

Xiaobo Zhou

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/2617944/publications.pdf>

Version: 2024-02-01

75
papers

1,017
citations

687363

13
h-index

677142

22
g-index

76
all docs

76
docs citations

76
times ranked

747
citing authors

#	ARTICLE	IF	CITATIONS
1	Improving Performance of Heterogeneous MapReduce Clusters with Adaptive Task Tuning. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 774-786.	5.6	71
2	iShuffle: Improving Hadoop Performance with Shuffle-on-Write. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 1649-1662.	5.6	60
3	Energy Efficiency Aware Task Assignment with DVFS in Heterogeneous Hadoop Clusters. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 70-82.	5.6	43
4	Autonomic Provisioning with Self-Adaptive Neural Fuzzy Control for End-to-end Delay Guarantee. , 2010, , .		42
5	Optimizing virtual machine scheduling in NUMA multicore systems. , 2013, , .		40
6	Efficient Server Provisioning with Control for End-to-End Response Time Guarantee on Multitier Clusters. IEEE Transactions on Parallel and Distributed Systems, 2012, 23, 78-86.	5.6	39
7	Cross-Platform Resource Scheduling for Spark and MapReduce on YARN. IEEE Transactions on Computers, 2017, 66, 1341-1353.	3.4	37
8	NINEPIN: Non-invasive and energy efficient performance isolation in virtualized servers. , 2012, , .		29
9	Adaptive Scheduling Parallel Jobs with Dynamic Batching in Spark Streaming. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 2672-2685.	5.6	29
10	Scalable Distributed DL Training: Batching Communication and Computation. Proceedings of the AAAI Conference on Artificial Intelligence, 2019, 33, 5289-5296.	4.9	28
11	Heterogeneity-Aware Workload Placement and Migration in Distributed Sustainable Datacenters. , 2014, , .		27
12	Adaptive scheduling of parallel jobs in spark streaming. , 2017, , .		27
13	Deadline-Aware MapReduce Job Scheduling with Dynamic Resource Availability. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 814-826.	5.6	26
14	Elastic Power-Aware Resource Provisioning of Heterogeneous Workloads in Self-Sustainable Datacenters. IEEE Transactions on Computers, 2016, 65, 508-521.	3.4	23
15	V-Cache: Towards Flexible Resource Provisioning for Multi-tier Applications in IaaS Clouds. , 2013, , .		22
16	Towards Energy Efficiency in Heterogeneous Hadoop Clusters by Adaptive Task Assignment. , 2015, , .		22
17	Resource Allocation for Session-Based Two-Dimensional Service Differentiation on e-Commerce Servers. IEEE Transactions on Parallel and Distributed Systems, 2006, 17, 838-850.	5.6	21
18	PERFUME: Power and performance guarantee with fuzzy MIMO control in virtualized servers. , 2011, , .		18

#	ARTICLE	IF	CITATIONS
19	Automated and Agile Server Parameter Tuning with Learning and Control. , 2012, , .		18
20	Autonomic performance and power control for co-located Web applications on virtualized servers. , 2013, , .		18
21	Efficient server provisioning with end-to-end delay guarantee on multi-tier clusters. , 2009, , .		17
22	An Adaptive Growing Hierarchical Self Organizing Map for Network Intrusion Detection. , 2010, , .		17
23	Variant: a malware similarity testing framework. , 2015, , .		16
24	FlexSlot: Moving Hadoop Into the Cloud with Flexible Slot Management. , 2014, , .		15
25	Performance Isolation of Data-Intensive Scale-out Applications in a Multi-tenant Cloud. , 2018, , .		15
26	pVOCL: Power-Aware Dynamic Placement and Migration in Virtualized GPU Environments. , 2013, , .		14
27	Coordinated Power and Performance Guarantee with Fuzzy MIMO Control in Virtualized Server Clusters. IEEE Transactions on Computers, 2015, 64, 97-111.	3.4	14
28	Autonomic Performance and Power Control for Co-Located Web Applications in Virtualized Datacenters. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 1289-1302.	5.6	14
29	Aggressive Synchronization with Partial Processing for Iterative ML Jobs on Clusters. , 2018, , .		13
30	Overlapping Communication With Computation in Parameter Server for Scalable DL Training. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 2144-2159.	5.6	13
31	Pufferfish. , 2019, , .		12
32	Preemptive and Low Latency Datacenter Scheduling via Lightweight Containers. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 2749-2762.	5.6	12
33	Coordinated VM Resizing and Server Tuning: Throughput, Power Efficiency and Scalability. , 2012, , .		11
34	Characterizing Scheduling Delay for Low-Latency Data Analytics Workloads. , 2018, , .		11
35	Dependency-Aware Network Adaptive Scheduling of Data-Intensive Parallel Jobs. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 515-529.	5.6	11
36	A Correlation-Aware Page-Level FTL to Exploit Semantic Links in Workloads. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 723-737.	5.6	11

#	ARTICLE	IF	CITATIONS
37	Elastic Parameter Server: Accelerating ML Training With Scalable Resource Scheduling. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 1128-1143.	5.6	11
38	An adaptive process allocation strategy for proportional responsiveness differentiation on Web servers. , 2004, , .		9
39	A robust packet scheduling algorithm for proportional delay differentiation services. , 0, , .		9
40	Blockchain-Enabled Collaborative Intrusion Detection in Software Defined Networks. , 2020, , .		9
41	A Machine Learning Approach to Anomaly Detection Based on Traffic Monitoring for Secure Blockchain Networking. IEEE Transactions on Network and Service Management, 2022, 19, 3619-3632.	4.9	9
42	Heterogeneity Aware Workload Management in Distributed Sustainable Datacenters. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 375-387.	5.6	8
43	OS-Augmented Oversubscription of Opportunistic Memory with a User-Assisted OOM Killer. , 2019, , .		8
44	CoSAC: Coordinated Session-Based Admission Control for Multi-Tier Internet Applications. , 2009, , .		7
45	Improving Utilization and Parallelism of Hadoop Cluster by Elastic Containers. , 2018, , .		7
46	Two-Tier Resource Allocation for Slowdown Differentiation on Server Clusters. , 0, , .		6
47	Regression based multi-tier resource provisioning for session slowdown guarantees. , 2010, , .		6
48	Self-Tuning Batching with DVFS for Improving Performance and Energy Efficiency in Servers. , 2013, , .		6
49	Addressing Skewness in Iterative ML Jobs with Parameter Partition. , 2019, , .		6
50	ODDS: Optimizing Data-Locality Access for Scientific Data Analysis. IEEE Transactions on Cloud Computing, 2020, 8, 220-231.	4.4	6
51	aMOSS: Automated Multi-objective Server Provisioning with Stress-Strain Curving. , 2011, , .		5
52	Network-Adaptive Scheduling of Data-Intensive Parallel Jobs with Dependencies in Clusters. , 2017, , .		5
53	Blockchain-based Distributed Banking for Permissioned and Accountable Financial Transaction Processing. , 2020, , .		5
54	It Can Understand the Logs, Literally. , 2019, , .		4

#	ARTICLE	IF	CITATIONS
55	Lightweight and Identifier-Oblivious Engine for Cryptocurrency Networking Anomaly Detection. IEEE Transactions on Dependable and Secure Computing, 2023, 20, 1302-1318.	5.4	4
56	Analysis of a bandwidth allocation strategy for proportional streaming services. , 2004, , .		3
57	Packet Scheduling with Buffer Management for Fair Bandwidth Sharing and Delay Differentiation. , 2007, , .		3
58	First byte: Force-based clustering of filtered block N-grams to detect code reuse in malicious software. , 2013, , .		3
59	Online Adaptive Anomaly Detection for Augmented Network Flows. , 2014, , .		3
60	Autonomic Performance and Power Control on Virtualized Servers: Survey, Practices, and Trends. Journal of Computer Science and Technology, 2014, 29, 631-645.	1.5	3
61	StoreApp: A shared storage appliance for efficient and scalable virtualized Hadoop clusters. , 2015, , .		3
62	PLS: Proactive Load Shifting for Distributed SDN Controllers. , 2018, , .		3
63	Multi-tier Service Differentiation: Coordinated Resource Provisioning and Admission Control. , 2012, , .		2
64	An Analysis of Open Ports and Port Pairs in EC2 Instances. , 2017, , .		2
65	Discover and Secure (DaS): An Automated Virtual Machine Security Management Framework. , 2018, , .		2
66	System Hardening for Infrastructure as a Service (IaaS). , 2020, , .		2
67	Rate avalanche: The performance degradation in multi-rate 802.11 WLANs. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	1
68	Introduction to the special issue on self-adaptive and self-organizing wireless networking systems. ACM Transactions on Autonomous and Adaptive Systems, 2009, 4, 1-4.	0.8	1
69	An Integrated Feedback Control Approach for Proportional Slowdown Differentiation on Internet Servers. , 0, , .		0
70	A robust application-level approach for responsiveness differentiation. , 2005, , .		0
71	DIBS: Dual interval bandwidth scheduling for short-term differentiation. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	0
72	RADAR: Rate-Alert Dynamic RTS/CTS Exchange for Performance Enhancement in Multi-Rate Wireless Networks. , 2008, , .		0

#	ARTICLE	IF	CITATIONS
73	DIBS: A dual interval bandwidth scheduling algorithm for short-term differentiation. International Journal of Parallel, Emergent and Distributed Systems, 2011, 26, 165-178.	1.0	0
74	An I/O Efficient Distributed Approximation Framework Using Cluster Sampling. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 1494-1511.	5.6	0
75	Preemptive Two-Stage Update Mechanism for Software-Defined Networks. , 2019, , .		0