

@INPROCEEDINGS{6012857, author={de O Sandes, E. F. and de Melo, A. C. M. A.}, booktitle={Parallel Distributed Processing Symposium (IPDPS), 2011 IEEE International}, title={Smith-Waterman Alignment of Huge Sequences with GPU in Linear Space}, year={2011}, month=may, pages={1199-1211}, keywords={bioinformatics;cellular biophysics;coprocessors;parallel algorithms;GPU;GTX 285 Board;Myers-Miller algorithm;Smith-Waterman alignment;ancestral relationships;bioinformatics;cross-species chromosome alignments;high performance computing platform;linear space complexity;parallel algorithm;species peculiarity identification;Bioinformatics;Computer architecture;Graphics processing unit;Heuristic algorithms;Instruction sets;Mathematical model;Microprocessors}, doi={10.1109/IPDPS.2011.114}, ISSN={1530-2075}, }

@article{Luporini:2015:COA:2695583.2687415, author = {Luporini, Fabio and Varbanescu, Ana Lucia and Rathgeber, Florian and Bercea, Gheorghe-Teodor and Ramanujam, J. and Ham, David A. and Kelly, Paul H. J.}, title = {Cross-Loop Optimization of Arithmetic Intensity for Finite Element Local Assembly}, journal = {ACM Trans. Archit. Code Optim.}, issue\_date = {January 2015}, volume = {11}, number = {4}, month = jan, year = {2015}, issn = {1544-3566}, pages = {57:1-57:25}, articleno = {57}, numpages = {25}, url = {<http://doi.acm.org/10.1145/2687415>}, doi = {10.1145/2687415}, acmid = {2687415}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Finite element integration, SIMD vectorization, compilers, local assembly, optimizations}, }

@article{Elango:2015:URM:2695583.2693656, author = {Elango, Venmugil and Sedaghati, Naser and Rastello, Fabrice and Pouchet, Louis-Noël and Ramanujam, J. and Teodorescu, Radu and Sadayappan, P.}, title = {On Using the Roofline Model with Lower Bounds on Data Movement}, journal = {ACM Trans. Archit. Code Optim.}, issue\_date = {January 2015}, volume = {11}, number = {4}, month = jan, year = {2015}, issn = {1544-3566}, pages = {67:1-67:23}, articleno = {67}, numpages = {23}, url = {<http://doi.acm.org/10.1145/2693656>}, doi = {10.1145/2693656}, acmid = {2693656}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {I/O lower bounds, Operational intensity upper bounds, algorithm-architecture codesign, architecture design space exploration}, }

@article{Kong:2015:CFD:2695583.2687652, author = {Kong, Martin and Pop, Antoniu and Pouchet, Louis-Noël and Govindarajan, R. and Cohen, Albert and Sadayappan, P.}, title = {Compiler/Runtime Framework for Dynamic Dataflow Parallelization of Tiled Programs}, journal = {ACM Trans. Archit. Code Optim.}, issue\_date = {January 2015}, volume = {11}, number = {4}, month = jan, year = {2015}, issn = {1544-3566}, pages = {61:1-61:30}, articleno = {61}, numpages = {30}, url = {<http://doi.acm.org/10.1145/2687652>}, doi = {10.1145/2687652}, acmid = {2687652}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Dataflow, auto-parallelization, dependence partitioning, dynamic wavefront, point-to-point synchronization, polyhedral compiler, polyhedral framework, tile dependences, tiling}, }

@article{Cilardo:2015:IMM:2695583.2675359, author = {Cilardo, Alessandro and Gallo, Luca}, title = {Improving Multibank Memory Access Parallelism with Lattice-Based Partitioning}, journal = {ACM Trans. Archit. Code Optim.}, issue\_date = {January 2015}, volume = {11}, number = {4}, month = jan, year = {2015}, issn = {1544-3566}, pages = {45:1-45:25}, articleno = {45}, numpages = {25}, url = {<http://doi.acm.org/10.1145/2675359>}, doi = {10.1145/2675359}, acmid = {2675359}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Memory partitioning, field-programmable gate arrays, fine-grained distributed shared memory, polyhedral model}, }

@article{tavarageri2013adaptive,

title={Adaptive parallel tiled code generation and accelerated auto-tuning},  
author={Tavarageri, Sanket and Ramanujam, J and Sadayappan, P},

```
journal={International Journal of High Performance Computing Applications},
volume={27},
number={4},
pages={412--425},
year={2013},
doi={10.1177/1094342013493939},
publisher={Sage Publications}
```

```
}
```

```
@inproceedings{Pouchet:2011:LTC:1926385.1926449, author = {Pouchet, Louis-Noël and Bondhugula, Uday and Bastoul, Cédric and Cohen, Albert and Ramanujam, J. and Sadayappan, P. and Vasilache, Nicolas}, title = {Loop Transformations: Convexity, Pruning and Optimization}, booktitle = {Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages}, series = {POPL '11}, year = {2011}, isbn = {978-1-4503-0490-0}, location = {Austin, Texas, USA}, pages = {549-562}, numpages = {14}, url = {http://doi.acm.org/10.1145/1926385.1926449}, doi = {10.1145/1926385.1926449}, acmid = {1926449}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {affine scheduling, compilation, compiler optimization, loop transformations, parallelism}, }
```

```
@incollection{Shirako:2012, year={2012}, isbn={978-3-642-28651-3}, booktitle={Compiler Construction}, volume={7210}, series={Lecture Notes in Computer Science}, editor={O'Boyle, Michael}, doi={10.1007/978-3-642-28652-0_6}, title={Analytical Bounds for Optimal Tile Size Selection}, url={http://dx.doi.org/10.1007/978-3-642-28652-0_6}, publisher={Springer Berlin Heidelberg}, author={Shirako, Jun and Sharma, Kamal and Fauzia, Naznin and Pouchet, Louis-Noël and Ramanujam, J. and Sadayappan, P. and Sarkar, Vivek}, pages={101-121}, language={English} }
```

```
@inproceedings{Henretty:2013:SCS:2464996.2467268, author = {Henretty, Tom and Veras, Richard and Franchetti, Franz and Pouchet, Louis-Noël and Ramanujam, J. and Sadayappan, P.}, title = {A Stencil Compiler for Short-vector SIMD Architectures}, booktitle = {Proceedings of the 27th International ACM Conference on International Conference on Supercomputing}, series = {ICS '13}, year = {2013}, isbn = {978-1-4503-2130-3}, location = {Eugene, Oregon, USA}, pages = {13-24}, numpages = {12}, url = {http://doi.acm.org/10.1145/2464996.2467268}, doi = {10.1145/2464996.2467268}, acmid = {2467268}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {dsl, multicore, simd, split tiling, stencils}, }
```

```
@inproceedings{Paul:2014:EHA:2616606.2616999, author = {Paul, Somnath and Karam, Robert and Bhunia, Swarup and Puri, Ruchir}, title = {Energy-efficient Hardware Acceleration Through Computing in the Memory}, booktitle = {Proceedings of the Conference on Design, Automation and Test in Europe}, series = {DATE '14}, year = {2014}, isbn = {978-3-9815370-2-4}, location = {Dresden, Germany}, pages = {266:1-266:6}, articleno = {266}, numpages = {6}, url = {http://dl.acm.org/citation.cfm?id=2616606.2616999}, acmid = {2616999}, publisher = {European Design and Automation Association}, address = {3001 Leuven, Belgium, Belgium}, }
```

```
@INPROCEEDINGS{6971829, author={Shafique, M. and Garg, S. and Mitra, T. and Parameswaran, S. and Henkel, J.}, booktitle={Hardware/Software Codesign and System Synthesis (CODES+ISSS), 2014 International Conference on}, title={Dark silicon as a challenge for hardware/software co-design}, year={2014}, month={Oct}, pages={1-10}, keywords={hardware-software codesign;power aware computing;automated synthesis;design space exploration;energy efficiency improvement;hardware-software co-design;heterogeneous computing;heterogeneous dark silicon processors;on-chip
```

resources;performance improvement;power budget;power constraint;reliability improvement;run-time management;safe thermal limits;thermal constraint;transistors;Computer architecture;Hardware;Program processors;Reliability;Silicon;Switches},  
doi={10.1145/2656075.2661645},}

@INPROCEEDINGS{6957228, author={Wang, H.H.-W. and Lin, L.Y.-Z. and Huang, R.H.-M. and Wen, C.H.-P.}, booktitle={Parallel Processing (ICPP), 2014 43rd International Conference on}, title={CASTA: CUDA-Accelerated Static Timing Analysis for VLSI Designs}, year={2014}, month={Sept}, pages={192-200}, keywords={VLSI;general purpose computers;graphics processing units;integrated circuit design;parallel architectures;table lookup;CASTA;CUDA-accelerated static timing analysis;VLSI designs;benchmark circuits;cell levelization type sorting;general-purpose computing;graphics processing unit;look-up tables;table-index remapping;texture-accelerated rendering;three-order speedup;timing-table restructuring;Circuit faults;Delays;Graphics processing units;Instruction sets;Libraries;Table lookup;CUDA;GPU;Parallel Computing;STA}, doi={10.1109/ICPP.2014.28}, ISSN={0190-3918},}

@INPROCEEDINGS{6957231, author={Rong Shi and Xiaoyi Lu and Potluri, S. and Hamidouche, K. and Jie Zhang and Panda, D.K.}, booktitle={Parallel Processing (ICPP), 2014 43rd International Conference on}, title={HAND: A Hybrid Approach to Accelerate Non-contiguous Data Movement Using MPI Datatypes on GPU Clusters}, year={2014}, month={Sept}, pages={221-230}, keywords={application program interfaces;graphics processing units;message passing;GPU clusters;HAND framework;MPI applications;MPI datatype handling;N-body particle simulation application;arbitrary noncontiguous data movement optimization;data patterns;datatype latency reduction;datatype-aware design;hybrid adaptive selection;hybrid approach-to-accelerate noncontiguous data movement;indexed datatypes;modified DDTBench suite;optimized techniques;performance improvement;struct datatypes;Arrays;Graphics processing units;Kernel;Shape;Three-dimensional displays;Tuning;Vectors;CUDA;Datatype;GPU;MPI}, doi={10.1109/ICPP.2014.31}, ISSN={0190-3918},}

@inproceedings{Mendonca:2013:BSC:2510648.2511034, author = {Mendonca, Fernando Machado and Melo, Alba Cristina Magalhaes Alves de}, title = {Biological Sequence Comparison on Hybrid Platforms with Dynamic Workload Adjustment}, booktitle = {Proceedings of the 2013 IEEE 27th International Symposium on Parallel and Distributed Processing Workshops and PhD Forum}, series = {IPDPSW '13}, year = {2013}, isbn = {978-0-7695-4979-8}, pages = {501-509}, numpages = {9}, url = {<http://dx.doi.org/10.1109/IPDPSW.2013.28>}, doi = {10.1109/IPDPSW.2013.28}, acmid = {2511034}, publisher = {IEEE Computer Society}, address = {Washington, DC, USA}, keywords = {bioinformatics, smith-waterman, GPUs, multicores}, }

@INPROCEEDINGS{6835930, author={Hechtman, B.A. and Shuai Che and Hower, D.R. and Yingying Tian and Beckmann, B.M. and Hill, M.D. and Reinhardt, S.K. and Wood, D.A.}, booktitle={High Performance Computer Architecture (HPCA), 2014 IEEE 20th International Symposium on}, title={QuickRelease: A throughput-oriented approach to release consistency on GPUs}, year={2014}, month={Feb}, pages={189-200}, keywords={graphics processing units;storage management;GPU memory system;QuickRelease;finer-grain synchronization;graphics processing unit;scratchpad memory;throughput-oriented memory system;Coherence;Graphics processing units;Instruction sets;Kernel;Protocols;Synchronization;Throughput}, doi={10.1109/HPCA.2014.6835930},}

@INPROCEEDINGS{6835968, author={Hayenga, M. and Naresh, V.R.K. and Lipasti, M.H.}, booktitle={High Performance Computer Architecture (HPCA), 2014 IEEE 20th International Symposium on}, title={Revolver: Processor architecture for power efficient loop execution}, year={2014}, month={Feb}, pages={591-602}, keywords={computer architecture;energy

conservation;instruction sets;power aware computing;Revolver architecture;branch prediction;dispatch logic;energy efficiency;frontend instruction dispatches;instruction fetch;loop buffers;loop execution;loop iterations;micro-op cache techniques;out-of-order execution core;out-of-order processor architecture;power efficient loop execution;processor core;processor frontend;static instruction instances;Arrays;Clocks;Out of order;Pipelines;Rain;Registers;Resource management}, doi={10.1109/HPCA.2014.6835968},}

@ARTICLE{6506838, author={Ilic, A. and Pratas, F. and Sousa, L.}, journal={Computer Architecture Letters}, title={Cache-aware Roofline model: Upgrading the loft}, year={2014}, month={Jan}, volume={13}, number={1}, pages={21-24}, keywords={cache storage;computer architecture;multiprocessing systems;application optimization;built-in hardware counters;cache-aware Roofline model;cache-awareness;computer architecture upper bound performance;curve fitness;Modeling;Multiprocessing systems;Performance evaluation;Simulation;Application optimization;C.0.d Modeling of computer architecture & C.0 General & C Computer Systems Organization;C.0.e System architectures;C.4.d Modeling techniques & C.4 Performance of Systems & C Computer Systems Organization;C.4.g Measurement;Multicore computer architectures;Performance modeling;evaluation;integration and modeling & C.0 General & C Computer Systems Organization;modeling;simulation of multiple-processor systems & C.4 Performance of Systems & C Computer Syst}, doi={10.1109/L-CA.2013.6}, ISSN={1556-6056},}

@INPROCEEDINGS{6800414, author={Torrellas, J.}, booktitle={Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014}, title={Extreme-scale computer architecture: Energy efficiency from the ground up}, year={2014}, month={March}, pages={1-5}, keywords={multiprocessing systems;parallel architectures;power aware computing;computing stack;core cluster hierarchy;energy consumption;energy efficiency;extreme-scale computer architecture;hardware techniques;on-chip memories;power consumption;software techniques;voltage guardbands;voltage regulation;Computer architecture;Hardware;Organizations;Power demand;System-on-chip;Transistors;Voltage control}, doi={10.7873/DATE.2014.213},}

@INPROCEEDINGS{7016212, author={Uy, R.L.}, booktitle={Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM), 2014 International Conference on}, title={Beyond multi-core: A survey of architectural innovations on microprocessor}, year={2014}, month={Nov}, pages={1-6}, keywords={microprocessor chips;multiprocessing systems;ILP;Intel processors;architectural innovations;heat dissipation;instruction level parallelism;microprocessor;multicore processor;single-core processors;Clocks;Microarchitecture;Microprocessors;Multicore processing;Program processors;Technological innovation;Transistors;Architectural innovations;Intel microarchitecture;computer architecture;multi-core architecture}, url={http://dx.doi.org/10.1109/HNICEM.2014.7016212}, doi={10.1109/HNICEM.2014.7016212},}

@INPROCEEDINGS{862383, author={Schreiber, R. and Aditya, S. and Rau, B.R. and Kathail, V. and Mahlke, S. and Abraham, S. and Snider, G.}, booktitle={Application-Specific Systems, Architectures, and Processors, 2000. Proceedings. IEEE International Conference on}, title={High-level synthesis of nonprogrammable hardware accelerators}, year={2000}, month={}, pages={113-124}, keywords={application specific integrated circuits;circuit CAD;coprocessors;digital signal processing chips;high level synthesis;integrated circuit design;parallel architectures;DSP chips;PICO-N system;RTL definition;controller local memory;coprocessors;customized VLIW processors;embedded ASIC;embedded nonprogrammable accelerator synthesis;high-level synthesis;initiation interval;interfaces;loop nests;nonprogrammable hardware accelerators;register transfer level;specified throughput;synchronous array;synthesizable VHDL;user application software

modification;very-long instruction word processors;Acceleration;Automatic generation control;Control system synthesis;Coprocessors;Hardware;High level synthesis;Process control;Registers;Synchronous generators;VLIW}, doi={10.1109/ASAP.2000.862383}, ISSN={2160-0511},}

@ARTICLE{595572, author={Andonov, R. and Rajopadhye, S.}, journal={Parallel and Distributed Systems, IEEE Transactions on}, title={Knapsack on VLSI: From algorithm to optimal circuit}, year={1997}, month={jun}, volume={8}, number={6}, pages={545-561}, keywords={VLSI;application specific integrated circuits;logic design;parallel algorithms;systolic arrays;NP-hard problem;application specific VLSI design;correctness preserving transformations;dynamic dependencies;linear systolic array;model of computation;nonlinear discrete optimization;parallel solution;recurrence equations;space-time transformations;systolic arrays;systolic synthesis;unbounded knapsack problem;Circuits;Computational modeling;Control system synthesis;Difference equations;NP-hard problem;Nonlinear equations;Phase change random access memory;Sufficient conditions;Systolic arrays;Very large scale integration}, doi={10.1109/71.595572}, ISSN={1045-9219},}

@inproceedings{Hoefer:2015:SBP:2807591.2807644, author = {Hoefer, Torsten and Belli, Roberto}, title = {Scientific Benchmarking of Parallel Computing Systems: Twelve Ways to Tell the Masses when Reporting Performance Results}, booktitle = {Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis}, series = {SC '15}, year = {2015}, isbn = {978-1-4503-3723-6}, location = {Austin, Texas}, pages = {73:1-73:12}, articleno = {73}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2807591.2807644>}, doi = {10.1145/2807591.2807644}, acmid = {2807644}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {benchmarking, data analysis, parallel computing, statistics}, }

@inproceedings{Bender:2016:CA:2935764.2935798, author = {Bender, Michael A. and Demaine, Erik D. and Ebrahimi, Roozbeh and Fineman, Jeremy T. and Johnson, Rob and Lincoln, Andrea and Lynch, Jayson and McCauley, Samuel}, title = {Cache-Adaptive Analysis}, booktitle = {Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures}, series = {SPAA '16}, year = {2016}, isbn = {978-1-4503-4210-0}, location = {Pacific Grove, California, USA}, pages = {135-144}, numpages = {10}, url = {<http://doi.acm.org/10.1145/2935764.2935798>}, doi = {10.1145/2935764.2935798}, acmid = {2935798}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {cache adaptive, cache oblivious, external memory, memory adaptive, paging algorithms}, }

@inproceedings{Bhaskaracharya:2016:SIA:2837614.2837636, author = {Bhaskaracharya, Somashekaracharya G. and Bondhugula, Uday and Cohen, Albert}, title = {SMO: An Integrated Approach to Intra-array and Inter-array Storage Optimization}, booktitle = {Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages}, series = {POPL '16}, year = {2016}, isbn = {978-1-4503-3549-2}, location = {St. Petersburg, FL, USA}, pages = {526-538}, numpages = {13}, url = {<http://doi.acm.org/10.1145/2837614.2837636>}, doi = {10.1145/2837614.2837636}, acmid = {2837636}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Compilers, array contraction, memory optimization, polyhedral framework, storage mapping optimization}, }

@INPROCEEDINGS{7429307, author={P. Chatarasi and J. Shirako and V. Sarkar}, booktitle={2015 International Conference on Parallel Architecture and Compilation (PACT)}, title={Polyhedral Optimizations of Explicitly Parallel Programs}, year={2015}, pages={213-226}, doi={10.1109/PACT.2015.44}, ISSN={1089-795X}, month={Oct},}

@INPROCEEDINGS{7429326, author={V. Porpodas and T. M. Jones}, booktitle={2015 International Conference on Parallel Architecture and Compilation (PACT)}, title={Throttling Automatic Vectorization: When Less is More}, year={2015}, pages={432-444}, doi={10.1109/PACT.2015.32},

Last update: 2016/08/30 12:59 melange:papers:spring2015 <https://www.cs.colostate.edu/AlphaZ/wiki/doku.php?id=melange:papers:spring2015&rev=1472583577>

---

ISSN={1089-795X}, month={Oct},}

From: <https://www.cs.colostate.edu/AlphaZ/wiki/> - **AlphaZ**

Permanent link: <https://www.cs.colostate.edu/AlphaZ/wiki/doku.php?id=melange:papers:spring2015&rev=1472583577>

Last update: **2016/08/30 12:59**

