FPGA for HPC - exploring the possibilities of an alternative computer architecture

A Snowmass Letter of Interest for the Computational Frontier

Piotr Korcyl[1], Grzegorz Korcyl[2] and Salvatore Calì[1]

[1]Institute for Theoretical Physics, Jagiellonian University, Kraków, Poland
[2]Institute for Applied Computer Science, Jagiellonian University, Kraków, Poland

August 11, 2020

1 Introduction

High Performance Computing (HPC) is nowadays a crucial resource for advancing science. It is required at all levels of discovery: starting at the design of new experiments, through simulations of their performance, triggering system to select events to be analyzed, to the offline data analysis as well as numerical evaluation of predictions of theory side. In case of elementary particle theoretical physics focusing on the internal structure of proton and neutron (having in mind the physical setup of the Electron-Ion Collider) the approach is based on Markov Chain Monte Carlo and is called Lattice QCD.

2 Current scope - application to Lattice QCD

In this work we propose to evaluate Field Programmable Gate Arrays devices as computing resource alternative to the CPU and GPU. In particular we concentrate on Xilinx Alveo accelerators which together with the Xilinx Vitis development environment present a set of features that can be beneficial for the acceleration of the CG algorithm applied to Lattice QCD. Among them the most promising are:

- **Complete system generation**: Xilinx Vitis development environment allows to design the complete solution with high-level C++ sources and OpenCL bindings, taking full advantage of C++ classes, templates, overloaded operators and specializations.
- **Vast amount of logic resources**: FPGA devices offer massive and natural parallelism. Multiple instances of the computing kernel can be created working in parallel. Additionally using data flow concepts and streaming data from memory kernels can be pipelined offering new levels of paralellization, i.e. consecutive solver's iterations can be executed simulatneously. In this way one can get a much better ratio between FLOPS and Watts than on GPGPUs, which often are used way below their theoretical peak performance.
- **High Bandwidth Memory**: modern FPGA accelerators incorporate large partitions of High Bandwidth Memory which is able to deliver the data to multiple running kernels.
- **UltraRAM blocks**: temporary kernel's variables and objects can be stored in large chunks of on-chip UltraRAM which works with a latency equal to that of CPU registers.
- **Arbitrary data types**: using an architecture which is not based on predefined instruction sets allows to implement and efficiently use arbitrary data types. Depending on the algorithm this can be used to accelerate computations and save storage space. We have implemented a mixed precision Conjugate Gradient algorithm following [1] for the FPGA platform.

- **Dynamic reconfiguration**: the FPGA device can be reconfigured at any time in less than 70 ms. This allows to load a different version of the kernels (for instance with a different floating point number format) or a completely new kernels executing the next step in the workflow.
- **Kernel-to-kernel communication**: multiple instances of the kernels can be interconnected in a computing pipeline if only the amount of resources allows. Input datasets are processed by the first instance and the products are passed to the second instance. Xilinx Vivado HLS provides a set of tools that include dataflow directives and hls_stream constructs used to create fully streamlined designs.

In the HPCG_FPGA project [2], we used all these features in order to get the most out of Xilinx Alveo. The project, originally developed for Lattice QCD, can be successfully used as a reference design that includes examples of the above-mentioned mechanisms. Our results related to the Dirac-Wilson were published and reported on multiple conferences, including Lattice2018, Supercomputing Frontiers 2018, APLAT2020: [3], [4], [5], [2].

2.1 Software development

A complete project implementing the Conjugate Gradient for Lattice QCD was published as an open-source git repository [2]. Several other elements of the Lattice QCD workflow are now in sight [6]. They include smearing routines, especially the simplest iterative APE smearing procedure or the porting of modern linear solvers such as the multi-grid solver.

2.2 Hardware development

In our HPCG_FPGA project we managed to harvest all the advantages of FPGAs and demonstrated that Xilinx Alveo offers unique possibility for acceleration of computationally-demanding applications. So far to test our solutions we have been using either cloud resources or access to a dedicated Xilinx FPGA cluster at ETH in Zurich. Now we consider building a prototype of multi-FPGA, multi-node system in order to prove the scalability of our proposal.

2.3 Previous results

Some of the aspects discussed above were covered in similar attempts reported by the group from University of Frankfurt and the Maxeler company [7], and by the Trinity College Dublin group [8], [9], [10].

3 Outreach to other scientific domains

We can use our expertise and know-how to propose, design and evaluate FPGA-based compute resources for other scientific communities, such as numerical cosmology simulations, online and offline experimental data processing, just to name a few.

4 Scientific links to other institutions and industry partners

We profit from collaboration with industrial partners such as

- Xilinx, US - we are member of their University Program, Accelerator Program and Adaptive Compute Clusters Program. We receive substantive and technical support as well as software and hardware donations.
- EBV, Germany - hardware supplier

and academic institutions

- National Center for Nuclear Research, Swierk, Poland - R&D of the computational infrastructure for the analysis of data from the European X-ray Free Electron Laser
- FAIR, GSI, Darmstadt, Germany - R&D of the PANDA experiment data acquisition and online processing system based on FPGA devices.

5 Acknowledgements

This work was in part supported by Deutsche Forschungsgemeinschaft under Grant No.SFB/TRR 55 and by the polish NCN grant No. UMO-2016/21/B/ ST2/01492, by the Foundation for Polish Science grant no.TEAM/2017-4/39 and by the Polish Ministry for Science and Higher Education grant no. 7150/E-338/M/2018. The project could be realized thanks to the support from Xilinx University Program and their donations.

References

[1] R. Strzodka, D. Goddeke, Pipelined mixed precision algorithms on fpgas for fast and accurate pde solvers from low precision components, in: Proceedings of the 14th Annual IEEE Symposium on Field-Programmable Custom Computing Machines, FCCM '06, IEEE Computer Society, Washington, DC, USA, 2006, pp. 259–270. doi:10.1109/FCCM.2006.57.
URL http://dx.doi.org/10.1109/FCCM.2006.57

[2] G. Korcyl, P. Korcyl, Optimized implementation of the conjugate gradient algorithm for FPGA-based platforms using the Dirac-Wilson operator as an example (1 2020). arXiv:2001.05218.

[3] G. Korcyl, P. Korcyl, Towards Lattice Quantum Chromodynamics on FPGA devices, Comput. Phys. Commun. 249 (2020) 107029. arXiv:1810.04201, doi:10.1016/j.cpc.2019.107029.

[4] P. Korcyl, G. Korcyl, Implementation of the conjugate gradient algorithm in Lattice QCD on FPGA devices, PoS LATTICE2018 (2018) 313. arXiv:1811.03683, doi:10.22323/1.334.0313.

[5] G. Korcyl, P. Korcyl, Investigating the Dirac operator evaluation with FPGAs (4 2019). arXiv:1904.08616, doi:10.14529/jsfi190204.

[6] S. Calì, G. Korcyl, P. Korcyl, Using xilinx alveo accelerators for lattice qcd, in: Asia-Pacific Intertional Symposium on Lattice Field Theory 2020, 2020.

[7] T. Janson, U. Kebschull, Highly Parallel Lattice QCD Wilson Dirac Operator with FPGAs, Parallel Computing is Everywhere 32 664 – 672. doi:10.3233/978-1-61499-843-3-664.

[8] O. Callanan, A. Nisbet, E. Özer, J. Sexton, D. Gregg, FPGA implementation of a lattice quantum chromodynamics algorithm using logarithmic arithmetic, in: 19th International Parallel and Distributed Processing Symposium (IPDPS 2005), CD-ROM / Abstracts Proceedings, 4-8 April 2005, Denver, CO, USA, 2005. doi:10.1109/IPDPS.2005.228.
URL https://doi.org/10.1109/IPDPS.2005.228

[9] O. Callanan, D. Gregg, A. Nisbet, M. Peardon, High performance scientific computing using fpgas with ieee floating point and logarithmic arithmetic for lattice qcd, in: 2006 International Conference on Field Programmable Logic and Applications, 2006, pp. 1–6. doi:10.1109/FPL.2006.311191.

[10] O. Callanan, High performance scientific computing using fpgas for lattice qcd, Ph.D. thesis, Trinity College (Dublin, Ireland), School of Computer Science and Statistics, pp 161 (2007).