

Accelerating an iterative Helmholtz solver with FPGAs

Art Petrenko¹, Tristan van Leeuwen², Diego Oriato³, Simon Tilbury³ and Felix J. Herrmann¹

¹Department of Earth, Ocean and Atmospheric Sciences, University of British Columbia

²Centrum Wiskunde & Informatica, formerly with UBC

³Maxeler Technologies

January 16, 2014

Abstract

We implement the Kaczmarz row-projection algorithm (Kaczmarz (1937)) on a CPU host + FPGA accelerator platform using techniques of dataflow programming. This algorithm is then used as the preconditioning step in CGMN, a modified version of the conjugate gradients method (Björck and Elfving (1979)) that we use to solve the time-harmonic acoustic isotropic constant density wave equation. Using one accelerator we achieve a speed-up of over $2\times$ compared with one Intel core.

Introduction

Using reconfigurable hardware (field programmable gate arrays, FPGAs) to handle the computationally intensive kernel of an algorithm is attractive. Unless the algorithm is limited by the bandwidth of an interconnect (such as a link to memory), speed of execution is directly proportional to the amount of data to be processed. As long as the kernel fits onto the FPGA its complexity is irrelevant, since all the operations on an FPGA happen in the same clock tick. In contrast, on a CPU core, more complex algorithms require more clock ticks as instructions are executed sequentially, or with only a small amount of parallelism. We demonstrate hardware acceleration with FPGAs applied to seismic wave propagation simulation, a large computational burden during full-waveform inversion (FWI).

Theory

In the frequency domain, and specializing to the acoustic isotropic constant density case, simulating wave propagation corresponds to solving a large linear system of the form

$$H(\mathbf{m}, \omega)\mathbf{u} = \mathbf{q}, \quad (1)$$

where H is the $N \times N$ Helmholtz operator that depends on the earth model \mathbf{m} and the angular frequency ω , \mathbf{u} is the (complex) Fourier transform of the pressure wavefield and \mathbf{q} is the time-harmonic source. We take \mathbf{m} to be the squared slowness ($1/v^2$), and use a 27-point cube stencil to calculate the entries of H for the 3D case, as in Operto et al. (2007). A perfectly matched layer is used at the boundaries of the domain to eliminate reflection artifacts.

Because the matrix H is large (up to $10^9 \times 10^9$) and sparse (only a few non-zero diagonals), an iterative algorithm like the method of conjugate gradients (CG) is well suited to solve Equation 1. However CG cannot be used directly because H is not symmetric positive semidefinite (i.e. it has both positive and negative eigenvalues). Instead we use the Kaczmarz algorithm as a preconditioner to transform Equation 1 into an equivalent system that does have these properties. The Kaczmarz algorithm (Kaczmarz (1937)) is an iterative method for solving generic linear systems $A\mathbf{x} = \mathbf{b}$ that projects its iterate \mathbf{x}_k onto the hyperplane defined by a row \mathbf{a}_i of the matrix A :

$$\mathbf{x}_{k+1} = \mathbf{x}_k + w(b_i - \langle \mathbf{a}_i, \mathbf{x}_k \rangle) \frac{\mathbf{a}_i^*}{\|\mathbf{a}_i\|^2}. \quad (2)$$

Here w is a relaxation factor we set equal to 1.5 and b_i is the i^{th} element of the right-hand side. We refer to one application of Equation 2 as a Kaczmarz iteration, and to the projections onto each row from first to last as a forward Kaczmarz sweep (FKSWP). Backward sweeps correspond to projecting onto the rows in reverse order, and a double sweep (DKSWP) is a forward sweep followed by a backward sweep ($k: 1 \rightarrow 2N; i: 1 \rightarrow N, N \rightarrow 1$). Equation 1 is transformed into the equivalent system

$$(I - Q)\mathbf{u} = R\mathbf{q}, \quad (3)$$

where the action of matrices Q and R is computed using a double Kaczmarz sweep: $\text{DKSWP}(H, \mathbf{u}, \mathbf{q}, w) = Q\mathbf{u} + R\mathbf{q}$. Equation 3 is then solved with CG, an algorithm known as CGMN and due to Björck and Elfving (1979).

Our contribution is implementing the double Kaczmarz sweep for a CPU host + FPGA accelerator platform. Related work has been done by Grull et al. (2012), who implement a related algorithm, the simultaneous algebraic reconstruction technique (SART), on the same platform. Furthermore, Elble et al. (2010) have examined the improvements that Kaczmarz and other algorithms experience on Graphical Processing Units (GPUs).

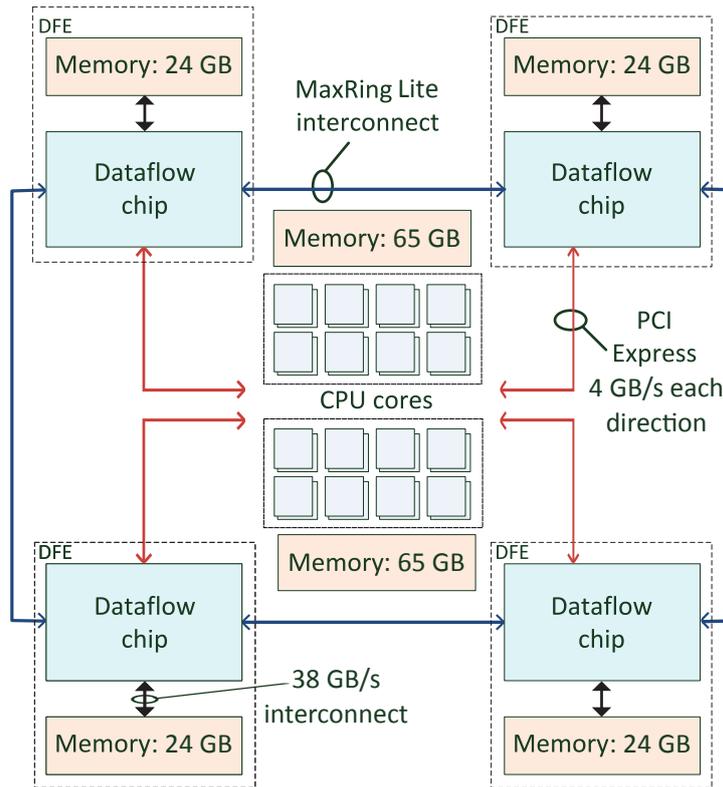


Figure 1 System of CPU host + FPGA accelerators; adapted from Pell et al. (2013).

Method

The target platform for our implementation is described in Pell et al. (2013) and shown in Figure 1. Four dataflow engines (DFEs), consisting of FPGAs and their associated large memory, are connected via a PCIe bus to the CPU. Each FPGA also has 4.6 MB of fast on-chip RAM (FMem). Coarse-grained parallelization by solving problems with different right-hand sides \mathbf{q} on each DFE is a straightforward extension but has not yet been attempted by the authors; currently only one accelerator is used. The DFE is operated via a compiled binary that is called from a high-level numerical computing environment using an automatically generated C intermediate layer.

Execution of the CGMN algorithm using the DFEs proceeds in three stages. First the matrix $H(\mathbf{m}, \omega)$ is copied to the FPGA's memory. It will be read twice during each CGMN iteration (for each forward and backward sweep), but will not be modified. The second stage consists of the forward Kaczmarz sweep, during which the current CGMN search direction \mathbf{p} is streamed to the FPGA from the CPU host, and the result $\text{FKSWP}(H, \mathbf{p}, 0, w)$ is stored to memory. The third stage is the backward Kaczmarz sweep, which reads the results of the forward sweep from memory and streams $\text{DKSWP}(H, \mathbf{p}, 0, w)$ back to the CPU. The CPU carries out the elementwise and reduction operations necessary, and stages 2 and 3 are repeated until CGMN converges.

If the rows of H are processed from 1 to N , each Kaczmarz iteration depends on the results of the last. Furthermore, each iteration's computation takes many ticks of the FPGA clock. This is the latency L , and is chiefly due to the pairwise summation of values in $\langle a_i, \mathbf{x}_k \rangle$ (Equation 2). To avoid having to wait for L ticks between iterations, L independent iterations used to fill the computational pipeline. Two sources of parallelism can be exploited. First, multiple Helmholtz problems sharing the same matrix H but different right-hand sides \mathbf{q} can be solved on the same DFE. Second, the order in which DKSWP accesses the rows of H can be changed such that consecutive Kaczmarz iterations update disjoint sets of iterate elements. At the present we use the latter method exclusively.

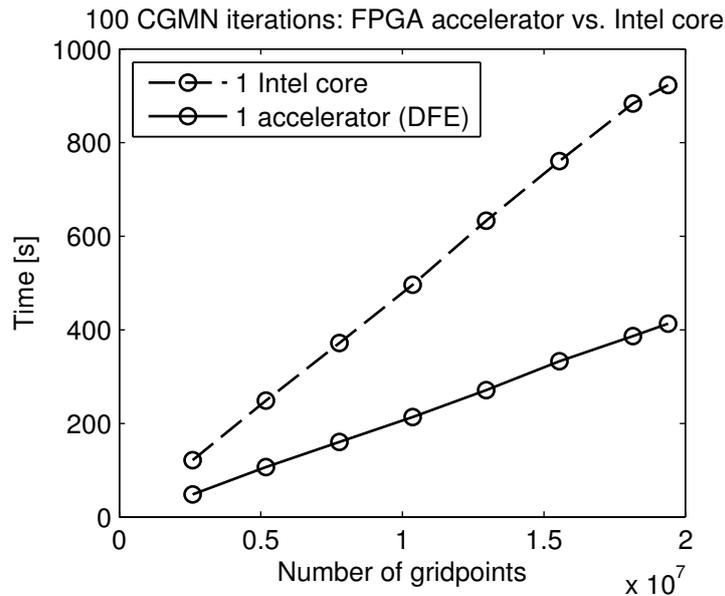


Figure 2 Comparison of CGMN execution time on one accelerator (DFE) and one Intel Xeon core.

Results

In Figure 2, computational times for 100 CGMN iterations running on the CPU + accelerator platform (Vectis model at 100 MHz) are compared with times on one Intel Xeon E5-2670 core. A speed-up of over $2\times$ is seen when CGMN is run on the accelerator. Figure 3 shows the velocity model used in the experiment, as well as the solution wavefield \mathbf{u} that results from a time-harmonic source at one point in the domain.

At the current stage of development there are several factors that limit the performance of the accelerator. Chief among these is the inefficient use of the FPGA's memory controller. To estimate the speed-up we can expect from a more optimized implementation we slow down the sweeps by having computation progress only once every 90 FPGA clock ticks. At this speed even inefficient use of the memory controller is sufficient to keep the FPGA supplied with data. We measure performance for a range of small systems (from 0.5 to 3×10^5 gridpoints) and then calculate theoretical performance of the full-speed sweeps, taking into account several planned improvements. These include: two passes over the data per CGMN iteration instead of three, getting rid of zero elements currently used to pad the rows of the matrix H to boundaries of addressable memory chunks, and increasing the FPGA operating frequency to 120 MHz. With a 100% efficient memory controller these improvements will translate to a speed-up of $13\times$ over the single Intel core times.

Speed-up of the pipelined accelerated Kaczmarz sweeps will further be limited by the fact that the sweeps (when implemented on an Intel core) take up only about 95% of the total time used by CGMN. Recall that offloading parts of an algorithm from the CPU onto the FPGA accelerator decreases the CPU execution time, but does not increase the execution time on the FPGA, since the amount of data to process (problem size) remains the same. Hence we will implement all of CGMN on the accelerator, from which we expect a speed-up of $36\times$ over a single Intel core. At that point the link between the FPGA and its dedicated memory that is used to stream the matrix H will become the bottleneck preventing further increases in FPGA operating frequency. To avoid this bandwidth limitation, the elements of the Helmholtz matrix will be generated on the FPGA from the earth model \mathbf{m} as they are needed, rather than being read in from memory.

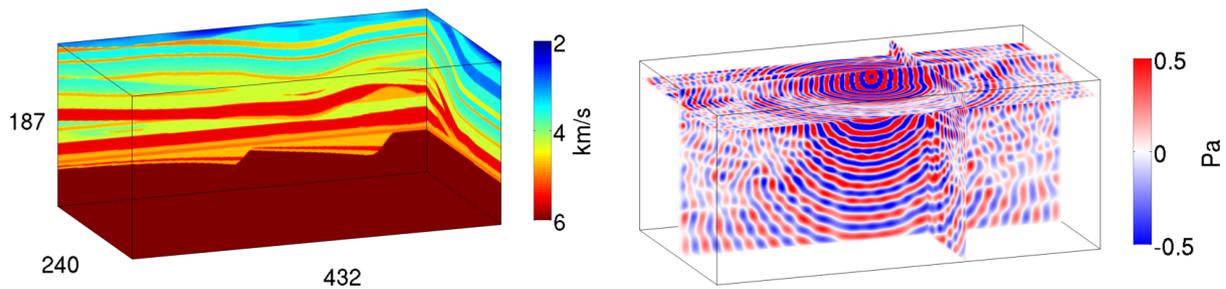


Figure 3 On the left is the largest part of the SEG/EAGE Overthrust acoustic velocity model (Aminzadeh et al. (1997)) used in the experiment. The axes are labelled with the number of grid points along that axis. The right plot shows the real part of the Fourier transform of the pressure wavefield u that results from a time-harmonic point source. The solution shown took 2703 CGMN iterations (9826 s) to converge to a relative residual norm of 10^{-6} .

Conclusions

We have demonstrated a time-harmonic iterative wave equation solver that off-loads its computational kernel onto a reconfigurable hardware accelerator. Work on this project is ongoing, and while further development is needed in order to better realize the potential for acceleration inherent in the platform, our preliminary results give us reason to be optimistic.

Acknowledgements

We would like to thank Henryk Modzelewski for his support of the hardware and software used by the three authors affiliated with UBC, Eddie Hung for stimulating discussions about implementation, Rafael Lago for guidance about Krylov solvers, Brendan Smithyman for editing this manuscript, and Maxeler Technologies for allowing one of the authors to visit for a two-week period. This work was in part financially supported by the Natural Sciences and Engineering Research Council of Canada Discovery Grant (22R81254) and the Collaborative Research and Development Grant DNOISE II (375142-08). This research was carried out as part of the SINBAD II project with support from the following organizations: BG Group, BGP, BP, CGG, Chevron, ConocoPhillips, ION, Petrobras, PGS, Total SA, WesternGeco, and Woodside.

References

- Aminzadeh, F., Jean, B. and Kunz, T. [1997] *3-D salt and overthrust models*. Society of Exploration Geophysicists.
- Björck, Å. and Elfving, T. [1979] Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations. *BIT Numerical Mathematics*, **19**(2), 145–163, ISSN 0006-3835, doi:10.1007/BF01930845.
- Elble, J.M., Sahinidis, N.V. and Vouzis, P. [2010] GPU computing with Kaczmarz's and other iterative algorithms for linear systems. *Parallel Computing*, **36**(5-6), 215–231, ISSN 0167-8191, doi:10.1016/j.parco.2009.12.003.
- Grüll, F., Kunz, M., Hausmann, M. and Kebschull, U. [2012] An implementation of 3d electron tomography on fpgas. *Reconfigurable Computing and FPGAs (ReConFig), 2012 International Conference on*, 1–5, doi: 10.1109/ReConFig.2012.6416732.
- Kaczmarz, S. [1937] Angenäherte auflösung von systemen linearer gleichungen. *Bulletin International de l'Academie Polonaise des Sciences et des Lettres*, **35**, 355–357.
- Operto, S., Virieux, J., Amestoy, P., L'Excellent, J.Y., Giraud, L. and Ali, H.B.H. [2007] 3D finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study. *Geophysics*, **72**(5), SM195–SM211, doi:10.1190/1.2759835.
- Pell, O., Bower, J., Dimond, R., Mencer, O. and Flynn, M.J. [2013] Finite-difference wave propagation modeling on special-purpose dataflow machines. *Parallel and Distributed Systems, IEEE Transactions on*, **24**(5), 906–915, ISSN 1045-9219, doi:10.1109/TPDS.2012.198.