Scalable and Efficient Simulation of Frequency-Domain Wave Propagation

Jacob Badger, Leszek Demkowicz, Andrey Bakulin, and Sergey Fomel, The University of Texas at Austin*

SUMMARY

Developing fast and scalable solvers for frequency-domain wave simulation is a notoriously difficult problem in mathematics and computer science. One challenge is that classical discretization techniques such as finite difference methods yield indefinite discrete systems that preclude use of classical scalable solvers. Frequency-domain simulation has thus been limited to problems with ca. 1 billion degrees of freedom (DOFs). This work summarizes a novel scalable and efficient multigrid solver for high-frequency wave propagation problems; solution of problems with over 150 billion DOFs is demonstrated. The proposed solver is based on the discontinuous Petrov–Galerkin (DPG) finite element method and is applicable—without modification—to a wide class of problems including acoustics, electromagnetics, elasticity, poroelasticity, and coupled multiphysics systems. When combined with adaptive mesh refinement, the solver enables accurate and efficient simulation of complex near-surface structures and other challenging features. Scalability and efficiency of the solver are demonstrated via visco-acoustic simulation on the GO_3D_OBS model (Górszczyk and Operto, 2021) and elastic simulation on the SEAM Arid model (Oristaglio, 2013).

INTRODUCTION

Frequency-domain (FD) simulation is useful in contexts with attenuation and other frequency dependent effects and in contexts with long offsets or high-contrast media, where time-domain methods require many timesteps. In the context of full waveform inversion (FWI), FD simulation enables use of frequency-continuation and mitigates frequency whitening. However, FD simulation requires solving a sparse linear system for each frequency (ω), and systems arising from three-dimensional wave propagation problems are often too large for current solvers. Development of increasingly fast, efficient, and scalable solvers is thus needed to enable FD simulation at scale.

Leading direct sparse solvers can factorize linear systems arising from three-dimensional problems in $O(\omega^6)$ operations and $O(\omega^4)$ memory. Once a system has been factored, it can be applied in $O(\omega^4)$ operations, direct sparse solvers can thus be competitive when the factorization can be amortized over a sufficiently large number of shots (Operto et al., 2023). Hierarchical semi-separable structure, rank-revealing factorizations, and other techniques (Wang et al., 2011; Kostin et al., 2017) can reduce the cost of factorization and application at the expense of incurring algebraic error. Still, this approach has so-far been limited to systems with $O(10^8)$ DOFs.

Iterative solvers are often more efficient for benign (*definite*) problems, enabling solution of systems with up to $O(10^{13})$ DOFs. However, high-frequency wave operators—and often the discrete systems they lead to—are notoriously *indefinite*, developing adequate preconditioners for these problems is challenging (Earnst and Gander, 2012) and has limited iterative solvers for FD simulation to $O(10^9)$ DOFs.

The discontinuous Petrov-Galerkin (DPG) (Demkowicz and Gopalakrishnan, 2010) is a minimum residual finite element method that always produces Hermitian positive-*definite* systems. DPG systems are thus amenable to more classical preconditioning techniques. A DPG multigrid solver (DPG-MG) with $O(\omega^4)$ compute and $O(\omega^3)$ memory complexity was proposed in (Petrides, 2019), but a shared memory implementation limited the solver to $O(10^7)$ DOFs. A scalable and performant version of the DPG-MG solver was recently detailed in the author's Ph.D. thesis (Badger, 2024) and shown to enable solution of problems with $O(10^{12})$ DOFs. Here we provide a brief overview of the scalable DPG-MG solver and potential applications. Additional details can be found in the cited theses and in papers (Petrides and Demkowicz, 2021) and (Badger et al., 2023).

APPROACH

DPG overview

Similar to hybridizable discontinuous Galerkin (HDG) methods, DPG introduces additional interface or *trace* variables defined on the mesh skeleton. DPG systems can then be written in the form of a weighted least squares system:

$$\begin{bmatrix} \mathsf{G} & \mathsf{B} & \hat{\mathsf{B}} \\ \mathsf{B}^* & 0 & 0 \\ \hat{\mathsf{B}}^* & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathsf{s} \\ \mathsf{u} \\ \hat{\mathsf{u}} \end{bmatrix} = \begin{bmatrix} \mathsf{f} \\ 0 \\ 0 \end{bmatrix}$$

where u corresponds to field trial DOFs, $\hat{\mathsf{u}}$ denotes interface trial DOFs, and s denotes the Riesz representation of the residual; the element-wise norm $||\mathsf{G}^{-1}\mathsf{s}||$ can be shown to provide a built-in *a posteriori* error indicator. The system is rarely solved in this form; instead, DPG utilizes discontinuous "broken" test functions that impart block-diagonal structure to Gram matrix G, allowing it to be condensed elementwise. The remaining "bubble" DOFs can further be condensed elementwise, resulting in a global system defined only on interface DOFs.

DPG multigrid solver

The DPG-MG solver leverages a multigrid preconditioned conjugate gradient (PCG) iteration. Multigrid methods are defined by three components: smoothing, prolongation, and coarse-grid correction operators, these are summarized.

Smoothing. A damped overlapping block-Jacobi smoother was adopted with blocks defined by vertex patches on the previous-grid.
Prolongation. A two-stage prolongation operator was leveraged; during the first stage Schur-complements condensed DOFs not supported on the previous mesh skeleton, the second stage then projected the result onto previous-grid traces.
Coarse-grid correction. The coarse-grid correction was defined by via an inner PCG iteration, preconditioned by a single smoothing step on the coarse grid. This necessitated use of the flexible PCG iteration on the outer loop; β was thus defined via the real part of the Polak–Ribièr formula:

$$\beta_k = \frac{\Re\{\mathbf{r}_{k+1}^*(\mathbf{z}_{k+1} - \mathbf{z}_k)\}}{\mathbf{r}_k^*\mathbf{z}_k}.$$

hp-adaptive mesh refinement

The DPG-MG solver supports general unstructured meshes composed of all element shapes and could thus be used to simulate non-trivial topography and other complex features. However, mesh generation is often cumbersome and difficult to automate, in the following results we thus start from simple hexahedral meshes and use *hp*-adaptive mesh refinement to generate a wavespeed-adapted mesh. An example is shown in Fig. 1, where the size (h) and order (p) of mesh elements is varied; h-refinements were used regions with high wavespeed variation to accurately resolve jumps, while p-refinements were used to minimize dispersion error. The wavespeed in Fig. 1 is visualized on the finite element mesh. In simulations the full-resolution model was stored separately, and a (p+3) order quadrature rule was used to form element matrices.

EXAMPLES

The following examples were computed on *Frontera* Cascade Lake nodes (56 CPU cores, 192 GB of RAM) at the Texas Advanced Computing Center (Stanzione et al., 2020). See (Badger, 2024) for additional details on problem setup.

Visco-acoustic simulation

The DPG-MG solver was used to simulate visco-acoustic wave propagation on a 100 km × 100 km × 30 km section of the GO_3D_OBS model at 15 Hz. Simulation was performed on 2048 *Frontera* nodes (114,688 cores) with an *hp*-adapted mesh with over 500 million hexahedral elements, a minimum of 8 points per wavelength, and 157 billion DOFs, a 5.3× reduction compared to a uniform mesh with a similar minimum resolution (Badger, 2024). The adaptive mesh was generated in parallel in 12 s. The simulation originally refactored patches during the solve to reduce memory usage and solved a single shot (to 10^{-4} relative residual) in 492 iterations and 3,891 s. A modified version of the solver (storing smoother patches) simulated 16 shots in 613 iterations and 2,361 s, a cost of 84 node-hours per shot. The solution is depicted in Fig. 2 and Fig. 3.

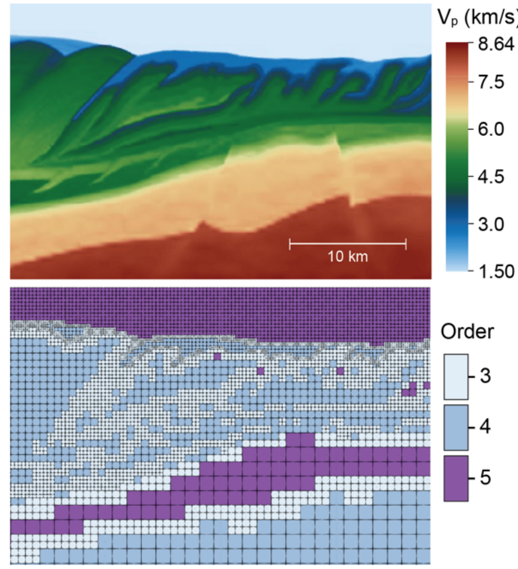

Figure 1: *hp*-adaptive meshing of a section of the GO_3D_OBS model; the size (h) and order (p) of mesh elements (bottom panel) is adapted to the wavespeed (top panel). The pictured mesh was adapted for simulation at 2 Hz with a minimum of 8 PPW.

Visco-elastic simulation

Horizontal transverse isotropic (HTI) visco-elastic simulation on the SEAM Arid model (7 km × 7 km patch) at 25 Hz was performed on 256 *Frontera* nodes (14,336 cores). The *hp*-adapted mesh had 18 million elements, with a minimum of ca. 7 points per shear wavelength and 7 billion DOFs. Adaptivity was particularly advantageous in this case due to the complex near-surface structure of the model, implying a 20× reduction in DOFs compared to a uniform mesh with similar minimum resolution. A batch of 32 shots was solved (to 10^{-4} relative residual) in 189 iterations and 797 s, implying a cost of 1.8 node-hours per shot. The resulting wavefield and a section of the h-adapted mesh are depicted in Fig. 4.

CONCLUSIONS

The DPG-MG solver was shown to enable simulation of visco-acoustic and HTI visco-elastic wave propagation at unprecedented scales. Use of *hp*-adaptivity enabled fast, automated meshing of complex problems, which could prove useful in the context of FWI. *hp*-adaptivity further enabled localization of low near-surface wavespeed and could be used to simulate small-scale heterogeneity. Support for unstructured meshes (in combination with *hp*-adaptivity) is expected to enable efficient simulation of non-trivial

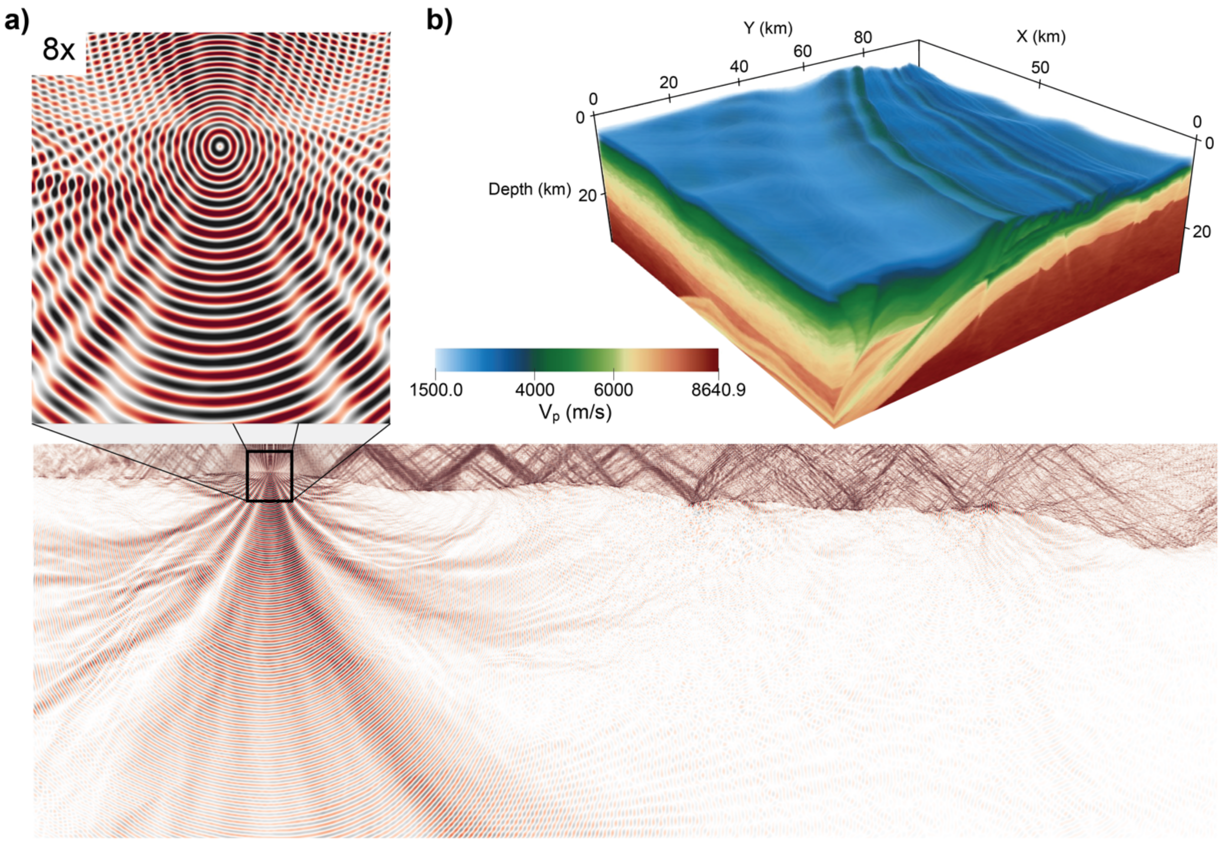

Figure 2: Visco-acoustic simulation of a 15 Hz time-harmonic shot in the GO_3D_OBS model; (a) real part of pressure wavefield at y = 50 km, (b) wavespeed (V_p) on a 100 km × 100 km section of the model. Figure adapted from "Scalable DPG multigrid solver with applications in high-frequency wave propagation," by J. Badger, 2024, Ph.D. Thesis at the University of Texas at Austin.

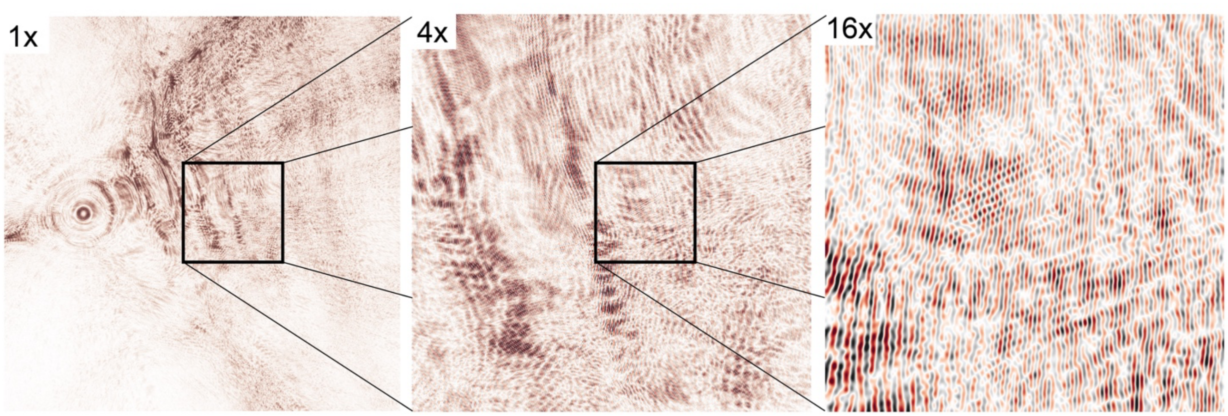

Figure 3: Visco-acoustic simulation of a 15 Hz time-harmonic shot in the GO_3D_OBS model; real part of pressure wavefield at z = 0 km. Figure adapted from "Scalable DPG multigrid solver with applications in high-frequency wave propagation," by J. Badger, 2024, Ph.D. Thesis, The University of Texas at Austin.

topography. Together, these unique capabilities may enable more accurate simulation of seismic speckle (Bakulin, 2023) and contribute to improved subsurface characterization in challenging geological contexts. Finally, The DPG-MG solver relies entirely on dense matrix block operations; a GPU implementation is thus expected to unlock significant further efficiency.

ACKNOWLEDGEMENTS

We thank Andrzej Górszczyk and Saudi Aramco for kindly providing the GO_3D_OBS and SEAM Arid model, respectively.

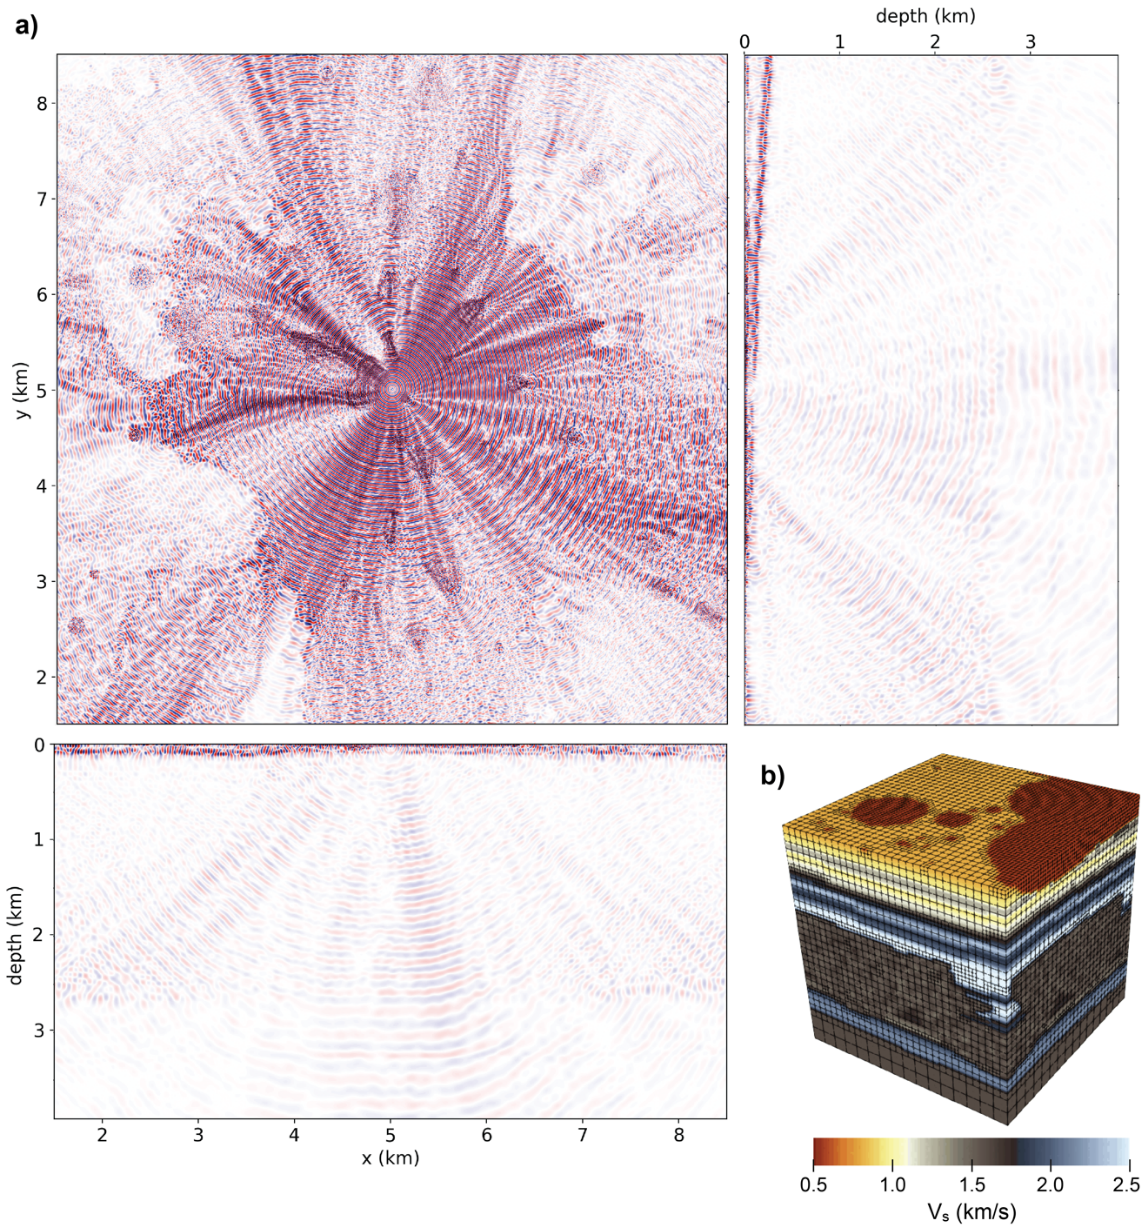

Figure 4: Visco-elastic simulation of a 25 Hz time-harmonic shot on a 7 km × 7 km patch of the SEAM Arid model, centered at (5 km, 5 km); (a) real part of the z-displacement, and (b) *hp*-adapted mesh on a 0.5 km× 0.5 km section of the model. Figure adapted from "Scalable DPG multigrid solver with applications in high-frequency wave propagation," by J. Badger, 2024, Ph.D. Thesis, The University of Texas at Austin.