

@inproceedings{Acharya:2015:PNM:2688500.2688512, author = {Acharya, Aravind and Bondhugula, Uday}, title = {PLUTO+: Near-complete Modeling of Affine Transformations for Parallelism and Locality}, booktitle = {Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming}, series = {PPoPP 2015}, year = {2015}, isbn = {978-1-4503-3205-7}, location = {San Francisco, CA, USA}, pages = {54-64}, numpages = {11}, url = {<http://doi.acm.org/10.1145/2688500.2688512>}, doi = {10.1145/2688500.2688512}, acmid = {2688512}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Affine transformations, affine scheduling, automatic parallelization, polyhedral model, stencil computations, tiling}, }

@INPROCEEDINGS{7161519, author={Tithi, J.J. and Ganapathi, P. and Talati, A. and Aggarwal, S. and Chowdhury, R.}, booktitle={Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International}, title={High-Performance Energy-Efficient Recursive Dynamic Programming with Matrix-Multiplication-Like Flexible Kernels}, year={2015}, month={May}, pages={303-312}, keywords={divide and conquer methods;dynamic programming;mathematics computing;matrix multiplication;parallel algorithms;DP problem;FW-APSP;Floyd-Warshall all-pairs shortest path;cache-oblivious recursive divide-and-conquer;dynamic programming;gap penalty;high-performing parallel implementation;matrix-multiplication-like flexible kernel;optimization;parallel CORDAC algorithm;cache-oblivious;divide-and-conquer;dynamic programming;flexible kernel;polyhedral compiler;recursive}, doi={10.1109/IPDPS.2015.107}, ISSN={1530-2075}, }

@inproceedings{Bondhugula:2014:TOT:2628071.2628106, author = {Bondhugula, Uday and Bandishti, Vinayaka and Cohen, Albert and Potron, Guillain and Vasilache, Nicolas}, title = {Tiling and Optimizing Time-iterated Computations on Periodic Domains}, booktitle = {Proceedings of the 23rd International Conference on Parallel Architectures and Compilation}, series = {PACT '14}, year = {2014}, isbn = {978-1-4503-2809-8}, location = {Edmonton, AB, Canada}, pages = {39-50}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2628071.2628106>}, doi = {10.1145/2628071.2628106}, acmid = {2628106}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {automatic parallelization, periodic, polyhedral model, stencils, tiling}, }

@inproceedings{Tang:2015:CWI:2688500.2688514, author = {Tang, Yuan and You, Ronghui and Kan, Haibin and Tithi, Jesmin Jahan and Ganapathi, Pramod and Chowdhury, Rezaul A.}, title = {Cache-oblivious Wavefront: Improving Parallelism of Recursive Dynamic Programming Algorithms Without Losing Cache-efficiency}, booktitle = {Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming}, series = {PPoPP 2015}, year = {2015}, isbn = {978-1-4503-3205-7}, location = {San Francisco, CA, USA}, pages = {205-214}, numpages = {10}, url = {<http://doi.acm.org/10.1145/2688500.2688514>}, doi = {10.1145/2688500.2688514}, acmid = {2688514}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Cilk, cache-oblivious parallel algorithms, cache-oblivious wavefront, dynamic programming, multi-core, nested parallel computation}, }

@inproceedings{Huang:2014:ARD:2628071.2628089, author = {Huang, Cheng-Chieh and Nagarajan, Vijay}, title = {ATCache: Reducing DRAM Cache Latency via a Small SRAM Tag Cache}, booktitle = {Proceedings of the 23rd International Conference on Parallel Architectures and Compilation}, series = {PACT '14}, year = {2014}, isbn = {978-1-4503-2809-8}, location = {Edmonton, AB, Canada}, pages = {51-60}, numpages = {10}, url = {<http://doi.acm.org/10.1145/2628071.2628089>}, doi = {10.1145/2628071.2628089}, acmid = {2628089}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {design, dram cache, performance}, }

@inproceedings{Fatehi:2014:ITS:2628071.2628093, author = {Fatehi, Ehsan and Gratz, Paul}, title = {ILP and TLP in Shared Memory Applications: A Limit Study}, booktitle = {Proceedings of the 23rd International Conference on Parallel Architectures and Compilation}, series = {PACT '14}, year =

{2014}, isbn = {978-1-4503-2809-8}, location = {Edmonton, AB, Canada}, pages = {113-126}, numpages = {14}, url = {<http://doi.acm.org/10.1145/2628071.2628093>}, doi = {10.1145/2628071.2628093}, acmid = {2628093}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {instruction-level parallelism (ilp), limits, pthreads, thread-level parallelism (tlp)}, }

@inproceedings{Cameron:2014:BDP:2628071.2628079, author = {Cameron, Robert D. and Shermer, Thomas C. and Shriraman, Arrvindh and Herdy, Kenneth S. and Lin, Dan and Hull, Benjamin R. and Lin, Meng}, title = {Bitwise Data Parallelism in Regular Expression Matching}, booktitle = {Proceedings of the 23rd International Conference on Parallel Architectures and Compilation}, series = {PACT '14}, year = {2014}, isbn = {978-1-4503-2809-8}, location = {Edmonton, AB, Canada}, pages = {139-150}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2628071.2628079>}, doi = {10.1145/2628071.2628079}, acmid = {2628079}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {parallel bit streams, regular expression matching}, }

@inproceedings{Ansel:2014:OEF:2628071.2628092, author = {Ansel, Jason and Kamil, Shoaib and Veeramachaneni, Kalyan and Ragan-Kelley, Jonathan and Bosboom, Jeffrey and O'Reilly, Una-May and Amarasinghe, Saman}, title = {OpenTuner: An Extensible Framework for Program Autotuning}, booktitle = {Proceedings of the 23rd International Conference on Parallel Architectures and Compilation}, series = {PACT '14}, year = {2014}, isbn = {978-1-4503-2809-8}, location = {Edmonton, AB, Canada}, pages = {303-316}, numpages = {14}, url = {<http://doi.acm.org/10.1145/2628071.2628092>}, doi = {10.1145/2628071.2628092}, acmid = {2628092}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {autotuner, optimization}, }

@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}, 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}, }

@INPROCEEDINGS{5377644, author={Shafiq, M. and Pericas, M. and de la Cruz, R. and Araya-Polo, M. and Navarro, N. and Ayguade, E.}, booktitle={Field-Programmable Technology, 2009. FPT 2009.}, title={Exploiting memory customization in FPGA for 3D stencil computations}, year={2009}, month={Dec}, pages={38-45}, keywords={field programmable gate arrays;signal processing;3D stencil computations;FPGA;IBM PowerXCell 8i;data reuse;memory customization;memory organization;memory-bound kernels;Bandwidth;Computer applications;Field programmable gate arrays;Finite difference methods;Finite impulse response filter;Hardware;Kernel;Nearest neighbor searches;Throughput;Time domain analysis}, doi={10.1109/FPT.2009.5377644}, }

@inproceedings{Wahib:2015:AGK:2749246.2749255, author = {Wahib, Mohamed and Maruyama, Naoya}, title = {Automated GPU Kernel Transformations in Large-Scale Production Stencil Applications}, booktitle = {Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing}, series = {HPDC '15}, year = {2015}, isbn = {978-1-4503-3550-8}, location = {Portland, Oregon, USA}, pages = {259-270}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2749246.2749255>}, doi = {10.1145/2749246.2749255}, acmid = {2749255}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {cuda, gpu, source-to-

source translation, stencil computations}, }

@inproceedings{Benson:2015:FPP:2688500.2688513, author = {Benson, Austin R. and Ballard, Grey}, title = {A Framework for Practical Parallel Fast Matrix Multiplication}, booktitle = {Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming}, series = {PPoPP 2015}, year = {2015}, isbn = {978-1-4503-3205-7}, location = {San Francisco, CA, USA}, pages = {42-53}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2688500.2688513>}, doi = {10.1145/2688500.2688513}, acmid = {2688513}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {dense linear algebra, fast matrix multiplication, parallel linear algebra, shared memory}, }

@inproceedings{Sukhija:2014:PSR:2672598.2672904, author = {Sukhija, Nitin and Malone, Brandon and Srivastava, Srishti and Banicescu, Ioana and Ciorba, Florina M.}, title = {Portfolio-Based Selection of Robust Dynamic Loop Scheduling Algorithms Using Machine Learning}, booktitle = {Proceedings of the 2014 IEEE International Parallel & Distributed Processing Symposium Workshops}, series = {IPDPSW '14}, year = {2014}, isbn = {978-1-4799-4116-2}, pages = {1638-1647}, numpages = {10}, url = {<http://dx.doi.org/10.1109/IPDPSW.2014.183>}, doi = {10.1109/IPDPSW.2014.183}, acmid = {2672904}, publisher = {IEEE Computer Society}, address = {Washington, DC, USA}, keywords = {Dynamic loop scheduling, robustness, algorithm selection, empirical robustness prediction models, machine learning techniques, variable system availability}, }

@incollection{Tiwari2014, year={2014}, isbn={978-3-319-09872-2}, booktitle={Euro-Par 2014 Parallel Processing}, series={Euro-Par 2014 Parallel Processing}, volume={8632}, editor={Silva, Fernando and Dutra, Ins and Santos Costa, Vtor}, doi={10.1007/978-3-319-09873-9\_6}, title={Modeling the Impact of Reduced Memory Bandwidth on HPC Applications}, url={[http://dx.doi.org/10.1007/978-3-319-09873-9\\_6](http://dx.doi.org/10.1007/978-3-319-09873-9_6)}, publisher={Springer International Publishing}, author={Tiwari, Ananta and Gamst, Anthony and Laurenzano, MichaelA. and Schulz, Martin and Carrington, Laura}, pages={63-74}, language={English} }

@INPROCEEDINGS{7056046, author={Agarwal, N. and Nellans, D. and O'Connor, M. and Keckler, S.W. and Wensch, T.F.}, booktitle={High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium}, title={Unlocking bandwidth for GPUs in CC-NUMA systems}, year={2015}, month={Feb}, pages={354-365}, keywords={cache storage;graphics processing units;parallel processing;storage management;CC-NUMA GPU-CPU systems;CPU memory bandwidth;GDDR memory;GPU kernel;GPU memory bandwidth;GPU relaxed memory semantics;GPU-based HPC applications;aggressive memory prefetching;bandwidth balancing;hardware cache-coherence;memory-intensive GPU workloads;minimal hardware support;on-demand software page migration;oracular page placement;software runtime system;virtual address-based program locality;Bandwidth;Graphics processing units;Hardware;Memory management;Random access memory;Runtime}, doi={10.1109/HPCA.2015.7056046}, }

@INPROCEEDINGS{6270616, author={Changyou Zhang and Kun Huang and Xiang Cui and Yifeng Chen}, booktitle={Parallel and Distributed Processing Symposium Workshops PhD Forum (IPDPSW), 2012 IEEE 26th International}, title={Power-aware Programming with GPU Accelerators}, year={2012}, month={May}, pages={2443-2449}, keywords={graphics processing units;multi-threading;multiprocessing systems;power aware computing;ubiquitous computing;GPU accelerators;high-level program development;manycore processor;multithreaded processor;on-chip parallelism;parallel processor;power consumption values;power efficiency;power estimation;power-aware programming;processor computational power;processor memory bandwidth;program statements;ubiquitous computing;Bandwidth;Graphics processing unit;Hardware;Memory management;Message systems;Power demand;Power measurement;GPU;Power-aware;Primitive;Programming}, doi={10.1109/IPDPSW.2012.301}, }

@inproceedings{Fang:2014:TIX:2568088.2576799, author = {Fang, Jianbin and Sips, Henk and Zhang, LiLun and Xu, Chuanfu and Che, Yonggang and Varbanescu, Ana Lucia}, title = {Test-driving Intel Xeon Phi}, booktitle = {Proceedings of the 5th ACM/SPEC International Conference on Performance Engineering}, series = {ICPE '14}, year = {2014}, isbn = {978-1-4503-2733-6}, location = {Dublin, Ireland}, pages = {137-148}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2568088.2576799>}, doi = {10.1145/2568088.2576799}, acmid = {2576799}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {experience with xeon phi, microbenchmarking, optimization, performance analysis}, }

@inproceedings{Ravishankar:2015:DMC:2688500.2688515, author = {Ravishankar, Mahesh and Dathathri, Roshan and Elango, Venmugil and Pouchet, Louis-Noël and Ramanujam, J. and Rountev, Atanas and Sadayappan, P.}, title = {Distributed Memory Code Generation for Mixed Irregular/Regular Computations}, booktitle = {Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming}, series = {PPoPP 2015}, year = {2015}, isbn = {978-1-4503-3205-7}, location = {San Francisco, CA, USA}, pages = {65-75}, numpages = {11}, url = {<http://doi.acm.org/10.1145/2688500.2688515>}, doi = {10.1145/2688500.2688515}, acmid = {2688515}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Distributed Memory, Inspector/Executor, Irregular Computation, Polyhedral Compilation}, }

@inproceedings{Cong:2014:OMS:2593069.2593090, author = {Cong, Jason and Li, Peng and Xiao, Bingjun and Zhang, Peng}, title = {An Optimal Microarchitecture for Stencil Computation Acceleration Based on Non-Uniform Partitioning of Data Reuse Buffers}, booktitle = {Proceedings of the 51st Annual Design Automation Conference}, series = {DAC '14}, year = {2014}, isbn = {978-1-4503-2730-5}, location = {San Francisco, CA, USA}, pages = {77:1-77:6}, articleno = {77}, numpages = {6}, url = {<http://doi.acm.org/10.1145/2593069.2593090>}, doi = {10.1145/2593069.2593090}, acmid = {2593090}, publisher = {ACM}, address = {New York, NY, USA}, }

@ARTICLE{6470606, author={Sano, K. and Hatsuda, Y. and Yamamoto, S.}, journal={Parallel and Distributed Systems, IEEE Transactions}, title={Multi-FPGA Accelerator for Scalable Stencil Computation with Constant Memory Bandwidth}, year={2014}, month={March}, volume={25}, number={3}, pages={695-705}, keywords={field programmable gate arrays;parallel processing;storage management;CCM;GPU;Jacobi computation;SSA;custom computing machine;deep pipelining approach;domain-specific programmable concept;field programmable gate array;graphics processing unit;high-performance stencil computations;memory bandwidth;multiFPGA accelerator;multicore microprocessors;scalable stencil computation;scalable streaming-array;scientific computations;Arrays;Bandwidth;Computational modeling;Field programmable gate arrays;Hardware;Scalability;FPGA;Scalable streaming-array;custom computing machine;high-performance computation;stencil computation}, doi={10.1109/TPDS.2013.51}, ISSN={1045-9219}, }

@inproceedings{Fan:2006:IHE:1176254.1176322, author = {Fan, Kevin and Kudlur, Manjunath and Park, Hyunchul and Mahlke, Scott}, title = {Increasing Hardware Efficiency with Multifunction Loop Accelerators}, booktitle = {Proceedings of the 4th International Conference on Hardware/Software Codesign and System Synthesis}, series = {CODES+ISSS '06}, year = {2006}, isbn = {1-59593-370-0}, location = {Seoul, Korea}, pages = {276-281}, numpages = {6}, url = {<http://doi.acm.org/10.1145/1176254.1176322>}, doi = {10.1145/1176254.1176322}, acmid = {1176322}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {application-specific hardware, high-level synthesis, loop accelerator, modulo scheduling, multifunction design}, }

@article{Meswani:2013:MPP:2493921.2493922, author = {Meswani, Mitesh R. and Carrington, Laura and Unat, Didem and Snively, Allan and Baden, Scott and Poole, Stephen}, title = {Modeling and Predicting Performance of High Performance Computing Applications on Hardware Accelerators},



```
journal = {Int. J. High Perform. Comput. Appl.}, issue_date = {May 2013}, volume = {27}, number =
{2}, month = may, year = {2013}, issn = {1094-3420}, pages = {89-108}, numpages = {20}, url =
{http://dx.doi.org/10.1177/1094342012468180}, doi = {10.1177/1094342012468180}, acmid =
{2493922}, publisher = {Sage Publications, Inc.}, address = {Thousand Oaks, CA, USA}, keywords =
{FPGA, GPU, HPC, accelerators, benchmarking, idioms, performance modeling, performance
prediction}, }
```

```
@article{Nery:2013:HRM:2537182.2537569, author = {Nery, Alexandre S. and Jozwiak, Lech and
Lindwer, Menno and Cocco, Mauro and Nedjah, Nadia and Franca, Felipe M. G.}, title = {Hardware
Reuse in Modern Application-specific Processors and Accelerators}, journal = {Microprocess.
Microsyst.}, issue_date = {August, 2013}, volume = {37}, number = {6-7}, month = aug, year =
{2013}, issn = {0141-9331}, pages = {684-692}, numpages = {9}, url =
{http://dx.doi.org/10.1016/j.micpro.2012.06.005}, doi = {10.1016/j.micpro.2012.06.005}, acmid =
{2537569}, publisher = {Elsevier Science Publishers B. V.}, address = {Amsterdam, The
Netherlands, The Netherlands}, keywords = {Application-specific processors, Area reduction,
Hardware accelerator, Power reduction, Resource sharing}, }
```

```
@inproceedings{Bandishti:2012:TSC:2388996.2389051, author = {Bandishti, Vinayaka and
Panailath, Irshad and Bondhugula, Uday}, title = {Tiling Stencil Computations to Maximize
Parallelism}, booktitle = {Proceedings of the International Conference on High Performance
Computing, Networking, Storage and Analysis}, series = {SC '12}, year = {2012}, isbn =
{978-1-4673-0804-5}, location = {Salt Lake City, Utah}, pages = {40:1-40:11}, articleno = {40},
numpages = {11}, url = {http://dl.acm.org/citation.cfm?id=2388996.2389051}, acmid = {2389051},
publisher = {IEEE Computer Society Press}, address = {Los Alamitos, CA, USA}, keywords =
{compilers, program transformation}, }
```

```
@inproceedings{Wonnacott13,
```

```
  Author = {Dave G. Wonnacott and Michelle Mills Strout},
  Booktitle = {Proceedings of the 3rd International Workshop on Polyhedral
Compilation Techniques (IMPACT)},
  series = {IMPACT 2013},
  Month = {January},
  Title = {On the Scalability of Loop Tiling Techniques},
  Year = {2013},
  url={http://impact.gforge.inria.fr/impact2013/papers/impact2013\_on\_the\_scala
bility\_of\_loop\_tiling\_techniques.pdf}
```

```
}
```

```
@inproceedings{Pugh:1991:OTF:125826.125848, author = {Pugh, William}, title = {The Omega Test:
A Fast and Practical Integer Programming Algorithm for Dependence Analysis}, booktitle =
{Proceedings of the 1991 ACM/IEEE Conference on Supercomputing}, series = {Supercomputing
'91}, year = {1991}, isbn = {0-89791-459-7}, location = {Albuquerque, New Mexico, USA}, pages =
{4-13}, numpages = {10}, url = {http://doi.acm.org/10.1145/125826.125848}, doi =
{10.1145/125826.125848}, acmid = {125848}, publisher = {ACM}, address = {New York, NY, USA},
}
```

```
@inproceedings{cohenautomatic,
```

```
  title={Automatic Intra-Array Storage Optimization},
  author={Bhaskaracharya, Somashekaracharya G and Bondhugula, Uday and Cohen,
```

```
Albert },
url={http://www.csa.iisc.ernet.in/TR/2014/3/paper.pdf},
series={IISc-CSA-TR-2014-3, Nov 2014 and submitted to ACM TOPLAS, Feb 2015},
year={2015},
}
```

@inproceedings{Stock:2014:FED:2594291.2594342, author = {Stock, Kevin and Kong, Martin and Grosser, Tobias and Pouchet, Louis-Noël and Rastello, Fabrice and Ramanujam, J. and Sadayappan, P.}, title = {A Framework for Enhancing Data Reuse via Associative Reordering}, booktitle = {Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation}, series = {PLDI '14}, year = {2014}, isbn = {978-1-4503-2784-8}, location = {Edinburgh, United Kingdom}, pages = {65-76}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2594291.2594342>}, doi = {10.1145/2594291.2594342}, acmid = {2594342}, publisher = {ACM}, address = {New York, NY, USA}, }

@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}, }

From:

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

Permanent link:

<https://www.cs.colostate.edu/AlphaZ/wiki/doku.php?id=melange:papers:fall2015>

Last update: **2015/11/10 09:40**

