Salvador Petit

List of Publications by Year
in descending order

Source: https://exaly.com/author-pdf/8342651/publications.pdf

Version: 2024-02-01

100 papers	640 citations	840585 11 h-index	996849 15 g-index
110 all docs	110 docs citations	110 times ranked	411 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	Segment Switching: A New Switching Strategy for Optical HPC Networks. IEEE Access, 2021, 9, 43095-43106.	2.6	1
6	Thread Isolation to Improve Symbiotic Scheduling on SMT Multicore Processors. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 359-373.	4.0	4
7	An efficient cache flat storage organization for multithreaded workloads for low power processors. Future Generation Computer Systems, 2020, 110, 1037-1054.	4.9	0
8	Impact of the Array Shape and Memory Bandwidth on the Execution Time of CNN Systolic Arrays. , 2020, , .		3
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	Modeling and analysis of the performance of exascale photonic networks. Concurrency Computation Practice and Experience, 2019, 31, e4773.	1.4	3
16	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
17	Accurately modeling the on-chip and off-chip GPU memory subsystem. Future Generation Computer Systems, 2018, 82, 510-519.	4.9	12
18	Workload Characterization for Exascale Computing Networks. , 2018, , .		1

#	Article	IF	Citations
19	A Workload Generator for Evaluating SMT Real-Time Systems. , 2018, , .		0
20	Improving GPU Cache Hierarchy Performance with a Fetch and Replacement Cache. Lecture Notes in Computer Science, 2018, , 235-248.	1.0	2
21	Improving System Turnaround Time with Intel CAT by Identifying LLC Critical Applications. Lecture Notes in Computer Science, 2018, , 603-615.	1.0	8
22	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
23	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
24	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
25	On Microarchitectural Mechanisms for Cache Wearout Reduction. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2017, 25, 857-871.	2.1	11
26	Improving IBM POWER8 Performance Through Symbiotic Job Scheduling. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 2838-2851.	4.0	4
27	Application Clustering Policies to Address System Fairness with Intelâ€™s Cache Allocation Technology. , 2017, , .		23
28	Modeling a Photonic Network for Exascale Computing. , 2017, , .		2
29	Exploiting Data Compression to Mitigate Aging in GPU Register Files. , 2017, , .		1
30	Student Research Poster. , 2016, , .		0
31	Symbiotic job scheduling on the IBM POWER8. , 2016, , .		11
32	Accurately modeling a photonic NoC in a detailed CMP simulation framework. , 2016, , .		3
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	Current challenges in simulations of HPC systems. , 2015, , .		0
38	A reuse-based refresh policy for energy-aware eDRAM caches. Microprocessors and Microsystems, 2015, 39, 37-48.	1.8	5
39	Addressing Fairness in SMT Multicores with a Progress-Aware Scheduler. , 2015, , .		12
40	Accurately modeling the GPU memory subsystem. , 2015, , .		3
41	A Research-Oriented Course on Advanced Multicore Architecture. , 2015, , .		1
42	Design of Hybrid Second-Level Caches. IEEE Transactions on Computers, 2015, 64, 1884-1897.	2.4	13
43	Addressing bandwidth contention in SMT multicores through scheduling. , 2014, , .		2
44	Analyzing the Optimal Voltage/Frequency Pair in Fault-Tolerant Caches. , 2014, , .		1
45	Dynamic WCET Estimation for Real-Time Multicore Embedded Systems Supporting DVFS. , 2014, , .		2
46	Euro-Par 2014: Parallel Processing Workshops. Lecture Notes in Computer Science, 2014, , .	1.0	1
47	Using peer-assessed returnables in multiple stages to improve learning in computer organization courses. , 2014, , .		0
48	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
49	Cache-Hierarchy Contention-Aware Scheduling in CMPs. IEEE Transactions on Parallel and Distributed Systems, 2014, 25, 581-590.	4.0	21
50	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
51	Hardware-based generation of independent subtraces of instructions in clustered processors. IEEE Transactions on Computers, 2013, 62, 944-955.	2.4	0
52	Using Huge Pages and Performance Counters to Determine the LLC Architecture. Procedia Computer Science, 2013, 18, 2557-2560.	1.2	1
53	Combining RAM Technologies for Hard-error Recovery in L1 Data Caches Working at Very-low Power Modes. , 2013, , .		3
54	An empirical model for predicting cross-core performance interference on multicore processors. , 2013, , .		2

#	Article	IF	Citations
55	Exploiting reuse information to reduce refresh energy in on-chip eDRAM caches. , 2013, , .		2
56	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
57	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
58	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
59	Page-Based Memory Allocation Policies of Local and Remote Memory in Cluster Computers. , 2012, , .		1
60	Understanding Cache Hierarchy Contention in CMPs to Improve Job Scheduling. , 2012, , .		16
61	Efficiently Handling Memory Accesses to Improve QoS in Multicore Systems under Real-Time Constraints. , 2012, , .		0
62	Analyzing the optimal ratio of SRAM banks in hybrid caches. , 2012, , .		3
63	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
64	Design, Performance, and Energy Consumption of eDRAM/SRAM Macrocells for L1 Data Caches. IEEE Transactions on Computers, 2012, 61, 1231-1242.	2.4	11
65	A cost-effective heuristic to schedule local and remote memory in cluster computers. Journal of Supercomputing, 2012, 59, 1533-1551.	2.4	4
66	MRU-Tour-based Replacement Algorithms for Last-Level Caches. , 2011, , .		1
67	Improving Last-Level Cache Performance by Exploiting the Concept of MRU-Tour. , 2011, , .		1
68	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
69	A Dynamic Power-Aware Partitioner with Task Migration for Multicore Embedded Systems. Lecture Notes in Computer Science, 2011, , 218-229.	1.0	5
70	A Cluster Computer Performance Predictor for Memory Scheduling. Lecture Notes in Computer Science, 2011, , 353-362.	1.0	1
71	Dynamic task set partitioning based on balancing resource requirements and utilization to reduce power consumption. , 2010, , .		0
72	Exploiting subtrace-level parallelism in clustered processors. , 2010, , .		0

#	Article	IF	Citations
73	Out-of-order retirement of instructions in sequentially consistent multiprocessors. , 2010, , .		0
74	A Scheduling Heuristic to Handle Local and Remote Memory in Cluster Computers. , 2010, , .		2
75	Balancing Task Resource Requirements in Embedded Multithreaded Multicore Processors to Reduce Power Consumption. , 2010, , .		1
76	Dynamic task set partitioning based on balancing memory requirements to reduce power consumption. , 2009, , .		0
77	Power Reduction In Advanced Embedded IPC Processors. Intelligent Automation and Soft Computing, 2009, 15, 495-507.	1.6	0
78	A Complexity-Effective Out-of-Order Retirement Microarchitecture. IEEE Transactions on Computers, 2009, 58, 1626-1639.	2.4	12
79	An Efficient Low-Complexity Alternative to the ROB for Out-of-Order Retirement of Instructions. , 2009, , .		1
80	A power-aware hybrid RAM-CAM renaming mechanism for fast recovery. , 2009, , .		2
81	An hybrid eDRAM/SRAM macrocell to implement first-level data caches. , 2009, , .		20
82	Paired ROBs: A Cost-Effective Reorder Buffer Sharing Strategy for SMT Processors. Lecture Notes in Computer Science, 2009, , 309-320.	1.0	1
83	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
84	The impact of out-of-order commit in coarse-grain, fine-grain and simultaneous multithreaded architectures. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	3
85	Reducing the Number of Bits in the BTB to Attack the Branch Predictor Hot-Spot. Lecture Notes in Computer Science, 2008, , 317-326.	1.0	3
86	Leakage Current Reduction in Data Caches on Embedded Systems. , 2007, , .		7
87	Multi2Sim: A Simulation Framework to Evaluate Multicore-Multithreaded Processors. , 2007, , .		129
88	Impacts of Multiprocessor Configurations on Workloads in Bioinformatics. , 2007, , .		0
89	Spim-Cache: A Pedagogical Tool for Teaching Cache Memories Through Code-Based Exercises. IEEE Transactions on Education, 2007, 50, 244-250.	2.0	9
90	Applying the zeros switch-off technique to reduce static energy in data caches. Proceedings Symposium on Computer Architecture and High Performance Computing, 2006, , .	0.0	2

#	Article	IF	Citations
91	Addressing a workload characterization study to the design of consistency protocols. Journal of Supercomputing, 2006, 38, 49-72.	2.4	0
92	<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
93	An execution-driven simulation tool for teaching cache memories in introductory computer organization courses. , 2006, , .		4
94	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
95	Exploiting temporal locality in drowsy cache policies. , 2005, , .		34
96	LIDE. Computer Architecture News, 2000, 28, 11-18.	2.5	8
97	About the sensitivity of the HLRC-DU protocol on diff and page sizes. , 0, , .		2
98	Characterizing parallel workloads to reduce multiple writer overhead in shared virtual memory systems. , 0, , .		4
99	Characterizing the Dynamic Behavior of Workload Execution in SVM systems. , 0, , .		1
100	A Comparison Study of the HLRC-DU Protocol versus a HLRC Hardware Assisted Protocol. , 0, , .		0