Joel S Emer

List of Publications by Year
in descending order

96 papers	11,924 citations	279798 23 h-index	302126 39 g-index
96 all docs	96 docs citations	96 times ranked	6128 citing authors

#	Article	IF	Citations
1	Freely scalable and reconfigurable optical hardware for deep learning. Scientific Reports, 2021, 11, 3144.	3.3	32
2	Architecture-Level Energy Estimation for Heterogeneous Computing Systems. , 2021, , .		2
3	Sparseloop: An Analytical, Energy-Focused Design Space Exploration Methodology for Sparse Tensor Accelerators. , 2021, , .		11
4	Gamma: leveraging Gustavsonâ€™s algorithm to accelerate sparse matrix multiplication. , 2021, , .		52
5	SpZip: Architectural Support for Effective Data Compression In Irregular Applications. , 2021, , .		14
6	How to Evaluate Deep Neural Network Processors: TOPS/W (Alone) Considered Harmful. IEEE Solid-State Circuits Magazine, 2020, 12, 28-41.	0.4	40
7	Thereâ€™s plenty of room at the Top: What will drive computer performance after Mooreâ€™s law?. Science, 2020, 368, .	12.6	171
8	Efficient Processing of Deep Neural Networks. Synthesis Lectures on Computer Architecture, 2020, 15, 1-341.	1.3	72
9	A 0.32â€“128 TOPS, Scalable Multi-Chip-Module-Based Deep Neural Network Inference Accelerator With Ground-Referenced Signaling in 16 nm. IEEE Journal of Solid-State Circuits, 2020, 55, 920-932.	5.4	57
10	ExTensor. , 2019, , .		121
11	Buffets. , 2019, , .		47
12	Timeloop: A Systematic Approach to DNN Accelerator Evaluation. , 2019, , .		251
13	Eyeriss v2: A Flexible Accelerator for Emerging Deep Neural Networks on Mobile Devices. IEEE Journal on Emerging and Selected Topics in Circuits and Systems, 2019, 9, 292-308.	3.6	609
14	DAWG: A Defense Against Cache Timing Attacks in Speculative Execution Processors. , 2018, , .		151
15	A modular digital VLSI flow for high-productivity SoC design. , 2018, , .		14
16	Automatic Construction of Program-Optimized FPGA Memory Networks. , 2017, , .		5
17	(FPL 2015) Scavenger. ACM Transactions on Reconfigurable Technology and Systems, 2017, 10, 1-23.	2.5	3
18	Using Dataflow to Optimize Energy Efficiency of Deep Neural Network Accelerators. IEEE Micro, 2017, 37, 12-21.	1.8	105

#	Article	IF	Citations
19	Efficient Processing of Deep Neural Networks: A Tutorial and Survey. Proceedings of the IEEE, 2017, 105, 2295-2329.	21.3	2,217
20	Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks. IEEE Journal of Solid-State Circuits, 2017, 52, 127-138.	5.4	1,877
21	SAM: Optimizing Multithreaded Cores for Speculative Parallelism. , 2017, , .		5
22	SCNN. , 2017, , .		550
23	Towards closing the energy gap between HOG and CNN features for embedded vision. , 2017, , .		35
24	SCNN. Computer Architecture News, 2017, 45, 27-40.	2.5	241
25	Unlocking Ordered Parallelism with the Swarm Architecture. IEEE Micro, 2016, 36, 105-117.	1.8	18
26	Eyeriss: A Spatial Architecture for Energy-Efficient Dataflow for Convolutional Neural Networks. , 2016, , .		258
27	LMC. , 2016, , .		7
28	Eyeriss. Computer Architecture News, 2016, 44, 367-379.	2.5	833
29	Scavenger: Automating the construction of application-optimized memory hierarchies. , 2015, , .		10
30	High performing cache hierarchies for server workloads: Relaxing inclusion to capture the latency benefits of exclusive caches. , 2015, , .		23
31	Efficient Control and Communication Paradigms for Coarse-Grained Spatial Architectures. ACM Transactions on Computer Systems, 2015, 33, 1-32.	0.8	56
32	The LEAP FPGA operating system. , 2014, , .		31
33	Exploiting spatial architectures for edit distance algorithms. , 2014, , .		11
34	LEAP Shared Memories: Automating the Construction of FPGA Coherent Memories. , 2014, , .		14
35	Efficient Spatial Processing Element Control via Triggered Instructions. IEEE Micro, 2014, 34, 120-137.	1.8	58
36	A Hierarchical Architectural Framework for Reconfigurable Logic Computing. , 2013, , .		0

#	Article	IF	Citations
37	Using in-flight chains to build a scalable cache coherence protocol. Transactions on Architecture and Code Optimization, 2013, 10, 1-24.	2.0	3
38	Triggered instructions. Computer Architecture News, 2013, 41, 142-153.	2.5	30
39	Triggered instructions. , 2013, , .		76
40	Optimizing under abstraction: Using prefetching to improve FPGA performance. , 2013, , .		7
41	The gradient-based cache partitioning algorithm. Transactions on Architecture and Code Optimization, 2012, 8, 1-21.	2.0	11
42	CRUISE. Computer Architecture News, 2012, 40, 249-260.	2.5	17
43	Leveraging latency-insensitivity to ease multiple FPGA design. , 2012, , .		25
44	ZIP-IO: Architecture for application-specific compression of Big Data. , 2012, , .		9
45	CRUISE. , 2012, , .		50
46	Scheduling heterogeneous multi-cores through performance impact estimation (PIE). , 2012, , .		120
47	Scheduling heterogeneous multi-cores through Performance Impact Estimation (PIE). Computer Architecture News, 2012, 40, 213-224.	2.5	89
48	CRUISE. ACM SIGPLAN Notices, 2012, 47, 249-260.	0.2	5
49	HAsim: FPGA-based high-detail multicore simulation using time-division multiplexing. , 2011, , .		64
50	PACMan. , 2011, , .		84
51	Leap scratchpads. , 2011, , .		74
52	SHiP. , 2011, , .		181
53	DEC Alpha. , 2011, , 535-545.		0
54	The Future of Architectural Simulation. IEEE Micro, 2010, 30, 8-18.	1.8	8

#	Article	IF	Citations
55	Design contest overview: Combined architecture for network stream categorization and intrusion detection (CANSCID). , 2010, , .		10
56	High performance cache replacement using re-reference interval prediction (RRIP). Computer Architecture News, 2010, 38, 60-71.	2.5	176
57	High performance cache replacement using re-reference interval prediction (RRIP). , 2010, , .		404
58	Achieving Non-Inclusive Cache Performance with Inclusive Caches: Temporal Locality Aware (TLA) Cache Management Policies. , 2010, , .		93
59	Accelerating Simulation with FPGAs. , 2010, , 107-126.		0
60	Accelerating architecture research. , 2009, , .		3
61	Soft connections. , 2009, , .		8
62	A-Port Networks. ACM Transactions on Reconfigurable Technology and Systems, 2009, 2, 1-26.	2.5	15
63	CAMP: A technique to estimate per-structure power at run-time using a few simple parameters. , 2009, , .		59
64	Top Picks from the 2008 Computer Architecture Conferences. IEEE Micro, 2009, 29, 6-9.	1.8	0
65	Computing Accurate AVFs using ACE Analysis on Performance Models: A Rebuttal. IEEE Computer Architecture Letters, 2008, 7, 21-24.	1.5	36
66	A-Ports. , 2008, , .		26
67	Set-Dueling-Controlled Adaptive Insertion for High-Performance Caching. IEEE Micro, 2008, 28, 91-98.	1.8	17
68	Quick Performance Models Quickly: Closely-Coupled Partitioned Simulation on FPGAs. , 2008, , .		26
69	Adaptive insertion policies for managing shared caches. , 2008, , .		239
70	Adaptive insertion policies for high performance caching. , 2007, , .		415
71	Adaptive insertion policies for high performance caching. Computer Architecture News, 2007, 35, 381-391.	2.5	129
72	Late-binding. Computer Architecture News, 2007, 35, 347-357.	2.5	3

#	Article	IF	Citations
73	Single-Threaded vs. Multithreaded: Where Should We Focus?. IEEE Micro, 2007, 27, 14-24.	1.8	8
74	Single-threaded vs. multi-threaded. IEEE Micro, 2007, 27, x6.	1.8	0
75	Computing Architectural Vulnerability Factors for Address-Based Structures. Computer Architecture News, 2005, 33, 532-543.	2.5	77
76	Techniques to Reduce the Soft Error Rate of a High-Performance Microprocessor. Computer Architecture News, 2004, 32, 264.	2.5	122
77	Performance Potential of Effective Address Prediction of Load Instructions. , 2004, , 227-246.		1
78	Tarantula. Computer Architecture News, 2002, 30, 281-292.	2.5	28
79	A comparative study of arbitration algorithms for the Alpha 21364 pipelined router. Computer Architecture News, 2002, 30, 223-234.	2.5	10
80	Reducing cache misses using hardware and software page placement. , 1999, , .		76
81	A characterization of processor performance in the VAX-11/780. , 1998, , .		30
82	Retrospective: characterization of processor performance in the VAX-11/780. , 1998, , .		2
83	Memory dependence prediction using store sets. Computer Architecture News, 1998, 26, 142-153.	2.5	23
84	A language for describing predictors and its application to automatic synthesis. , 1997, , .		20
85	Converting thread-level parallelism to instruction-level parallelism via simultaneous multithreading. ACM Transactions on Computer Systems, 1997, 15, 322-354.	0.8	161
86	<title>Architecture of a flexible real-time video encoder/decoder: the DECchip 21230</title>. , 1997, , .		0
87	A language for describing predictors and its application to automatic synthesis. Computer Architecture News, 1997, 25, 304-314.	2.5	0
88	Exploiting choice. Computer Architecture News, 1996, 24, 191-202.	2.5	53
89	Exploiting choice. , 1996, , .		547
90	Incremental versus revolutionary research. ACM Computing Surveys, 1996, 28, 27.	23.0	0

#	Article	IF	Citations
91	Instruction fetching. Computer Architecture News, 1995, 23, 345-356.	2.5	4
92	Instruction fetching. , 1995, , .		49
93	Design analysis of a heterogeneous distributed system. , 1986, , .		0
94	Performance of the VAX-11/780 translation buffer. ACM Transactions on Computer Systems, 1985, 3, 31-62.	0.8	122
95	A Characterization of Processor Performance in the vax-11/780. , 1984, , .		75
96	A Characterization of Processor Performance in the vax-11/780. Computer Architecture News, 1984, 12, 301-310.	2.5	3