

Jack Dongarra

List of Publications by Citations

Source: <https://exaly.com/author-pdf/8735190/jack-dongarra-publications-by-citations.pdf>

Version: 2024-04-28

This document has been generated based on the publications and citations recorded by exaly.com. For the latest version of this publication list, visit the link given above.

The third column is the impact factor (IF) of the journal, and the fourth column is the number of citations of the article.

522
papers

15,969
citations

53
h-index

112
g-index

576
ext. papers

19,023
ext. citations

2
avg, IF

6.59
L-index

#	Paper	IF	Citations
522	Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods 1994 ,		1383
521	PVM 1994 ,		1120
520	A set of level 3 basic linear algebra subprograms. <i>ACM Transactions on Mathematical Software</i> , 1990 , 16, 1-17	2.3	1039
519	Automated empirical optimizations of software and the ATLAS project. <i>Parallel Computing</i> , 2001 , 27, 3-35	1	704
518	An extended set of FORTRAN basic linear algebra subprograms. <i>ACM Transactions on Mathematical Software</i> , 1988 , 14, 1-17	2.3	514
517	The International Exascale Software Project roadmap. <i>International Journal of High Performance Computing Applications</i> , 2011 , 25, 3-60	1.8	392
516	An updated set of basic linear algebra subprograms (BLAS). <i>ACM Transactions on Mathematical Software</i> , 2002 , 28, 135-151	2.3	356
515	A class of parallel tiled linear algebra algorithms for multicore architectures. <i>Parallel Computing</i> , 2009 , 35, 38-53	1	248
514	Exascale computing and big data. <i>Communications of the ACM</i> , 2015 , 58, 56-68	2.5	236
513	Towards dense linear algebra for hybrid GPU accelerated manycore systems. <i>Parallel Computing</i> , 2010 , 36, 232-240	1	219
512	Netsolve: a Network-Enabled Server for Solving Computational Science Problems. <i>International Journal of High Performance Computing Applications</i> , 1997 , 11, 212-223		207
511	Chebyshev tau-QZ algorithm methods for calculating spectra of hydrodynamic stability problems. <i>Applied Numerical Mathematics</i> , 1996 , 22, 399-434	2.5	189
510	Distribution of mathematical software via electronic mail. <i>Communications of the ACM</i> , 1987 , 30, 403-407	2.5	187
509	Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects. <i>Journal of Physics: Conference Series</i> , 2009 , 180, 012037	0.3	185
508	The GrADS Project: Software Support for High-Level Grid Application Development. <i>International Journal of High Performance Computing Applications</i> , 2001 , 15, 327-344	1.8	181
507	Implementing Linear Algebra Algorithms for Dense Matrices on a Vector Pipeline Machine. <i>SIAM Review</i> , 1984 , 26, 91-112	7.4	181
506	The PVM concurrent computing system: Evolution, experiences, and trends. <i>Parallel Computing</i> , 1994 , 20, 531-545	1	158

505	From CUDA to OpenCL: Towards a performance-portable solution for multi-platform GPU programming. <i>Parallel Computing</i> , 2012 , 38, 391-407	1	145
504	DAGuE: A generic distributed DAG engine for High Performance Computing. <i>Parallel Computing</i> , 2012 , 38, 37-51	1	139
503	Algorithm-based fault tolerance applied to high performance computing. <i>Journal of Parallel and Distributed Computing</i> , 2009 , 69, 410-416	4.4	133
502	Algorithm 679: A set of level 3 basic linear algebra subprograms: model implementation and test programs. <i>ACM Transactions on Mathematical Software</i> , 1990 , 16, 18-28	2.3	133
501	Dense linear algebra solvers for multicore with GPU accelerators 2010 ,		109
500	PaRSEC: Exploiting Heterogeneity to Enhance Scalability. <i>Computing in Science and Engineering</i> , 2013 , 15, 36-45	1.5	107
499	An Improved Magma Gemm For Fermi Graphics Processing Units. <i>International Journal of High Performance Computing Applications</i> , 2010 , 24, 511-515	1.8	97
498	ScaLAPACK: a portable linear algebra library for distributed memory computers [Design issues and performance. <i>Computer Physics Communications</i> , 1996 , 97, 1-15	4.2	95
497	Accelerating scientific computations with mixed precision algorithms. <i>Computer Physics Communications</i> , 2009 , 180, 2526-2533	4.2	92
496	Post-failure recovery of MPI communication capability: Design and rationale. <i>International Journal of High Performance Computing Applications</i> , 2013 , 27, 244-254	1.8	91
495	Condition Numbers of Gaussian Random Matrices. <i>SIAM Journal on Matrix Analysis and Applications</i> , 2005 , 27, 603-620	1.5	91
494	Software Libraries for Linear Algebra Computations on High Performance Computers. <i>SIAM Review</i> , 1995 , 37, 151-180	7.4	90
493	New Grid Scheduling and Rescheduling Methods in the GrADS Project. <i>International Journal of Parallel Programming</i> , 2005 , 33, 209-229	1.5	88
492	Collecting Performance Data with PAPI-C 2010 , 157-173		88
491	Performance analysis of MPI collective operations. <i>Cluster Computing</i> , 2007 , 10, 127-143	2.1	87
490	Block reduction of matrices to condensed forms for eigenvalue computations. <i>Journal of Computational and Applied Mathematics</i> , 1989 , 27, 215-227	2.4	81
489	Mixed Precision Iterative Refinement Techniques for the Solution of Dense Linear Systems. <i>International Journal of High Performance Computing Applications</i> , 2007 , 21, 457-466	1.8	80
488	Keeneland: Bringing Heterogeneous GPU Computing to the Computational Science Community. <i>Computing in Science and Engineering</i> , 2011 , 13, 90-95	1.5	79

487	Algorithm-Based Fault Tolerance for Fail-Stop Failures. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2008 , 19, 1628-1641	3.7	72
486	Parallel tiled QR factorization for multicore architectures. <i>Concurrency Computation Practice and Experience</i> , 2008 , 20, 1573-1590	1.4	69
485	Integrated Pvm Framework Supports Heterogeneous Network Computing. <i>Computers in Physics</i> , 1993 , 7, 166		68
484	Pumma: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers. <i>Concurrency and Computation: Practice and Experience</i> , 1994 , 6, 543-570		63
483	Big data and extreme-scale computing: Pathways to Convergence-Toward a shaping strategy for a future software and data ecosystem for scientific inquiry. <i>International Journal of High Performance Computing Applications</i> , 2018 , 32, 435-479	1.8	62
482	Flexible Development of Dense Linear Algebra Algorithms on Massively Parallel Architectures with DPLASMA 2011 ,		62
481	Scheduling workflow applications on processors with different capabilities. <i>Future Generation Computer Systems</i> , 2006 , 22, 665-675	7.5	62
480	Improving the Accuracy of Computed Eigenvalues and Eigenvectors. <i>SIAM Journal on Numerical Analysis</i> , 1983 , 20, 23-45	2.4	62
479	A Note on Auto-tuning GEMM for GPUs. <i>Lecture Notes in Computer Science</i> , 2009 , 884-892	0.9	62
478	Overview of GridRPC: A Remote Procedure Call API for Grid Computing. <i>Lecture Notes in Computer Science</i> , 2002 , 274-278	0.9	61
477	High-performance bidiagonal reduction using tile algorithms on homogeneous multicore architectures. <i>ACM Transactions on Mathematical Software</i> , 2013 , 39, 1-22	2.3	59
476	. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2012 , 23, 2045-2057	3.7	59
475	On some parallel banded system solvers. <i>Parallel Computing</i> , 1984 , 1, 223-235	1	59
474	Unrolling loops in fortran. <i>Software - Practice and Experience</i> , 1979 , 9, 219-226	2.5	59
473	QR Factorization on a Multicore Node Enhanced with Multiple GPU Accelerators 2011 ,		54
472	Dynamic task scheduling for linear algebra algorithms on distributed-memory multicore systems 2009 ,		54
471	Exploiting the Performance of 32 bit Floating Point Arithmetic in Obtaining 64 bit Accuracy (Revisiting Iterative Refinement for Linear Systems) 2006 ,		54
470	Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers 2018 ,		54

469	. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2008 , 19, 1175-1186	3.7	53
468	A Parallel Divide and Conquer Algorithm for the Symmetric Eigenvalue Problem on Distributed Memory Architectures. <i>SIAM Journal of Scientific Computing</i> , 1999 , 20, 2223-2236	2.6	53
467	Applying NetSolve's network-enabled server. <i>IEEE Computational Science and Engineering</i> , 1998 , 5, 57-67		52
466	Enabling and scaling matrix computations on heterogeneous multi-core and multi-GPU systems 2012 ,		51
465	High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems. <i>International Journal of High Performance Computing Applications</i> , 2016 , 30, 3-10	1.8	49
464	The International Exascale Software Project: a Call To Cooperative Action By the Global High-Performance Community. <i>International Journal of High Performance Computing Applications</i> , 2009 , 23, 309-322	1.8	49
463	Accelerating the reduction to upper Hessenberg, tridiagonal, and bidiagonal forms through hybrid GPU-based computing. <i>Parallel Computing</i> , 2010 , 36, 645-654	1	49
462	Preface: Basic Linear Algebra Subprograms Technical (Blast) Forum Standard. <i>International Journal of High Performance Computing Applications</i> , 2002 , 16, 1-1	1.8	49
461	Algorithm-based fault tolerance for dense matrix factorizations 2012 ,		48
460	Using Mixed Precision for Sparse Matrix Computations to Enhance the Performance while Achieving 64-bit Accuracy. <i>ACM Transactions on Mathematical Software</i> , 2008 , 34, 1-22	2.3	48
459	Numerical Libraries and the Grid. <i>International Journal of High Performance Computing Applications</i> , 2001 , 15, 359-374	1.8	46
458	Scalability Issues Affecting the Design of a Dense Linear Algebra Library. <i>Journal of Parallel and Distributed Computing</i> , 1994 , 22, 523-537	4.4	46
457	. <i>Computer</i> , 1993 , 26, 88-95	1.6	46
456	Performance, Design, and Autotuning of Batched GEMM for GPUs. <i>Lecture Notes in Computer Science</i> , 2016 , 21-38	0.9	45
455	Scheduling dense linear algebra operations on multicore processors. <i>Concurrency Computation Practice and Experience</i> , 2010 , 22, 15-44	1.4	44
454	NetSolve 1996 ,		44
453	HARNESS: a next generation distributed virtual machine. <i>Future Generation Computer Systems</i> , 1999 , 15, 571-582	7.5	44
452	Implementation of some concurrent algorithms for matrix factorization. <i>Parallel Computing</i> , 1986 , 3, 25-34	1	43

451	Optimizing matrix multiplication for a short-vector SIMD architecture [CELL processor. <i>Parallel Computing</i> , 2009 , 35, 138-150	1	42
450	A tool to aid in the design, implementation, and understanding of matrix algorithms for parallel processors. <i>Journal of Parallel and Distributed Computing</i> , 1990 , 9, 185-202	4-4	42
449	Standards for graph algorithm primitives 2013 ,		41
448	DAGuE: A Generic Distributed DAG Engine for High Performance Computing 2011 ,		41
447	Scheduling block-cyclic array redistribution. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 1998 , 9, 192-205	3-7	41
446	. <i>IEEE Parallel and Distributed Technology</i> , 1995 , 3, 75-83		41
445	Toward a new metric for ranking high performance computing systems. 2013 ,		40
444	Fault tolerant high performance computing by a coding approach 2005 ,		39
443	HARNESS and fault tolerant MPI. <i>Parallel Computing</i> , 2001 , 27, 1479-1495	1	38
442	. <i>Proceedings of the IEEE</i> , 2018 , 106, 2068-2083	14-3	37
441	Reduction to condensed form for the Eigenvalue problem on distributed memory architectures. <i>Parallel Computing</i> , 1992 , 18, 973-982	1	37
440	The Impact of Multicore on Math Software 2006 , 1-10		37
439	Towards Efficient MapReduce Using MPI. <i>Lecture Notes in Computer Science</i> , 2009 , 240-249	0-9	37
438	Fault-Tolerant Matrix Operations for Networks of Workstations Using Diskless Checkpointing. <i>Journal of Parallel and Distributed Computing</i> , 1997 , 43, 125-138	4-4	35
437	Self-adapting software for numerical linear algebra and LAPACK for clusters. <i>Parallel Computing</i> , 2003 , 29, 1723-1743	1	35
436	Telescoping Languages: A Strategy for Automatic Generation of Scientific Problem-Solving Systems from Annotated Libraries. <i>Journal of Parallel and Distributed Computing</i> , 2001 , 61, 1803-1826	4-4	35
435	. <i>IEEE Spectrum</i> , 1987 , 24, 38-43	1-7	35
434	Implementation of mixed precision in solving systems of linear equations on the Cell processor. <i>Concurrency Computation Practice and Experience</i> , 2007 , 19, 1371-1385	1-4	34

433	A proposal for a set of parallel basic linear algebra subprograms. <i>Lecture Notes in Computer Science</i> , 1996 , 107-114	0.9	34
432	Solving banded systems on a parallel processor. <i>Parallel Computing</i> , 1987 , 5, 219-246	1	34
431	Batched matrix computations on hardware accelerators based on GPUs. <i>International Journal of High Performance Computing Applications</i> , 2015 , 29, 193-208	1.8	33
430	High-performance computing systems: Status and outlook*. <i>Acta Numerica</i> , 2012 , 21, 379-474	15.1	33
429	The Design and Performance of Batched BLAS on Modern High-Performance Computing Systems. <i>Procedia Computer Science</i> , 2017 , 108, 495-504	1.6	32
428	Unified model for assessing checkpointing protocols at extreme-scale. <i>Concurrency Computation Practice and Experience</i> , 2014 , 26, 2772-2791	1.4	32
427	Multiprocessing linear algebra algorithms on the CRAY X-MP-2: Experiences with small granularity. <i>Journal of Parallel and Distributed Computing</i> , 1984 , 1, 22-31	4.4	32
426	Accelerating Numerical Dense Linear Algebra Calculations with GPUs 2014 , 3-28		32
425	Iterative Sparse Triangular Solves for Preconditioning. <i>Lecture Notes in Computer Science</i> , 2015 , 650-661	0.9	31
424	Recovery Patterns for Iterative Methods in a Parallel Unstable Environment. <i>SIAM Journal of Scientific Computing</i> , 2008 , 30, 102-116	2.6	30
423	Mixed-Precision Cholesky QR Factorization and Its Case Studies on Multicore CPU with Multiple GPUs. <i>SIAM Journal of Scientific Computing</i> , 2015 , 37, C307-C330	2.6	29
422	Squeezing the most out of eigenvalue solvers on high-performance computers. <i>Linear Algebra and Its Applications</i> , 1986 , 77, 113-136	0.9	28
421	Improving the Performance of CA-GMRES on Multicores with Multiple GPUs 2014 ,		27
420	Energy Footprint of Advanced Dense Numerical Linear Algebra Using Tile Algorithms on Multicore Architectures 2012 ,		27
419	Algorithm-based fault tolerance for dense matrix factorizations. <i>ACM SIGPLAN Notices</i> , 2012 , 47, 225-234	0.2	27
418	The Singular Value Decomposition: Anatomy of Optimizing an Algorithm for Extreme Scale. <i>SIAM Review</i> , 2018 , 60, 808-865	7.4	27
417	A Step towards Energy Efficient Computing: Redesigning a Hydrodynamic Application on CPU-GPU 2014 ,		26
416	. <i>IEEE Transactions on Computers</i> , 2009 , 58, 1512-1524	2.5	26

415	Redesigning the message logging model for high performance. <i>Concurrency Computation Practice and Experience</i> , 2010 , 22, 2196-2211	1.4	26
414	A comparison of search heuristics for empirical code optimization 2008 ,		26
413	The design and implementation of the parallel out-of-core ScaLAPACK LU, QR, and Cholesky factorization routines. <i>Concurrency and Computation: Practice and Experience</i> , 2000 , 12, 1481-1493		26
412	A portable environment for developing parallel FORTRAN programs. <i>Parallel Computing</i> , 1987 , 5, 175-186		26
411	LU factorization for accelerator-based systems 2011 ,		25
410	A scalable framework for heterogeneous GPU-based clusters 2012 ,		25
409	The Spectral Decomposition of Nonsymmetric Matrices on Distributed Memory Parallel Computers. <i>SIAM Journal of Scientific Computing</i> , 1997 , 18, 1446-1461	2.6	25
408	. <i>IEEE Annals of the History of Computing</i> , 2008 , 30, 30-41	0.2	25
407	A comparative study of automatic vectorizing compilers. <i>Parallel Computing</i> , 1991 , 17, 1223-1244	1	25
406	Implementation and Tuning of Batched Cholesky Factorization and Solve for NVIDIA GPUs. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2016 , 27, 2036-2048	3.7	24
405	High-performance high-resolution semi-Lagrangian tracer transport on a sphere. <i>Journal of Computational Physics</i> , 2011 , 230, 6778-6799	4.1	24
404	Innovations of the NetSolve Grid Computing System. <i>Concurrency Computation Practice and Experience</i> , 2002 , 14, 1457-1479	1.4	24
403	A parallel algorithm for the reduction of a nonsymmetric matrix to block upper-Hessenberg form. <i>Parallel Computing</i> , 1995 , 21, 1189-1211	1	24
402	Incomplete Sparse Approximate Inverses for Parallel Preconditioning. <i>Parallel Computing</i> , 2018 , 71, 1-221		23
401	LU Factorization of Small Matrices: Accelerating Batched DGETRF on the GPU 2014 ,		23
400	QR factorization of tall and skinny matrices in a grid computing environment 2010 ,		23
399	Numerical Considerations in Computing Invariant Subspaces. <i>SIAM Journal on Matrix Analysis and Applications</i> , 1992 , 13, 145-161	1.5	23
398	A collection of parallel linear equations routines for the Denelcor HEP. <i>Parallel Computing</i> , 1984 , 1, 133-142		23

397	Adaptive precision in block-Jacobi preconditioning for iterative sparse linear system solvers. <i>Concurrency Computation Practice and Experience</i> , 2019 , 31, e4460	1.4	23
396	Investigating half precision arithmetic to accelerate dense linear system solvers 2017 ,		22
395	Optimizing symmetric dense matrix-vector multiplication on GPUs 2011 ,		22
394	State-of-the-art eigensolvers for electronic structure calculations of large scale nano-systems. <i>Journal of Computational Physics</i> , 2008 , 227, 7113-7124	4.1	22
393	The marketplace of high-performance computing. <i>Parallel Computing</i> , 1999 , 25, 1517-1544	1	22
392	. <i>Computing in Science and Engineering</i> , 2019 , 21, 4-5	1.5	21
391	Dynamic task discovery in PaRSEC 2017 ,		21
390	Parallel reduction to condensed forms for symmetric eigenvalue problems using aggregated fine-grained and memory-aware kernels 2011 ,		21
389	Two-Stage Tridiagonal Reduction for Dense Symmetric Matrices Using Tile Algorithms on Multicore Architectures 2011 ,		21
388	Recent Developments in Gridsolve. <i>International Journal of High Performance Computing Applications</i> , 2006 , 20, 131-141	1.8	21
387	Recent trends in the marketplace of high performance computing. <i>Parallel Computing</i> , 2005 , 31, 261-273		21
386	Algorithmic redistribution methods for block-cyclic decompositions. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 1999 , 10, 1201-1216	3.7	21
385	Preconditioned Krylov solvers on GPUs. <i>Parallel Computing</i> , 2017 , 68, 32-44	1	20
384	Hierarchical DAG Scheduling for Hybrid Distributed Systems 2015 ,		20
383	Tridiagonalization of a dense symmetric matrix on multiple GPUs and its application to symmetric eigenvalue problems. <i>Concurrency Computation Practice and Experience</i> , 2014 , 26, 2652-2666	1.4	20
382	An evaluation of User-Level Failure Mitigation support in MPI. <i>Computing (Vienna/New York)</i> , 2013 , 95, 1171-1184	2.2	20
381	Accelerating collaborative filtering using concepts from high performance computing 2015 ,		20
380	ALGORITHMIC ISSUES ON HETEROGENEOUS COMPUTING PLATFORMS. <i>Parallel Processing Letters</i> , 1999 , 09, 197-213	0.3	20

379	Algorithmic bombardment for the iterative solution of linear systems: A poly-iterative approach. <i>Journal of Computational and Applied Mathematics</i> , 1996 , 74, 91-109	2.4	20
378	Parallel loops \square A test suite for parallelizing compilers: Description and example results. <i>Parallel Computing</i> , 1991 , 17, 1247-1255	1	20
377	The eigenvalue problem for Hermitian matrices with time reversal symmetry. <i>Linear Algebra and Its Applications</i> , 1984 , 60, 27-42	0.9	20
376	Divide and Conquer on Hybrid GPU-Accelerated Multicore Systems. <i>SIAM Journal of Scientific Computing</i> , 2012 , 34, C70-C82	2.6	19
375	A Class of Hybrid LAPACK Algorithms for Multicore and GPU Architectures 2011 ,		19
374	MPI collective algorithm selection and quadtree encoding. <i>Parallel Computing</i> , 2007 , 33, 613-623	1	19
373	Static tiling for heterogeneous computing platforms. <i>Parallel Computing</i> , 1999 , 25, 547-568	1	19
372	Parallel matrix transpose algorithms on distributed memory concurrent computers. <i>Parallel Computing</i> , 1995 , 21, 1387-1405	1	19
371	Binomial Graph: A Scalable and Fault-Tolerant Logical Network Topology. <i>Lecture Notes in Computer Science</i> , 2007 , 471-482	0.9	19
370	Using Jacobi iterations and blocking for solving sparse triangular systems in incomplete factorization preconditioning. <i>Journal of Parallel and Distributed Computing</i> , 2018 , 119, 219-230	4.4	19
369	Accelerating Linear System Solutions Using Randomization Techniques. <i>ACM Transactions on Mathematical Software</i> , 2013 , 39, 1-13	2.3	18
368	. <i>Computer</i> , 2015 , 48, 42-49	1.6	18
367	Profiling high performance dense linear algebra algorithms on multicore architectures for power and energy efficiency. <i>Computer Science - Research and Development</i> , 2012 , 27, 277-287		18
366	Scalable Tile Communication-Avoiding QR Factorization on Multicore Cluster Systems 2010 ,		18
365	QR Factorization for the Cell Broadband Engine. <i>Scientific Programming</i> , 2009 , 17, 31-42	1.4	18
364	HierKNEM: An Adaptive Framework for Kernel-Assisted and Topology-Aware Collective Communications on Many-core Clusters 2012 ,		18
363	Numerical linear algebra algorithms and software. <i>Journal of Computational and Applied Mathematics</i> , 2000 , 123, 489-514	2.4	18
362	The Netlib Mathematical Software Repository. <i>D-Lib Magazine</i> , 1995 , 1,		18

361	Investigating power capping toward energy-efficient scientific applications. <i>Concurrency Computation Practice and Experience</i> , 2019 , 31, e4485	1.4	18
360	Porting the PLASMA Numerical Library to the OpenMP Standard. <i>International Journal of Parallel Programming</i> , 2017 , 45, 612-633	1.5	17
359	Sunway TaihuLight supercomputer makes its appearance. <i>National Science Review</i> , 2016 , 3, 265-266	10.8	17
358	A block-asynchronous relaxation method for graphics processing units. <i>Journal of Parallel and Distributed Computing</i> , 2013 , 73, 1613-1626	4.4	17
357	Hierarchical QR factorization algorithms for multi-core clusters. <i>Parallel Computing</i> , 2013 , 39, 212-232	1	17
356	Tile QR factorization with parallel panel processing for multicore architectures 2010 ,		17
355	The PlayStation 3 for High-Performance Scientific Computing. <i>Computing in Science and Engineering</i> , 2008 , 10, 84-87	1.5	17
354	DARPA's HPCS Program: History, Models, Tools, Languages. <i>Advances in Computers</i> , 2008 , 1-100	2.9	17
353	A Parallel Implementation of the Nonsymmetric QR Algorithm for Distributed Memory Architectures. <i>SIAM Journal of Scientific Computing</i> , 2002 , 24, 284-311	2.6	17
352	The Design of Fast and Energy-Efficient Linear Solvers: On the Potential of Half-Precision Arithmetic and Iterative Refinement Techniques. <i>Lecture Notes in Computer Science</i> , 2018 , 586-600	0.9	16
351	Unified Development for Mixed Multi-GPU and Multi-coprocessor Environments Using a Lightweight Runtime Environment 2014 ,		16
350	PTG: An Abstraction for Unhindered Parallelism 2014 ,		16
349	A Fast Batched Cholesky Factorization on a GPU 2014 ,		16
348	Achieving numerical accuracy and high performance using recursive tile LU factorization with partial pivoting. <i>Concurrency Computation Practice and Experience</i> , 2014 , 26, 1408-1431	1.4	16
347	A novel hybrid CPU/GPU generalized eigensolver for electronic structure calculations based on fine-grained memory aware tasks. <i>International Journal of High Performance Computing Applications</i> , 2014 , 28, 196-209	1.8	16
346	One-sided Dense Matrix Factorizations on a Multicore with Multiple GPU Accelerators*. <i>Procedia Computer Science</i> , 2012 , 9, 37-46	1.6	16
345	Java access to numerical libraries. <i>Concurrency and Computation: Practice and Experience</i> , 1997 , 9, 1279-1291		16
344	Exploring New Architectures in Accelerating CFD for Air Force Applications 2008 ,		16

343	. <i>Computing in Science and Engineering</i> , 2001 , 3, 32-39	1.5	16
342	A Parallel Algorithm for the Nonsymmetric Eigenvalue Problem. <i>SIAM Journal of Scientific Computing</i> , 1993 , 14, 542-569	2.6	16
341	Solving the secular equation including spin orbit coupling for systems with inversion and time reversal symmetry. <i>Journal of Computational Physics</i> , 1984 , 54, 278-288	4.1	16
340	Asynchronous Iterative Algorithm for Computing Incomplete Factorizations on GPUs. <i>Lecture Notes in Computer Science</i> , 2015 , 1-16	0.9	15
339	Accelerating GPU Kernels for Dense Linear Algebra. <i>Lecture Notes in Computer Science</i> , 2011 , 83-92	0.9	15
338	High Performance Dense Linear System Solver with Soft Error Resilience 2011 ,		15
337	High performance matrix inversion based on LU factorization for multicore architectures 2011 ,		15
336	Improvement of parallelization efficiency of batch pattern BP training algorithm using Open MPI. <i>Procedia Computer Science</i> , 2010 , 1, 525-533	1.6	15
335	Self-healing network for scalable fault-tolerant runtime environments. <i>Future Generation Computer Systems</i> , 2010 , 26, 479-485	7.5	15
334	Deploying fault tolerance and tasks migration with NetSolve. <i>Future Generation Computer Systems</i> , 1999 , 15, 745-755	7.5	15
333	A Checkpoint-on-Failure Protocol for Algorithm-Based Recovery in Standard MPI. <i>Lecture Notes in Computer Science</i> , 2012 , 477-488	0.9	15
332	Towards Achieving Performance Portability Using Directives for Accelerators 2016 ,		15
331	Fault Tolerance Techniques for High-Performance Computing. <i>Computer Communications and Networks</i> , 2015 , 3-85	0.5	14
330	Analysis of dynamically scheduled tile algorithms for dense linear algebra on multicore architectures. <i>Concurrency Computation Practice and Experience</i> , 2012 , 24, 305-321	1.4	14
329	Parallel Two-Sided Matrix Reduction to Band Bidiagonal Form on Multicore Architectures. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2010 , 21, 417-423	3.7	14
328	Analytical modeling and optimization for affinity based thread scheduling on multicore systems 2009 ,		14
327	Computing the conditioning of the components of a linear least-squares solution. <i>Numerical Linear Algebra With Applications</i> , 2009 , 16, 517-533	1.6	14
326	THE VIRTUAL INSTRUMENT: SUPPORT FOR GRID-ENABLED MCELL SIMULATIONS. <i>International Journal of High Performance Computing Applications</i> , 2004 , 18, 3-17	1.8	14

325	Tools to aid in the analysis of memory access patterns for FORTRAN programs. <i>Parallel Computing</i> , 1988 , 9, 25-35	1	14
324	Composing resilience techniques: ABFT, periodic and incremental checkpointing. <i>International Journal of Networking and Computing</i> , 2015 , 5, 2-25	0.2	14
323	Dynamic Reconfiguration and Virtual Machine Management in the Harness Metacomputing System. <i>Lecture Notes in Computer Science</i> , 1998 , 127-134	0.9	14
322	Design for a Soft Error Resilient Dynamic Task-Based Runtime 2015 ,		13
321	Soft error resilient QR factorization for hybrid system with GPGPU. <i>Journal of Computational Science</i> , 2013 , 4, 457-464	3.4	13
320	A survey of recent developments in parallel implementations of Gaussian elimination. <i>Concurrency Computation Practice and Experience</i> , 2015 , 27, 1292-1309	1.4	13
319	Efficient Pattern Search in Large Traces Through Successive Refinement. <i>Lecture Notes in Computer Science</i> , 2004 , 47-54	0.9	13
318	NetSolve: Grid enabling scientific computing environments. <i>Advances in Parallel Computing</i> , 2005 , 14, 33-51	1.1	13
317	Biological sequence alignment on the computational grid using the GrADS framework. <i>Future Generation Computer Systems</i> , 2005 , 21, 980-986	7.5	13
316	Recent Enhancements To Pvm. <i>International Journal of High Performance Computing Applications</i> , 1995 , 9, 108-127		13
315	A Framework for Batched and GPU-Resident Factorization Algorithms Applied to Block Householder Transformations. <i>Lecture Notes in Computer Science</i> , 2015 , 31-47	0.9	13
314	A survey of numerical linear algebra methods utilizing mixed-precision arithmetic. <i>International Journal of High Performance Computing Applications</i> , 2021 , 35, 344-369	1.8	13
313	ADAPT 2018 ,		13
312	Request Sequencing: Optimizing Communication for the Grid. <i>Lecture Notes in Computer Science</i> , 2000 , 1213-1222	0.9	13
311	Acceleration of GPU-based Krylov solvers via data transfer reduction. <i>International Journal of High Performance Computing Applications</i> , 2015 , 29, 366-383	1.8	12
310	A new metric for ranking high-performance computing systems. <i>National Science Review</i> , 2016 , 3, 30-35	10.8	12
309	PaRSEC in Practice: Optimizing a Legacy Chemistry Application through Distributed Task-Based Execution 2015 ,		12
308	Power monitoring with PAPI for extreme scale architectures and dataflow-based programming models 2014 ,		12

307	Block-asynchronous Multigrid Smoothers for GPU-accelerated Systems. <i>Procedia Computer Science</i> , 2012 , 9, 7-16	1.6	12
306	NetSolve: Past, Present, and Future – A Look at a Grid Enabled Server	615-624	12
305	Numerically Stable Real Number Codes Based on Random Matrices. <i>Lecture Notes in Computer Science</i> , 2005 , 115-122	0.9	12
304	. <i>IEEE Computational Science and Engineering</i> , 1995 , 2, 62-69		12
303	SLATE 2019 ,		12
302	Retrospect: Deterministic Replay of MPI Applications for Interactive Distributed Debugging. <i>Lecture Notes in Computer Science</i> , 2007 , 297-306	0.9	12
301	. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2018 , 29, 973-984	3.7	12
300	With Extreme Computing, the Rules Have Changed. <i>Computing in Science and Engineering</i> , 2017 , 19, 52-62.5		11
299	Algorithm-Based Fault Tolerance for Dense Matrix Factorizations, Multiple Failures and Accuracy. <i>ACM Transactions on Parallel Computing</i> , 2015 , 1, 1-28	1.4	11
298	Updating incomplete factorization preconditioners for model order reduction. <i>Numerical Algorithms</i> , 2016 , 73, 611-630	2.1	11
297	Kernel-assisted and topology-aware MPI collective communications on multicore/many-core platforms. <i>Journal of Parallel and Distributed Computing</i> , 2013 , 73, 1000-1010	4.4	11
296	Experiences in autotuning matrix multiplication for energy minimization on GPUs. <i>Concurrency Computation Practice and Experience</i> , 2015 , 27, 5096-5113	1.4	11
295	A Parallel Tiled Solver for Dense Symmetric Indefinite Systems on Multicore Architectures 2012 ,		11
294	Key concepts for parallel out-of-core LU factorization. <i>Parallel Computing</i> , 1997 , 23, 49-70	1	11
293	A Scalable Checkpoint Encoding Algorithm for Diskless Checkpointing 2008 ,		11
292	Accurate Cache and TLB Characterization Using Hardware Counters. <i>Lecture Notes in Computer Science</i> , 2004 , 432-439	0.9	11
291	Recursive Approach in Sparse Matrix LU Factorization. <i>Scientific Programming</i> , 2001 , 9, 51-60	1.4	11
290	JLAPACK – Compiling LAPACK FORTRAN to Java. <i>Scientific Programming</i> , 1999 , 7, 111-138	1.4	11

289	Batched Gauss-Jordan Elimination for Block-Jacobi Preconditioner Generation on GPUs 2017 ,		10
288	Towards batched linear solvers on accelerated hardware platforms 2015 ,		10
287	Power-aware computing: Measurement, control, and performance analysis for Intel Xeon Phi 2017 ,		10
286	Optimizing Krylov Subspace Solvers on Graphics Processing Units 2014 ,		10
285	Portable HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi. <i>Lecture Notes in Computer Science</i> , 2014 , 571-581	0.9	10
284	Toward a High Performance Tile Divide and Conquer Algorithm for the Dense Symmetric Eigenvalue Problem. <i>SIAM Journal of Scientific Computing</i> , 2012 , 34, C249-C274	2.6	10
283	High Performance Dense Linear System Solver with Resilience to Multiple Soft Errors. <i>Procedia Computer Science</i> , 2012 , 9, 216-225	1.6	10
282	A Comprehensive Study of Task Coalescing for Selecting Parallelism Granularity in a Two-Stage Bidiagonal Reduction 2012 ,		10
281	Matrix product on heterogeneous master-worker platforms 2008 ,		10
280	Automatic analysis of inefficiency patterns in parallel applications. <i>Concurrency Computation Practice and Experience</i> , 2007 , 19, 1481-1496	1.4	10
279	Automatic blocking of QR and LU factorizations for locality 2004 ,		10
278	Numerical libraries and the grid 2001 ,		10
277	Programming methodology and performance issues for advanced computer architectures. <i>Parallel Computing</i> , 1988 , 8, 41-58	1	10
276	Load-balancing Sparse Matrix Vector Product Kernels on GPUs. <i>ACM Transactions on Parallel Computing</i> , 2020 , 7, 1-26	1.4	10
275	On Using Incremental Profiling for the Performance Analysis of Shared Memory Parallel Applications. <i>Lecture Notes in Computer Science</i> , 2007 , 62-71	0.9	10
274	Evaluation of the HPC Challenge Benchmarks in Virtualized Environments. <i>Lecture Notes in Computer Science</i> , 2012 , 436-445	0.9	10
273	Leading Edge Hybrid Multi-GPU Algorithms for Generalized Eigenproblems in Electronic Structure Calculations. <i>Lecture Notes in Computer Science</i> , 2013 , 67-80	0.9	10
272	Mixed-precision iterative refinement using tensor cores on GPUs to accelerate solution of linear systems. <i>Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences</i> , 2020 , 476, 20200110	2.4	10

271	Fast Cholesky factorization on GPUs for batch and native modes in MAGMA. <i>Journal of Computational Science</i> , 2017 , 20, 85-93	3.4	9
270	Factorization and Inversion of a Million Matrices using GPUs: Challenges and Countermeasures. <i>Procedia Computer Science</i> , 2017 , 108, 606-615	1.6	9
269	An efficient distributed randomized algorithm for solving large dense symmetric indefinite linear systems. <i>Parallel Computing</i> , 2014 , 40, 213-223	1	9
268	A Parallel Solver for Incompressible Fluid Flows. <i>Procedia Computer Science</i> , 2013 , 18, 439-448	1.6	9
267	A Class of Communication-avoiding Algorithms for Solving General Dense Linear Systems on CPU/GPU Parallel Machines. <i>Procedia Computer Science</i> , 2012 , 9, 17-26	1.6	9
266	QCG-OMPI: MPI applications on grids. <i>Future Generation Computer Systems</i> , 2011 , 27, 357-369	7.5	9
265	L2 Cache Modeling for Scientific Applications on Chip Multi-Processors. <i>Parallel Processing (ICPP), Proceedings of the International Symposium</i> , 2007 ,		9
264	Enabling interactive and collaborative oil reservoir simulations on the Grid. <i>Concurrency Computation Practice and Experience</i> , 2005 , 17, 1387-1414	1.4	9
263	Linear algebra on high performance computers. <i>Applied Mathematics and Computation</i> , 1986 , 20, 57-88	2.7	9
262	Implementing Linear Algebra Routines on Multi-core Processors with Pipelining and a Look Ahead 2006 , 147-156		9
261	Failure Detection and Propagation in HPC systems 2016 ,		9
260	ParILUT---A New Parallel Threshold ILU Factorization. <i>SIAM Journal of Scientific Computing</i> , 2018 , 40, C503-C519	2.6	9
259	Distributed Probabilistic Model-Building Genetic Algorithm. <i>Lecture Notes in Computer Science</i> , 2003 , 1015-1028	0.9	9
258	A Comparison of Parallel Solvers for Diagonally Dominant and General Narrow-Banded Linear Systems II. <i>Lecture Notes in Computer Science</i> , 1999 , 1078-1087	0.9	9
257	Performance of asynchronous optimized Schwarz with one-sided communication. <i>Parallel Computing</i> , 2019 , 86, 66-81	1	8
256	Energy efficiency and performance frontiers for sparse computations on GPU supercomputers 2015 ,		8
255	Heterogeneous Streaming 2016 ,		8
254	Communication-Avoiding Symmetric-Indefinite Factorization. <i>SIAM Journal on Matrix Analysis and Applications</i> , 2014 , 35, 1364-1406	1.5	8

253	Novel HPC techniques to batch execution of many variable size BLAS computations on GPUs 2017 ,		8
252	Practical scalable consensus for pseudo-synchronous distributed systems 2015 ,		8
251	Efficient implementation of quantum materials simulations on distributed CPU-GPU systems 2015 ,		8
250	Strengthening compute and data intensive capacities of Armenia 2015 ,		8
249	MAGMA embedded: Towards a dense linear algebra library for energy efficient extreme computing 2015 ,		8
248	Assessing the Impact of ABFT and Checkpoint Composite Strategies 2014 ,		8
247	Power profiling of Cholesky and QR factorizations on distributed memory systems. <i>Computer Science - Research and Development</i> , 2014 , 29, 139-147		8
246	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 ,		8
245	Feedback-directed thread scheduling with memory considerations 2007 ,		8
244	Multithreading for synchronization tolerance in matrix factorization. <i>Journal of Physics: Conference Series</i> , 2007 , 78, 012028	0.3	8
243	An asynchronous algorithm on the NetSolve global computing system. <i>Future Generation Computer Systems</i> , 2006 , 22, 279-290	7.5	8
242	Extreme-Scale Task-Based Cholesky Factorization Toward Climate and Weather Prediction Applications 2020 ,		8
241	Power Management and Event Verification in PAPI 2016 , 41-51		8
240	Scalability Analysis of the SPEC OpenMP Benchmarks on Large-Scale Shared Memory Multiprocessors. <i>Lecture Notes in Computer Science</i> , 2007 , 815-822	0.9	8
239	Fast Batched Matrix Multiplication for Small Sizes Using Half-Precision Arithmetic on GPUs 2019 ,		8
238	Accelerating the SVD two stage bidiagonal reduction and divide and conquer using GPUs. <i>Parallel Computing</i> , 2018 , 74, 3-18	1	8
237	GridSolve: The Evolution of A Network Enabled Solver 2007 , 215-224		8
236	Distributed-memory lattice H-matrix factorization. <i>International Journal of High Performance Computing Applications</i> , 2019 , 33, 1046-1063	1.8	7

235	Comparing the performance of rigid, moldable and grid-shaped applications on failure-prone HPC platforms. <i>Parallel Computing</i> , 2019 , 85, 1-12	1	7
234	Looking back at dense linear algebra software. <i>Journal of Parallel and Distributed Computing</i> , 2014 , 74, 2548-2560	4.4	7
233	clMAGMA 2014 ,		7
232	Revisiting the Double Checkpointing Algorithm 2013 ,		7
231	A look back on 30 years of the Gordon Bell Prize. <i>International Journal of High Performance Computing Applications</i> , 2017 , 31, 469-484	1.8	7
230	HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi. <i>Scientific Programming</i> , 2015 , 2015, 1-11	1.4	7
229	A scalable approach to solving dense linear algebra problems on hybrid CPU-GPU systems. <i>Concurrency Computation Practice and Experience</i> , 2015 , 27, 3702-3723	1.4	7
228	Anatomy of a globally recursive embedded LINPACK benchmark 2012 ,		7
227	Reducing the Amount of Pivoting in Symmetric Indefinite Systems. <i>Lecture Notes in Computer Science</i> , 2012 , 133-142	0.9	7
226	Soft error resilient QR factorization for hybrid system with GPGPU 2011 ,		7
225	Paravirtualization effect on single- and multi-threaded memory-intensive linear algebra software. <i>Cluster Computing</i> , 2009 , 12, 101-122	2.1	7
224	Conjugate-gradient eigenvalue solvers in computing electronic properties of nanostructure architectures. <i>International Journal of Computational Science and Engineering</i> , 2006 , 2, 205	0.4	7
223	Trends in High Performance Computing. <i>Computer Journal</i> , 2004 , 47, 399-403	1.3	7
222	GrADSolve ^π grid-based RPC system for parallel computing with application-level scheduling. <i>Journal of Parallel and Distributed Computing</i> , 2004 , 64, 774-783	4.4	7
221	Middleware for the use of storage in communication. <i>Parallel Computing</i> , 2002 , 28, 1773-1787	1	7
220	Software distribution using Xnetlib. <i>ACM Transactions on Mathematical Software</i> , 1995 , 21, 79-88	2.3	7
219	Randomized algorithms to update partial singular value decomposition on a hybrid CPU/GPU cluster 2015 ,		7
218	A Scalable High Performant Cholesky Factorization for Multicore with GPU Accelerators. <i>Lecture Notes in Computer Science</i> , 2011 , 93-101	0.9	7

217	Beyond the CPU: Hardware Performance Counter Monitoring on Blue Gene/Q. <i>Lecture Notes in Computer Science</i> , 2013 , 213-225	0.9	7
216	A Framework for Out of Memory SVD Algorithms. <i>Lecture Notes in Computer Science</i> , 2017 , 158-178	0.9	7
215	Efficiency of General Krylov Methods on GPUs -- An Experimental Study 2016 ,		7
214	Generic Matrix Multiplication for Multi-GPU Accelerated Distributed-Memory Platforms over PaRSEC 2019 ,		7
213	A Scalable Approach to MPI Application Performance Analysis. <i>Lecture Notes in Computer Science</i> , 2005 , 309-316	0.9	7
212	Performance Instrumentation and Measurement for Terascale Systems. <i>Lecture Notes in Computer Science</i> , 2003 , 53-62	0.9	7
211	Self-Adapting Numerical Software and Automatic Tuning of Heuristics. <i>Lecture Notes in Computer Science</i> , 2003 , 759-767	0.9	7
210	High-performance Cholesky factorization for GPU-only execution 2017 ,		6
209	Optimization for performance and energy for batched matrix computations on GPUs 2015 ,		6
208	Numerical algorithms for high-performance computational science. <i>Philosophical Transactions Series A, Mathematical, Physical, and Engineering Sciences</i> , 2020 , 378, 20190066	3	6
207	Accelerating the SVD bi-diagonalization of a batch of small matrices using GPUs. <i>Journal of Computational Science</i> , 2018 , 26, 237-245	3.4	6
206	A failure detector for HPC platforms. <i>International Journal of High Performance Computing Applications</i> , 2018 , 32, 139-158	1.8	6
205	Search Space Generation and Pruning System for Autotuners 2016 ,		6
204	Improving Performance of GMRES by Reducing Communication and Pipelining Global Collectives 2017 ,		6
203	Adaptive precision solvers for sparse linear systems 2015 ,		6
202	Scaling up matrix computations on shared-memory manycore systems with 1000 CPU cores 2014 ,		6
201	Multi-GPU Implementation of LU Factorization. <i>Procedia Computer Science</i> , 2012 , 9, 106-115	1.6	6
200	A Hybridization Methodology for High-Performance Linear Algebra Software for GPUs 2012 , 473-484		6

199	An Implementation of the Tile QR Factorization for a GPU and Multiple CPUs. <i>Lecture Notes in Computer Science</i> , 2012 , 248-257	0.9	6
198	Using agent-based software for scientific computing in the NetSolve system. <i>Parallel Computing</i> , 1998 , 24, 1777-1790	1	6
197	The use of bulk states to accelerate the band edge state calculation of a semiconductor quantum dot. <i>Journal of Computational Physics</i> , 2007 , 223, 774-782	4.1	6
196	HARNESS fault tolerant MPI design, usage and performance issues. <i>Future Generation Computer Systems</i> , 2002 , 18, 1127-1142	7.5	6
195	Logistical quality of service in NetSolve. <i>Computer Communications</i> , 1999 , 22, 1034-1044	5.1	6
194	PB-BLAS: a set of parallel block basic linear algebra subprograms. <i>Concurrency and Computation: Practice and Experience</i> , 1996 , 8, 517-535		6
193	Parallel Tiled QR Factorization for Multicore Architectures 2007 , 639-648		6
192	A Holistic Approach for Performance Measurement and Analysis for Petascale Applications. <i>Lecture Notes in Computer Science</i> , 2009 , 686-695	0.9	6
191	Enhancing Parallelism of Tile Bidiagonal Transformation on Multicore Architectures Using Tree Reduction. <i>Lecture Notes in Computer Science</i> , 2012 , 661-670	0.9	6
190	Stability and Performance of Various Singular Value QR Implementations on Multicore CPU with a GPU. <i>ACM Transactions on Mathematical Software</i> , 2016 , 43, 1-18	2.3	6
189	Why is it Hard to Describe Properties of Algorithms?. <i>Procedia Computer Science</i> , 2016 , 101, 4-7	1.6	6
188	Performance Tuning and Optimization Techniques of Fixed and Variable Size Batched Cholesky Factorization on GPUs. <i>Procedia Computer Science</i> , 2016 , 80, 119-130	1.6	6
187	Performance Analysis of Tile Low-Rank Cholesky Factorization Using PaRSEC Instrumentation Tools 2019 ,		6
186	PAPI software-defined events for in-depth performance analysis. <i>International Journal of High Performance Computing Applications</i> , 2019 , 33, 1113-1127	1.8	5
185	Batched one-sided factorizations of tiny matrices using GPUs: Challenges and countermeasures. <i>Journal of Computational Science</i> , 2018 , 26, 226-236	3.4	5
184	MIAMI: A framework for application performance diagnosis 2014 ,		5
183	On the performance and energy efficiency of sparse linear algebra on GPUs. <i>International Journal of High Performance Computing Applications</i> , 2017 , 31, 375-390	1.8	5
182	Non-GPU-resident symmetric indefinite factorization. <i>Concurrency Computation Practice and Experience</i> , 2017 , 29, e4012	1.4	5

181	2017,		5
180	Tuning stationary iterative solvers for fault resilience 2015,		5
179	Hierarchical QR Factorization Algorithms for Multi-core Cluster Systems 2012,		5
178	Scheduling Two-Sided Transformations Using Tile Algorithms on Multicore Architectures. <i>Scientific Programming</i> , 2010 , 18, 35-50	1.4	5
177	Processes Distribution of Homogeneous Parallel Linear Algebra Routines on Heterogeneous Clusters 2005,		5
176	An Iterative Solver Benchmark. <i>Scientific Programming</i> , 2001 , 9, 223-231	1.4	5
175	Programming tools and environments. <i>Communications of the ACM</i> , 1998 , 41, 64-73	2.5	5
174	Stochastic Performance Prediction for Iterative Algorithms in Distributed Environments. <i>Journal of Parallel and Distributed Computing</i> , 1999 , 58, 68-91	4.4	5
173	Tiling on systems with communication/computation overlap. <i>Concurrency and Computation: Practice and Experience</i> , 1999 , 11, 139-153		5
172	Preliminary Results of Autotuning GEMM Kernels for the NVIDIA Kepler Architecture- GeForce GTX 680		5
171	Optimal Checkpointing Period: Time vs. Energy. <i>Lecture Notes in Computer Science</i> , 2014 , 203-214	0.9	5
170	Domain Overlap for Iterative Sparse Triangular Solves on GPUs. <i>Lecture Notes in Computational Science and Engineering</i> , 2016 , 527-545	0.3	5
169	Performance Instrumentation and Compiler Optimizations for MPI/OpenMP Applications. <i>Lecture Notes in Computer Science</i> , 2008 , 267-278	0.9	5
168	Towards an Efficient Tile Matrix Inversion of Symmetric Positive Definite Matrices on Multicore Architectures. <i>Lecture Notes in Computer Science</i> , 2011 , 129-138	0.9	5
167	Optimizing Memory-Bound SYMV Kernel on GPU Hardware Accelerators. <i>Lecture Notes in Computer Science</i> , 2013 , 72-79	0.9	5
166	Efficient exascale discretizations: High-order finite element methods. <i>International Journal of High Performance Computing Applications</i> , 109434202110208	1.8	5
165	LU, QR, and Cholesky factorizations: Programming model, performance analysis and optimization techniques for the Intel Knights Landing Xeon Phi 2016,		5
164	Linear algebra software for large-scale accelerated multicore computing*. <i>Acta Numerica</i> , 2016 , 25, 1-160	5.1	5

163	Impacts of Multi-GPU MPI Collective Communications on Large FFT Computation 2019 ,		5
162	. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2019 , 30, 1158-1169	3.7	5
161	Variable-size batched Gauss-Jordan elimination for block-Jacobi preconditioning on graphics processors. <i>Parallel Computing</i> , 2019 , 81, 131-146	1	5
160	. <i>Proceedings of the IEEE</i> , 2018 , 106, 2040-2055	14.3	5
159	ParILUT - A Parallel Threshold ILU for GPUs 2019 ,		4
158	2013 ,		4
157	Optimizing the SVD Bidiagonalization Process for a Batch of Small Matrices. <i>Procedia Computer Science</i> , 2017 , 108, 1008-1018	1.6	4
156	Mixed-precision block gram Schmidt orthogonalization 2015 ,		4
155	Dynamically Balanced Synchronization-Avoiding LU Factorization with Multicore and GPUs 2014 ,		4
154	2011 ,		4
153	Changing technologies of HPC. <i>Future Generation Computer Systems</i> , 1997 , 12, 461-474	7.5	4
152	GrADSolve [RPC for High Performance Computing on the Grid. <i>Lecture Notes in Computer Science</i> , 2003 , 394-403	0.9	4
151	Automatic translation of Fortran to JVM bytecode. <i>Concurrency Computation Practice and Experience</i> , 2003 , 15, 207-222	1.4	4
150	High performance linear algebra package for FORTRAN 90. <i>Lecture Notes in Computer Science</i> , 1998 , 579-583	0.9	4
149	Performance of random sampling for computing low-rank approximations of a dense matrix on GPUs 2015 ,		4
148	Dense Linear Algebra for Hybrid GPU-Based Systems. <i>Chapman & Hall/CRC Computational Science</i> , 2010 , 37-55		4
147	A Block-Asynchronous Relaxation Method for Graphics Processing Units		4
146	MagmaDNN: Towards High-Performance Data Analytics and Machine Learning for Data-Driven Scientific Computing. <i>Lecture Notes in Computer Science</i> , 2019 , 490-503	0.9	4

145	Performance optimization of Sparse Matrix-Vector Multiplication for multi-component PDE-based applications using GPUs. <i>Concurrency Computation Practice and Experience</i> , 2016 , 28, 3447-3465	1.4	4
144	Performance-Portable Autotuning of OpenCL Kernels for Convolutional Layers of Deep Neural Networks 2016 ,		4
143	Batched Generation of Incomplete Sparse Approximate Inverses on GPUs 2016 ,		4
142	Data through the Computational Lens, Preface for ICCS 2016. <i>Procedia Computer Science</i> , 2016 , 80, 1-7	1.6	4
141	Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling. <i>Journal of Advances in Modeling Earth Systems</i> , 2018 , 10, 1952-1969	7.1	4
140	Hash Functions for Datatype Signatures in MPI. <i>Lecture Notes in Computer Science</i> , 2005 , 76-83	0.9	4
139	On the Design, Development, and Analysis of Optimized Matrix-Vector Multiplication Routines for Coprocessors. <i>Lecture Notes in Computer Science</i> , 2015 , 58-73	0.9	3
138	Mixing LU and QR factorization algorithms to design high-performance dense linear algebra solvers. <i>Journal of Parallel and Distributed Computing</i> , 2015 , 85, 32-46	4.4	3
137	. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2018 , 29, 1879-1892	3.7	3
136	Optimization and performance evaluation of the IDR iterative Krylov solver on GPUs. <i>International Journal of High Performance Computing Applications</i> , 2018 , 32, 220-230	1.8	3
135	Assessing the cost of redistribution followed by a computational kernel: Complexity and performance results. <i>Parallel Computing</i> , 2016 , 52, 22-41	1	3
134	Performance of Hierarchical-matrix BiCGStab Solver on GPU Clusters 2018 ,		3
133	Analysis and Design Techniques towards High-Performance and Energy-Efficient Dense Linear Solvers on GPUs. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2018 , 29, 2700-2712	3.7	3
132	Optimized Batched Linear Algebra for Modern Architectures. <i>Lecture Notes in Computer Science</i> , 2017 , 511-522	0.9	3
131	Variable-Size Batched LU for Small Matrices and Its Integration into Block-Jacobi Preconditioning 2017 ,		3
130	Variable-Size Batched Gauss-Huard for Block-Jacobi Preconditioning. <i>Procedia Computer Science</i> , 2017 , 108, 1783-1792	1.6	3
129	Towards numerical benchmark for half-precision floating point arithmetic 2017 ,		3
128	Towards batched linear solvers on accelerated hardware platforms. <i>ACM SIGPLAN Notices</i> , 2015 , 50, 261-262		3

127	Performance and reliability trade-offs for the double checkpointing algorithm. <i>International Journal of Networking and Computing</i> , 2014 , 4, 23-41	0.2	3
126	POSTER: Utilizing dataflow-based execution for coupled cluster methods 2014 ,		3
125	Performance and Portability with OpenCL for Throughput-Oriented HPC Workloads across Accelerators, Coprocessors, and Multicore Processors 2014 ,		3
124	BlackjackBench: Portable Hardware Characterization with Automated Results' Analysis. <i>Computer Journal</i> , 2014 , 57, 1002-1016	1.3	3
123	From Serial Loops to Parallel Execution on Distributed Systems. <i>Lecture Notes in Computer Science</i> , 2012 , 246-257	0.9	3
122	Toward a scalable multi-GPU eigensolver via compute-intensive kernels and efficient communication 2013 ,		3
121	Self Adaptive Application Level Fault Tolerance for Parallel and Distributed Computing 2007 ,		3
120	High Performance Development for High End Computing With Python Language Wrapper (PLW). <i>International Journal of High Performance Computing Applications</i> , 2007 , 21, 360-369	1.8	3
119	Design of Interactive Environment for Numerically Intensive Parallel Linear Algebra Calculations. <i>Lecture Notes in Computer Science</i> , 2004 , 270-277	0.9	3
118	Comparison of Nonlinear Conjugate-Gradient Methods for Computing the Electronic Properties of Nanostructure Architectures. <i>Lecture Notes in Computer Science</i> , 2005 , 317-325	0.9	3
117	Numerical Libraries and Tools for Scalable Parallel Cluster Computing. <i>International Journal of High Performance Computing Applications</i> , 2001 , 15, 175-180	1.8	3
116	. <i>IEEE Software</i> , 1989 , 6, 78-85	1.5	3
115	Design, Optimization, and Benchmarking of Dense Linear Algebra Algorithms on AMD GPUs 2020 ,		3
114	Improving the Performance of the GMRES Method Using Mixed-Precision Techniques. <i>Communications in Computer and Information Science</i> , 2020 , 51-66	0.3	3
113	TOP500 Supercomputers for June 2002		3
112	Truss Structural Optimization using NetSolve System. <i>The Proceedings of OPTIS</i> , 2002 , 2002.5, 141-146	0	3
111	Mixed-Precision Orthogonalization Scheme and Adaptive Step Size for Improving the Stability and Performance of CA-GMRES on GPUs. <i>Lecture Notes in Computer Science</i> , 2015 , 17-30	0.9	3
110	Accelerating NWChem Coupled Cluster Through Dataflow-Based Execution. <i>Lecture Notes in Computer Science</i> , 2016 , 366-376	0.9	3

109	Dense Symmetric Indefinite Factorization on GPU Accelerated Architectures. <i>Lecture Notes in Computer Science</i> , 2016 , 86-95	0.9	3
108	Prospectus for the Next LAPACK and ScaLAPACK Libraries 2006 , 11-23		3
107	A Fully Empirical Autotuned Dense QR Factorization for Multicore Architectures. <i>Lecture Notes in Computer Science</i> , 2011 , 194-205	0.9	3
106	GPU-Accelerated Asynchronous Error Correction for Mixed Precision Iterative Refinement. <i>Lecture Notes in Computer Science</i> , 2012 , 908-919	0.9	3
105	Reducing the Time to Tune Parallel Dense Linear Algebra Routines with Partial Execution and Performance Modeling. <i>Lecture Notes in Computer Science</i> , 2012 , 730-739	0.9	3
104	Multi-criteria Checkpointing Strategies: Response-Time versus Resource Utilization. <i>Lecture Notes in Computer Science</i> , 2013 , 420-431	0.9	3
103	A Set of Batched Basic Linear Algebra Subprograms and LAPACK Routines. <i>ACM Transactions on Mathematical Software</i> , 2021 , 47, 1-23	2.3	3
102	On the Development of Variable Size Batched Computation for Heterogeneous Parallel Architectures 2016 ,		3
101	Massively Parallel Automated Software Tuning 2019 ,		2
100	Accelerating NWChem Coupled Cluster through dataflow-based execution. <i>International Journal of High Performance Computing Applications</i> , 2018 , 32, 540-551	1.8	2
99	Task-Based Cholesky Decomposition on Knights Corner Using OpenMP. <i>Lecture Notes in Computer Science</i> , 2016 , 544-562	0.9	2
98	Virtual Systolic Array for QR Decomposition 2013 ,		2
97	Tridiagonalization of a Symmetric Dense Matrix on a GPU Cluster 2013 ,		2
96	Flexible batched sparse matrix-vector product on GPUs 2017 ,		2
95	Autotuning batch Cholesky factorization in CUDA with interleaved layout of matrices 2017 ,		2
94	GPU-accelerated co-design of induced dimension reduction 2015 ,		2
93	Weighted dynamic scheduling with many parallelism grains for offloading of numerical workloads to multiple varied accelerators 2015 ,		2
92	Access-averse framework for computing low-rank matrix approximations 2014 ,		2

91	Designing LU-QR Hybrid Solvers for Performance and Stability 2014 ,		2
90	Trace-based performance analysis for the petascale simulation code FLASH. <i>International Journal of High Performance Computing Applications</i> , 2011 , 25, 428-439	1.8	2
89	BlackjackBench. <i>Performance Evaluation Review</i> , 2012 , 40, 74-79	0.4	2
88	SmartGridRPC: The new RPC model for high performance Grid computing. <i>Concurrency Computation Practice and Experience</i> , 2010 , 22, 2467-2487	1.4	2
87	Request Sequencing: Enabling Workflow for Efficient Problem Solving in GridSolve 2008 ,		2
86	Special section: Grid computing and the message passing interface. <i>Future Generation Computer Systems</i> , 2008 , 24, 119-120	7.5	2
85	NetBuild: transparent cross-platform access to computational software libraries. <i>Concurrency Computation Practice and Experience</i> , 2002 , 14, 1445-1456	1.4	2
84	NanoPSE: Nanoscience Problem Solving Environment for atomistic electronic structure of semiconductor nanostructures. <i>Journal of Physics: Conference Series</i> , 2005 , 16, 277-282	0.3	2
83	The Component Structure of a Self-Adapting Numerical Software System. <i>International Journal of Parallel Programming</i> , 2005 , 33, 137-143	1.5	2
82	Comparison of the CRAY X-MP-4, Fujitsu VP-200, and Hitachi S-810/20. <i>Simulation</i> , 1986 , 47, 93-107	1.2	2
81	BLAS for GPUs. <i>Chapman & Hall/CRC Computational Science</i> , 2010 , 57-80		2
80	Flexible Data Redistribution in a Task-Based Runtime System 2020 ,		2
79	Linear Systems Solvers for Distributed-Memory Machines with GPU Accelerators. <i>Lecture Notes in Computer Science</i> , 2019 , 495-506	0.9	2
78	HAN: a Hierarchical Autotuned Collective Communication Framework 2020 ,		2
77	Investigating the Benefit of FP16-Enabled Mixed-Precision Solvers for Symmetric Positive Definite Matrices Using GPUs. <i>Lecture Notes in Computer Science</i> , 2020 , 237-250	0.9	2
76	Accelerating Computation of Eigenvectors in the Dense Nonsymmetric Eigenvalue Problem. <i>Lecture Notes in Computer Science</i> , 2015 , 182-191	0.9	2
75	Fast and Small Short Vector SIMD Matrix Multiplication Kernels for the Synergistic Processing Element of the CELL Processor. <i>Lecture Notes in Computer Science</i> , 2008 , 935-944	0.9	2
74	Matrix multiplication on batches of small matrices in half and half-complex precisions. <i>Journal of Parallel and Distributed Computing</i> , 2020 , 145, 188-201	4.4	2

73	Checkpointing Strategies for Shared High-Performance Computing Platforms. <i>International Journal of Networking and Computing</i> , 2019 , 9, 28-52	0.2	2
72	Increasing Accuracy of Iterative Refinement in Limited Floating-Point Arithmetic on Half-Precision Accelerators 2019 ,		2
71	Fine-grained bit-flip protection for relaxation methods. <i>Journal of Computational Science</i> , 2019 , 36, 1005-1013	0.2	2
70	Dense Linear Algebra on Accelerated Multicore Hardware 2012 , 123-146		2
69	Solving dense symmetric indefinite systems using GPUs. <i>Concurrency Computation Practice and Experience</i> , 2017 , 29, e4055	1.4	1
68	Structure-Aware Linear Solver for Realtime Convex Optimization for Embedded Systems. <i>IEEE Embedded Systems Letters</i> , 2017 , 9, 61-64	1	1
67	Reducing the amount of out-of-core data access for GPU-accelerated randomized SVD. <i>Concurrency Computation Practice and Experience</i> , 2020 , 32, e5754	1.4	1
66	Least squares solvers for distributed-memory machines with GPU accelerators 2019 ,		1
65	Sampling algorithms to update truncated SVD 2017 ,		1
64	Performance Analysis and Optimisation of Two-sided Factorization Algorithms for Heterogeneous Platform. <i>Procedia Computer Science</i> , 2015 , 51, 180-190	1.6	1
63	Hybrid Multi-elimination ILU Preconditioners on GPUs 2014 ,		1
62	Deflation Strategies to Improve the Convergence of Communication-Avoiding GMRES 2014 ,		1
61	Parallel Simulation of Superscalar Scheduling 2014 ,		1
60	2012 ,		1
59	Mixed-Tool Performance Analysis on Hybrid Multicore Architectures 2010 ,		1
58	REVISITING MATRIX PRODUCT ON MASTER-WORKER PLATFORMS. <i>International Journal of Foundations of Computer Science</i> , 2008 , 19, 1317-1336	0.6	1
57	13. Parallel Linear Algebra Software 2006 , 233-247		1
56	Optimal Routing in Binomial Graph Networks 2007 ,		1

55	Location-independent naming for virtual distributed software repositories. <i>Software Engineering Notes: an Informal Newsletter of the Special Interest Committee on Software Engineering / ACM</i> , 1995 , 20, 179-185	0.4	1
54	Templates for linear algebra problems. <i>Lecture Notes in Computer Science</i> , 1995 , 115-140	0.9	1
53	Using Advanced Vector Extensions AVX-512 for MPI Reductions 2020 ,		1
52	Integrating Deep Learning in Domain Sciences at Exascale. <i>Communications in Computer and Information Science</i> , 2020 , 35-50	0.3	1
51	Self-healing in Binomial Graph Networks 2007 , 1032-1041		1
50	Evaluating Data Redistribution in PaRSEC. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2022 , 33, 1856-1872	3.7	1
49	Self-Adapting Software for Numerical Linear Algebra Library Routines on Clusters. <i>Lecture Notes in Computer Science</i> , 2003 , 665-672	0.9	1
48	Counter Inspection Toolkit: Making Sense Out of Hardware Performance Events 2019 , 17-37		1
47	Hands-On Research and Training in High Performance Data Sciences, Data Analytics, and Machine Learning for Emerging Environments. <i>Lecture Notes in Computer Science</i> , 2019 , 643-655	0.9	1
46	Accelerating the Conjugate Gradient Algorithm with GPUs in CFD Simulations. <i>Lecture Notes in Computer Science</i> , 2017 , 35-43	0.9	1
45	Implementing Matrix Multiplication on the Cell B. E.. <i>Chapman & Hall/CRC Computational Science</i> , 2010 , 3-20		1
44	Weighted Block-Asynchronous Iteration on GPU-Accelerated Systems. <i>Lecture Notes in Computer Science</i> , 2013 , 145-154	0.9	1
43	MAGMA templates for scalable linear algebra on emerging architectures. <i>International Journal of High Performance Computing Applications</i> , 2020 , 34, 645-658	1.8	1
42	Translational process: Mathematical software perspective. <i>Journal of Computational Science</i> , 2021 , 52, 101216	3.4	1
41	Performance analysis and acceleration of explicit integration for large kinetic networks using batched GPU computations 2016 ,		1
40	Towards Half-Precision Computation for Complex Matrices: A Case Study for Mixed Precision Solvers on GPUs 2019 ,		1
39	Progressive Optimization of Batched LU Factorization on GPUs 2019 ,		1
38	Accelerating Restarted GMRES With Mixed Precision Arithmetic. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2021 , 1-1	3.7	1

37	Accelerating Geostatistical Modeling and Prediction With Mixed-Precision Computations: A High-Productivity Approach with PaRSEC. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2021 , 1-1	3.7	1
36	Eigenvalue Computation with NetSolve Global Computing System. <i>Lecture Notes in Computer Science</i> , 2006 , 446-453	0.9	1
35	Evaluation of dataflow programming models for electronic structure theory. <i>Concurrency Computation Practice and Experience</i> , 2018 , 30, e4490	1.4	0
34	Using long vector extensions for MPI reductions. <i>Parallel Computing</i> , 2022 , 109, 102871	1	0
33	Transparent Cross-Platform Access to Software Services Using GridSolve and GridRPC 2010 , 253-274		0
32	Implementing a Systolic Algorithm for QR Factorization on Multicore Clusters with PaRSEC. <i>Lecture Notes in Computer Science</i> , 2014 , 657-667	0.9	0
31	An Introduction to High Performance Computing and Its Intersection with Advances in Modeling Rare Earth Elements and Actinides. <i>ACS Symposium Series</i> , 3-53	0.4	0
30	Batch QR Factorization on GPUs: Design, Optimization, and Tuning. <i>Lecture Notes in Computer Science</i> , 2022 , 60-74	0.9	0
29	Guest editors Note: Special issue on clusters, clouds, and data for scientific computing. <i>International Journal of High Performance Computing Applications</i> , 2019 , 33, 1067-1068	1.8	
28	Guest editors Note: Special issue on clusters, clouds and data for scientific computing. <i>International Journal of High Performance Computing Applications</i> , 2018 , 32, 3-3	1.8	
27	Enabling workflows in GridSolve: request sequencing and service trading. <i>Journal of Supercomputing</i> , 2013 , 64, 1133-1152	2.5	
26	Guest Editors Note: Special Issue on Clusters, Clouds and Data for Scientific Computing. <i>Parallel Processing Letters</i> , 2015 , 25, 1502002	0.3	
25	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. <i>Scientific Programming</i> , 2015 , 2015, 1-17	1.4	
24	Design and Implementation of a Large Scale Tree-Based QR Decomposition Using a 3D Virtual Systolic Array and a Lightweight Runtime. <i>Parallel Processing Letters</i> , 2014 , 24, 1442004	0.3	
23	Introduction for August Special Issue CCDSC. <i>International Journal of High Performance Computing Applications</i> , 2013 , 27, 231-231	1.8	
22	GUEST EDITORS' NOTE: SPECIAL ISSUE ON CLUSTERS, CLOUDS, AND DATA FOR SCIENTIFIC COMPUTING. <i>Parallel Processing Letters</i> , 2013 , 23, 1302001	0.3	
21	Selected papers of the Workshop on Clusters, Clouds and Grids for Scientific Computing (CCGSC). <i>International Journal of High Performance Computing Applications</i> , 2011 , 25, 259-260	1.8	
20	GUEST EDITORS NOTE. <i>Parallel Processing Letters</i> , 2011 , 21, 109-109	0.3	

- 19 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 0.6
- 18 IMPROVED RUNTIME AND TRANSFER TIME PREDICTION MECHANISMS IN A NETWORK ENABLED SERVERS MIDDLEWARE. *Parallel Processing Letters*, **2007**, 17, 47-59 0.3
- 17 High Performance Computing Trends and Self Adapting Numerical Software. *Lecture Notes in Computer Science*, **2003**, 1-9 0.9
- 16 Deploying fault-tolerance and task migration with NetSolve. *Lecture Notes in Computer Science*, **1998**, 418-432 0.9
- 15 Comparing Distributed Termination Detection Algorithms for Modern HPC Platforms. *International Journal of Networking and Computing*, **2022**, 12, 26-46 0.2
- 14 Disaster Survival Guide in Petascale Computing. *Chapman & Hall/CRC Computational Science*, **2007**, 263-288
- 13 Do Moldable Applications Perform Better on Failure-Prone HPC Platforms?. *Lecture Notes in Computer Science*, **2019**, 787-799 0.9
- 12 An update notice on the level 3 BLAS. *ACM SIGNUM Newsletter*, **1989**, 24, 9-10
- 11 Providing access to high performance computing technologies. *Lecture Notes in Computer Science*, **1997**, 24-34 0.9
- 10 Block-cyclic array redistribution on networks of workstations. *Lecture Notes in Computer Science*, **1997**, 343-350 0.9
- 9 High performance linear algebra package LAPACK90. *Lecture Notes in Computer Science*, **1998**, 387-391 0.9
- 8 Self-adaptive Multiprecision Preconditioners on Multicore and Manycore Architectures. *Lecture Notes in Computer Science*, **2015**, 115-123 0.9
- 7 A Scalable Non-blocking Multicast Scheme for Distributed DAG Scheduling. *Lecture Notes in Computer Science*, **2009**, 195-204 0.9
- 6 Implementing Matrix Factorizations on the Cell B. E.. *Chapman & Hall/CRC Computational Science*, **2010**, 21-35
- 5 Programming the LU Factorization for a Multicore System with Accelerators. *Lecture Notes in Computer Science*, **2013**, 28-35 0.9
- 4 BLAS. *Discrete Mathematics and Its Applications*, **2013**, 1697-1704
- 3 Techniques for Solving Large-Scale Graph Problems on Heterogeneous Platforms. *Communications in Computer and Information Science*, **2016**, 318-332 0.3
- 2 Exploiting Block Structures of KKT Matrices for Efficient Solution of Convex Optimization Problems. *IEEE Access*, **2021**, 9, 116604-116611 3.5

- 1 Task based Cholesky decomposition on Xeon Phi architectures using OpenMP. *International Journal of Computational Science and Engineering*, **2018**, 17, 310 0.4