JosÃ© Duato

List of Publications by Year
in descending order

238	4,602	147801 31	168389 53
papers	citations	h-index	g-index

240	240	240	1409
all docs	docs citations	times ranked	citing authors

#	Article	IF	Citations
1	Analyzing the impact of the MPI allreduce in distributed training of convolutional neural networks. Computing (Vienna/New York), 2023, 105, 1101-1119.	4.8	2
2	Enforcing Predictability of Many-Cores With DCFNoC. IEEE Transactions on Computers, 2021, 70, 270-283.	3.4	5
3	DVL-Lossy: Isolating Congesting Flows to Optimize Packet Dropping in Lossy Data-Center Networks. IEEE Micro, 2021, 41, 37-44.	1.8	1
4	Performance Modeling for Distributed Training of Convolutional Neural Networks. , 2021, , .		3
5	Evaluation of MPI Allreduce for Distributed Training of Convolutional Neural Networks. , 2021, , .		2
6	UPR: deadlock-free dynamic network reconfiguration by exploiting channel dependency graph compatibility. Journal of Supercomputing, 2021, 77, 12826-12856.	3.6	1
7	Accelerating distributed deep neural network training with pipelined MPI allreduce. Cluster Computing, 2021, 24, 3797-3813.	5.0	8
8	Path2SL: Leveraging InfiniBand Resources to Reduce Head-of-Line Blocking in Fat Trees. IEEE Micro, 2020, 40, 8-14.	1.8	0
9	Optimizing Packet Dropping by Efficient Congesting-Flow Isolation in Lossy Data-Center Networks. , 2020, , .		3
10	HP-DCFNoC: High Performance Distributed Dynamic TDM Scheduler Based on DCFNoC Theory. IEEE Access, 2020, 8, 194836-194849.	4.2	1
11	Theoretical Scalability Analysis of Distributed Deep Convolutional Neural Networks. , 2019, , .		10
12	Efficient Dynamic Isolation of Congestion in Lossless DataCenter Networks. , 2019, , .		2
13	DCFNoC. , 2019, , .		10
14	Analysis of model parallelism for distributed neural networks. , 2019, , .		8
15	Combining Source-adaptive and Oblivious Routing with Congestion Control in High-performance Interconnects using Hybrid and Direct Topologies. Transactions on Architecture and Code Optimization, 2019, 16, 1-26.	2.0	0
16	Path2SL: Optimizing Head-of-Line Blocking Reduction in InfiniBand-Based Fat-Tree Networks. , 2019, , .		2
17	Modeling Traffic Workloads in Data-center Network Simulation Tools. , 2019, , .		1
18	Constructing virtual 5â€dimensional tori out of lowerâ€dimensional network cards. Concurrency Computation Practice and Experience, 2019, 31, e4361.	2.2	1

#	Article	IF	Citations
19	Accurately modeling the on-chip and off-chip GPU memory subsystem. Future Generation Computer Systems, 2018, 82, 510-519.	7.5	12
20	Feasible enhancements to congestion control in InfiniBand-based networks. Journal of Parallel and Distributed Computing, 2018, 112, 35-52.	4.1	4
21	TLB-Based Temporality-Aware Classification in CMPs with Multilevel TLBs. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 2401-2413.	5.6	7
22	A Case Study on Implementing Virtual 5D Torus Networks Using Network Components of Lower Dimensionality. , 2017, , .		1
23	Perf&Fair: A Progress-Aware Scheduler to Enhance Performance and Fairness in SMT Multicores. IEEE Transactions on Computers, 2017, 66, 905-911.	3.4	20
24	The k-ary n-direct s-indirect family of topologies for large-scale interconnection networks. Journal of Supercomputing, 2016, 72, 1035-1062.	3.6	11
25	TokenTLB. , 2016, , .		6
26	Adaptive routing for n-Dimensional Twin torus. IEEE Transactions on Computers, 2016, , 1-1.	3.4	1
27	Impact of Memory-Level Parallelism on the Performance of GPU Coherence Protocols. , 2016, , .		0
28	Bandwidth-Aware On-Line Scheduling in SMT Multicores. IEEE Transactions on Computers, 2016, 65, 422-434.	3.4	14
29	A dynamic execution time estimation model to save energy in heterogeneous multicores running periodic tasks. Future Generation Computer Systems, 2016, 56, 211-219.	7.5	8
30	A Family of Fault-Tolerant Efficient Indirect Topologies. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 927-940.	5.6	8
31	Efficient TLB-Based Detection of Private Pages in Chip Multiprocessors. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 748-761.	5.6	16
32	A reuse-based refresh policy for energy-aware eDRAM caches. Microprocessors and Microsystems, 2015, 39, 37-48.	2.8	5
33	A parallel and sensitive software tool for methylation analysis on multicore platforms. Bioinformatics, 2015, 31, 3130-3138.	4.1	9
34	On the design of a new dynamic credit-based end-to-end flow control mechanism for HPC clusters. Parallel Computing, 2015, 46, 32-59.	2.1	1
35	Optimizing the configuration of combined high-radix switches. Journal of Supercomputing, 2015, 71, 2614-2643.	3.6	0
36	Addressing Fairness in SMT Multicores with a Progress-Aware Scheduler. , 2015, , .		12

#	Article	IF	Citations
37	Accurately modeling the GPU memory subsystem. , 2015, , .		3
38	N-Dimensional Twin Torus Topology. IEEE Transactions on Computers, 2015, 64, 2847-2861.	3.4	16
39	A HoL-blocking aware mechanism for selecting the upward path in fat-tree topologies. Journal of Supercomputing, 2015, 71, 2339-2364.	3.6	2
40	Improving the user experience of the rCUDA remote GPU virtualization framework. Concurrency Computation Practice and Experience, 2015, 27, 3746-3770.	2.2	12
41	Efficient and Cost-Effective Hybrid Congestion Control for HPC Interconnection Networks. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 107-119.	5.6	22
42	Design of Hybrid Second-Level Caches. IEEE Transactions on Computers, 2015, 64, 1884-1897.	3.4	13
43	Design of an ICT Tool for Decision Making in Social and Health Policies. , 2015, , 997-1014.		0
44	Achieving balanced buffer utilization with a proper co-design of flow control and routing algorithm. , 2014, , .		10
45	SLURM Support for Remote GPU Virtualization: Implementation and Performance Study. , 2014, , .		19
46	Addressing bandwidth contention in SMT multicores through scheduling. , 2014, , .		2
47	Deadlock-free routing mechanism for 3D twin torus networks. , 2014, , .		0
48	Dynamic WCET Estimation for Real-Time Multicore Embedded Systems Supporting DVFS. , 2014, , .		2
49	Optimal Configuration for N-Dimensional Twin Torus Networks. , 2014, , .		0
50	Combining HoL-blocking avoidance and differentiated services in high-speed interconnects. , 2014, , .		4
51	Building 3D Torus Using Low-Profile Expansion Cards. IEEE Transactions on Computers, 2014, 63, 2701-2715.	3.4	10
52	A complete and efficient CUDA-sharing solution for HPC clusters. Parallel Computing, 2014, 40, 574-588.	2.1	64
53	Formalization and configuration methodology for high-radix combined switches. Journal of Supercomputing, 2014, 69, 1410-1444.	3.6	2
54	Cache-Hierarchy Contention-Aware Scheduling in CMPs. IEEE Transactions on Parallel and Distributed Systems, 2014, 25, 581-590.	5.6	21

#	Article	IF	Citations
55	A new proposal to deal with congestion in InfiniBand-based fat-trees. Journal of Parallel and Distributed Computing, 2014, 74, 1802-1819.	4.1	36
56	Powerâ€aware scheduling with effective task migration for realâ€time multicore embedded systems. Concurrency Computation Practice and Experience, 2013, 25, 1987-2001.	2.2	19
57	Silicon-aware distributed switch architecture for on-chip networks. Journal of Systems Architecture, 2013, 59, 505-515.	4.3	8
58	Obtaining the optimal configuration of high-radix Combined switches. Journal of Parallel and Distributed Computing, 2013, 73, 1239-1250.	4.1	8
59	Deterministic Routing with HoL-Blocking-Awareness for Direct Topologies. Procedia Computer Science, 2013, 18, 2521-2524.	2.0	1
60	Hardware-based generation of independent subtraces of instructions in clustered processors. IEEE Transactions on Computers, 2013, 62, 944-955.	3.4	0
61	Temporal-Aware Mechanism to Detect Private Data in Chip Multiprocessors. , 2013, , .		15
62	Using Huge Pages and Performance Counters to Determine the LLC Architecture. Procedia Computer Science, 2013, 18, 2557-2560.	2.0	1
63	Combining RAM Technologies for Hard-error Recovery in L1 Data Caches Working at Very-low Power Modes. , 2013, , .		3
64	An empirical model for predicting cross-core performance interference on multicore processors. , 2013, , .		2
65	An Effective and Feasible Congestion Management Technique for High-Performance MINs with Tag-Based Distributed Routing. IEEE Transactions on Parallel and Distributed Systems, 2013, 24, 1918-1929.	5.6	19
66	Exploiting reuse information to reduce refresh energy in on-chip eDRAM caches. , 2013, , .		2
67	Increasing the Effectiveness of Directory Caches by Avoiding the Tracking of Noncoherent Memory Blocks. IEEE Transactions on Computers, 2013, 62, 482-495.	3.4	25
68	BBQ: A Straightforward Queuing Scheme to Reduce HoL-Blocking in High-Performance Hybrid Networks. Lecture Notes in Computer Science, 2013, , 699-712.	1.3	16
69	Design of an ICT Tool for Decision Making in Social and Health Policies. , 2013, , 802-819.		2
70	Addressing Link Degradation in NoC-Based ULSI Designs. Lecture Notes in Computer Science, 2013, , 327-336.	1.3	0
71	Combining recency of information with selective random and a victim cache in last-level caches. Transactions on Architecture and Code Optimization, 2012, 9, 1-20.	2.0	8
72	PS-Dir. , 2012, , .		10

#	Article	IF	Citations
73	A New Family of Hybrid Topologies for Large-Scale Interconnection Networks. , 2012, , .		8
74	A New End-to-End Flow-Control Mechanism for High Performance Computing Clusters. , 2012, , .		2
75	Optimal Configuration of High-Radix Combined Switches. , 2012, , .		1
76	A Survey and Evaluation of Topology-Agnostic Deterministic Routing Algorithms. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 405-425.	5.6	89
77	Enabling High-Performance Crossbars through a Floorplan-Aware Design. , 2012, , .		6
78	Cache Miss Characterization in Hierarchical Large-Scale Cache-Coherent Systems. , 2012, , .		4
79	Page-Based Memory Allocation Policies of Local and Remote Memory in Cluster Computers. , 2012, , .		1
80	Understanding Cache Hierarchy Contention in CMPs to Improve Job Scheduling. , 2012, , .		16
81	Efficiently Handling Memory Accesses to Improve QoS in Multicore Systems under Real-Time Constraints. , 2012, , .		0
82	Analyzing the optimal ratio of SRAM banks in hybrid caches. , 2012, , .		3
83	Impact on Performance and Energy of the Retention Time and Processor Frequency in L1 Macrocell-Based Data Caches. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2012, 20, 1108-1117.	3.1	4
84	A new degree of freedom for memory allocation in clusters. Cluster Computing, 2012, 15, 101-123.	5.0	6
85	Switch-based packing technique to reduce traffic and latency in token coherence. Journal of Parallel and Distributed Computing, 2012, 72, 409-423.	4.1	0
86	Design, Performance, and Energy Consumption of eDRAM/SRAM Macrocells for L1 Data Caches. IEEE Transactions on Computers, 2012, 61, 1231-1242.	3.4	11
87	Progressive Congestion Management Based on Packet Marking and Validation Techniques. IEEE Transactions on Computers, 2012, 61, 1296-1310.	3.4	4
88	Extending Magny-Cours Cache Coherence. IEEE Transactions on Computers, 2012, 61, 593-606.	3.4	2
89	On the Impact of Within-Die Process Variation in GALS-Based NoC Performance. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2012, 31, 294-307.	2.7	4
90	A cost-effective heuristic to schedule local and remote memory in cluster computers. Journal of Supercomputing, 2012, 59, 1533-1551.	3.6	4

#	Article	IF	Citations
91	Exploiting SIMD Instructions in Current Processors to Improve Classical String Algorithms. Lecture Notes in Computer Science, 2012, , 254-267.	1.3	8
92	Towards an Efficient NoC Topology through Multiple Injection Ports. , 2011, , .		3
93	Enabling CUDA acceleration within virtual machines using rCUDA. , 2011, , .		55
94	A Distributed Switch Architecture for On-Chip Networks. , 2011, , .		2
95	MRU-Tour-based Replacement Algorithms for Last-Level Caches. , 2011, , .		1
96	Evaluation of an Alternative for Increasing Switch Radix. , 2011, , .		3
97	Performance of CUDA Virtualized Remote GPUs in High Performance Clusters. , 2011, , .		27
98	Highly scalable barriers for future high-performance computing clusters. , 2011, , .		2
99	Improving Last-Level Cache Performance by Exploiting the Concept of MRU-Tour. , 2011, , .		1
100	MEMSCALE™: A Scalable Environment for Databases. , 2011, , .		7
101	Combining Congested-Flow Isolation and Injection Throttling in HPC Interconnection Networks. , 2011, , .		11
102	Efficient and Scalable Starvation Prevention Mechanism for Token Coherence. IEEE Transactions on Parallel and Distributed Systems, 2011, 22, 1610-1623.	5.6	2
103	Dynamic Fault Tolerance in Fat Trees. IEEE Transactions on Computers, 2011, 60, 508-525.	3.4	31
104	Cost-Efficient On-Chip Routing Implementations for CMP and MPSoC Systems. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2011, 30, 534-547.	2.7	44
105	Efficient routing implementation in complex systems-on-chip designs. , 2011, , .		4
106	A low-latency modular switch for CMP systems. Microprocessors and Microsystems, 2011, 35, 742-754.	2.8	11
107	OBQA: Smart and cost-efficient queue scheme for Head-of-Line blocking elimination in fat-trees. Journal of Parallel and Distributed Computing, 2011, 71, 1460-1472.	4.1	15
108	Fault-Tolerant Vertical Link Design for Effective 3D Stacking. IEEE Computer Architecture Letters, 2011, 10, 41-44.	1.5	7

#	Article	IF	Citations
109	A Communication-Driven Routing Technique for Application-Specific NoCs. International Journal of Parallel Programming, 2011, 39, 357-374.	1.5	3
110	How to reduce packet dropping in a bufferless NoC. Concurrency Computation Practice and Experience, 2011, 23, 86-99.	2.2	4
111	Characterizing the impact of process variation on 45 nm NoC-based CMPs. Journal of Parallel and Distributed Computing, 2011, 71, 651-663.	4.1	7
112	MEMSCALE. , 2011, , .		3
113	C-Switches: Increasing Switch Radix with Current Integration Scale. , 2011, , .		5
114	A New Energy-Aware Dynamic Task Set Partitioning Algorithm for Soft and Hard Embedded Real-Time Systems. Computer Journal, 2011, 54, 1282-1294.	2.4	11
115	Unleash Your Memory-Constrained Applications: A 32-Node Non-coherent Distributed-Memory Prototype Cluster. , 2011, , .		2
116	A Dynamic Power-Aware Partitioner with Task Migration for Multicore Embedded Systems. Lecture Notes in Computer Science, 2011, , 218-229.	1.3	5
117	A Cluster Computer Performance Predictor for Memory Scheduling. Lecture Notes in Computer Science, 2011, , 353-362.	1.3	1
118	Ensuring the performance and scalability of peer-to-peer distributed virtual environments. Future Generation Computer Systems, 2010, 26, 905-915.	7.5	8
119	Power saving in regular interconnection networks. Parallel Computing, 2010, 36, 696-712.	2.1	15
120	A methodology for the characterization of process variation in NoC links. , 2010, , .		19
121	Dynamic task set partitioning based on balancing resource requirements and utilization to reduce power consumption. , 2010, , .		0
122	Scalable hardware support for conditional parallelization. , 2010, , .		2
123	Exploiting subtrace-level parallelism in clustered processors. , 2010, , .		0
124	EMC²: Extending Magny-Cours coherence for large-scale servers. , 2010, , .		9
125	VCT<inf>lite</inf>: Towards an efficient implementation of virtual cut-through switching in on-chip networks. , 2010, , .		4
126	Buffer Management Strategies to Reduce HoL Blocking. IEEE Transactions on Parallel and Distributed Systems, 2010, 21, 739-753.	5.6	39

#	Article	IF	Citations
127	A practical way to extend shared memory support beyond a motherboard at low cost. , 2010, , .		6
128	An Efficient Implementation of GPU Virtualization in High Performance Clusters. Lecture Notes in Computer Science, 2010, , 385-394.	1.3	22
129	Getting Rid of Coherency Overhead for Memory-Hungry Applications. , 2010, , .		5
130	A Scalable and Early Congestion Management Mechanism for MINs. , 2010, , .		7
131	A Scheduling Heuristic to Handle Local and Remote Memory in Cluster Computers. , 2010, , .		2
132	rCUDA: Reducing the number of GPU-based accelerators in high performance clusters. , 2010, , .		179
133	Balancing Task Resource Requirements in Embedded Multithreaded Multicore Processors to Reduce Power Consumption. , 2010, , .		1
134	Dealing with Transient Faults in the Interconnection Network of CMPs at the Cache Coherence Level. IEEE Transactions on Parallel and Distributed Systems, 2010, 21, 1117-1131.	5.6	3
135	A Latency-Efficient Router Architecture for CMP Systems. , 2010, , .		6
136	An Efficient Strategy for Reducing Head-of-Line Blocking in Fat-Trees. Lecture Notes in Computer Science, 2010, , 413-427.	1.3	11
137	A Switch Architecture Guaranteeing QoS Provision and HOL Blocking Elimination. IEEE Transactions on Parallel and Distributed Systems, 2009, 20, 13-24.	5.6	5
138	A new mechanism to deal with process variability in NoC links. , 2009, , .		9
139	Region-Based Routing: A Mechanism to Support Efficient Routing Algorithms in NoCs. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2009, 17, 356-369.	3.1	50
140	A performance evaluation of 2D-mesh, ring, and crossbar interconnects for chip multi-processors. , 2009, , .		7
141	Dynamic task set partitioning based on balancing memory requirements to reduce power consumption. , 2009, , .		0
142	A new strategy to manage the InfiniBand arbitration tables. Journal of Parallel and Distributed Computing, 2009, 69, 508-520.	4.1	0
143	Efficient and Scalable Hardware-Based Multicast in Fat-Tree Networks. IEEE Transactions on Parallel and Distributed Systems, 2009, 20, 1285-1298.	5.6	9
144	M-GRASP: A GRASP With Memory for Latency-Aware Partitioning Methods in DVE Systems. IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, 2009, 39, 1214-1223.	2.9	4

#	Article	IF	Citations
145	Efficient implementation of distributed routing algorithms for NoCs. IET Computers and Digital Techniques, 2009, 3, 460.	1.2	39
146	An Efficient Low-Complexity Alternative to the ROB for Out-of-Order Retirement of Instructions. , 2009, , .		1
147	An hybrid eDRAM/SRAM macrocell to implement first-level data caches. , 2009, , .		20
148	Dependability Analysis of a Fault-Tolerant Network Reconfiguring Strategy. Lecture Notes in Computer Science, 2009, , 1040-1051.	1.3	1
149	A proposal for managing ASI fabrics. Journal of Systems Architecture, 2008, 54, 664-678.	4.3	6
150	Logic-Based Distributed Routing for NoCs. IEEE Computer Architecture Letters, 2008, 7, 13-16.	1.5	73
151	Beyond Fat--tree: Unidirectional Load--Balanced Multistage Interconnection Network. IEEE Computer Architecture Letters, 2008, 7, 49-52.	1.5	16
152	Extending the TokenCMP Cache Coherence Protocol for Low Overhead Fault Tolerance in CMP Architectures. IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 1044-1056.	5.6	1
153	A simple power-aware scheduling for multicore systems when running real-time applications. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	39
154	An Efficient Implementation of Distributed Routing Algorithms for NoCs. , 2008, , .		54
155	Exploring High-Dimensional Topologies for NoC Design Through an Integrated Analysis and Synthesis Framework. , 2008, , .		25
156	On the Potentials of Segment-Based Routing for NoCs. , 2008, , .		25
157	An Efficient and Deadlock-Free Network Reconfiguration Protocol. IEEE Transactions on Computers, 2008, 57, 762-779.	3.4	50
158	Efficient Deadline-Based QoS Algorithms for High-Performance Networks. IEEE Transactions on Computers, 2008, 57, 928-939.	3.4	6
159	Reducing Packet Dropping in a Bufferless NoC. Lecture Notes in Computer Science, 2008, , 899-909.	1.3	26
160	A fault-tolerant directory-based cache coherence protocol for CMP architectures. , 2008, , .		6
161	High-radix crossbar switches enabled by Proximity Communication. , 2008, , .		9
162	Efficient unicast and multicast support for CMPs. , 2008, , .		79

#	Article	IF	Citations
163	Network Reconfiguration Suitability for Scientific Applications. , 2008, , .		1
164	Switch-Based Packing Technique for Improving Token Coherence Scalability. , 2008, , .		3
165	CART: Communication-Aware Routing Technique for Application-Specific NoCs. , 2008, , .		4
166	On the Potential of NoC Virtualization for Multicore Chips. , 2008, , .		23
167	On the Influence of the Packet Marking and Injection Control Schemes in Congestion Management for MINs. Lecture Notes in Computer Science, 2008, , 930-939.	1.3	10
168	A Communication-Aware Topological Mapping Technique for NoCs. Lecture Notes in Computer Science, 2008, , 910-919.	1.3	21
169	FBICM: Efficient Congestion Management for High-Performance Networks Using Distributed Deterministic Routing. Lecture Notes in Computer Science, 2008, , 503-517.	1.3	10
170	Fault-Tolerant Cache Coherence Protocols for CMPs: Evaluation and Trade-Offs. Lecture Notes in Computer Science, 2008, , 555-568.	1.3	1
171	A New Cost-Effective Technique for QoS Support in Clusters. IEEE Transactions on Parallel and Distributed Systems, 2007, 18, 1714-1726.	5.6	14
172	Efficient Switches with QoS Support for Clusters. , 2007, , .		0
173	Handling Topology Changes in InfiniBand. IEEE Transactions on Parallel and Distributed Systems, 2007, 18, 172-185.	5.6	6
174	Deadline-based QoS Algorithms for High-performance Networks. , 2007, , .		1
175	RECN-IQ: A Cost-Effective Input-Queued Switch Architecture with Congestion Management. , 2007, , .		18
176	Decongestants for clogged networks. IEEE Potentials, 2007, 26, 36-41.	0.3	0
177	A Low Overhead Fault Tolerant Coherence Protocol for CMP Architectures. , 2007, , .		24
178	A Formal Model to Manage the InfiniBand Arbitration Tables Providing QoS. IEEE Transactions on Computers, 2007, 56, 1024-1039.	3.4	4
179	A Latency-Aware Partitioning Method for Distributed Virtual Environment Systems. IEEE Transactions on Parallel and Distributed Systems, 2007, 18, 1215-1226.	5.6	30
180	Exploring IBA Design Space for Improved Performance. IEEE Transactions on Parallel and Distributed Systems, 2007, 18, 498-510.	5.6	0

#	Article	IF	Citations
181	An Effective Starvation Avoidance Mechanism to Enhance the Token Coherence Protocol. Parallel, Distributed and Network-based Processing, Proceedings of the Euromicro Workshop on, 2007, , .	0.0	10
182	Congestion Management in MINs through Marked and Validated Packets. , 2007, , .		15
183	A genetic approach for adding QoS to distributed virtual environments. Computer Communications, 2007, 30, 731-739.	5.1	3
184	Fast routing computation on InfiniBand networks. IEEE Transactions on Parallel and Distributed Systems, 2006, 17, 215-226.	5.6	16
185	Efficient, Scalable Congestion Management for Interconnection Networks. IEEE Micro, 2006, 26, 52-66.	1.8	43
186	A routing methodology for achieving fault tolerance in direct networks. IEEE Transactions on Computers, 2006, 55, 400-415.	3.4	76
187	MMR: A MultiMedia Router architecture to support hybrid workloads. Journal of Parallel and Distributed Computing, 2006, 66, 307-321.	4.1	0
188	FIR: An efficient routing strategy for tori and meshes. Journal of Parallel and Distributed Computing, 2006, 66, 907-921.	4.1	2
189	Enforcing in-order packet delivery in system area networks with adaptive routing. Journal of Parallel and Distributed Computing, 2005, 65, 1223-1236.	4.1	6
190	On-chip interconnects and instruction steering schemes for clustered microarchitectures. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 130-144.	5.6	12
191	A methodology for developing deadlock-free dynamic network reconfiguration processes. Part II. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 428-443.	5.6	27
192	A family of mechanisms for congestion control in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 772-784.	5.6	64
193	Traffic scheduling solutions with QoS support for an input-buffered multimedia router. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 1009-1021.	5.6	5
194	A theory for deadlock-free dynamic network reconfiguration. Part I. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 412-427.	5.6	39
195	A two-level directory architecture for highly scalable cc-NUMA multiprocessors. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 67-79.	5.6	37
196	Improving the performance of distributed virtual environment systems. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 637-649.	5.6	73
197	Deadlock-free dynamic reconfiguration over InfiniBandâ„¢ NETWORKS. International Journal of Parallel, Emergent and Distributed Systems, 2004, 19, 127-143.	0.4	9
198	An Efficient Fault-Tolerant Routing Methodology for Meshes and Tori. IEEE Computer Architecture Letters, 2004, 3, 3-3.	1.5	50

#	Article	IF	Citations
199	On the development of a communication-aware task mapping technique. Journal of Systems Architecture, 2004, 50, 207-220.	4.3	23
200	An effective methodology to improve the performance of the up*/down* routing algorithm. IEEE Transactions on Parallel and Distributed Systems, 2004, 15, 740-754.	5.6	49
201	An architecture for high-performance scalable shared-memory multiprocessors exploiting on-chip integration. IEEE Transactions on Parallel and Distributed Systems, 2004, 15, 755-768.	5.6	17
202	QoS in InfiniBand subnetworks. IEEE Transactions on Parallel and Distributed Systems, 2004, 15, 810-823.	5.6	38
203	Simple Deadlock-Free Dynamic Network Reconfiguration. Lecture Notes in Computer Science, 2004, , 504-515.	1.3	18
204	Supporting adaptive routing in IBA switches. Journal of Systems Architecture, 2003, 49, 441-456.	4.3	14
205	Applying in-transit buffers to boost the performance of networks with source routing. IEEE Transactions on Computers, 2003, 52, 1134-1153.	3.4	3
206	Fc3d: flow control-based distributed deadlock detection mechanism for true fully adaptive routing in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 2003, 14, 765-779.	5.6	29
207	Deadlock-free dynamic reconfiguration schemes for increased network dependability. IEEE Transactions on Parallel and Distributed Systems, 2003, 14, 780-794.	5.6	63
208	Boosting the performance of Myrinet networks. IEEE Transactions on Parallel and Distributed Systems, 2002, 13, 1166-1182.	5.6	3
209	A protocol for deadlock-free dynamic reconfiguration in high-speed local area networks. IEEE Transactions on Parallel and Distributed Systems, 2001, 12, 115-132.	5.6	62
210	A cost-effective approach to deadlock handling in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 2001, 12, 716-729.	5.6	19
211	A Comparison of Router Architectures for Virtual Cut-Through and Wormhole Switching in a NOW Environment. Journal of Parallel and Distributed Computing, 2001, 61, 224-253.	4.1	26
212	A general theory for deadlock-free adaptive routing using a mixed set of resources. IEEE Transactions on Parallel and Distributed Systems, 2001, 12, 1219-1235.	5.6	74
213	An efficient implementation of tree-based multicast routing for distributed shared-memory multiprocessors. Journal of Systems Architecture, 2000, 46, 1019-1032.	4.3	38
214	High-performance routing in networks of workstations with irregular topology. IEEE Transactions on Parallel and Distributed Systems, 2000, 11, 699-719.	5.6	63
215	A Flexible Routing Scheme for Networks of Workstations. Lecture Notes in Computer Science, 2000, , 260-267.	1.3	22
216	On the use of virtual channels in networks of workstations with irregular topology. IEEE Transactions on Parallel and Distributed Systems, 2000, 11, 813-828.	5.6	10

#	Article	IF	Citations
217	Dynamic Reconfiguration in High Speed Local Area Networks. , 2000, , 207-231.		5
218	Dynamically configurable message flow control for fault-tolerant routing. IEEE Transactions on Parallel and Distributed Systems, 1999, 10, 7-22.	5.6	21
219	Deadlock-Free Routing in Irregular Networks with Dynamic Reconfiguration. Lecture Notes in Computer Science, 1999, , 165-180.	1.3	4
220	A cost-effective methodology for the evaluation of interconnection networks. Journal of Systems Architecture, 1998, 44, 815-830.	4.3	2
221	SUBOPTIMAL-OPTIMAL ROUTING FOR LAN INTERNETWORKING USING TRANSPARENT BRIDGES. International Journal of Foundations of Computer Science, 1998, 09, 139-156.	1.1	1
222	A lab course on computer architecture. , 1998, , .		0
223	A tool for the analysis of reconfiguration and routing algorithms in irregular networks. Lecture Notes in Computer Science, 1998, , 159-173.	1.3	4
224	A theory of fault-tolerant routing in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 1997, 8, 790-802.	5.6	72
225	A necessary and sufficient condition for deadlock-free routing in cut-through and store-and-forward networks. IEEE Transactions on Parallel and Distributed Systems, 1996, 7, 841-854.	5.6	114
226	A theory of deadlock-free adaptive multicast routing in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 1995, 6, 976-987.	5.6	48
227	Configurable flow control mechanisms for fault-tolerant routing. Computer Architecture News, 1995, 23, 220-229.	2.5	4
228	A necessary and sufficient condition for deadlock-free adaptive routing in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 1995, 6, 1055-1067.	5.6	292
229	Highly adaptive wormhole routing algorithms for n-dimensional torus. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 1995, , 87-104.	0.0	1
230	A Necessary and Sufficient Condition for Deadlock-Free Adaptive Routing in Wormhole Networks. , 1994, , .		37
231	Performance evaluation of adaptive routing algorithms for k-ary n-cubes. Lecture Notes in Computer Science, 1994, , 45-59.	1.3	62
232	A THEORY TO INCREASE THE EFFECTIVE REDUNDANCY IN WORMHOLE NETWORKS. Parallel Processing Letters, 1994, 04, 125-138.	0.6	15
233	Improving the efficiency of virtual channels with time-dependent selection functions. Future Generation Computer Systems, 1994, 10, 45-58.	7.5	13
234	A new theory of deadlock-free adaptive routing in wormhole networks. IEEE Transactions on Parallel and Distributed Systems, 1993, 4, 1320-1331.	5.6	669

#	Article	IF	Citations
235	ON THE DESIGN OF DEADLOCK-FREE ADAPTIVE MULTICAST ROUTING ALGORITHMS. Parallel Processing Letters, 1993, 03, 321-333.	0.6	1
236	Deadlock-free adaptive routing algorithms for the 3D-torus: Limitations and solutions. Lecture Notes in Computer Science, 1993, , 684-687.	1.3	38
237	CHANNEL CLASSES: A NEW CONCEPT FOR DEADLOCK AVOIDANCE IN WORMHOLE NETWORKS. Parallel Processing Letters, 1992, 02, 347-354.	0.6	4
238	Dynamic Fault Tolerance with Misrouting in Fat Trees. , 0, , .		7