDIA GPUs using CUDA. Advances in Engineering Software, 2011, 42, 266-272.	3.8	42
71	Parallel Sparse Matrix-Vector Multiplication as a Test Case for Hybrid MPI+OpenMP Programming. , 2011, , .		7
72	HYBRID-PARALLEL SPARSE MATRIX-VECTOR MULTIPLICATION WITH EXPLICIT COMMUNICATION OVERLAP ON CURRENT MULTICORE-BASED SYSTEMS. Parallel Processing Letters, 2011, 21, 339-358.	0.6	23

#	ARTICLE	IF	CITATIONS
73	LIKWID: Lightweight Performance Tools. , 2011, , 165-175.		12
74	LIKWID: A Lightweight Performance-Oriented Tool Suite for x86 Multicore Environments. , 2010, , .		340
75	Multicore-aware parallel temporal blocking of stencil codes for shared and distributed memory. , 2010, , .		33
76	LEVERAGING SHARED CACHES FOR PARALLEL TEMPORAL BLOCKING OF STENCIL CODES ON MULTICORE PROCESSORS AND CLUSTERS. Parallel Processing Letters, 2010, 20, 359-376.	0.6	17
77	Luttinger, Peierls or Mott? Quantum Phase Transitions in Strongly Correlated 1D Electronâ€“Phonon Systems. Springer Series in Materials Science, 2010, , 1-21.	0.6	0
78	Introducing a Performance Model for Bandwidth-Limited Loop Kernels. Lecture Notes in Computer Science, 2010, , 615-624.	1.3	28
79	Performance Limitations for Sparse Matrix-Vector Multiplications on Current Multi-Core Environments. , 2010, , 13-26.		5
80	The world's fastest CPU and SMP node: Some performance results from the NEC SX-9. , 2009, , .		7
81	Quantum Phase Transition in a 1D Transport Model with Boson-Affected Hopping: Luttinger Liquid versus Charge-Density-Wave Behavior. Physical Review Letters, 2009, 102, 106404.	7.8	19
82	BENCHMARK ANALYSIS AND APPLICATION RESULTS FOR LATTICE BOLTZMANN SIMULATIONS ON NEC SX VECTOR AND INTEL NEHALEM SYSTEMS. Parallel Processing Letters, 2009, 19, 491-511.	0.6	14
83	Hybrid MPI/OpenMP Parallel Programming on Clusters of Multi-Core SMP Nodes. , 2009, , .		245
84	Efficient Temporal Blocking for Stencil Computations by Multicore-Aware Wavefront Parallelization. , 2009, , .		77
85	Challenges and Potentials of Emerging Multicore Architectures. , 2009, , 551-566.		8
86	Vector Computers in a World of Commodity Clusters, Massively Parallel Systems and Many-Core Many-Threaded CPUs: Recent Experience Based on an Advanced Lattice Boltzmann Flow Solver. , 2009, , 333-347.		3
87	RZBENCH: Performance Evaluation of Current HPC Architectures Using Low-Level and Application Benchmarks. , 2009, , 485-501.		1
88	Architecture and Performance Characteristics of Modern High Performance Computers. , 2008, , 681-730.		2
89	Data access optimizations for highly threaded multi-core CPUs with multiple memory controllers. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	19
90	DATA ACCESS CHARACTERISTICS AND OPTIMIZATIONS FOR SUN ULTRASPARC T2 AND T2+ SYSTEMS. Parallel Processing Letters, 2008, 18, 471-490.	0.6	0

#	ARTICLE	IF	CITATIONS
91	Metallicity in the half-filled Holstein-Hubbard model. Europhysics Letters, 2008, 84, 57001.	2.0	74
92	Introducing a parallel cache oblivious blocking approach for the lattice Boltzmann method. Progress in Computational Fluid Dynamics, 2008, 8, 179.	0.2	28
93	Direct Numerical Simulation of Turbulent Flow Over Dimples – Code Optimization for NEC SX-8 plus Flow Results. , 2008, , 303-318.		5
94	Carrier-density effects in many-polaron systems. Journal of Physics Condensed Matter, 2007, 19, 255202.	1.8	3
95	The spin-Peierls chain revisited. Journal of Magnetism and Magnetic Materials, 2007, 310, 1380-1382.	2.3	5
96	On the single processor performance of simple lattice Boltzmann kernels. Computers and Fluids, 2006, 35, 910-919.	2.5	167
97	Hole-doped Hubbard ladders. Physica B: Condensed Matter, 2006, 378-380, 319-320.	2.7	2
98	Phase diagram of the spin-Peierls chain with local coupling: Density-matrix renormalization-group calculations and unitary transformations. Physical Review B, 2006, 74, .	3.2	15
99	Have the Vectors the Continuing Ability to Parry the Attack of the Killer Micros?. , 2006, , 25-37.		5
100	Optimizing performance on modern HPC systems: learning from simple kernel benchmarks. , 2006, , 273-287.		5
101	DMRG Investigation of Stripe Formation in Doped Hubbard Ladders. , 2005, , 339-347.		0
102	Luttinger liquid versus charge density wave behaviour in the one-dimensional spinless fermion Holstein model. Physica B: Condensed Matter, 2005, 359-361, 699-701.	2.7	11
103	Performance of Scientific Applications on Modern Supercomputers. , 2005, , 3-25.		3
104	Exact Numerical Treatment of Finite Quantum Systems Using Leading-Edge Supercomputers. , 2005, , 165-177.		0
105	Stripe formation in doped Hubbard ladders. Physical Review B, 2005, 71, .	3.2	62
106	cxHPC: Setting up ByGRID – First Steps Towards an e-Science Infrastructure in Bavaria. , 2005, , 97-102.		0
107	Quantum lattice dynamical effects on single-particle excitations in one-dimensional Mott and Peierls insulators. Physical Review B, 2004, 69, .	3.2	57
108	Parallelization strategies for density matrix renormalization group algorithms on shared-memory systems. Journal of Computational Physics, 2004, 194, 795-808.	3.8	37

#	ARTICLE	IF	CITATIONS
109	Fast Sparse Matrix-Vector Multiplication for TeraFlop/s Computers. Lecture Notes in Computer Science, 2003, , 287-301.	1.3	7
110	One-Dimensional Electron-Phonon Systems: Mott- Versus Peierls-Insulators. , 2003, , 339-349.		1
111	Pseudo-Vectorization and RISC Optimization Techniques for the Hitachi SR8000 Architecture. , 2003, , 425-442.		1
112	Introduction to High Performance Computing for Scientists and Engineers. , 0, , .		173