Rami Melhem

List of Publications by Year
in descending order

86 papers	1,183 citations	759233 12 h-index	677142 22 g-index
86 all docs	86 docs citations	86 times ranked	760 citing authors

#	Article	IF	Citations
1	Increasing PCM main memory lifetime. , 2010, , .		119
2	Fault-Tolerant Rate-Monotonic Scheduling. Real-Time Systems, 1998, 15, 149-181.	1.3	81
3	Corollaries to Amdahl's Law for Energy. IEEE Computer Architecture Letters, 2008, 7, 25-28.	1.5	80
4	Energy-Aware Scheduling for Streaming Applications on Chip Multiprocessors. , 2007, , .		51
5	CAFO: Cost aware flip optimization for asymmetric memories. , 2015, , .		47
6	Supporting superpages in non-contiguous physical memory. , 2015, , .		41
7	Counter-Based Tree Structure for Row Hammering Mitigation in DRAM. IEEE Computer Architecture Letters, 2017, 16, 18-21.	1.5	39
8	Refresh Now and Then. IEEE Transactions on Computers, 2014, 63, 3114-3126.	3.4	36
9	RDIS: A recursively defined invertible set scheme to tolerate multiple stuck-at faults in resistive memory. , 2012, , .		29
10	Tolerating process variations in nanophotonic on-chip networks. , 2012, , .		29
11	Quality of Service Support for Fine-Grained Sharing on GPUs. , 2017, , .		28
12	Compilation Techniques for Optimizing Communication on Distributed-Memory Systems. , 1993, , .		25
13	Oblivious Routing in Fat-Tree Based System Area Networks With Uncertain Traffic Demands. IEEE/ACM Transactions on Networking, 2009, 17, 1439-1452.	3.8	25
14	Adaptive Page Migration for Irregular Data-intensive Applications under GPU Memory Oversubscription. , 2020, , .		25
15	Parallel Gauss-Jordan elimination for the solution of dense linear systems. Parallel Computing, 1987, 4, 339-343.	2.1	23
16	CA-RAM: A High-Performance Memory Substrate for Search-Intensive Applications. , 2007, , .		21
17	Analyzing the impact of useless write-backs on the endurance and energy consumption of PCM main memory. , 2011, , .		20
18	FusedCache: A Naturally Inclusive, Racetrack Memory, Dual-Level Private Cache. IEEE Transactions on Multi-Scale Computing Systems, 2016, 2, 69-82.	2.4	20

#	Article	IF	Citations
19	Shadow Replication: An Energy-Aware, Fault-Tolerant Computational Model for Green Cloud Computing. Energies, 2014, 7, 5151-5176.	3.1	19
20	Shadow Computing: An energy-aware fault tolerant computing model. , 2014, , .		19
21	Improving Bit Flip Reduction for Biased and Random Data. IEEE Transactions on Computers, 2016, 65, 3345-3356.	3.4	19
22	Simultaneous Multikernel: Fine-Grained Sharing of GPUs. IEEE Computer Architecture Letters, 2016, 15, 113-116.	1.5	18
23	Energy Consumption of Resilience Mechanisms in Large Scale Systems. , 2014, , .		17
24	RDIS: Tolerating Many Stuck-At Faults in Resistive Memory. IEEE Transactions on Computers, 2015, 64, 847-861.	3.4	17
25	Dynamic partitioning to mitigate stuck-at faults in emerging memories. , 2017, , .		15
26	An Efficient Hardware-Based Multi-hash Scheme for High Speed IP Lookup. , 2008, , .		14
27	Demand-Driven Data Flow Analysis for Communication Optimization. Parallel Processing Letters, 1997, 07, 359-370.	0.6	13
28	Winning with Pinning in NoC. , 2009, , .		13
29	Proactive circuit allocation in multiplane NoCs. , 2013, , .		13
30	ContextPreRF: Enhancing the Performance and Energy of GPUs With Nonuniform Register Access. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2016, 24, 343-347.	3.1	13
31	Reconfiguration in 3D meshes. , 0, , .		12
32	Adaptive and Power-Aware Resilience for Extreme-Scale Computing. , 2016, , .		12
33	Codesign of NoC and Cache Organization for Reducing Access Latency in Chip Multiprocessors. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 1038-1046.	5.6	11
34	Thread Assignment Optimization with Real-Time Performance and Memory Bandwidth Guarantees for Energy-Efficient Heterogeneous Multi-core Systems. , 2012, , .		11
35	Partial Redundancy in HPC Systems with Non-Uniform Node Reliabilities. , 2018, , .		11
36	Yoda: Judge Me by My Size, Do You?. , 2017, , .		10

#	Article	IF	Citations
37	Enabling Fine-Grain Restricted Coset Coding Through Word-Level Compression for PCM. , 2018, , .		10
38	A Dynamic Pressure-Aware Associative Placement Strategy for Large Scale Chip Multiprocessors. IEEE Computer Architecture Letters, 2010, 9, 29-32.	1.5	9
39	Power of One Bit: Increasing Error Correction Capability with Data Inversion. , 2013, , .		9
40	Formal analysis of a systolic system for finite element stiffness matrices. Journal of Computer and System Sciences, 1985, 31, 1-27.	1.2	8
41	Energy efficient redundant configurations for real-time parallel reliable servers. Real-Time Systems, 2009, 41, 195-221.	1.3	8
42	Compiler Techniques for Efficient Communications in Circuit Switched Networks for Multiprocessor Systems. IEEE Transactions on Parallel and Distributed Systems, 2009, 20, 331-345.	5.6	8
43	Energy-aware checkpointing of divisible tasks with soft or hard deadlines. , 2013, , .		8
44	HMMSim: a simulator for hardware-software co-design of hybrid main memory. , 2015, , .		8
45	FLOWER and FaME: A Low Overhead Bit-Level Fault-map and Fault-Tolerance Approach for Deeply Scaled Memories. , 2020, , .		8
46	Quality of Service Support for Fine-Grained Sharing on GPUs. Computer Architecture News, 2017, 45, 269-281.	2.5	8
47	Low Diameter Interconnections for Routing in High-Performance Parallel Systems. IEEE Transactions on Computers, 2007, 56, 502-510.	3.4	7
48	Real-Time Scheduling for Phase Change Main Memory Systems. , 2011, , .		7
49	Compiler-Assisted Data Distribution and Network Configuration for Chip Multiprocessors. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 2058-2066.	5.6	7
50	A language for the simulation of systolic architectures. Computer Architecture News, 1985, 13, 310-314.	2.5	6
51	GASOLIN: Global Arbitration for Streams of Data in Optical Links. , 2015, , .		6
52	Symbol Shifting: Tolerating More Faults in PCM Blocks. IEEE Transactions on Computers, 2016, 65, 2270-2283.	3.4	5
53	Effect of scheduling jitter on end-to-end delay in TDMA protocols. , 0, , .		4
54	Scalable Multi-cache Simulation Using GPUs. , 2011, , .		4

#	Article	IF	Citations
55	Power-aware Manhattan Routing on Chip Multiprocessors. , 2012, , .		4
56	Data Dependent Sparing to Manage Better-Than-Bad Blocks. IEEE Computer Architecture Letters, 2013, 12, 43-46.	1.5	4
57	A Practical Data Classification Framework for Scalable and High Performance Chip-Multiprocessors. IEEE Transactions on Computers, 2014, 63, 2905-2918.	3.4	4
58	Characterizing the Overhead of Software-Managed Hybrid Main Memory. , 2015, , .		4
59	Synthesis of systolic algorithm design. Parallel Computing, 1989, 12, 195-207.	2.1	3
60	Minimizing Wavelength Conversions in WDM Path Establishment*. Photonic Network Communications, 2001, 3, 197-211.	2.7	3
61	SYMBOLIC EXPRESSION ANALYSIS FOR COMPILED COMMUNICATION. Parallel Processing Letters, 2008, 18, 567-587.	0.6	3
62	ASSESSING THE PERFORMANCE OF ENERGY-AWARE MAPPINGS. Parallel Processing Letters, 2013, 23, 1340003.	0.6	3
63	MSCS. , 2015, , .		3
64	Symmetry-Agnostic Coordinated Management of the Memory Hierarchy in Multicore Systems. Transactions on Architecture and Code Optimization, 2016, 12, 1-26.	2.0	3
65	The application of a sequence notation to the design of systolic computations. BIT Numerical Mathematics, 1989, 29, 409-427.	2.0	2
66	Minimizing expected energy consumption for streaming applications with linear dependencies on chip multiprocessors. , 2009, , .		2
67	Writeback-aware bandwidth partitioning for multi-core systems with PCM. , 2013, , .		2
68	Weighted-Tuple Synchronization for Parallel Architecture Simulators. , 2014, , .		2
69	Reciprocal abstraction for computer architecture co-simulation. , 2015, , .		2
70	Harvesting Underutilized Resources to Improve Responsiveness and Tolerance to Crash and Silent Faults for Data-Intensive Applications. , 2017, , .		2
71	CoLoR: Co-Located Rescuers for Fault Tolerance in HPC Systems. , 2018, , .		2
72	Revolver: Vertex-Centric Graph Partitioning Using Reinforcement Learning. , 2018, , .		2

#	Article	IF	Citations
73	Optimal Placement of In-memory Checkpoints Under Heterogeneous Failure Likelihoods. , 2019, , .		2
74	Verification of a class of self-timed computational networks. BIT Numerical Mathematics, 1987, 27, 480-500.	2.0	1
75	RECONFIGURATION IN FAULT-TOLERANT 3D MESHES. Parallel Processing Letters, 1995, 05, 387-399.	0.6	1
76	Supporting Loss Guarantees in Buffer-Limited Networks. , 2006, , .		1
77	Empirical, Analytical Study of Hardware-Based Page Swap in Hybrid Main Memory System. , 2016, , .		1
78	PREMSim: A Resilience Framework for Modeling Traditional and Emerging Memory Reliability. , 2019, , .		1
79	Scheduling to Minimize theWorst-Case Loss Rate. , 2007, , .		0
80	GroupBeat: Wireless sensor networks made reliable. , 2008, , .		0
81	Guest Editor's Introduction: Special Section on Power-Aware Parallel and Distributed Computing (PAPADS). IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 1441-1444.	5.6	0
82	Guest Editors' Introduction: Special Section on Energy Efficient Computing. IEEE Transactions on Computers, 2012, 61, 1666-1667.	3.4	0
83	Deterministic Multiplexing of NoC on Grid CMPs. , 2013, , .		0
84	Ordering circuit establishment in multiplane NoCs. ACM Transactions on Design Automation of Electronic Systems, 2013, 18, 1-33.	2.6	0
85	A systematic fault-tolerant computational model for both crash failures and silent data corruption. , 2018, , .		0
86	Predicting Single Event Effects in DRAM. , 2019, , .		0