Enrique S Quintana-Ort

List of Publications by Citations

234 papers	2,436 citations	25 h-index	41 g-index
256 ext. papers	2,887 ext. citations	1.8 avg, IF	5.04 L-index

#	Paper	IF	Citations
234	iMODS: internal coordinates normal mode analysis server. *Nucleic Acids Research*, **2014**, 42, W271-6	20.1	219
233	rCUDA: Reducing the number of GPU-based accelerators in high performance clusters **2010**,		117
232	The science of deriving dense linear algebra algorithms. *ACM Transactions on Mathematical Software*, **2005**, 31, 1-26	2.3	88
231	Solving stable generalized Lyapunov equations with the matrix sign function. *Numerical Algorithms*, **1999**, 20, 75-100	2.1	88
230	Programming matrix algorithms-by-blocks for thread-level parallelism. *ACM Transactions on Mathematical Software*, **2009**, 36, 1-26	2.3	79
229	Supermatrix out-of-order scheduling of matrix operations for SMP and multi-core architectures **2007**,		64
228	An Extension of the StarSs Programming Model for Platforms with Multiple GPUs. *Lecture Notes in Computer Science*, **2009**, 851-862	0.9	55
227	A complete and efficient CUDA-sharing solution for HPC clusters. *Parallel Computing*, **2014**, 40, 574-588	1	48
226	Parallelizing dense and banded linear algebra libraries using SMPSs. *Concurrency Computation Practice and Experience*, **2009**, 21, 2438-2456	1.4	48
225	Extending OpenMP to Survive the Heterogeneous Multi-Core Era. *International Journal of Parallel Programming*, **2010**, 38, 440-459	1.5	45
224	Evaluation and tuning of the Level 3 CUBLAS for graphics processors. *Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on*, **2008**,		43
223	Enabling CUDA acceleration within virtual machines using rCUDA **2011**,		40
222	Exploiting the capabilities of modern GPUs for dense matrix computations. *Concurrency Computation Practice and Experience*, **2009**, 21, 2457-2477	1.4	40
221	Updating an LU Factorization with Pivoting. *ACM Transactions on Mathematical Software*, **2008**, 35, 1-16	2.3	40
220	The libflame Library for Dense Matrix Computations. *Computing in Science and Engineering*, **2009**, 11, 56-63	1.5	39
219	Tools for Power-Energy Modelling and Analysis of Parallel Scientific Applications **2012**,		39
218	Analytical Modeling Is Enough for High-Performance BLIS. *ACM Transactions on Mathematical Software*, **2016**, 43, 1-18	2.3	38

217	Balanced Truncation Model Reduction of Large-Scale Dense Systems on Parallel Computers. *Mathematical and Computer Modelling of Dynamical Systems*, **2000**, 6, 383-405	1	36
216	Representing linear algebra algorithms in code: the FLAME application program interfaces. *ACM Transactions on Mathematical Software*, **2005**, 31, 27-59	2.3	34
215	Solving Dense Linear Systems on Graphics Processors. *Lecture Notes in Computer Science*, **2008**, 739-748	0.9	34
214	Solving dense linear systems on platforms with multiple hardware accelerators. *ACM SIGPLAN Notices*, **2009**, 44, 121-130	0.2	33
213	A Proposal to Extend the OpenMP Tasking Model for Heterogeneous Architectures. *Lecture Notes in Computer Science*, **2009**, 154-167	0.9	33
212	SuperMatrix **2008**,		32
211	Parallel Computation of 3-D Soil-Structure Interaction in Time Domain with a Coupled FEM/SBFEM Approach. *Journal of Scientific Computing*, **2012**, 52, 446-467	2.3	31
210	Model Reduction Based on Spectral Projection Methods **2005**, 5-48		29
209	Performance of CUDA Virtualized Remote GPUs in High Performance Clusters **2011**,		23
208	Solving Stable Sylvester Equations via Rational Iterative Schemes. *Journal of Scientific Computing*, **2006**, 28, 51-83	2.3	23
207	Formal derivation of algorithms. *ACM Transactions on Mathematical Software*, **2003**, 29, 218-243	2.3	23
206	Adaptive precision in block-Jacobi preconditioning for iterative sparse linear system solvers. *Concurrency Computation Practice and Experience*, **2019**, 31, e4460	1.4	23
205	Spectral division methods for block generalized Schur decompositions. *Mathematics of Computation*, **2004**, 73, 1827-1848	1.6	22
204	Accumulating Householder transformations, revisited. *ACM Transactions on Mathematical Software*, **2006**, 32, 169-179	2.3	21
203	NUMERICAL SOLUTION OF DISCRETE STABLE LINEAR MATRIX EQUATIONS ON MULTICOMPUTERS. *International Journal of Parallel, Emergent and Distributed Systems*, **2002**, 17, 127-146		21
202	Efficient Implementation of Hyperspectral Anomaly Detection Techniques on GPUs and Multicore Processors. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **2014**, 7, 2256-2266	4.7	20
201	An Efficient Implementation of GPU Virtualization in High Performance Clusters. *Lecture Notes in Computer Science*, **2010**, 385-394	0.9	19
200	Solving dense linear systems on platforms with multiple hardware accelerators **2008**,		19

199	Efficient algorithms for generalized algebraic Bernoulli equations based on the matrix sign function. *Numerical Algorithms*, **2007**, 46, 351-368	2.1	18
198	Solving linear-quadratic optimal control problems on parallel computers. *Optimization Methods and Software*, **2008**, 23, 879-909	1.3	17
197	Satisfying your dependencies with SuperMatrix **2007**,		17
196	The FLAME approach: From dense linear algebra algorithms to high-performance multi-accelerator implementations. *Journal of Parallel and Distributed Computing*, **2012**, 72, 1134-1143	4.4	16
195	Condensed forms for the symmetric eigenvalue problem on multi-threaded architectures. *Concurrency Computation Practice and Experience*, **2011**, 23, 694-707	1.4	15
194	Scheduling of QR Factorization Algorithms on SMP and Multi-Core Architectures **2008**,		15
193	PARALLEL DISTRIBUTED SOLVERS FOR LARGE STABLE GENERALIZED LYAPUNOV EQUATIONS. *Parallel Processing Letters*, **1999**, 09, 147-158	0.3	15
192	SLURM Support for Remote GPU Virtualization: Implementation and Performance Study **2014**,		14
191	Analysis and optimization of power consumption in the iterative solution of sparse linear systems on multi-core and many-core platforms **2011**,		14
190	LARGE SCALE SIMULATION OF WAVE PROPAGATION IN SOILS INTERACTING WITH STRUCTURES USING FEM AND SBFEM. *Journal of Computational Acoustics*, **2011**, 19, 75-93		14
189	Parallelization of Multilevel ILU Preconditioners on Distributed-Memory Multiprocessors. *Lecture Notes in Computer Science*, **2012**, 162-172	0.9	14
188	Real-Time Endmember Extraction on Multicore Processors. *IEEE Geoscience and Remote Sensing Letters*, **2011**, 8, 924-928	4.1	13
187	Assessing Power Monitoring Approaches for Energy and Power Analysis of Computers. *Sustainable Computing: Informatics and Systems*, **2014**, 4, 68-82	3	12
186	Solving Some Mysteries in Power Monitoring of Servers: Take Care of Your Wattmeters!. *Lecture Notes in Computer Science*, **2013**, 3-18	0.9	12
185	Exploring large macromolecular functional motions on clusters of multicore processors. *Journal of Computational Physics*, **2013**, 246, 275-288	4.1	11
184	Reducing Energy Consumption of Dense Linear Algebra Operations on Hybrid CPU-GPU Platforms **2012**,		11
183	Improving power efficiency of dense linear algebra algorithms on multi-core processors via slack control **2011**,		11
182	The Generalized Newton Iteration forthe Matrix Sign Function. *SIAM Journal of Scientific Computing* , **2002**, 24, 669-683	2.6	11

181	Using Hybrid CPU-GPU Platforms to Accelerate the Computation of the Matrix Sign Function. *Lecture Notes in Computer Science*, **2010**, 132-139	0.9	11
180	Reduction to Condensed Forms for Symmetric Eigenvalue Problems on Multi-core Architectures. *Lecture Notes in Computer Science*, **2010**, 387-395	0.9	11
179	Batched Gauss-Jordan Elimination for Block-Jacobi Preconditioner Generation on GPUs **2017**,		10
178	Improving the user experience of the rCUDA remote GPU virtualization framework. *Concurrency Computation Practice and Experience*, **2015**, 27, 3746-3770	1.4	10
177	Performance versus energy consumption of hyperspectral unmixing algorithms on multi-core platforms. *Eurasip Journal on Advances in Signal Processing*, **2013**, 2013,	1.9	10
176	Modeling power and energy of the task-parallel Cholesky factorization on multicore processors. *Computer Science - Research and Development*, **2014**, 29, 105-112		10
175	DVFS-control techniques for dense linear algebra operations on multi-core processors. *Computer Science - Research and Development*, **2012**, 27, 289-298		10
174	Parallel codes for computing the numerical rank. *Linear Algebra and Its Applications*, **1998**, 275-276, 451-470	0.9	10
173	Architecture-aware configuration and scheduling of matrix multiplication on asymmetric multicore processors. *Cluster Computing*, **2016**, 19, 1037-1051	2.1	10
172	Reformulated Conjugate Gradient for the Energy-Aware Solution of Linear Systems on GPUs **2013**,		9
171	Unveiling the performance-energy trade-off in iterative linear system solvers for multithreaded processors. *Concurrency Computation Practice and Experience*, **2015**, 27, 885-904	1.4	9
170	Assessing the impact of the CPU power-saving modes on the task-parallel solution of sparse linear systems. *Cluster Computing*, **2014**, 17, 1335-1348	2.1	9
169	A Review of Lightweight Thread Approaches for High Performance Computing **2016**,		9
168	Balancing task- and data-level parallelism to improve performance and energy consumption of matrix computations on the Intel Xeon Phi. *Computers and Electrical Engineering*, **2015**, 46, 95-111	4.3	8
167	Accelerating the Lyapack library using GPUs. *Journal of Supercomputing*, **2013**, 65, 1114-1124	2.5	8
166	Solving Matrix Equations on Multi-Core and Many-Core Architectures. *Algorithms*, **2013**, 6, 857-870	1.8	8
165	Exploiting task and data parallelism in ILUPACKS preconditioned CG solver on NUMA architectures and many-core accelerators. *Parallel Computing*, **2016**, 54, 97-107	1	8
164	Toward a modular precision ecosystem for high-performance computing. *International Journal of High Performance Computing Applications*, **2019**, 33, 1069-1078	1.8	7

163	DMR API: Improving cluster productivity by turning applications into malleable. *Parallel Computing*, **2018**, 78, 54-66	1	7
162	Accelerating the SRP-PHAT algorithm on multi- and many-core platforms using OpenCL. *Journal of Supercomputing*, **2019**, 75, 1284-1297	2.5	7
161	Balanced CSR Sparse Matrix-Vector Product on Graphics Processors. *Lecture Notes in Computer Science*, **2017**, 697-709	0.9	7
160	Efficient Scalable Computing through Flexible Applications and Adaptive Workloads **2017**,		7
159	Optimization of power consumption in the iterative solution of sparse linear systems on graphics processors. *Computer Science - Research and Development*, **2012**, 27, 299-307		7
158	Solving dense generalized eigenproblems on multi-threaded architectures. *Applied Mathematics and Computation*, **2012**, 218, 11279-11289	2.7	7
157	Saving Energy in the LU Factorization with Partial Pivoting on Multi-core Processors **2012**,		7
156	A Runtime System for Programming Out-of-Core Matrix Algorithms-by-Tiles on Multithreaded Architectures. *ACM Transactions on Mathematical Software*, **2012**, 38, 1-25	2.3	7
155	Efficient Solution Of The Rank-Deficient Linear Least Squares Problem. *SIAM Journal of Scientific Computing*, **1998**, 20, 1155-1163	2.6	7
154	Rapid Development of High-Performance Out-of-Core Solvers. *Lecture Notes in Computer Science*, **2006**, 413-422	0.9	7
153	A portable subroutine library for solving linear control problems on distributed memory computers **1999**, 61-87		7
152	GPU-Based Dynamic Wave Field Synthesis Using Fractional Delay Filters and Room Compensation. *IEEE/ACM Transactions on Audio Speech and Language Processing*, **2017**, 25, 435-447	3.6	6
151	. *IEEE Access*, **2019**, 7, 17617-17633	3.5	6
150	Task-Parallel LU Factorization of Hierarchical Matrices Using OmpSs **2017**,		6
149	. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **2014**, 7, 2297-2304	4.7	6
148	Energy-efficient execution of dense linear algebra algorithms on multi-core processors. *Cluster Computing*, **2013**, 16, 497-509	2.1	6
147	Adaptive precision solvers for sparse linear systems **2015**,		6
146	Systematic Fusion of CUDA Kernels for Iterative Sparse Linear System Solvers. *Lecture Notes in Computer Science*, **2015**, 675-686	0.9	6

145	High Performance Matrix Inversion on a Multi-core Platform with Several GPUs **2011**,		6
144	Parallel Model Reduction of Large Linear Descriptor Systems via Balanced Truncation. *Lecture Notes in Computer Science*, **2005**, 340-353	0.9	6
143	EnergySaving Cluster Roll: Power Saving System for Clusters. *Lecture Notes in Computer Science*, **2010**, 162-173	0.9	6
142	GLTO: On the Adequacy of Lightweight Thread Approaches for OpenMP Implementations **2017**,		5
141	Leveraging Data-Parallelism in ILUPACK using Graphics Processors **2014**,		5
140	Assessing the Performance-Energy Balance of Graphics Processors for Spectral Unmixing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, **2014**, 7, 2305-2316	4.7	5
139	Tuning stationary iterative solvers for fault resilience **2015**,		5
138	A dynamic pipeline for RNA sequencing on multicore processors **2013**,		5
137	Power Consumption of Mixed Precision in the Iterative Solution of Sparse Linear Systems **2011**,		5
136	GLT: A Unified API for Lightweight Thread Libraries. *Lecture Notes in Computer Science*, **2017**, 470-481	0.9	5
135	Out-of-Core Computation of the QR Factorization on Multi-core Processors. *Lecture Notes in Computer Science*, **2009**, 809-820	0.9	5
134	Variable-size batched GaussJordan elimination for block-Jacobi preconditioning on graphics processors. *Parallel Computing*, **2019**, 81, 131-146	1	5
133	Programming parallel dense matrix factorizations with look-ahead and OpenMP. *Cluster Computing*, **2020**, 23, 359-375	2.1	5
132	Communication in task-parallel ILU-preconditioned CG solvers using MPI+OmpSs. *Concurrency Computation Practice and Experience*, **2017**, 29, e4280	1.4	4
131	Fast and Reliable Noise Estimation for Hyperspectral Subspace Identification. *IEEE Geoscience and Remote Sensing Letters*, **2015**, 12, 1199-1203	4.1	4
130	Energy balance between voltage-frequency scaling and resilience for linear algebra routines on low-power multicore architectures. *Parallel Computing*, **2018**, 73, 28-39	1	4
129	Deriving dense linear algebra libraries. *Formal Aspects of Computing*, **2013**, 25, 933-945	1.2	4
128	Modeling power consumption of 3D MPDATA and the CG method on ARM and Intel multicore architectures. *Journal of Supercomputing*, **2017**, 73, 4373-4389	2.5	4

127	Automatic detection of power bottlenecks in parallel scientific applications. *Computer Science - Research and Development*, **2014**, 29, 221-229		4
126	Applying OOC Techniques in the Reduction to Condensed Form for Very Large Symmetric Eigenproblems on GPUs **2012**,		4
125	Analysis of Strategies to Save Energy for Message-Passing Dense Linear Algebra Kernels **2012**,		4
124	Retargeting PLAPACK to clusters with hardware accelerators **2010**,		4
123	Evaluation of the Energy Performance of Dense Linear Algebra Kernels on Multi-core and Many-Core Processors **2011**,		4
122	Parallel solution of Riccati matrix equations with the matrix sign function. *Automatica*, **1998**, 34, 151-156	5.7	4
121	**2020**,		4
120	Parallel Solution of Hierarchical Symmetric Positive Definite Linear Systems. *Applied Mathematics and Nonlinear Sciences*, **2017**, 2, 201-212	4	4
119	DMRlib: Easy-Coding and Efficient Resource Management for Job Malleability. *IEEE Transactions on Computers*, **2021**, 70, 1443-1457	2.5	4
118	Solving Stable Stein Equations on Distributed Memory Computers?. *Lecture Notes in Computer Science*, **1999**, 1120-1123	0.9	4
117	Accelerating multi-channel filtering of audio signal on ARM processors. *Journal of Supercomputing*, **2017**, 73, 203-214	2.5	3
116	Exploiting nested task-parallelism in the H-LU factorization. *Journal of Computational Science*, **2019**, 33, 20-33	3.4	3
115	Are our dense linear algebra libraries energy-friendly?. *Computer Science - Research and Development*, **2015**, 30, 187-196		3
114	Reducing the cost of power monitoring with DC wattmeters. *Computer Science - Research and Development*, **2015**, 30, 107-114		3
113	Optimized Fundamental Signal Processing Operations For Energy Minimization on Heterogeneous Mobile Devices. *IEEE Transactions on Circuits and Systems I: Regular Papers*, **2018**, 65, 1614-1627	3.9	3
112	Multi-threaded dense linear algebra libraries for low-power asymmetric multicore processors. *Journal of Computational Science*, **2018**, 25, 140-151	3.4	3
111	A fast band-Krylov eigensolver for macromolecular functional motion simulation on multicore architectures and graphics processors. *Journal of Computational Physics*, **2016**, 309, 314-323	4.1	3
110	FaST-LMM for Two-Way Epistasis Tests on High-Performance Clusters. *Journal of Computational Biology*, **2018**, 25, 862-870	1.7	3

109	Variable-Size Batched LU for Small Matrices and Its Integration into Block-Jacobi Preconditioning **2017**,		3
108	Variable-Size Batched Gauss-Huard for Block-Jacobi Preconditioning. *Procedia Computer Science*, **2017**, 108, 1783-1792	1.6	3
107	**2017**,		3
106	Vectorization of binaural sound virtualization on the ARM Cortex-A15 architecture **2015**,		3
105	Exploring the Suitability of Remote GPGPU Virtualization for the OpenACC Programming Model Using rCUDA **2015**,		3
104	Modeling power and energy consumption of dense matrix factorizations on multicore processors. *Concurrency Computation Practice and Experience*, **2014**, 26, 2743-2757	1.4	3
103	Parallel Algorithms for Computing Rank-Revealing QR Factorizations. *Lecture Notes in Control and Information Sciences*, **1997**, 122-137	0.5	3
102	Design, Tuning and Evaluation of Parallel Multilevel ILU Preconditioners. *Lecture Notes in Computer Science*, **2008**, 314-327	0.9	3
101	Accelerating Band Linear Algebra Operations on GPUs with Application in Model Reduction. *Lecture Notes in Computer Science*, **2014**, 386-400	0.9	3
100	Revisiting the Gauss-Huard Algorithm for the Solution of Linear Systems on Graphics Accelerators. *Lecture Notes in Computer Science*, **2016**, 505-514	0.9	3
99	Attaining High Performance in General-Purpose Computations on Current Graphics Processors. *Lecture Notes in Computer Science*, **2008**, 406-419	0.9	3
98	Accelerating Model Reduction of Large Linear Systems with Graphics Processors. *Lecture Notes in Computer Science*, **2012**, 88-97	0.9	3
97	Leveraging Task-Parallelism in Energy-Efficient ILU Preconditioners. *Lecture Notes in Computer Science*, **2012**, 55-63	0.9	3
96	Integration and exploitation of intra-routine malleability in BLIS. *Journal of Supercomputing*, **2020**, 76, 2860-2875	2.5	3
95	Adaptive Precision Block-Jacobi for High Performance Preconditioning in the Ginkgo Linear Algebra Software. *ACM Transactions on Mathematical Software*, **2021**, 47, 1-28	2.3	3
94	High-Performance GPU Implementation of PageRank with Reduced Precision Based on Mantissa Segmentation **2018**,		3
93	Adapting concurrency throttling and voltageErequency scaling for dense eigensolvers. *Journal of Supercomputing*, **2017**, 73, 29-43	2.5	2
92	Time and energy modeling of a high-performance multi-threaded Cholesky factorization. *Journal of Supercomputing*, **2017**, 73, 139-151	2.5	2

91	Solving Weighted Least Squares (WLS) problems on ARM-based architectures. *Journal of Supercomputing*, **2017**, 73, 530-542	2.5	2
90	Dynamic look-ahead in the reduction to band form for the singular value decomposition. *Parallel Computing*, **2019**, 81, 22-31	1	2
89	Hierarchical approach for deriving a reproducible unblocked LU factorization. *International Journal of High Performance Computing Applications*, **2019**, 33, 791-803	1.8	2
88	Accelerating the task/data-parallel version of ILUPACK's BiCG in multi-CPU/GPU configurations. *Parallel Computing*, **2019**, 85, 79-87	1	2
87	Real-time Sound Source Localization on an Embedded GPU Using a Spherical Microphone Array. *Procedia Computer Science*, **2015**, 51, 201-210	1.6	2
86	Time and energy modeling of high-performance Level-3 BLAS on x86 architectures. *Simulation Modelling Practice and Theory*, **2015**, 55, 77-94	3.9	2
85	Systematic derivation of time and power models for linear algebra kernels on multicore architectures. *Sustainable Computing: Informatics and Systems*, **2015**, 7, 24-40	3	2
84	Scalable RNA Sequencing on Clusters of Multicore Processors **2015**,		2
83	Binding Performance and Power of Dense Linear Algebra Operations **2012**,		2
82	A simulator to assess energy saving strategies and policies in HPC workloads. *Operating Systems Review (ACM)*, **2012**, 46, 2-9	0.8	2
81	Acceleration of PageRank with Customized Precision Based on Mantissa Segmentation. *ACM Transactions on Parallel Computing*, **2020**, 7, 1-19	1.4	2
80	Solving Linear Systems on the Intel Xeon-Phi Accelerator via the Gauss-Huard Algorithm. *Communications in Computer and Information Science*, **2015**, 107-117	0.3	2
79	Exploiting Task-Parallelism in Message-Passing Sparse Linear System Solvers Using OmpSs. *Lecture Notes in Computer Science*, **2016**, 631-643	0.9	2
78	Accelerating FaST-LMM for Epistasis Tests. *Lecture Notes in Computer Science*, **2017**, 548-557	0.9	2
77	An Algorithm-by-Blocks for SuperMatrix Band Cholesky Factorization. *Lecture Notes in Computer Science*, **2008**, 228-239	0.9	2
76	Refactoring Conventional Task Schedulers to Exploit Asymmetric ARM big.LITTLE Architectures in Dense Linear Algebra **2016**,		2
75	An efficient GPU version of the preconditioned GMRES method. *Journal of Supercomputing*, **2019**, 75, 1455-1469	2.5	2
74	The libflame Library for Dense Matrix Computations. *Computing in Science and Engineering*, **2019**, 1-1	1.5	2

73	Fine-grained bit-flip protection for relaxation methods. *Journal of Computational Science*, **2019**, 36, 100583	3.1	2
72	Exploring the interoperability of remote GPGPU virtualization using rCUDA and directive-based programming models. *Journal of Supercomputing*, **2018**, 74, 5628-5642	2.5	2
71	Ginkgo : A Modern Linear Operator Algebra Framework for High Performance Computing. *ACM Transactions on Mathematical Software*, **2022**, 48, 1-33	2.3	2
70	A BLIS-like matrix multiplication for machine learning in the RISC-V ISA-based GAP8 processor. *Journal of Supercomputing*,	2.5	2
69	Extending the Gauss-Huard method for the solution of Lyapunov matrix equations and matrix inversion. *Concurrency Computation Practice and Experience*, **2017**, 29, e4076	1.4	1
68	Evaluating the Potential of Low Power Systems for Headphone-based Spatial Audio Applications. *Procedia Computer Science*, **2015**, 51, 191-200	1.6	1
67	Extending lyapack for the solution of band Lyapunov equations on hybrid CPU-GPU platforms. *Journal of Supercomputing*, **2015**, 71, 740-750	2.5	1
66	A framework for genomic sequencing on clusters of multicore and manycore processors. *International Journal of High Performance Computing Applications*, **2018**, 32, 393-406	1.8	1
65	Evaluating fault tolerance on asymmetric multicore systems-on-chip using iso-metrics. *IET Computers and Digital Techniques*, **2016**, 10, 85-92	0.9	1
64	Performance and Energy Analysis of the Iterative Solution of Sparse Linear Systems on Multicore and Manycore Architectures. *Lecture Notes in Computer Science*, **2014**, 772-782	0.9	1
63	On the Use of a GPU-Accelerated Mobile Device Processor for Sound Source Localization. *Procedia Computer Science*, **2017**, 108, 586-595	1.6	1
62	Reduction to Tridiagonal Form for Symmetric Eigenproblems on Asymmetric Multicore Processors **2017**,		1
61	Exploiting Task-Parallelism on GPU Clusters via OmpSs and rCUDA Virtualization **2015**,		1
60	Solving dense linear systems with hybrid ARM+GPU platforms **2015**,		1
59	Concurrent and Accurate Short Read Mapping on Multicore Processors. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, **2015**, 12, 995-1007	3	1
58	Adaptive Downtime for Live Migration of Virtual Machines **2014**,		1
57	Enhancing performance and energy consumption of runtime schedulers for dense linear algebra. *Concurrency Computation Practice and Experience*, **2014**, 26, 2591-2611	1.4	1
56	Leveraging task-parallelism in message-passing dense matrix factorizations using SMPSs. *Parallel Computing*, **2014**, 40, 113-128	1	1

55	A factored variant of the Newton iteration for the solution of algebraic Riccati equations via the matrix sign function. *Numerical Algorithms*, **2014**, 66, 363-377	2.1	1
54	High Performance Implementations of the BST Method on Hybrid CPU-GPU Platforms **2012**,		1
53	Increasing data locality and introducing Level-3 BLAS in the Neville elimination. *Applied Mathematics and Computation*, **2011**, 218, 3348-3358	2.7	1
52	Solving large dense matrix problems on multi-core processors **2009**,		1
51	Out-of-core solution of linear systems on graphics processors. *International Journal of Parallel, Emergent and Distributed Systems*, **2009**, 24, 521-538	1	1
50	Using Graphics Processors to Accelerate the Solution of Out-of-Core Linear Systems **2009**,		1
49	Fast development of dense linear algebra codes on graphics processors **2009**,		1
48	Design of scalable dense linear algebra libraries for multithreaded architectures: the LU factorization. *Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on*, **2008**,		1
47	Specialized Spectral Division Algorithms for Generalized Eigenproblems Via the Inverse-Free Iteration **2007**, 157-166		1
46	Rapid Development of High-Performance Linear Algebra Libraries. *Lecture Notes in Computer Science*, **2006**, 376-384	0.9	1
45	Multiprecision Block-Jacobi for Iterative Triangular Solves. *Lecture Notes in Computer Science*, **2020**, 546-560	0.9	1
44	The Impact of Voltage-Frequency Scaling for the Matrix-Vector Product on the IBM POWER8. *Lecture Notes in Computer Science*, **2016**, 103-116	0.9	1
43	Design of a Task-Parallel Version of ILUPACK for Graphics Processors. *Communications in Computer and Information Science*, **2017**, 91-103	0.3	1
42	Solving Sparse Differential Riccati Equations on Hybrid CPU-GPU Platforms. *Lecture Notes in Computer Science*, **2017**, 116-132	0.9	1
41	Parallelizing Dense Linear Algebra Operations with Task Queues in llc. *Lecture Notes in Computer Science*, **2007**, 89-96	0.9	1
40	Runtime Scheduling of the LU Factorization: Performance and Energy. *Lecture Notes in Computer Science*, **2013**, 153-167	0.9	1
39	Efficient Symmetric Band Matrix-Matrix Multiplication on GPUs. *Communications in Computer and Information Science*, **2014**, 1-12	0.3	1
38	A Data-Parallel ILUPACK for Sparse General and Symmetric Indefinite Linear Systems. *Lecture Notes in Computer Science*, **2017**, 121-133	0.9	1

37	Accelerating BST Methods for Model Reduction with Graphics Processors. *Lecture Notes in Computer Science*, **2012**, 549-558	0.9	1
36	Unleashing CPU-GPU Acceleration for Control Theory Applications. *Lecture Notes in Computer Science*, **2013**, 102-111	0.9	1
35	Fast block QR update in digital signal processing. *Journal of Supercomputing*, **2019**, 75, 1051-1064	2.5	1
34	Factorized solution of generalized stable Sylvester equations using many-core GPU accelerators. *Journal of Supercomputing*, **2021**, 77, 10152-10164	2.5	1
33	Resiliency in numerical algorithm design for extreme scale simulations. *International Journal of High Performance Computing Applications*, **2022**, 36, 251-285	1.8	1
32	Reproducibility of parallel preconditioned conjugate gradient in hybrid programming environments. *International Journal of High Performance Computing Applications*, **2020**, 34, 502-518	1.8	0
31	Revisiting conventional task schedulers to exploit asymmetry in multi-core architectures for dense linear algebra operations. *Parallel Computing*, **2017**, 68, 59-76	1	0
30	Out-of-core macromolecular simulations on multithreaded architectures. *Concurrency Computation Practice and Experience*, **2015**, 27, 1540-1550	1.4	0
29	Parallel Implementation of LQG Balanced Truncation for Large-Scale Systems. *Lecture Notes in Computer Science*, **2008**, 227-234	0.9	0
28	Cholesky and Gram-Schmidt Orthogonalization for Tall-and-Skinny QR Factorizations on Graphics Processors. *Lecture Notes in Computer Science*, **2019**, 469-480	0.9	0
27	Compression and load balancing for efficient sparse matrix-vector product on multicore processors and graphics processing units. *Concurrency Computation Practice and Experience*, e6515	1.4	0
26	Look-ahead in the two-sided reduction to compact band forms for symmetric eigenvalue problems and the SVD. *Numerical Algorithms*, **2019**, 80, 635-660	2.1	0
25	Low precision matrix multiplication for efficient deep learning in NVIDIA Carmel processors. *Journal of Supercomputing*, **2021**, 77, 11257-11269	2.5	0
24	High performance and energy efficient inference for deep learning on multicore ARM processors using general optimization techniques and BLIS. *Journal of Systems Architecture*, **2022**, 125, 102459	5.5	0
23	Architecture-aware optimization of an HEVC decoder on asymmetric multicore processors. *Journal of Real-Time Image Processing*, **2017**, 13, 25-38	1.9	
22	Unleashing GPU acceleration for symmetric band linear algebra kernels and model reduction. *Cluster Computing*, **2015**, 18, 1351-1362	2.1	
21	Tall-and-skinny QR factorization with approximate Householder reflectors on graphics processors. *Journal of Supercomputing*, **2020**, 76, 8771-8786	2.5	
20	Two-sided orthogonal reductions to condensed forms on asymmetric multicore processors. *Parallel Computing*, **2018**, 78, 85-100	1	

19 · Evaluating the NVIDIA Tegra Processor as a Low-Power Alternative for Sparse GPU Computations. *Communications in Computer and Information Science*, **2018**, 111-122 · 0.3

18 · Static scheduling of the LU factorization with look-ahead on asymmetric multicore processors. *Parallel Computing*, **2018**, 76, 18-27 · 1

17 · Characterizing the efficiency of multicore and manycore processors for the solution of sparse linear systems. *Computer Science - Research and Development*, **2016**, 31, 175-183

16 · Noise estimation for hyperspectral subspace identification on FPGAs. *Journal of Supercomputing*, **2019**, 75, 1323-1335 · 2.5

15 · Solution of Few-Body Coulomb Problems with Latent Matrices on Multicore Processors. *Procedia Computer Science*, **2017**, 108, 1743-1752 · 1.6

14 · The Implementation of BLAS for Band Matrices **2007**, 668-677

13 · Stabilizing large-scale generalized systems on parallel computers using multithreading and message-passing. *Concurrency Computation Practice and Experience*, **2007**, 19, 531-542 · 1.4

12 · Parallel solvers for discrete-time algebric Riccati equations. *Concurrency Computation Practice and Experience*, **2001**, 13, 153-162 · 1.4

11 · Efficient Algorithms for the Block Hessenberg Form. *Journal of Supercomputing*, **2001**, 20, 55-66 · 2.5

10 · Parallel Solution of Band Linear Systems in Model Reduction **2008**, 678-687

9 · Structure-Aware Calculation of Many-Electron Wave Function Overlaps on Multicore Processors. *Lecture Notes in Computer Science*, **2020**, 13-24 · 0.9

8 · Stabilizing large control linear systems on multicomputers. *Lecture Notes in Computer Science*, **1997**, 338-364 · 0.3

7 · A Parallel Multi-threaded Solver for Symmetric Positive Definite Bordered-Band Linear Systems. *Lecture Notes in Computer Science*, **2016**, 96-105 · 0.9

6 · The Impact of Panel Factorization on the Gauss-Huard Algorithm for the Solution of Linear Systems on Modern Architectures. *Lecture Notes in Computer Science*, **2016**, 405-416 · 0.9

5 · Tuning the Blocksize for Dense Linear Algebra Factorization Routines with the Roofline Model. *Lecture Notes in Computer Science*, **2016**, 18-29 · 0.9

4 · Out-of-Core Solution of Eigenproblems for Macromolecular Simulations. *Lecture Notes in Computer Science*, **2014**, 490-499 · 0.9

3 · Balanced and Compressed Coordinate Layout for the Sparse Matrix-Vector Product on GPUs. *Lecture Notes in Computer Science*, **2021**, 83-95 · 0.9

2 · On the performance of a GPU-based SoC in a distributed spatial audio system. *Journal of Supercomputing*, **2021**, 77, 6920-6935 · 2.5

List of Publications

1 Approximate Computing for Scientific Applications **2022**, 415-465