LFA-tuned matrix-free multigrid for the elastic Helmholtz equations

Racheli Yovel, Eran Treister

Computer Science Department,
Ben Gurion University of the Negev,
Be'er Sheva, Israel.

אוניברסיטת בן-גוריון בנגב
Ben-Gurion University of the Negev

The Helmholtz equations

The acoustic Helmholtz equation in heterogeneous media

$$Hp = \rho\nabla \cdot \rho^{-1}\nabla p + \omega^2\kappa^2(1-\gamma\imath)p = q$$

The Helmholtz equations

The acoustic Helmholtz equation in heterogeneous media

$$Hp = \rho\nabla \cdot \rho^{-1}\nabla p + \omega^2\kappa^2(1-\gamma\imath)p = q$$

The elastic Helmholtz equation in isotropic heterogeneous media

$$\mathcal{H}\vec{u} = \nabla(\lambda+\mu)\nabla\cdot\vec{u} + (\vec{\nabla}\cdot\mu\vec{\nabla})\vec{u} + \omega^2\rho(1-\gamma\imath)\vec{u} = \vec{q}$$

Pressure and shear wave velocities: $V_p = \sqrt{\frac{\lambda+2\mu}{\rho}}$ and $V_s = \sqrt{\frac{\mu}{\rho}}$.

The Helmholtz equations

The acoustic Helmholtz equation in heterogeneous media

$$Hp = \rho\nabla \cdot \rho^{-1}\nabla p + \omega^2\kappa^2(1-\gamma\imath)p = q$$

The elastic Helmholtz equation in isotropic heterogeneous media

$$\mathcal{H}\vec{u} = \nabla(\lambda+\mu)\nabla\cdot\vec{u} + (\vec{\nabla}\cdot\mu\vec{\nabla})\vec{u} + \omega^2\rho(1-\gamma\imath)\vec{u} = \vec{q}$$

Pressure and shear wave velocities: $V_p = \sqrt{\frac{\lambda+2\mu}{\rho}}$ and $V_s = \sqrt{\frac{\mu}{\rho}}$.

Boundary Conditions

- Neumann / Dirichlet: cause reflections.

The Helmholtz equations

The acoustic Helmholtz equation in heterogeneous media

$$Hp = \rho \nabla \cdot \rho^{-1} \nabla p + \omega^2 \kappa^2 (1 - \gamma \imath) p = q$$

The elastic Helmholtz equation in isotropic heterogeneous media

$$\mathcal{H}\vec{u} = \nabla(\lambda + \mu)\nabla \cdot \vec{u} + (\vec{\nabla} \cdot \mu \vec{\nabla})\vec{u} + \omega^2 \rho (1 - \gamma \imath)\vec{u} = \vec{q}$$

Pressure and shear wave velocities: $V_p = \sqrt{\frac{\lambda+2\mu}{\rho}}$ and $V_s = \sqrt{\frac{\mu}{\rho}}$.

Boundary Conditions

- Neumann / Dirichlet: cause reflections.
- Absorbing alternatives: Sommerfeld, **ABC**, PML.

The Helmholtz equations

Difficulties:

- Complex, indefinite and ill-conditioned

The Helmholtz equations

Difficulties:

- Complex, indefinite and ill-conditioned
- High frequencies $\Rightarrow$ fine meshes $\Rightarrow$ large matrix

The Helmholtz equations

Difficulties:

- Complex, indefinite and ill-conditioned
- High frequencies $\Rightarrow$ fine meshes $\Rightarrow$ large matrix
- Nearly incompressible case: the grad-div dominates when

$$\sigma = \frac{\lambda}{2(\lambda + \mu)} \to \frac{1}{2}$$

The Helmholtz equations

Difficulties:

- Complex, indefinite and ill-conditioned
- High frequencies ⇒ fine meshes ⇒ large matrix
- Nearly incompressible case: the grad-div dominates when

$$\sigma = \frac{\lambda}{2(\lambda + \mu)} \rightarrow \frac{1}{2}$$

- Numerical dispersion

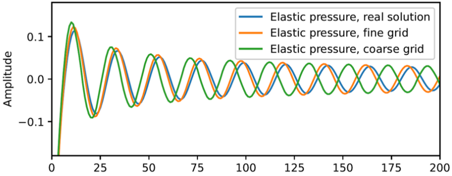

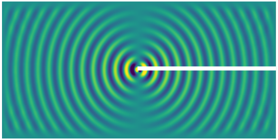

The Helmholtz equations

FD Discretizations for acoustic Helmholtz

- 4-th order [Harari & Turkel 1994], [Singer & Turkel 1998]
- 6-th order [Turkel et. al. 2013]
- Non-compact [Dastur & Liau 2020]
- Lower dispersion compact [Liu 2015], [Chen et. al. 2012]
- Rotated grids [Alghamiry et. al. 2022]
- Wavelength adaptive [Xu & Gao 2018]

Solvers for acoustic Helmholtz

- Domain decomposition [Gander & Zhang 2013], [Hu & Li 2016], [Tausa et. al. 2020], [Daia et. al. 2022], [Claeys, MS2], [Dolean, MS2], [Gong, MS2], [Rees, IP6]
- Shifted Laplacian multigrid [Erlangga et. al. 2006], [Umetani et. al 2009], [Chen et. al. 2012], [Dwarka, MS2], [Chen, MS1]
- FFT [Beylkin 2009], [Osnabrugge 2016], [Wang et. al. 2020]

Finite difference discretization and available solvers

FD Discretizations for elastic Helmholtz

- 2^{nd} order staggered (rotated grids) [Virieux 1986]
- Compact nodal [Sketl & Pratt 1998], [Gosselin-Cliche & Giroux 2014]
- 4^{th} order staggered [Levander 1988]
- 4^{th} order staggered with mass spreading [Li et. al. 2016]

[Gap]

Multigrid adaptations for elastic Helmholtz

- Shifted Laplacian "as is" [Airaksinen et. al. 2009]
- Shifted Laplacian using line-smoothers. [Rizzuti & Mulder 2016]

The idea of multigrid

Multigrid: iterative methods for elliptic PDEs [Brandt 1977]

- Smoothers (i.e. w-Jacobi, GS) reduce oscillatory error modes
- Coarse grid correction (CGC) reduce the smooth modes

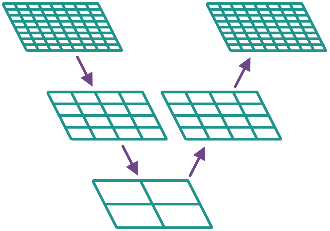

The idea of multigrid

Multigrid: iterative methods for elliptic PDEs [Brandt 1977]

- Smoothers (i.e. w-Jacobi, GS) reduce oscillatory error modes
- Coarse grid correction (CGC) reduce the smooth modes
- Isn't CGC enough? Why smoothers?

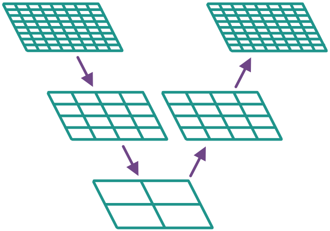

The idea of multigrid

Multigrid: iterative methods for elliptic PDEs [Brandt 1977]

- Smoothers (i.e. w-Jacobi, GS) reduce oscillatory error modes
- Coarse grid correction (CGC) reduce the smooth modes
- Isn't CGC enough? Why smoothers?

Aliasing:

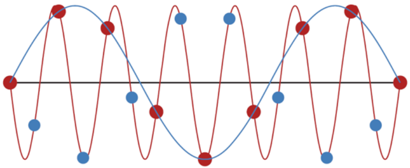

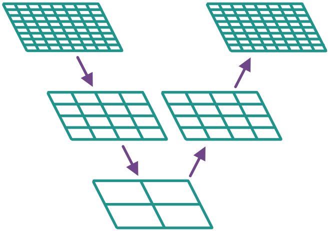

Left image from:
"Why Multigrid Methods Are So Efficient" [Yavneh, 2006]

The idea of multigrid

Multigrid: iterative methods for elliptic PDEs [Brandt 1977]

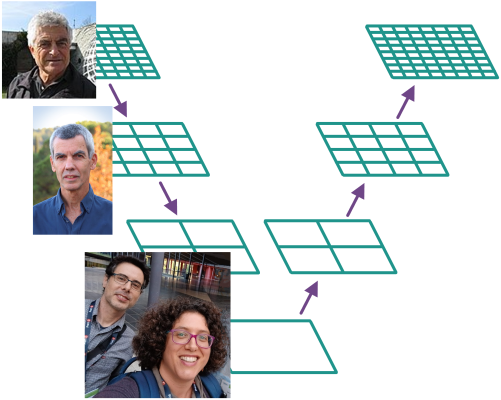

"Why Multigrid Methods Are So Efficient" [Yavneh, 2006]

Matrix-free multigrid

Two ways to build the coarse operator:

Matrix-free multigrid

Two ways to build the coarse operator:

- Galerkin coarsening: $A_H = RA_hP$.

Matrix-free multigrid

Two ways to build the coarse operator:

- Galerkin coarsening: $A_H = RA_hP$.
- Re-discretization: using the same stencil on coarse grid.

Matrix-free multigrid

Two ways to build the coarse operator:

- Galerkin coarsening: $A_H = RA_hP$.
- Re-discretization: using the same stencil on coarse grid.

Being matrix-free is not (only) about implementation, but about improving complexity [Pazner, IP2].

Local Fourier analysis

LFA is a predictive tool for analysis of multigrid cycles [Brandt 1977].

Local Fourier analysis

LFA is a predictive tool for analysis of multigrid cycles [Brandt 1977].

- **Two grid operator:**

$$TG = S(I - PA_H^{-1}RA_h)S$$

Local Fourier analysis

LFA is a predictive tool for analysis of multigrid cycles [Brandt 1977].

- **Two grid operator:**

$$TG = S(I - PA_H^{-1}RA_h)S$$

- **Smoothing factor:**

$$\mu_{loc} = \max_{\text{high } \theta}(\rho(\tilde{S}(\theta)))$$

Local Fourier analysis

LFA is a predictive tool for analysis of multigrid cycles [Brandt 1977].

- **Two grid operator:**

$$TG = S(I - PA_H^{-1}RA_h)S$$

- **Smoothing factor:**

$$\mu_{loc} = \max_{\text{high } \theta}(\rho(\tilde{S}(\theta)))$$

- **Two-grid factor:**

$$\rho_{loc} = \max_{\text{low } \theta}(\rho(\widetilde{TG}(\theta)))$$

Local Fourier analysis

LFA is a predictive tool for analysis of multigrid cycles [Brandt 1977].

- **Two grid operator:**

$$TG = S(I - PA_H^{-1}RA_h)S$$

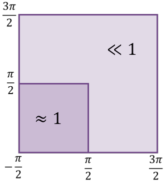

- **Smoothing factor:**

$$\mu_{loc} = \max_{\text{high } \theta}(\rho(\tilde{S}(\theta)))$$

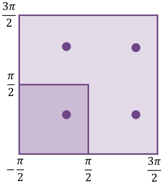

- **Two-grid factor:**

$$\rho_{loc} = \max_{\text{low } \theta}(\rho(\widetilde{TG}(\theta))$$

LFA for multiplicative overlapping smoothers

- The amplification factor $\approx \frac{|\text{error after}|}{|\text{error before}|}$
- We need "middle". "before" and "after" are not enough!

[Sivaloganathan 1991], [Maclachlan & Oosterlee 2011], [Rodrigo et. al. 2016], [Treister & Y. 2024]

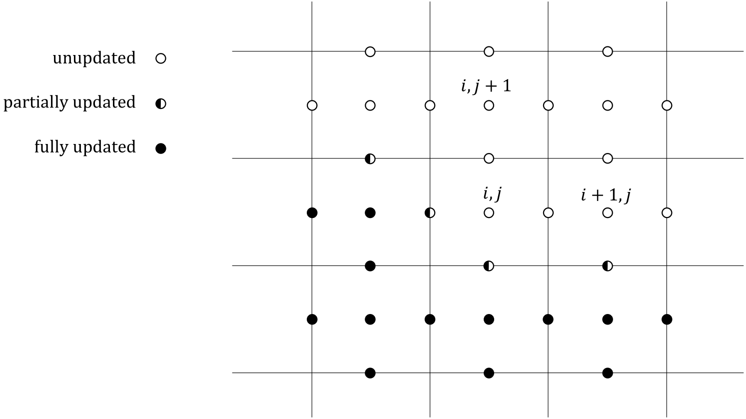

Shifted Laplacian multigrid

Standard smoothers: unstable for indefinite problems

Shifted Laplacian multigrid

Standard smoothers: unstable for indefinite problems

Coarse grid correction: might amplify error [Elman et. al. 2001]

Shifted Laplacian multigrid

Standard smoothers: unstable for indefinite problems

Coarse grid correction: might amplify error [Elman et. al. 2001]

Shifted Laplacian [Erlangga et. al. 2006]

Use an attenuated matrix $H_s = H + i\alpha\omega^2 M$ as a preconditioner.

The attenuated system $H_s\mathbf{x} = \mathbf{r}$ is solvable by multigrid.

Shifted Laplacian multigrid

Standard smoothers: unstable for indefinite problems

Coarse grid correction: might amplify error [Elman et. al. 2001]

Shifted Laplacian [Erlangga et. al. 2006]

Use an attenuated matrix $H_s = H + i\alpha\omega^2 M$ as a preconditioner.

The attenuated system $H_s\mathbf{x} = \mathbf{r}$ is solvable by multigrid.

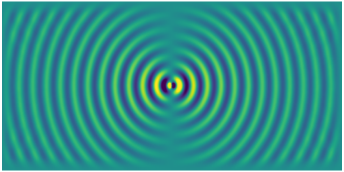

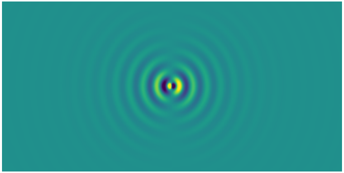

Multigrid for elastic Helmholtz

Problem: shifted Laplacian is not efficient for elastic Helmholtz.

Multigrid for elastic Helmholtz

Problem: shifted Laplacian is not efficient for elastic Helmholtz.

Our adaptations: a new variable $p = -(\lambda + \mu)\nabla \cdot \vec{u}$.

Elastic Helmholtz - mixed formulation

$$\begin{pmatrix} \vec{\nabla}\cdot\mu\vec{\nabla} + \omega^2 M & \nabla \\ \nabla\cdot & -\frac{1}{\lambda+\mu} Id \end{pmatrix} \begin{pmatrix} \vec{u} \\ p \end{pmatrix} = \begin{pmatrix} \vec{q} \\ 0 \end{pmatrix}.$$

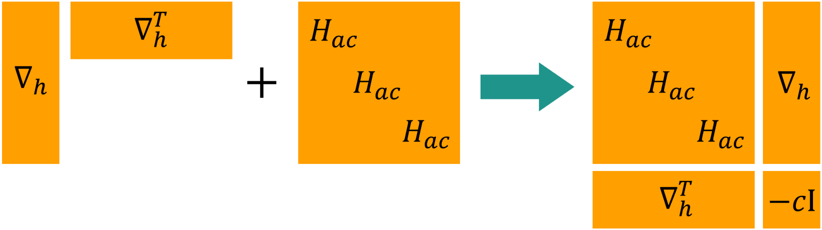

Multigrid for elastic Helmholtz

- MAC staggered grid finite differences discretization:
  [McKee et. al. 2008]

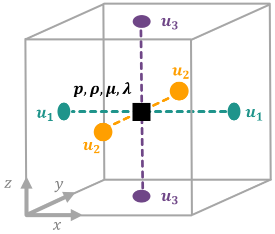

- MAC staggered grid finite differences discretization: [McKee et. al. 2008]

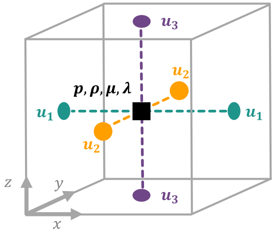

- Cell-wise Vanka smoother: invert a 5×5 matrix at every cell. [Vanka 1986]

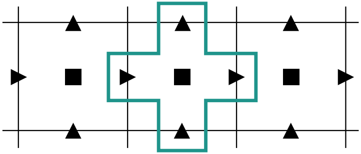

Multigrid for elastic Helmholtz: results

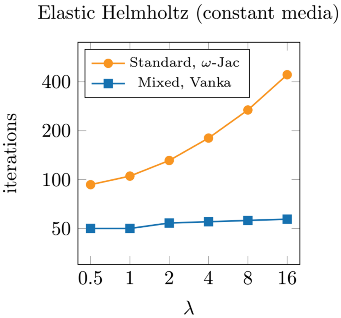

- $\mu = \rho = 1$, shift $\alpha = 0.2$. Second-order discretization.
- Scaling: largest λ corresponds to Poisson's ratio $\sigma = 0.47$
- Grid size 512×256, 10 grid points per wavelength
- GMRES(5) + 3-level W-cycle as a preconditioner
- Galerkin coarsening

Multigrid for elastic Helmholtz: results

Shifted Laplacian: acoustic (shear) vs. elastic

Grid size	400×128		800×256		1600×512	
ω	2.4π	3.5π	4.7π	7.1π	9.4π	14.2π
Acoustic	25	40	45	86	75	196
Elastic	27	37	47	78	79	148

- Linear heterogeneous media
- Acoustic: Jacobi W(2,2) cycles
- Elastic: Vanka W(1,1) cycles
- Added shift of 0.2 for both
- GMRES(5) + 3-level W-cycle as a preconditioner
- Galerkin coarsening

Matrix-free geometric MG

The standard 5 point stencil

$$-\mu\Delta_h - \omega^2 M = \frac{\mu}{h^2}\begin{bmatrix} & -1 & \\ -1 & 4-\frac{h^2}{\mu}\rho\omega^2(1-\imath\gamma) & -1 \\ & -1 & \end{bmatrix}$$

is not efficient enough for re-discretization.

Matrix-free geometric MG

The standard 5 point stencil

$$-\mu\Delta_h - \omega^2 M = \frac{\mu}{h^2}\begin{bmatrix} & -1 & \\ -1 & 4-\frac{h^2}{\mu}\rho\omega^2(1-\imath\gamma) & -1 \\ & -1 & \end{bmatrix}$$

is not efficient enough for re-discretization.
Acoustic: [Singer & Turkel 1998], [Umetani et. al. 2009]

$$\frac{1}{h^2}\begin{bmatrix} -1/6 & -2/3 & -1/6 \\ -2/3 & 10/3 & -2/3 \\ -1/6 & -2/3 & -1/6 \end{bmatrix} - \begin{bmatrix} & -1/12 & \\ -1/12 & -2/3 & -1/12 \\ & -1/12 & \end{bmatrix}\kappa^2\omega^2(1-\gamma\imath)$$

We adapt it to the elastic mixed formulation.

Matrix-free geometric MG: discretization

β-spread mass matrix:

$$M^{\beta} = \rho\omega^2(1-\gamma\imath)\left(\beta\begin{bmatrix}1\end{bmatrix} + (1-\beta)\cdot\frac{1}{4}\begin{bmatrix} & 1 & \\ 1 & & 1 \\ & 1 & \end{bmatrix}\right).$$

Matrix-free geometric MG: discretization

β-spread mass matrix:

$$M^{\beta} = \rho\omega^2(1-\gamma\imath)\left(\beta\begin{bmatrix}1\end{bmatrix} + (1-\beta)\cdot\frac{1}{4}\begin{bmatrix} & 1 & \\ 1 & & 1 \\ & 1 & \end{bmatrix}\right).$$

β-spread first derivative:

$$(\partial_{x_1})^{\beta}_{h/2} = \frac{1}{h}\left(\beta\begin{bmatrix}-1 & * & 1\end{bmatrix} + (1-\beta)\cdot\frac{1}{4}\begin{bmatrix}-1 & & 1 \\ -2 & * & 2 \\ -1 & & 1\end{bmatrix}\right),$$

Matrix-free geometric MG: discretization

β-spread mass matrix:

$$M^{\beta} = \rho\omega^2(1-\gamma\imath)\left(\beta\begin{bmatrix}1\end{bmatrix} + (1-\beta)\cdot\frac{1}{4}\begin{bmatrix} & 1 & \\ 1 & & 1 \\ & 1 & \end{bmatrix}\right).$$

β-spread first derivative:

$$(\partial_{x_1})^{\beta}_{h/2} = \frac{1}{h}\left(\beta\begin{bmatrix}-1 & * & 1\end{bmatrix} + (1-\beta)\cdot\frac{1}{4}\begin{bmatrix}-1 & & 1 \\ -2 & * & 2 \\ -1 & & 1\end{bmatrix}\right),$$

Divergence and gradient:

$$(\nabla\cdot)^{\beta}_{h} = \begin{pmatrix}(\partial_{x_1})^{\beta}_{h/2} & (\partial_{x_2})^{\beta}_{h/2}\end{pmatrix}, \quad \nabla^{\beta}_{h} = \begin{pmatrix}(\partial_{x_1})^{\beta}_{h/2} \\ (\partial_{x_2})^{\beta}_{h/2}\end{pmatrix}$$

Matrix-free geometric MG: discretization

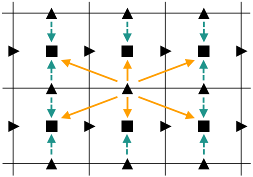

Matrix-free geometric MG: discretization

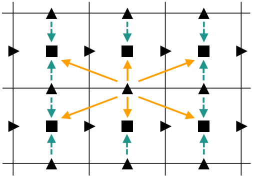

Singer & Turkel's acoustic stencil ($\beta = 2/3$):

$$\nabla_h^T \nabla_h^\beta + M^\beta = (\nabla\cdot)_h^\beta (\nabla\cdot)_h^T + M^\beta = -\Delta_h^\beta + M^\beta.$$

Matrix-free geometric MG: discretization

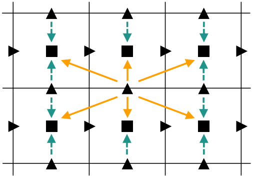

Singer & Turkel's acoustic stencil ($\beta = 2/3$):

$$\nabla_h^T \nabla_h^\beta + M^\beta = (\nabla\cdot)_h^\beta (\nabla\cdot)_h^T + M^\beta = -\Delta_h^\beta + M^\beta.$$

Discretization for the elastic equation in mixed formulation:

$$\begin{pmatrix} \vec{\nabla}_h^T A_e(\boldsymbol{\mu}) \vec{\nabla}_h^\beta - M^\beta A_f(\boldsymbol{\rho}) & (\nabla\cdot)_h^T \\ (\nabla\cdot)_h^\beta & \mathrm{diag}\left(\frac{1}{\boldsymbol{\lambda}+\boldsymbol{\mu}}\right) \end{pmatrix} \begin{pmatrix} \vec{\mathbf{u}} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \vec{\mathbf{q}}_s \\ \mathbf{0} \end{pmatrix}$$

where A_e is edge averaging, A_f is face averaging.

Matrix-free geometric MG: discretization

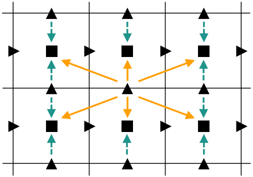

Singer & Turkel's acoustic stencil ($\beta = 2/3$):

$$\nabla_h^T \nabla_h^\beta + M^\beta = (\nabla\cdot)_h^\beta (\nabla\cdot)_h^T + M^\beta = -\Delta_h^\beta + M^\beta.$$

Discretization for the elastic equation in mixed formulation:

$$\begin{pmatrix} \vec{\nabla}_h^T A_e(\boldsymbol{\mu}) \vec{\nabla}_h^\beta - M^\beta A_f(\boldsymbol{\rho}) & (\nabla\cdot)_h^T \\ (\nabla\cdot)_h^\beta & \operatorname{diag}\left(\frac{1}{\boldsymbol{\lambda}+\boldsymbol{\mu}}\right) \end{pmatrix} \begin{pmatrix} \vec{\mathbf{u}} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \vec{\mathbf{q}}_s \\ \mathbf{0} \end{pmatrix}$$

where A_e is edge averaging, A_f is face averaging.

Matrix-free geometric MG: results

Tuning the stencil (with $G_s = 10$ and $\sigma = 0.499$):

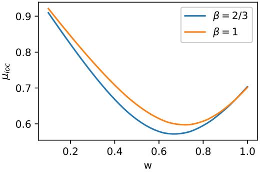

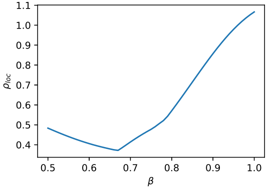

Damping $w = 0.7$ near optimal, $\beta = 2/3$ optimal

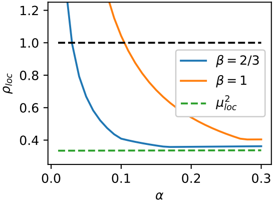

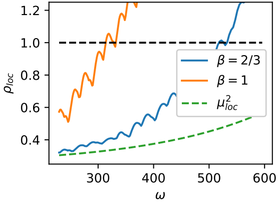

Minimal shift 0.03 for the spread discretization, 0.11 for standard

Numerical dispersion:

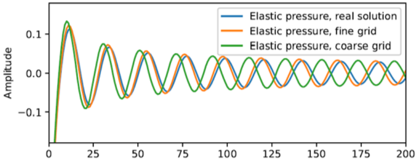

(a) $\beta = 1$

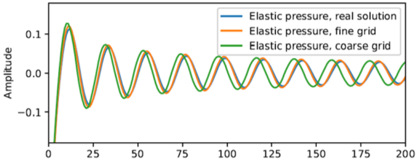

(b) $\beta = 0.8$

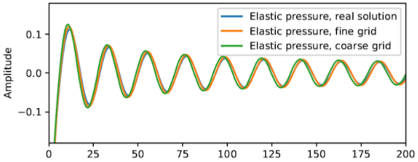

(c) $\beta = 2/3$

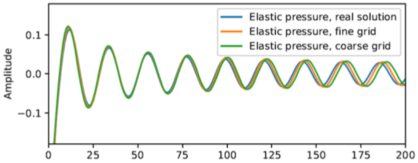

(d) $\beta = 0.5$

Matrix-free geometric MG: results

Multigrid convergence on the shifted problem:

LFA two-grid factor and convergence factor in practice									
	$\omega = \frac{\pi}{5h}$, $\alpha = 0.15$			$\omega = \frac{\pi}{4h}$, $\alpha = 0.2$			$\omega = \frac{\pi}{3.3h}$, $\alpha = 0.3$		
discretization	ρ_{loc}	c_f	μ^2_{loc}	ρ_{loc}	c_f	μ^2_{loc}	ρ_{loc}	c_f	μ^2_{loc}
$\beta = 1$	0.73	0.6	0.35	0.81	0.74	0.38	0.75	0.7	0.44
$\beta = 2/3$	0.37	0.24	0.31	0.4	0.27	0.34	0.47	0.35	0.39

- The convergence factor is defined by

$$c_f^{(k)} = \left(\frac{\|r_k\|}{\|r_0\|} \right)^{1/k}$$

- μ^2_{loc} serves as a best-case lower bound.

Matrix-free geometric MG: results

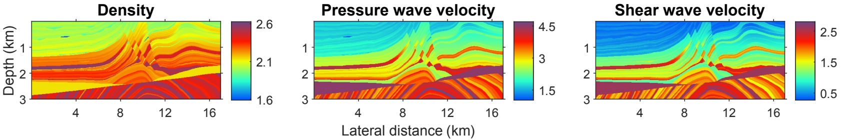

Results for Marmousi geophysical model:

Iteration count for Marmousi-2 elastic media

	$\beta = 1,\ G_s = 10$			$\beta = 2/3,\ G_s = 10\,(8)$		
Grid size	2-level $\alpha = 0.1$	3-level $\alpha = 0.4$	4-level $\alpha = 0.5$	2-level $\alpha = 0.1$	3-level $\alpha = 0.3\,(0.4)$	4-level $\alpha = 0.4\,(0.5)$
544×112	31	107	147	28 (32)	67 (106)	100 (194)
1088×224	48	172	238	39 (51)	102 (174)	157 (338)
2176×448	70	243	374	53 (63)	137 (242)	238 (572)

- α is the added attenuation
- G_s grid points per shear wavelength

Conclusion

- Mixed formulation enables elastic shifted Laplacian MG.
- Our shifted Laplacian MG scales w.r.t. Poisson's ratio.
- Our spread discretization enables a matrix-free method.
- Less shift is needed and less grid points per wavelength.

Future: Better 3D, Block-preconditioning, better acoustic solvers

Thanks

Funding:

- The U.S. Department of Energy (travel grant)
- Israel Science foundation (grant No. 1589/19)
- Ariane de Rothschild woman doctoral program
- Kreitman High-tech scholarship, BGU

References:

- *LFA-tuned matrix-free multigrid method for the elastic Helmholtz equation*, SISC, [Y. & Treister, 2024]
- *A hybrid shifted Laplacian multigrid and domain decomposition preconditioner for the elastic Helmholtz equations*, JCP, [Triester & Y., 2024]

`yovelr@bgu.ac.il`