Paul V Gratz

List of Publications by Year
in descending order

Source: https://exaly.com/author-pdf/1015039/publications.pdf

Version: 2024-02-01

77	1,540	933447 10	580821 25
papers	citations	h-index	g-index

77	77	77	861
all docs	docs citations	times ranked	citing authors

#	Article	IF	Citations
1	Regional congestion awareness for load balance in networks-on-chip. High Performance Computer Architecture (HPCA), Proceedings of the IEEE International Symposium on, 2008, , .	0.0	266
2	Distributed Microarchitectural Protocols in the TRIPS Prototype Processor. Microarchitecture (MICRO), Proceedings of the Annual International Symposium on, 2006, , .	0.0	112
3	Implementation and Evaluation of On-Chip Network Architectures. Proceedings - IEEE International Conference on Computer Design: VLSI in Computers and Processors, 2006, , .	0.0	99
4	Path confidence based lookahead prefetching. , 2016, , .		75
5	Bandwidth-efficient on-chip interconnect designs for GPGPUs. , 2015, , .		60
6	LumiNOC: A Power-Efficient, High-Performance, Photonic Network-on-Chip. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2014, 33, 826-838.	2.7	55
7	An evaluation of the TRIPS computer system. , 2009, , .		48
8	Dynamic voltage and frequency scaling for shared resources in multicore processor designs. , 2013, , .		47
9	Up by their bootstraps: Online learning in Artificial Neural Networks for CMP uncore power management. , 2014, , .		44
10	Perceptron-based prefetch filtering. , 2019, , .		44
11	Implementation and Evaluation of a Dynamically Routed Processor Operand Network. , 2007, , .		41
12	A control-theoretic approach for energy efficient CPU-GPU subsystem in mobile platforms. , 2015, , .		36
13	Use it or lose it. , 2013, , .		35
14	GCA:Global Congestion Awareness for Load Balance in Networks-on-Chip. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 2022-2035.	5.6	33
15	GCA: Global congestion awareness for load balance in Networks-on-Chip. , 2013, , .		32
16	In-network Monitoring and Control Policy for DVFS of CMP Networks-on-Chip and Last Level Caches. , 2012, , .		30
17	Reducing Network-on-Chip energy consumption through spatial locality speculation. , 2011, , .		29
18	B-Fetch: Branch Prediction Directed Prefetching for Chip-Multiprocessors. , 2014, , .		29

#	Article	IF	Citations
19	AcENoCs: A Configurable HW/SW Platform for FPGA Accelerated NoC Emulation. , 2011, , .		27
20	Synchronized Progress in Interconnection Networks (SPIN): A New Theory for Deadlock Freedom. , 2018, , .		24
21	In-network monitoring and control policy for DVFS of CMP networks-on-chip and last level caches. ACM Transactions on Design Automation of Electronic Systems, 2013, 18, 1-21.	2.6	22
22	DRAIN: Deadlock Removal for Arbitrary Irregular Networks. , 2020, , .		19
23	Kill the Program Counter. , 2017, , .		18
24	LumiNOC. , 2012, , .		17
25	ARI. Transactions on Architecture and Code Optimization, 2013, 10, 1-19.	2.0	17
26	SWAP. , 2019, , .		16
27	B-Fetch: Branch Prediction Directed Prefetching for In-Order Processors. IEEE Computer Architecture Letters, 2012, 11, 41-44.	1.5	15
28	Asynchronous Bypass Channels: Improving Performance for Multi-synchronous NoCs. , 2010, , .		14
29	Power gating with block migration in chip-multiprocessor last-level caches. , 2013, , .		14
30	Hardware Memory Management for Future Mobile Hybrid Memory Systems. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2020, 39, 3627-3637.	2.7	14
31	Asynchronous Bypass Channels for Multi-Synchronous NoCs: A Router Microarchitecture, Topology, and Routing Algorithm. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2011, 30, 1663-1676.	2.7	13
32	Speculative paging for future NVM storage. , 2017, , .		13
33	Towards platform level power management in mobile systems. , 2014, , .		12
34	Kill the Program Counter. ACM SIGPLAN Notices, 2017, 52, 737-749.	0.2	11
35	Pitstop: Enabling a Virtual Network Free Network-on-Chip. , 2021, , .		10
36	MTB-Fetch: Multithreading Aware Hardware Prefetching for Chip Multiprocessors. IEEE Computer Architecture Letters, 2018, 17, 175-178.	1.5	9

#	Article	IF	Citations
37	WaveSync: A low-latency source synchronous bypass network-on-chip architecture. , 2012, , .		8
38	ILP and TLP in shared memory applications. , 2014, , .		8
39	Dynamic Memory Pressure Aware Ballooning. , 2015, , .		8
40	An evaluation of the TRIPS computer system. ACM SIGPLAN Notices, 2009, 44, 1-12.	0.2	7
41	Stochastic Pre-classification for SDN Data Plane Matching. , 2014, , .		7
42	The design space of ultra-low energy asymmetric cryptography. , 2014, , .		7
43	Use It or Lose It. ACM Transactions on Design Automation of Electronic Systems, 2015, 20, 1-26.	2.6	7
44	SDPR: Improving Latency and Bandwidth in On-Chip Interconnect Through Simultaneous Dual-Path Routing. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2018, 37, 545-558.	2.7	7
45	Energy-efficient optical broadcast for nanophotonic networks-on-chip. , 2012, , .		6
46	Having your cake and eating it too: Energy savings without performance loss through resource sharing driven power management. , 2015, , .		6
47	Software Hint-Driven Data Management for Hybrid Memory in Mobile Systems. Transactions on Embedded Computing Systems, 2022, 21, 1-18.	2.9	6
48	Shared Last-Level Caches and The Case for Longer Timeslices. , 2015, , .		5
49	Optimizing Post-Copy Live Migration with System-Level Checkpoint Using Fabric-Attached Memory. , 2019, , .		5
50	Energy-efficient implementations of GF (p) and GF(2m) elliptic curve cryptography. , 2015, , .		4
51	Kill the Program Counter. Operating Systems Review (ACM), 2017, 51, 737-749.	1.9	4
52	OpenMem: Hardware/Software Cooperative Management for Mobile Memory System. , 2021, , .		4
53	Leveraging Unused Cache Block Words to Reduce Power in CMP Interconnect. IEEE Computer Architecture Letters, 2010, 9, 33-36.	1.5	3
54	Exploiting path diversity for low-latency and high-bandwidth with the dual-path NoC router. , 2012, , .		3

#	Article	IF	Citations
55	Stochastic Pre-Classification for Software Defined Firewalls. , 2013, , .		3
56	WaveSync. ACM Transactions on Design Automation of Electronic Systems, 2014, 19, 1-22.	2.6	3
57	Virtualize and share non-volatile memories in user space. CCF Transactions on High Performance Computing, 2020, 2, 16-35.	1.7	3
58	Automatic Microprocessor Performance Bug Detection. , 2021, , .		3
59	An evaluation of the TRIPS computer system. Computer Architecture News, 2009, 37, 1-12.	2.5	3
60	Stay in your Lane: A NoC with Low-overhead Multi-packet Bypassing. , 2022, , .		3
61	Clotho: Proactive wearout deceleration in Chip-Multiprocessor interconnects. , 2015, , .		2
62	Resource Sharing Centric Dynamic Voltage and Frequency Scaling for CMP Cores, Uncore, and Memory. ACM Transactions on Design Automation of Electronic Systems, 2016, 21, 1-25.	2.6	2
63	Synchronized Progress in Interconnection Networks (SPIN): A New Theory for Deadlock Freedom. IEEE Micro, 2019, 39, 110-117.	1.8	2
64	Nano-Photonic Networks-on-Chip for Future Chip Multiprocessors. , 2015, , 155-186.		2
65	SB-Fetch. , 2020, , .		2
66	LumiNOC: A low-latency, high-bandwidth per Watt, photonic Network-on-Chip. , 2013, , .		1
67	STORM: A Simple Traffic-Optimized Router Microarchitecture for Networks-on-Chip. , 2014, , .		1
68	Kill the Program Counter. Computer Architecture News, 2017, 45, 737-749.	2.5	1
69	Guest Editorial: Emerging Technologies and Architectures for Manycore Computing Part 1: Hardware Techniques. IEEE Transactions on Multi-Scale Computing Systems, 2018, 4, 97-98.	2.4	1
70	GenMatcher. Transactions on Architecture and Code Optimization, 2019, 15, 1-22.	2.0	1
71	SIMD-Matcher: A SIMD-based Arbitrary Matching Framework. Transactions on Architecture and Code Optimization, 0, , .	2.0	1
72	Reducing Minor Page Fault Overheads through Enhanced Page Walker. Transactions on Architecture and Code Optimization, 2022, 19, 1-26.	2.0	1

#	Article	IF	Citations
73	Wear-Aware Adaptive Routing for Networks-on-Chips. , 2015, , .		0
74	Minimal exercise vector generation for reliability improvement. , 2017, , .		0
75	SpecLock: Speculative Lock Forwarding. , 2019, , .		0
76	vNVML: An Efficient User Space Library for Virtualizing and Sharing Non-Volatile Memories. , 2019, , .		0
77	A Generic FPGA Accelerator for Minimum Storage Regenerating Codes. , 2020, , .		0