Jack Dongarra

List of Publications by Year
in descending order

Source: https://exaly.com/author-pdf/8735190/publications.pdf

Version: 2024-02-01

535 papers	22,357 citations	29994 54 h-index	24179 110 g-index
576 all docs	576 docs citations	576 times ranked	8930 citing authors

#	Article	IF	Citations
1	A set of level 3 basic linear algebra subprograms. ACM Transactions on Mathematical Software, 1990, 16, 1-17.	1.6	1,328
2	PVM. , 1994, , .		1,313
3	Automated empirical optimizations of software and the ATLAS project. Parallel Computing, 2001, 27, 3-35.	1.3	928
4	An extended set of FORTRAN basic linear algebra subprograms. ACM Transactions on Mathematical Software, 1988, 14, 1-17.	1.6	655
5	An updated set of basic linear algebra subprograms (BLAS). ACM Transactions on Mathematical Software, 2002, 28, 135-151.	1.6	504
6	The International Exascale Software Project roadmap. International Journal of High Performance Computing Applications, 2011, 25, 3-60.	2.4	495
7	A class of parallel tiled linear algebra algorithms for multicore architectures. Parallel Computing, 2009, 35, 38-53.	1.3	327
8	Exascale computing and big data. Communications of the ACM, 2015, 58, 56-68.	3.3	322
9	Towards dense linear algebra for hybrid GPU accelerated manycore systems. Parallel Computing, 2010, 36, 232-240.	1.3	295
10	Netsolve: a Network-Enabled Server for Solving Computational Science Problems. International Journal of High Performance Computing Applications, 1997, 11, 212-223.	1.6	265
11	Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects. Journal of Physics: Conference Series, 2009, 180, 012037.	0.3	245
12	The GrADS Project: Software Support for High-Level Grid Application Development. International Journal of High Performance Computing Applications, 2001, 15, 327-344.	2.4	239
13	Chebyshev tau-QZ algorithm methods for calculating spectra of hydrodynamic stability problems. Applied Numerical Mathematics, 1996, 22, 399-434.	1.2	236
14	Implementing Linear Algebra Algorithms for Dense Matrices on a Vector Pipeline Machine. SIAM Review, 1984, 26, 91-112.	4.2	235
15	Distribution of mathematical software via electronic mail. Communications of the ACM, 1987, 30, 403-407.	3.3	230
16	The PVM concurrent computing system: Evolution, experiences, and trends. Parallel Computing, 1994, 20, 531-545.	1.3	207
17	From CUDA to OpenCL: Towards a performance-portable solution for multi-platform GPU programming. Parallel Computing, 2012, 38, 391-407.	1.3	198
18	DAGuE: A generic distributed DAG engine for High Performance Computing. Parallel Computing, 2012, 38, 37-51.	1.3	196

#	Article	IF	Citations
19	PaRSEC: Exploiting Heterogeneity to Enhance Scalability. Computing in Science and Engineering, 2013, 15, 36-45.	1.2	192
20	NetSolve. , 1996, , .		191
21	Algorithm 679: A set of level 3 basic linear algebra subprograms: model implementation and test programs. ACM Transactions on Mathematical Software, 1990, 16, 18-28.	1.6	175
22	Collecting Performance Data with PAPI-C. , 2010, , 157-173.		174
23	Algorithm-based fault tolerance applied to high performance computing. Journal of Parallel and Distributed Computing, 2009, 69, 410-416.	2.7	172
24	Dense linear algebra solvers for multicore with GPU accelerators. , 2010, , .		147
25	Post-failure recovery of MPI communication capability. International Journal of High Performance Computing Applications, 2013, 27, 244-254.	2.4	136
26	Accelerating scientific computations with mixed precision algorithms. Computer Physics Communications, 2009, 180, 2526-2533.	3.0	134
27	An Improved Magma Gemm For Fermi Graphics Processing Units. International Journal of High Performance Computing Applications, 2010, 24, 511-515.	2.4	132
28	New Grid Scheduling and Rescheduling Methods in the GrADS Project. International Journal of Parallel Programming, 2005, 33, 209-229.	1.1	127
29	Performance analysis of MPI collective operations. Cluster Computing, 2007, 10, 127-143.	3.5	126
30	Condition Numbers of Gaussian Random Matrices. SIAM Journal on Matrix Analysis and Applications, 2005, 27, 603-620.	0.7	116
31	Overview of GridRPC: A Remote Procedure Call API for Grid Computing. Lecture Notes in Computer Science, 2002, , 274-278.	1.0	114
32	ScaLAPACK: a portable linear algebra library for distributed memory computers â€" design issues and performance. Computer Physics Communications, 1996, 97, 1-15.	3.0	112
33	A Note on Auto-tuning GEMM for GPUs. Lecture Notes in Computer Science, 2009, , 884-892.	1.0	107
34	Software Libraries for Linear Algebra Computations on High Performance Computers. SIAM Review, 1995, 37, 151-180.	4.2	104
35	Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers. , 2018, , .		104
36	Block reduction of matrices to condensed forms for eigenvalue computations. Journal of Computational and Applied Mathematics, 1989, 27, 215-227.	1.1	100

#	Article	IF	Citations
37	Integrated Pvm Framework Supports Heterogeneous Network Computing. Computers in Physics, 1993, 7, 166-175.	0.6	99
38	Mixed Precision Iterative Refinement Techniques for the Solution of Dense Linear Systems. International Journal of High Performance Computing Applications, 2007, 21, 457-466.	2.4	99
39	Pumma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers. Concurrency and Computation: Practice and Experience, 1994, 6, 543-570.	0.6	98
40	Flexible Development of Dense Linear Algebra Algorithms on Massively Parallel Architectures with DPLASMA. , 2011, , .		95
41	Keeneland: Bringing Heterogeneous GPU Computing to the Computational Science Community. Computing in Science and Engineering, 2011, 13, 90-95.	1.2	93
42	High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems. International Journal of High Performance Computing Applications, 2016, 30, 3-10.	2.4	93
43	Parallel tiled QR factorization for multicore architectures. Concurrency Computation Practice and Experience, 2008, 20, 1573-1590.	1.4	92
44	Algorithm-Based Fault Tolerance for Fail-Stop Failures. IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 1628-1641.	4.0	92
45	Big data and extreme-scale computing. International Journal of High Performance Computing Applications, 2018, 32, 435-479.	2.4	90
46	Unrolling loops in fortran. Software - Practice and Experience, 1979, 9, 219-226.	2.5	88
47	Autotuning GEMM Kernels for the Fermi GPU. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 2045-2057.	4.0	85
48	Scheduling workflow applications on processors with different capabilities. Future Generation Computer Systems, 2006, 22, 665-675.	4.9	79
49	Solving Systems of Linear Equations on the CELL Processor Using Cholesky Factorization. IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 1175-1186.	4.0	78
50	High-performance bidiagonal reduction using tile algorithms on homogeneous multicore architectures. ACM Transactions on Mathematical Software, 2013, 39, 1-22.	1.6	74
51	Improving the Accuracy of Computed Eigenvalues and Eigenvectors. SIAM Journal on Numerical Analysis, 1983, 20, 23-45.	1.1	73
52	Autotuning in High-Performance Computing Applications. Proceedings of the IEEE, 2018, 106, 2068-2083.	16.4	72
53	Visualization and debugging in a heterogeneous environment. Computer, 1993, 26, 88-95.	1.2	70
54	Performance, Design, and Autotuning of Batched GEMM for GPUs. Lecture Notes in Computer Science, 2016, , 21-38.	1.0	70

#	Article	IF	Citations
55	On some parallel banded system solvers. Parallel Computing, 1984, 1, 223-235.	1.3	69
56	Preface: Basic Linear Algebra Subprograms Technical (Blast) Forum Standard. International Journal of High Performance Computing Applications, 2002, 16, 1-1.	2.4	69
57	Exploiting the Performance of 32 bit Floating Point Arithmetic in Obtaining 64 bit Accuracy (Revisiting) Tj ETQq1 1 0.784314 rgBT /Ov		69
58	Applying NetSolve's network-enabled server. IEEE Computational Science and Engineering, 1998, 5, 57-67.	0.6	68
59	Algorithm-based fault tolerance for dense matrix factorizations. , 2012, , .		67
60	Using Mixed Precision for Sparse Matrix Computations to Enhance the Performance while Achieving 64-bit Accuracy. ACM Transactions on Mathematical Software, 2008, 34, 1-22.	1.6	65
61	QR Factorization on a Multicore Node Enhanced with Multiple GPU Accelerators. , 2011, , .		65
62	Enabling and scaling matrix computations on heterogeneous multi-core and multi-GPU systems. , 2012, , .		64
63	A Parallel Divide and Conquer Algorithm for the Symmetric Eigenvalue Problem on Distributed Memory Architectures. SIAM Journal of Scientific Computing, 1999, 20, 2223-2236.	1.3	63
64	Dynamic task scheduling for linear algebra algorithms on distributed-memory multicore systems. , 2009, , .		63
65	Standards for graph algorithm primitives. , 2013, , .		63
66	Numerical Libraries and the Grid. International Journal of High Performance Computing Applications, 2001, 15, 359-374.	2.4	61
67	Fault tolerant high performance computing by a coding approach. , 2005, , .		61
68	A survey of numerical linear algebra methods utilizing mixed-precision arithmetic. International Journal of High Performance Computing Applications, 2021, 35, 344-369.	2.4	61
69	A proposal for a set of parallel basic linear algebra subprograms. Lecture Notes in Computer Science, 1996, , 107-114.	1.0	59
70	Scheduling block-cyclic array redistribution. IEEE Transactions on Parallel and Distributed Systems, 1998, 9, 192-205.	4.0	59
71	DAGuE: A Generic Distributed DAG Engine for High Performance Computing. , 2011, , .		57
72	The Singular Value Decomposition: Anatomy of Optimizing an Algorithm for Extreme Scale. SIAM Review, 2018, 60, 808-865.	4.2	57

#	Article	IF	Citations
73	The Impact of Multicore on Math Software. , 2006, , 1-10.		57
74	Visual programming and debugging for parallel computing. IEEE Parallel and Distributed Technology, 1995, 3, 75-83.	0.7	56
75	HARNESS: a next generation distributed virtual machine. Future Generation Computer Systems, 1999, 15, 571-582.	4.9	56
76	The International Exascale Software Project: a Call To Cooperative Action By the Global High-Performance Community. International Journal of High Performance Computing Applications, 2009, 23, 309-322.	2.4	56
77	A tool to aid in the design, implementation, and understanding of matrix algorithms for parallel processors. Journal of Parallel and Distributed Computing, 1990, 9, 185-202.	2.7	55
78	Scalability Issues Affecting the Design of a Dense Linear Algebra Library. Journal of Parallel and Distributed Computing, 1994, 22, 523-537.	2.7	55
79	Optimizing matrix multiplication for a short-vector SIMD architecture â€“ CELL processor. Parallel Computing, 2009, 35, 138-150.	1.3	53
80	Scheduling dense linear algebra operations on multicore processors. Concurrency Computation Practice and Experience, 2010, 22, 15-44.	1.4	53
81	Accelerating the reduction to upper Hessenberg, tridiagonal, and bidiagonal forms through hybrid GPU-based computing. Parallel Computing, 2010, 36, 645-654.	1.3	53
82	Towards Efficient MapReduce Using MPI. Lecture Notes in Computer Science, 2009, , 240-249.	1.0	53
83	Accelerating Numerical Dense Linear Algebra Calculations with GPUs. , 2014, , 3-28.		52
84	Implementation of some concurrent algorithms for matrix factorization. Parallel Computing, 1986, 3, 25-34.	1.3	51
85	HARNESS and fault tolerant MPI. Parallel Computing, 2001, 27, 1479-1495.	1.3	50
86	The Design and Performance of Batched BLAS on Modern High-Performance Computing Systems. Procedia Computer Science, 2017, 108, 495-504.	1.2	50
87	The design and implementation of the parallel out-of-core ScaLAPACK LU, QR, and Cholesky factorization routines. Concurrency and Computation: Practice and Experience, 2000, 12, 1481-1493.	0.6	49
88	High-performance computing systems: Status and outlook. Acta Numerica, 2012, 21, 379-474.	6.3	48
89	Batched matrix computations on hardware accelerators based on GPUs. International Journal of High Performance Computing Applications, 2015, 29, 193-208.	2.4	48
90	Implementation of mixed precision in solving systems of linear equations on the Cell processor. Concurrency Computation Practice and Experience, 2007, 19, 1371-1385.	1.4	47

#	Article	IF	Citations
91	Solving banded systems on a parallel processor. Parallel Computing, 1987, 5, 219-246.	1.3	46
92	Iterative Sparse Triangular Solves for Preconditioning. Lecture Notes in Computer Science, 2015, , 650-661.	1.0	46
93	Reduction to condensed form for the Eigenvalue problem on distributed memory architectures. Parallel Computing, 1992, 18, 973-982.	1.3	45
94	Fault-Tolerant Matrix Operations for Networks of Workstations Using Diskless Checkpointing. Journal of Parallel and Distributed Computing, 1997, 43, 125-138.	2.7	44
95	Telescoping Languages: A Strategy for Automatic Generation of Scientific Problem-Solving Systems from Annotated Libraries. Journal of Parallel and Distributed Computing, 2001, 61, 1803-1826.	2.7	43
96	Dynamic task discovery in PaRSEC. , 2017, , .		43
97	Multiprocessing linear algebra algorithms on the CRAY X-MP-2: Experiences with small granularity. Journal of Parallel and Distributed Computing, 1984, 1, 22-31.	2.7	42
98	SLATE. , 2019, , .		42
99	Computer benchmarking: Paths and pitfalls: The most popular way of rating computer performance can confuse as well as inform; avoid misunderstanding by asking just what the benchmark is measuring. IEEE Spectrum, 1987, 24, 38-43.	0.5	41
100	Self-adapting software for numerical linear algebra and LAPACK for clusters. Parallel Computing, 2003, 29, 1723-1743.	1.3	41
101	Adaptive precision in blockâ€Jacobi preconditioning for iterative sparse linear system solvers. Concurrency Computation Practice and Experience, 2019, 31, e4460.	1.4	41
102	Unified model for assessing checkpointing protocols at extremeâ€scale. Concurrency Computation Practice and Experience, 2014, 26, 2772-2791.	1.4	40
103	Squeezing the most out of eigenvalue solvers on high-performance computers. Linear Algebra and Its Applications, 1986, 77, 113-136.	0.4	39
104	Recovery Patterns for Iterative Methods in a Parallel Unstable Environment. SIAM Journal of Scientific Computing, 2008, 30, 102-116.	1.3	38
105	A Parallel Algorithm for the Nonsymmetric Eigenvalue Problem. SIAM Journal of Scientific Computing, 1993, 14, 542-569.	1.3	37
106	Algorithmic redistribution methods for block-cyclic decompositions. IEEE Transactions on Parallel and Distributed Systems, 1999, 10, 1201-1216.	4.0	37
107	Energy Footprint of Advanced Dense Numerical Linear Algebra Using Tile Algorithms on Multicore Architectures. , 2012, , .		37
108	Investigating half precision arithmetic to accelerate dense linear system solvers. , 2017, , .		37

#	Article	IF	Citations
109	Innovations of the NetSolve Grid Computing System. Concurrency Computation Practice and Experience, 2002, 14, 1457-1479.	1.4	36
110	Mixed-Precision Cholesky QR Factorization and Its Case Studies on Multicore CPU with Multiple GPUs. SIAM Journal of Scientific Computing, 2015, 37, C307-C330.	1.3	36
111	Fault Tolerance Techniques for High-Performance Computing. Computer Communications and Networks, 2015, , 3-85.	0.8	36
112	Improving the Performance of CA-GMRES on Multicores with Multiple GPUs. , 2014, , .		35
113	A Step towards Energy Efficient Computing: Redesigning a Hydrodynamic Application on CPU-GPU. , 2014, , .		35
114	Incomplete Sparse Approximate Inverses for Parallel Preconditioning. Parallel Computing, 2018, 71, 1-22.	1.3	35
115	Hierarchical DAG Scheduling for Hybrid Distributed Systems. , 2015, , .		34
116	A portable environment for developing parallel FORTRAN programs. Parallel Computing, 1987, 5, 175-186.	1.3	33
117	MPI collective algorithm selection and quadtree encoding. Parallel Computing, 2007, 33, 613-623.	1.3	33
118	LU factorization for accelerator-based systems. , 2011, , .		33
119	A comparison of search heuristics for empirical code optimization. , 2008, , .		32
120	Netlib and NA-Net: Building a Scientific Computing Community. IEEE Annals of the History of Computing, 2008, 30, 30-41.	0.2	32
121	Highly Scalable Self-Healing Algorithms for High Performance Scientific Computing. IEEE Transactions on Computers, 2009, 58, 1512-1524.	2.4	32
122	Optimizing symmetric dense matrix-vector multiplication on GPUs. , 2011, , .		32
123	LU Factorization of Small Matrices: Accelerating Batched DGETRF on the GPU. , 2014, , .		32
124	NetSolve: Past, Present, and Future â€“ A Look at a Grid Enabled Server. , 0, , 615-624.		31
125	Redesigning the message logging model for high performance. Concurrency Computation Practice and Experience, 2010, 22, 2196-2211.	1.4	31
126	QR factorization of tall and skinny matrices in a grid computing environment. , 2010, , .		31

#	Article	IF	Citations
127	Parallel reduction to condensed forms for symmetric eigenvalue problems using aggregated fine-grained and memory-aware kernels. , 2011, , .		31
128	A scalable framework for heterogeneous GPU-based clusters. , 2012, , .		31
129	Algorithm-based fault tolerance for dense matrix factorizations. ACM SIGPLAN Notices, 2012, 47, 225-234.	0.2	31
130	PTG: An Abstraction for Unhindered Parallelism. , 2014, , .		31
131	The TOP500 List and Progress in High-Performance Computing. Computer, 2015, 48, 42-49.	1.2	31
132	Preconditioned Krylov solvers on GPUs. Parallel Computing, 2017, 68, 32-44.	1.3	31
133	Recent Developments in Gridsolve. International Journal of High Performance Computing Applications, 2006, 20, 131-141.	2.4	30
134	A parallel algorithm for the reduction of a nonsymmetric matrix to block upper-Hessenberg form. Parallel Computing, 1995, 21, 1189-1211.	1.3	29
135	Parallel matrix transpose algorithms on distributed memory concurrent computers. Parallel Computing, 1995, 21, 1387-1405.	1.3	29
136	The Spectral Decomposition of Nonsymmetric Matrices on Distributed Memory Parallel Computers. SIAM Journal of Scientific Computing, 1997, 18, 1446-1461.	1.3	29
137	ALGORITHMIC ISSUES ON HETEROGENEOUS COMPUTING PLATFORMS. Parallel Processing Letters, 1999, 09, 197-213.	0.4	29
138	High-performance high-resolution semi-Lagrangian tracer transport on a sphere. Journal of Computational Physics, 2011, 230, 6778-6799.	1.9	29
139	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
140	The Design of Fast and Energy-Efficient Linear Solvers: On the Potential of Half-Precision Arithmetic and Iterative Refinement Techniques. Lecture Notes in Computer Science, 2018, , 586-600.	1.0	29
141	A comparative study of automatic vectorizing compilers. Parallel Computing, 1991, 17, 1223-1244.	1.3	28
142	Numerical Considerations in Computing Invariant Subspaces. SIAM Journal on Matrix Analysis and Applications, 1992, 13, 145-161.	0.7	28
143	Two-Stage Tridiagonal Reduction for Dense Symmetric Matrices Using Tile Algorithms on Multicore Architectures. , 2011, , .		28
144	The Netlib Mathematical Software Repository. D-Lib Magazine, 1995, 1, .	0.5	28

#	Article	IF	Citations
145	A collection of parallel linear equations routines for the Denelcor HEP. Parallel Computing, 1984, 1, 133-142.	1.3	27
146	Java access to numerical libraries. Concurrency and Computation: Practice and Experience, 1997, 9, 1279-1291.	0.6	27
147	The marketplace of high-performance computing. Parallel Computing, 1999, 25, 1517-1544.	1.3	27
148	The quest for petascale computing. Computing in Science and Engineering, 2001, 3, 32-39.	1.2	27
149	A Parallel Implementation of the Nonsymmetric QR Algorithm for Distributed Memory Architectures. SIAM Journal of Scientific Computing, 2002, 24, 284-311.	1.3	27
150	Recent trends in the marketplace of high performance computing. Parallel Computing, 2005, 31, 261-273.	1.3	27
151	Sunway TaihuLight supercomputer makes its appearance. National Science Review, 2016, 3, 265-266.	4.6	27
152	Using Jacobi iterations and blocking for solving sparse triangular systems in incomplete factorization preconditioning. Journal of Parallel and Distributed Computing, 2018, 119, 219-230.	2.7	27
153	Investigating power capping toward energyâ€efficient scientific applications. Concurrency Computation Practice and Experience, 2019, 31, e4485.	1.4	27
154	Algorithmic bombardment for the iterative solution of linear systems: A poly-iterative approach. Journal of Computational and Applied Mathematics, 1996, 74, 91-109.	1.1	26
155	Accelerating GPU Kernels for Dense Linear Algebra. Lecture Notes in Computer Science, 2011, , 83-92.	1.0	26
156	An evaluation of User-Level Failure Mitigation support in MPI. Computing (Vienna/New York), 2013, 95, 1171-1184.	3.2	26
157	Race to Exascale. Computing in Science and Engineering, 2019, 21, 4-5.	1.2	26
158	Request Sequencing: Optimizing Communication for the Grid. Lecture Notes in Computer Science, 2000, , 1213-1222.	1.0	26
159	Mixed-precision iterative refinement using tensor cores on GPUs to accelerate solution of linear systems. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 2020, 476, 20200110.	1.0	26
160	Numerical linear algebra algorithms and software. Journal of Computational and Applied Mathematics, 2000, 123, 489-514.	1.1	25
161	State-of-the-art eigensolvers for electronic structure calculations of large scale nano-systems. Journal of Computational Physics, 2008, 227, 7113-7124.	1.9	25
162	Hierarchical QR factorization algorithms for multi-core clusters. Parallel Computing, 2013, 39, 212-232.	1.3	25

#	Article	IF	Citations
163	Achieving numerical accuracy and high performance using recursive tile LU factorization with partial pivoting. Concurrency Computation Practice and Experience, 2014, 26, 1408-1431.	1.4	25
164	Accelerating collaborative filtering using concepts from high performance computing. , 2015, , .		25
165	Porting the PLASMA Numerical Library to the OpenMP Standard. International Journal of Parallel Programming, 2017, 45, 612-633.	1.1	25
166	The eigenvalue problem for Hermitian matrices with time reversal symmetry. Linear Algebra and Its Applications, 1984, 60, 27-42.	0.4	24
167	Linear algebra on high performance computers. Applied Mathematics and Computation, 1986, 20, 57-88.	1.4	24
168	Parallel loops â€" A test suite for parallelizing compilers: Description and example results. Parallel Computing, 1991, 17, 1247-1255.	1.3	24
169	Tools and techniques for performance---Exploiting the performance of 32 bit floating point arithmetic in obtaining 64 bit accuracy (revisiting iterative refinement for linear systems). , 2006, , .		24
170	The PlayStation 3 for High-Performance Scientific Computing. Computing in Science and Engineering, 2008, 10, 84-87.	1.2	24
171	DARPA's HPCS Program: History, Models, Tools, Languages. Advances in Computers, 2008, , 1-100.	1.2	24
172	Scalable Tile Communication-Avoiding QR Factorization on Multicore Cluster Systems. , 2010, , .		24
173	A new metric for ranking high-performance computing systems. National Science Review, 2016, 3, 30-35.	4.6	24
174	Efficient exascale discretizations: High-order finite element methods. International Journal of High Performance Computing Applications, 2021, 35, 527-552.	2.4	24
175	Extreme-Scale Task-Based Cholesky Factorization Toward Climate and Weather Prediction Applications. , 2020, , .		24
176	Numerical libraries and the grid. , 2001, , .		23
177	QR Factorization for the Cell Broadband Engine. Scientific Programming, 2009, 17, 31-42.	0.5	23
178	HierKNEM: An Adaptive Framework for Kernel-Assisted and Topology-Aware Collective Communications on Many-core Clusters. , 2012, , .		23
179	A Fast Batched Cholesky Factorization on a GPU. , 2014, , .		23
180	A novel hybrid CPUâ€"GPU generalized eigensolver for electronic structure calculations based on fine-grained memory aware tasks. International Journal of High Performance Computing Applications, 2014, 28, 196-209.	2.4	23

#	Article	IF	Citations
181	Unified Development for Mixed Multi-GPU and Multi-coprocessor Environments Using a Lightweight Runtime Environment. , 2014, , .		23
182	Recent Enhancements To Pvm. International Journal of High Performance Computing Applications, 1995, 9, 108-127.	1.6	22
183	Static tiling for heterogeneous computing platforms. Parallel Computing, 1999, 25, 547-568.	1.3	22
184	The Virtual Instrument: Support for Grid-Enabled Mcell Simulations. International Journal of High Performance Computing Applications, 2004, 18, 3-17.	2.4	22
185	NetSolve: Grid enabling scientific computing environments. Advances in Parallel Computing, 2005, 14, 33-51.	0.3	22
186	Tile QR factorization with parallel panel processing for multicore architectures. , 2010, , .		22
187	High Performance Dense Linear System Solver with Soft Error Resilience. , 2011, , .		22
188	Tridiagonalization of a dense symmetric matrix on multiple GPUs and its application to symmetric eigenvalue problems. Concurrency Computation Practice and Experience, 2014, 26, 2652-2666.	1.4	22
189	Binomial Graph: A Scalable and Fault-Tolerant Logical Network Topology. Lecture Notes in Computer Science, 2007, , 471-482.	1.0	22
190	Load-balancing Sparse Matrix Vector Product Kernels on GPUs. ACM Transactions on Parallel Computing, 2020, 7, 1-26.	1.2	22
191	Exploring New Architectures in Accelerating CFD for Air Force Applications. , 2008, , .		21
192	A Class of Hybrid LAPACK Algorithms for Multicore and GPU Architectures. , 2011, , .		21
193	Analysis of dynamically scheduled tile algorithms for dense linear algebra on multicore architectures. Concurrency Computation Practice and Experience, 2012, 24, 305-321.	1.4	21
194	With Extreme Computing, the Rules Have Changed. Computing in Science and Engineering, 2017, 19, 52-62.	1.2	21
195	Dynamic Reconfiguration and Virtual Machine Management in the Harness Metacomputing System. Lecture Notes in Computer Science, 1998, , 127-134.	1.0	21
196	Key concepts for parallel out-of-core LU factorization. Parallel Computing, 1997, 23, 49-70.	1.3	20
197	Deploying fault tolerance and taks migration with NetSolve. Future Generation Computer Systems, 1999, 15, 745-755.	4.9	20
198	Numerically Stable Real Number Codes Based on Random Matrices. Lecture Notes in Computer Science, 2005, , 115-122.	1.0	20

#	Article	IF	Citations
199	High performance matrix inversion based on LU factorization for multicore architectures. , 2011, , .		20
200	Divide and Conquer on Hybrid GPU-Accelerated Multicore Systems. SIAM Journal of Scientific Computing, 2012, 34, C70-C82.	1.3	20
201	Accelerating Linear System Solutions Using Randomization Techniques. ACM Transactions on Mathematical Software, 2013, 39, 1-13.	1.6	20
202	Asynchronous Iterative Algorithm for Computing Incomplete Factorizations on GPUs. Lecture Notes in Computer Science, 2015, , 1-16.	1.0	20
203	ParILUT---A New Parallel Threshold ILU Factorization. SIAM Journal of Scientific Computing, 2018, 40, C503-C519.	1.3	20
204	Fast Batched Matrix Multiplication for Small Sizes Using Half-Precision Arithmetic on GPUs. , 2019, , .		20
205	A Framework for Batched and GPU-Resident Factorization Algorithms Applied to Block Householder Transformations. Lecture Notes in Computer Science, 2015, , 31-47.	1.0	20
206	The National HPCC Software Exchange. IEEE Computational Science and Engineering, 1995, 2, 62-69.	0.6	19
207	Improvement of parallelization efficiency of batch pattern BP training algorithm using Open MPI. Procedia Computer Science, 2010, 1, 525-533.	1.2	19
208	Profiling high performance dense linear algebra algorithms on multicore architectures for power and energy efficiency. Computer Science - Research and Development, 2012, 27, 277-287.	2.7	19
209	One-sided Dense Matrix Factorizations on a Multicore with Multiple GPU Accelerators*. Procedia Computer Science, 2012, 9, 37-46.	1.2	19
210	A block-asynchronous relaxation method for graphics processing units. Journal of Parallel and Distributed Computing, 2013, 73, 1613-1626.	2.7	19
211	Power monitoring with PAPI for extreme scale architectures and dataflow-based programming models. , 2014, , .		19
212	A survey of recent developments in parallel implementations of Gaussian elimination. Concurrency Computation Practice and Experience, 2015, 27, 1292-1309.	1.4	19
213	Design for a Soft Error Resilient Dynamic Task-Based Runtime. , 2015, , .		19
214	Towards Achieving Performance Portability Using Directives for Accelerators. , 2016, , .		19
215	Solving the secular equation including spin orbit coupling for systems with inversion and time reversal symmetry. Journal of Computational Physics, 1984, 54, 278-288.	1.9	18
216	Analytical modeling and optimization for affinity based thread scheduling on multicore systems. , 2009, , .		18

#	Article	IF	Citations
217	Self-healing network for scalable fault-tolerant runtime environments. Future Generation Computer Systems, 2010, 26, 479-485.	4.9	18
218	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
219	Implementing Linear Algebra Routines on Multi-core Processors with Pipelining and a Look Ahead. , 2006, , 147-156.		18
220	A Checkpoint-on-Failure Protocol for Algorithm-Based Recovery in Standard MPI. Lecture Notes in Computer Science, 2012, , 477-488.	1.0	18
221	Composing resilience techniques: ABFT, periodic and incremental checkpointing. International Journal of Networking and Computing, 2015, 5, 2-25.	0.3	18
222	Accurate Cache and TLB Characterization Using Hardware Counters. Lecture Notes in Computer Science, 2004, , 432-439.	1.0	17
223	Biological sequence alignment on the computational grid using the GrADS framework. Future Generation Computer Systems, 2005, 21, 980-986.	4.9	17
224	Acceleration of GPU-based Krylov solvers via data transfer reduction. International Journal of High Performance Computing Applications, 2015, 29, 366-383.	2.4	17
225	Algorithm-Based Fault Tolerance for Dense Matrix Factorizations, Multiple Failures and Accuracy. ACM Transactions on Parallel Computing, 2015, 1, 1-28.	1.2	17
226	Batched Gauss-Jordan Elimination for Block-Jacobi Preconditioner Generation on GPUs. , 2017, , .		17
227	On the performance and energy efficiency of sparse linear algebra on GPUs. International Journal of High Performance Computing Applications, 2017, 31, 375-390.	2.4	17
228	Retrospect: Deterministic Replay of MPI Applications for Interactive Distributed Debugging. Lecture Notes in Computer Science, 2007, , 297-306.	1.0	17
229	Evaluation of the HPC Challenge Benchmarks in Virtualized Environments. Lecture Notes in Computer Science, 2012, , 436-445.	1.0	17
230	Tools to aid in the analysis of memory access patterns for FORTRAN programs. Parallel Computing, 1988, 9, 25-35.	1.3	16
231	Programming methodology and performance issues for advanced computer architectures. Parallel Computing, 1988, 8, 41-58.	1.3	16
232	A Scalable Checkpoint Encoding Algorithm for Diskless Checkpointing. , 2008, , .		16
233	Matrix product on heterogeneous master-worker platforms. , 2008, , .		16
234	High Performance Dense Linear System Solver with Resilience to Multiple Soft Errors. Procedia Computer Science, 2012, 9, 216-225.	1.2	16

#	Article	IF	Citations
235	A Comprehensive Study of Task Coalescing for Selecting Parallelism Granularity in a Two-Stage Bidiagonal Reduction. , 2012, , .		16
236	Experiences in autotuning matrix multiplication for energy minimization on GPUs. Concurrency Computation Practice and Experience, 2015, 27, 5096-5113.	1.4	16
237	A Guide for Achieving High Performance with Very Small Matrices on GPU: A Case Study of Batched LU and Cholesky Factorizations. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 973-984.	4.0	16
238	ADAPT. , 2018, , .		16
239	A Set of Batched Basic Linear Algebra Subprograms and LAPACK Routines. ACM Transactions on Mathematical Software, 2021, 47, 1-23.	1.6	16
240	L2 Cache Modeling for Scientific Applications on Chip Multi-Processors. Parallel Processing (ICPP), Proceedings of the International Symposium, 2007, , .	0.0	15
241	Computing the conditioning of the components of a linear leastâ€squares solution. Numerical Linear Algebra With Applications, 2009, 16, 517-533.	0.9	15
242	Kernel-assisted and topology-aware MPI collective communications on multicore/many-core platforms. Journal of Parallel and Distributed Computing, 2013, 73, 1000-1010.	2.7	15
243	Soft error resilient QR factorization for hybrid system with GPGPU. Journal of Computational Science, 2013, 4, 457-464.	1.5	15
244	Portable HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi. Lecture Notes in Computer Science, 2014, , 571-581.	1.0	15
245	PaRSEC in Practice: Optimizing a Legacy Chemistry Application through Distributed Task-Based Execution. , 2015, , .		15
246	Factorization and Inversion of a Million Matrices using GPUs: Challenges and Countermeasures. Procedia Computer Science, 2017, 108, 606-615.	1.2	15
247	Novel HPC techniques to batch execution of many variable size BLAS computations on GPUs. , 2017, , .		15
248	Performance of asynchronous optimized Schwarz with one-sided communication. Parallel Computing, 2019, 86, 66-81.	1.3	15
249	Numerical algorithms for high-performance computational science. Philosophical Transactions Series A, Mathematical, Physical, and Engineering Sciences, 2020, 378, 20190066.	1.6	15
250	JLAPACK â€“ Compiling LAPACK FORTRAN to Java. Scientific Programming, 1999, 7, 111-138.	0.5	14
251	Block-asynchronous Multigrid Smoothers for GPU-accelerated Systems. Procedia Computer Science, 2012, 9, 7-16.	1.2	14
252	Updating incomplete factorization preconditioners for model order reduction. Numerical Algorithms, 2016, 73, 611-630.	1.1	14

#	Article	IF	Citations
253	Power-aware computing: Measurement, control, and performance analysis for Intel Xeon Phi. , 2017, , .		14
254	Impacts of Multi-GPU MPI Collective Communications on Large FFT Computation. , 2019, , .		14
255	Generic Matrix Multiplication for Multi-GPU Accelerated Distributed-Memory Platforms over PaRSEC. , 2019, , .		14
256	Accelerating Geostatistical Modeling and Prediction With Mixed-Precision Computations: A High-Productivity Approach With PaRSEC. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 964-976.	4.0	14
257	GridSolve: The Evolution of A Network Enabled Solver. , 2007, , 215-224.		14
258	Recursive Approach in Sparse Matrix LU Factorization. Scientific Programming, 2001, 9, 51-60.	0.5	13
259	Automatic analysis of inefficiency patterns in parallel applications. Concurrency Computation Practice and Experience, 2007, 19, 1481-1496.	1.4	13
260	Toward a High Performance Tile Divide and Conquer Algorithm for the Dense Symmetric Eigenvalue Problem. SIAM Journal of Scientific Computing, 2012, 34, C249-C274.	1.3	13
261	A Class of Communication-avoiding Algorithms for Solving General Dense Linear Systems on CPU/GPU Parallel Machines. Procedia Computer Science, 2012, 9, 17-26.	1.2	13
262	A Parallel Tiled Solver for Dense Symmetric Indefinite Systems on Multicore Architectures. , 2012, , .		13
263	A look back on 30 years of the Gordon Bell Prize. International Journal of High Performance Computing Applications, 2017, 31, 469-484.	2.4	13
264	Improving Performance of GMRES by Reducing Communication and Pipelining Global Collectives. , 2017, , .		13
265	Distributed-memory lattice H-matrix factorization. International Journal of High Performance Computing Applications, 2019, 33, 1046-1063.	2.4	13
266	Solving Linear Diophantine Systems on Parallel Architectures. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 1158-1169.	4.0	13
267	Accelerating Restarted GMRES With Mixed Precision Arithmetic. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 1027-1037.	4.0	13
268	Automatic blocking of QR and LU factorizations for locality. , 2004, , .		12
269	An asynchronous algorithm on the NetSolve global computing system. Future Generation Computer Systems, 2006, 22, 279-290.	4.9	12
270	Multithreading for synchronization tolerance in matrix factorization. Journal of Physics: Conference Series, 2007, 78, 012028.	0.3	12

#	Article	IF	Citations
271	A Parallel Solver for Incompressible Fluid Flows. Procedia Computer Science, 2013, 18, 439-448.	1.2	12
272	Optimizing Krylov Subspace Solvers on Graphics Processing Units. , 2014, , .		12
273	An efficient distributed randomized algorithm for solving large dense symmetric indefinite linear systems. Parallel Computing, 2014, 40, 213-223.	1.3	12
274	Practical scalable consensus for pseudo-synchronous distributed systems. , 2015, , .		12
275	HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi. Scientific Programming, 2015, 2015, 1-11.	0.5	12
276	Strengthening compute and data intensive capacities of Armenia. , 2015, , .		12
277	Towards batched linear solvers on accelerated hardware platforms. , 2015, , .		12
278	Failure Detection and Propagation in HPC systems. , 2016, , .		12
279	Heterogeneous Streaming. , 2016, , .		12
280	Accelerating the SVD bi-diagonalization of a batch of small matrices using GPUs. Journal of Computational Science, 2018, 26, 237-245.	1.5	12
281	A failure detector for HPC platforms. International Journal of High Performance Computing Applications, 2018, 32, 139-158.	2.4	12
282	Performance Analysis of Tile Low-Rank Cholesky Factorization Using PaRSEC Instrumentation Tools. , 2019, , .		12
283	A Scalable Approach to MPI Application Performance Analysis. Lecture Notes in Computer Science, 2005, , 309-316.	1.0	12
284	Performance Instrumentation and Measurement for Terascale Systems. Lecture Notes in Computer Science, 2003, , 53-62.	1.0	12
285	A Scalable High Performant Cholesky Factorization for Multicore with GPU Accelerators. Lecture Notes in Computer Science, 2011, , 93-101.	1.0	12
286	HAN: a Hierarchical AutotuNed Collective Communication Framework. , 2020, , .		12
287	Software distribution using Xnetlib. ACM Transactions on Mathematical Software, 1995, 21, 79-88.	1.6	11
288	Enabling interactive and collaborative oil reservoir simulations on the Grid. Concurrency Computation Practice and Experience, 2005, 17, 1387-1414.	1.4	11

#	Article	IF	Citations
289	Anatomy of a globally recursive embedded LINPACK benchmark. , 2012, , .		11
290	Multi-GPU Implementation of LU Factorization. Procedia Computer Science, 2012, 9, 106-115.	1.2	11
291	Hierarchical QR Factorization Algorithms for Multi-core Cluster Systems. , 2012, , .		11
292	Power profiling of Cholesky and QR factorizations on distributed memory systems. Computer Science - Research and Development, 2014, 29, 139-147.	2.7	11
293	Communication-Avoiding Symmetric-Indefinite Factorization. SIAM Journal on Matrix Analysis and Applications, 2014, 35, 1364-1406.	0.7	11
294	clMAGMA. , 2014, , .		11
295	MAGMA embedded: Towards a dense linear algebra library for energy efficient extreme computing. , 2015, , .		11
296	Linear algebra software for large-scale accelerated multicore computing. Acta Numerica, 2016, 25, 1-160.	6.3	11
297	Batched Generation of Incomplete Sparse Approximate Inverses on GPUs. , 2016, , .		11
298	Fast Cholesky factorization on GPUs for batch and native modes in MAGMA. Journal of Computational Science, 2017, 20, 85-93.	1.5	11
299	High-performance Cholesky factorization for GPU-only execution. , 2017, , .		11
300	PAPI software-defined events for in-depth performance analysis. International Journal of High Performance Computing Applications, 2019, 33, 1113-1127.	2.4	11
301	On Using Incremental Profiling for the Performance Analysis of Shared Memory Parallel Applications. Lecture Notes in Computer Science, 2007, , 62-71.	1.0	11
302	A Framework for Out of Memory SVD Algorithms. Lecture Notes in Computer Science, 2017, , 158-178.	1.0	11
303	Assessing the Impact of ABFT and Checkpoint Composite Strategies. , 2014, , .		10
304	Efficient implementation of quantum materials simulations on distributed CPU-GPU systems. , 2015, , .		10
305	Variable-size batched Gaussâ€“Jordan elimination for block-Jacobi preconditioning on graphics processors. Parallel Computing, 2019, 81, 131-146.	1.3	10
306	Matrix multiplication on batches of small matrices in half and half-complex precisions. Journal of Parallel and Distributed Computing, 2020, 145, 188-201.	2.7	10

#	Article	IF	Citations
307	Power Management and Event Verification in PAPI. , 2016, , 41-51.		10
308	Enhancing Parallelism of Tile Bidiagonal Transformation on Multicore Architectures Using Tree Reduction. Lecture Notes in Computer Science, 2012, , 661-670.	1.0	10
309	Leading Edge Hybrid Multi-GPU Algorithms for Generalized Eigenproblems in Electronic Structure Calculations. Lecture Notes in Computer Science, 2013, , 67-80.	1.0	10
310	Design, Optimization, and Benchmarking of Dense Linear Algebra Algorithms on AMD GPUs. , 2020, , .		10
311	PB-BLAS: a set of parallel block basic linear algebra subprograms. Concurrency and Computation: Practice and Experience, 1996, 8, 517-535.	0.6	9
312	NetBuild: transparent cross-platform access to computational software libraries. Concurrency Computation Practice and Experience, 2002, 14, 1445-1456.	1.4	9
313	Middleware for the use of storage in communication. Parallel Computing, 2002, 28, 1773-1787.	1.3	9
314	Conjugate-gradient eigenvalue solvers in computing electronic properties of nanostructure architectures. International Journal of Computational Science and Engineering, 2006, 2, 205.	0.4	9
315	Feedback-directed thread scheduling with memory considerations. , 2007, , .		9
316	Soft error resilient QR factorization for hybrid system with GPGPU. , 2011, , .		9
317	QCG-OMPI: MPI applications on grids. Future Generation Computer Systems, 2011, 27, 357-369.	4.9	9
318	A Hybridization Methodology for High-Performance Linear Algebra Software for GPUs. , 2012, , 473-484.		9
319	An Implementation of the Tile QR Factorization for a GPU and Multiple CPUs. Lecture Notes in Computer Science, 2012, , 248-257.	1.0	9
320	Reducing the Amount of Pivoting in Symmetric Indefinite Systems. Lecture Notes in Computer Science, 2012, , 133-142.	1.0	9
321	MIAMI: A framework for application performance diagnosis. , 2014, , .		9
322	Energy efficiency and performance frontiers for sparse computations on GPU supercomputers. , 2015, , .		9
323	Stability and Performance of Various Singular Value QR Implementations on Multicore CPU with a GPU. ACM Transactions on Mathematical Software, 2017, 43, 1-18.	1.6	9
324	Why is it Hard to Describe Properties of Algorithms?. Procedia Computer Science, 2016, 101, 4-7.	1.2	9

#	Article	IF	Citations
325	Performance Tuning and Optimization Techniques of Fixed and Variable Size Batched Cholesky Factorization on GPUs. Procedia Computer Science, 2016, 80, 119-130.	1.2	9
326	Search Space Generation and Pruning System for Autotuners. , 2016, , .		9
327	Batched one-sided factorizations of tiny matrices using GPUs: Challenges and countermeasures. Journal of Computational Science, 2018, 26, 226-236.	1.5	9
328	Accelerating the SVD two stage bidiagonal reduction and divide and conquer using GPUs. Parallel Computing, 2018, 74, 3-18.	1.3	9
329	Autotuning Numerical Dense Linear Algebra for Batched Computation With GPU Hardware Accelerators. Proceedings of the IEEE, 2018, 106, 2040-2055.	16.4	9
330	Beyond the CPU: Hardware Performance Counter Monitoring on Blue Gene/Q. Lecture Notes in Computer Science, 2013, , 213-225.	1.0	9
331	Using agent-based software for scientific computing in the NetSolve system. Parallel Computing, 1998, 24, 1777-1790.	1.3	8
332	Programming tools and environments. Communications of the ACM, 1998, 41, 64-73.	3.3	8
333	Logistical quality of service in NetSolve. Computer Communications, 1999, 22, 1034-1044.	3.1	8
334	Trends in High Performance Computing. Computer Journal, 2004, 47, 399-403.	1.5	8
335	GrADSolveâ€”a grid-based RPC system for parallel computing with application-level scheduling. Journal of Parallel and Distributed Computing, 2004, 64, 774-783.	2.7	8
336	Looking back at dense linear algebra software. Journal of Parallel and Distributed Computing, 2014, 74, 2548-2560.	2.7	8
337	Adaptive precision solvers for sparse linear systems. , 2015, , .		8
338	Performance optimization of Sparse Matrixâ€Vector Multiplication for multiâ€component PDEâ€based applications using GPUs. Concurrency Computation Practice and Experience, 2016, 28, 3447-3465.	1.4	8
339	Performance-Portable Autotuning of OpenCL Kernels for Convolutional Layers of Deep Neural Networks. , 2016, , .		8
340	Comparing the performance of rigid, moldable and grid-shaped applications on failure-prone HPC platforms. Parallel Computing, 2019, 85, 1-12.	1.3	8
341	Reducing the amount of outâ€ofâ€core data access for GPUâ€accelerated randomized SVD. Concurrency Computation Practice and Experience, 2020, 32, e5754.	1.4	8
342	Optimal Checkpointing Period: Time vs. Energy. Lecture Notes in Computer Science, 2014, , 203-214.	1.0	8

#	Article	IF	Citations
343	Scalability Analysis of the SPEC OpenMP Benchmarks on Large-Scale Shared Memory Multiprocessors. Lecture Notes in Computer Science, 2007, , 815-822.	1.0	8
344	An Iterative Solver Benchmark. Scientific Programming, 2001, 9, 223-231.	0.5	7
345	HARNESS fault tolerant MPI design, usage and performance issues. Future Generation Computer Systems, 2002, 18, 1127-1142.	4.9	7
346	GrADSolve â€“ RPC for High Performance Computing on the Grid. Lecture Notes in Computer Science, 2003, , 394-403.	1.0	7
347	Paravirtualization effect on single- and multi-threaded memory-intensive linear algebra software. Cluster Computing, 2009, 12, 101-122.	3.5	7
348	Scheduling Two-Sided Transformations Using Tile Algorithms on Multicore Architectures. Scientific Programming, 2010, 18, 35-50.	0.5	7
349	Revisiting the Double Checkpointing Algorithm. , 2013, , .		7
350	Scaling up matrix computations on shared-memory manycore systems with 1000 CPU cores. , 2014, , .		7
351	A scalable approach to solving dense linear algebra problems on hybrid CPUâ€GPU systems. Concurrency Computation Practice and Experience, 2015, 27, 3702-3723.	1.4	7
352	Optimization for performance and energy for batched matrix computations on GPUs. , 2015, , .		7
353	Efficiency of General Krylov Methods on GPUs -- An Experimental Study. , 2016, , .		7
354	LU, QR, and Cholesky factorizations: Programming model, performance analysis and optimization techniques for the Intel Knights Landing Xeon Phi. , 2016, , .		7
355	Variable-Size Batched LU for Small Matrices and Its Integration into Block-Jacobi Preconditioning. , 2017, , .		7
356	Towards numerical benchmark for half-precision floating point arithmetic. , 2017, , .		7
357	Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling. Journal of Advances in Modeling Earth Systems, 2018, 10, 1952-1969.	1.3	7
358	Analysis and Design Techniques towards High-Performance and Energy-Efficient Dense Linear Solvers on GPUs. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 2700-2712.	4.0	7
359	ParILUT - A Parallel Threshold ILU for GPUs. , 2019, , .		7
360	Self-Adapting Numerical Software and Automatic Tuning of Heuristics. Lecture Notes in Computer Science, 2003, , 759-767.	1.0	7

#	Article	IF	Citations
361	Performance Instrumentation and Compiler Optimizations for MPI/OpenMP Applications. Lecture Notes in Computer Science, 2008, , 267-278.	1.0	7
362	Fast and Small Short Vector SIMD Matrix Multiplication Kernels for the Synergistic Processing Element of the CELL Processor. Lecture Notes in Computer Science, 2008, , 935-944.	1.0	7
363	Randomized algorithms to update partial singular value decomposition on a hybrid CPU/GPU cluster. , 2015, , .		7
364	Changing technologies of HPC. Future Generation Computer Systems, 1997, 12, 461-474.	4.9	6
365	Tiling on systems with communication/computation overlap. Concurrency and Computation: Practice and Experience, 1999, 11, 139-153.	0.6	6
366	Automatic translation of Fortran to JVM bytecode. Concurrency Computation Practice and Experience, 2003, 15, 207-222.	1.4	6
367	The Component Structure of a Self-Adapting Numerical Software System. International Journal of Parallel Programming, 2005, 33, 137-143.	1.1	6
368	Processes Distribution of Homogeneous Parallel Linear Algebra Routines on Heterogeneous Clusters. , 2005, , .		6
369	High Performance Development for High End Computing With Python Language Wrapper (PLW). International Journal of High Performance Computing Applications, 2007, 21, 360-369.	2.4	6
370	The use of bulk states to accelerate the band edge state calculation of a semiconductor quantum dot. Journal of Computational Physics, 2007, 223, 774-782.	1.9	6
371	From Serial Loops to Parallel Execution on Distributed Systems. Lecture Notes in Computer Science, 2012, , 246-257.	1.0	6
372	Dynamically Balanced Synchronization-Avoiding LU Factorization with Multicore and GPUs. , 2014, , .		6
373	Tuning stationary iterative solvers for fault resilience. , 2015, , .		6
374	Task-Based Cholesky Decomposition on Knights Corner Using OpenMP. Lecture Notes in Computer Science, 2016, , 544-562.	1.0	6
375	Assessing the cost of redistribution followed by a computational kernel: Complexity and performance results. Parallel Computing, 2016, 52, 22-41.	1.3	6
376	Optimized Batched Linear Algebra for Modern Architectures. Lecture Notes in Computer Science, 2017, , 511-522.	1.0	6
377	Flexible batched sparse matrix-vector product on GPUs. , 2017, , .		6
378	Non-GPU-resident symmetric indefinite factorization. Concurrency Computation Practice and Experience, 2017, 29, e4012.	1.4	6

#	Article	IF	Citations
379	Out of memory SVD solver for big data. , 2017, , .		6
380	Symmetric Indefinite Linear Solver Using OpenMP Task on Multicore Architectures. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 1879-1892.	4.0	6
381	Towards Half-Precision Computation for Complex Matrices: A Case Study for Mixed Precision Solvers on GPUs. , 2019, , .		6
382	MAGMA templates for scalable linear algebra on emerging architectures. International Journal of High Performance Computing Applications, 2020, 34, 645-658.	2.4	6
383	Leveraging PaRSEC Runtime Support to Tackle Challenging 3D Data-Sparse Matrix Problems. , 2021, , .		6
384	Investigating the Benefit of FP16-Enabled Mixed-Precision Solvers for Symmetric Positive Definite Matrices Using GPUs. Lecture Notes in Computer Science, 2020, , 237-250.	1.0	6
385	Mixed-Precision Orthogonalization Scheme and Adaptive Step Size for Improving the Stability and Performance of CA-GMRES on GPUs. Lecture Notes in Computer Science, 2015, , 17-30.	1.0	6
386	Parallel Tiled QR Factorization for Multicore Architectures. , 2007, , 639-648.		6
387	A Holistic Approach for Performance Measurement and Analysis for Petascale Applications. Lecture Notes in Computer Science, 2009, , 686-695.	1.0	6
388	Multi-criteria Checkpointing Strategies: Response-Time versus Resource Utilization. Lecture Notes in Computer Science, 2013, , 420-431.	1.0	6
389	Stochastic Performance Prediction for Iterative Algorithms in Distributed Environments. Journal of Parallel and Distributed Computing, 1999, 58, 68-91.	2.7	5
390	Numerical Libraries and Tools for Scalable Parallel Cluster Computing. International Journal of High Performance Computing Applications, 2001, 15, 175-180.	2.4	5
391	On Scalability for MPI Runtime Systems. , 2011, , .		5
392	Implementing a Blocked Aasen's Algorithm with a Dynamic Scheduler on Multicore Architectures. , 2013, , .		5
393	Towards batched linear solvers on accelerated hardware platforms. ACM SIGPLAN Notices, 2015, 50, 261-262.	0.2	5
394	Mixed-precision block gram Schmidt orthogonalization. , 2015, , .		5
395	Mixing LU and QR factorization algorithms to design high-performance dense linear algebra solvers. Journal of Parallel and Distributed Computing, 2015, 85, 32-46.	2.7	5
396	On the Development of Variable Size Batched Computation for Heterogeneous Parallel Architectures. , 2016, , .		5

#	Article	IF	Citations
397	Data through the Computational Lens, Preface for ICCS 2016. Procedia Computer Science, 2016, 80, 1-7.	1.2	5
398	Performance of Hierarchical-matrix BiCGStab Solver on GPU Clusters. , 2018, , .		5
399	Evaluation of Programming Models to Address Load Imbalance on Distributed Multi-Core CPUs: A Case Study with Block Low-Rank Factorization. , 2019, , .		5
400	Accelerating NWChem Coupled Cluster Through Dataflow-Based Execution. Lecture Notes in Computer Science, 2016, , 366-376.	1.0	5
401	Domain Overlap for Iterative Sparse Triangular Solves on GPUs. Lecture Notes in Computational Science and Engineering, 2016, , 527-545.	0.1	5
402	Prospectus for the Next LAPACK and ScaLAPACK Libraries. , 2006, , 11-23.		5
403	Reducing the Time to Tune Parallel Dense Linear Algebra Routines with Partial Execution and Performance Modeling. Lecture Notes in Computer Science, 2012, , 730-739.	1.0	5
404	Linear Systems Solvers for Distributed-Memory Machines with GPU Accelerators. Lecture Notes in Computer Science, 2019, , 495-506.	1.0	5
405	Improving the Performance of the GMRES Method Using Mixed-Precision Techniques. Communications in Computer and Information Science, 2020, , 51-66.	0.4	5
406	Using long vector extensions for MPI reductions. Parallel Computing, 2022, 109, 102871.	1.3	5
407	Batch QR Factorization onÂGPUs: Design, Optimization, andÂTuning. Lecture Notes in Computer Science, 2022, , 60-74.	1.0	5
408	A Framework to Exploit Data Sparsity in Tile Low-Rank Cholesky Factorization. , 2022, , .		5
409	Comparison of the CRAY X-MP-4, Fujitsu VP-200, and Hitachi S-810/20. Simulation, 1986, 47, 93-107.	1.1	4
410	1988 Gordon Bell Prize. IEEE Software, 1989, 6, 78-85.	2.1	4
411	High performance linear algebra package for FORTRAN 90. Lecture Notes in Computer Science, 1998, , 579-583.	1.0	4
412	Self Adaptive Application Level Fault Tolerance for Parallel and Distributed Computing. , 2007, , .		4
413	Request Sequencing: Enabling Workflow for Efficient Problem Solving in GridSolve. , 2008, , .		4
414	SmartGridRPC: The new RPC model for high performance Grid computing. Concurrency Computation Practice and Experience, 2010, 22, 2467-2487.	1.4	4

#	Article	IF	Citations
415	Virtual Systolic Array for QR Decomposition. , 2013, , .		4
416	Performance and reliability trade-offs for the double checkpointing algorithm. International Journal of Networking and Computing, 2014, 4, 23-41.	0.3	4
417	POSTER: Utilizing dataflow-based execution for coupled cluster methods. , 2014, , .		4
418	Performance and Portability with OpenCL for Throughput-Oriented HPC Workloads across Accelerators, Coprocessors, and Multicore Processors. , 2014, , .		4
419	Flexible Linear Algebra Development and Scheduling with Cholesky Factorization. , 2015, , .		4
420	Variable-Size Batched Gauss-Huard for Block-Jacobi Preconditioning. Procedia Computer Science, 2017, 108, 1783-1792.	1.2	4
421	Optimizing the SVD Bidiagonalization Process for a Batch of Small Matrices. Procedia Computer Science, 2017, 108, 1008-1018.	1.2	4
422	Optimization and performance evaluation of the IDR iterative Krylov solver on GPUs. International Journal of High Performance Computing Applications, 2018, 32, 220-230.	2.4	4
423	Massively Parallel Automated Software Tuning. , 2019, , .		4
424	Checkpointing Strategies for Shared High-Performance Computing Platforms. International Journal of Networking and Computing, 2019, 9, 28-52.	0.3	4
425	Computational Science in the Interconnected World: Selected papers from 2019 International Conference on Computational Science. Journal of Computational Science, 2020, 47, 101222.	1.5	4
426	Hash Functions for Datatype Signatures in MPI. Lecture Notes in Computer Science, 2005, , 76-83.	1.0	4
427	Dense Symmetric Indefinite Factorization on GPU Accelerated Architectures. Lecture Notes in Computer Science, 2016, , 86-95.	1.0	4
428	High-Order Finite Element Method using Standard and Device-Level Batch GEMM on GPUs. , 2020, , .		4
429	Performance of random sampling for computing low-rank approximations of a dense matrix on GPUs. , 2015, , .		4
430	Dense Linear Algebra for Hybrid GPU-Based Systems. Chapman & Hall/CRC Computational Science, 2010, , 37-55.	0.5	4
431	Implementing a Systolic Algorithm for QR Factorization on Multicore Clusters with PaRSEC. Lecture Notes in Computer Science, 2014, , 657-667.	1.0	4
432	MagmaDNN: Towards High-Performance Data Analytics and Machine Learning for Data-Driven Scientific Computing. Lecture Notes in Computer Science, 2019, , 490-503.	1.0	4

#	Article	IF	Citations
433	Templates for linear algebra problems. Lecture Notes in Computer Science, 1995, , 115-140.	1.0	3
434	NanoPSE: Nanoscience Problem Solving Environment for atomistic electronic structure of semiconductor nanostructures. Journal of Physics: Conference Series, 2005, 16, 277-282.	0.3	3
435	Comparison of Nonlinear Conjugate-Gradient Methods for Computing the Electronic Properties of Nanostructure Architectures. Lecture Notes in Computer Science, 2005, , 317-325.	1.0	3
436	Toward a scalable multi-GPU eigensolver via compute-intensive kernels and efficient communication. , 2013, , .		3
437	Access-averse framework for computing low-rank matrix approximations. , 2014, , .		3
438	Designing LU-QR Hybrid Solvers for Performance and Stability. , 2014, , .		3
439	Parallel Simulation of Superscalar Scheduling. , 2014, , .		3
440	BlackjackBench: Portable Hardware Characterization with Automated Results' Analysis. Computer Journal, 2014, 57, 1002-1016.	1.5	3
441	GPU-accelerated co-design of induced dimension reduction. , 2015, , .		3
442	Solving dense symmetric indefinite systems using GPUs. Concurrency Computation Practice and Experience, 2017, 29, e4055.	1.4	3
443	Structure-Aware Linear Solver for Realtime Convex Optimization for Embedded Systems. IEEE Embedded Systems Letters, 2017, 9, 61-64.	1.3	3
444	Autotuning batch Cholesky factorization in CUDA with interleaved layout of matrices. , 2017, , .		3
445	Accelerating NWChem Coupled Cluster through dataflow-based execution. International Journal of High Performance Computing Applications, 2018, 32, 540-551.	2.4	3
446	Post-exascale supercomputing: research opportunities abound. Frontiers of Information Technology and Electronic Engineering, 2018, 19, 1203-1208.	1.5	3
447	Increasing Accuracy of Iterative Refinement in Limited Floating-Point Arithmetic on Half-Precision Accelerators. , 2019, , .		3
448	Progressive Optimization of Batched LU Factorization on GPUs. , 2019, , .		3
449	Fine-grained bit-flip protection for relaxation methods. Journal of Computational Science, 2019, 36, 100583.	1.5	3
450	Accelerating Computation of Eigenvectors in the Dense Nonsymmetric Eigenvalue Problem. Lecture Notes in Computer Science, 2015, , 182-191.	1.0	3

#	Article	IF	Citations
451	GPU-Accelerated Asynchronous Error Correction for Mixed Precision Iterative Refinement. Lecture Notes in Computer Science, 2012, , 908-919.	1.0	3
452	Using Advanced Vector Extensions AVX-512 for MPI Reductions. , 2020, , .		3
453	Truss Structual Optimization using NetSolve System. The Proceedings of OPTIS, 2002, 2002.5, 141-146.	0.0	3
454	With Extreme Scale Computing the Rules Have Changed. Lecture Notes in Computer Science, 2016, , 3-6.	1.0	3
455	Integrating Deep Learning in Domain Sciences at Exascale. Communications in Computer and Information Science, 2020, , 35-50.	0.4	3
456	Flexible Data Redistribution in a Task-Based Runtime System. , 2020, , .		3
457	An Introduction to High Performance Computing and Its Intersection with Advances in Modeling Rare Earth Elements and Actinides. ACS Symposium Series, 0, , 3-53.	0.5	3
458	Special section: Applications of distributed and grid computing. Future Generation Computer Systems, 2008, 24, 582-584.	4.9	2
459	Special section: Grid computing and the message passing interface. Future Generation Computer Systems, 2008, 24, 119-120.	4.9	2
460	Trace-based performance analysis for the petascale simulation code FLASH. International Journal of High Performance Computing Applications, 2011, 25, 428-439.	2.4	2
461	BlackjackBench. Performance Evaluation Review, 2012, 40, 74-79.	0.4	2
462	Tridiagonalization of a Symmetric Dense Matrix on a GPU Cluster. , 2013, , .		2
463	New Algorithm for Computing Eigenvectors of the Symmetric Eigenvalue Problem. , 2014, , .		2
464	Hybrid Multi-elimination ILU Preconditioners on GPUs. , 2014, , .		2
465	Deflation Strategies to Improve the Convergence of Communication-Avoiding GMRES. , 2014, , .		2
466	Weighted dynamic scheduling with many parallelism grains for offloading of numerical workloads to multiple varied accelerators. , 2015, , .		2
467	Performance analysis and acceleration of explicit integration for large kinetic networks using batched GPU computations. , 2016, , .		2
468	Hessenberg Reduction with Transient Error Resilience on GPU-Based Hybrid Architectures. , 2016, , .		2

#	Article	IF	Citations
469	The Art of Computational Science, Bridging Gaps â€“ Forming Alloys. Preface for ICCS 2017. Procedia Computer Science, 2017, 108, 1-6.	1.2	2
470	Sampling algorithms to update truncated SVD. , 2017, , .		2
471	Least squares solvers for distributed-memory machines with GPU accelerators. , 2019, , .		2
472	Dense Linear Algebra on Accelerated Multicore Hardware. , 2012, , 123-146.		2
473	Weighted Block-Asynchronous Iteration on GPU-Accelerated Systems. Lecture Notes in Computer Science, 2013, , 145-154.	1.0	2
474	Providing access to high performance computing technologies. Lecture Notes in Computer Science, 1997, , 24-34.	1.0	2
475	Evaluating Data Redistribution in PaRSEC. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 1856-1872.	4.0	2
476	Replacing Pivoting in Distributed Gaussian Elimination with Randomized Techniques. , 2020, , .		2
477	Location-independent naming for virtual distributed software repositories. Software Engineering Notes: an Informal Newsletter of the Special Interest Committee on Software Engineering / ACM, 1995, 20, 179-185.	0.5	1
478	Selected numerical algorithms. Future Generation Computer Systems, 2004, 20, 349-351.	4.9	1
479	13. Parallel Linear Algebra Software. , 2006, , 233-247.		1
480	Optimal Routing in Binomial Graph Networks. , 2007, , .		1
481	REVISITING MATRIX PRODUCT ON MASTER-WORKER PLATFORMS. International Journal of Foundations of Computer Science, 2008, 19, 1317-1336.	0.8	1
482	Mixed-Tool Performance Analysis on Hybrid Multicore Architectures. , 2010, , .		1
483	A Block-Asynchronous Relaxation Method for Graphics Processing Units. , 2012, , .		1
484	Performance Analysis and Optimisation of Two-sided Factorization Algorithms for Heterogeneous Platform. Procedia Computer Science, 2015, 51, 180-190.	1.2	1
485	Bidiagonalization and R-Bidiagonalization: Parallel Tiled Algorithms, Critical Paths and Distributed-Memory Implementation. , 2017, , .		1
486	Scaling point set registration in 3D across thread counts on multicore and hardware accelerator platforms through autotuning for large scale analysis of scientific point clouds. , 2017, , .		1

#	Article	IF	Citations
487	Evaluation of dataflow programming models for electronic structure theory. Concurrency Computation Practice and Experience, 2018, 30, e4490.	1.4	1
488	Variable-Size Batched Condition Number Calculation on GPUs. , 2018, , .		1
489	Optimizing GPU Kernels for Irregular Batch Workloads: A Case Study for Cholesky Factorization. , 2018, , .		1
490	Translational process: Mathematical software perspective. Journal of Computational Science, 2021, 52, 101216.	1.5	1
491	Eigenvalue Computation with NetSolve Global Computing System. Lecture Notes in Computer Science, 2006, , 446-453.	1.0	1
492	Hands-On Research and Training in High Performance Data Sciences, Data Analytics, and Machine Learning for Emerging Environments. Lecture Notes in Computer Science, 2019, , 643-655.	1.0	1
493	Accelerating the Conjugate Gradient Algorithm with GPUs in CFD Simulations. Lecture Notes in Computer Science, 2017, , 35-43.	1.0	1
494	Self-Adapting Software for Numerical Linear Algebra Library Routines on Clusters. Lecture Notes in Computer Science, 2003, , 665-672.	1.0	1
495	Transparent Cross-Platform Access to Software Services Using GridSolve and GridRPC. , 2010, , 253-274.		1
496	Implementing Matrix Multiplication on the Cell B. E.. Chapman & Hall/CRC Computational Science, 2010, , 3-20.	0.5	1
497	An update notice on the level 3 BLAS. ACM SIGNUM Newsletter, 1989, 24, 9-10.	0.2	1
498	LAPACK is now available. ACM SIGNUM Newsletter, 1992, 27, 3-4.	0.2	1
499	Heterogenous Acceleration for Linear Algebra in Multi-coprocessor Environments. Lecture Notes in Computer Science, 2015, , 31-42.	1.0	1
500	Self-healing in Binomial Graph Networks. , 2007, , 1032-1041.		1
501	Preface To the Special Issue. International Journal of High Performance Computing Applications, 1997, 11, 83-83.	1.6	0
502	Deploying fault-tolerance and task migration with NetSolve. Lecture Notes in Computer Science, 1998, , 418-432.	1.0	0
503	High Performance Computing Trends and Self Adapting Numerical Software. Lecture Notes in Computer Science, 2003, , 1-9.	1.0	0
504	IMPROVED RUNTIME AND TRANSFER TIME PREDICTION MECHANISMS IN A NETWORK ENABLED SERVERS MIDDLEWARE. Parallel Processing Letters, 2007, 17, 47-59.	0.4	0

#	Article	IF	Citations
505	Editorial introduction to the special issue on computational linear algebra and sparse matrix computations. Applicable Algebra in Engineering, Communications and Computing, 2007, 18, 205-207.	0.3	0
506	Special section: Cluster and computational grids for scientific computing. Future Generation Computer Systems, 2008, 24, 30.	4.9	0
507	Selected papers of the Workshop on Clusters, Clouds and Grids for Scientific Computing (CCGSC). International Journal of High Performance Computing Applications, 2011, 25, 259-260.	2.4	0
508	GUEST EDITORS NOTE. Parallel Processing Letters, 2011, 21, 109-109.	0.4	0
509	Poster: A Novel Hybrid CPU-GPU Generalized Eigensolver for Electronic Structure Calculations Based on Fine Grained Memory Aware Tasks. , 2012, , .		0
510	Empowering Science through Computing, Preface for ICCS 2012. Procedia Computer Science, 2012, 9, 1-6.	1.2	0
511	Enabling workflows in GridSolve: request sequencing and service trading. Journal of Supercomputing, 2013, 64, 1133-1152.	2.4	0
512	Introduction for August Special Issue CCDSC. International Journal of High Performance Computing Applications, 2013, 27, 231-231.	2.4	0
513	GUEST EDITORS' NOTE: SPECIAL ISSUE ON CLUSTERS, CLOUDS, AND DATA FOR SCIENTIFIC COMPUTING. Parallel Processing Letters, 2013, 23, 1302001.	0.4	0
514	Design and Implementation of a Large Scale Tree-Based QR Decomposition Using a 3D Virtual Systolic Array and a Lightweight Runtime. Parallel Processing Letters, 2014, 24, 1442004.	0.4	0
515	Guest Editorsâ€™ Note: Special Issue on Clusters, Clouds and Data for Scientific Computing. Parallel Processing Letters, 2015, 25, 1502002.	0.4	0
516	Computing Low-Rank Approximation of a Dense Matrix on Multicore CPUs with a GPU and Its Application to Solving a Hierarchically Semiseparable Linear System of Equations. Scientific Programming, 2015, 2015, 1-17.	0.5	0
517	Techniques for Solving Large-Scale Graph Problems on Heterogeneous Platforms. Communications in Computer and Information Science, 2016, , 318-332.	0.4	0
518	Guest editorsâ€™ note. International Journal of High Performance Computing Applications, 2018, 32, 3-3.	2.4	0
519	Task based Cholesky decomposition on Xeon Phi architectures using OpenMP. International Journal of Computational Science and Engineering, 2018, 17, 310.	0.4	0
520	Guest editorsâ€™ note: Special issue on clusters, clouds, and data for scientific computing. International Journal of High Performance Computing Applications, 2019, 33, 1067-1068.	2.4	0
521	Exploiting Block Structures of KKT Matrices for Efficient Solution of Convex Optimization Problems. IEEE Access, 2021, 9, 116604-116611.	2.6	0
522	The Semantic Conference Organizer. , 2003, , .		0

#	Article	IF	Citations
523	Prospectus for a Dense Linear Algebra Software Library. Chapman & Hall/CRC Computer and Information Science Series, 2007, , 29-1-29-21.	0.4	0
524	Disaster Survival Guide in Petascale Computing. Chapman & Hall/CRC Computational Science, 2007, , 263-288.	0.5	0
525	A Scalable Non-blocking Multicast Scheme for Distributed DAG Scheduling. Lecture Notes in Computer Science, 2009, , 195-204.	1.0	0
526	Implementing Matrix Factorizations on the Cell B. E.. Chapman & Hall/CRC Computational Science, 2010, , 21-35.	0.5	0
527	Programming the LU Factorization for a Multicore System with Accelerators. Lecture Notes in Computer Science, 2013, , 28-35.	1.0	0
528	BLAS. Discrete Mathematics and Its Applications, 2013, , 1697-1704.	0.1	0
529	Distributed information management in the National HPCC Software Exchange. , 1995, , .		0
530	Block-cyclic array redistribution on networks of workstations. Lecture Notes in Computer Science, 1997, , 343-350.	1.0	0
531	High performance linear algebra package LAPACK90. Lecture Notes in Computer Science, 1998, , 387-391.	1.0	0
532	Self-adaptive Multiprecision Preconditioners on Multicore and Manycore Architectures. Lecture Notes in Computer Science, 2015, , 115-123.	1.0	0
533	Scalable Data Generation for Evaluating Mixed-Precision Solvers. , 2020, , .		0
534	Comparing Distributed Termination Detection Algorithms for Modern HPC Platforms. International Journal of Networking and Computing, 2022, 12, 26-46.	0.3	0
535	Optimal Routing in Binomial Graph Networks. , 2007, , .		0