Julio Sahuquillo

List of Publications by Year
in descending order

114	922	687220 13	752573 20
papers	citations	h-index	g-index

116	116	116	534
all docs	docs citations	times ranked	citing authors

#	Article	IF	Citations
1	VMT: Virtualized Multi-Threading for Accelerating Graph Workloads on Commodity Processors. IEEE Transactions on Computers, 2022, 71, 1386-1398.	2.4	0
2	DeepP: Deep Learning Multi-Program Prefetch Configuration for the IBM POWER 8. IEEE Transactions on Computers, 2022, 71, 2646-2658.	2.4	1
3	Effect of Hyper-Threading in Latency-Critical Multithreaded Cloud Applications and Utilization Analysis of the Major System Resources. Future Generation Computer Systems, 2022, 131, 194-208.	4.9	4
4	A Neural Network to Estimate Isolated Performance from Multi-Program Execution. , 2022, , .		0
5	Hy-Sched: A Simple Hyperthreading-Aware Thread to Core Allocation Strategy. IEEE Computer Architecture Letters, 2021, 20, 26-29.	1.0	0
6	Segment Switching: A New Switching Strategy for Optical HPC Networks. IEEE Access, 2021, 9, 43095-43106.	2.6	1
7	Thread Isolation to Improve Symbiotic Scheduling on SMT Multicore Processors. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 359-373.	4.0	4
8	An efficient cache flat storage organization for multithreaded workloads for low power processors. Future Generation Computer Systems, 2020, 110, 1037-1054.	4.9	0
9	Phase-Aware Cache Partitioning to Target Both Turnaround Time and System Performance. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 2556-2568.	4.0	12
10	Bandwidth-Aware Dynamic Prefetch Configuration for IBM POWER8. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 1970-1982.	4.0	6
11	An Aging-Aware GPU Register File Design Based on Data Redundancy. IEEE Transactions on Computers, 2019, 68, 4-20.	2.4	6
12	Efficient Management of Cache Accesses to Boost GPGPU Memory Subsystem Performance. IEEE Transactions on Computers, 2019, 68, 1442-1454.	2.4	0
13	Way Combination for an Adaptive and Scalable Coherence Directory. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 2608-2623.	4.0	2
14	FOS: a low-power cache organization for multicores. Journal of Supercomputing, 2019, 75, 6542-6573.	2.4	0
15	Foreword to the Special Issue on Processors, Interconnects, Storage, and Caches for Exascale Systems. Concurrency Computation Practice and Experience, 2019, 31, e5408.	1.4	0
16	Modeling and analysis of the performance of exascale photonic networks. Concurrency Computation Practice and Experience, 2019, 31, e4773.	1.4	3
17	Designing lab sessions focusing on real processors for computer architecture courses: A practical perspective. Journal of Parallel and Distributed Computing, 2018, 118, 128-139.	2.7	1
18	Accurately modeling the on-chip and off-chip GPU memory subsystem. Future Generation Computer Systems, 2018, 82, 510-519.	4.9	12

#	Article	IF	Citations
19	Improving GPU Cache Hierarchy Performance with a Fetch and Replacement Cache. Lecture Notes in Computer Science, 2018, , 235-248.	1.0	2
20	Improving System Turnaround Time with Intel CAT by Identifying LLC Critical Applications. Lecture Notes in Computer Science, 2018, , 603-615.	1.0	8
21	Efficient selective multicore prefetching under limited memory bandwidth. Journal of Parallel and Distributed Computing, 2018, 120, 32-43.	2.7	3
22	The Tag Filter Architecture: An energy-efficient cache and directory design. Journal of Parallel and Distributed Computing, 2017, 100, 193-202.	2.7	5
23	A research-oriented course on Advanced Multicore Architecture: Contents and active learning methodologies. Journal of Parallel and Distributed Computing, 2017, 105, 63-72.	2.7	6
24	A Hardware Approach to Fairly Balance the Inter-Thread Interference in Shared Caches. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 3021-3032.	4.0	7
25	Perf&Fair: A Progress-Aware Scheduler to Enhance Performance and Fairness in SMT Multicores. IEEE Transactions on Computers, 2017, 66, 905-911.	2.4	20
26	On Microarchitectural Mechanisms for Cache Wearout Reduction. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2017, 25, 857-871.	2.1	11
27	Improving IBM POWER8 Performance Through Symbiotic Job Scheduling. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 2838-2851.	4.0	4
28	Application Clustering Policies to Address System Fairness with Intelâ€™s Cache Allocation Technology. , 2017, , .		23
29	Exploiting Data Compression to Mitigate Aging in GPU Register Files. , 2017, , .		1
30	A Directory Cache with Dynamic Private-Shared Partitioning. , 2016, , .		2
31	Symbiotic job scheduling on the IBM POWER8. , 2016, , .		11
32	A Simple Activation/Deactivation Prefetching Scheme for Chip Multiprocessors. , 2016, , .		1
33	Impact of Memory-Level Parallelism on the Performance of GPU Coherence Protocols. , 2016, , .		0
34	Bandwidth-Aware On-Line Scheduling in SMT Multicores. IEEE Transactions on Computers, 2016, 65, 422-434.	2.4	14
35	Enhancing the L1 Data Cache Design to Mitigate HCI. IEEE Computer Architecture Letters, 2016, 15, 93-96.	1.0	2
36	A dynamic execution time estimation model to save energy in heterogeneous multicores running periodic tasks. Future Generation Computer Systems, 2016, 56, 211-219.	4.9	8

#	Article	IF	Citations
37	PS-Cache: an energy-efficient cache design for chip multiprocessors. Journal of Supercomputing, 2015, 71, 67-86.	2.4	9
38	A reuse-based refresh policy for energy-aware eDRAM caches. Microprocessors and Microsystems, 2015, 39, 37-48.	1.8	5
39	PS directory: a scalable multilevel directory cache for CMPs. Journal of Supercomputing, 2015, 71, 2847-2876.	2.4	7
40	Generating realistic workload for web performance studies. , 2015, , 157-186.		1
41	A new testbed for web performance evaluation. , 2015, , 225-251.		0
42	The impact of dynamic user workloads on web performance. , 2015, , 253-282.		0
43	Addressing Fairness in SMT Multicores with a Progress-Aware Scheduler. , 2015, , .		12
44	Accurately modeling the GPU memory subsystem. , 2015, , .		3
45	The Tag Filter Cache: An Energy-Efficient Approach. , 2015, , .		2
46	Impact of Partitioning Cache Schemes on the Cache Hierarchy of SMT Processors. , 2015, , .		0
47	Design of Hybrid Second-Level Caches. IEEE Transactions on Computers, 2015, 64, 1884-1897.	2.4	13
48	Addressing bandwidth contention in SMT multicores through scheduling. , 2014, , .		2
49	Analyzing the Optimal Voltage/Frequency Pair in Fault-Tolerant Caches. , 2014, , .		1
50	Dynamic WCET Estimation for Real-Time Multicore Embedded Systems Supporting DVFS. , 2014, , .		2
51	Efficient Register Renaming and Recovery for High-Performance Processors. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2014, 22, 1506-1514.	2.1	6
52	Cache-Hierarchy Contention-Aware Scheduling in CMPs. IEEE Transactions on Parallel and Distributed Systems, 2014, 25, 581-590.	4.0	21
53	Powerâ€aware scheduling with effective task migration for realâ€time multicore embedded systems. Concurrency Computation Practice and Experience, 2013, 25, 1987-2001.	1.4	19
54	Analyzing web server performance under dynamic user workloads. Computer Communications, 2013, 36, 386-395.	3.1	14

#	Article	IF	Citations
55	Hardware-based generation of independent subtraces of instructions in clustered processors. IEEE Transactions on Computers, 2013, 62, 944-955.	2.4	0
56	PS-cache: An energy-efficient cache design for chip multiprocessors. , 2013, , .		2
57	Using Huge Pages and Performance Counters to Determine the LLC Architecture. Procedia Computer Science, 2013, 18, 2557-2560.	1.2	1
58	Referrer Graph: A cost-effective algorithm and pruning method for predicting web accesses. Computer Communications, 2013, 36, 881-894.	3.1	46
59	Combining RAM Technologies for Hard-error Recovery in L1 Data Caches Working at Very-low Power Modes. , 2013, , .		3
60	An empirical model for predicting cross-core performance interference on multicore processors. , 2013, , .		2
61	The impact of user-browser interaction on web performance. , 2013, , .		3
62	Exploiting reuse information to reduce refresh energy in on-chip eDRAM caches. , 2013, , .		2
63	A New Methodology for Studying Realistic Processors in Computer Science Degrees. , 2013, , .		1
64	OMHI 2012: First International Workshop on On-chip Memory Hierarchies and Interconnects: Organization, Management and Implementation. Lecture Notes in Computer Science, 2013, , 305-306.	1.0	0
65	Effects of Process Variation on the Access Time in SRAM Cells. Lecture Notes in Computer Science, 2013, , 347-356.	1.0	0
66	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.	1.6	8
67	PS-Dir. , 2012, , .		10
68	The Impact of User's Dynamic Behavior on Web Performance. , 2012, , .		2
69	A Sequentially Consistent Multiprocessor Architecture for Out-of-Order Retirement of Instructions. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 1361-1368.	4.0	1
70	Page-Based Memory Allocation Policies of Local and Remote Memory in Cluster Computers. , 2012, , .		1
71	Understanding Cache Hierarchy Contention in CMPs to Improve Job Scheduling. , 2012, , .		16
72	Efficiently Handling Memory Accesses to Improve QoS in Multicore Systems under Real-Time Constraints. , 2012, , .		0

#	Article	IF	Citations
73	Analyzing the optimal ratio of SRAM banks in hybrid caches. , 2012, , .		3
74	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.	2.1	4
75	Key factors in web latency savings in an experimental prefetching system. Journal of Intelligent Information Systems, 2012, 39, 187-207.	2.8	3
76	A taxonomy of web prediction algorithms. Expert Systems With Applications, 2012, 39, 8496-8502.	4.4	10
77	Design, Performance, and Energy Consumption of eDRAM/SRAM Macrocells for L1 Data Caches. IEEE Transactions on Computers, 2012, 61, 1231-1242.	2.4	11
78	A cost-effective heuristic to schedule local and remote memory in cluster computers. Journal of Supercomputing, 2012, 59, 1533-1551.	2.4	4
79	Energy Behaviour of NUCA Caches in CMPs. , 2011, , .		5
80	MRU-Tour-based Replacement Algorithms for Last-Level Caches. , 2011, , .		1
81	Improving Last-Level Cache Performance by Exploiting the Concept of MRU-Tour. , 2011, , .		1
82	A New Energy-Aware Dynamic Task Set Partitioning Algorithm for Soft and Hard Embedded Real-Time Systems. Computer Journal, 2011, 54, 1282-1294.	1.5	11
83	A Dynamic Power-Aware Partitioner with Task Migration for Multicore Embedded Systems. Lecture Notes in Computer Science, 2011, , 218-229.	1.0	5
84	A Cluster Computer Performance Predictor for Memory Scheduling. Lecture Notes in Computer Science, 2011, , 353-362.	1.0	1
85	Using current web page structure to improve prefetching performance. Computer Networks, 2010, 54, 1404-1417.	3.2	17
86	Dynamic task set partitioning based on balancing resource requirements and utilization to reduce power consumption. , 2010, , .		0
87	Exploiting subtrace-level parallelism in clustered processors. , 2010, , .		0
88	A Scheduling Heuristic to Handle Local and Remote Memory in Cluster Computers. , 2010, , .		2
89	Evaluation, Analysis and Adaptation of Web Prefetching Techniques in Current Web. Advanced Information and Knowledge Processing, 2010, , 239-271.	0.2	4
90	Speculative Validation of Web Objects for Further Reducing the User-Perceived Latency. Lecture Notes in Computer Science, 2010, , 239-250.	1.0	1

#	Article	IF	Citations
91	Dynamic task set partitioning based on balancing memory requirements to reduce power consumption. , 2009, , .		0
92	Dweb model: Representing Web 2.0 dynamism. Computer Communications, 2009, 32, 1118-1128.	3.1	27
93	An Efficient Low-Complexity Alternative to the ROB for Out-of-Order Retirement of Instructions. , 2009, , .		1
94	An hybrid eDRAM/SRAM macrocell to implement first-level data caches. , 2009, , .		20
95	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
96	Understanding cache hierarchy interactions with a program-driven simulator. , 2007, , .		3
97	Leakage Current Reduction in Data Caches on Embedded Systems. , 2007, , .		7
98	Multi2Sim: A Simulation Framework to Evaluate Multicore-Multithreaded Processors. , 2007, , .		129
99	A user-focused evaluation of web prefetching algorithms. Computer Communications, 2007, 30, 2213-2224.	3.1	33
100	Spim-Cache: A Pedagogical Tool for Teaching Cache Memories Through Code-Based Exercises. IEEE Transactions on Education, 2007, 50, 244-250.	2.0	9
101	The Impact of the Web Prefetching Architecture on the Limits of Reducing User's Perceived Latency. , 2006, , .		18
102	Addressing a workload characterization study to the design of consistency protocols. Journal of Supercomputing, 2006, 38, 49-72.	2.4	0
103	Web prefetching performance metrics: A survey. Performance Evaluation, 2006, 63, 988-1004.	0.9	35
104	<tex>$hbox RAC_rm FP$</tex>: A Training Tool to Work With Floating-Point Representation, Algorithms, and Circuits in Undergraduate Courses. IEEE Transactions on Education, 2006, 49, 321-331.	2.0	3
105	An execution-driven simulation tool for teaching cache memories in introductory computer organization courses. , 2006, , .		4
106	DDG: An Efficient Prefetching Algorithm for Current Web Generation. , 2006, , .		23
107	Cost-Benefit Analysis of Web Prefetching Algorithms from the Userâ€™s Point of View. Lecture Notes in Computer Science, 2006, , 1113-1118.	1.0	4
108	Modelling users' dynamic behaviour in e-business environments using navigations. International Journal of Electronic Business, 2005, 3, 225.	0.2	0

#	Article	IF	Citations
109	Exploring the performance of split data cache schemes on superscalar processors and symmetric multiprocessors. Journal of Systems Architecture, 2005, 51, 451-469.	2.5	2
110	Modeling continuous changes of the user's dynamic behavior in the WWW. , 2005, , .		8
111	Exploiting temporal locality in drowsy cache policies. , 2005, , .		34
112	On-chip interconnects and instruction steering schemes for clustered microarchitectures. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 130-144.	4.0	12
113	Splitting the data cache: a survey. IEEE Concurrency, 2000, 8, 30-35.	0.8	18
114	The differences between distributed shared memory caching and proxy caching. IEEE Concurrency, 2000, 8, 45-47.	0.8	4