

Accelerating an Iterative Helmholtz Solver with FPGAs

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



a place of mind
THE UNIVERSITY OF BRITISH COLUMBIA

Introduction

Seismic exploration aims to characterize the Earth's subsurface to assist the discovery and production of hydrocarbon resources. One algorithm used in this context is full-waveform inversion (FWI) (Herrmann et al. (2013)). FWI compares real wavefield data measured in a seismic survey with synthetic data generated by numerical wave simulations. Based on the difference, it updates its current guess at the subsurface parameters and repeats the process until convergence. The high computational burden of simulating seismic waves limits the number of FWI iterations that can be performed and hence the accuracy of the subsurface parameter estimate. Our contribution is acceleration of a wave equation solver via its implementation on a computing platform consisting of a CPU host and a reconfigurable hardware accelerator.

Solving the Wave Equation: Mathematical Background

$$A \quad u = q$$

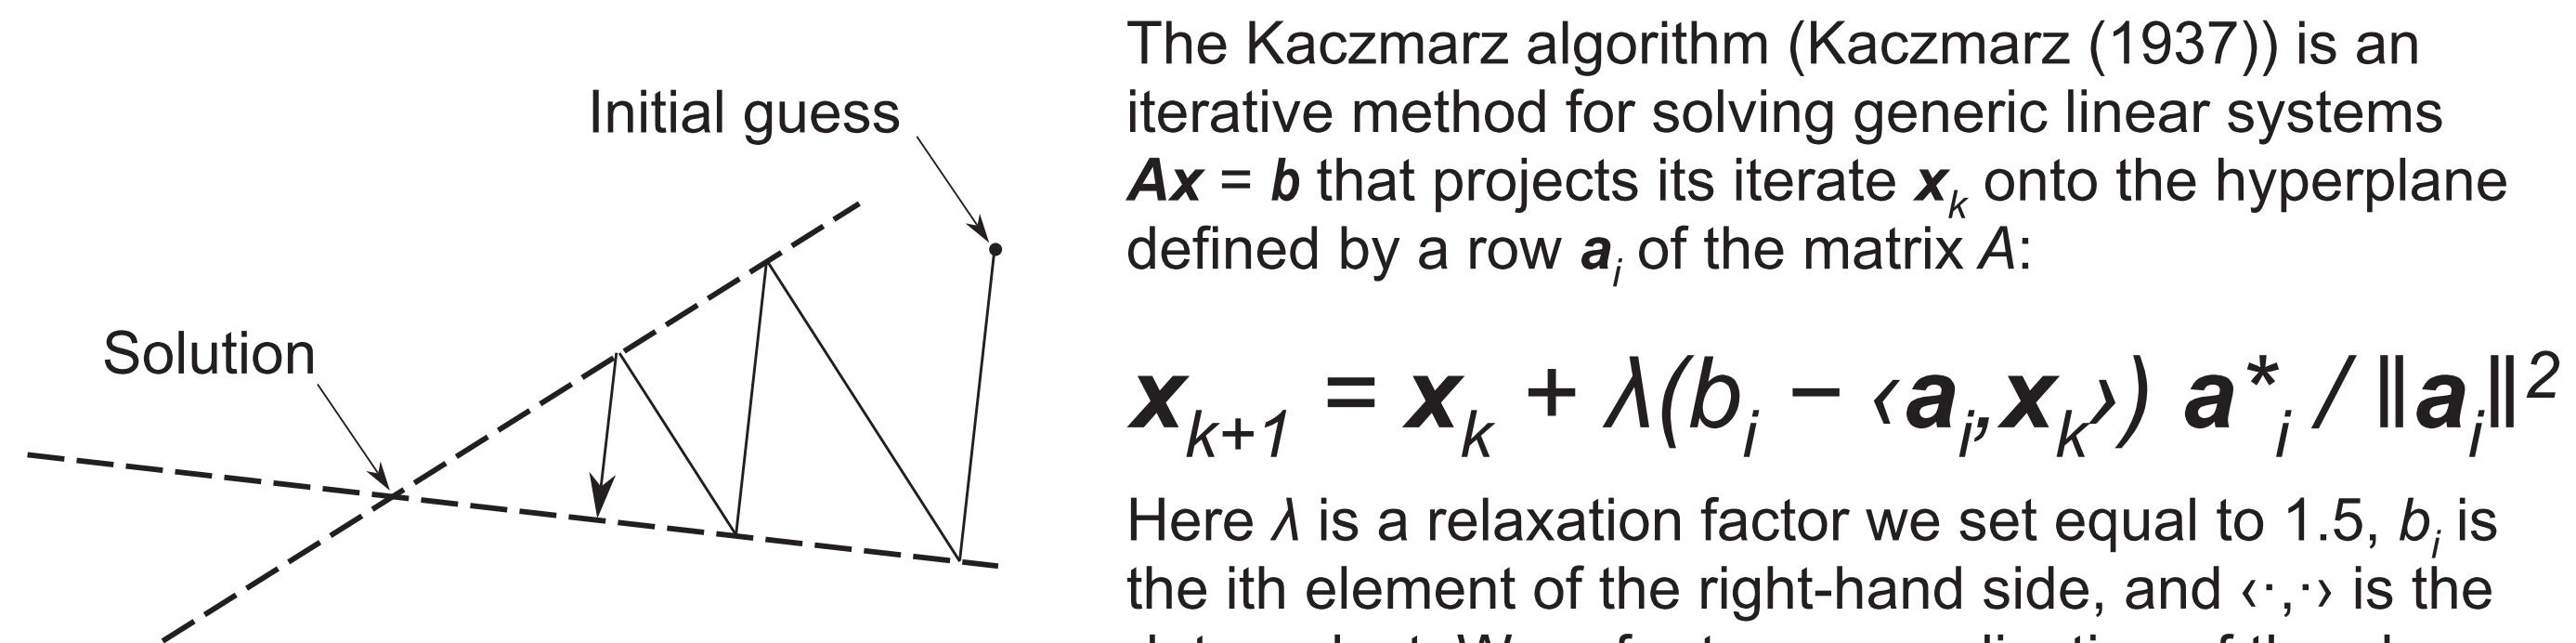
In the frequency domain, and specializing to the acoustic isotropic constant density case, simulating wave propagation corresponds to solving a large linear system, schematically illustrated at left. (Only 9 of 27 non-zero diagonals of A are shown).

$A = (\omega m + \Delta)$ is the $N \times N$ Helmholtz operator that depends on the earth model m and the angular frequency ω .

u is the (complex) Fourier transform of the pressure wavefield, q is the time-harmonic source and Δ is the Laplacian.

We take m to be the squared slowness ($1/v^2$), and use a 27-point cube stencil to calculate the entries of A for the 3D case, as in Operto et al. (2007). A perfectly matched layer (PML) is used at the boundaries of the domain to eliminate reflection artifacts.

Because the matrix A 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 the wave equation above. However CG cannot be used directly because A is not symmetric positive semidefinite (SPSD, i.e. it has both positive and negative eigenvalues). Instead we use the Kaczmarz algorithm as a preconditioner to transform the wave equation into an equivalent system that does have these properties.



The Kaczmarz algorithm (Kaczmarz (1937)) is an iterative method for solving generic linear systems $Ax = b$ that projects its iterate x_k onto the hyperplane defined by a row a_i of the matrix A :

$$x_{k+1} = x_k + \lambda(b_i - \langle a_i, x_k \rangle) a_i^* / \|a_i\|^2$$

Here λ is a relaxation factor we set equal to 1.5, b_i is the i th element of the right-hand side, and $\langle \cdot, \cdot \rangle$ is the dot product. We refer to one application of the above equation as a Kaczmarz iteration, and to the

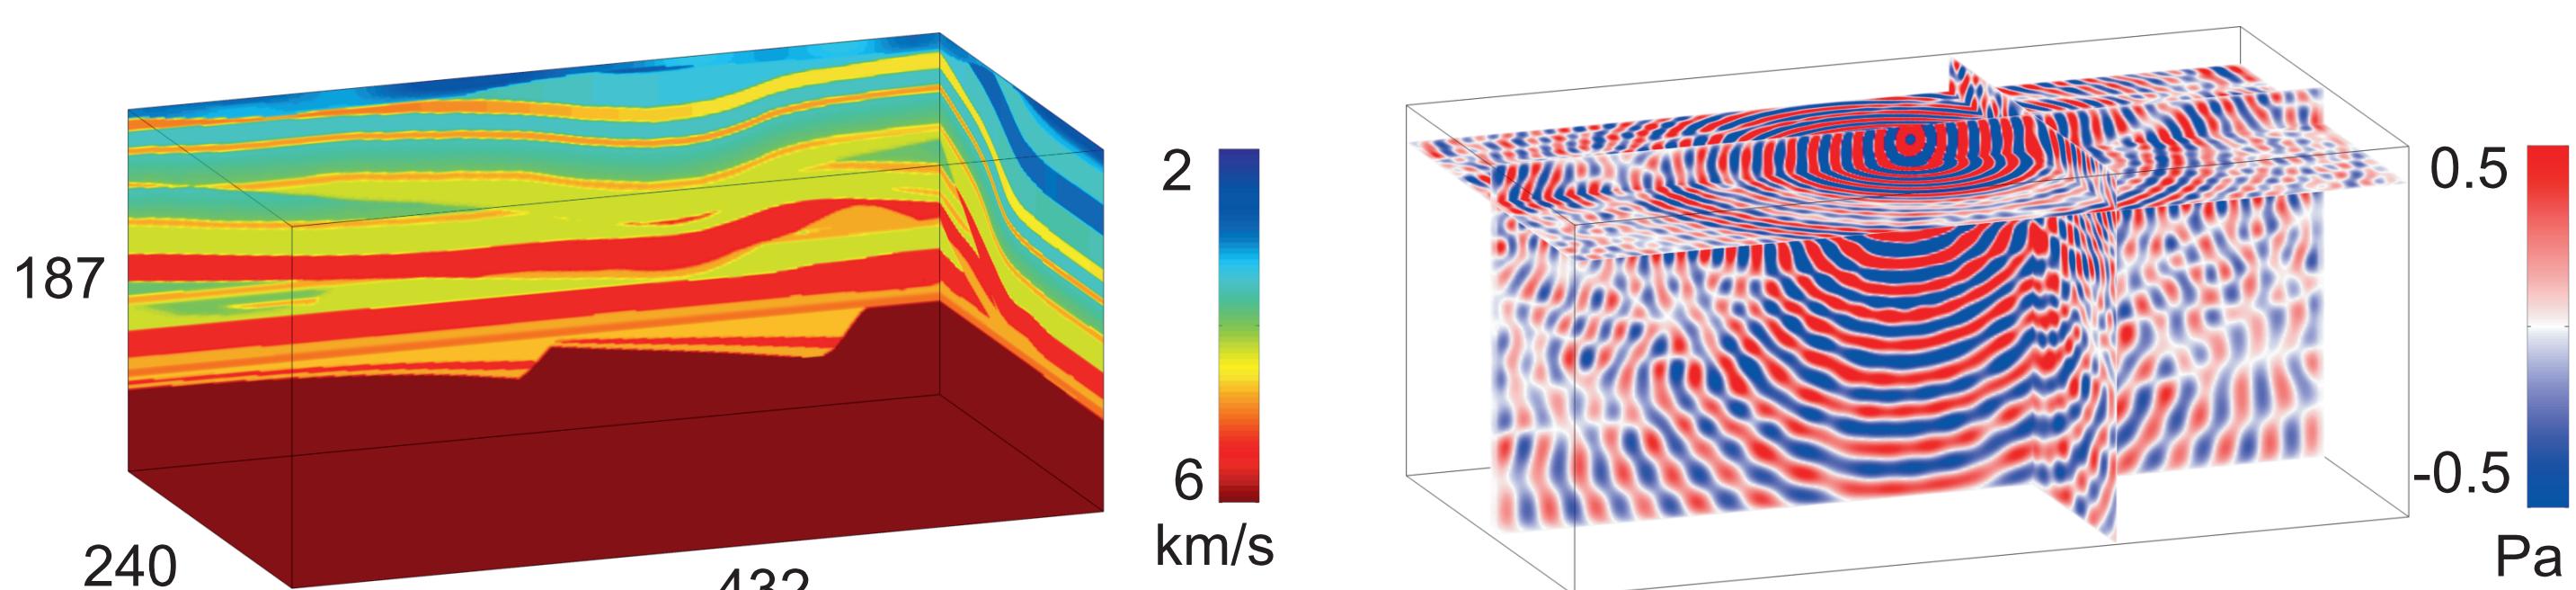
projections onto each row from first to last as a forward Kaczmarz sweep (FKSWP). Backward sweeps correspond to projecting onto the rows in reverse order, and a double sweep (DKSWP) is a forward sweep followed by a backward sweep ($k: 1 \rightarrow 2N; i: 1 \rightarrow N, N \rightarrow 1$). The double Kaczmarz sweep can be represented in matrix form:

$$\text{DKSWP}(A, u, q, \lambda) = Qu + Rq$$

The wave equation is transformed into an equivalent system for the fixed point of DKSWP:

$$(I - Q)u = Rq$$

This equivalent system is SPSD and is solved with CG, an algorithm known as CGMN and due to Björck and Elfving (1979). A parallel version of CGMN (dubbed CARP-CG) has recently been used by van Leeuwen et al. (2012).



Subsurface Model

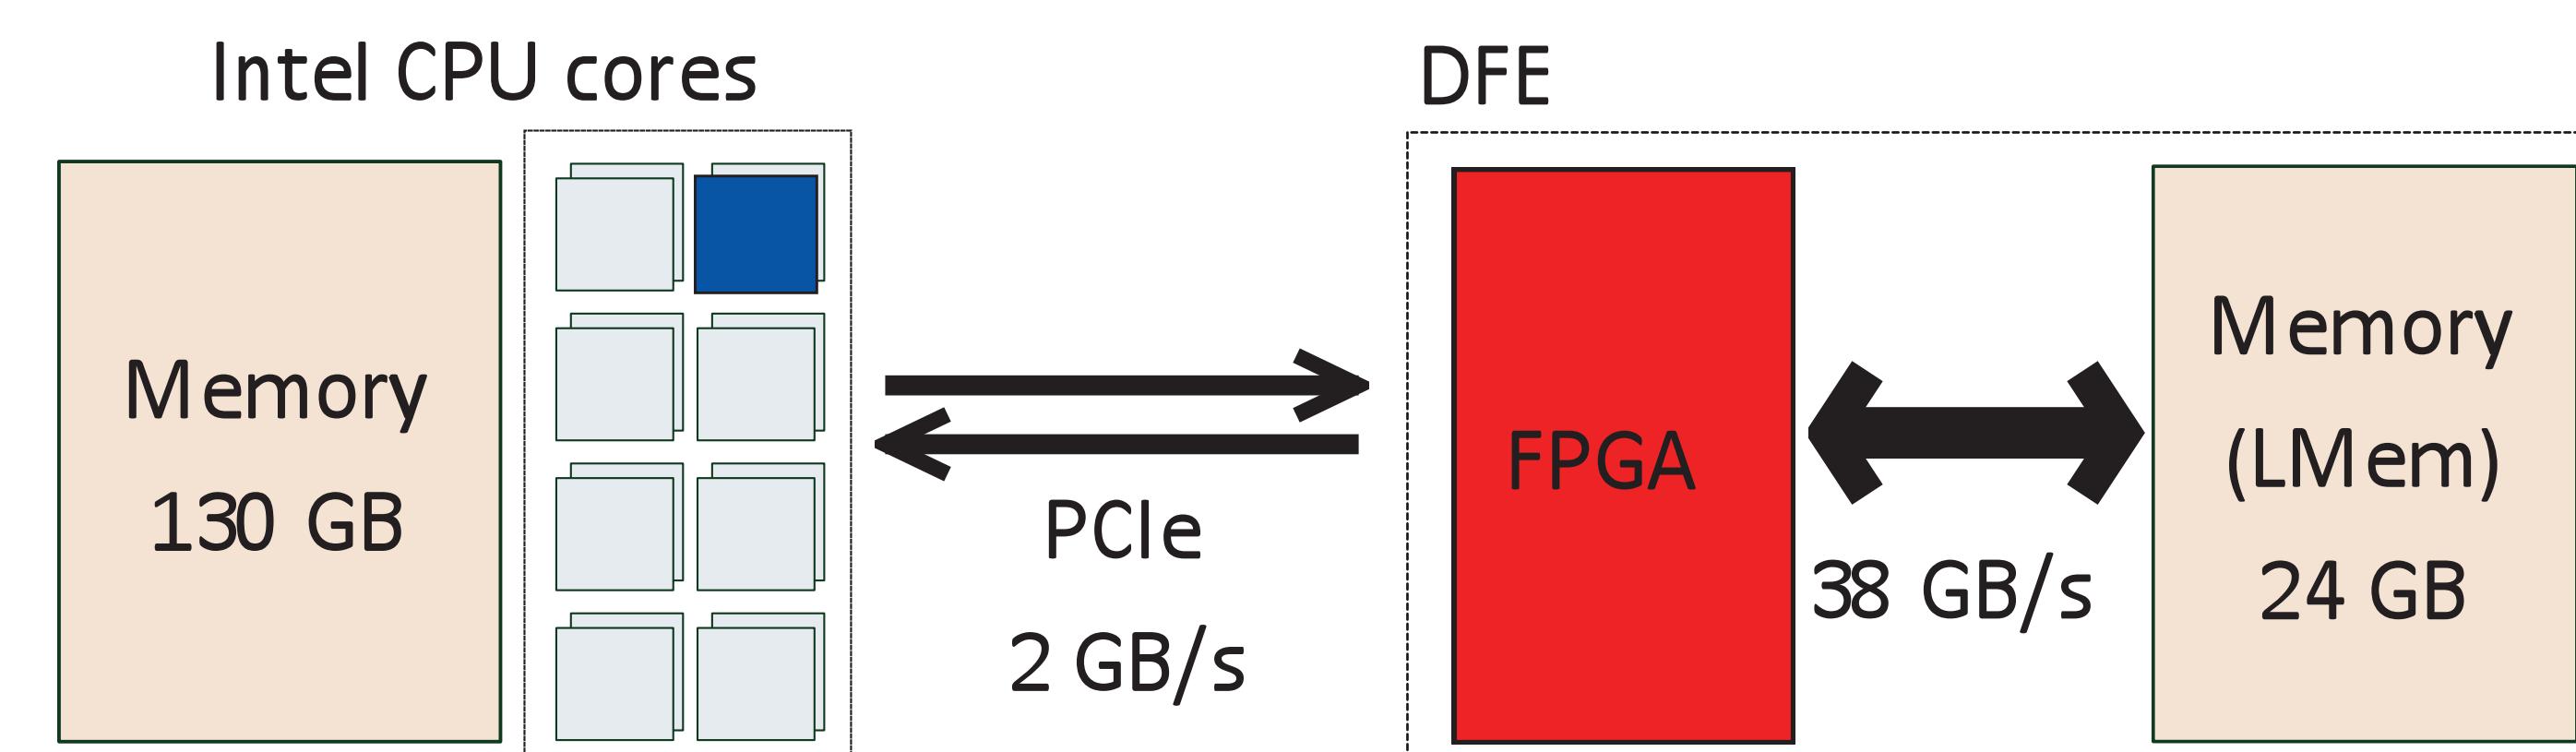
The largest part of the SEG/EAGE overthrust velocity model (Aminzadeh (1997)) used in the experiments. The axes are labelled with the corresponding number of grid points.

Example Wavefield

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 using one accelerator) to converge to a relative residual norm of 10^{-6} .

What are FPGAs?

Field-programmable gate arrays (FPGAs) are computer chips that consist of electronic building blocks that can be configured and re-configured by the user to represent algorithms in hardware. A conventional CPU can be thought of as a pipeline for processing instructions: the more complex the algorithm, the more instructions there are, the more clock ticks are required to execute it. In contrast, an FPGA can be thought of as a pipeline for processing data: 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 algorithm fits onto the FPGA, its complexity is irrelevant, since all the operations on an FPGA are hard-wired and happen in the same clock tick.



The target platform for our implementation is described in Pell et al. (2013) and shown above. A dataflow engine (DFE), consisting of an FPGA and its associated large memory, is connected via a PCIe bus to the CPU. The FPGA also has 4.6 MB of fast on-chip RAM (FMem). Although the machine the first author has access to has four DFES, currently only one accelerator is used. The DFE is operated via a compiled binary that is called from a high-level numerical computing environment (MATLAB) using an automatically generated C intermediate layer. Related work has been done by Grüll et al. (2012), who implement an algorithm related to the Kaczmarz sweep (the simultaneous algebraic reconstruction technique, SART) on the same platform.

Dataflow During a Kaczmarz Row Projection

Execution of the CGMN algorithm using the DFES proceeds in three stages. First the matrix $A(m, \omega)$ is copied to LMem. 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 p is streamed to the FPGA from the CPU host, and the result FKSWP($A, p, 0, \lambda$) is stored to LMem. The third stage is the backward Kaczmarz sweep, which reads the results of the forward sweep from memory and streams DKSWP($A, p, 0, \lambda$) back to the CPU. The CPU carries out the element-wise and reduction operations necessary, and stages 2 and 3 are repeated until CGMN converges.

Overcoming Dependency of Sequential Projections

If the rows of A are processed from 1 to N , each Kaczmarz row projection depends on the results of the last. Furthermore, each projection'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 $\langle a_i, x_k \rangle$ in the Kaczmarz row projection equation. To avoid having to wait for L ticks between iterations, L independent iterations are used to fill the computational pipeline. Two sources of parallelism can be exploited. First, multiple Helmholtz problems sharing the same matrix A but different right-hand sides q can be solved on the same DFE. Second, the order in which DKSWP accesses the rows of A can be changed such that consecutive Kaczmarz iterations update disjoint sets of iterate elements. Access to the rows (and to the elements of the Kaczmarz iterate and right-hand side) is strided by 3 grid-points in the fast dimension, as illustrated at right. When the end of the fast dimension is reached the index wraps around to rows skipped the first time around. At the present we use the latter pipelining method exclusively.

