

Mithuna Thottethodi

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/9175195/publications.pdf>

Version: 2024-02-01

41
papers

1,187
citations

1307594

7
h-index

1372567

10
g-index

41
all docs

41
docs citations

41
times ranked

821
citing authors

#	ARTICLE	IF	CITATIONS
1	Karma: Cost-Effective Geo-Replicated Cloud Storage with Dynamic Enforcement of Causal Consistency. IEEE Transactions on Cloud Computing, 2021, 9, 197-211.	4.4	8
2	Dart: Divide and Specialize for Fast Response to Congestion in RDMA-Based Datacenter Networks. IEEE/ACM Transactions on Networking, 2020, 28, 322-335.	3.8	11
3	Secure automatic bounds checking: prevention is simpler than cure. , 2020, , .		3
4	Network Interface Architecture for Remote Indirect Memory Access (RIMA) in Datacenters. Transactions on Architecture and Code Optimization, 2020, 17, 1-22.	2.0	2
5	Newton: A DRAM-makerâ€™s Accelerator-in-Memory (AiM) Architecture for Machine Learning. , 2020, , .		48
6	SparTen. , 2019, , .		144
7	ACCORD: Automated Change Coordination across Independently Administered Cloud Services. , 2018, , .		0
8	Fast Congestion Control in RDMA-based Datacenter Networks. , 2018, , .		3
9	Millipede: Die-Stacked Memory Optimizations for Big Data Machine Learning Analytics. , 2018, , .		1
10	NutShell. , 2017, , .		4
11	Efficient Collaborative Approximation in MapReduce without Missing Rare Keys. , 2017, , .		0
12	Extended Task Queuing: Active Messages for Heterogeneous Systems. , 2016, , .		4
13	Scalable, Global, Optimal-bandwidth, Application-Specific Routing. , 2016, , .		1
14	RAHTM: Routing Algorithm Aware Hierarchical Task Mapping. , 2014, , .		13
15	MorphStore: A local file system for Big Data with utility-driven replication and load-adaptive access scheduling. , 2014, , .		1
16	PreTrans: Reducing TLB CAM-search via page number prediction and speculative pre-translation. , 2013, , .		3
17	MapReduce with communication overlap (MaRCO). Journal of Parallel and Distributed Computing, 2013, 73, 608-620.	4.1	44
18	Software-programmable continuous-flow multi-purpose lab-on-a-chip. Microfluidics and Nanofluidics, 2013, 15, 647-659.	2.2	24

#	ARTICLE	IF	CITATIONS
19	Understanding and mitigating the impact of load imbalance in the memory caching tier. , 2013, , .		49
20	A Mostly-Clean DRAM Cache for Effective Hit Speculation and Self-Balancing Dispatch. , 2012, , .		74
21	Selective commitment and selective margin: Techniques to minimize cost in an IaaS cloud. , 2012, , .		9
22	Dynamic server provisioning to minimize cost in an IaaS cloud. , 2011, , .		53
23	TransCom. , 2011, , .		4
24	SieveStore. , 2010, , .		85
25	Trifecta: A Nonspeculative Scheme to Exploit Common, Data-Dependent Subcritical Paths. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2010, 18, 53-65.	3.1	24
26	LiteTM: Reducing transactional state overhead. , 2010, , .		5
27	Adaptive Flow Control for Robust Performance and Energy. , 2010, , .		28
28	Disjoint-path routing: Efficient communication for streaming applications. , 2009, , .		4
29	Undergraduate dual-core prototyping and analysis of factors influencing student success on dual-core designs. , 2009, , .		1
30	Automatic volume management for programmable microfluidics. ACM SIGPLAN Notices, 2008, 43, 56-67.	0.2	0
31	Automatic volume management for programmable microfluidics. , 2008, , .		8
32	Aquacore. , 2007, , .		23
33	Table-lookup based Crossbar Arbitration for Minimal-Routed, 2D Mesh and Torus Networks. , 2007, , .		2
34	Aquacore. Computer Architecture News, 2007, 35, 254-265.	2.5	6
35	Effective Management of DRAM Bandwidth in Multicore Processors. Parallel Architecture and Compilation Techniques (PACT), Proceedings of the International Conference on, 2007, , .	0.0	39
36	Evaluating ISA Support and Hardware Support for Recursive Data Layouts. , 2007, , 95-106.		2

#	ARTICLE	IF	CITATIONS
37	Architectural support for operating system-driven CMP cache management. , 2006, , .		175
38	Exploiting global knowledge to achieve self-tuned congestion control for k-ary n-cube networks. IEEE Transactions on Parallel and Distributed Systems, 2004, 15, 257-272.	5.6	30
39	Recursive array layouts and fast matrix multiplication. IEEE Transactions on Parallel and Distributed Systems, 2002, 13, 1105-1123.	5.6	71
40	Nonlinear array layouts for hierarchical memory systems. , 1999, , .		116
41	Recursive array layouts and fast parallel matrix multiplication. , 1999, , .		65