

@Article{Griebel2000, author="Griebel, Martin and Feautrier, Paul and Lengauer, Christian", title="Index Set Splitting", journal="International Journal of Parallel Programming", year="2000", month="Dec", day="01", volume="28", number="6", pages="607-631", abstract="There are many algorithms for the space-time mapping of nested loops. Some of them even make the optimal choices within their framework. We propose a preprocessing phase for algorithms in the polytope model, which extends the model and yields space-time mappings whose schedule is, in some cases, orders of magnitude faster. These are cases in which the dependence graph has small irregularities. The basic idea is to split the index set of the loop nests into parts with a regular dependence structure and apply the existing space-time mapping algorithms to these parts individually. This work is based on a seminal idea in the more limited context of loop parallelization at the code level. We elevate the idea to the model level (our model is the polytope model), which increases its applicability by providing a clearer and wider range of choices at an acceptable analysis cost. Index set splitting is one facet in the effort to extend the power of the polytope model and to enable the generation of competitive target code.", issn="1573-7640", doi="10.1023/A:1007516818651", url="https://doi.org/10.1023/A:1007516818651" }

@inproceedings{Irigoin:1988:SP:73560.73588, author = {Irigoin, F. and Triolet, R.}, title = {Supernode Partitioning}, booktitle = {Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages}, series = {POPL '88}, year = {1988}, isbn = {0-89791-252-7}, location = {San Diego, California, USA}, pages = {319-329}, numpages = {11}, url = {http://doi.acm.org/10.1145/73560.73588}, doi = {10.1145/73560.73588}, acmid = {73588}, publisher = {ACM}, address = {New York, NY, USA}, }

@ARTICLE{7738524, author={Y. H. Chen and T. Krishna and J. S. Emer and V. Sze}, journal={IEEE Journal of Solid-State Circuits}, title={Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks}, year={2017}, volume={52}, number={1}, pages={127-138}, keywords={DRAM chips;data flow computing;energy conservation;feedforward neural nets;learning (artificial intelligence);neural net architecture;power aware computing;reconfigurable architectures;AI systems;AlexNet;CNN shapes;DRAM accesses;Eyeriss;MAC;RS dataflow reconfiguration;accelerator chip;convolutional layers;data movement energy cost;dataflow processing;deep convolutional neural networks;energy efficiency;energy-efficient reconfigurable accelerator;multiply and accumulation;off-chip DRAM;reconfiguring architecture;row stationary;spatial architecture;Clocks;Computer architecture;Hardware;Neural networks;Random access memory;Shape;Throughput;Convolutional neural networks (CNNs);dataflow processing;deep learning;energy-efficient accelerators;spatial architecture}, doi={10.1109/JSSC.2016.2616357}, ISSN={0018-9200}, month={Jan},}

@ARTICLE{88484, author={J. H. Saltz and R. Mirchandaney and K. Crowley}, journal={IEEE Transactions on Computers}, title={Run-time parallelization and scheduling of loops}, year={1991}, volume={40}, number={5}, pages={603-612}, keywords={parallel programming;scheduling;Encore Multimax;automatic parallelization;automatic scheduling;compile-time information;concurrently executable loop iterations;do loop;execution time preprocessing;executors;inspector procedures;loop dependency analysis;loop indexes;run-time methods;run-time reordering;source code loop structures;symbolic transformation rules;transformed versions;wavefronts;Computer science;Costs;Failure analysis;Level set;NASA;Parallel processing;Performance analysis;Processor scheduling;Runtime}, doi={10.1109/12.88484}, ISSN={0018-9340}, month={May},}

BibTeX | EndNote | ACM Ref @article{Feautrier:2006:SSS:1217445.1217447, author = {Feautrier, Paul}, title = {Scalable and Structured Scheduling}, journal = {Int. J. Parallel Program.}, issue_date = {October 2006}, volume = {34}, number = {5}, month = oct, year = {2006}, issn = {0885-7458}, pages = {459-487}, numpages = {29}, url = {http://dx.doi.org/10.1007/s10766-006-0011-4}, doi =

{10.1007/s10766-006-0011-4}, acmid = {1217447}, publisher = {Kluwer Academic Publishers}, address = {Norwell, MA, USA}, keywords = {automatic parallelization, scalability, structured scheduling}, }

@article{Verdoolaege:2012:ECS:2362389.2362390, author = {Verdoolaege, Sven and Janssens, Gerda and Bruynooghe, Maurice}, title = {Equivalence Checking of Static Affine Programs Using Widening to Handle Recurrences}, journal = {ACM Trans. Program. Lang. Syst.}, issue_date = {October 2012}, volume = {34}, number = {3}, month = nov, year = {2012}, issn = {0164-0925}, pages = {11:1-11:35}, articleno = {11}, numpages = {35}, url = {<http://doi.acm.org/10.1145/2362389.2362390>}, doi = {10.1145/2362389.2362390}, acmid = {2362390}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Commutativity, equivalence checking, polytope model, recurrences, widening}, }

@inproceedings{Kulkarni:2007:OPR:1250734.1250759, author = {Kulkarni, Milind and Pingali, Keshav and Walter, Bruce and Ramanarayanan, Ganesh and Bala, Kavita and Chew, L. Paul}, title = {Optimistic Parallelism Requires Abstractions}, booktitle = {Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation}, series = {PLDI '07}, year = {2007}, isbn = {978-1-59593-633-2}, location = {San Diego, California, USA}, pages = {211-222}, numpages = {12}, url = {<http://doi.acm.org/10.1145/1250734.1250759>}, doi = {10.1145/1250734.1250759}, acmid = {1250759}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {abstractions, irregular programs, optimistic parallelism}, }

@inproceedings{sanjay-fsttcs86,

Address = {New Delhi, India},
Author = {Rajopadhye, S. V. and Purushothaman, S. and Fujimoto, R. M.},
Booktitle = {Proceedings, Sixth Conference on Foundations of Software Technology and Theoretical Computer Science},
Key = {Rajopadhye86b},
Month = {December},
Pages = {488-503},
Publisher = {Springer Verlag, LNCS~241},
Title = {On Synthesizing Systolic Arrays from Recurrence Equations with Linear Dependencies},
Year = {1986}}

@article{quinton-jvsp89,

Author = {Quinton, P. and {Van Dongen}, V.},
Journal = {Journal of {VLSI} Signal Processing},
Number = 2,
Pages = {95-113},
Publisher = {Kluwer Academic Publishers, Boston},
Title = {The Mapping of Linear Recurrence Equations on Regular Arrays},
Volume = 1,
Year = 1989}

```
@Article{feautrier91, author= "Feautrier, P.", title= "Dataflow analysis of array and scalar
references", journal= "International Journal of Parallel Programming", year= 1991, volume= 20,
number= 1, pages= "23-53", month= "Feb", annote= "This article explains how a simple imperative
language
```

```
program (consisting only of assignments, for loops with affine loop
limits, and arrays with affine index expressions), can be statically
analyzed to find the flow dependencies."
```

```
}
```

```
@article{feautrier92a,
```

```
author    = {Paul Feautrier},
title     = {Some Efficient Solutions to the Affine Scheduling Problem.
             {Part I}. One-dimensional Time},
journal   = {International Journal of Parallel Programming},
volume    = {21},
number    = {5},
year      = {1992},
pages     = {313-347},
ee        = {http://dx.doi.org/10.1007/BF01407835}
```

```
}
```

```
@article{feautrier92b,
```

```
author    = {Paul Feautrier},
title     = {Some Efficient Solutions to the Affine Scheduling Problem.
             {Part II}. Multidimensional Time},
journal   = {International Journal of Parallel Programming},
volume    = {21},
number    = {6},
year      = {1992},
pages     = {389-420},
ee        = {http://dx.doi.org/10.1007/BF01379404}
```

```
}
```

```
@Article{collard-et-al-fuzzy-jpdc-1997,
```

```
author =      {Collard, J-F. and Barthou, D. and Feautrier, P.},
title =      {Fuzzy Array Data Flow Analysis},
journal =    {Journal of Parallel and Distributed Computing},
year =      1997,
volume =    40,
number =    2,
pages =     {210-226},
month =     {Feb}}
```

```
@InProceedings{BPCB10,
```

```
author = {Benabderrahmane, M.-W. and Pouchet, L.-N. and Cohen A. and Bastoul, C.},
title = {The Polyhedral Model Is More Widely Applicable Than You Think},
booktitle = {Proceedings of the International Conference on Compiler Construction ({ETAPS CC'10})},
year = 2010,
series = {LNCS},
address = {Paphos, Cyprus},
pages = {283--303},
month = Mar,
publisher = {Springer-Verlag},
```

@ARTICAL{123, author={Sanket Tavarageri, Albert Hartono, Muthu Baskaran, Louis-Noel Pouchet, J. Ramanujam and P. Sadayappan}, title={Parametric Tiling of Affine Loop Nests}, year={2010}, doi = {<http://web.cse.ohio-state.edu/~pouchet.2/doc/cpc-article.10.pdf>}, }

@INPROCEEDINGS{6043234, author={A. Pedram and A. Gerstlauer and R. A. v. d. Geijn}, booktitle={ASAP 2011 - 22nd IEEE International Conference on Application-specific Systems, Architectures and Processors}, title={A high-performance, low-power linear algebra core}, year={2011}, pages={35-42}, keywords={floating point arithmetic;matrix multiplication;GFLOPS-W;application-specific custom hardware;floating point operations per second;linear algebra core;matrix computations;matrix-matrix multiplication;power consumption reduction;technology scaling;Bandwidth;Computer architecture;Hardware;Kernel;Linear algebra;Program processors;Registers}, doi={10.1109/ASAP.2011.6043234}, ISSN={1063-6862}, month={Sept}, }

@inproceedings{Bandishti:2012:TSC:2388996.2389051, author = {Bandishti, Vinayaka and Pananilath, 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}, }

@ARTICLE{7582549, author={U. Bondhugula and V. Bandishti and I. Pananilath}, journal={IEEE Transactions on Parallel and Distributed Systems}, title={Diamond Tiling: Tiling Techniques to Maximize Parallelism for Stencil Computations}, year={2016}, url={<http://ieeexplore.ieee.org/document/7582549/>}, volume={PP}, number={99}, pages={1-1}, keywords={Diamond;Face;Indexes;Optimization;Parallel processing;Shape;Silicon;Compilers;locality;loop tiling;parallelism;program transformation;stencils}, doi={10.1109/TPDS.2016.2615094}, ISSN={1045-9219}, month={}, }

@ARTICLE{7155440, author={T. Nowatzki and J. Menon and C. H. Ho and K. Sankaralingam}, journal={IEEE Micro}, title={Architectural Simulators Considered Harmful}, year={2015}, url={<http://ieeexplore.ieee.org/document/7155440/>}, volume={35}, number={6}, pages={4-12}, keywords={computer architecture;digital simulation;architectural layers;architectural simulators;black boxes;evaluation standard recalibration;quantitative simulators;Analytical models;Benchmark testing;Computer architecture;Market research;Mathematical model;Simulation;architecture;benchmarks;evaluation

standards;footprint;modeling;simulators;validation}, doi={10.1109/MM.2015.74},
ISSN={0272-1732}, month={Nov},}

@INPROCEEDINGS{7349585, author={J. D. Garvey and T. S. Abdelrahman}, booktitle={2015 44th International Conference on Parallel Processing}, title={Automatic Performance Tuning of Stencil Computations on GPUs}, year={2015}, pages={300-309}, keywords={graphics processing units;learning (artificial intelligence);parallel processing;storage management;Nvidia GTX Titan GPU;OpenCL stencil kernel;automatic performance tuning;graphics processing unit;machine learning;optimization;random sampling;stencil computation;Graphics processing units;Instruction sets;Kernel;Merging;Optimization;Parallel processing;Yttrium;GPGPU;auto-tuning;machine learning;stencil}, doi={10.1109/ICPP.2015.39}, ISSN={0190-3918}, month={Sept},}

@article{DBLPSteve,

```
author      = {Sharan Chetlur and
               Cliff Woolley and
               Philippe Vandermersch and
               Jonathan Cohen and
               John Tran and
               Bryan Catanzaro and
               Evan Shelhamer},
title       = {cuDNN: Efficient Primitives for Deep Learning},
journal     = {CoRR},
volume     = {abs/1410.0759},
year       = {2014},
url        = {http://arxiv.org/abs/1410.0759},
timestamp  = {Sun, 02 Nov 2014 11:25:59 +0100},
biburl     =
{http://dblp.uni-trier.de/rec/bib/journals/corr/ChetlurWVCTCS14},
bibsource  = {dblp computer science bibliography, http://dblp.org}
```

}

@article{Bao:2016:SDF:3012405.3011017, author = {Bao, Wenlei and Hong, Changwan and Chunduri, Sudheer and Krishnamoorthy, Sriram and Pouchet, Louis-Noel and Rastello, Fabrice and Sadayappan, P.}, title = {Static and Dynamic Frequency Scaling on Multicore CPUs}, journal = {ACM Trans. Archit. Code Optim.}, issue_date = {December 2016}, volume = {13}, number = {4}, month = dec, year = {2016}, issn = {1544-3566}, pages = {51:1-51:26}, articleno = {51}, numpages = {26}, url = {<http://doi.acm.org/10.1145/3011017>}, doi = {10.1145/3011017}, acmid = {3011017}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {Affine Programs, CPU Energy, Static Analysis, Voltage and Frequency Scaling}, }

@inproceedings{Pouchet:2013:PDR:2435264.2435273, author = {Pouchet, Louis-Noel and Zhang, Peng and Sadayappan, P. and Cong, Jason}, title = {Polyhedral-based Data Reuse Optimization for Configurable Computing}, booktitle = {Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays}, series = {FPGA '13}, year = {2013}, isbn = {978-1-4503-1887-7}, location = {Monterey, California, USA}, pages = {29-38}, numpages = {10}, url = {<http://doi.acm.org/10.1145/2435264.2435273>}, doi = {10.1145/2435264.2435273}, acmid = {2435273}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {compilation, data reuse, high-level synthesis, program transformations}, }

@article{Kong:2013:PTM:2499370.2462187, author = {Kong, Martin and Veras, Richard and Stock,

Kevin and Franchetti, Franz and Pouchet, Louis-No^el and Sadayappan, P.}, title = {When Polyhedral Transformations Meet SIMD Code Generation}, journal = {SIGPLAN Not.}, issue_date = {June 2013}, volume = {48}, number = {6}, month = jun, year = {2013}, issn = {0362-1340}, pages = {127-138}, numpages = {12}, url = {<http://doi.acm.org/10.1145/2499370.2462187>}, doi = {10.1145/2499370.2462187}, acmid = {2462187}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {affine scheduling, autotuning, compiler optimization, loop transformations, program synthesis}, }

@article{cummins2017synthesizing,

title={Synthesizing benchmarks for predictive modeling},

author={Cummins, Chris and Petoumenos, Pavlos and Wang, Zheng and Leather, Hugh},

year={2017},

url={<http://homepages.inf.ed.ac.uk/hleather/publications/2017-benchsynth-cgo.pdf>}

}

@article{optimistic2017,

title={Optimistic Loop Optimization},

author={Doerfert, Johannes and Grosser, Tobias and Hack, Sebastian},

year={2017}

}

@inbook{e0cc7363fd684a529d1ba82b8195d530,

title = "Minimizing the cost of iterative compilation with active learning",
keywords = "Active Learning, Compilers, Iterative Compilation, Machine Learning, Sequential Analysis;",
author = "William Ogilvie and Pavlos Petoumenos and Zheng Wang and Hugh Leather",
note = "Date of Acceptance: 25/10/2016",
year = "2016",
month = "10",
booktitle = "The International Symposium on Code Generation and Optimization (CGO) 2017",

}

BibTeX | EndNote | ACM Ref

@inproceedings{Putnam:2014:RFA:2665671.2665678, author = {Putnam, Andrew and Caulfield, Adrian M. and Chung, Eric S. and Chiou, Derek and Constantinides, Kypros and Demme, John and Esmailzadeh, Hadi and Fowers, Jeremy and Gopal, Gopi Prashanth and Gray, Jan and Haselman, Michael and Hauck, Scott and Heil, Stephen and Hormati, Amir and Kim, Joo-Young and Lanka, Sitaram and Larus, James and Peterson, Eric and Pope, Simon and Smith, Aaron and Thong, Jason and Xiao, Phillip Yi and Burger, Doug}, title = {A Reconfigurable Fabric for Accelerating Large-scale Datacenter Services}, booktitle = {Proceeding of the 41st Annual International Symposium on Computer Architecture}, series = {ISCA '14}, year = {2014}, isbn = {978-1-4799-4394-4}, location = {Minneapolis, Minnesota, USA}, pages = {13-24}, numpages = {12}, url = {<http://dl.acm.org/citation.cfm?id=2665671.2665678>}, acmid = {2665678}, publisher = {IEEE Press}, address = {Piscataway, NJ, USA}, }

[download]

@miscellaneous{accelerating-deep-convolutional-neural-networks-using-specialized-hardware, author = {Kalin Ovtcharov, Olatunji Ruwase, Joo-Young Kim, Jeremy Fowers, Karin Strauss, Eric Chung}, title = {Accelerating Deep Convolutional Neural Networks Using Specialized Hardware}, booktitle = {}, year = {2015}, month = {February}, abstract = {

We describe the design of a convolutional neural network accelerator running on a Stratix V FPGA. The design runs at three times the throughput of previous FPGA CNN accelerator designs. We show that the throughput/watt is significantly higher than for a GPU, and project the performance when ported to an Arria 10 FPGA.

}, publisher = {Microsoft Research}, url = {<https://www.microsoft.com/en-us/research/publication/accelerating-deep-convolutional-neural-networks-using-specialized-hardware/>}, address = {}, pages = {}, journal = {}, volume = {}, chapter = {}, isbn = {}, }

@inproceedings{Deitz:2001:ERS:377792.377807, author = {Deitz, Steven J. and Chamberlain, Bradford L. and Snyder, Lawrence}, title = {Eliminating Redundancies in Sum-of-product Array Computations}, booktitle = {Proceedings of the 15th International Conference on Supercomputing}, series = {ICS '01}, year = {2001}, isbn = {1-58113-410-X}, location = {Sorrento, Italy}, pages = {65-77}, numpages = {13}, url = {<http://doi.acm.org/10.1145/377792.377807>}, doi = {10.1145/377792.377807}, acmid = {377807}, publisher = {ACM}, address = {New York, NY, USA}, }

@inproceedings{Basu:2015:CTH:2863692.2863932, author = {Basu, Protonu and Hall, Mary and Williams, Samuel and Straalen, Brian Van and Olikier, Leonid and Colella, Phillip}, title = {Compiler-Directed Transformation for Higher-Order Stencils}, booktitle = {Proceedings of the 2015 IEEE International Parallel and Distributed Processing Symposium}, series = {IPDPS '15}, year = {2015}, isbn = {978-1-4799-8649-1}, pages = {313-323}, numpages = {11}, url = {<http://dx.doi.org/10.1109/IPDPS.2015.103>}, doi = {10.1109/IPDPS.2015.103}, acmid = {2863932}, publisher = {IEEE Computer Society}, address = {Washington, DC, USA}, keywords = {Compiler Optimization, Stencil, High-Order, Multigrid, Mehrstellen}, }

@inproceedings{Putnam:2008:CHC:1344671.1344720, author = {Putnam, Andrew R. and Bennett, Dave and Dellinger, Eric and Mason, Jeff and Sundararajan, Prasanna}, title = {CHiMPS: A High-level Compilation Flow for Hybrid CPU-FPGA Architectures}, booktitle = {Proceedings of the 16th International ACM/SIGDA Symposium on Field Programmable Gate Arrays}, series = {FPGA '08}, year = {2008}, isbn = {978-1-59593-934-0}, location = {Monterey, California, USA}, pages = {261-261}, numpages = {1}, url = {<http://doi.acm.org/10.1145/1344671.1344720>}, doi =

```
{10.1145/1344671.1344720}, acmid = {1344720}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {FPGA, FPGA accelerators, c-to-gates, high-performance computing, reconfigurable computing}, } @inproceedings{Wong:2011:CFV:1950413.1950419, author = {Wong, Henry and Betz, Vaughn and Rose, Jonathan}, title = {Comparing FPGA vs. Custom Cmos and the Impact on Processor Microarchitecture}, booktitle = {Proceedings of the 19th ACM/SIGDA International Symposium on Field Programmable Gate Arrays}, series = {FPGA '11}, year = {2011}, isbn = {978-1-4503-0554-9}, location = {Monterey, CA, USA}, pages = {5-14}, numpages = {10}, url = {http://doi.acm.org/10.1145/1950413.1950419}, doi = {10.1145/1950413.1950419}, acmid = {1950419}, publisher = {ACM}, address = {New York, NY, USA}, keywords = {area, cmos, delay, fpga, soft processor}, }
```

```
@article{DBLP:journals/corr/GruslysMDLG16,
```

```
author      = {Audrunas Gruslys and  
              R{\'}mi Munos and  
              Ivo Danihelka and  
              Marc Lanctot and  
              Alex Graves},  
title       = {Memory-Efficient Backpropagation Through Time},  
journal     = {CoRR},  
volume     = {abs/1606.03401},  
year       = {2016},  
url        = {http://arxiv.org/abs/1606.03401},  
timestamp  = {Fri, 01 Jul 2016 17:39:49 +0200},  
biburl     = {http://dblp.uni-trier.de/rec/bib/journals/corr/GruslysMDLG16},  
bibsource  = {dblp computer science bibliography, http://dblp.org}
```

```
}
```

```
@inproceedings{FlowMap1994, author = {J. Cong and Ding, Yuzheng}, title = {FlowMap: an optimal technology mapping algorithm for delay optimization in lookup-table based FPGA designs}, booktitle = {IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems}, year = {1994}, isbn = {1937-4151}, pages = {1-12}, url = {http://ieeexplore.ieee.org/document/273754/}, doi = {10.1109/43.273754}, publisher = {IEEE} }
```

```
@article{MILROY20161589, title = "Towards Characterizing the Variability of Statistically Consistent Community Earth System Model Simulations", journal = "Procedia Computer Science", volume = "80", number = "", pages = "1589 - 1600", year = "2016", note = "", issn = "1877-0509", doi = "http://dx.doi.org/10.1016/j.procs.2016.05.489", url = "http://www.sciencedirect.com/science/article/pii/S1877050916309759", author = "Daniel J. Milroy and Allison H. Baker and Dorit M. Hammerling and John M. Dennis and Sheri A. Mickelson and Elizabeth R. Jessup", keywords = "Community Earth System Model", keywords = "CESM Ensemble Consistency Test", keywords = "statistical consistency", keywords = "code modification as source of variability", keywords = "compiler as source of variability", keywords = "Community Atmosphere Model", keywords = "non-bit-for-bit", keywords = "Fused Multiply-Add" }
```

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

Permanent link:
<https://www.cs.colostate.edu/AlphaZ/wiki/doku.php?id=melange:papers:spring2017&rev=1518045397>

Last update: **2018/02/07 16:16**

