# Saman P Amarasinghe

List of Publications by Year
in descending order



| 125 | 7,847 | 361413 | 206112 |
|---|---|---|---|
| papers | citations | 20 h-index | 48 g-index |

| 127 | 127 | 127 | 2707 |
|---|---|---|---|
| all docs | docs citations | times ranked | citing authors |

| # | Article | IF | Citations |
|---|---|---|---|
| 1 | StreamIt: A Language for Streaming Applications. Lecture Notes in Computer Science, 2002, , 179-196. | 1.3 | 508 |
| 2 | Halide. , 2013, , . | | 469 |
| 3 | SUIF. ACM SIGPLAN Notices, 1994, 29, 31-37. | 0.2 | 382 |
| 4 | Halide. ACM SIGPLAN Notices, 2013, 48, 519-530. | 0.2 | 342 |
| 5 | OpenTuner. , 2014, , . | | 334 |
| 6 | Exploiting coarse-grained task, data, and pipeline parallelism in stream programs. , 2006, , . | | 332 |
| 7 | Automatically patching errors in deployed software. , 2009, , . | | 282 |
| 8 | Kendo. , 2009, , . | | 265 |
| 9 | Exploiting superword level parallelism with multimedia instruction sets. , 2000, , . | | 251 |
| 10 | PetaBricks. , 2009, , . | | 232 |
| 11 | A stream compiler for communication-exposed architectures. , 2002, , . | | 223 |
| 12 | Decoupling algorithms from schedules for easy optimization of image processing pipelines. ACM Transactions on Graphics, 2012, 31, 1-12. | 7.2 | 192 |
| 13 | The tensor algebra compiler. , 2017, 1, 1-29. | | 183 |
| 14 | Evaluation of the Raw Microprocessor. Computer Architecture News, 2004, 32, 2. | 2.5 | 179 |
| 15 | A Practical Approach to Exploiting Coarse-Grained Pipeline Parallelism in C Programs. , 2007, , . | | 159 |
| 16 | Communication optimization and code generation for distributed memory machines. , 1993, , . | | 153 |
| 17 | Abstraction layers for scalable microfluidic biocomputing. Natural Computing, 2008, 7, 255-275. | 3.0 | 151 |
| 18 | Data and computation transformations for multiprocessors. , 1995, , . | | 148 |

| # | Article | IF | Citations |
|---|---|---|---|
| 19 | Space-time scheduling of instruction-level parallelism on a raw machine. , 1998, , . | | 144 |
| 20 | Bidwidth analysis with application to silicon compilation. , 2000, , . | | 144 |
| 21 | An empirical characterization of stream programs and its implications for language and compiler design. , 2010, , . | | 125 |
| 22 | Array-data flow analysis and its use in array privatization. , 1993, , . | | 115 |
| 23 | GraphIt: a high-performance graph DSL. , 2018, 2, 1-30. | | 108 |
| 24 | Meta optimization. , 2003, , . | | 106 |
| 25 | Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code. , 2019, , . | | 96 |
| 26 | Detecting coarse-grain parallelism using an interprocedural parallelizing compiler. , 1995, , . | | 90 |
| 27 | Meta optimization. ACM SIGPLAN Notices, 2003, 38, 77-90. | 0.2 | 81 |
| 28 | Making caches work for graph analytics. , 2017, , . | | 71 |
| 29 | Transparent dynamic instrumentation. , 2012, , . | | 70 |
| 30 | Exploiting superword level parallelism with multimedia instruction sets. ACM SIGPLAN Notices, 2000, 35, 145-156. | 0.2 | 67 |
| 31 | Exploiting coarse-grained task, data, and pipeline parallelism in stream programs. Operating Systems Review (ACM), 2006, 40, 151-162. | 1.9 | 66 |
| 32 | Format abstraction for sparse tensor algebra compilers. , 2018, 2, 1-30. | | 60 |
| 33 | Halide. Communications of the ACM, 2017, 61, 106-115. | 4.5 | 60 |
| 34 | An overview of a compiler for scalable parallel machines. Lecture Notes in Computer Science, 1994, , 253-272. | 1.3 | 56 |
| 35 | PetaBricks. ACM SIGPLAN Notices, 2009, 44, 38-49. | 0.2 | 53 |
| 36 | Dynamic cache contention detection in multi-threaded applications. , 2011, , . | | 52 |

| # | Article | IF | Citations |
|---|---|---|---|
| 37 | Portable performance on heterogeneous architectures. , 2013, , . | | 52 |
| 38 | Autotuning algorithmic choice for input sensitivity. , 2015, , . | | 50 |
| 39 | Language and compiler support for auto-tuning variable-accuracy algorithms. , 2011, , . | | 49 |
| 40 | Kendo. ACM SIGPLAN Notices, 2009, 44, 97-108. | 0.2 | 46 |
| 41 | Umbra. , 2010, , . | | 45 |
| 42 | Dynamic native optimization of interpreters. , 2003, , . | | 43 |
| 43 | Cache aware optimization of stream programs. , 2005, , . | | 40 |
| 44 | Simit. ACM Transactions on Graphics, 2016, 35, 1-21. | 7.2 | 39 |
| 45 | A lightweight streaming layer for multicore execution. Computer Architecture News, 2008, 36, 18-27. | 2.5 | 38 |
| 46 | A unified framework for schedule and storage optimization. , 2001, , . | | 37 |
| 47 | Language and Compiler Design for Streaming Applications. International Journal of Parallel Programming, 2005, 33, 261-278. | 1.5 | 37 |
| 48 | Phased scheduling of stream programs. , 2003, , . | | 32 |
| 49 | Exploiting coarse-grained task, data, and pipeline parallelism in stream programs. Computer Architecture News, 2006, 34, 151-162. | 2.5 | 30 |
| 50 | Evaluation of IVR data collection UIs for untrained rural users. , 2010, , . | | 30 |
| 51 | A sparse iteration space transformation framework for sparse tensor algebra. , 2020, 4, 1-30. | | 30 |
| 52 | Siblingrivalry. , 2012, , . | | 29 |
| 53 | Teleport messaging for distributed stream programs. , 2005, , . | | 28 |
| 54 | Optimizing Indirect Memory References with milk. , 2016, , . | | 28 |

| # | Article | IF | Citations |
|---|---|---|---|
| 55 | Exploiting coarse-grained task, data, and pipeline parallelism in stream programs. ACM SIGPLAN Notices, 2006, 41, 151-162. | 0.2 | 26 |
| 56 | Taco: A tool to generate tensor algebra kernels. , 2017, , . | | 25 |
| 57 | Tensor Algebra Compilation with Workspaces. , 2019, , . | | 25 |
| 58 | Interprocedural analysis for parallelization. Lecture Notes in Computer Science, 1996, , 61-80. | 1.3 | 24 |
| 59 | Autotuning algorithmic choice for input sensitivity. ACM SIGPLAN Notices, 2015, 50, 379-390. | 0.2 | 24 |
| 60 | Dynamic expressivity with static optimization for streaming languages. , 2013, , . | | 23 |
| 61 | Helium: lifting high-performance stencil kernels from stripped x86 binaries to halide DSL code. , 2015, , . | | 23 |
| 62 | Optimizing stream programs using linear state space analysis. , 2005, , . | | 23 |
| 63 | Detection of false sharing using machine learning. , 2013, , . | | 22 |
| 64 | Linear analysis and optimization of stream programs. , 2003, , . | | 21 |
| 65 | Optimizing ordered graph algorithms with GraphIt. , 2020, , . | | 21 |
| 66 | Autotuning multigrid with PetaBricks. , 2009, , . | | 20 |
| 67 | goSLP: globally optimized superword level parallelism framework. , 2018, 2, 1-28. | | 20 |
| 68 | The three pillars of machine programming. , 2018, , . | | 20 |
| 69 | Data and computation transformations for multiprocessors. ACM SIGPLAN Notices, 1995, 30, 166-178. | 0.2 | 20 |
| 70 | A stream compiler for communication-exposed architectures. Computer Architecture News, 2002, 30, 291-303. | 2.5 | 18 |
| 71 | Cache aware optimization of stream programs. ACM SIGPLAN Notices, 2005, 40, 115-126. | 0.2 | 18 |
| 72 | Bidwidth analysis with application to silicon compilation. ACM SIGPLAN Notices, 2000, 35, 108-120. | 0.2 | 18 |

| # | Article | IF | Citations |
|---|---|---|---|
| 73 | Transparent dynamic instrumentation. ACM SIGPLAN Notices, 2012, 47, 133-144. | 0.2 | 17 |
| 74 | Distributed Halide. , 2016, , . | | 17 |
| 75 | VeGen: a vectorizer generator for SIMD and beyond. , 2021, , . | | 17 |
| 76 | Communication optimization and code generation for distributed memory machines. ACM SIGPLAN Notices, 1993, 28, 126-138. | 0.2 | 16 |
| 77 | A stream compiler for communication-exposed architectures. ACM SIGPLAN Notices, 2002, 37, 291-303. | 0.2 | 16 |
| 78 | BHive: A Benchmark Suite and Measurement Framework for Validating x86-64 Basic Block Performance Models. , 2019, , . | | 16 |
| 79 | Kendo. Computer Architecture News, 2009, 37, 97-108. | 2.5 | 15 |
| 80 | Efficient memory shadowing for 64-bit architectures. , 2010, , . | | 15 |
| 81 | Automatic generation of efficient sparse tensor format conversion routines. , 2020, , . | | 15 |
| 82 | Aikido. , 2012, , . | | 14 |
| 83 | Compilation of sparse array programming models. , 2021, 5, 1-29. | | 13 |
| 84 | Seq: a high-performance language for bioinformatics. , 2019, 3, 1-29. | | 12 |
| 85 | A unified framework for schedule and storage optimization. ACM SIGPLAN Notices, 2001, 36, 232-242. | 0.2 | 10 |
| 86 | StreamJIT. , 2014, , . | | 10 |
| 87 | A stream compiler for communication-exposed architectures. Operating Systems Review (ACM), 2002, 36, 291-303. | 1.9 | 10 |
| 88 | A Unified Backend for Targeting FPGAs from DSLs. , 2018, , . | | 9 |
| 89 | Compiling Graph Applications for GPU s with GraphIt. , 2021, , . | | 9 |
| 90 | A Python-based programming language for high-performance computational genomics. Nature Biotechnology, 2021, 39, 1062-1064. | 17.5 | 9 |

| # | Article | IF | Citations |
|---|---|---|---|
| 91 | PetaBricks. , 2011, , . | | 9 |
| 92 | Autoscheduling for sparse tensor algebra with an asymptotic cost model. , 2022, , . | | 9 |
| 93 | Space-time scheduling of instruction-level parallelism on a raw machine. Operating Systems Review (ACM), 1998, 32, 46-57. | 1.9 | 8 |
| 94 | Manipulating lossless video in the compressed domain. , 2009, , . | | 8 |
| 95 | Dynamic cache contention detection in multi-threaded applications. ACM SIGPLAN Notices, 2011, 46, 27-38. | 0.2 | 8 |
| 96 | Cimple: instruction and memory level parallelism. , 2018, , . | | 8 |
| 97 | Revec: program rejuvenation through revectorization. , 2019, , . | | 8 |
| 98 | Unified compilation techniques for shared and distributed address space machines. , 1995, , . | | 7 |
| 99 | Phased scheduling of stream programs. ACM SIGPLAN Notices, 2003, 38, 103-112. | 0.2 | 7 |
| 100 | A step towards unifying schedule and storage optimization. ACM Transactions on Programming Languages and Systems, 2007, 29, 34. | 2.1 | 7 |
| 101 | Portable performance on heterogeneous architectures. Computer Architecture News, 2013, 41, 431-444. | 2.5 | 7 |
| 102 | Taming the Zoo: The Unified GraphIt Compiler Framework for Novel Architectures. , 2021, , . | | 7 |
| 103 | Linear analysis and optimization of stream programs. ACM SIGPLAN Notices, 2003, 38, 12-25. | 0.2 | 7 |
| 104 | An efficient evolutionary algorithm for solving incrementally structured problems. , 2011, , . | | 6 |
| 105 | Distributed Halide. ACM SIGPLAN Notices, 2016, 51, 1-12. | 0.2 | 6 |
| 106 | All you need is superword-level parallelism: systematic control-flow vectorization with SLP. , 2022, , . | | 6 |
| 107 | Space-time scheduling of instruction-level parallelism on a raw machine. ACM SIGPLAN Notices, 1998, 33, 46-57. | 0.2 | 5 |
| 108 | BuildIt: A Type-Based Multi-stage Programming Framework for Code Generation in C++. , 2021, , . | | 5 |

| # | Article | IF | Citations |
|---|---|---|---|
| 109 | Hyperparameter Tuning in Bandit-Based Adaptive Operator Selection. Lecture Notes in Computer Science, 2012, , 73-82. | 1.3 | 5 |
| 110 | Gloss. ACM SIGPLAN Notices, 2018, 53, 98-112. | 0.2 | 5 |
| 111 | Sparse Tensor Transpositions. , 2020, , . | | 5 |
| 112 | A Practical Approach to Exploiting Coarse-Grained Pipeline Parallelism in C Programs. Microarchitecture (MICRO), Proceedings of the Annual International Symposium on, 2007, , . | 0.0 | 4 |
| 113 | StreamJIT. ACM SIGPLAN Notices, 2014, 49, 177-195. | 0.2 | 4 |
| 114 | Gloss. , 2018, , . | | 4 |
| 115 | Strength Reduction of Integer Division and Modulo Operations. Lecture Notes in Computer Science, 2003, , 254-273. | 1.3 | 4 |
| 116 | Portable performance on heterogeneous architectures. ACM SIGPLAN Notices, 2013, 48, 431-444. | 0.2 | 4 |
| 117 | A Deep Dive Into Understanding The Random Walk-Based Temporal Graph Learning. , 2021, , . | | 4 |
| 118 | Helium: lifting high-performance stencil kernels from stripped x86 binaries to halide DSL code. ACM SIGPLAN Notices, 2015, 50, 391-402. | 0.2 | 2 |
| 119 | Aikido. Computer Architecture News, 2012, 40, 173-184. | 2.5 | 1 |
| 120 | Domain-Specific Language Abstractions for Compression. , 2021, , . | | 1 |
| 121 | Compiler 2.0. , 2020, , . | | 1 |
| 122 | GraphIt to CUDA Compiler in 2021 LOC: A Case for High-Performance DSL Implementation via Staging with BuilDSL. , 2022, , . | | 1 |
| 123 | Tiled Multicore Processors. Integrated Circuits and Systems, 2009, , 1-33. | 0.2 | 0 |
| 124 | Efficient memory shadowing for 64-bit architectures. ACM SIGPLAN Notices, 2010, 45, 93-102. | 0.2 | 0 |
| 125 | Aikido. ACM SIGPLAN Notices, 2012, 47, 173-184. | 0.2 | 0 |