

# Ji-Dong Zhai

## List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/6059683/publications.pdf>

Version: 2024-02-01

47  
papers

691  
citations

777949

13  
h-index

759306

22  
g-index

48  
all docs

48  
docs citations

48  
times ranked

415  
citing authors

#	ARTICLE	IF	CITATIONS
1	Periodic Weather-Aware LSTM With Event Mechanism for Parking Behavior Prediction. IEEE Transactions on Knowledge and Data Engineering, 2022, 34, 5896-5909.	4.0	8
2	POCLib: A High-Performance Framework for Enabling Near Orthogonal Processing on Compression. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 459-475.	4.0	61
3	Exploring Data Analytics Without Decompression on Embedded GPU Systems. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 1553-1568.	4.0	18
4	Payment behavior prediction on shared parking lots with TR-GCN. VLDB Journal, 2022, 31, 1035-1058.	2.7	3
5	Leveraging Code Snippets to Detect Variations in the Performance of HPC Systems. IEEE Transactions on Parallel and Distributed Systems, 2022, , 1-1.	4.0	1
6	AStitch: enabling a new multi-dimensional optimization space for memory-intensive ML training and inference on modern SMT architectures. , 2022, , .		21
7	Zoro: A robotic middleware combining high performance and high reliability. Journal of Parallel and Distributed Computing, 2022, 166, 126-138.	2.7	3
8	Exploring Query Processing on CPU-GPU Integrated Edge Device. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 4057-4070.	4.0	7
9	Detecting Performance Variance for Parallel Applications Without Source Code. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 4239-4255.	4.0	1
10	CompressDB: Enabling Efficient Compressed Data Direct Processing for Various Databases. , 2022, , .		19
11	TADOC: Text analytics directly on compression. VLDB Journal, 2021, 30, 163-188.	2.7	37
12	G-TADOC: Enabling Efficient GPU-Based Text Analytics without Decompression. , 2021, , .		17
13	Collaborative Heterogeneity-Aware OS Scheduler for Asymmetric Multicore Processors. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 1224-1237.	4.0	12
14	An Efficient Parallel Secure Machine Learning Framework on GPUs. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 2262-2276.	4.0	22
15	AIPerf: Automated machine learning as an AI-HPC benchmark. Big Data Mining and Analytics, 2021, 4, 208-220.	7.5	11
16	Critique of "Planetary Normal Mode Computation: Parallel Algorithms, Performance, and Reproducibility" by SCC Team From Tsinghua University. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 2631-2634.	4.0	2
17	Critique of "MemXCT: memory-centric X-ray CT reconstruction with massive parallelization" by SCC Team from Tsinghua University. IEEE Transactions on Parallel and Distributed Systems, 2021, , 1-1.	4.0	0
18	Message Passing Optimization in Robot Operating System. International Journal of Parallel Programming, 2020, 48, 119-136.	1.1	14

#	ARTICLE	IF	CITATIONS
19	A Fast Lock for Explicit Message Passing Architectures. IEEE Transactions on Computers, 2020, , 1-1.	2.4	0
20	Enabling Efficient Random Access to Hierarchically-Compressed Data. , 2020, , .		14
21	PewLSTM: Periodic LSTM with Weather-Aware Gating Mechanism for Parking Behavior Prediction. , 2020, , .		16
22	ParSecureML: An Efficient Parallel Secure Machine Learning Framework on GPUs. , 2020, , .		4
23	Performance evaluation and analysis of sparse matrix and graph kernels on heterogeneous processors. CCF Transactions on High Performance Computing, 2019, 1, 131-143.	1.1	10
24	Automatic Irregularity-Aware Fine-Grained Workload Partitioning on Integrated Architectures. IEEE Transactions on Knowledge and Data Engineering, 2019, , 1-1.	4.0	19
25	Student Cluster Competition 2018, Team Tsinghua University: Reproducing performance of multi-physics simulations of the Tsunamigenic 2004 Sumatra megathrust earthquake on the Intel Skylake Architecture. Parallel Computing, 2019, 90, 102570.	1.3	0
26	Statistical Analysis and Prediction of Parking Behavior. Lecture Notes in Computer Science, 2019, , 93-104.	1.0	8
27	An Efficient In-Memory Checkpoint Method and its Practice on Fault-Tolerant HPL. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 758-771.	4.0	9
28	A vision of post-exascale programming. Frontiers of Information Technology and Electronic Engineering, 2018, 19, 1261-1266.	1.5	6
29	Zwift. , 2018, , .		20
30	Student cluster competition 2017, team Tsinghua University: Reproducing vectorization of the tersoff multi-body potential on the Intel Skylake and NVIDIA Volta architectures. Parallel Computing, 2018, 78, 47-53.	1.3	0
31	An adaptive breadth-first search algorithm on integrated architectures. Journal of Supercomputing, 2018, 74, 6135-6155.	2.4	17
32	BitFlow: Exploiting Vector Parallelism for Binary Neural Networks on CPU. , 2018, , .		18
33	vS <scp>ensor</scp>. , 2018, , .		2
34	Efficient document analytics on compressed data. Proceedings of the VLDB Endowment, 2018, 11, 1522-1535.	2.1	30
35	Understanding Co-Running Behaviors on Integrated CPU/GPU Architectures. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 905-918.	4.0	73
36	FinePar: Irregularity-aware fine-grained workload partitioning on integrated architectures. , 2017, , .		24

#	ARTICLE	IF	CITATIONS
37	Scalable Graph Traversal on Sunway TaihuLight with Ten Million Cores. , 2017, , .		27
38	Characterizing and optimizing TPC-C workloads on large-scale systems using SSD arrays. Science China Information Sciences, 2016, 59, 1.	2.7	0
39	Performance Prediction for Large-Scale Parallel Applications Using Representative Replay. IEEE Transactions on Computers, 2016, 65, 2184-2198.	2.4	19
40	Building Semi-Elastic Virtual Clusters for Cost-Effective HPC Cloud Resource Provisioning. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 1915-1928.	4.0	23
41	To Co-run, or Not to Co-run: A Performance Study on Integrated Architectures. , 2015, , .		8
42	Optimizing seam carving on multi-GPU systems for real-time content-aware image resizing. Journal of Supercomputing, 2015, 71, 3500-3524.	2.4	6
43	Automatic Cloud I/O Configurator for I/O Intensive Parallel Applications. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 3275-3288.	4.0	0
44	Optimizing Seam Carving on multi-GPU systems for real-time image resizing. , 2014, , .		0
45	Efficiently Acquiring Communication Traces for Large-Scale Parallel Applications. IEEE Transactions on Parallel and Distributed Systems, 2011, 22, 1862-1870.	4.0	13
46	PHANTOM. ACM SIGPLAN Notices, 2010, 45, 305-314.	0.2	49
47	LogGPO: An accurate communication model for performance prediction of MPI programs. Science in China Series F: Information Sciences, 2009, 52, 1785-1791.	1.1	18