Conjugate gradient algorithm on FPGA devices

Grzegorz Korcyl[†] and Piotr Korcyl[*]

[†] Department of Information Technologies, Jagiellonian University, ul. Łojasiewicza 11, 30-348 Kraków, Poland

[*] Department of Discrete Field Theory, Jagiellonian University, ul. Łojasiewicza 11, 30-348 Kraków, Poland

Institut für Theoretische Physik, Universität Regensburg, 93040 Regensburg, Germany

grzegorz.korcyl@uj.edu.pl, piotr.korcyl@uj.edu.pl

UNIWERSYTET JAGIELLOŃSKI W KRAKOWIE

1. Introduction

Results of porting parts of the Lattice Quantum Chromodynamics code on modern FPGA devices are presented. A single-node, **double precision** implementation of the Conjugate Gradient algorithm is used to invert numerically the Dirac-Wilson operator on a 4-dimensional grid on a **Xilinx Zynq evaluation board**. The code is divided into **two software/hardware parts** in such a way that the entire multiplication by the Dirac operator is performed in **programmable logic**, and the rest of the algorithm runs on the ARM cores. **Optimized data blocks** are used to efficiently use data movement infrastructure enabling single stencil evaluation every **1 clock cycle**. We show that the FPGA implementation can offer superior performance compared to that obtained using Intel Xeon Phi KNL while consuming **6.4 W (logic: 3.1 W, ARM: 3.3 W)**.

2. FPGA and SDx environment

Modern Field Programmable Gate Arrays (FPGA) present a number of game-changing features that can make a strong impact on how computing resources for HPC are being designed. Most significant features are:

- Natural parallelization
- Streamlined processing
- Dynamic reconfiguration
- Direct access to network transceivers
- No OS overhead

Xilinx Zynq example:

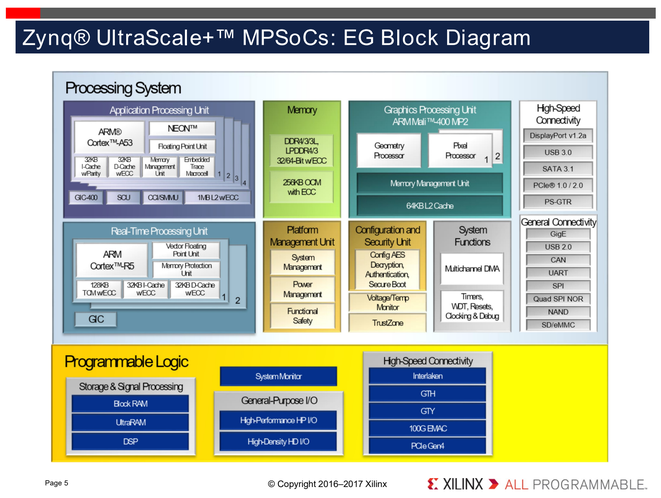

Software Defined X programming environment allows to compile C/C++ code for Multi-Processor devices: processing system (PS) and programmable logic (PL).

Xilinx SDx workflow example:

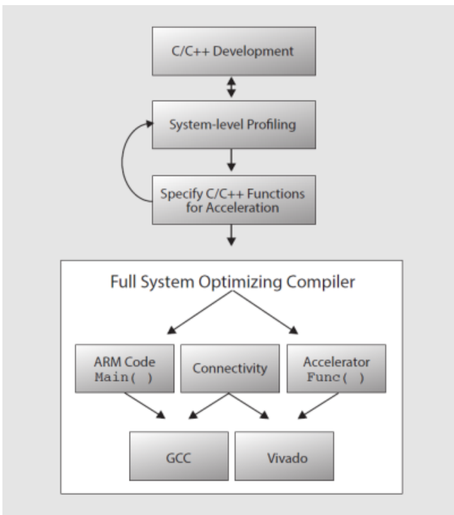

3. Zynq MPSoC internal structure

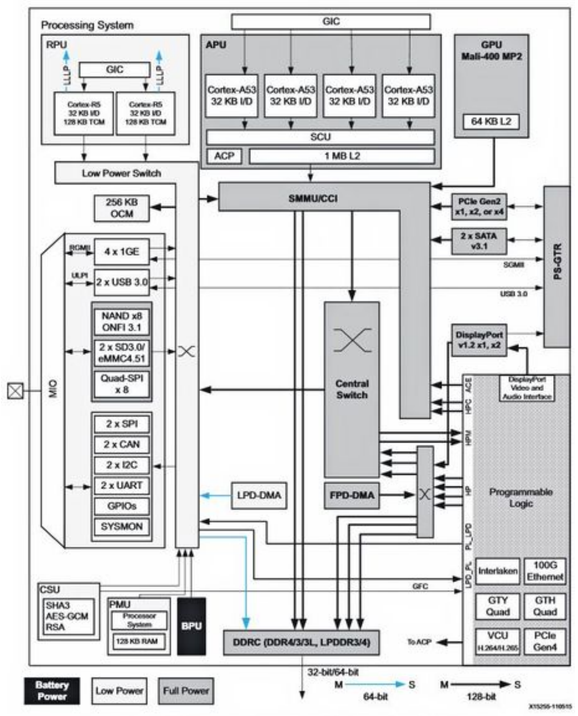

Figure 1-1: Zynq UltraScale+ MPSoC Top-Level Block Diagram

4. Conjugate gradient algorithm

Dirac equation to be solved reads

$$D^{AB}_{\alpha\beta}(n,m)\psi^{\beta}_{B}(m) = \eta^{\alpha}_{A}(n).$$

The simplest conjugate gradient algorithm written in pseudocode

$\psi \leftarrow \psi_0$
$r \leftarrow \eta - DD^\dagger\psi$
$p \leftarrow r$
while $|r| \geq r_{min}$ **do**
  $r_{old} \leftarrow |r|$
  $\alpha \leftarrow \frac{r_{old}}{\langle p|DD^\dagger|p\rangle}$
  $\psi \leftarrow \psi + \alpha p$
  $r \leftarrow r - \alpha DD^\dagger p$
  $\beta \leftarrow \frac{|r|}{r_{old}}$
  $p \leftarrow r + \beta p$
end while

5. Details

Dirac matrix multiplication: example compiled on Zynq XCZU9EG device

1. a block $6^3 \times 8$ of gauge field matrices and fermion vector is copied from DDR at the beginning of each function call
2. the block is separated into two independent $6^3 \times 4$ subblocks, including doubling the boundary data, increasing data locality
3. pipelined stencil multiplication for each site is performed thanks to optimized data availability
4. the resulting fermion vector is returned to DDR

6. Timings

Execution timeline (PL cycles @ 150 MHz captured on Zynq XCZU9EG device): data import, evaluation, result export. Achieved bandwidth DDR → PL using 4 channels: 1.2GBs / channel.

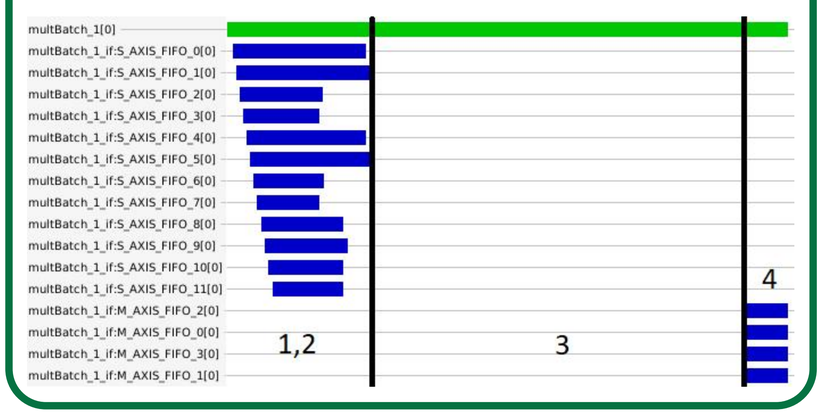

7. Resource consumption versus parallel execution

limit	latency	interval	bram	dsp	ff [10^6]	lut [10^6]	uram	(based on DSP)
4	149	1	508	6960	1.58	0.99	696	xcvu11p (≈ 32k$)
3	151	2	448	4320	0.97	0.64	696	xcku115 (≈ 8k$)
2	151	2	428	3480	0.83	0.57	696	xczu15 (≈ 4.5k$)
1	162	4	412	1740	0.47	0.35	696	xc7k480 (≈ 4.5k$)
	250	120	1388	546	0.13	0.09	–	xczu9eg (≈ 2.5k$)
			(76%)	(21%)	(24%)	(34%)		accelerator
			(95%)	(21%)	(31%)	(40%)		full project

8. Performance in GFLOPs

Estimates with * are based on the Zynq XCZU9EG board transfer infrastructure and assume 500MHz clock.

	only compute	with transfer
4	752	32*
3	376	30*
2	376	30*
1	188	28*
	2.1	1.3

9. Further steps

Extension to multiple nodes:

- openAMP allows to address memory in a heterogeneous system (many FPGA devices)
- ARM hosts not needed on compute nodes
- benefit from multilane, multigigabit transceivers (network layers implemented in logic)

Work supported by Deutsche Forschungsgemeinschaft under Grant No. SFB/TRR 55, 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