Evgenia Smirni

List of Publications by Year
in descending order

35 papers	672 citations	1163117 8 h-index	1199594 12 g-index
36 all docs	36 docs citations	36 times ranked	407 citing authors

#	Article	IF	Citations
1	A large-scale study of soft-errors on GPUs in the field. , 2016, , .		59
2	PRACTISE: Robust prediction of data center time series. , 2015, , .		54
3	KPC-Toolbox: Simple Yet Effective Trace Fitting Using Markovian Arrival Processes. , 2008, , .		51
4	Machine Learning Models for GPU Error Prediction in a Large Scale HPC System. , 2018, , .		49
5	Data Centers in the Cloud: A Large Scale Performance Study. , 2012, , .		44
6	State-of-the-practice in data center virtualization: Toward a better understanding of VM usage. , 2013, , .		44
7	Characterizing Temperature, Power, and Soft-Error Behaviors in Data Center Systems: Insights, Challenges, and Opportunities. , 2017, , .		37
8	Dealing with Burstiness in Multi-Tier Applications: Models and Their Parameterization. IEEE Transactions on Software Engineering, 2012, 38, 1040-1053.	5.6	35
9	Characterizing the BMAP/MAP/1 Departure Process via the ETAQA Truncation. Stochastic Models, 2005, 21, 821-846.	0.5	34
10	Model-Driven System Capacity Planning under Workload Burstiness. IEEE Transactions on Computers, 2010, 59, 66-80.	3.4	31
11	Fault Site Pruning for Practical Reliability Analysis of GPGPU Applications. , 2018, , .		27
12	Application-driven dynamic vertical scaling of virtual machines in resource pools. , 2014, , .		25
13	Practical Resilience Analysis of GPGPU Applications in the Presence of Single- and Multi-Bit Faults. IEEE Transactions on Computers, 2021, 70, 30-44.	3.4	22
14	Less Can Be More: Micro-managing VMs in Amazon EC2. , 2015, , .		19
15	Spatial–Temporal Prediction Models for Active Ticket Managing in Data Centers. IEEE Transactions on Network and Service Management, 2018, 15, 39-52.	4.9	17
16	Fastrack for taming burstiness and saving power in multi-tiered systems. , 2010, , .		16
17	CEDULE: A Scheduling Framework for Burstable Performance in Cloud Computing. , 2018, , .		13
18	The impact of I/O on program behavior and parallel scheduling. Performance Evaluation Review, 1998, 26, 56-65.	0.6	11

#	Article	IF	Citations
19	Efficient Deep Neural Network Serving: Fast and Furious. IEEE Transactions on Network and Service Management, 2018, 15, 112-126.	4.9	11
20	DyScale: A MapReduce Job Scheduler for Heterogeneous Multicore Processors. IEEE Transactions on Cloud Computing, 2017, 5, 317-330.	4.4	9
21	Enabling Software Resilience in GPGPU Applications via Partial Thread Protection. , 2021, , .		9
22	How to Supercharge the Amazon T2: Observations and Suggestions. , 2017, , .		8
23	Characterizing Accuracy-Aware Resilience of GPGPU Applications. , 2020, , .		7
24	CEDULE+: Resource Management for Burstable Cloud Instances Using Predictive Analytics. IEEE Transactions on Network and Service Management, 2021, 18, 945-957.	4.9	7
25	It's not a Sprint, it's a Marathon. , 2019, , .		7
26	Enhancing data availability in disk drives through background activities. , 2008, , .		6
27	Scheduling data analytics work with performance guarantees: queuing and machine learning models in synergy. Cluster Computing, 2016, 19, 849-864.	5.0	5
28	Scheduling for performance and availability in systems with temporal dependent workloads. , 2008, , .		3
29	SUGAR: Speeding Up GPGPU Application Resilience Estimation with Input Sizing. , 2021, , .		3
30	Lifespan and Failures of SSDs and HDDs: Similarities, Differences, and Prediction Models. IEEE Transactions on Dependable and Secure Computing, 2023, 20, 256-272.	5.4	3
31	New Results on the Performance Effects of Autocorrelated Flows in Systems. , 2007, , .		1
32	Versatile models of systems using map queueing networks. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	0
33	Data-centric Reliability Management in GPUs. , 2021, , .		0
34	Machine Learning for Reliability Analysis of Large Scale Systems. Lecture Notes in Computer Science, 2020, , 3-7.	1.3	0
35	SUGAR: Speeding Up GPGPU Application Resilience Estimation with Input Sizing. Performance Evaluation Review, 2021, 49, 45-46.	0.6	0