

Advancing Subsurface Modeling with a Scalable Frequency-Domain Wave Solver

Allocation: “Efficient seismic simulation in challenging geological environments with a scalable frequency-domain solver”

Team:

Sergey Fomel	(UT Austin, Bureau of Economic Geology; PI)
Andrey Bakulin	(UT Austin, Bureau of Economic Geology)
Jacob Badger	(UT Austin, Institute for Fusion Studies)
Lars Koesterke	(TACC, Collaborator)

Frontera User Meeting

August 5-6, 2024


COI Disclosure

The presenter has a financial interest in Cillatory, LLC (defunct) and FrequenSol, LLC.


Motivation

Wave Applications


Medical ultrasound (acoustics)



Seismic simulation (elasticity)



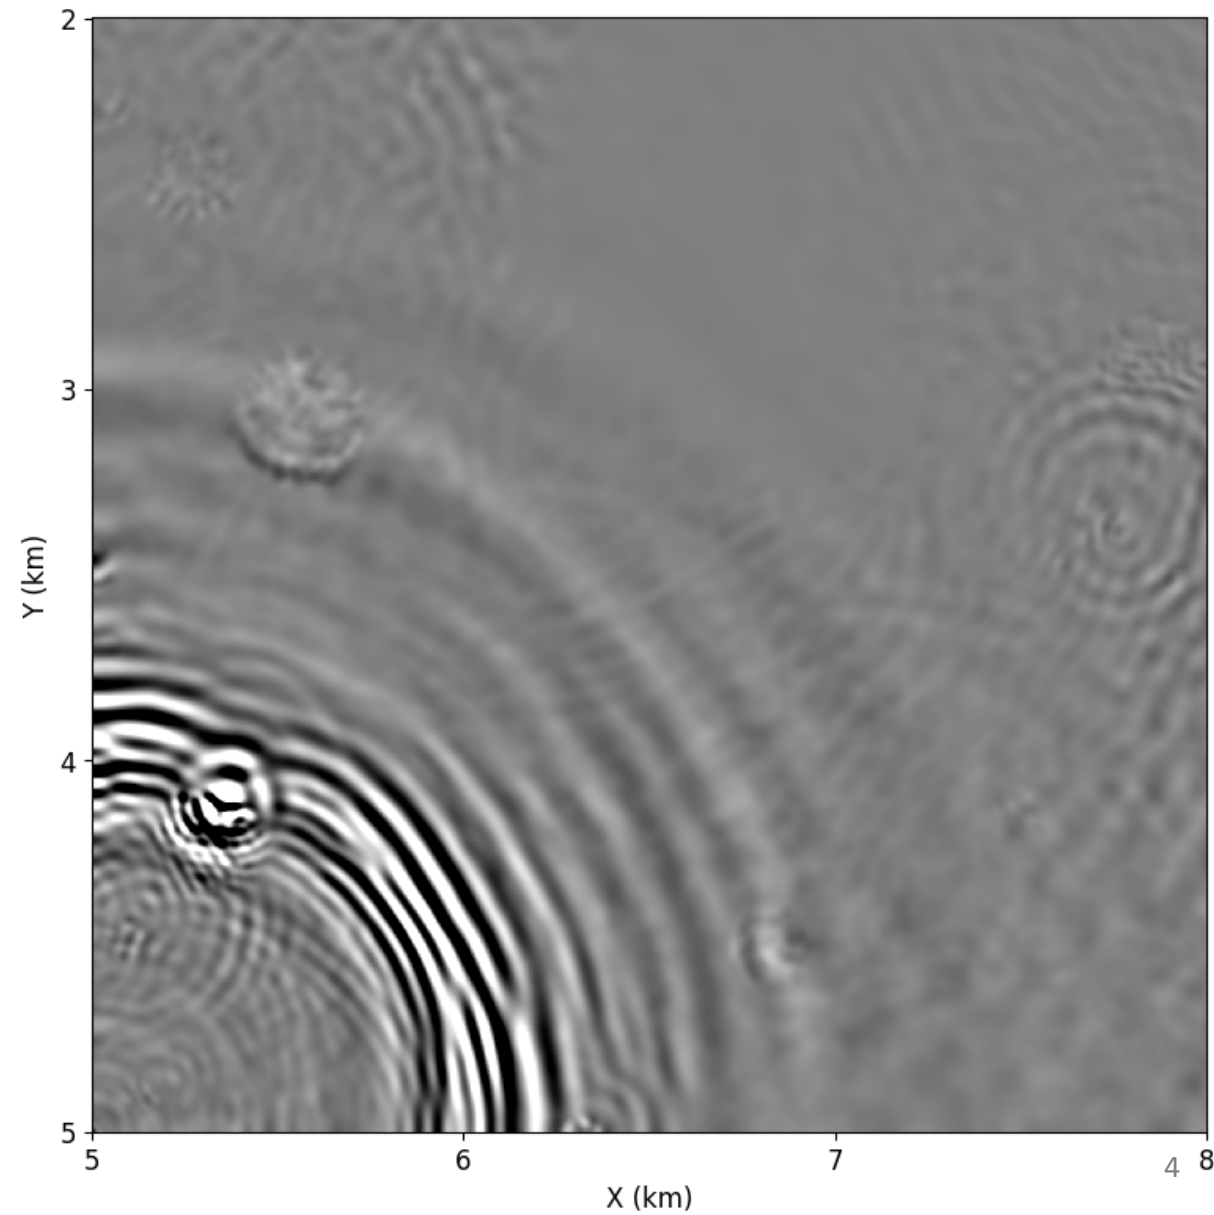
Tokamak RF heating (electromagnetics)



Wave Simulation (Time-Domain)

Time Domain Simulation

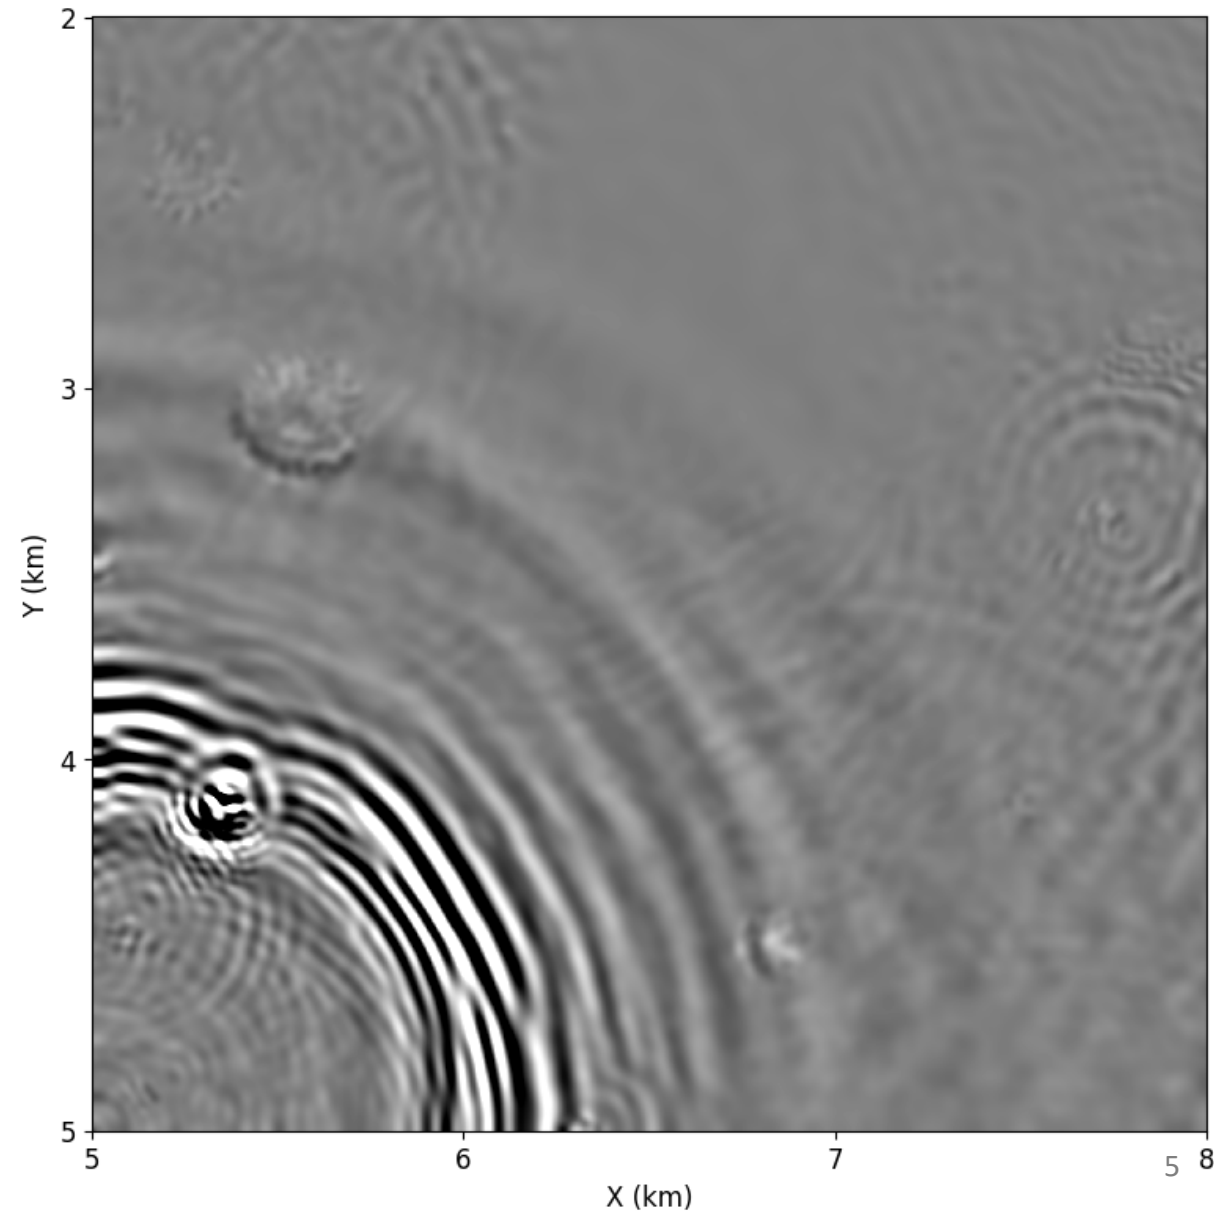
- Starts from known initial state



Wave Simulation (Time-Domain)

Time Domain Simulation


- Starts from known initial state
- “Steps” forward by small time increments



Wave Simulation (Time-Domain)

Time Domain Simulation


- Starts from known initial state
- “Steps” forward by small time increments
- Stepping via:
 - Matrix multiplication (explicit)
 - Elliptic linear solve (implicit)



Wave Simulation (Time-Domain)

Time Domain Simulation

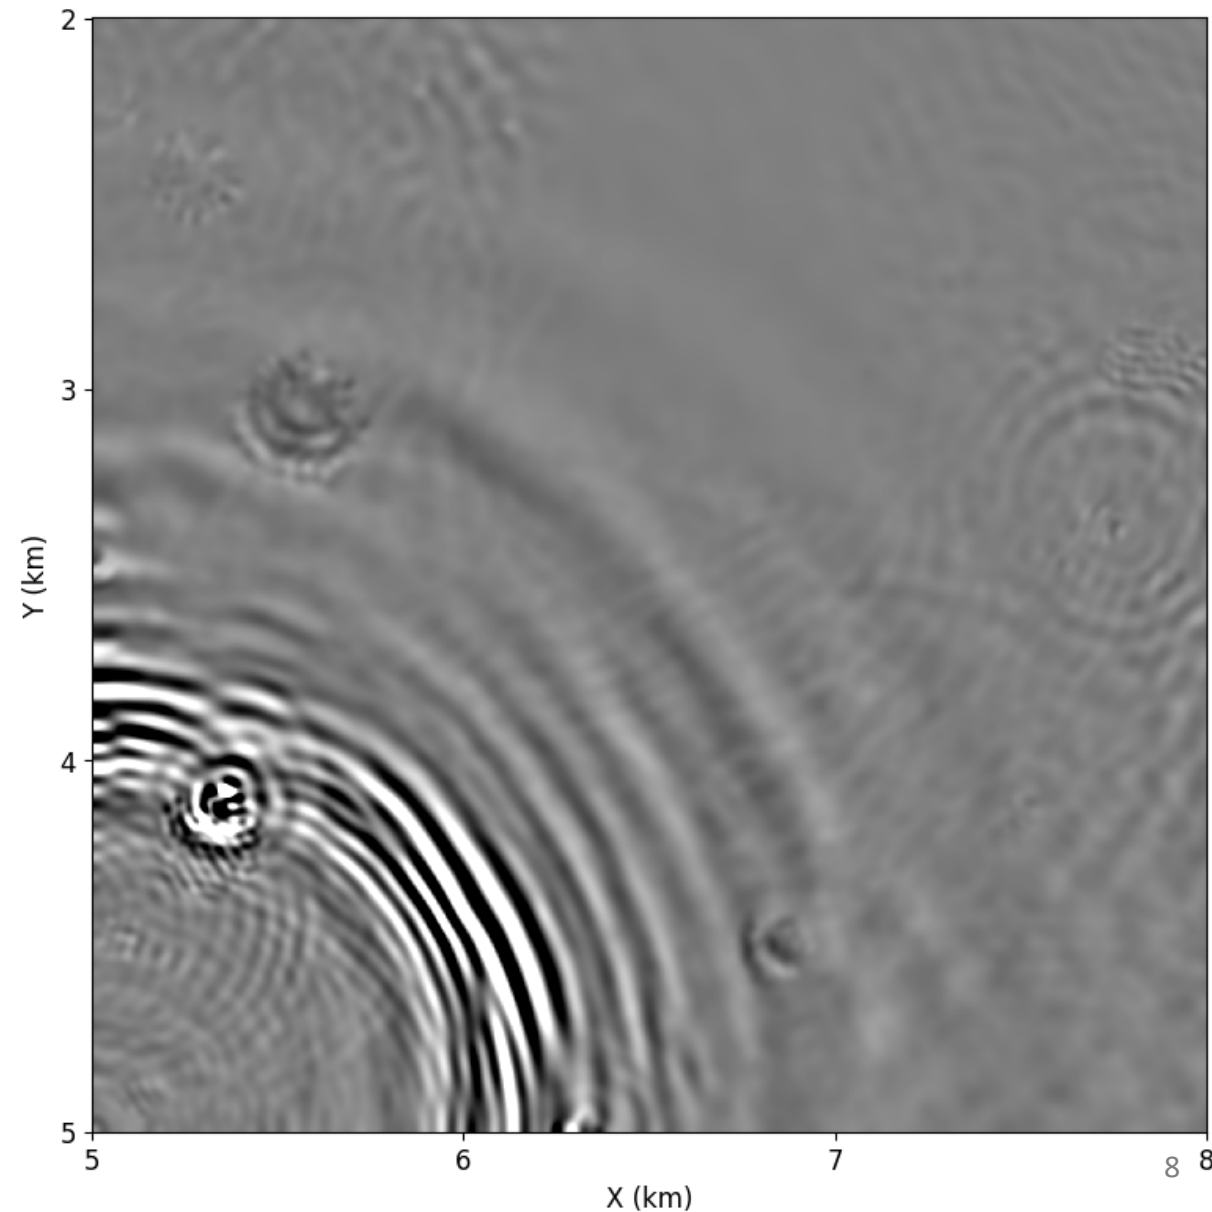
- Starts from known initial state
- “Steps” forward by small time increments
- Stepping via:
 - Matrix multiplication (explicit)
 - Elliptic linear solve (implicit)



Wave Simulation (Time-Domain)

Time Domain Simulation


- Starts from known initial state
- “Steps” forward by small time increments
- Stepping via:
 - Matrix multiplication (explicit)
 - Elliptic linear solve (implicit)



Wave Simulation (Time-Domain)

Time Domain Simulation


- Starts from known initial state
- “Steps” forward by small time increments
- Stepping via:
 - Matrix multiplication (explicit)
 - Elliptic linear solve (implicit)



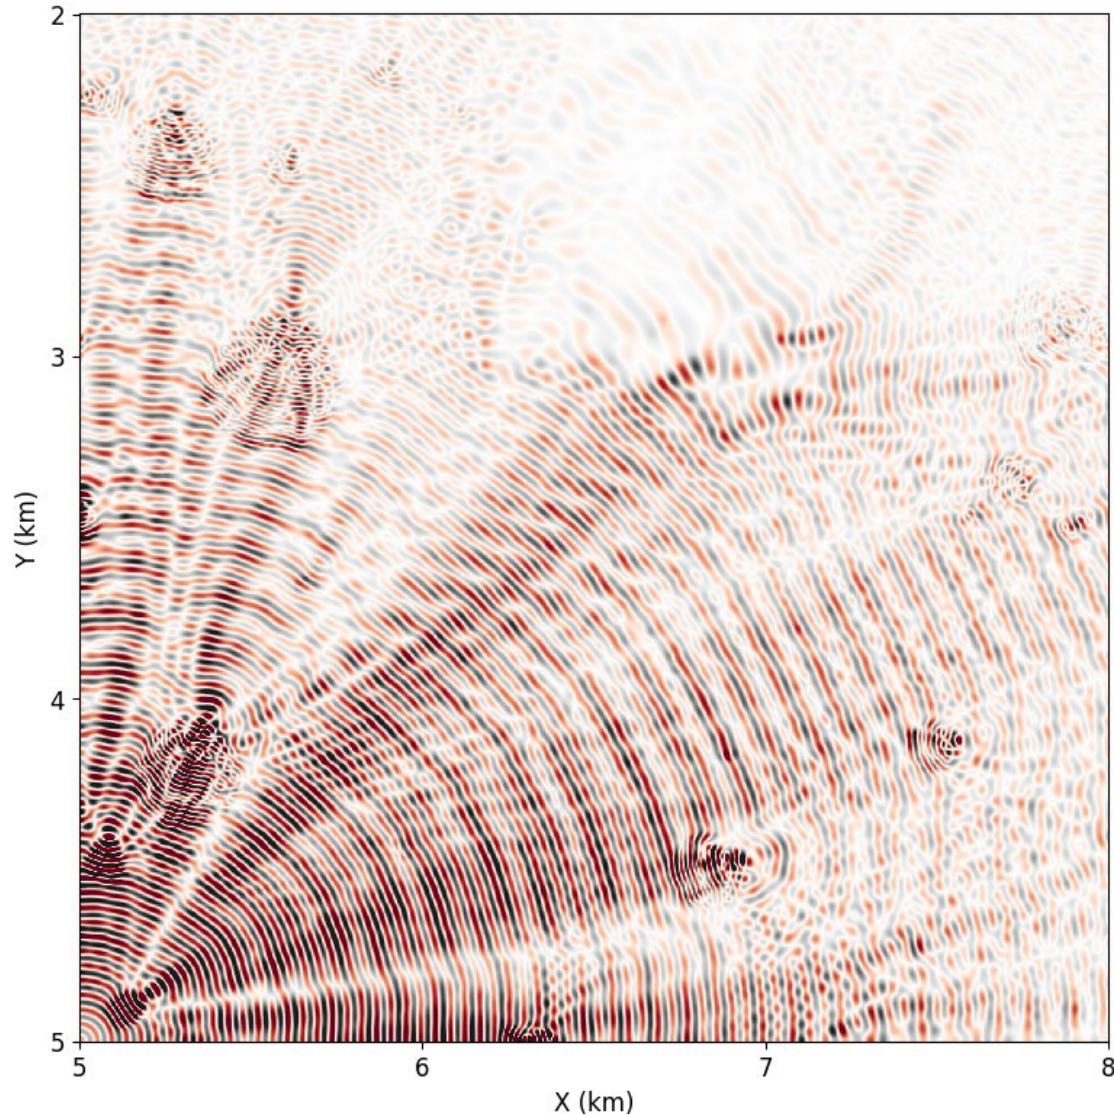
Wave Simulation (Time-Domain)

Time Domain Simulation

- Starts from known initial state
- “Steps” forward by small time increments
- Stepping via:
 - Matrix multiplication (explicit)
 - Elliptic linear solve (implicit)



Wave Simulation (Frequency-Domain)




Frequency Domain Simulation

- Separates problem into independent frequencies (ω)

$$u_\omega(\mathbf{x}, t) = \Re\{\tilde{u}_\omega(\mathbf{x})e^{i\omega t}\} \quad (1)$$

- Each frequency requires an *indefinite* linear solve
- Produces a complex-valued field (\tilde{u}_ω) for each frequency

Wave Simulation (Frequency-Domain)



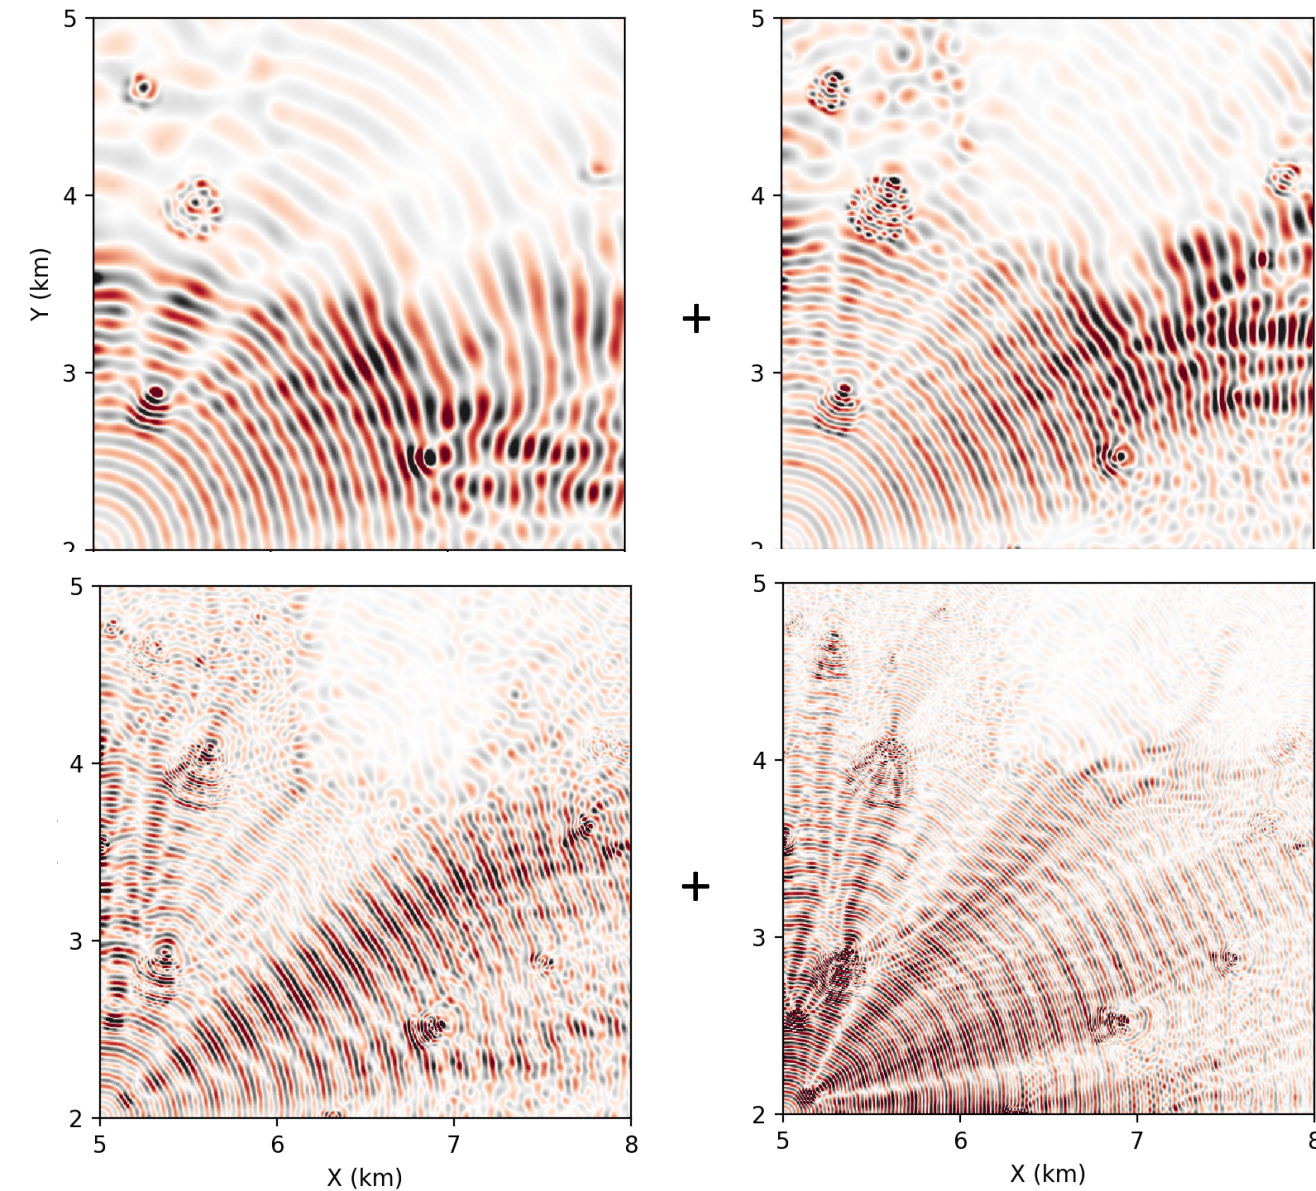
Frequency Domain Simulation

- Separates problem into independent frequencies (ω)

$$u_\omega(\mathbf{x}, t) = \Re\{\tilde{u}_\omega(\mathbf{x})e^{i\omega t}\} \quad (1)$$

- Each frequency requires an *indefinite* linear solve
- Produces a complex-valued field (\tilde{u}_ω) for each frequency
- Time-harmonic field recovered via (1)

Wave Simulation (Frequency-Domain)




Time-domain solution can be reconstructed by summing over frequencies:

$$u(\mathbf{x}, t) = \Re \left\{ \sum_{\omega} c_{\omega} \tilde{u}_{\omega}(\mathbf{x}) e^{i\omega t} \right\}$$

+

+ ... =



Challenges

Leading methods for all approaches have same $O(\omega^4)$ complexity (in 3D)

Time domain (explicit):

- Stability
 - Timestep depends on wavespeed contrast, minimum element size (grid spacing), etc.
 - Often requires many timesteps, especially for high-contrast problems

Time domain (implicit):

- Efficiency
 - Elliptic solve per timestep
 - Often relies on (pre-computed) sparse matrix factorizations

Frequency domain:


- Scalability
 - Direct solvers: factored in $O(\omega^6)$, applied in $O(\omega^4)$
 - Leading scalable iterative solvers achieve $O(\omega^4)$ complexity
 - Previously limited to ~ 1 billion degrees of freedom

Fast Frequency-Domain Solver

Fast Frequency-Domain Solver¹

Key Points:

- Iterative solver (PCG iteration)
- Multigrid preconditioner
- Based on discontinuous Petrov-Galerkin (DPG) finite element method of Demkowicz and Gopalakrishnan²



¹J. Badger. Scalable DPG multigrid solver with applications in high-frequency wave propagation. Ph.D. Thesis, The University of Texas at Austin, 2024.

²Leszek Demkowicz and Jay Gopalakrishnan. A class of discontinuous Petrov–Galerkin methods. Part I: the transport equation. *Comput. Methods Appl. Mech. Engrg.*, 199(23–24):1558–1572, 2010.

DPG (Linear Acoustics)

DPG:

$$\mathbf{u} = (p, \mathbf{u}) \in L^2(\Omega_h) \times \mathbf{L}^2(\Omega_h)$$

$$\hat{\mathbf{u}} = (\hat{p}, \hat{u}_n) \in H_{p_0, \Gamma_1}^{1/2}(\Gamma_h) \times H_{u_0, \Gamma_2}^{-1/2}(\Gamma_h)$$

$$\mathbf{v} = (q, \mathbf{v}) \in H^1(\Omega_h) \times H(\text{div}, \Omega_h)$$

$$\|\mathbf{v}\|_V^2 = \alpha \|q\|^2 + \alpha \|\mathbf{v}\|^2 + \|i\omega \bar{c}^{-2} q + \text{div}_h \mathbf{v}\|^2 + \|i\omega \bar{\mathbf{I}} \mathbf{v} + \nabla_h q\|^2$$

$$b(\mathbf{u}, \mathbf{v}) = i\omega(c^{-2} p, q) - (\mathbf{u}, \nabla_h q) + i\omega(\mathbf{I} \mathbf{u}, \mathbf{v}) - (p, \text{div}_h \mathbf{v})$$

$$\hat{b}(\hat{\mathbf{u}}, \mathbf{v}) = \langle \hat{u}_n, q \rangle_{\Gamma_h} + \langle \hat{p}, v_n \rangle_{\Gamma_h}$$

$$l(\mathbf{v}) = (\mathbf{f}_p, q) + (\mathbf{f}_u, \mathbf{v})$$

Galerkin:

$$\mathbf{u} = p \in H_0^1(\Omega_h)$$

$$\mathbf{v} = q \in H_0^1(\Omega_h)$$

$$b(\mathbf{u}, \mathbf{v}) = (\nabla p, \nabla q) - \omega^2 c^{-2} (p, q)$$

$$l(\mathbf{v}) = (\mathbf{f}_p, q)$$

Discrete System:

$$[\mathbf{B}][\mathbf{p}] = [\mathbf{f}]$$

Discrete System:

$$\begin{bmatrix} \mathbf{G}_{qq} & \mathbf{G}_{qv} & \mathbf{B}_{qp} & \mathbf{B}_{qu} & & \hat{\mathbf{B}}_{q\hat{v}_n} \\ \mathbf{G}_{vq} & \mathbf{G}_{vv} & \mathbf{B}_{vp} & \mathbf{B}_{vu} & \hat{\mathbf{B}}_{v\hat{p}} & \\ \mathbf{B}_{qp}^* & \mathbf{B}_{vp}^* & & & & \\ \mathbf{B}_{qu}^* & \mathbf{B}_{vu}^* & & & & \\ & \hat{\mathbf{B}}_{v\hat{p}}^* & & & & \\ \hat{\mathbf{B}}_{q\hat{v}_n}^* & & & & & \end{bmatrix} \begin{bmatrix} q \\ \mathbf{v} \\ p \\ \mathbf{u} \\ \hat{p} \\ \hat{u}_n \end{bmatrix} = \begin{bmatrix} \mathbf{f}_p \\ \mathbf{f}_u \end{bmatrix}$$

- DPG element matrices have ~10x as many DOFs (as Galerkin)
- Can be condensed onto 'trace' DOFs on element level:
 - Requires assembling and inverting large dense matrices
 - Condensed system 2x has many unknowns (4x as many non-zeros)
 - But resulting system is Hermitian positive definite

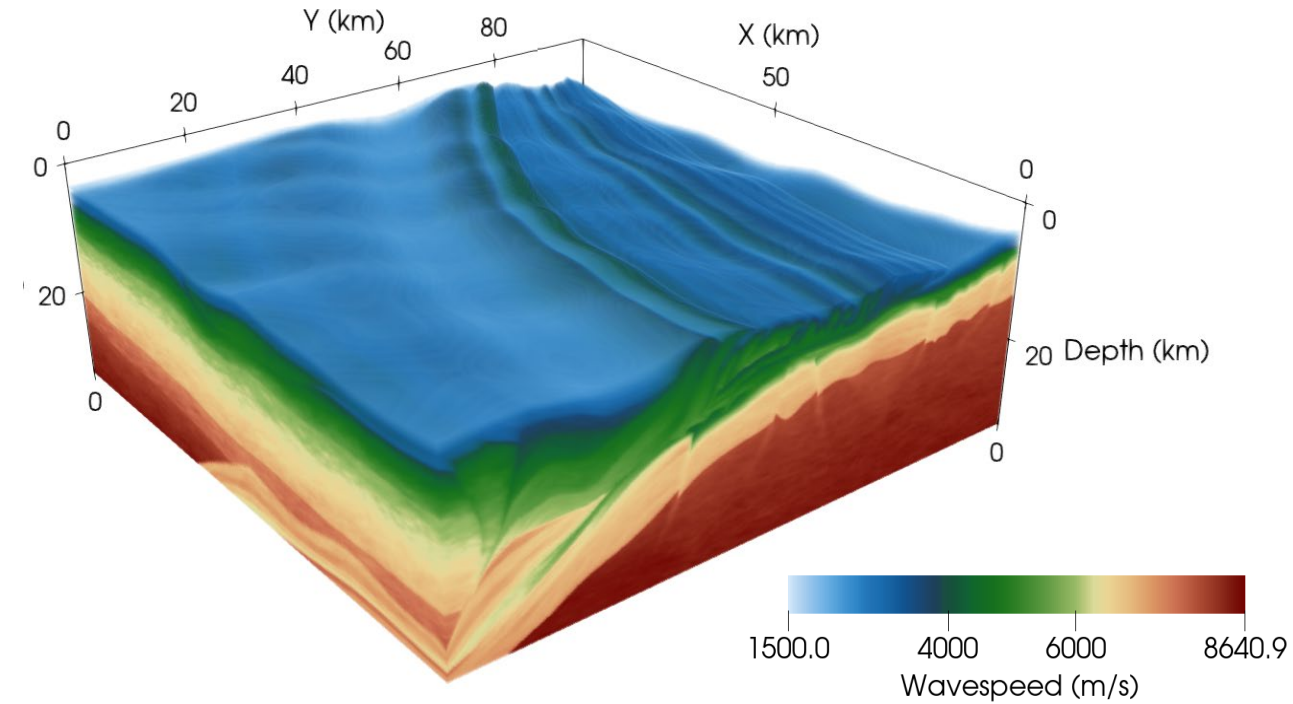
Example: 4th order elasticity

- Assembling blocks of full DPG system: ~0.2s
- Condensing onto trace DOFs: ~0.6s


**Assembling a single 4th order element can take ~1s
(and scales as p^9)**

Frontera-Enabled Outcomes

Scale (GO_3D_OBS Model¹)



- Visco-acoustic simulation (~6x contrast)
- 15 Hz (1,000 wavelengths)
- 800 billion DOFs (4th order)



¹ Andrzej Górszczyk and Stéphane Operto. GO_3D_OBS: the multi-parameter benchmark geomodel for seismic imaging method assessment and next-generation 3D survey design (version 1.0). *Geosci. Model Dev.*, 14(3):1773–1799, 2021.



Scale (GO_3D_OBS Model¹)



- 460,000 cores (8,192 Frontera nodes)
- ~500x larger than any previous result

¹ Andrzej Górszczyk and Stéphane Operto. GO_3D_OBS: the multi-parameter benchmark geomodel for seismic imaging method assessment and next-generation 3D survey design (version 1.0). *Geosci. Model Dev.*, 14(3):1773–1799, 2021.

Scale (SEAM Arid Model¹)

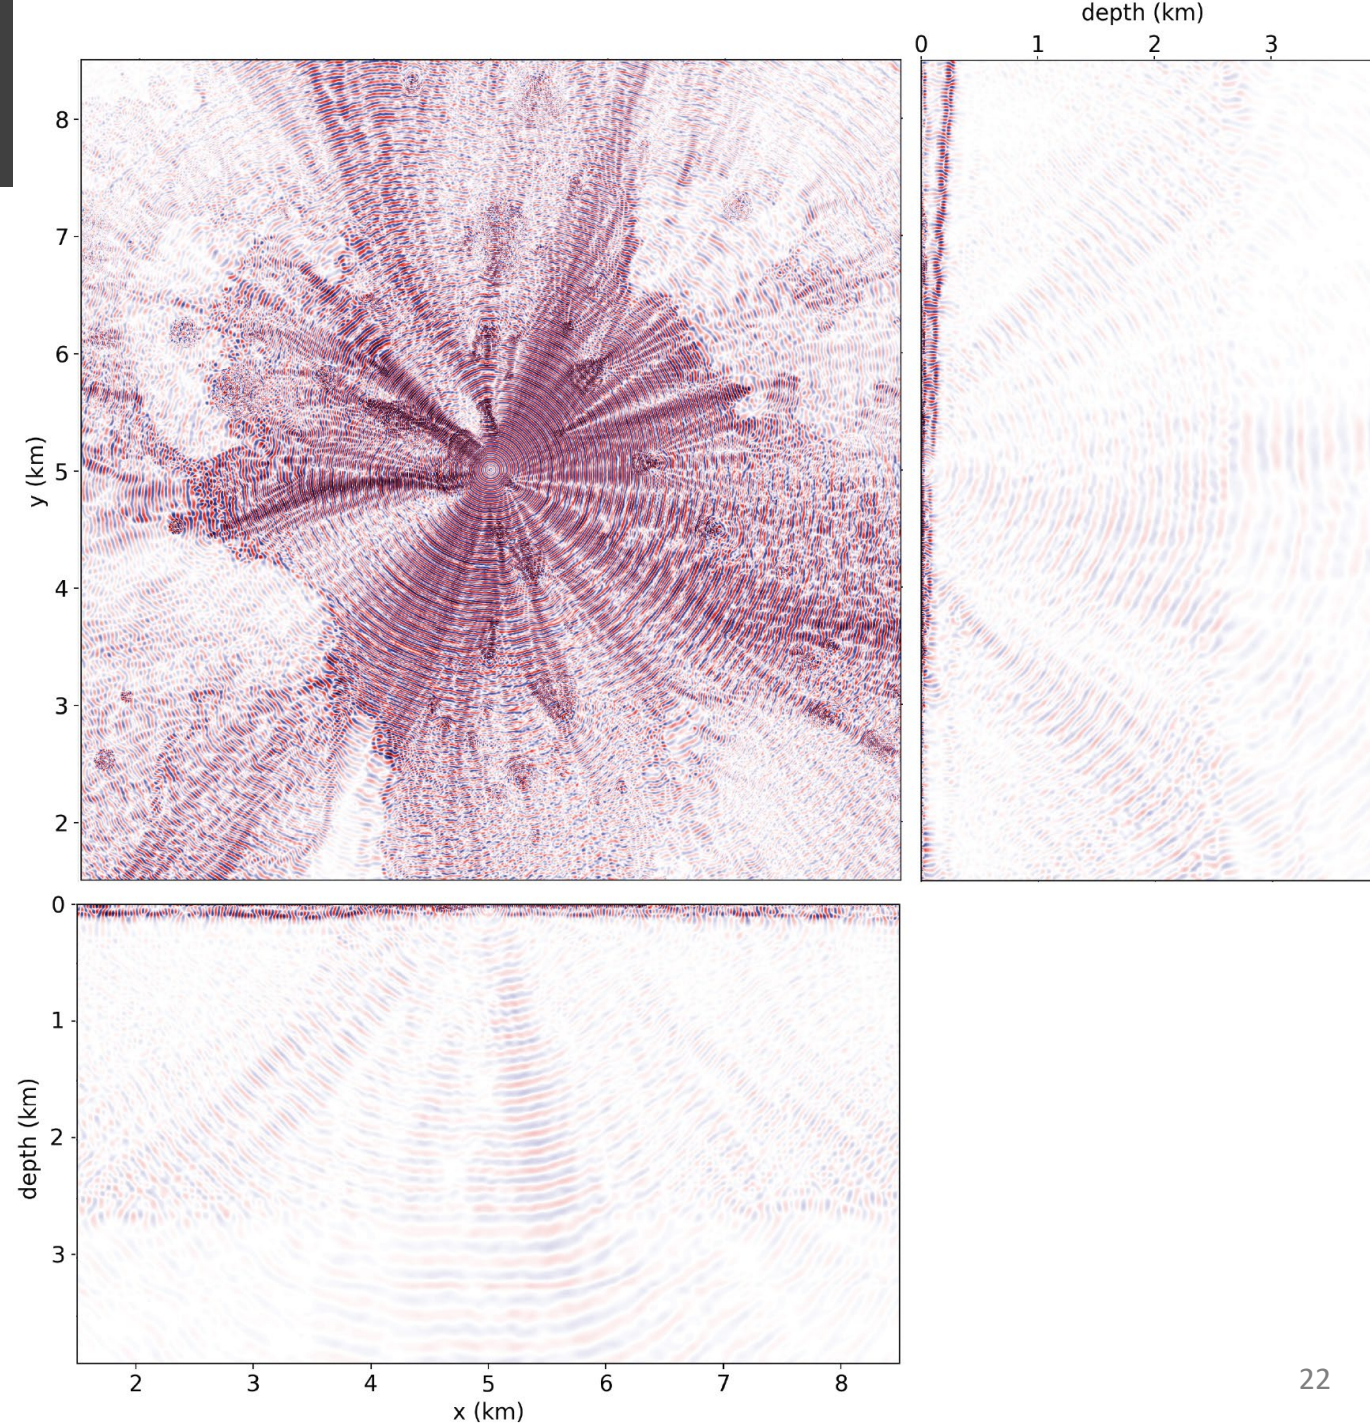


- HTI visco-elastic model (10x contrast)
- 25 Hz (~300 shear wavelengths)
- 7.4 Billion DOFs (~20x reduction vs. uniform mesh)

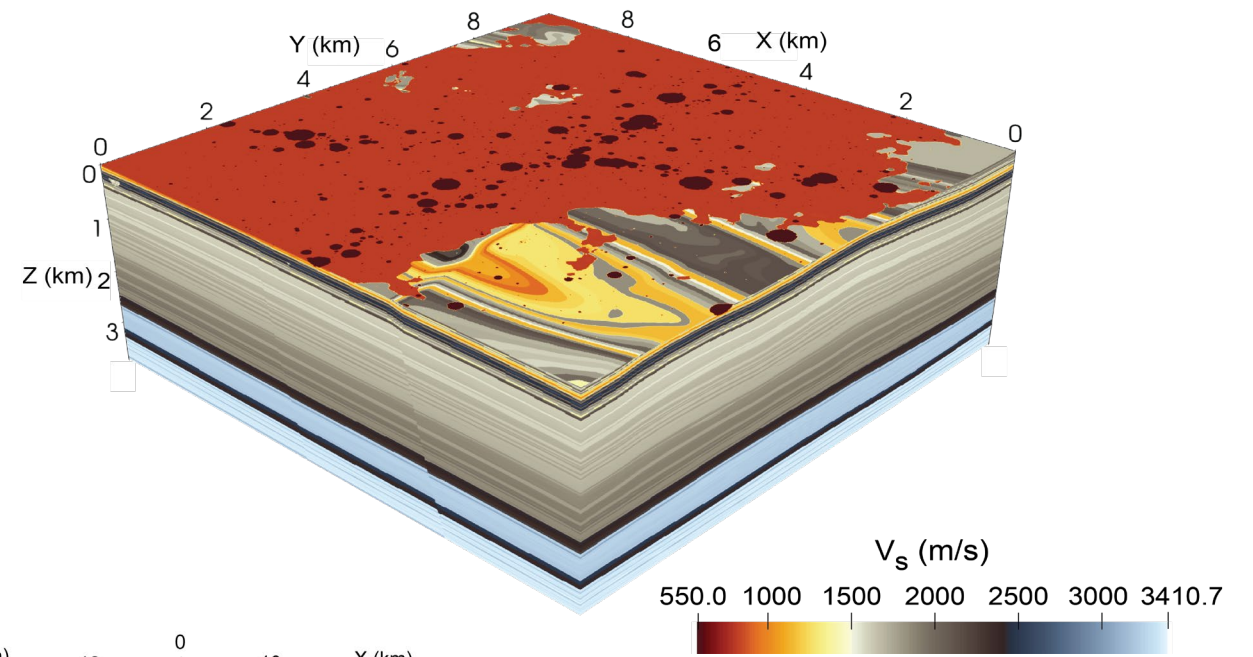

¹C. Regone, J. Stefani, P. Wang, C. Gerea, G. Gonzalez, and M. Oristaglio. Geologic model building in SEAM Phase II — Land seismic challenges. *The Leading Edge*, 36(9), 738–749, 2017.

Scale (SEAM Arid Model)

- 14,336 cores (256 Frontera nodes)
- ~20x larger than any previous result
(on basis of DOFs, for 3D frequency-domain elasticity)

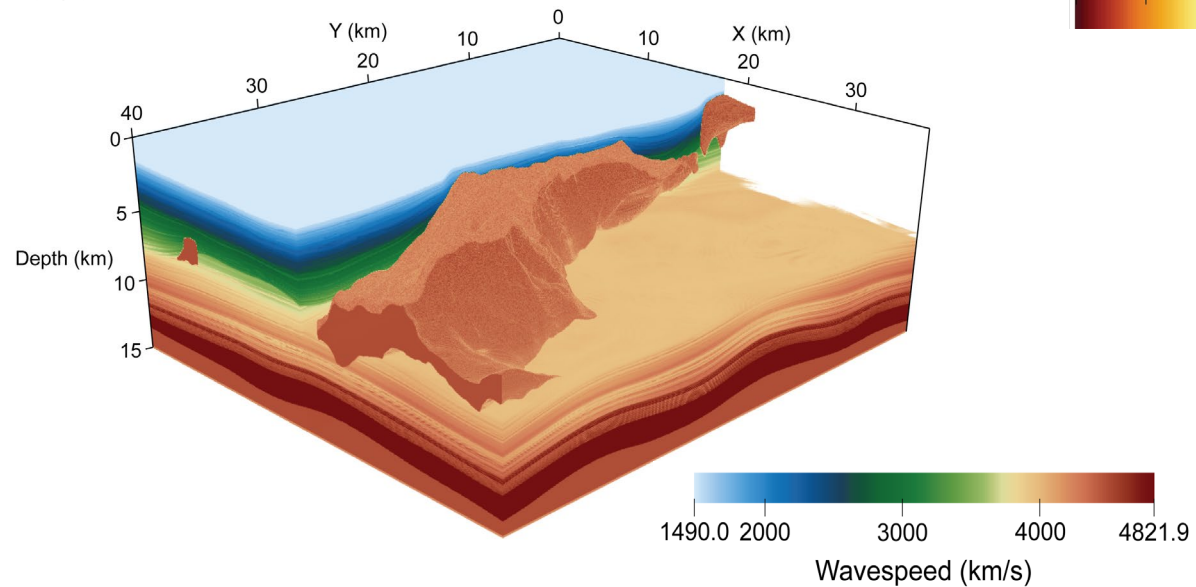


Acknowledgement



Special thanks to

Franchesca Samsel (TACC),
fantastic colormaps make a huge
difference in visualization




“Efficient seismic simulation in challenging geological environments with a scalable frequency-domain solver”

Directions



1. Efficient & accurate simulation of topography/small-scale heterogeneity
2. Design of distributed acoustic sensing (DAS) installations in complex geological settings
3. In-situ fracture imaging (applications in CCS, engineered geothermal, etc.)

Desert Environment with Buried Topography




- Visco-acoustic model
- High contrast (ca. 10x)
- Buried topography (flat surface)
- Attenuating near surface ($Q_p=40$)
- Frequencies 2-32 Hz simulated, then Fourier transformed (0.133 Hz spacing, 225 frequencies in total)

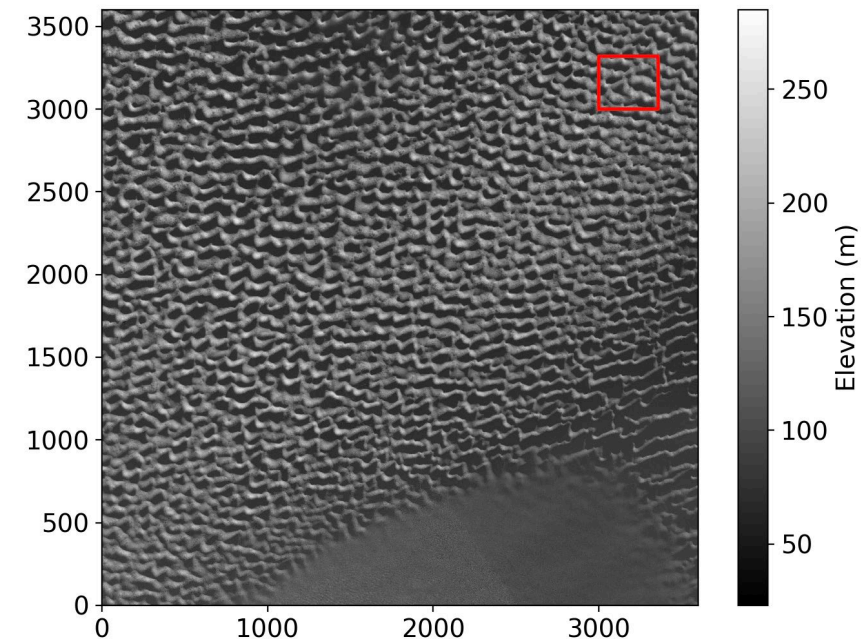
Flat Sand Model



Desert Environment with Dunes





- Dune topography added to sand layer (from NASA shuttle radar topography mission)
- Remainder of model unchanged



NASA radar tomography mission N22E054.hgt,
selected region boxed

Dunes



Insights

Insights

- Other promising frequency domain solvers include:
 - Domain decomposition (e.g. ORAS)
 - Complex-shifted Laplacian + multigrid
 - Many others
- Leading scalable solvers all have the same $O(\omega^4)$ complexity
- The rest is a game of constants (and DPG starts with a 4x disadvantage)

Insights

- Other promising frequency domain solvers include:
 - Domain decomposition (e.g. ORAS)
 - Complex-shifted Laplacian + multigrid
 - Many others
- Leading scalable solvers all have the same $O(\omega^4)$ complexity
- The rest is a game of constants (and DPG starts with a 4x disadvantage)

Points of Differentiation:

- Dense matrix operations (no sparse matrices)
- Adaptive mesh refinement
- Scale (Frontera + fully distributed data structures)

Sparse is Slow

Example: 4th order DPG acoustics (~200 x 200 matrix blocks)

- **Method 1:** MKL spBLAS
 - 64-bit indices
 - Complex single-precision
- **Method 2:** Element-wise block (dense) operations
 - Complex single-precision

```
for group in elem_groups:  
  
    # gather local block arrays from global array  
    x_exp = group.expand(x)  
  
    for block in group:  
        off = block.offset()  
        n = block.size()  
  
        # Apply block  
        Ax_loc[off:off+n] = block.apply(x_loc[off:off+n])  
  
    # Collapse (reduce) onto global array  
    Ax = group.collapse(Ap_loc)
```

Method 2: Sparse matrix multiplication via element-wise block operations

Sparse is Slow

Example: 4th order DPG acoustics (~200 x 200 matrix blocks)

- **Method 1:** MKL spBLAS
 - 64-bit indices
 - Complex single-precision
- **Method 2:** Element-wise block (dense) operations
 - Complex single-precision

Results (Frontera CLX node)	
1 core	1.4x
28 cores (OpenMP)	9.1x
28 cores + load batching (16 loads)	21x

Speed up of block-wise matrix multiplication relative to sparse

Reasons:

- Dense blocks are *Hermitian* (Stored in packed or RFP format, then unpacked)
- Sparse indices further double bandwidth per entry (for 64-bit indices w/ complex single precision)
- Less vectorizable (dot vs. outer product)
- Often higher cache miss rate

Sparse is Slow

Sparse Alternatives:

- Matrix-free (partial assembly; CEED, MFEM, etc.)
 - Memory efficient
 - Amenable to modern accelerators
 - Restrictive (tensor-product elements, uniform p , etc.)
 - Does ~ 2 - 5 x as many operations
 - Requires point smoothers
- Storing unassembled dense element matrices (single RHS)
 - More general (hybrid & p -adaptive meshes, etc.)
 - Block smoothers
 - Memory inefficient
 - Fewer operations but bandwidth limited (GEMV)

Sparse is Slow

Sparse Alternatives:

- Matrix-free (partial assembly; CEED, MFEM, etc.)
 - Memory efficient
 - Amenable to modern accelerators
 - Restrictive (tensor-product elements, uniform p , etc.)
 - Does $\sim 2\text{-}5x$ as many operations
 - Requires point smoothers
- Storing unassembled dense element matrices (batched RHS)
 - More general (hybrid & p -adaptive meshes, etc.)
 - Block smoothers
 - Fewer operations (GEMM)
 - System memory amortized over RHS's
 - Flexible PCG iteration requires at least 6 arrays per RHS
 - GMRES often requires many more arrays per RHS
- (Variable sized) block sparse matrices
 - In context of finite elements, blocks via nodal interactions

Future Directions & Summary

Summary

- **Scalable solver + Frontera** enabled simulation of frequency domain at unprecedented scale (~500x larger than any previous result we are aware of)
- **Adaptive meshing + performant implementation** can make the frequency-domain approach competitive (even in time domain data where hundreds of frequencies required)
- LRAC allocation enabling high-impact research on contemporary challenges:
 - 4D seismic noise and repeatability in on-shore monitoring for carbon capture and storage (CCS)
 - Design of distributed acoustic sensing (DAS) installations
 - In-situ fracture imaging (applications in CCS, engineered geothermal)

Open-source release of dissertation code intended soon
(disclosure completed in April 2024, awaiting signatures & decisions)

Final Thoughts

- **Frequency-domain** approach will be hard to beat in **frequency-sparse** contexts
 - Narrow bandwidth vibratory seismic sources
 - Plasma RF heating
 - Design of micro-optical filters, etc.
 - (and contexts with attenuation)
- Alternative discretizations may further improve efficiency of frequency-domain approach
 - DPG enables use of PCG iteration but is expensive (>100x slower assembly, 4x more non-zero entries vs. continuous Galerkin)
 - Block-wise dense operations and adaptivity largely responsible for surprising performance (neither are limited to DPG)
 - May require re-thinking the role of sparse matrices/factorizations in existing solvers
- **Implicit** time-domain methods + adaptivity may have potential in challenging contexts
 - Current practice dominated by explicit finite difference & spectral element methods:
 1. Less amenable to adaptivity
 2. Time-step limited by smallest element size
 - May require re-thinking the role of sparse matrices/factorizations
- Performant code requires work that is often less-incentivized by academic model
 - Implementation and optimization work can be difficult to publish; often happens in industry/at national labs
 - TACC has been an invaluable resource in bridging this gap