

Bingsheng He, $\ddot{a}^{1/2} \cdot \ddot{a}, \grave{e} \ddot{o}$

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/1733777/publications.pdf>

Version: 2024-02-01

267
papers

6,424
citations

168829

31
h-index

190340

53
g-index

270
all docs

270
docs citations

270
times ranked

4188
citing authors

#	ARTICLE	IF	CITATIONS
1	Practical Vertical Federated Learning with Unsupervised Representation Learning. IEEE Transactions on Big Data, 2024, , 1-1.	4.4	5
2	A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. IEEE Transactions on Knowledge and Data Engineering, 2023, 35, 3347-3366.	4.0	209
3	Efficient Decomposition Selection for Multi-class Classification. IEEE Transactions on Knowledge and Data Engineering, 2023, 35, 3751-3764.	4.0	0
4	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
5	Taming System Dynamics on Resource Optimization for Data Processing Workflows: A Probabilistic Approach. IEEE Transactions on Parallel and Distributed Systems, 2022, 33, 231-248.	4.0	2
6	A Structure-Aware Storage Optimization for Out-of-Core Concurrent Graph Processing. IEEE Transactions on Computers, 2022, 71, 1612-1625.	2.4	4
7	Privacy-preserving workflow scheduling in geo-distributed data centers. Future Generation Computer Systems, 2022, 130, 46-58.	4.9	7
8	The Serverless Computing Survey: A Technical Primer for Design Architecture. ACM Computing Surveys, 2022, 54, 1-34.	16.1	34
9	Payment behavior prediction on shared parking lots with TR-GCN. VLDB Journal, 2022, 31, 1035-1058.	2.7	3
10	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
11	ThunderGP: Resource-Efficient Graph Processing Framework on FPGAs with HLS. ACM Transactions on Reconfigurable Technology and Systems, 2022, 15, 1-31.	1.9	4
12	An in-depth study of continuous subgraph matching. Proceedings of the VLDB Endowment, 2022, 15, 1403-1416.	2.1	4
13	The OARF Benchmark Suite: Characterization and Implications for Federated Learning Systems. ACM Transactions on Intelligent Systems and Technology, 2022, 13, 1-32.	2.9	8
14	Understanding and Optimizing Conjunctive Predicates Under Memory-Efficient Storage Layouts. IEEE Transactions on Knowledge and Data Engineering, 2021, 33, 2803-2817.	4.0	2
15	Optimizing in-memory database engine for AI-powered on-line decision augmentation using persistent memory. Proceedings of the VLDB Endowment, 2021, 14, 799-812.	2.1	11
16	ThunderGP. , 2021, , .		48
17	DepGraph: A Dependency-Driven Accelerator for Efficient Iterative Graph Processing. , 2021, , .		12
18	VColor*: a practical approach for coloring large graphs. Frontiers of Computer Science, 2021, 15, 1.	1.6	4

#	ARTICLE	IF	CITATIONS
19	HGP4CNN: an efficient parallelization framework for training convolutional neural networks on modern GPUs. Journal of Supercomputing, 2021, 77, 12741-12770.	2.4	0
20	MG-Join: A Scalable Join for Massively Parallel Multi-GPU Architectures. , 2021, , .		9
21	PathEnum: Towards Real-Time Hop-Constrained s-t Path Enumeration. , 2021, , .		11
22	GPU-Accelerated Graph Label Propagation for Real-Time Fraud Detection. , 2021, , .		7
23	Efficient Deep Learning Pipelines for Accurate Cost Estimations Over Large Scale Query Workload. , 2021, , .		9
24	Parallelizing Intra-Window Join on Multicores. , 2021, , .		5
25	iMLBench: A Machine Learning Benchmark Suite for CPU-GPU Integrated Architectures. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 1740-1752.	4.0	11
26	Fine-Grained Multi-Query Stream Processing on Integrated Architectures. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 2303-2320.	4.0	15
27	YuenyeungSpTRSV: A Thread-Level and Warp-Level Fusion Synchronization-Free Sparse Triangular Solve. IEEE Transactions on Parallel and Distributed Systems, 2021, 32, 2321-2337.	4.0	4
28	LargeGraph. Transactions on Architecture and Code Optimization, 2021, 18, 1-24.	1.6	3
29	Database Systems on GPUs. Foundations and Trends in Databases, 2021, 11, 1-108.	4.1	2
30	Gengar: An RDMA-based Distributed Hybrid Memory Pool. , 2021, , .		1
31	ThunderRW. Proceedings of the VLDB Endowment, 2021, 14, 1992-2005.	2.1	13
32	Skew-Oblivious Data Routing for Data Intensive Applications on FPGAs with HLS. , 2021, , .		3
33	Performance Modeling and Directives Optimization for High-Level Synthesis on FPGA. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2020, 39, 1428-1441.	1.9	25
34	Revisiting hash join on graphics processors: a decade later. Distributed and Parallel Databases, 2020, 38, 771-793.	1.0	5
35	Cost-Aware Partitioning for Efficient Large Graph Processing in Geo-Distributed Datacenters. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 1707-1723.	4.0	11
36	gMig: Efficient vGPU Live Migration with Overlapped Software-Based Dirty Page Verification. IEEE Transactions on Parallel and Distributed Systems, 2020, 31, 1209-1222.	4.0	4

#	ARTICLE	IF	CITATIONS
37	Accelerating Generative Neural Networks on Unmodified Deep Learning Processors - A Software Approach. IEEE Transactions on Computers, 2020, , 1-1.	2.4	6
38	A High-Performance Index for Real-Time Matrix Retrieval. IEEE Transactions on Knowledge and Data Engineering, 2020, , 1-1.	4.0	1
39	Towards Concurrent Stateful Stream Processing on Multicore Processors. , 2020, , .		6
40	Practical Federated Gradient Boosting Decision Trees. Proceedings of the AAAI Conference on Artificial Intelligence, 2020, 34, 4642-4649.	3.6	83
41	Privacy-Preserving Gradient Boosting Decision Trees. Proceedings of the AAAI Conference on Artificial Intelligence, 2020, 34, 784-791.	3.6	23
42	PA-Tree: Polled-Mode Asynchronous B+ Tree for NVMe. , 2020, , .		4
43	Maxson: Reduce Duplicate Parsing Overhead on Raw Data. , 2020, , .		3
44	Object-Level Memory Allocation and Migration in Hybrid Memory Systems. IEEE Transactions on Computers, 2020, 69, 1401-1413.	2.4	15
45	A Survey on Spark Ecosystem: Big Data Processing Infrastructure, Machine Learning, and Applications. IEEE Transactions on Knowledge and Data Engineering, 2020, , 1-1.	4.0	23
46	GPU-Accelerated Subgraph Enumeration on Partitioned Graphs. , 2020, , .		29
47	Hardware-Conscious Stream Processing. SIGMOD Record, 2020, 48, 18-29.	0.7	19
48	CapelliniSpTRSV: A Thread-Level Synchronization-Free Sparse Triangular Solve on GPUs. , 2020, , .		7
49	AsynGraph. Transactions on Architecture and Code Optimization, 2020, 17, 1-21.	1.6	9
50	G ³. Proceedings of the VLDB Endowment, 2020, 13, 2813-2816.	2.1	23
51	RapidMatch. Proceedings of the VLDB Endowment, 2020, 14, 176-188.	2.1	22
52	Improving execution efficiency of just-in-time compilation based query processing on GPUs. Proceedings of the VLDB Endowment, 2020, 14, 202-214.	2.1	11
53	Accelerating exact constrained shortest paths on GPUs. Proceedings of the VLDB Endowment, 2020, 14, 547-559.	2.1	10
54	GradSA: Gradient Sparsification and Accumulation for Communication-Efficient Distributed Deep Learning. Lecture Notes in Computer Science, 2020, , 77-91.	1.0	1

#	ARTICLE	IF	CITATIONS
55	ByteSeries. , 2020, , .		8
56	Poet. , 2020, , .		2
57	Energy Efficient In-memory Integer Multiplication Based on Racetrack Memory. , 2020, , .		0
58	Efficient Multi-Class Probabilistic SVMs on GPUs. IEEE Transactions on Knowledge and Data Engineering, 2019, 31, 1693-1706.	4.0	10
59	BriskStream. , 2019, , .		31
60	Aucher: Multi-modal Queries on Live Audio Streams in Real-Time. , 2019, , .		2
61	Exploiting GPUs for Efficient Gradient Boosting Decision Tree Training. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 2706-2717.	4.0	32
62	Automatic Irregularity-Aware Fine-Grained Workload Partitioning on Integrated Architectures. IEEE Transactions on Knowledge and Data Engineering, 2019, , 1-1.	4.0	19
63	Incorporating Probabilistic Optimizations for Resource Provisioning of Data Processing Workflows. , 2019, , .		2
64	Privacy Regulation Aware Process Mapping in Geo-Distributed Cloud Data Centers. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 1872-1888.	4.0	23
65	Supporting Superpages and Lightweight Page Migration in Hybrid Memory Systems. Transactions on Architecture and Code Optimization, 2019, 16, 1-26.	1.6	19
66	Efficient Multi-Class Probabilistic SVMs on GPUs. , 2019, , .		0
67	DiGraph. , 2019, , .		30
68	CGraph. ACM Transactions on Storage, 2019, 15, 1-26.	1.4	11
69	Guest Editorsâ€™ Introduction: Special Issue on Big Data Systems on Emerging Architectures. IEEE Transactions on Big Data, 2019, 5, 2-3.	4.4	2
70	A Survey on Graph Processing Accelerators: Challenges and Opportunities. Journal of Computer Science and Technology, 2019, 34, 339-371.	0.9	53
71	On-The-Fly Parallel Data Shuffling for Graph Processing on OpenCL-Based FPGAs. , 2019, , .		16
72	TraV: An Interactive Exploration System for Massive Trajectory Data. , 2019, , .		3

#	ARTICLE	IF	CITATIONS
73	OBFS: OpenCL Based BFS Optimizations on Software Programmable FPGAs. , 2019, , .		5
74	Adaptive Kernel Value Caching for SVM Training. IEEE Transactions on Neural Networks and Learning Systems, 2019, 31, 1-11.	7.2	9
75	GraphM. , 2019, , .		17
76	Revisiting Hash Join on Graphics Processors: A Decade Later. , 2019, , .		2
77	Towards Declarative and Data-Centric Virtual Machine Image Management in IaaS Clouds. IEEE Transactions on Cloud Computing, 2019, 7, 1124-1138.	3.1	4
78	An Adaptive Efficiency-Fairness Meta-Scheduler for Data-Intensive Computing. IEEE Transactions on Services Computing, 2019, 12, 865-879.	3.2	8
79	Fairness-Efficiency Allocation of CPU-GPU Heterogeneous Resources. IEEE Transactions on Services Computing, 2019, 12, 474-488.	3.2	11
80	Search and Query Accelerators. , 2019, , 1470-1475.		0
81	Big Data and Exascale Computing. , 2019, , 184-187.		1
82	Deploying Hash Tables on Die-Stacked High Bandwidth Memory. , 2019, , .		5
83	Towards Efficient Resource Allocation for Heterogeneous Workloads in IaaS Clouds. IEEE Transactions on Cloud Computing, 2018, 6, 264-275.	3.1	56
84	Fair Resource Allocation for Data-Intensive Computing in the Cloud. IEEE Transactions on Services Computing, 2018, 11, 20-33.	3.2	34
85	Big Data and Exascale Computing. , 2018, , 1-4.		0
86	Scalable GPU Virtualization with Dynamic Sharing of Graphics Memory Space. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 1823-1836.	4.0	14
87	Frog: Asynchronous Graph Processing on GPU with Hybrid Coloring Model. IEEE Transactions on Knowledge and Data Engineering, 2018, 30, 29-42.	4.0	29
88	Long-Term Multi-Resource Fairness for Pay-as-you Use Computing Systems. IEEE Transactions on Parallel and Distributed Systems, 2018, 29, 1147-1160.	4.0	16
89	gMig. , 2018, , .		10
90	JouleMR: Towards Cost-Effective and Green-Aware Data Processing Frameworks. IEEE Transactions on Big Data, 2018, 4, 258-272.	4.4	15

#	ARTICLE	IF	CITATIONS
91	Many-core needs fine-grained scheduling: A case study of query processing on Intel Xeon Phi processors. <i>Journal of Parallel and Distributed Computing</i> , 2018, 120, 395-404.	2.7	4
92	Efficient Disk-Based Directed Graph Processing: A Strongly Connected Component Approach. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2018, 29, 830-842.	4.0	15
93	Query Processing on OpenCL-Based FPGAs: Challenges and Opportunities. , 2018, , .		0
94	RTSI: An Index Structure for Multi-Modal Real-Time Search on Live Audio Streaming Services. , 2018, , .		6
95	Towards concurrency race debugging. , 2018, , .		3
96	Hebe: An Order-Oblivious and High-Performance Execution Scheme for Conjunctive Predicates. , 2018, , .		2
97	FCN-engine. , 2018, , .		19
98	An efficient graph accelerator with parallel data conflict management. , 2018, , .		30
99	Layer-Centric Memory Reuse and Data Migration for Extreme-Scale Deep Learning on Many-Core Architectures. <i>Transactions on Architecture and Code Optimization</i> , 2018, 15, 1-26.	1.6	18
100	GLP4NN. , 2018, , .		4
101	Energy-Efficient Speculative Execution using Advanced Reservation for Heterogeneous Clusters. , 2018, , .		8
102	Efficient Gradient Boosted Decision Tree Training on GPUs. , 2018, , .		26
103	A Stack-Centric Processing Model For Iterative Processing. <i>IEEE Transactions on Big Data</i> , 2018, , 1-1.	4.4	0
104	vS <sc>ensor</sc>. , 2018, , .		2
105	gMig. <i>ACM SIGPLAN Notices</i> , 2018, 53, 31-44.	0.2	3
106	vS <sc>ensor</sc>. <i>ACM SIGPLAN Notices</i> , 2018, 53, 124-136.	0.2	4
107	A Hybrid Logic Block Architecture in FPGA for Holistic Efficiency. <i>IEEE Transactions on Circuits and Systems II: Express Briefs</i> , 2017, 64, 71-75.	2.2	3
108	Understanding Co-Running Behaviors on Integrated CPU/GPU Architectures. <i>IEEE Transactions on Parallel and Distributed Systems</i> , 2017, 28, 905-918.	4.0	73

#	ARTICLE	IF	CITATIONS
109	Multikernel Data Partitioning With Channel on OpenCL-Based FPGAs. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2017, 25, 1906-1918.	2.1	9
110	FinePar: Irregularity-aware fine-grained workload partitioning on integrated architectures. , 2017, , .		24
111	Revisiting the Design of Data Stream Processing Systems on Multi-Core Processors. , 2017, , .		31
112	Multi-Query Optimization for Complex Event Processing in SAP ESP. , 2017, , .		15
113	A distributed in-memory key-value store system on heterogeneous CPU-GPU cluster. VLDB Journal, 2017, 26, 729-750.	2.7	6
114	Hardware/software cooperative caching for hybrid DRAM/NVM memory architectures. , 2017, , .		39
115	On Achieving Efficient Data Transfer for Graph Processing in Geo-Distributed Datacenters. , 2017, , .		23
116	AdaStorm: Resource Efficient Storm with Adaptive Configuration. , 2017, , .		9
117	Dynamic Module Partitioning for Library Based Placement on Heterogeneous FPGAs. , 2017, , .		1
118	DIDO: Dynamic Pipelines for In-Memory Key-Value Stores on Coupled CPU-GPU Architectures. , 2017, , .		10
119	Data Management Systems on Future Hardware: Challenges and Opportunities. , 2017, , .		0
120	A Declarative Optimization Engine for Resource Provisioning of Scientific Workflows in Geo-Distributed Clouds. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 647-661.	4.0	13
121	QoS-Aware Resource Allocation for Video Transcoding in Clouds. IEEE Transactions on Circuits and Systems for Video Technology, 2017, 27, 49-61.	5.6	56
122	A Variation-Aware Adaptive Fuzzy Control System for Thermal Management of Microprocessors. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2017, 25, 683-695.	2.1	3
123	Building an Efficient Put-Intensive Key-Value Store with Skip-Tree. IEEE Transactions on Parallel and Distributed Systems, 2017, 28, 961-973.	4.0	29
124	Network Performance Aware Optimizations on IaaS Clouds. IEEE Transactions on Computers, 2017, 66, 672-687.	2.4	4
125	COMBA: A comprehensive model-based analysis framework for high level synthesis of real applications. , 2017, , .		67
126	Efficient process mapping in geo-distributed cloud data centers. , 2017, , .		10

#	ARTICLE	IF	CITATIONS
127	Dynamic module partitioning for library based placement on heterogeneous FPGAs. , 2017, , .		2
128	A novel two-stage modular multiplier based on racetrack memory for asymmetric cryptography. , 2017, , .		6
129	A Study of Main-Memory Hash Joins on Many-core Processor. , 2017, , .		15
130	Accelerating dynamic graph analytics on GPUs. Proceedings of the VLDB Endowment, 2017, 11, 107-120.	2.1	54
131	Multi-objective Optimizations in Geo-Distributed Data Analytics Systems. , 2017, , .		4
132	VColor: A practical vertex-cut based approach for coloring large graphs. , 2016, , .		10
133	Elastic Multi-resource Fairness: Balancing Fairness and Efficiency in Coupled CPU-GPU Architectures. , 2016, , .		11
134	GPL. , 2016, , .		53
135	A performance analysis framework for optimizing OpenCL applications on FPGAs. , 2016, , .		59
136	Relational query processing on OpenCL-based FPGAs. , 2016, , .		27
137	A Study of Sorting Algorithms on Approximate Memory. , 2016, , .		3
138	Efficient Query Processing on Many-core Architectures. , 2016, , .		5
139	A Study of Big Data Computing Platforms: Fairness and Energy Consumption. , 2016, , .		3
140	Modular Placement for Interposer based Multi-FPGA Systems. , 2016, , .		11
141	Not All Joules are Equal: Towards Energy-Efficient and Green-Aware Data Processing Frameworks. , 2016, , .		9
142	Library-Based Placement and Routing in FPGAs with Support of Partial Reconfiguration. ACM Transactions on Design Automation of Electronic Systems, 2016, 21, 1-26.	1.9	3
143	Thermal-Aware Task Scheduling for 3D-Network-on-Chip: A Bottom to Top Scheme. Journal of Circuits, Systems and Computers, 2016, 25, 1640003.	1.0	8
144	Analysis of Minimum Interaction Time for Continuous Distributed Interactive Computing. IEEE Transactions on Parallel and Distributed Systems, 2016, , 1-1.	4.0	7

#	ARTICLE	IF	CITATIONS
145	Monetary Cost Optimizations for Hosting Workflow-as-a-Service in IaaS Clouds. IEEE Transactions on Cloud Computing, 2016, 4, 34-48.	3.1	67
146	NV-Tree: A Consistent and Workload-Adaptive Tree Structure for Non-Volatile Memory. IEEE Transactions on Computers, 2016, 65, 2169-2183.	2.4	20
147	Melia: A MapReduce Framework on OpenCL-Based FPGAs. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 3547-3560.	4.0	37
148	Decentralized Thermal-Aware Task Scheduling for Large-Scale Many-Core Systems. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2016, 24, 2075-2088.	2.1	5
149	A Performance Debugging Framework for Unnecessary Lock Contentions with Record/Replay Techniques. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 1889-1901.	4.0	2
150	A racetrack memory based in-memory booth multiplier for cryptography application. , 2016, , .		10
151	F2C: Enabling Fair and Fine-Grained Resource Sharing in Multi-Tenant IaaS Clouds. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 2589-2602.	4.0	20
152	Rank-Aware Dynamic Migrations and Adaptive Demotions for DRAM Power Management. IEEE Transactions on Computers, 2016, 65, 187-202.	2.4	13
153	Dynamic Job Ordering and Slot Configurations for MapReduce Workloads. IEEE Transactions on Services Computing, 2016, 9, 4-17.	3.2	26
154	Rotated Logging Storage Architectures for Data Centers: Models and Optimizations. IEEE Transactions on Computers, 2016, 65, 203-215.	2.4	4
155	Access Control in Cloud Computing. , 2016, , 1914-1936.		0
156	A Discrete Thermal Controller for Chip-Multiprocessors. , 2016, , .		0
157	Accelerating Database Query Processing on OpenCL-based FPGAs (Abstract Only). , 2016, , .		3
158	Hierarchical Library Based Power Estimator for Versatile FPGAs. , 2015, , .		5
159	Monetary cost optimizations for MPI-based HPC applications on Amazon clouds. , 2015, , .		22
160	Optimization of asynchronous graph processing on GPU with hybrid coloring model. ACM SIGPLAN Notices, 2015, 50, 271-272.	0.2	1
161	MrPhi: An Optimized MapReduce Framework on Intel Xeon Phi Coprocessors. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 3066-3078.	4.0	23
162	Fast Subgraph Matching on Large Graphs using Graphics Processors. Lecture Notes in Computer Science, 2015, , 299-315.	1.0	39

#	ARTICLE	IF	CITATIONS
163	Gemini: An Adaptive Performance-Fairness Scheduler for Data-Intensive Cluster Computing. , 2015, , .		13
164	Guest Editors' Introduction: Special Issue on Economics and Market Mechanisms for Cloud Computing. IEEE Transactions on Cloud Computing, 2015, 3, 245-246.	3.1	2
165	A study of data partitioning on OpenCL-based FPGAs. , 2015, , .		30
166	To Co-run, or Not to Co-run: A Performance Study on Integrated Architectures. , 2015, , .		8
167	Improving main memory hash joins on Intel Xeon Phi processors. Proceedings of the VLDB Endowment, 2015, 8, 642-653.	2.1	59
168	Improving Update-Intensive Workloads on Flash Disks through Exploiting Multi-Chip Parallelism. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 152-162.	4.0	16
169	Sensor Placement and Measurement of Wind for Water Quality Studies in Urban Reservoirs. ACM Transactions on Sensor Networks, 2015, 11, 1-27.	2.3	33
170	Understanding the Behavior of Solid State Disk. Proceedings in Adaptation, Learning and Optimization, 2015, , 341-355.	1.5	2
171	Willow: Saving Data Center Network Energy for Network-Limited Flows. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 2610-2620.	4.0	18
172	Hotplug or Ballooning: A Comparative Study on Dynamic Memory Management Techniques for Virtual Machines. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 1350-1363.	4.0	33
173	PCMLogging: Optimizing Transaction Logging and Recovery Performance with PCM. IEEE Transactions on Knowledge and Data Engineering, 2015, 27, 3332-3346.	4.0	17
174	Synergy of Dynamic Frequency Scaling and Demotion on DRAM Power Management: Models and Optimizations. IEEE Transactions on Computers, 2015, 64, 2367-2381.	2.4	7
175	Optimization of asynchronous graph processing on GPU with hybrid coloring model. , 2015, , .		13
176	A Declarative Optimization Engine for Resource Provisioning of Scientific Workflows in IaaS Clouds. , 2015, , .		14
177	Energy-Efficient Query Processing on Embedded CPU-GPU Architectures. , 2015, , .		4
178	Improving Data Partitioning Performance on OpenCL-Based FPGAs. , 2015, , .		3
179	Real-Time In-Memory Checkpointing for Future Hybrid Memory Systems. , 2015, , .		24
180	On performance debugging of unnecessary lock contentions on multicore processors: A replay-based approach. , 2015, , .		7

#	ARTICLE	IF	CITATIONS
181	Network Performance Aware MPI Collective Communication Operations in the Cloud. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 3079-3089.	4.0	40
182	A Combined SDC-SDF Architecture for Normal I/O Pipelined Radix-2 FFT. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2015, 23, 973-977.	2.1	53
183	VMbuddies: Coordinating Live Migration of Multi-Tier Applications in Cloud Environments. IEEE Transactions on Parallel and Distributed Systems, 2015, 26, 1192-1205.	4.0	68
184	Access Control in Cloud Computing. Advances in Systems Analysis, Software Engineering, and High Performance Computing Book Series, 2015, , 340-362.	0.5	6
185	Towards Economic Fairness for Big Data Processing in Pay-as-You-Go Cloud Computing. , 2014, , .		13
186	Long-term resource fairness. , 2014, , .		24
187	Medusa. SIGMOD Record, 2014, 43, 35-40.	0.7	27
188	In-cache query co-processing on coupled CPU-GPU architectures. Proceedings of the VLDB Endowment, 2014, 8, 329-340.	2.1	58
189	Towards automatic partial reconfiguration in FPGAs. , 2014, , .		3
190	Demo Abstract: Wind measurements for water quality studies in urban reservoirs. , 2014, , .		1
191	Reciprocal Resource Fairness: Towards Cooperative Multiple-Resource Fair Sharing in IaaS Clouds. , 2014, , .		30
192	Finding Constant from Change: Revisiting Network Performance Aware Optimizations on IaaS Clouds. , 2014, , .		16
193	Simplified Resource Provisioning for Workflows in IaaS Clouds. , 2014, , .		6
194	A novel authenticated multi-party key agreement for private cloud. , 2014, , .		1
195	Thermal-aware task scheduling for 3D-network-on-chip: A Bottom-to-Top scheme. , 2014, , .		2
196	Towards multi-resource physical machine provisioning for IaaS clouds. , 2014, , .		7
197	Medusa: Simplified Graph Processing on GPUs. IEEE Transactions on Parallel and Distributed Systems, 2014, 25, 1543-1552.	4.0	186
198	Transformation-Based Monetary Cost Optimizations for Workflows in the Cloud. IEEE Transactions on Cloud Computing, 2014, 2, 85-98.	3.1	48

#	ARTICLE	IF	CITATIONS
199	FD-Buffer: A Cost-Based Adaptive Buffer Replacement Algorithm for FlashMemory Devices. IEEE Transactions on Computers, 2014, 63, 2288-2301.	2.4	12
200	Optimal sensor placement and measurement of wind for water quality studies in urban reservoirs. , 2014, , .		38
201	Pipelined Compaction for the LSM-Tree. , 2014, , .		23
202	DynamicMR: A Dynamic Slot Allocation Optimization Framework for MapReduce Clusters. IEEE Transactions on Cloud Computing, 2014, 2, 333-347.	3.1	57
203	A Survey of Resource Management in Multi-Tier Web Applications. IEEE Communications Surveys and Tutorials, 2014, 16, 1574-1590.	24.8	64
204	Kernelet: High-Throughput GPU Kernel Executions with Dynamic Slicing and Scheduling. IEEE Transactions on Parallel and Distributed Systems, 2014, 25, 1522-1532.	4.0	86
205	When data management systems meet approximate hardware. Proceedings of the VLDB Endowment, 2014, 7, 877-880.	2.1	3
206	GPU-Accelerated Cloud Computing for Data-Intensive Applications. , 2014, , 105-129.		4
207	MROrder: Flexible Job Ordering Optimization for Online MapReduce Workloads. Lecture Notes in Computer Science, 2013, , 291-304.	1.0	8
208	Handling partitioning skew in MapReduce using LEEN. Peer-to-Peer Networking and Applications, 2013, 6, 409-424.	2.6	40
209	Dynamic slot allocation technique for MapReduce clusters. , 2013, , .		17
210	A Framework for Analyzing Monetary Cost of Database Systems in the Cloud. Lecture Notes in Computer Science, 2013, , 118-129.	1.0	1
211	Optimizing the MapReduce framework on Intel Xeon Phi coprocessor. , 2013, , .		34
212	Revisiting co-processing for hash joins on the coupled CPU-GPU architecture. Proceedings of the VLDB Endowment, 2013, 6, 889-900.	2.1	87
213	Simulation of Information Propagation over Complex Networks: Performance Studies on Multi-GPU. , 2013, , .		2
214	Parallel graph processing on graphics processors made easy. Proceedings of the VLDB Endowment, 2013, 6, 1270-1273.	2.1	12
215	Towards GPU-Accelerated Large-Scale Graph Processing in the Cloud. , 2013, , .		11
216	OmniDB. Proceedings of the VLDB Endowment, 2013, 6, 1374-1377.	2.1	49

#	ARTICLE	IF	CITATIONS
217	Simulation studies of viral advertisement diffusion on multi-GPU. , 2013, , .		1
218	Spectral Decomposition for Optimal Graph Index Prediction. Lecture Notes in Computer Science, 2013, , 187-200.	1.0	2
219	An overview of Medusa. , 2012, , .		7
220	An overview of Medusa. ACM SIGPLAN Notices, 2012, 47, 283-284.	0.2	5
221	Flag Commit: Supporting Efficient Transaction Recovery in Flash-Based DBMSs. IEEE Transactions on Knowledge and Data Engineering, 2012, 24, 1624-1639.	4.0	18
222	A Map-Reduce Based Framework for Heterogeneous Processing Element Cluster Environments. , 2012, , .		10
223	An overview of CMPI. ACM SIGPLAN Notices, 2012, 47, 297-298.	0.2	1
224	GPGPU for real-time data analytics. , 2012, , .		0
225	RAMZzz: Rank-aware DRAM power management with dynamic migrations and demotions. , 2012, , .		14
226	HPC Simulations of Information Propagation Over Social Networks. Procedia Computer Science, 2012, 9, 292-301.	1.2	11
227	An overview of CMPI. , 2012, , .		7
228	Improving large graph processing on partitioned graphs in the cloud. , 2012, , .		67
229	Speedup for Multi-Level Parallel Computing. , 2012, , .		6
230	Green-aware workload scheduling in geographically distributed data centers. , 2012, , .		59
231	Maestro: Replica-Aware Map Scheduling for MapReduce. , 2012, , .		70
232	Adaptive Disk I/O Scheduling for MapReduce in Virtualized Environment. , 2011, , .		35
233	Mars: Accelerating MapReduce with Graphics Processors. IEEE Transactions on Parallel and Distributed Systems, 2011, 22, 608-620.	4.0	127
234	Towards Pay-As-You-Consume Cloud Computing. , 2011, , .		88

#	ARTICLE	IF	CITATIONS
235	GPU-Assisted Buffer Management. <i>Procedia Computer Science</i> , 2011, 4, 362-371.	1.2	1
236	High-throughput transaction executions on graphics processors. <i>Proceedings of the VLDB Endowment</i> , 2011, 4, 314-325.	2.1	46
237	Operation-aware buffer management in flash-based systems. , 2011, , .		31
238	PCMLogging. , 2011, , .		34
239	GViewer: GPU-Accelerated Graph Visualization and Mining. <i>Lecture Notes in Computer Science</i> , 2011, , 304-307.	1.0	2
240	Database compression on graphics processors. <i>Proceedings of the VLDB Endowment</i> , 2010, 3, 670-680.	2.1	87
241	Tree indexing on solid state drives. <i>Proceedings of the VLDB Endowment</i> , 2010, 3, 1195-1206.	2.1	126
242	Comet. , 2010, , .		114
243	FD-buffer. , 2010, , .		9
244	Large graph processing in the cloud. , 2010, , .		75
245	Supporting extended precision on graphics processors. , 2010, , .		38
246	LEEN: Locality/Fairness-Aware Key Partitioning for MapReduce in the Cloud. , 2010, , .		128
247	Stack-based parallel recursion on graphics processors. , 2009, , .		6
248	Tree Indexing on Flash Disks. <i>Proceedings - International Conference on Data Engineering</i> , 2009, , .	0.0	62
249	Relational query coprocessing on graphics processors. <i>ACM Transactions on Database Systems</i> , 2009, 34, 1-39.	1.5	170
250	Stack-based parallel recursion on graphics processors. <i>ACM SIGPLAN Notices</i> , 2009, 44, 299-300.	0.2	0
251	Frequent itemset mining on graphics processors. , 2009, , .		92
252	A Uniform Framework for Ad-Hoc Indexes to Answer Reachability Queries on Large Graphs. <i>Lecture Notes in Computer Science</i> , 2009, , 138-152.	1.0	10

#	ARTICLE	IF	CITATIONS
253	Relational joins on graphics processors. , 2008, , .		229
254	Cache-oblivious databases. ACM Transactions on Database Systems, 2008, 33, 1-42.	1.5	24
255	In-memory grid files on graphics processors. , 2007, , .		9
256	Adaptive Index Utilization in Memory-Resident Structural Joins. IEEE Transactions on Knowledge and Data Engineering, 2007, 19, 772-788.	4.0	2
257	Efficient gather and scatter operations on graphics processors. , 2007, , .		81
258	A general framework for improving query processing performance on multi-level memory hierarchies. , 2007, , .		0
259	GPUQP. , 2007, , .		46
260	EaseDB. , 2007, , .		3
261	Cache-Conscious Automata for XML Filtering. IEEE Transactions on Knowledge and Data Engineering, 2006, 18, 1629-1644.	4.0	15
262	Cache-oblivious nested-loop joins. , 2006, , .		9
263	A Quantitative Summary of XML Structures. Lecture Notes in Computer Science, 2006, , 228-240.	1.0	5
264	Accurate Emulation of Wireless Sensor Networks. Lecture Notes in Computer Science, 2004, , 576-583.	1.0	7
265	The HKUST Frog Pond â€“ A Case Study of Sensory Data Analysis. Lecture Notes in Computer Science, 2004, , 551-558.	1.0	0
266	MEADOWS. , 2004, , .		8
267	Cache-Conscious Automata for XML Filtering. , 0, , .		3