

@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}, 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-Noel 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
(CG0) 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-netwo
rks-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 Oliker, 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/GruslyMDLG16,
```

```
author = {Audrunas Gruslys and R{\e}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/GruslyMDLG16}, 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} }

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=1487019229>Last update: **2017/02/13 13:53**