

Antonio Gonzalez

List of Publications by Year in descending order

Source: <https://exaly.com/author-pdf/2398757/publications.pdf>

Version: 2024-02-01

256
papers

4,161
citations

430754

18
h-index

414303

32
g-index

261
all docs

261
docs citations

261
times ranked

1578
citing authors

#	ARTICLE	IF	CITATIONS
1	Multiple-banked register file architectures. , 2000, , .		161
2	Energy-effective issue logic. , 2001, , .		154
3	Clustered speculative multithreaded processors. , 1999, , .		138
4	Penelope: The NBTI-Aware Processor. , 2007, , .		138
5	KAM theory without action-angle variables. Nonlinearity, 2005, 18, 855-895.	0.6	110
6	Speculative execution via address prediction and data prefetching. , 1997, , .		107
7	Speculative multithreaded processors. , 1998, , .		105
8	Eliminating cache conflict misses through XOR-based placement functions. , 1997, , .		104
9	Understanding the Thermal Implications of Multi-Core Architectures. IEEE Transactions on Parallel and Distributed Systems, 2007, 18, 1055-1065.	4.0	104
10	Mitosis compiler. , 2005, , .		94
11	Impact of Parameter Variations on Circuits and Microarchitecture. IEEE Micro, 2006, 26, 30-39.	1.8	89
12	A data cache with multiple caching strategies tuned to different types of locality. , 2014, , .		82
13	IATAC: a smart predictor to turn-off L2 cache lines. Transactions on Architecture and Code Optimization, 2005, 2, 55-77.	1.6	76
14	Swing module scheduling: a lifetime-sensitive approach. , 0, , .		74
15	Meeting points. , 2008, , .		70
16	Computation Reuse in DNNs by Exploiting Input Similarity. , 2018, , .		70
17	Low Vccmin fault-tolerant cache with highly predictable performance. , 2009, , .		63
18	Design space exploration for multicore architectures. , 2006, , .		59

#	ARTICLE	IF	CITATIONS
19	Power/Performance/Thermal Design-Space Exploration for Multicore Architectures. IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 666-681.	4.0	57
20	Lifetime-sensitive modulo scheduling in a production environment. IEEE Transactions on Computers, 2001, 50, 234-249.	2.4	52
21	Reducing the complexity of the issue logic. , 2001, , .		52
22	Mitosis compiler. ACM SIGPLAN Notices, 2005, 40, 269-279.	0.2	50
23	A comparative study of modulo scheduling techniques. , 2002, , .		47
24	Randomized cache placement for eliminating conflicts. IEEE Transactions on Computers, 1999, 48, 185-192.	2.4	46
25	Mitosis: A Speculative Multithreaded Processor Based on Precomputation Slices. IEEE Transactions on Parallel and Distributed Systems, 2008, 19, 914-925.	4.0	44
26	A Dynamically Adaptable Hardware Transactional Memory. , 2010, , .		42
27	FASTM: A Log-based Hardware Transactional Memory with Fast Abort Recovery. , 2009, , .		41
28	Scalability of Broadcast Performance in Wireless Network-on-Chip. IEEE Transactions on Parallel and Distributed Systems, 2016, 27, 3631-3645.	4.0	38
29	DDGacc. ACM SIGPLAN Notices, 2012, 47, 159-168.	0.2	37
30	Dynamic removal of redundant computations. , 1999, , .		33
31	Power-aware control speculation through selective throttling. , 0, , .		33
32	MeRLiN. , 2017, , .		33
33	Energy-effective issue logic. Computer Architecture News, 2001, 29, 230-239.	2.5	32
34	Modulo scheduling for a fully-distributed clustered VLIW architecture. , 2000, , .		31
35	The effectiveness of loop unrolling for modulo scheduling in clustered VLIW architectures. , 0, , .		30
36	Distributing the Frontend for Temperature Reduction. , 0, , .		30

#	ARTICLE	IF	CITATIONS
37	A fast and accurate framework to analyze and optimize cache memory behavior. ACM Transactions on Programming Languages and Systems, 2004, 26, 263-300.	1.7	29
38	Exploiting Narrow Values for Soft Error Tolerance. IEEE Computer Architecture Letters, 2006, 5, 12-12.	1.0	29
39	Multiple-banked register file architectures. Computer Architecture News, 2000, 28, 316-325.	2.5	27
40	Accelerating microprocessor silicon validation by exposing ISA diversity. , 2011, , .		27
41	Thread-spawning schemes for speculative multithreading. , 0, , .		26
42	Compiler directed early register release. , 2005, , .		26
43	Back-end assignment schemes for clustered multithreaded processors. , 2004, , .		25
44	Selective replication. ACM Transactions on Computer Systems, 2009, 27, 1-30.	0.6	25
45	Efficient interconnects for clustered microarchitectures. , 0, , .		24
46	Power- and complexity-aware issue queue designs. IEEE Micro, 2003, 23, 50-58.	1.8	22
47	Processor Microarchitecture: An Implementation Perspective. Synthesis Lectures on Computer Architecture, 2010, 5, 1-116.	1.3	22
48	An ultra low-power hardware accelerator for automatic speech recognition. , 2016, , .		22
49	Hardware schemes for early register release. , 0, , .		21
50	Heterogeneous way-size cache. , 2006, , .		21
51	Thermal-aware clustered microarchitectures. , 0, , .		20
52	Cross-layer system reliability assessment framework for hardware faults. , 2016, , .		20
53	Graph-partitioning based instruction scheduling for clustered processors. , 0, , .		19
54	Late allocation and early release of physical registers. IEEE Transactions on Computers, 2004, 53, 1244-1259.	2.4	19

#	ARTICLE	IF	CITATIONS
55	Refueling: Preventing Wire Degradation due to Electromigration. IEEE Micro, 2008, 28, 37-46.	1.8	19
56	HK-NUCA: Boosting Data Searches in Dynamic Non-Uniform Cache Architectures for Chip Multiprocessors. , 2011, , .		19
57	Independent front-end and back-end dynamic voltage scaling for a GALS microarchitecture. , 2006, , .		18
58	Energy-efficient register caching with compiler assistance. Transactions on Architecture and Code Optimization, 2009, 6, 1-23.	1.6	18
59	The Dark Side of DNN Pruning. , 2018, , .		18
60	Modulo scheduling with reduced register pressure. IEEE Transactions on Computers, 1998, 47, 625-638.	2.4	17
61	Dual path instruction processing. , 2002, , .		17
62	MT-SBST: Self-test optimization in multithreaded multicore architectures. , 2010, , .		17
63	Boosting single-thread performance in multi-core systems through fine-grain multi-threading. , 2009, , .		17
64	Circuit propagation delay estimation through multivariate regression-based modeling under spatio-temporal variability. , 2010, , .		16
65	Instruction scheduling for clustered VLIW architectures. , 0, , .		15
66	On reducing register pressure and energy in multiple-banked register files. , 0, , .		15
67	Non redundant data cache. , 2003, , .		15
68	Thread partitioning and value prediction for exploiting speculative thread-level parallelism. IEEE Transactions on Computers, 2004, 53, 114-125.	2.4	15
69	On-Line Failure Detection and Confinement in Caches. , 2008, , .		15
70	Thread shuffling: Combining DVFS and thread migration to reduce energy consumptions for multi-core systems. , 2011, , .		15
71	DDGacc. , 2012, , .		15
72	A Case for Acoustic Wave Detectors for Soft-Errors. IEEE Transactions on Computers, 2016, 65, 5-18.	2.4	15

#	ARTICLE	IF	CITATIONS
73	Executing algorithms with hypercube topology on torus multicomputers. IEEE Transactions on Parallel and Distributed Systems, 1995, 6, 803-814.	4.0	14
74	Heterogeneous Clustered VLIW Microarchitectures. , 2007, , .		14
75	AGAMOS: A Graph-Based Approach to Modulo Scheduling for Clustered Microarchitectures. IEEE Transactions on Computers, 2009, 58, 770-783.	2.4	14
76	An interleaved cache clustered VLIW processor. , 2002, , .		13
77	Dynamic Code Partitioning for Clustered Architectures. International Journal of Parallel Programming, 2001, 29, 59-79.	1.1	12
78	On-chip interconnects and instruction steering schemes for clustered microarchitectures. IEEE Transactions on Parallel and Distributed Systems, 2005, 16, 130-144.	4.0	12
79	Reducing Soft Errors through Operand Width Aware Policies. IEEE Transactions on Dependable and Secure Computing, 2009, 6, 217-230.	3.7	12
80	SoftHV. , 2011, , .		12
81	Reducing DUE-FIT of caches by exploiting acoustic wave detectors for error recovery. , 2013, , .		12
82	Low-Power Automatic Speech Recognition Through a Mobile GPU and a Viterbi Accelerator. IEEE Micro, 2017, 37, 22-29.	1.8	12
83	An Ultra Low-Power Hardware Accelerator for Acoustic Scoring in Speech Recognition. , 2017, , .		12
84	A Low-Power, High-Performance Speech Recognition Accelerator. IEEE Transactions on Computers, 2019, 68, 1817-1831.	2.4	12
85	End-to-end register data-flow continuous self-test. , 2009, , .		12
86	Improving latency tolerance of multithreading through decoupling. IEEE Transactions on Computers, 2001, 50, 1084-1094.	2.4	11
87	Exploiting pseudo-schedules to guide data dependence graph partitioning. , 0, , .		11
88	Cache organizations for clustered microarchitectures. , 2004, , .		11
89	Low-Complexity Distributed Issue Queue. , 0, , .		11
90	VCTA: A Via-Configurable Transistor Array regular fabric. , 2010, , .		11

#	ARTICLE	IF	CITATIONS
91	Setting an error detection infrastructure with low cost acoustic wave detectors. , 2012, , .		11
92	Analysis and Optimization of Engines for Dynamically Typed Languages. , 2015, , .		11
93	A Novel Register Renaming Technique for Out-of-Order Processors. , 2018, , .		11
94	Visibility Rendering Order: Improving Energy Efficiency on Mobile GPUs through Frame Coherence. IEEE Transactions on Parallel and Distributed Systems, 2019, 30, 473-485.	4.0	11
95	Reducing memory traffic via redundant store instructions. Lecture Notes in Computer Science, 1999, , 1246-1249.	1.0	10
96	An efficient solver for Cache Miss Equations. , 0, , .		10
97	Instruction replication for clustered microarchitectures. , 0, , .		10
98	Software Directed Issue Queue Power Reduction. , 0, , .		10
99	Improving Branch Prediction and Predicated Execution in Out-of-Order Processors. , 2007, , .		10
100	Version management alternatives for hardware transactional memory. , 2008, , .		10
101	Efficient Power Gating of SIMD Accelerators Through Dynamic Selective Devectorization in an HW/SW Codesigned Environment. Transactions on Architecture and Code Optimization, 2014, 11, 1-23.	1.6	10
102	Rendering Elimination: Early Discard of Redundant Tiles in the Graphics Pipeline. , 2019, , .		10
103	Ecological consumer neuroscience for competitive advantage and business or organizational differentiation. European Research on Management and Business Economics, 2020, 26, 174-180.	3.4	10
104	Trace-level speculative multithreaded architecture. , 0, , .		9
105	Virtual Cluster Scheduling Through the Scheduling Graph. , 2007, , .		9
106	LRU-PEA: A smart replacement policy for non-uniform cache architectures on chip multiprocessors. , 2009, , .		9
107	Exploring the limits of early register release. Transactions on Architecture and Code Optimization, 2009, 6, 1-30.	1.6	9
108	Anaphase: A Fine-Grain Thread Decomposition Scheme for Speculative Multithreading. , 2009, , .		9

#	ARTICLE	IF	CITATIONS
109	A novel variation-tolerant 4T-DRAM cell with enhanced soft-error tolerance. , 2012, , .		9
110	UNFOLD. , 2017, , .		9
111	Near-optimal loop tiling by means of cache miss equations and genetic algorithms. , 0, , .		8
112	Power efficient data cache designs. , 0, , .		8
113	Replacement techniques for dynamic NUCA cache designs on CMPs. Journal of Supercomputing, 2013, 64, 548-579.	2.4	8
114	Avoiding core's DUE & SDC via acoustic wave detectors and tailored error containment and recovery. Computer Architecture News, 2014, 42, 37-48.	2.5	8
115	Chryso. , 2015, , .		8
116	Shared resource aware scheduling on power-constrained tiled many-core processors. Journal of Parallel and Distributed Computing, 2017, 100, 30-41.	2.7	8
117	Limits of Instruction Level Parallelism with Data Value Speculation. Lecture Notes in Computer Science, 1999, , 452-465.	1.0	8
118	The auction. , 2010, , .		8
119	A Detailed Methodology to Compute Soft Error Rates in Advanced Technologies. , 2016, , .		8
120	A Fast and Accurate Approach to Analyze Cache Memory Behavior. Lecture Notes in Computer Science, 2000, , 194-198.	1.0	8
121	Software-controlled operand-gating. , 0, , .		7
122	Reliability: Fallacy or Reality?. IEEE Micro, 2007, 27, 36-45.	1.8	7
123	Thread fusion. , 2008, , .		7
124	Avoiding core's DUE & SDC via acoustic wave detectors and tailored error containment and recovery. , 2014, , .		7
125	MASKIt: Soft error rate estimation for combinational circuits. , 2016, , .		7
126	Warm-Up Simulation Methodology for HW/SW Co-Designed Processors. , 2014, , .		7

#	ARTICLE	IF	CITATIONS
127	A mechanism for reducing the cost of branches in RISC architectures. <i>Microprocessing and Microprogramming</i> , 1988, 24, 565-572.	0.3	6
128	Effective instruction scheduling techniques for an interleaved cache clustered VLIW processor. , 0, , .		6
129	Flexible compiler-managed L0 buffers for clustered VLIW processors. , 0, , .		6
130	Memory bank predictors. , 0, , .		6
131	Selective predicate prediction for out-of-order processors. , 2006, , .		6
132	Fuse: A Technique to Anticipate Failures due to Degradation in ALUs. , 2007, , .		6
133	Boosting single-thread performance in multi-core systems through fine-grain multi-threading. <i>Computer Architecture News</i> , 2009, 37, 474-483.	2.5	6
134	A Performance and Area Efficient Architecture for Intrusion Detection Systems. , 2011, , .		6
135	Implementing a hybrid SRAM / eDRAM NUCA architecture. , 2011, , .		6
136	Dynamic Selective Devectorization for Efficient Power Gating of SIMD Units in a HW/SW Co-Designed Environment. , 2013, , .		6
137	Speculative dynamic vectorization to assist static vectorization in a HW/SW co-designed environment. , 2013, , .		6
138	Framework for economical error recovery in embedded cores. , 2014, , .		6
139	Performance Analysis and Optimization of Automatic Speech Recognition. <i>IEEE Transactions on Multi-Scale Computing Systems</i> , 2018, 4, 847-860.	2.5	6
140	Fast and Accurate SER Estimation for Large Combinational Blocks in Early Stages of the Design. <i>IEEE Transactions on Sustainable Computing</i> , 2021, 6, 427-440.	2.2	6
141	Near-Optimal Padding for Removing Conflict Misses. <i>Lecture Notes in Computer Science</i> , 2005, , 329-343.	1.0	6
142	Reducing branch delay to zero in pipelined processors. <i>IEEE Transactions on Computers</i> , 1993, 42, 363-371.	2.4	5
143	A method for exploiting communication/computation overlap in hypercubes. <i>Parallel Computing</i> , 1998, 24, 221-245.	1.3	5
144	Analyzing data locality in numeric applications. <i>IEEE Micro</i> , 2000, 20, 58-66.	1.8	5

#	ARTICLE	IF	CITATIONS
145	Hypercube algorithms on mesh connected multicomputers. IEEE Transactions on Parallel and Distributed Systems, 2002, 13, 1247-1260.	4.0	5
146	Dynamic cluster resizing. , 0, , .		5
147	Using Coherence Information and Decay Techniques to Optimize L2 Cache Leakage in CMPs. , 2009, , .		5
148	Thread-management techniques to maximize efficiency in multicore and simultaneous multithreaded microprocessors. Transactions on Architecture and Code Optimization, 2010, 7, 1-25.	1.6	5
149	Effectiveness of hybrid recovery techniques on parametric failures. , 2013, , .		5
150	Deconfigurable microprocessor architectures for silicon debug acceleration. , 2013, , .		5
151	iRMW: A low-cost technique to reduce NBTI-dependent parametric failures in L1 data caches. , 2014, , .		5
152	Power-Aware Adaptive Issue Queue and Register File. Lecture Notes in Computer Science, 2003, , 34-43.	1.0	5
153	MEM: A new execution model for Prolog. Microprocessing and Microprogramming, 1993, 39, 83-86.	0.3	4
154	COMMUNICATION PIPELINING IN HYPERCUBES. Parallel Processing Letters, 1996, 06, 507-523.	0.4	4
155	Data value speculation in superscalar processors. Microprocessors and Microsystems, 1998, 22, 293-301.	1.8	4
156	Control-flow speculation through value prediction. IEEE Transactions on Computers, 2001, 50, 1362-1376.	2.4	4
157	Local scheduling techniques for memory coherence in a clustered VLIW processor with a distributed data cache. , 0, , .		4
158	Optimizing program locality through CMEs and GAs. , 0, , .		4
159	Control-Flow Independence Reuse via Dynamic Vectorization. , 0, , .		4
160	Online error detection and correction of erratic bits in register files. , 2009, , .		4
161	Energy efficiency via thread fusion and value reuse. IET Computers and Digital Techniques, 2010, 4, 114-125.	0.9	4
162	Hardware/software-based diagnosis of load-store queues using expandable activity logs. , 2011, , .		4

#	ARTICLE	IF	CITATIONS
163	New reliability mechanisms in memory design for sub-22nm technologies. , 2011, , .		4
164	Improving the Performance Efficiency of an IDS by Exploiting Temporal Locality in Network Traffic. , 2012, , .		4
165	Performance analysis and predictability of the software layer in dynamic binary translators/optimizers. , 2013, , .		4
166	Vectorizing for Wider Vector Units in a HW/SW Co-designed Environment. , 2013, , .		4
167	Cross-layer early reliability evaluation: Challenges and promises. , 2014, , .		4
168	CROB: Implementing a Large Instruction Window through Compression. Lecture Notes in Computer Science, 2011, , 115-134.	1.0	4
169	MeRLiN. Computer Architecture News, 2017, 45, 241-254.	2.5	4
170	Independent Front-end and Back-end Dynamic Voltage Scaling for a GALS Microarchitecture. , 2006, , .		4
171	Low Communication Overhead Jacobi Algorithms for Eigenvalues Computation on Hypercubes. Journal of Supercomputing, 1999, 14, 171-193.	2.4	3
172	Value Compression to Reduce Power in Data Caches. Lecture Notes in Computer Science, 2003, , 616-622.	1.0	3
173	Frontend frequency-voltage adaptation for optimal energy-delay. , 0, , .		3
174	An accurate cost model for guiding data locality transformations. ACM Transactions on Programming Languages and Systems, 2005, 27, 946-987.	1.7	3
175	Demystifying on-the-fly spill code. ACM SIGPLAN Notices, 2005, 40, 180-189.	0.2	3
176	Compiler analysis for trace-level speculative multithreaded architectures. , 2005, , .		3
177	Variable-based multi-module data caches for clustered VLIW processors. , 2005, , .		3
178	Control speculation for energy-efficient next-generation superscalar processors. IEEE Transactions on Computers, 2006, 55, 281-291.	2.4	3
179	Using MCD-DVS For Dynamic Thermal Management Performance Improvement. , 0, , .		3
180	Early Register Release for Out-of-Order Processors with RegisterWindows. Parallel Architecture and Compilation Techniques (PACT), Proceedings of the International Conference on, 2007, , .	0.0	3

#	ARTICLE	IF	CITATIONS
181	End-to-end register data-flow continuous self-test. Computer Architecture News, 2009, 37, 105-115.	2.5	3
182	Leveraging Register Windows to Reduce Physical Registers to the Bare Minimum. IEEE Transactions on Computers, 2010, 59, 1598-1610.	2.4	3
183	High-Performance low-vcc in-order core. , 2010, , .		3
184	Fg-STP: Fine-Grain Single Thread Partitioning on Multicores. , 2011, , .		3
185	Fast time-to-market with via-configurable transistor array regular fabric: A delay-locked loop design case study. , 2011, , .		3
186	Impact of positive bias temperature instability (PBTI) on 3T1D-DRAM cells. , 2011, , .		3
187	Speculative dynamic vectorization for HW/SW co-designed processors. , 2012, , .		3
188	Impact of positive bias temperature instability (PBTI) on 3T1D-DRAM cells. The Integration VLSI Journal, 2012, 45, 246-252.	1.3	3
189	A HW/SW Co-designed Programmable Functional Unit. IEEE Computer Architecture Letters, 2012, 11, 9-12.	1.0	3
190	Ultra-low power render-based collision detection for CPU/GPU systems. , 2015, , .		3
191	Assisting Static Compiler Vectorization with a Speculative Dynamic Vectorizer in an HW/SW Codesigned Environment. ACM Transactions on Computer Systems, 2016, 33, 1-33.	0.6	3
192	Removing checks in dynamically typed languages through efficient profiling. , 2017, , .		3
193	DNN pruning with principal component analysis and connection importance estimation. Journal of Systems Architecture, 2022, 122, 102336.	2.5	3
194	A Jacobi-based algorithm for computing symmetric eigenvalues and eigenvectors in a two-dimensional mesh. , 0, , .		2
195	Software Data Prefetching for Software Pipelined Loops. Journal of Parallel and Distributed Computing, 1999, 58, 236-259.	2.7	2
196	Implementing the one-sided Jacobi method on a 2D/3D mesh multicomputer. Parallel Computing, 2001, 27, 1253-1271.	1.3	2
197	Selective branch prediction reversal by correlating with data values and control flow. , 0, , .		2
198	Distributed Data Cache Designs for Clustered VLIW Processors. IEEE Transactions on Computers, 2005, 54, 1227-1241.	2.4	2

#	ARTICLE	IF	CITATIONS
199	Inherently Workload-Balanced Clustered Microarchitecture. , 0, , .		2
200	SAMIE-LSQ: set-associative multiple-instruction entry load/store queue. , 2006, , .		2
201	A software-hardware hybrid steering mechanism for clustered microarchitectures. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	2
202	P-slice based efficient speculative multithreading. , 2009, , .		2
203	MODEST. , 2010, , .		2
204	Implementing End-to-End Register Data-Flow Continuous Self-Test. IEEE Transactions on Computers, 2011, 60, 1194-1206.	2.4	2
205	Beforehand Migration on D-NUCA Caches. , 2011, , .		2
206	Hardware/Software Mechanisms for Protecting an IDS against Algorithmic Complexity Attacks. , 2012, , .		2
207	Exploiting temporal locality in network traffic using commodity multi-cores. , 2012, , .		2
208	Warm-Up Simulation Methodology for HW/SW Co-Designed Processors. , 2014, , .		2
209	Accurate off-line phase classification for HW/SW co-designed processors. , 2014, , .		2
210	INFORMER: An integrated framework for early-stage memory robustness analysis. , 2014, , .		2
211	ERICO: Effective Removal of Inline Caching Overhead in Dynamic Typed Languages. , 2016, , .		2
212	LAWS: Locality-Aware Scheme for Automatic Speech Recognition. IEEE Transactions on Computers, 2020, , 1-1.	2.4	2
213	Energy-Efficient Stream Compaction Through Filtering and Coalescing Accesses in GPGPU Memory Partitions. IEEE Transactions on Computers, 2022, 71, 1711-1723.	2.4	2
214	Overlapping communication and computation in hypercubes. Lecture Notes in Computer Science, 1996, , 253-257.	1.0	2
215	Errata on "Measuring Experimental Error in Microprocessor Simulation". Computer Architecture News, 2002, 30, 2-4.	2.5	2
216	Improving the Resilience of an IDS against Performance Throttling Attacks. Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, 2013, , 167-184.	0.2	2

#	ARTICLE	IF	CITATIONS
217	Design and Evaluation of an Ultra Low-power Human-quality Speech Recognition System. Transactions on Architecture and Code Optimization, 2020, 17, 1-19.	1.6	2
218	A Survey of Near-Data Processing Architectures for Neural Networks. Machine Learning and Knowledge Extraction, 2022, 4, 66-102.	3.2	2
219	Triangle Dropping: An Occluded-geometry Predictor for Energy-efficient Mobile GPUs. Transactions on Architecture and Code Optimization, 2022, 19, 1-20.	1.6	2
220	A partial breadth-first execution model for Prolog. , 0, , .		1
221	Optimizing cache miss equations polyhedra. Computer Architecture News, 2000, 28, 43-52.	2.5	1
222	A framework for modeling and optimization of prescient instruction prefetch. Performance Evaluation Review, 2003, 31, 13-24.	0.4	1
223	Empowering a helper cluster through data-width aware instruction selection policies. , 2006, , .		1
224	Efficient resources assignment schemes for clustered multithreaded processors. Parallel and Distributed Processing Symposium (IPDPS), Proceedings of the International Conference on, 2008, , .	1.0	1
225	Dynamic fine-grain body biasing of caches with latency and leakage 3T1D-based monitors. , 2011, , .		1
226	The migration prefetcher. Transactions on Architecture and Code Optimization, 2012, 8, 1-20.	1.6	1
227	Quantitative characterization of the software layer of a HW/SW co-designed processor. , 2016, , .		1
228	HW/SW co-designed processors: Challenges, design choices and a simulation infrastructure for evaluation. , 2017, , .		1
229	POSTER: Leveraging Run-Time Feedback for Efficient ASR Acceleration. , 2019, , .		1
230	Value Compression for Efficient Computation. Lecture Notes in Computer Science, 2005, , 519-529.	1.0	1
231	A survey of branch techniques in pipelined processors. Microprocessing and Microprogramming, 1993, 36, 243-257.	0.3	0
232	Design and evaluation of an instruction cache for reducing the cost of branches. Performance Evaluation, 1994, 20, 83-96.	0.9	0
233	The Multipath Architecture for Prolog Programs. Computer Journal, 1996, 39, 780-792.	1.5	0
234	Hardware support for early register release. International Journal of High Performance Computing and Networking, 2005, 3, 83.	0.4	0

#	ARTICLE	IF	CITATIONS
235	Instruction scheduling for a clustered VLIW processor with a word-interleaved cache. Concurrency Computation Practice and Experience, 2006, 18, 1391-1411.	1.4	0
236	Building a large instruction window through ROB compression. , 2007, , .		0
237	Guest Editors' Introduction: Micro's Top Picks from the Microarchitecture Conferences. IEEE Micro, 2007, 27, 8-11.	1.8	0
238	Message from the General and Program Chairs. , 2008, , .		0
239	A Power-Efficient Co-designed Out-of-Order Processor. , 2011, , .		0
240	TRAMS Project: Variability and Reliability of SRAM Memories in sub-22nm Bulk-CMOS Technologies. Procedia Computer Science, 2011, 7, 148-149.	1.2	0
241	Global productiveness propagation. , 2011, , .		0
242	Design of complex circuits using the Via-Configurable transistor array regular layout fabric. , 2011, , .		0
243	Global productiveness propagation. ACM SIGPLAN Notices, 2011, 46, 161-170.	0.2	0
244	Guest Editors' Introduction: Special Issue on Variability and Aging. IEEE Design and Test, 2013, 30, 5-7.	1.1	0
245	Reconfigurable microprocessor architectures for silicon debug acceleration. Computer Architecture News, 2013, 41, 631-642.	2.5	0
246	Author retrospective for the dual data cache. , 2014, , .		0
247	INFORMER: An integrated framework for early-stage memory robustness analysis. , 2014, , .		0
248	Message from the general chairs. , 2016, , .		0
249	Shared resource aware scheduling on power-constrained tiled many-core processors. , 2016, , .		0
250	An Energy-Efficient Memory Unit for Clustered Microarchitectures. IEEE Transactions on Computers, 2016, 65, 2631-2637.	2.4	0
251	2018 International Symposium on Computer Architecture Influential Paper Award. IEEE Micro, 2018, 38, 76-77.	1.8	0
252	Complete Exchange Algorithms for Meshes and Tori Using a Systematic Approach. Lecture Notes in Computer Science, 2000, , 591-594.	1.0	0

#	ARTICLE	IF	CITATIONS
253	Speculative execution for hiding memory latency. , 2004, , .		0
254	A New Jacobi Ordering for Multiple-Port Hypercubes. , 1999, , 77-88.		0
255	Reducing Misspeculation Penalty in Trace-Level Speculative Multithreaded Architectures. , 2005, , 43-55.		0
256	CREW: Computation reuse and efficient weight storage for hardware-accelerated MLPs and RNNs. Journal of Systems Architecture, 2022, 129, 102604.	2.5	0