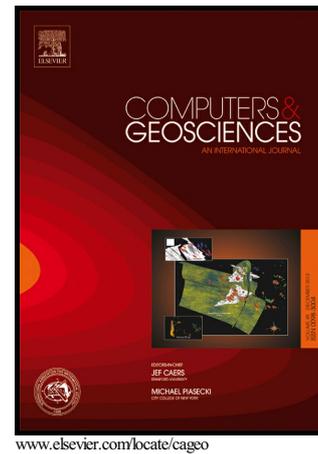


Author's Accepted Manuscript

Time-domain seismic modeling in viscoelastic media for full waveform inversion on heterogeneous computing platforms with OpenCL

Gabriel Fabien-Ouellet, Erwan Gloaguen, Bernard Giroux



PII: S0098-3004(16)30781-6

DOI: <http://dx.doi.org/10.1016/j.cageo.2016.12.004>

Reference: CAGEO3882

To appear in: *Computers and Geosciences*

Cite this article as: Gabriel Fabien-Ouellet, Erwan Gloaguen and Bernard Giroux Time-domain seismic modeling in viscoelastic media for full waveform inversion on heterogeneous computing platforms with OpenCL, *Computers and Geosciences*, <http://dx.doi.org/10.1016/j.cageo.2016.12.004>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting galley proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1 Time-domain seismic modeling in viscoelastic media for full waveform inversion on 2 heterogeneous computing platforms with OpenCL

3 Gabriel Fabien-Ouellet*, Erwan Gloaguen, Bernard Giroux

4 INRS-ETE, 490, Rue de la Couronne, Québec, QC, Canada G1K 9A9

5 *Email address: Gabriel.Fabien-Ouellet@ete.inrs.ca

6 **Abstract**

7 Full Waveform Inversion (FWI) aims at recovering the elastic parameters of the Earth by matching recordings of the ground motion
8 with the direct solution of the wave equation. Modeling the wave propagation for realistic scenarios is computationally intensive,
9 which limits the applicability of FWI. The current hardware evolution brings increasing parallel computing power that can speed up
10 the computations in FWI. However, to take advantage of the diversity of parallel architectures presently available, new programming
11 approaches are required. In this work, we explore the use of OpenCL to develop a portable code that can take advantage of the
12 many parallel processor architectures now available. We present a program called SeisCL for 2D and 3D viscoelastic FWI in the time
13 domain. The code computes the forward and adjoint wavefields using finite-difference and outputs the gradient of the misfit function
14 given by the adjoint state method. To demonstrate the code portability on different architectures, the performance of SeisCL is tested
15 on three different devices: Intel CPUs, NVidia GPUs and Intel Xeon PHI. Results show that the use of GPUs with OpenCL can speed
16 up the computations by nearly two orders of magnitudes over a single threaded application on the CPU. Although OpenCL allows
17 code portability, we show that some device-specific optimization is still required to get the best performance out of a specific
18 architecture. Using OpenCL in conjunction with MPI allows the domain decomposition of large models on several devices located on
19 different nodes of a cluster. For large enough models, the speedup of the domain decomposition varies quasi-linearly with the
20 number of devices. Finally, we investigate two different approaches to compute the gradient by the adjoint state method and show
21 the significant advantages of using OpenCL for FWI.

22 **Keywords**

23 OpenCL; GPU; Seismic; Viscoelasticity; Full waveform Inversion; Adjoint state method;

24 1. Introduction

25 In recent years, parallel computing has become ubiquitous due to a conjunction of both hardware and software availability.
26 Manifestations are seen at all scales, from high-performance computing with the use of large clusters, to mobile devices such as
27 smartphones that are built with multicore Central Processing Units (CPU) (Abdullah and Al-Hafidh, 2013). Graphics processing units
28 (GPU) bring this trend to the next level, packing now up to several thousand cores in a single device. Scientific simulations have
29 benefited from this technology through general-purpose processing on graphics processing units and, for certain applications, GPUs
30 can speed up calculation over one or two orders of magnitude over its CPU counterpart. This has caused a surge in the use of GPUs
31 in the scientific community (Nickolls and Dally, 2010, Owens, et al., 2008), with applications ranging from computational biology to
32 large-scale astrophysics. Furthermore, GPUs are increasingly used in large clusters (Zhe, et al., 2004), and now several of the
33 fastest supercomputers on earth integrate GPUs or accelerators (Dongarra, et al., 2015).

34
35 Nevertheless, GPUs are not fit for all kinds of scientific computations (Vuduc, et al., 2010). Potential gains from adopting GPUs
36 should be studied carefully before implementation. In particular, the algorithm should follow the logic of the single-program multiple-
37 data (SPMD) programming scheme, i.e. many elements are processed in parallel with the same instructions. In geophysics, and
38 more precisely in the seismic community, GPU computing has been applied most successfully in modeling wave propagation with
39 Finite-Difference Time-Domain (FDTD) schemes. Indeed, the finite-difference method is well suited to GPUs because the solution is
40 obtained by independent computations on a regular grid of elements and follows closely the SPMD model (Micikevicius, 2009). For
41 example, Michéa and Komatitsch (2010) show an acceleration by a factor between 20 to 60 between the single-core implementation
42 of the FDTD elastic wave propagation and a single GPU implementation. Okamoto (2011) shows a 45 times speed-up with a single
43 GPU implementation and presents a multi-GPU implementation that successfully parallelizes the calculation, although with a sub-
44 linear scaling. Both Rubio, et al. (2014) and Weiss and Shragge (2013) present multi-GPU FDTD programs for anisotropic elastic
45 wave propagation that shows the same unfavorable scaling behavior. Sharing computation through domain decomposition can be
46 problematic mainly because the memory transfers between GPUs and between nodes are usually too slow compared to the
47 computation on GPUs. GPU computing has also been applied successfully to the spectral element method (Komatitsch, et al., 2010),
48 the discontinuous Galerkin method (Mu, et al., 2013) and reverse time migration (Abdelkhalek, et al., 2009), among others.

49

50 Nearly all of the seismic modeling codes written for GPUs have been implemented with the CUDA standard (Nvidia, 2007). CUDA
51 allows easy programming on NVidia GPUs; however a CUDA program cannot run on devices other than NVidia GPUs. This can be
52 problematic and is a leap of faith that NVidia devices are and will remain the most efficient devices for seismic modeling. Also,
53 several clusters offer different types of GPU or, at least, a mix of GPU devices. Hence, the choice of CUDA limits the access to the
54 full power of a cluster. On the other hand, OpenCL (Stone, et al., 2010) is an open programming standard capable of using most
55 existing types of processors and is supported by the majority of manufacturers like Intel, AMD and Nvidia. On NVidia' GPUs,
56 OpenCL performance is comparable to CUDA (Du, et al., 2012). Despite this advantage over CUDA, few published seismic modeling
57 codes use OpenCL: Iturrarán-Viveros and Molero (2013) uses OpenCL in a 2.5D sonic seismic simulation, Kloc and Danek (2013)
58 uses OpenCL for Monte-Carlo full waveform inversion and Molero and Iturrarán-Viveros (2013) perform 2D anisotropic seismic
59 simulations with OpenCL.

60

61 Efficient seismic modeling is more and more needed because of the advent of full waveform inversion (FWI), see Virieux and Operto
62 (2009) for an extensive review. FWI is the process of recovering the Earth (visco)-elastic parameters by directly comparing raw
63 seismic records to the solution of the wave equation (Tarantola, 1984). Its main bottleneck is the numerical resolution of the wave
64 equation that must be repeatedly computed for hundreds if not thousands of shot points for a typical survey. For large-scale multi-
65 parameter waveform inversion, FDTD remains the most plausible solution for seismic modeling (Fichtner, 2011). In addition to
66 forward seismic modeling, FWI requires the computation of the misfit gradient. It can be obtained by the adjoint state method
67 (Plessix, 2006), which requires only an additional forward modeling of the residuals. Hence, it is based on the same modeling
68 algorithm and the benefit of a faster FDTD code would be twofold.

69

70 In this study, we investigate the use of OpenCL for modeling wave propagation in the context of time domain FWI. The main
71 objective is to present a scalable, multi-device portable code for the resolution of the 2D and 3D viscoelastic wave equation that can
72 additionally compute the gradient of the objective function used in FWI by the adjoint state method. This paper does not go into
73 specifics about the inversion process, as the gradient calculations calculated by our algorithm is general and can be used in any
74 gradient-based optimization approach. The seismic modeling program, called SeisCL, is available under a GNU General Public
75 License and is distributed over GitHub. The paper is organized in three parts. First, the equations for viscoelastic wave propagation,
76 its finite-difference solution and the adjoint state method for the calculation of the misfit gradient are briefly discussed. In the second
77 part, different algorithmic aspects of the program are presented in detail. The last section presents numerical results performed on

78 clusters with nodes containing three types of processors: Intel CPUs, NVidia GPUs and Intel Xeon PHI. The numerical results show
79 the validation of the code, the computational speedup over a single threaded implementation and the scaling over several nodes.

80

81 2. Theory

82 2.1 Finite difference viscoelastic wave propagation

83 FWI requires the solution of the heterogeneous wave equation. In this study, we consider the wave equation for an isotropic
84 viscoelastic medium in two and three dimensions. We adopt the velocity-stress formulation in which the viscoelastic effects are
85 modeled by L generalized standard linear solid (Liu, et al., 1976). The symbols used in this article and their meaning are summarized
86 in table 1. The forward problem in 3D is a set of $9 + 6L$ simultaneous equations with their boundary conditions:

87

$$88 \quad \partial_t v_i - \frac{1}{\rho} \partial_j \sigma_{ij} = f_{v_i}, \quad (1a)$$

$$89 \quad \begin{aligned} & \partial_t \sigma_{ij} - \left[M \frac{(1 + \tau_p)}{(1 + \alpha \tau_p)} - 2\mu \frac{(1 + \tau_s)}{(1 + \alpha \tau_s)} \right] \partial_k v_k \delta_{ij} \\ & - \mu \frac{(1 + \tau_s)}{(1 + \alpha \tau_s)} (\partial_j v_i + \partial_i v_j) - r_{ijl} \delta_l = f_{\sigma_{ij}}, \end{aligned} \quad (1b)$$

$$90 \quad \partial_t r_{ijl} + \frac{1}{\tau_{\sigma l}} \left[\left(M \frac{\tau_p}{(1 + \alpha \tau_p)} - 2\mu \frac{\tau_s}{(1 + \alpha \tau_s)} \right) \partial_k v_k \delta_{ij} + \mu \frac{\tau_s}{(1 + \alpha \tau_s)} (\partial_j v_i + \partial_i v_j) + r_{ijl} \right] = 0, \quad (1c)$$

$$91 \quad v_i|_{t=0} = 0, \quad (1d)$$

$$92 \quad \sigma_{ij}|_{t=0} = 0, \quad (1e)$$

$$93 \quad r_{ij}|_{t=0} = 0, \quad (1f)$$

$$94 \quad n_j(\mathbf{s}) \sigma_{ij} = 0, \quad (1g)$$

95 where Einstein summation convention is used over spatial indices i, j, k and the Maxwell body indice l . Equation 1a comes from
96 Newton's second law of motion. Equation 1b is the stress-strain relationship for the generalized standard linear solid model with L
97 Maxwell bodies, which becomes the generalized Hooke's law when the attenuation is nil, i.e. when the attenuation levels τ_s and τ_p
98 are set to zero. Equation 1c gives the variation of the so-called memory variables. Finally, the last four equations are the boundary
99 conditions, respectively a quiescent past for velocities, stresses and memory variables and a free surface. Those equations are

100 discussed in more details in several papers, see for example Carcione, et al. (1988), Robertsson, et al. (1994) and Blanch, et al.
101 (1995).

102

103 The attenuation of seismic waves is often described by the quality factor, defined as the ratio between the real and imaginary parts of
104 the seismic modulus (O'Connell and Budiansky, 1978). In the case of a generalized standard linear solid, it is given by:

$$105 \quad Q(\omega, \tau_{\sigma l}, \tau) = \frac{1 + \sum_{l=1}^L \frac{\omega^2 \tau_{\sigma l}^2 \tau}{1 + \omega^2 \tau_{\sigma l}^2}}{\sum_{l=1}^L \frac{\omega \tau_{\sigma l} \tau}{1 + \omega^2 \tau_{\sigma l}^2}}. \quad (2)$$

106 An arbitrary profile in frequency of the quality factor can be obtained by a least squares minimization over the relaxation times $\tau_{\sigma l}$
107 and the attenuation level τ . Usually, two or three Maxwell bodies are sufficient to obtain a relatively flat quality factor profile over the
108 frequency band of a typical seismic source (Bohlen, 2002). The two variables involved have different influences on the frequency
109 profile of the quality factor: $\tau_{\sigma l}$ controls the frequency peak location of the l^{th} Maxwell body, whereas τ controls the overall quality
110 factor magnitude. In FWI, an attenuation profile in frequency is usually imposed on the whole domain (Askan, et al., 2007, Bai, et al.,
111 2014, Malinowski, et al., 2011) and it is the magnitude of this profile that is sought. For this reason, $\tau_{\sigma l}$ is taken here as a spatially
112 constant variable that is fixed before inversion, whereas τ is let to vary spatially and should be updated through inversion.

113

114 To solve numerically equation 1, we use a finite-difference time-domain approach similar to (Levander, 1988, Virieux, 1986). In time,
115 derivatives are approximated by finite-difference of order 2 on a staggered grid, in which velocities are updated at integer time steps
116 Δt and stresses and memory variables are updated at half-time steps in a classic leapfrog fashion. In space, the standard staggered
117 grid is used. An elementary cell of the standard staggered grid is shown in Figure 1, summarizing the location of each seismic
118 variable. The forward D_i^+ and backward D_i^- differential operators of order $2N$ are given by:

$$119 \quad D_i^+ f(i) = \frac{1}{\Delta x} \sum_{n=1}^N h_n [f(i+n) - f(i-n+1)], \quad (3a)$$

$$120 \quad D_i^- f(i) = \frac{1}{\Delta x} \sum_{n=1}^N h_n [f(i+n-1) - f(i-n)], \quad (3b)$$

121 where Δx is the spatial step and the h_n coefficients are obtained by Holberg's method (Holberg, 1987) which reduces dispersion
122 compared to the Taylor coefficients. The choice of the forward or backward operator obeys the following simple rule: in the update
123 equations (1a, 1b and 1c) of a variable "a", to estimate the derivative of a variable "b", the forward operator is used if variable "b" is

124 located before variable “a” in the elementary cell (Figure 1) along the derivative direction. The backward operator is used otherwise.

125 For example, the update formula for v_x is:

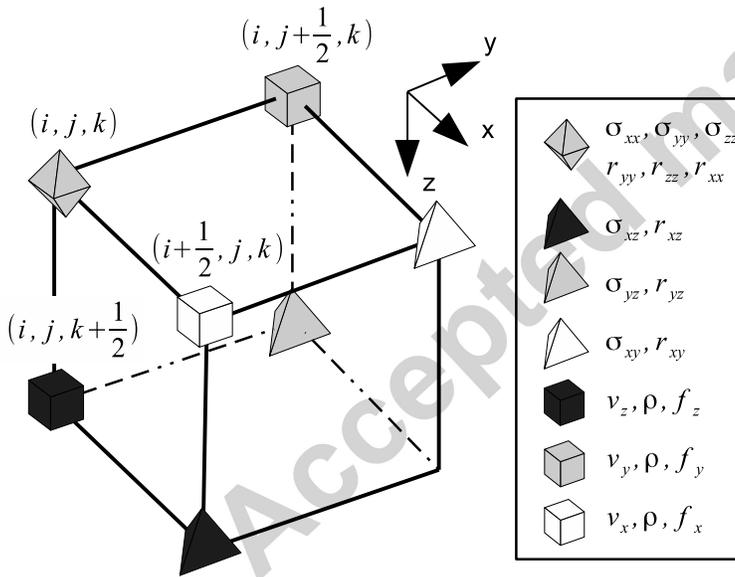
$$v_x^t = v_x^{t-1} + \frac{\Delta t}{\Delta x} \frac{1}{\rho} (D_x^+ \sigma_{xx}^{t-1/2} + D_y^- \sigma_{xy}^{t-1/2} + D_z^- \sigma_{xz}^{t-1/2}). \quad (4)$$

126 The complete set of equations can be obtained with equations 1 and 3 and Figure 1. The reader is referred to the work of Bohlen
127 (2002) for the complete list.

128

129 Finally, to emulate a semi-infinite half-space, artificial reflections caused by the edge of the model must be minimized. For this
130 purpose, two types of absorbing boundaries are implemented: the convolutional perfectly matched layer (CPML) (Roden and
131 Gedney, 2000) as formulated by Komatitsch and Martin (2007) for viscoelastic media and the dissipative layer of (Cerjan, et al.,
132 1985). On the top of the model, a free surface condition is implemented by the imaging method of (Levander, 1988).

133



134

135 Figure 1 An elementary cell showing the node location for each seismic variable.

136

137 Table 1 Symbols used in this article

Symbol	Meaning
$v(\mathbf{x}, t)$	Particle velocity
$\sigma(\mathbf{x}, t)$	Stress
$f(\mathbf{x}, t)$	Source term
$r(\mathbf{x}, t)$	Memory variable

$\overleftarrow{\cdot}$	Adjoint variable
$\rho(\mathbf{x})$	Density
$M(\mathbf{x})$	P-wave modulus
$\mu(\mathbf{x})$	Shear modulus
$Q(\mathbf{x})$	Quality factor
$\tau_p(\mathbf{x})$	P-wave attenuation level
$\tau_s(\mathbf{x})$	S-wave attenuation level
$\tau_{\sigma l}$	Stress relaxation time of the l^{th} Maxwell body
\mathbf{d}	Recorded particle velocities
T	Recording time
N_t	Number of time steps
J	Cost function

138

139

140

2.2 Full waveform inversion

141 The goal of full waveform inversion is to estimate the elastic parameters of the Earth based on a finite set of records of the ground
 142 motion \mathbf{d}_i , in the form of particle velocities or pressure. This is performed by the minimization of a cost function. For example, the
 143 conventional least-squares misfit function for particle velocity measurements is:

$$144 \quad J(\rho, M, \mu, \tau_p, \tau_s) = \frac{1}{2} (\mathbf{S}(v_i) - \mathbf{d}_i)^T (\mathbf{S}(v_i) - \mathbf{d}_i) + \frac{1}{2} (\mathbf{S}(\sigma_{ij}) - \mathbf{d}_{ij})^T (\mathbf{S}(\sigma_{ij}) - \mathbf{d}_{ij}), \quad (5)$$

145 where $\mathbf{S}(\cdot)$ is the restriction operator that samples the wavefield at the recorders' location in space and time. As 3D viscoelastic full
 146 waveform inversion may involve billions of model parameters, the cost function is usually minimized with a local gradient-based
 147 method. However, due to the sheer size of the problem, the computation of the gradient by finite difference is prohibitive. Lailly
 148 (1983) and Tarantola (1984) have shown that the misfit gradient can be obtained by the cross-correlation of the seismic wavefield
 149 with the residuals back propagated in time (see Fichtner, et al. (2006) for a more recent development). This method, called the
 150 adjoint state method, only requires one additional forward modeling. Based on the method of (Plessix, 2006), it can be shown
 151 (Fabien-Ouellet, et al., 2016) that the adjoint state equation for the viscoelastic wave equation of equation 1 is given by:

$$152 \quad \partial_{t'} \tilde{v}_i + \frac{1}{\rho} \partial_j \tilde{\sigma}_{ij} = \frac{1}{\rho} \frac{\partial J}{\partial v_i'}, \quad (6a)$$

$$153 \quad \left[M \frac{(1+\tau_p)}{(1+\alpha\tau_p)} - 2\mu \frac{(1+\tau_s)}{(1+\alpha\tau_s)} \right] \frac{\partial J}{\partial \sigma_{kk}} \delta_{ij} + \mu \frac{(1+\tau_s)}{(1+\alpha\tau_s)} (1 + \delta_{ij}) \frac{\partial J}{\partial \sigma_{ij}}, \quad (6b)$$

$$154 \quad \partial_{t'} \tilde{r}_{ijl} + \frac{1}{\tau_{\sigma l}} \left[\left(M \frac{\tau_p}{(1+\alpha\tau_p)} - 2\mu \frac{\tau_s}{(1+\alpha\tau_s)} \right) \partial_k \tilde{v}_k \delta_{ij} + \mu \frac{\tau_s}{(1+\alpha\tau_s)} (\partial_j \tilde{v}_i + \partial_i \tilde{v}_j) + \tilde{r}_{ijl} \right] = 0, \quad (6c)$$

$$155 \quad \tilde{v}_i|_{t'=0} = 0, \quad (6d)$$

$$156 \quad \tilde{\sigma}_{ij}|_{t'=0} = 0, \quad (6e)$$

$$157 \quad \tilde{r}_{ij}|_{t'=0} = 0, \quad (6f)$$

$$158 \quad n_j(\mathbf{s}) \tilde{\sigma}_{ij} = 0, \quad (6g)$$

159 with $t' = T - t$. Comparing equations 1 and 6, we see that both sets of equations are nearly identical, the only difference being the
 160 sign of the spatial derivatives and the source terms (the terms involving the misfit function derivative). Hence, the adjoint solution for
 161 the viscoelastic wave equation can be computed with the same forward modeling code, with the source term taken as the data
 162 residuals reversed in time and with an opposite sign for the spatial derivatives. This allows using the same modeling code for the
 163 forward and adjoint problem, with only minor changes to store or recompute the forward and residual wavefields. Once both
 164 wavefields are computed, the gradient can be obtained by calculating their scalar product, noted here $\langle \cdot, \cdot \rangle$. The misfit gradient for
 165 density, the P-wave modulus, the P-wave attenuation level, the shear modulus and the S-wave attenuation level are given
 166 respectively by:

167

$$168 \quad \frac{\partial J}{\partial \rho} = \langle \tilde{v}_x, \partial_t v_x \rangle + \langle \tilde{v}_y, \partial_t v_y \rangle + \langle \tilde{v}_z, \partial_t v_z \rangle, \quad (7a)$$

$$169 \quad \frac{\partial J}{\partial M} = -c_M^1 P_1 + c_M^2 P_2, \quad (7b)$$

$$170 \quad \frac{\partial J}{\partial \tau_p} = -c_{\tau_p}^1 P_1 + c_{\tau_p}^2 P_2, \quad (7c)$$

$$171 \quad \frac{\partial J}{\partial \mu} = -c_\mu^1 P_3 + c_\mu^2 P_1 - c_\mu^3 P_4 + c_\mu^4 P_5 - c_\mu^5 P_2 + c_\mu^6 P_6, \quad (7d)$$

$$172 \quad \frac{\partial J}{\partial \tau_s} = -c_{\tau_s}^1 P_3 + c_{\tau_s}^2 P_1 - c_{\tau_s}^3 P_4 + c_{\tau_s}^4 P_5 - c_{\tau_s}^5 P_2 + c_{\tau_s}^6 P_6, \quad (7e)$$

173 with

$$174 \quad P_1 = \langle \tilde{\sigma}'_{xx} + \tilde{\sigma}'_{yy} + \tilde{\sigma}'_{zz}, \partial_t (\sigma'_{xx} + \sigma'_{yy} + \sigma'_{zz}) \rangle,$$

$$175 \quad P_2 = \langle \tilde{R}_{xxl} + \tilde{R}_{yy l} + \tilde{R}_{zz l}, (1 + \tau_{\sigma l} \partial_t) (r_{xxl} + r_{yy l} + r_{zz l}) \rangle,$$

$$176 \quad P_3 = \langle \tilde{\sigma}'_{xy}, \partial_t \sigma'_{xy} \rangle + \langle \tilde{\sigma}'_{xz}, \partial_t \sigma'_{xz} \rangle + \langle \tilde{\sigma}'_{yz}, \partial_t \sigma'_{yz} \rangle,$$

$$\begin{aligned}
P_4 &= \left\langle \tilde{\sigma}'_{xx}, \partial_t \left((N_d - 1)\sigma'_{xx} - \sigma'_{yy} - \sigma'_{zz} \right) \right\rangle \\
177 \quad &+ \left\langle \tilde{\sigma}'_{yy}, \partial_t \left((N_d - 1)\sigma'_{yy} - \sigma'_{xx} - \sigma'_{zz} \right) \right\rangle \\
178 \quad &+ \left\langle \tilde{\sigma}'_{zz}, \partial_t \left((N_d - 1)\sigma'_{zz} - \sigma'_{yy} - \sigma'_{xx} \right) \right\rangle,
\end{aligned}$$

$$\begin{aligned}
P_5 &= \langle \tilde{R}_{xyl}, (1 + \tau_{\sigma l} \partial_t) r_{xyl} \rangle \\
179 \quad &+ \langle \tilde{R}_{xzl}, (1 + \tau_{\sigma l} \partial_t) r_{xzl} \rangle \\
180 \quad &+ \langle \tilde{R}_{yzl}, (1 + \tau_{\sigma l} \partial_t) r_{yzl} \rangle,
\end{aligned}$$

$$\begin{aligned}
P_6 &= \left\langle \tilde{R}_{xxl}, (1 + \tau_{\sigma l} \partial_t) \left((N_d - 1)r_{xxl} - r_{yyl} - r_{zzl} \right) \right\rangle \\
181 \quad &+ \left\langle \tilde{R}_{yy l}, (1 + \tau_{\sigma l} \partial_t) \left((N_d - 1)r_{yy l} - r_{xxl} - r_{zzl} \right) \right\rangle \\
182 \quad &+ \left\langle \tilde{R}_{zzl}, (1 + \tau_{\sigma l} \partial_t) \left((N_d - 1)r_{zzl} - r_{yyl} - r_{xxl} \right) \right\rangle \quad (7f)
\end{aligned}$$

183 where $\tilde{R}_{ijl} = \int_0^T \tilde{r}_{ijl} dt$, $\sigma'_{ij} = \sigma_{ij} - \sum_l R_{ijl}$ and N_d is the number of dimensions (2 or 3). Coefficients c are given in the
184 appendix. The misfit gradients for the P-wave modulus M and the P-wave attenuation level τ_p have the same structure and differ
185 only by the coefficients that weight the scalar products. The same relationship exists between μ and τ_s .

186 In the time domain, the scalar product takes the form:

$$187 \quad \langle a(t), b(t) \rangle = \int_T a(t)b(t) dt, \quad (8)$$

188 which is the zero-lag cross-correlation in time of the real-valued functions $a(x)$ and $b(x)$. When discretized in time, it is the sum of
189 the product of each sample. Using Parseval's formula, the last equation can also be expressed in the frequency domain:

$$190 \quad \langle a(t), b(t) \rangle = \frac{1}{2\pi} \int_{\omega} A^*(\omega)B(\omega) d\omega, \quad (9)$$

191 where $A(\omega)$ and $B(\omega)$ are the Fourier transform of the functions $a(t)$ and $b(t)$ and $*$ indicates complex conjugation. The
192 formulation in frequency can be used to perform frequency domain FWI (Pratt and Worthington, 1990) with a time-domain forward
193 modeling code as done by Nihei and Li (2007). The frequency components of the seismic variables can be obtained with the discrete
194 Fourier transform:

$$195 \quad A(m\Delta f) = \Delta t \sum_{n=0}^{N_t-1} a(n\Delta t) \exp \left[-\frac{i2\pi mn}{N_t} \right], \quad (10)$$

196 where a is the discrete function in time, A is the discrete function in the Fourier domain, Δt is the time interval, Δf is the frequency
197 interval and m is the frequency label. The calculation of a frequency component with the discrete Fourier transform involves the sum
198 of all the time samples weighted by a time varying function given by the complex exponential. In the FDTD scheme, the running sum

199 can be updated at each time step for all or a selected number of frequencies (Furse, 2000). Because FDTD must be oversampled to
200 remain stable (CFL condition), the discrete Fourier transform can be performed at a higher time interval to mitigate its computational
201 cost, e.g. several time steps can be skipped in equation 10, up to the Nyquist frequency of the highest selected frequency. Also, to
202 save memory and reduce computing time, only a handful of frequencies can be used during the inversion (Sirgue and Pratt, 2004).

203
204 Once the gradient is computed, different algorithms can be used to solve the inversion system, from the steepest descent to the full
205 Newton method (Virieux and Operto, 2009). This issue is not the focus of this study. However, all of these local methods need at
206 least the computation of the forward model and the misfit gradient, both of which are the main computational bottlenecks. Hence, a
207 faster forward/adjoint program should benefit all of the local approaches of FWI.

209 2.3 Background on heterogeneous computing

210 Heterogeneous computing platforms have become the norm in the high-performance computing industry. Clusters generally include
211 different kinds of processors (Dongarra, et al., 2015): the most common being CPUs, GPUs and Many Integrated Core (MIC), also
212 known as accelerators. Those devices may possess different architecture and usually codes written for one type of device is not
213 compatible with others. Writing a specific code for each type of processor can be tedious and non-productive. One solution is given
214 by OpenCL (Stone, et al., 2010), an open standard cross-platform for parallel programming. OpenCL allows the same code to use
215 one or a combination of processors available on a local machine. This portability is the main strength of OpenCL, especially with the
216 actual trend of massively parallel processors. For the moment, it cannot be used for parallelization on a cluster, but can be used in
217 conjunction with MPI.

218
219 Even though OpenCL allows the same code to be compatible with different devices, the programmer always has to make a choice
220 with the initial design because code optimization can be very different for CPUs, GPUs or MICs architectures. The program
221 presented in this study was written for the GPU architecture, which is arguably the most efficient type of processor available today for
222 finite-difference algorithms. For a good summary of the concepts of GPU computing applied to seismic finite-difference, see (Michéa
223 and Komatitsch, 2010). Essential elements to understand the rest of the article are presented in this section, using the OpenCL
224 nomenclature.

226 A GPU is a device designed to accelerate the creation and manipulation of images, or large matrices, intended primarily for output to
227 a display. It is separated from the CPU (host) and usually does not directly share memory. The set of instructions that can be
228 accomplished on a GPU is different than on the CPU, and classical programming languages cannot be used. A popular application
229 programming interface for GPUs is CUDA (Nvidia, 2007). However, CUDA is a closed standard owned by NVIDIA that can only be
230 used with Nvidia GPUs. It is the main reason why OpenCL was favored over CUDA in this work.

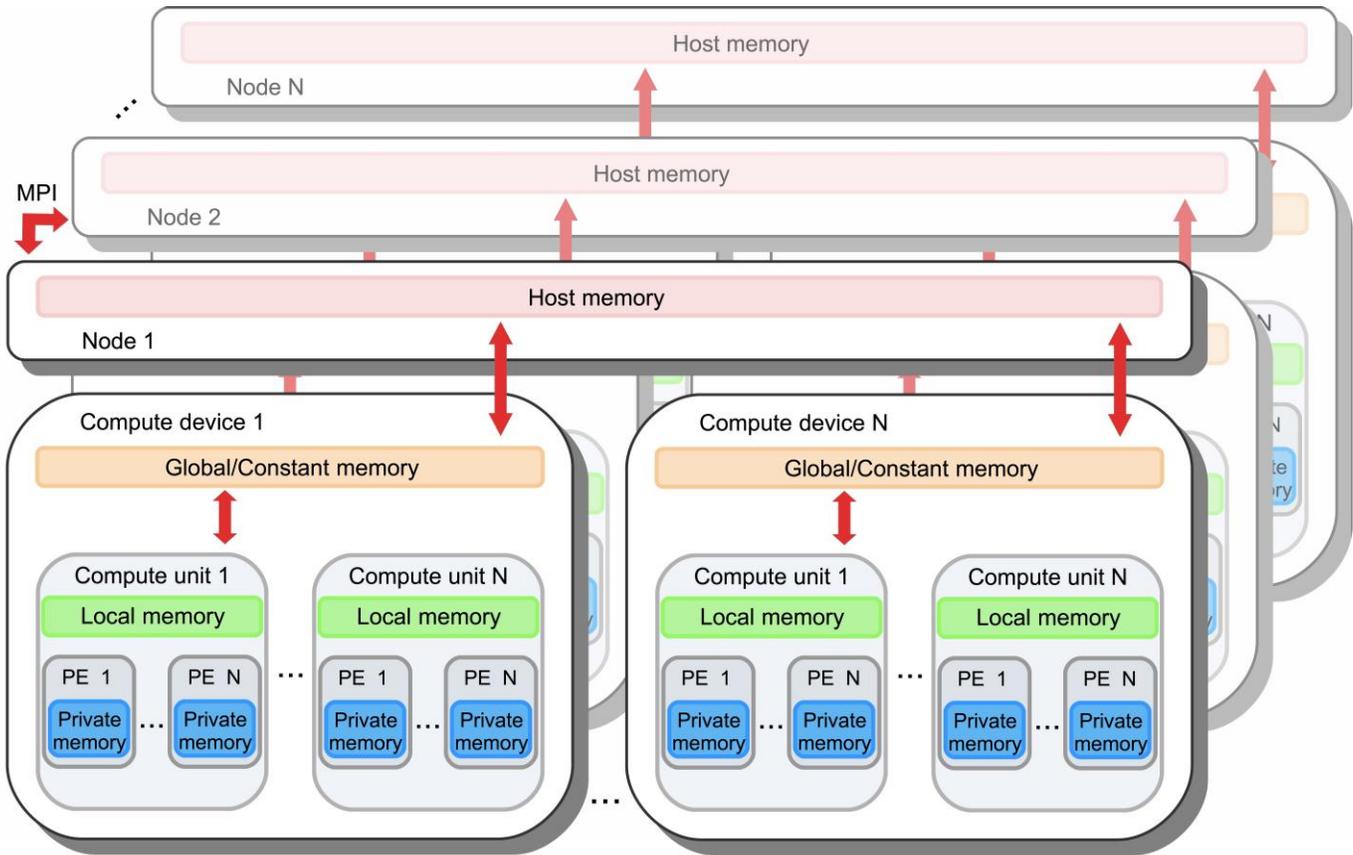
231

232 In order to code efficiently for GPUs, it is important to understand their architecture. The smallest unit of computation is a work item
233 (a thread in CUDA) and is executed by the processing elements (CUDA cores in the NVidia nomenclature). A single device can
234 contain thousands of processing elements that execute the same control flow (instructions) in parallel on different data in the single
235 instruction, multiple thread fashion. The processing elements are part of groups that are called compute units (thread blocks in
236 CUDA). In NVidia devices, the compute units contain 32 consecutive processing elements. In OpenCL, the programmer sends the
237 work items, organized into work groups, to be computed by the processing elements of a compute unit, located in a given device.

238

239 Several levels of memory exist in a GPU. This is schematized in Figure 2, in the context of a GPU cluster. First, each processing
240 element has its own register (private memory), a limited in size but very fast memory. Second, inside each compute unit, threads
241 share a low-latency memory, called the local memory. This memory is small, usually in the order of several kilobytes. The main
242 memory, called global memory, is shared between all processing elements and is the place where the memory needed for the
243 different kernels is located. Usually, this memory is not cached and is very slow compared to the local or private memory.

244 One of the most important aspects of GPU programming is the access to the global memory. Depending on the memory access
245 pattern, read/write operations can be performed in a serial or a parallel fashion by the compute units. Parallel (coalesced) memory
246 access is possible when a number of consecutive work items inside a work group performing the same instructions are accessing
247 consecutive memory addresses. For most NVidia devices, consecutive work items, or what is called a warp, can read 32 floats in a
248 single instruction when memory access is coalesced. With finite-difference codes, the number of instructions performed between the
249 read/write cycles in global memory is fairly low, which means that kernels are bandwidth limited. The memory access pattern is then
250 one of the main areas that should be targeted for optimization.



251

252 Figure 2 OpenCL memory diagram used in conjunction with MPI in the context of a cluster, inspired by (Howes and Munshi, 2014).

253

254 In practice, a program based on OpenCL is organized as follows, regardless of the type of processor used. First, instructions are
 255 given to the host to detect the available devices (GPUs, CPUs or accelerators) and connect them in a single computing context.
 256 Once the context is established, memory buffers used to transfer data between the host and the devices must be created. Then, the
 257 kernels are loaded and compiled for each device. This compilation is performed at runtime. The kernels are pieces of code written in
 258 C that contain the instruction to be computed on the devices. After that, the main part of the program can be executed, in which
 259 several kernels and memory transfers occur, managed on the host side by a queuing system. Finally, buffers must be released
 260 before the end of the program. Several OpenCL instances can be synchronized with the help of MPI, as shown in Figure 2.

261

262 3. Program structure

263 This section describes the implementation of the finite-difference algorithm for viscoelastic modeling and the calculation of the adjoint
 264 wavefield in an OpenCL/MPI environment. The program contains many kernels, and its simplified structure is shown in Algorithm 1.

265 This algorithm presents a typical gradient calculation over several seismic shots, on a parallel cluster where each node contains
 266 several devices. Its main features are discussed in the following sections.

267

268 **Algorithm 1 Pseudo-code for the parallel computation of the gradient with the adjoint state method.**

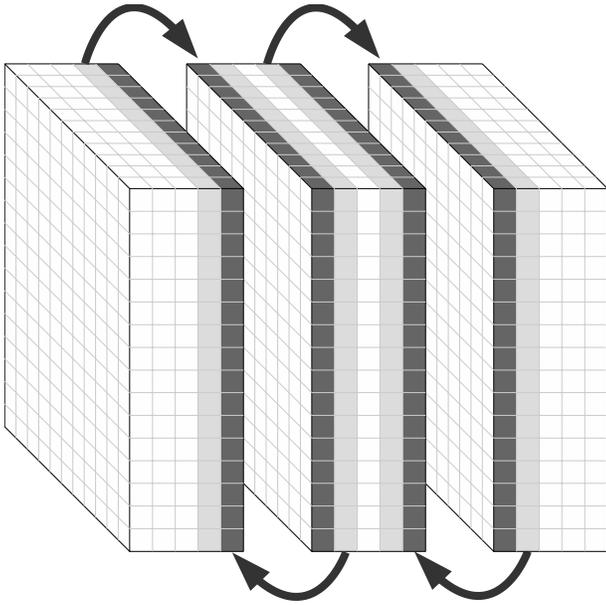
269 **Initialize MPI**
 270 **Initialize OpenCL**
 271 **Initialize model grid**
 272 1. **for all groups in MPI do**
 273 2. **for all shots in group do**
 274 3. **for all nodes in group do**
 275 4. **for all devices in node do**
 276 5. **Initialize seismic grid** ($v_i, \sigma_{ij}, r_{ij}, \tilde{v}_i, \tilde{\sigma}_{ij}, \tilde{r}_{ij}$)
 277 6. **Execute time stepping on shot**
 278 7. **Compute residuals**
 279 8. **Execute time stepping on residuals**
 280 9. **Compute gradient**
 281 10. **end for**
 282 11. **end for**
 283 12. **end for**
 284 13. **end for**

285

286 3.1 Node and device parallelism

287 In order to take advantage of large clusters, we use the MPI interface to parallelize computations between the nodes of a cluster. A
 288 popular approach to parallelizing finite-difference seismic modeling is domain decomposition (Mattson, et al., 2004). It consists of
 289 dividing the model grid into subdomains that can reside on different machines. At each time step, each machine updates its own
 290 velocity and stress sub-grids. As the finite-difference update of a variable at a given grid point requires the values of other variables
 291 at neighboring grid points (see equations 3 and 4), values defined at grid points on the domain boundary must be transferred
 292 between adjacent domains at each time steps. This is depicted in Figure 3.

293



294

295 **Figure 3 Domain decomposition for three devices for a finite-difference order of 2. Light gray cells are updated inside the device and transferred to the**
 296 **dark gray cells of the adjacent device.**

297

298

299

300

301

302

303

304

305

306

307

308

309

310

311

312

313

Fast interconnects are needed for this memory transfer that occurs at each time step, otherwise the scaling behavior can become unfavorable. For example, Bohlen (2002) observes super-linear scaling for up to 350 nodes on a cluster with 450 Mb/s interconnects, but only linear scaling with up to 12 nodes on a cluster with 100 Mb/s interconnects. When using GPUs, not only transfers are needed between nodes, but also between the devices and the host. This dramatically worsens performance. For example, Okamoto (2011) observes a scalability between N and $N^{2/3}$. For this reason, we chose to implement two different parallelism schemes in addition to the inherent OpenCL parallelization: domain decomposition and shot parallelization.

Nodes of a cluster are first divided into different groups: within each group, we perform domain decomposition and each group is assigned a subset of shots. Shot parallelism best corresponds to a task-parallel decomposition, and is illustrated in Algorithm 1 by the loop on all the groups of nodes that starts at line 1, and by the loop on all shots assigned to the groups at line 2. Parallelizing shots does not require communication between nodes and should show a linear scaling. Let's mention that a typical seismic survey involves hundreds if not thousands of shot points. This should be at least on par with the number of available nodes on large clusters. On the other hand, domain decomposition is required to enable computations for models exceeding the memory capacity of a single device. For this level of parallelism, MPI manages communications between nodes and the OpenCL host thread manages the local devices. The communications managed by MPI and OpenCL are illustrated respectively by the loop on all nodes belonging to the same group starting at line 3 of Algorithm 1 and by the loop on all devices found on the node starting at line 4.

314

315 To further mitigate the communication time required in domain decomposition, the seismic updates are divided between grid points
 316 on the domain boundary that needs to be transferred and interior grid points that are only needed locally. This is described in
 317 Algorithm 2. The grid points on the boundary are first updated, which allows overlapping the communication and the computation of
 318 the remaining grid points, i.e. lines 3 and 4 and lines 7 and 8 of Algorithm 2 are performed simultaneously for devices supporting
 319 overlapped communications. This is allowed in OpenCL by having two different queues for each device: one for buffer
 320 communication and the other for kernel calls.

321

322 **Algorithm 2 Pseudo code showing the overlapping computation and memory transfer for domain decomposition.**

```

323 1. while  $t < N_t$ 
324 2.   Call kernel_updatev on domain boundary
325 3.   Transfer  $v_i$  in boundary of devices, nodes
326 4.   Call kernel_updatev on domain interior
327 5.   Store  $S(v_i)$  in seismo at  $t$ 
328 6.   Call kernel_updates on domain boundary
329 7.   Transfer  $\sigma_{ij}$  in boundary of devices, nodes
330 8.   Call kernel_updates on domain interior
331 9.   Increment  $t$ 
332 10. end while
333

```

334

3.2 GPU kernels

335 The main elements of the kernels used to update stresses and velocities are shown in Algorithm 3. For better readability, the
 336 algorithm is simplified and does not include viscoelastic computations or CPML corrections. Note that the “for” loops in this pseudo-
 337 code are implicitly computed by OpenCL. The most important features of this algorithm are steps 3 and 4, where seismic variables
 338 needed in the computation of the spatial derivatives are loaded from the global memory to the local memory. As the computation of
 339 the spatial gradient of adjacent grid elements repeatedly uses the same grid points, this saves numerous reads from global memory.
 340 To be effective, those reads must be coalesced. This is achieved by setting the local working size in the z dimension, which is the
 341 fast dimension of the arrays, to a multiple of 32 for NVidias’ GPUs. Hence, seismic variables are updated in blocks of 32 in the z
 342 dimension. In the x and y dimensions, the size of the local working size does not impact coalesced memory reading. They are set
 343 equal to a magnitude that allows all the seismic variables needed in the update to fit in the local memory. This is illustrated in Figure

344 4.

345

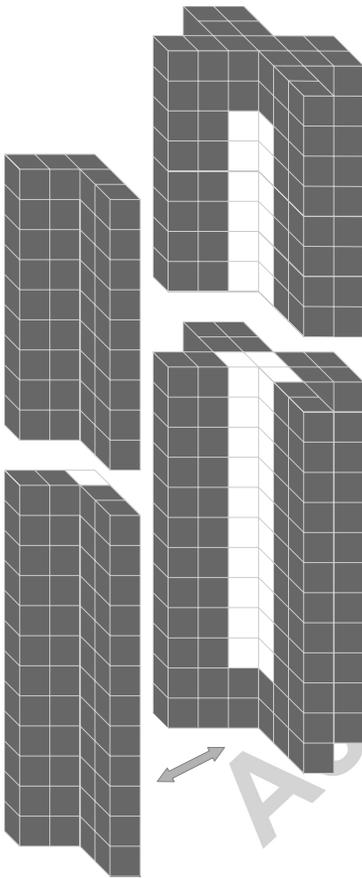
346 Algorithm 3 Pseudo code for the seismic update kernels showing how local memory is used.

```

347 1. for all local_domains in global_domain do
348 2.   for all grid point in local_domain do
349 3.     Load  $v_i$  ( $\sigma_{ij}$ ) from global to local memory
350 4.     Compute  $\partial_i v_i$  ( $\partial_i \sigma_{ij}$ ) from local memory
351 5.     Update  $\sigma_{ij}$  ( $v_i$ ) in global memory
352 6.   end for
353 7. end for
354

```

355



356

357 Figure 4 Exploded view of the local memory containing a seismic variable during update (equations 1a and 1b), for the 2nd order scheme. White cells
358 are cells being updated and gray cells are loaded into local memory only to update white cells.

359

360 3.3 Misfit gradient computation

361 The cross-correlation of the direct and residual fields requires both fields to be computed at the same time step (see equation 8).

362 This is challenging because forward computations are performed from time zero, whereas adjoint computations are performed from

363 final time. Several strategies can be employed to achieve this task (see (Dussaud, et al., 2008, Nguyen and McMechan, 2015) for
364 comparisons between methods).

- 365 1. When propagating the direct field, the whole grid for the particle velocities and stresses at each time step or a subset of the
366 time steps can be saved into memory. When the residual field is propagated from final time, the direct field is read from
367 memory for all grid points and the scalar product is evaluated iteratively, time step per time step.
- 368 2. In the so-called the backpropagation scheme (Clapp, 2008, Yang, et al., 2014), only the outside boundary of the grid that is
369 not in the absorbing layer is saved at each time step. The direct field is recovered during the residual propagation by
370 propagating back in time the direct field from the final state, injecting the saved wavefield on the outside boundary at each
371 time step. As both the forward and adjoint wavefields are available at the same time step, the scalar products can be
372 computed directly with equation 8.
- 373 3. A selected number of frequencies of the direct and residual field can be stored. This is performed by applying the discrete
374 Fourier transform incrementally at each time step (equations 9 and 10), as done by (Sirgue, et al., 2008). The scalar product
375 is evaluated at the end of the adjoint modeling in the frequency domain with equation 9. An alternative way of computing the
376 chosen frequencies (Furse, 2000) seems to be advantageous over the discrete Fourier transform, but has not been tested
377 in this study.
- 378 4. In the optimal checkpointing method proposed by (Griewank, 1992, Griewank and Walther, 2000), and applied by (Symes,
379 2007), the whole forward wavefield is stored for a limited number of time steps or checkpoints. To perform the scalar
380 product, the forward wavefield is recomputed for each time step during the backpropagation of the residuals from the
381 nearest checkpoint. For a fixed number of checkpoints, an optimal distribution that minimizes the number of forward
382 wavefield that has to be recomputed can be determined. For this optimal distribution, the number of checkpoints and the
383 number of recomputed time steps evolve logarithmically with the number of total time steps. Further improvements of the
384 method have been proposed by (Anderson, et al., 2012) and by (Komatitsch, et al., 2016) in the viscoelastic case.

385
386 The first option is usually impractical, as it requires a huge amount of memory even for problems of modest size. In 3D, it requires on
387 the order of $O(N_t N^3)$ elements to be stored, which becomes quickly intractable. Let's mention that the use of compression and
388 subsampling can be used to mitigate these high memory requirements (Boehm, et al., 2016, Sun and Fu, 2013). The
389 backpropagation scheme requires far less memory, on the order $O(N_t N^2)$ in 3D, but doubles the computation task for the direct
390 field. Also, it is not applicable in the viscoelastic case. Indeed, in order to back-propagate the wavefield, the time must be reversed

391 $t \rightarrow -t$ and, doing so, the memory variable differential equation (equation 1c) becomes unstable. Hence, when dealing with
 392 viscoelasticity, the frequency scheme and the optimal checkpointing scheme are the only viable options. The memory requirement of
 393 the frequency scheme grows with the number of computed frequencies on the order of $O(N_f N^3)$. However, as is common in FWI,
 394 only a selected number of frequencies can be used (Virieux and Operto, 2009). The optimal checkpointing method requires
 395 $O(N_c N^3)$ where N_c is the number of checkpoints. Because of the logarithmic relationship between the number of time steps, the
 396 number of checkpoints and the number of additional computations, the required memory should stay tractable. For example, for 10
 397 000 time steps, with only 30 buffers, the computing cost of this option is 3.4 times that of the forward modeling. In this work, we
 398 implemented the backpropagation scheme for elastic propagation and the frequency scheme using the discrete Fourier transform for
 399 both elastic and viscoelastic propagation. The implementation of the optimal checkpointing scheme or the hybrid
 400 backpropagation/checkpointing scheme of (Yang, et al., 2016) is left for future work.

401

402 The gradient computation involving the backpropagation of the direct field is illustrated in Algorithm 4. At each time step of the direct
 403 field propagation, the wavefield value at grid points on the outer edge of the model is stored. Because of the limited memory capacity
 404 of GPUs, this memory is transferred to the host. As mentioned before, this communication can be overlapped with other
 405 computations with the use of a second queue for communication. After obtaining the residuals, the residual wavefield is propagated
 406 forward in time using the same kernel as the direct wavefield. The back-propagation of the direct wavefield is calculated using the
 407 same kernel, the only difference being the sign of the time step $\Delta t \rightarrow -\Delta t$. Also, at each time step, the stored wavefield on the
 408 model edges is injected back. With this scheme, both the residual and the direct fields are available at each time step and can be
 409 multiplied to perform on the fly the scalar products needed to compute the gradient.

410 **Algorithm 4 Pseudo code for the backpropagation scheme.**

- 411 1. **while** $t < N_t$
- 412 2. **Call** kernel_updatev
- 413 3. **Store** v_i in *boundary of model*
- 414 4. **Call** kernel_updates
- 415 5. **Store** σ_{ij} in *boundary of model*
- 416 6. **Increment** t
- 417 7. **end while**
- 418 8. **Calculate** residuals
- 419 9. **while** $t < N_t$
- 420 10. **Call** kernel_updatev on v_i, \bar{v}_i
- 421 11. **Inject** v_i in *boundary of model*
- 422 12. **Call** kernel_updates on $\sigma_{ij}, \bar{\sigma}_{ij}$

- 423 13. **Inject** σ_{ij} in *boundary of model*
 424 14. **Compute** *gradient*
 425 15. **Increment** t
 426 16. **end while**
 427

428 The frequency scheme is illustrated in Algorithm 5. It first involves computing the direct wavefield and its discrete Fourier transform
 429 on the fly at each time step, for each desired frequency (equation 10). Afterward, the residual wavefield is obtained in exactly the
 430 same fashion. At the end, the scalar product needed for the gradients can be computed with the selected frequencies.

431

432 **Algorithm 5 Pseudo code for the frequency scheme.**

- 433 1. **while** $t < N_t$
 434 2. **Call** *kernel_updatev* for v_i
 435 3. **Call** *kernel_updates* for σ_{ij}
 436 4. **Call DFT for** v_i, σ_{ij} for *freqs*
 437 5. **Increment** t
 438 6. **end while**
 439 7. **Compute** *residuals*
 440 8. **while** $t < N_t$
 441 9. **Call** *kernel_updatev* for \tilde{v}_i
 442 10. **Call** *kernel_updates* for $\tilde{\sigma}_{ij}$
 443 11. **Call DFT for** $\tilde{v}_i, \tilde{\sigma}_{ij}$ for *freqs*
 444 12. **Increment** t
 445 13. **end while**
 446 14. **Compute** *gradients*
 447

448 4. Results and discussion

449 This section shows several numerical results obtained with SeisCL. The following tests were chosen to verify the performance of
 450 OpenCL in the context of FWI on heterogeneous clusters containing three different types of processors: Intel CPUs, Intel Xeon PHI
 451 (MIC) and NVidia GPUs.

452 4.1 Modeling validation

453 In order to test the accuracy of our forward/adjoint modeling algorithm, two synthetic cases are presented. First, the finite-difference
 454 solution of the viscoelastic wave equation is compared to the analytic solution. The analytic solution for the viscoelastic wave
 455 propagation of a point source derived by Pilant (2012) is used here in the form given by Gosselin-Cliche and Giroux (2014) for a

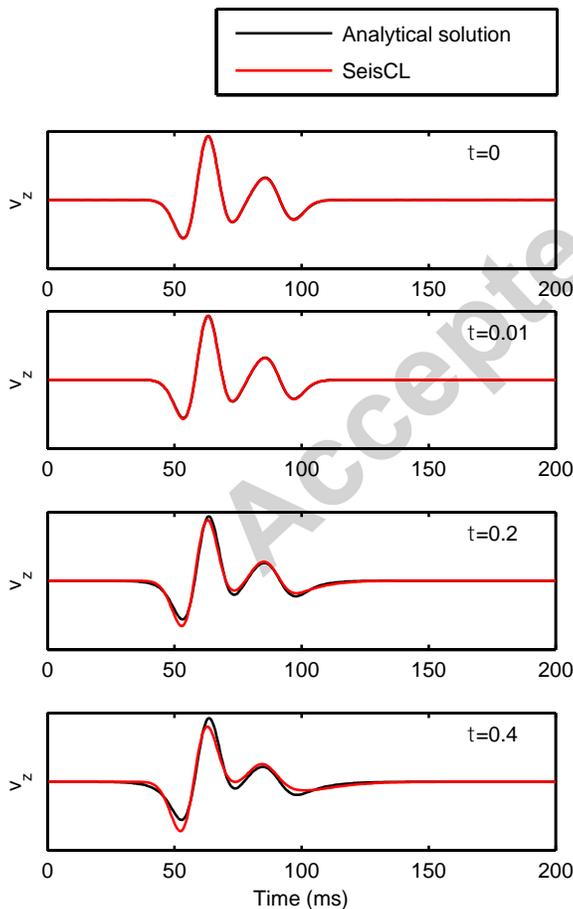
456 quality factor profile corresponding to a single Maxwell body. The source is a Ricker wavelet with a center frequency of 40 Hz,
457 oriented in the z direction. The viscoelastic model is homogeneous with $V_p=3500$ m/s, $V_s=2000$ m/s, $\rho=2000$ m/s with a single
458 Maxwell body. We tested 4 attenuation levels $\tau_p = \tau_s = \{0, 0.01, 0.2, 0.4\}$, i.e. $Q = \{\infty, 200, 10, 5\}$ at the center frequency of
459 40 Hz. Using a finite-difference stencil of order 4, a 6 m (8.33 points per wavelength) spatial discretization is used to avoid numerical
460 dispersion with a 0.5 ms time step for numerical stability. Figure 5 shows the comparison between the analytical solution and the
461 solution obtained with SeisCL. The traces represent the particle velocities in the z direction for an offset of 132 m in the z direction.
462 For the elastic case ($\tau=0$), the analytical solution is perfectly recovered by SeisCL. Using higher attenuation levels does, however,
463 introduce some errors in the solution. This error increases with τ and for an attenuation level of 0.4, the discrepancy becomes
464 obvious for the offset used herein. It is, however, the expected drawback of using an explicit time domain solution and similar time-
465 domain algorithms show the same behavior, see (Gosselin-Cliche and Giroux, 2014). Also, for reasonable attenuation levels, the
466 errors appear negligible and will not impact FWI results much. Accuracy could become an issue for very high attenuating media and
467 long propagation distances.

468

469 The second test aims at validating the misfit gradient output of SeisCL. For this test, a synthetic 2D cross-well tomographic survey is
470 simulated, where a model perturbation between two wells is to be imaged. The well separation is 250 m and the source and receiver
471 spacing are respectively 60 m and 12 m (Figure 5). Circular perturbations of a 60 m radius for the five viscoelastic parameters were
472 considered at five different locations. The same homogeneous model as the first experiment is used with $\tau=0.2$ and with
473 perturbations of 5 % of the constant value. Because significant crosstalk can exist between parameters, especially between the
474 velocities and the viscous parameters (Kamei and Pratt, 2013), we computed the gradient for one type of perturbation at a time. For
475 example, the P-wave velocity gradient is computed with constant models for all other parameters other than V_p . This eliminates any
476 crosstalk between parameters and allows a better appraisal of the match between the gradient update and the given perturbations.
477 Note that because the goal of the experiment is to test the validity of the approach, geological plausibility was not considered. As no
478 analytical solution exists for the gradient, the adjoint state gradient was compared to the gradient computed by finite-difference. The
479 finite-difference solution was obtained by perturbing each parameter of the grid sequentially by 2%, for all the grid position between
480 the two wells. The adjoint state gradient was computed with the frequency scheme using all frequencies of the discrete Fourier
481 transform between 0 and 125 Hz.

482

483 The results of this second experiment are shown in Figure 6. In this figure, each column represents a different perturbed parameter.
 484 The first row shows the perturbation, the second the steepest descent direction (minus the misfit gradient) obtained by finite-
 485 difference and the third the steepest descent direction given by the adjoint state model. Note that the gradients were normalized in
 486 this figure. As can be visually appraised, an excellent agreement is obtained between both methods, for all parameters. Although the
 487 inversion has not been performed here, it should converge to the right solution in the five different cases, the update correction being
 488 already in the right direction. This is expected considering the small value of the perturbation used in this experiment; the inverse
 489 problem is more or less linear in such circumstances. The good agreement between the finite-difference and the adjoint state
 490 gradients shows that the latter could be used in any gradient-based inversion approach. However, the adjoint approach is orders of
 491 magnitude faster than the finite-difference approach: the first grows proportionally to the number of frequencies (see next section)
 492 while the second grows linearly with the number of parameters. For this particular experiment, the adjoint approach required minutes
 493 to complete whereas the finite-difference approach required days.
 494



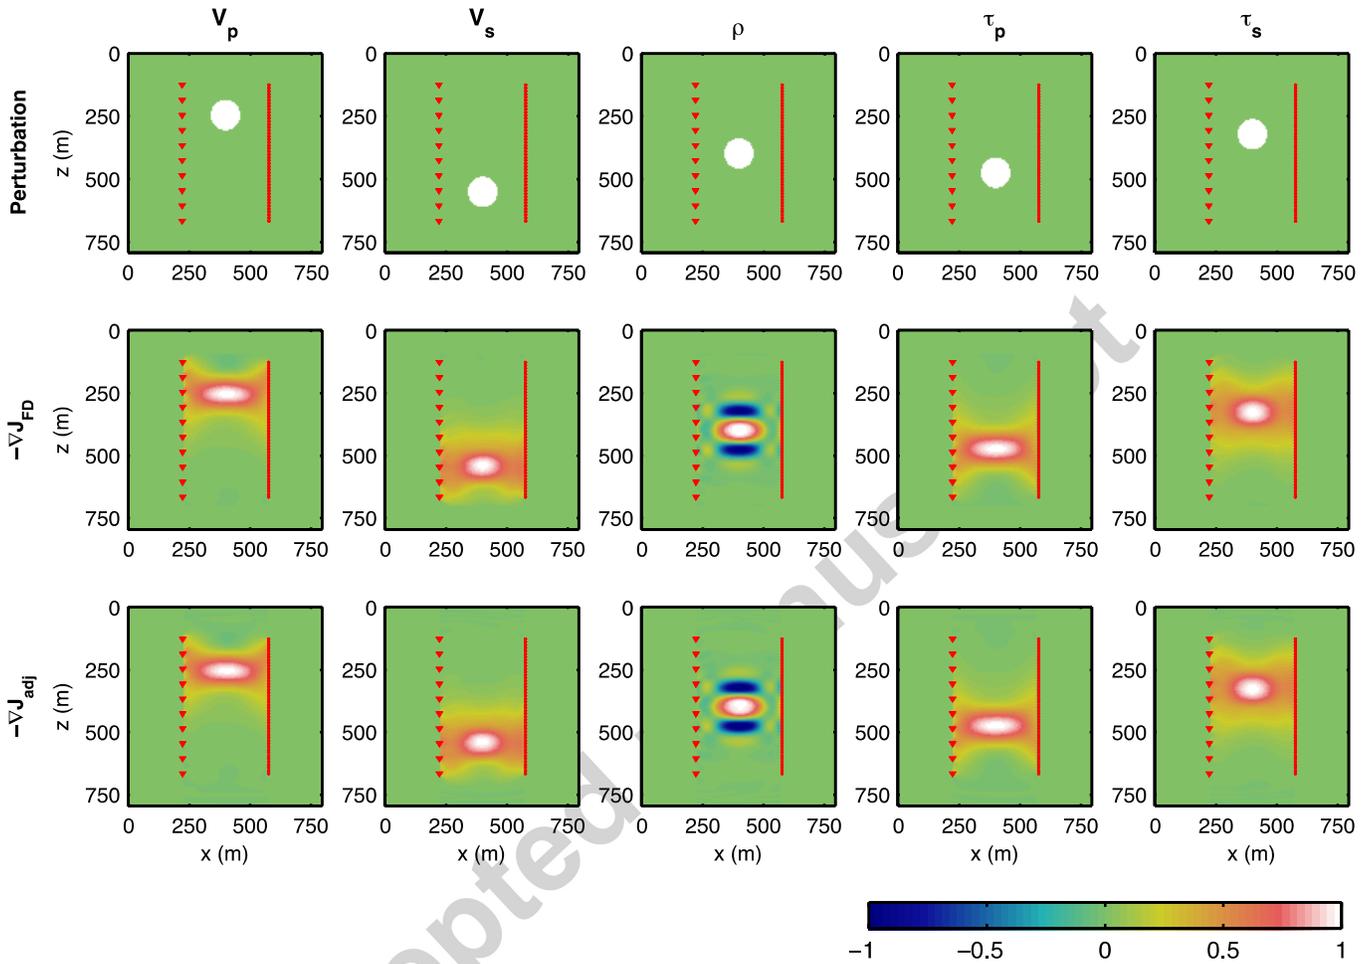
495

496 **Figure 5 Comparison between the analytical solution and SeisCL results for different attenuation levels, from the elastic case ($\tau=0$) to strong**
 497 **viscoelasticity ($\tau=0.4$).**

498

499

500



501

502 Figure 6 A cross-well experiment to test the validity of the misfit gradient. The red triangles represent the sources position and the red dots the
 503 receiver positions. Each column represents a different parameter. The first row shows the location of the perturbation, the second row represents the
 504 opposite of the misfit gradient obtained by finite-difference and the third row represents the opposite of the misfit gradient by the adjoint
 505 state method.

506

507 4.2 Performance comparison

508 The effort required to program with the OpenCL standard would be vain without a significant gain in the computing performance. In

509 the following, several tests are presented to measure the performance of SeisCL. As a measure, one can compute the speedup,

510 defined here as:

$$511 \text{ Speedup} = \frac{T_{\text{baseline}}}{T_{\text{SeisCL}}} \quad (11)$$

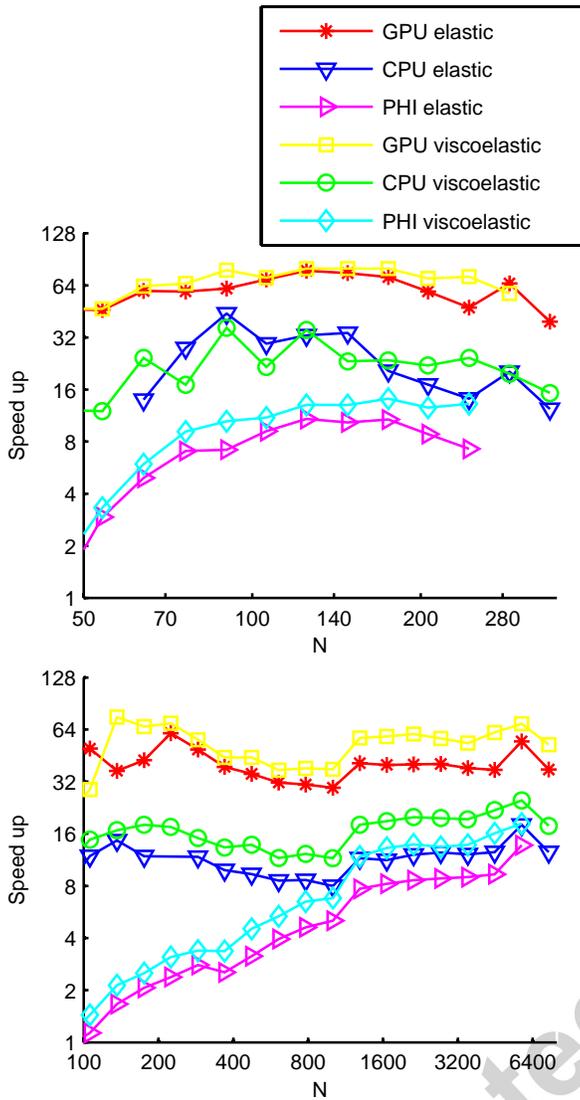
512 Different baselines are used depending on the test. In order to show the OpenCL compatibility of different devices, all tests are
513 performed on three types of processors: Intel CPUs, Intel Xeon PHI (MIC) and NVidia GPUs. Unless stated otherwise, the CPU
514 device consists of 2 Intel Xeon E5-2680 v2 processors with 10 cores each at a frequency of 2.8 GHz and with 25 MB of cache. The
515 GPU is an NVidia Tesla K40 with 2880 cores and 12 GB of memory, and the MIC is an Intel Xeon Phi 5110P.

516

517

518 4.2.1- Speedup using SeisCL over a single threaded CPU implementation

519 As a baseline, *SOFI2D* and *SOFI3D*, the 2D and 3D implementations of Bohlen (2002) are used with a single core. This baseline can
520 be compared to SeisCL as both codes use the same algorithm. It is also representative of the speed that can be achieved for a
521 FDTD code written in C, arguably one of the fastest high level languages for the CPU. In Figure 7, the speedup is measured as a
522 function of the model size for the 3D and 2D cases, where the model size is a cube and a square respectively with edges of N grid
523 points. The speed-up varies significantly with the model size. The highest speedups are attained with the GPU, which ranges
524 between 50 to more than 80 in 3D and between 30 and 75 in 2D. Significant speedups are also obtained with CPUs, as high as 35
525 times faster. This is higher than the number of cores (20) available. We make the hypothesis that this is caused by a better cache
526 usage of the OpenCL implementation, i.e. usage of local memory increases significantly the cache hits during computation compared
527 to the C implementation of Bohlen (2002). The 2D implementation seems less impacted by this phenomenon and speedups are in a
528 more normal range, between 11 and 25. We also noted that the time stepping computation can be very slow in the first several
529 hundred time steps for the C implementation. This is the source of the strong variations in speedups observed in Figure 7. Finally,
530 the Xeon Phi speedups are disappointing compared to their theoretical computing capacity. However, SeisCL has been optimized for
531 GPUs, not for the Xeon Phi. Even if we have not tested it, it is possible that with small modifications of the code, improved
532 performance could be attained. This shows, however, the limits of device portability with OpenCL: code optimization is paramount to
533 achieve high performances and this optimization can be quite different for different devices.



534

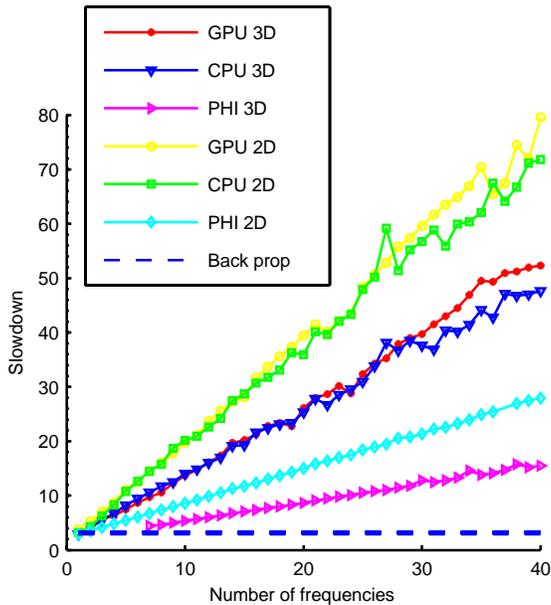
535 **Figure 7** Speedup of SeisCL over a single threaded CPU implementation for different model sizes in 3D (top) and 2D (bottom), for different processor
 536 types.

537

538 4.2.2- Performance of the gradient calculation

539 The next test aims at assessing the performance of the two different gradient schemes. For this experiment, the baseline is the time
 540 required to perform one forward modeling run, without the gradient calculations. The computing times are measured for the
 541 backpropagation scheme and the frequency scheme, for model sizes of 100x100x100 and 1000x1000 grid nodes in 3D and 2D
 542 respectively. The results are shown in Figure 8. For the frequency scheme, the computing time increases linearly with the number of
 543 frequencies. The cost rises faster in 3D than in 2D, which can be explained by the higher number of variables needed to be
 544 transformed in 3D. Surprisingly, the computation time for the Xeon PHI seems to increase much slower than for the CPU or the GPU.
 545 It is to be noticed that for testing purposes, the discrete Fourier transform was computed at every time step. However, significant

546 savings could be achieved if it was computed near the Nyquist frequency. Nevertheless, this test shows that the cost of computing
 547 the discrete Fourier transform during time stepping is not trivial but remains tractable. Finally, the backpropagation scheme has a
 548 cost that is roughly 3 times the cost of a single forward modeling for all devices. Hence, in the elastic case, the backpropagation
 549 scheme outperforms the frequency scheme no matter the number of frequencies. It also has the added benefit of containing all
 550 frequencies.



551

552 **Figure 8** Ratio of the computing time between the forward modeling and the adjoint modeling in the frequency scheme for an increasing number of
 553 frequencies. The dashed line denotes the back-propagation scheme for all devices.

554

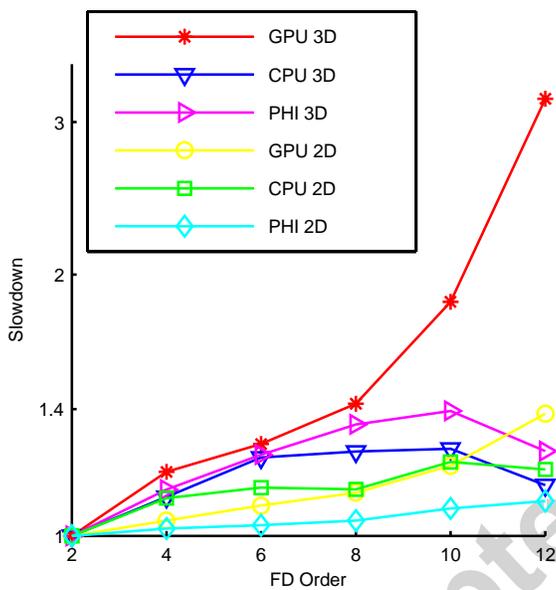
555

556 4.2.3- Measure of the cost of using higher order finite-difference stencils on different devices

557 The baseline for this test is the computation time of the 2nd order stencil for each device. The slowdown is used here as a measure,
 558 i.e. the inverse of the speedup. The same spatial and temporal step lengths were used for each order. As can be seen in Figure 9,
 559 for all three types of device, the slowdown is quite low and does not exceed 1.5 for the highest order of 12 considered here, except
 560 for the GPU in 3D where it exceeds 3 for an order of 12. Note that up to the 8th order, the GPU performance is comparable to the
 561 other device types. The higher cost for the GPU in 3D for orders 10 and 12 is caused by the limited amount of local memory. Indeed,
 562 for those orders, the amount of local memory required to compute the derivative of a single variable exceeds the device capacity. In
 563 those circumstances, SeisCL turns off the usage of local memory and uses global memory directly. The abrupt slowdown is manifest
 564 of the importance of using local memory. The reason why higher order stencils do not affect significantly the computing time of
 565 SeisCL is that it is bandwidth limited: access to the memory takes more time than the actual computations. As memory access is

566 locally shared, the higher number of reads required for higher finite-difference order does not increase significantly. The impact on
 567 computation at each time step is thus marginal. In most cases, the advantages of using higher orders outweigh the computational
 568 costs, because it allows reducing the grid size. For example, using a 6th order over a 2nd order stencil allows reducing the grid point
 569 per wavelength from around 22 to 3, i.e. it reduces the number of grid elements by a factor of 400 in 3D. However, in some
 570 situations, for instance in the presence of a free surface, topography or strong material discontinuities, higher order stencils introduce
 571 inaccuracies (Bohlen and Saenger, 2006, Kristek, et al., 2002, Robertsson, 1996). Hence, the choice of the order should be
 572 evaluated on a case-by-case basis.

573



574

575 **Figure 9 Slowdown of the computation using higher finite-difference order compared to the 2nd order for different devices.**

576

577 4.2.4- Tests on heterogeneous clusters

578 To evaluate the scalability of our code over large clusters, a strong scaling test was performed. Here, strong scaling refers to the
 579 variation of the computational time for a model of fixed sized for an increasing number of processors. The following results were
 580 obtained for a grid size of 96x96x9000 elements and an increasing number of devices for the domain decomposition. This test was
 581 performed on two different clusters: *Helios* of Laval University, Canada and *Guillimin* from McGill University, Canada. The *Helios*
 582 nodes contain 8 NVidia K80 GPUs (16 devices). This cluster was used to test strong scaling for GPUs on a single node of a cluster,
 583 which does not involve MPI. Two types of nodes were used on *Guillimin*: nodes containing two Intel Xeon X5650 with 6 cores each at

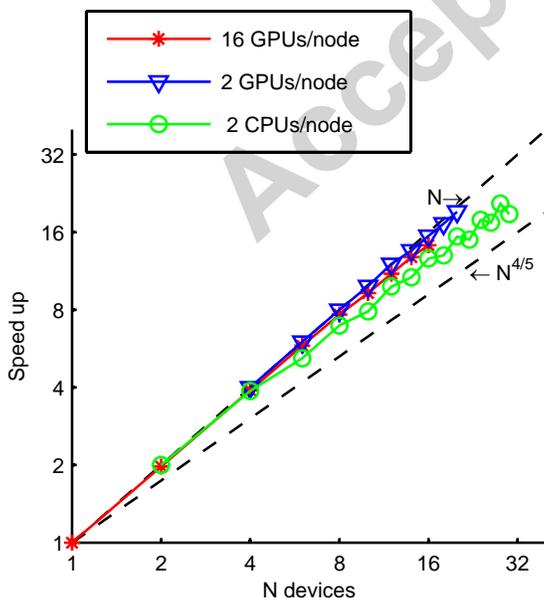
584 2.66 GHz and 12 MB of cache and nodes containing 2 NVidia K20 GPUs in addition to the same two Xeon CPUs. This cluster was
 585 used to test strong scaling across several nodes, which requires MPI communication.

586

587 Results are shown in Figure 10. The best scaling behavior is shown by the nodes on *Guillimin* with two GPUs, which is very nearly
 588 linear over the tested number of devices (blue triangles on Figure 10). Surprisingly, the scaling is slightly worse for many devices
 589 located on the same node (*Helios* nodes, red stars in Figure 10). We interpret this result as being caused by the increasing burden
 590 on the processor when a higher number of GPUs must be scheduled on the same nodes: at some point, the CPU becomes too slow
 591 to keep all GPUs busy. Compared to *Guillimin* nodes using CPUs, *Guillimin* nodes using GPUs also scale better. Still, the CPU
 592 scaling remains quite favorable and is higher than $N^{4/5}$. Those results are better than the results reported by Okamoto (2011), Rubio,
 593 et al. (2014), Weiss and Shragge (2013). We explain this favorable behavior by the separate computation of grid elements inside and
 594 outside of the communication zone in our code.

595 The strong scaling tests show that for large models that fit only on multiple nodes and devices, SeisCL can efficiently parallelize the
 596 computation domains with a minimal performance cost. Still, parallelization over shots should be favored when models fit in the
 597 memory of a single device because no fast interconnects are needed in this situation, and because SeisCL is somewhat more
 598 efficient when memory usage attains a certain level, as shown in Figure 7. In short, having both types of parallelization allows a
 599 greater flexibility over the type of cluster that can be used with SeisCL.

600



601

602 Figure 10 Strong scaling tests for a grid size of 96x96x1000. Red corresponds to results from *Helios*. Green and blue to *Guillimin*.

603

604 **5. Conclusion**

605 In this article, we presented a program called SeisCL for viscoelastic FWI on heterogeneous clusters. The algorithm solves the
606 viscoelastic wave equation by the Finite-Difference Time-Domain approach and uses the adjoint state method to output the gradient
607 of the misfit function. Two approaches were implemented for the gradient computation by the adjoint method: the backpropagation
608 approach and the frequency approach. The backpropagation approach was shown to be the most efficient in the elastic case, having
609 roughly the cost of 3 forward computations. It is, however, not applicable when viscoelasticity is introduced. The frequency approach
610 has an acceptable cost when a small number of frequencies is selected, but becomes quite prohibitive when all frequencies are
611 needed. Future work should focus on the implementation of the optimal checkpointing strategy, which is applicable to both elastic
612 and viscoelastic FWI and strikes a balance between computational costs and memory usage.

613 It was shown that using OpenCL speeds up the computations compared to a single-threaded implementation and allows the usage
614 of different processor architectures. To highlight the code portability, three types of processors were tested: Intel CPUs, Nvidia
615 GPUs and Intel Xeon PHI. The best performances were achieved with the GPUs: a speedup of nearly two orders of magnitude over
616 the single-threaded code was attained. On the other hand, code optimization was shown to be suboptimal on the Xeon PHI, which
617 shows that some efforts must still be spent on device-specific optimization. For the GPU, memory usage was the main area of code
618 optimization. In particular, the use of OpenCL local memory is paramount and coalesced access to global memory must be
619 embedded in the algorithm.

620 When using domain decomposition across devices and nodes of a cluster, overlapping communications and computations allowed
621 hiding the cost of memory transfers. Domain decomposition parallelization was shown to be nearly linear on clusters with fast
622 interconnects using different kinds of processors. Hence, SeisCL can be used to compute the misfit gradient efficiently for large 3D
623 models on a cluster. Furthermore, the task-parallel scheme of distributing shots allows flexibility when the speed of interconnects
624 between the nodes limits the computational gain. Together, both parallelization schemes allow a more efficient usage of large cluster
625 resources.

626 In summary, the very good performance of SeisCL on heterogeneous clusters containing different processor architectures,
627 particularly GPUs, is very promising to speed up full waveform inversion. Presently, the most efficient devices for SeisCL are GPUs,
628 but this can change in the future. The open nature and the flexibility of OpenCL will most probably allow SeisCL to use new hardware
629 developments. SeisCL is distributed with an open license over Github.

630

631

Appendix A

632

This section lists the misfit gradient coefficients. First, some constants are defined:

$$633 \quad \alpha = \sum_{l=1}^L \frac{\omega_0^2 \tau_{\sigma l}^2}{1 + \omega_0^2 \tau_{\sigma l}^2}, \quad (\text{A-1})$$

$$634 \quad b_1 = \left(N_d M \frac{(1 + \tau_p)}{(1 + \alpha \tau_p)} - 2(N_d - 1) \mu \frac{(1 + \tau_s)}{(1 + \alpha \tau_s)} \right)^{-2}, \quad (\text{A-2})$$

$$635 \quad b_2 = \left(N_d M \frac{\tau_p}{(1 + \alpha \tau_p)} - 2(N_d - 1) \mu \frac{\tau_s}{(1 + \alpha \tau_s)} \right)^{-2}. \quad (\text{A-3})$$

636

637

The misfit gradient coefficients are given by:

$$638 \quad c_M^1 = \frac{(1 + \tau_p)}{(1 + \alpha \tau_p)} b_1 \quad (\text{A-4})$$

$$639 \quad c_M^2 = \frac{\tau_p}{(1 + \alpha \tau_p)} b_2, \quad (\text{A-5})$$

$$640 \quad c_\mu^1 = \frac{1}{\mu^2} \frac{(1 + \alpha \tau_s)}{(1 + \tau_s)}, \quad (\text{A-6})$$

$$641 \quad c_\mu^2 = \frac{(N_d + 1)}{3} \frac{(1 + \tau_s)}{(1 + \alpha \tau_s)} b_1 \quad (\text{A-7})$$

$$642 \quad c_\mu^3 = \frac{1}{2N_d} \frac{1}{\mu^2} \frac{(1 + \alpha \tau_s)}{(1 + \tau_s)}, \quad (\text{A-8})$$

$$643 \quad c_\mu^4 = \frac{1}{\mu^2} \frac{(1 + \alpha \tau_s)}{\tau_s}, \quad (\text{A-9})$$

$$644 \quad c_\mu^5 = \frac{(N_d + 1)}{3} \frac{\tau_s}{(1 + \alpha \tau_s)} b_2, \quad (\text{A-10})$$

$$645 \quad c_\mu^6 = \frac{1}{2N_d} \frac{1}{\mu^2} \frac{(1 + \alpha \tau_s)}{\tau_s}, \quad (\text{A-11})$$

$$646 \quad c_{\tau_p}^1 = (1 - \alpha) \frac{M}{(1 + \alpha \tau_p)^2} b_1, \quad (\text{A-12})$$

$$647 \quad c_{\tau_p}^2 = \frac{M}{(1 + \alpha \tau_p)^2} b_2, \quad (\text{A-13})$$

$$648 \quad c_{\tau_s}^1 = \frac{1}{\mu} \frac{(1 - \alpha)}{(1 + \tau_s)^2}, \quad (\text{A-14})$$

$$649 \quad c_{\tau_s}^2 = \frac{(N_d+1)}{3} (1 - \alpha) \frac{\mu}{(1+\alpha\tau_s)^2} b_1 \quad (\text{A-15})$$

$$650 \quad c_{\tau_s}^3 = \frac{(1-\alpha)}{2N_d\mu(1+\tau_s)^2}, \quad (\text{A-16})$$

$$651 \quad c_{\tau_s}^4 = \frac{1}{\mu\tau_s^2}, \quad (\text{A-17})$$

$$652 \quad c_{\tau_s}^5 = \frac{(N_d+1)}{3} \frac{\mu}{(1+\alpha\tau_s)^2} b_2, \quad (\text{A-18})$$

$$653 \quad c_{\tau_s}^6 = \frac{1}{2N_d\mu\tau_s^2}, \quad (\text{A-19})$$

654 Acknowledgements

655 This work was funded by a Vanier Canada Graduate Scholarship and supported by the Canada Chair in geological and geophysical
656 data assimilation for stochastic geological modeling.

657

658

659 References

- 660 Abdelkhalik, R., Calandra, H., Coulaud, O., Roman, J. and Latu, G., 2009, Fast seismic modeling and Reverse Time Migration on a
661 GPU cluster, p. 36-43, 10.1109/hpcsim.2009.5192786.
- 662 Abdullah, D. and Al-Hafidh, M. M., 2013, Developing Parallel Application on Multi-core Mobile Phone: Editorial Preface, v. 4, no. 11.
- 663 Anderson, J. E., Tan, L. and Wang, D., 2012, Time-reversal checkpointing methods for RTM and FWI: Geophysics, v. 77, no. 4, p.
664 S93-S103, 10.1190/geo2011-0114.1.
- 665 Askan, A., Akcelik, V., Bielak, J. and Ghattas, O., 2007, Full Waveform Inversion for Seismic Velocity and Anelastic Losses in
666 Heterogeneous Structures: Bulletin of the Seismological Society of America, v. 97, no. 6, p. 1990-2008, 10.1785/0120070079.
- 667 Bai, J., Yingst, D., Bloor, R. and Leveille, J., 2014, Viscoacoustic waveform inversion of velocity structures in the time domain:
668 Geophysics, v. 79, no. 3, p. R103-R119, 10.1190/geo2013-0030.1.
- 669 Blanch, J. O., Robertsson, J. O. A. and Symes, W. W., 1995, Modeling of a constantQ: Methodology and algorithm for an efficient
670 and optimally inexpensive viscoelastic technique: Geophysics, v. 60, no. 1, p. 176-184, 10.1190/1.1443744.
- 671 Boehm, C., Hanzich, M., de la Puente, J. and Fichtner, A., 2016, Wavefield compression for adjoint methods in full-waveform
672 inversion: Geophysics, v. 81, no. 6, p. R385-R397, 10.1190/geo2015-0653.1.
- 673 Bohlen, T., 2002, Parallel 3-D viscoelastic finite difference seismic modelling: Computers & Geosciences, v. 28, no. 8, p. 887-899,
674 10.1016/s0098-3004(02)00006-7.
- 675 Bohlen, T. and Saenger, E. H., 2006, Accuracy of heterogeneous staggered-grid finite-difference modeling of Rayleigh waves:
676 Geophysics, v. 71, no. 4, p. T109-T115, 10.1190/1.2213051.
- 677 Carcione, J. M., Kosloff, D. and Kosloff, R., 1988, Viscoacoustic wave propagation simulation in the earth: Geophysics, v. 53, no. 6,
678 p. 769-777, 10.1190/1.1442512.
- 679 Cerjan, C., Kosloff, D., Kosloff, R. and Reshef, M., 1985, A nonreflecting boundary condition for discrete acoustic and elastic wave
680 equations: Geophysics, v. 50, no. 4, p. 705-708, 10.1190/1.1441945.
- 681 Clapp, R. G., 2008, Reverse time migration: saving the boundaries, Technical Report SEP-136, Stanford Exploration Project, p. 137.
- 682 Dongarra, J. J., Meuer, H. W. and Strohmaier, E., 2015, Top500 supercomputer sites, <http://www.top500.org/> (Accessed on
683 01/12/2015)

- 684 Du, P., Weber, R., Luszczek, P., Tomov, S., Peterson, G. and Dongarra, J., 2012, From CUDA to OpenCL: Towards a performance-
685 portable solution for multi-platform GPU programming: *Parallel Computing*, v. 38, no. 8, p. 391-407, 10.1016/j.parco.2011.10.002.
- 686 Dussaud, E., Symes, W. W., Williamson, P., Lemaistre, L., Singer, P., Denel, B. and Cherrett, A., 2008, Computational strategies for
687 reverse- time migration, p. 2267-2271, 10.1190/1.3059336.
- 688 Fabien-Ouellet, G., Gloaguen, E. and Giroux, B., The Adjoint State Method for the Viscoelastic Wave Equation in the Velocity-stress
689 Formulation, in *Proceedings 78th EAGE Conference and Exhibition 2016*, May 2016, 10.3997/2214-4609.201600826.
- 690 Fichtner, A., 2011, Full Seismic Waveform Modelling and Inversion.
- 691 Fichtner, A., Bunge, H. P. and Igel, H., 2006, The adjoint method in seismology - I. Theory: *Physics of the Earth and Planetary*
692 *Interiors*, v. 157, no. 1-2, p. 86-104, 10.1016/j.pepi.2006.03.016.
- 693 Furse, C. M., 2000, Faster than Fourier: ultra-efficient time-to-frequency-domain conversions for FDTD simulations: *IEEE Antennas*
694 *and Propagation Magazine*, v. 42, no. 6, p. 24-34, 10.1109/74.894179.
- 695 Gosselin-Cliche, B. and Giroux, B., 2014, 3D frequency-domain finite-difference viscoelastic-wave modeling using weighted average
696 27-point operators with optimal coefficients: *Geophysics*, v. 79, no. 3, p. T169-T188, 10.1190/geo2013-0368.1.
- 697 Griewank, A., 1992, Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation:
698 *Optimization Methods and Software*, v. 1, no. 1, p. 35-54, 10.1080/10556789208805505.
- 699 Griewank, A. and Walther, A., 2000, Algorithm 799: revolve: an implementation of checkpointing for the reverse or adjoint mode of
700 computational differentiation: *ACM Transactions on Mathematical Software*, v. 26, no. 1, p. 19-45, 10.1145/347837.347846.
- 701 Holberg, O., 1987, Computational Aspects of the Choice of Operator and Sampling Interval for Numerical Differentiation in Large-
702 Scale Simulation of Wave Phenomena*: *Geophysical Prospecting*, v. 35, no. 6, p. 629-655, 10.1111/j.1365-2478.1987.tb00841.x.
- 703 Howes, L. and Munshi, A., 2014, The OpenCL Specification: *Khronos OpenCL*.
- 704 Iturrarán-Viveros, U. and Molero, M., 2013, Simulation of sonic waves along a borehole in a heterogeneous formation: *Accelerating*
705 *2.5-D finite differences using [Py]OpenCL: Computers & Geosciences*, v. 56, p. 161-169, 10.1016/j.cageo.2013.03.014.
- 706 Kamei, R. and Pratt, R. G., 2013, Inversion strategies for visco-acoustic waveform inversion: *Geophysical Journal International*, v.
707 194, no. 2, p. 859-884, 10.1093/gji/ggt109.
- 708 Kloc, M. and Danek, T., 2013, The Multi GPU Accelerated Waveform Inversion in Distributed OpenCL Environment, *Emerging*
709 *Trends in Computing, Informatics, Systems Sciences, and Engineering*, Springer, p. 237-244.
- 710 Komatitsch, D., Erlebacher, G., Góddeke, D. and Michéa, D., 2010, High-order finite-element seismic wave propagation modeling
711 with MPI on a large GPU cluster: *Journal of Computational Physics*, v. 229, no. 20, p. 7692-7714, 10.1016/j.jcp.2010.06.024.
- 712 Komatitsch, D. and Martin, R., 2007, An unsplit convolutional perfectly matched layer improved at grazing incidence for the seismic
713 wave equation: *Geophysics*, v. 72, no. 5, p. SM155-SM167, 10.1190/1.2757586.
- 714 Komatitsch, D., Xie, Z., Bozdağ, E., Sales de Andrade, E., Peter, D., Liu, Q. and Tromp, J., 2016, Anelastic sensitivity kernels with
715 parsimonious storage for adjoint tomography and full waveform inversion: *Geophysical Journal International*, v. 206, no. 3, p. 1467-
716 1478, 10.1093/gji/ggw224.
- 717 Kristek, J., Moczo, P. and Archuleta, R. J., 2002, Efficient Methods to Simulate Planar Free Surface in the 3D 4th-Order Staggered-
718 Grid Finite-Difference Schemes: *Studia Geophysica et Geodaetica*, v. 46, no. 2, p. 355-381, 10.1023/a:1019866422821.
- 719 Lailly, P., The seismic inverse problem as a sequence of before stack migrations, in *Proceedings Conference on inverse scattering:*
720 *theory and application*, 1983, Society for Industrial and Applied Mathematics, Philadelphia, PA, p. 206-220.
- 721 Levander, A. R., 1988, Fourth- order finite- difference P-SV seismograms: *Geophysics*, v. 53, no. 11, p. 1425-1436,
722 10.1190/1.1442422.
- 723 Liu, H.-P., Anderson, D. L. and Kanamori, H., 1976, Velocity dispersion due to anelasticity; implications for seismology and mantle
724 composition: *Geophysical Journal International*, v. 47, no. 1, p. 41-58.
- 725 Malinowski, M., Operto, S. and Ribodetti, A., 2011, High-resolution seismic attenuation imaging from wide-aperture onshore data by
726 visco-acoustic frequency-domain full-waveform inversion: *Geophysical Journal International*, v. 186, no. 3, p. 1179-1204,
727 10.1111/j.1365-246X.2011.05098.x.
- 728 Mattson, T. G., Sanders, B. A. and Massingill, B. L., 2004, *Patterns for parallel programming*, Pearson Education.
- 729 Michéa, D. and Komatitsch, D., 2010, Accelerating a three-dimensional finite-difference wave propagation code using GPU graphics
730 cards: *Geophysical Journal International*, p. no-no, 10.1111/j.1365-246X.2010.04616.x.
- 731 Micikevicius, P., 2009, 3D finite difference computation on GPUs using CUDA, p. 79-84, 10.1145/1513895.1513905.
- 732 Molero, M. and Iturrarán-Viveros, U., 2013, Accelerating numerical modeling of wave propagation through 2-D anisotropic materials
733 using OpenCL: *Ultrasonics*, v. 53, no. 3, p. 815-822.
- 734 Mu, D., Chen, P. and Wang, L., 2013, Accelerating the discontinuous Galerkin method for seismic wave propagation simulations
735 using the graphic processing unit (GPU)—single-GPU implementation: *Computers & Geosciences*, v. 51, p. 282-292.
- 736 Nguyen, B. D. and McMechan, G. A., 2015, Five ways to avoid storing source wavefield snapshots in 2D elastic prestack reverse
737 time migration: *Geophysics*, v. 80, no. 1, p. S1-S18, 10.1190/geo2014-0014.1.
- 738 Nickolls, J. and Dally, W. J., 2010, The GPU Computing Era: *IEEE Micro*, v. 30, no. 2, p. 56-69, 10.1109/mm.2010.41.

- 739 Nihei, K. T. and Li, X., 2007, Frequency response modelling of seismic waves using finite difference time domain with phase
 740 sensitive detection (TD-PSD): *Geophysical Journal International*, v. 169, no. 3, p. 1069-1078, 10.1111/j.1365-246X.2006.03262.x.
- 741 Nvidia, C., 2007, Compute unified device architecture programming guide.
- 742 O'connell, R. and Budiansky, B., 1978, Measures of dissipation in viscoelastic media: *Geophysical Research Letters*, v. 5, no. 1, p. 5-
 743 8.
- 744 Okamoto, 2011, Accelerating large-scale simulation of seismic wave propagation by multi-GPUs and three-dimensional domain
 745 decomposition: *Earth, Planets and Space*, v. 62, no. 12, p. 939-942, 10.5047/eps.2010.11.009.
- 746 Owens, J. D., Houston, M., Luebke, D., Green, S., Stone, J. E. and Phillips, J. C., 2008, GPU Computing: *Proceedings of the IEEE*,
 747 v. 96, no. 5, p. 879-899, 10.1109/jproc.2008.917757.
- 748 Pilant, W. L., 2012, *Elastic waves in the earth*, Volume 11, Elsevier.
- 749 Plessix, R. E., 2006, A review of the adjoint-state method for computing the gradient of a functional with geophysical applications:
 750 *Geophysical Journal International*, v. 167, no. 2, p. 495-503, 10.1111/j.1365-246X.2006.02978.x.
- 751 Pratt, R. G. and Worthington, M., 1990, Inverse theory applied to multi-source cross-hole tomography. Part 1: acoustic wave-
 752 equation method: *Geophysical Prospecting*, v. 38, no. 3, p. 287-310, 10.1111/j.1365-2478.1990.tb01846.x
- 753 Robertsson, J. O. A., 1996, A numerical free- surface condition for elastic/viscoelastic finite- difference modeling in the presence of
 754 topography: *Geophysics*, v. 61, no. 6, p. 1921-1934, 10.1190/1.1444107.
- 755 Robertsson, J. O. A., Blanch, J. O. and Symes, W. W., 1994, Viscoelastic finite- difference modeling: *Geophysics*, v. 59, no. 9, p.
 756 1444-1456, 10.1190/1.1443701.
- 757 Roden, J. A. and Gedney, S. D., 2000, Convolution PML (CPML): An efficient FDTD implementation of the CFS-PML for arbitrary
 758 media: *Microwave and Optical Technology Letters*, v. 27, no. 5, p. 334-339, 10.1002/1098-2760(20001205)27:5<334::aid-
 759 mop14>3.0.co;2-a.
- 760 Rubio, F., Hanzich, M., Farrés, A., de la Puente, J. and María Cela, J., 2014, Finite-difference staggered grids in GPUs for
 761 anisotropic elastic wave propagation simulation: *Computers & Geosciences*, v. 70, p. 181-189, 10.1016/j.cageo.2014.06.003.
- 762 Sirgue, L., Etgen, J. and Albertin, U., 3D frequency domain waveform inversion using time domain finite difference methods, in
 763 *Proceedings 70th EAGE Conference and Exhibition incorporating SPE EUROPEC 2008*, 2008.
- 764 Sirgue, L. and Pratt, R. G., 2004, Efficient waveform inversion and imaging: A strategy for selecting temporal frequencies:
 765 *Geophysics*, v. 69, no. 1, p. 231-248, 10.1190/1.1649391.
- 766 Stone, J. E., Gohara, D. and Shi, G., 2010, OpenCL: A parallel programming standard for heterogeneous computing systems:
 767 *Computing in science & engineering*, v. 12, no. 1-3, p. 66-73.
- 768 Sun, W. and Fu, L.-Y., 2013, Two effective approaches to reduce data storage in reverse time migration: *Computers & Geosciences*,
 769 v. 56, p. 69-75, 10.1016/j.cageo.2013.03.013.
- 770 Symes, W. W., 2007, Reverse time migration with optimal checkpointing: *Geophysics*, v. 72, no. 5, p. SM213-SM221,
 771 10.1190/1.2742686.
- 772 Tarantola, A., 1984, Inversion of seismic reflection data in the acoustic approximation: *Geophysics*, v. 49, no. 8, p. 1259-1266,
 773 10.1190/1.1441754.
- 774 Virieux, J., 1986, P-SVwave propagation in heterogeneous media: Velocity- stress finite- difference method: *Geophysics*, v. 51, no.
 775 4, p. 889-901, 10.1190/1.1442147.
- 776 Virieux, J. and Operto, S., 2009, An overview of full-waveform inversion in exploration geophysics: *Geophysics*, v. 74, no. 6, p.
 777 WCC1-WCC26, 10.1190/1.3238367.
- 778 Vuduc, R., Chandramowlishwaran, A., Choi, J., Guney, M. and Shringarpure, A., On the limits of GPU acceleration, in *Proceedings*
 779 *Proceedings of the 2nd USENIX conference on Hot topics in parallelism*, 2010, USENIX Association, p. 13-13.
- 780 Weiss, R. M. and Shragge, J., 2013, Solving 3D anisotropic elastic wave equations on parallel GPU devices: *Geophysics*, v. 78, no.
 781 2, p. F7-F15, 10.1190/geo2012-0063.1.
- 782 Yang, P., Brossier, R., Metivier, L. and Virieux, J., 2016, Checkpointing-assisted reverse-forward simulation: An optimal
 783 recomputation method for FWI and RTM, p. 1089-1093, 10.1190/segam2016-13685603.1.
- 784 Yang, P., Gao, J. and Wang, B., 2014, RTM using effective boundary saving: A staggered grid GPU implementation: *Computers &*
 785 *Geosciences*, v. 68, p. 64-72, 10.1016/j.cageo.2014.04.004.
- 786 Zhe, F., Feng, Q., Kaufman, A. and Yoakum-Stover, S., 2004, GPU Cluster for High Performance Computing, p. 47-47,
 787 10.1109/sc.2004.26.
- 788
- 789
- 790

791

792

793 **Highlights**

- 794 • An open source software for viscoelastic full waveform inversion is presented.
- 795 • This software is based on OpenCL and can run on CPUs, GPUs and accelerators.
- 796 • On large clusters, MPI is used and a nearly linear scaling is achieved.
- 797 • Using GPUs, we obtain a speed-up of up to 80x over a single threaded CPU code.
- 798

Accepted manuscript