

# Towards a Multigrid Preconditioner Interpretation of Hierarchical Poincaré–Steklov Solvers

José Pablo Lucero Lorca<sup>[0000–0002–9005–4146]</sup>

## 1 Introduction

The Hierarchical Poincaré–Steklov (HPS) method, introduced by Martinsson [16, 17], is a direct solver for elliptic boundary-value problems that combines nested dissection with spectral element discretizations on tensor-product grids. Subsequent work extended this framework to variable-coefficient Helmholtz equations, showing high accuracy and efficiency at scale [1, 8, 9]. In the impedance-to-impedance (ItI) formulation—based on the discretization of Després [6]—Dirichlet and Neumann traces are replaced by local impedance maps, yielding a closed interface representation well suited to high-frequency and heterogeneous media.

The present work places HPS in a preconditioned iterative setting, where the hierarchical merging tree provides the multilevel organization. Viewed equivalently as a nested-dissection solver for a spectral element discretization, HPS naturally induces a multilevel preconditioner, thereby unifying the direct and iterative perspectives.

Related hierarchical and multilevel strategies for elliptic and Helmholtz-type problems are numerous. On the multilevel side, classical frameworks include methods such as [3]. For hierarchical discretizations and direct/hybridization-style solvers, composite spectral collocation and related multidomain spectral approaches are surveyed in [8] and include, for example, [20]. For Helmholtz problems, iterative approaches including shifted-Laplace preconditioning, sweeping and domain decomposition methods, and multigrid variants are reviewed in [7].

The present work targets preconditioning of the skeleton (trace) system and leverages the modular HPS construction of [19]. A key practical drawback of HPS pipelines is their memory footprint and communication volume in distributed-memory settings, due to the exchange of dense interface operators (see [18] and references therein). Our aim is to *add* flexibility by avoiding construction of the most expensive coarse spaces: on the final level, the exact solve is replaced by a small, fixed number of iterations of an iterative method applied to the assembled last-level

---

e-mail: mail@pablo.world

system. A broader comparison with other iterative methods is deferred to future work.

## 2 Model problem

We consider the variable-coefficient Helmholtz equation with impedance boundary conditions

$$-\Delta u - \underbrace{\kappa^2 (1 - b(\mathbf{x}))}_{:=c(\mathbf{x})} u = s(\mathbf{x}), \quad \mathbf{x} \in \Omega \quad \text{and} \quad \frac{\partial u}{\partial n} + i\eta u = t(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega, \quad (1)$$

where  $\Omega = (0, 1)^2 \subset \mathbb{R}^2$  and  $u : \Omega \rightarrow \mathbb{C}$  is the unknown field,  $\eta \in \mathbb{R}$  chosen equal to  $\kappa \in \mathbb{R}$  the wavenumber,  $b(\mathbf{x})$  a smooth coefficient, and  $s(\mathbf{x}), t(\mathbf{x})$  smooth source and boundary data. Impedance boundary conditions of this form are widely used in diffraction, acoustics, and electromagnetic scattering [4, 11–13]; see also [10, §1.1, §1.2] for an overview.

## 3 Discretization

Consider a structured spectral element mesh,  $\Omega = (0, 1)^2$  is divided into a square grid of square elements, each with a tensor-product Gauss–Legendre–Lobatto (GLL) grid of order  $N$ . This construction allows high-order local operators from the tensor product of 1D differentiation and mass matrices while preserving continuity of impedance data on shared edges (see [6]). Local ItI maps are assembled element-by-element and coupled through interface conditions as described in the following sections (more detailed expositions can be found in [2, 14] and references therein).

### 3.1 Local discretization

Each element problem is represented by

$$\tilde{L} = K_x \otimes M_y + M_x \otimes K_y + \text{diag}(c(x_i, y_j))(M_x \otimes M_y), \quad (2)$$

where  $w_i, w_j$  are GLL quadrature weights,  $M_x = \text{diag}(w_i)$  and  $M_y = \text{diag}(w_j)$  are 1D GLL mass matrices,  $K_x = D_x^\top M_x D_x$  and  $K_y = D_y^\top M_y D_y$  are stiffness matrices, and  $D_x, D_y$  are the 1D differentiation matrices. The diagonal operator contains the coefficient  $c$  evaluated on the tensor grid  $\{(x_i, y_j)\}_{i,j=1}^{N+1}$ .

Following [14], the corner nodes are removed from the discretization, since they can be recalculated later in post-processing — this is a property of tensor-

product spectral methods. The boundary index sets are denoted  $t_l, t_r, t_b, t_t$  for the left, right, bottom, and top edges, and their union is  $t_\Gamma$ . The inner index set, denoted  $t_i$ , contains all remaining nodes strictly inside the element. The outgoing and incoming impedance operators are  $\mathcal{I}_o = \begin{bmatrix} -D_x \otimes I \\ D_x \otimes I \\ -I \otimes D_y \\ I \otimes D_y \end{bmatrix} (t_\Gamma, :) - \eta I(t_\Gamma, :)$  and  $\mathcal{I}_i = \begin{bmatrix} -D_x \otimes I \\ D_x \otimes I \\ -I \otimes D_y \\ I \otimes D_y \end{bmatrix} (t_\Gamma, :) + \eta I(t_\Gamma, :)$ , where  $I$  is the identity of appropriate size.

To apply incoming impedance conditions, the boundary rows of  $\tilde{L}$  are replaced by  $\mathcal{I}_i$  to define  $L$  as  $L(t_\Gamma, :) := \mathcal{I}_i$  and  $L(t_i, :) := \tilde{L}(t_i, :)$ . The local Impedance-to-Impedance operator and interior contribution are

$$T = \mathcal{I}_o L^{-1} I(:, t_\Gamma), \quad H = \mathcal{I}_o L^{-1} I(:, t_i) \tilde{b}(t_i), \quad (3)$$

where  $\tilde{b}$  contains the local right-hand-side values.

The operators  $T$  and  $H$  yield the closed impedance relation  $\mathcal{I}_o u(t_\Gamma) = T \mathcal{I}_i u(t_\Gamma) + H$ , from which the full element solution follows by

$$Lu = b, \quad (4)$$

where  $b(t_\Gamma) = \mathcal{I}_i u(t_\Gamma)$  and  $b(t_i) = \tilde{b}(t_i)$ .

### 3.2 Global discretization

For each element  $e$ , let  $\alpha, \beta \in \{l, r, b, t\}$  denote its sides. The local relation between outgoing and incoming impedance data is

$$(\mathcal{I}_o u)_\alpha^{(e)} = \sum_{\beta \in \{l, r, b, t\}} T_{\alpha\beta}^{(e)} (\mathcal{I}_i u)_\beta^{(e)} + H_\alpha^{(e)}, \quad (5)$$

where  $T_{\alpha\beta}^{(e)} \in \mathbb{C}^{(N-1) \times (N-1)}$  and  $H_\alpha^{(e)} \in \mathbb{C}^{(N-1)}$  are the operators defined in eq. (3).

Let two elements  $e_1$  and  $e_2$  share an interior face. Transmission conditions enforce continuity of impedance data across shared faces:

$$(\mathcal{I}_i u)_\beta^{(e_2)} = (\mathcal{I}_o u)_\alpha^{(e_1)}, \quad (\mathcal{I}_i u)_\alpha^{(e_1)} = (\mathcal{I}_o u)_\beta^{(e_2)}. \quad (6)$$

Combining these with the local ItI maps gives the face system

$$\begin{bmatrix} I & -T_{\alpha\alpha}^{(e_1)} \\ -T_{\beta\beta}^{(e_2)} & I \end{bmatrix} \begin{bmatrix} (\mathcal{I}_i u)_\beta^{(e_2)} \\ (\mathcal{I}_i u)_\alpha^{(e_1)} \end{bmatrix} - \begin{bmatrix} \sum_{\gamma \neq \alpha} T_{\alpha\gamma}^{(e_1)} (\mathcal{I}_i u)_\gamma^{(e_1)} \\ \sum_{\gamma \neq \beta} T_{\beta\gamma}^{(e_2)} (\mathcal{I}_i u)_\gamma^{(e_2)} \end{bmatrix} = \begin{bmatrix} H_\alpha^{(e_1)} \\ H_\beta^{(e_2)} \end{bmatrix}. \quad (7)$$

Assembling all face equations yields the sparse non-hermitian global skeleton system

$$Mg = \text{RHS},$$

where  $g$  collects all interior incoming impedances and RHS stacks the local  $H_\alpha^{(e)}$  contributions. Physical boundary sides contribute directly to the right-hand side.

### 4 Solver

The HPS solver applies the nested-dissection procedure to the spectral element system described above. This section details the face ordering that enables its direct solution, later recasted as a relaxation scheme.

#### 4.1 Nested dissection: Local scheme

Let two elements  $e_1$  and  $e_2$  share an interior face  $f$  through sides  $\alpha$  of  $e_1$  and  $\beta$  of  $e_2$ . Their face equations (from (7)) are

$$(\mathcal{I}u)_\beta^{(e_2)} - \sum_{\gamma \in \{l,r,b,t\}} T_{\alpha\gamma}^{(e_1)} (\mathcal{I}u)_\gamma^{(e_1)} = H_\alpha^{(e_1)}, \tag{8}$$

$$(\mathcal{I}u)_\alpha^{(e_1)} - \sum_{\gamma \in \{l,r,b,t\}} T_{\beta\gamma}^{(e_2)} (\mathcal{I}u)_\gamma^{(e_2)} = H_\beta^{(e_2)}. \tag{9}$$

We now group the internal and external incoming impedances as  $x = [(\mathcal{I}u)_\beta^{(e_2)} (\mathcal{I}u)_\alpha^{(e_1)}]^\top$  and  $y = [(\mathcal{I}u)_{E_1}^{(e_1)} (\mathcal{I}u)_{E_2}^{(e_2)}]^\top$  where  $E_1 = \{l, r, b, t\} \setminus \{\alpha\}$  and  $E_2 = \{l, r, b, t\} \setminus \{\beta\}$ . With this notation the system becomes

$$\underbrace{\begin{bmatrix} I & -T_{\alpha\alpha}^{(e_1)} \\ -T_{\beta\beta}^{(e_2)} & I \end{bmatrix}}_A x = \underbrace{\begin{bmatrix} T_{\alpha E_1}^{(e_1)} & 0 \\ 0 & T_{\beta E_2}^{(e_2)} \end{bmatrix}}_B y + \underbrace{\begin{bmatrix} H_\alpha^{(e_1)} \\ H_\beta^{(e_2)} \end{bmatrix}}_h. \tag{10}$$

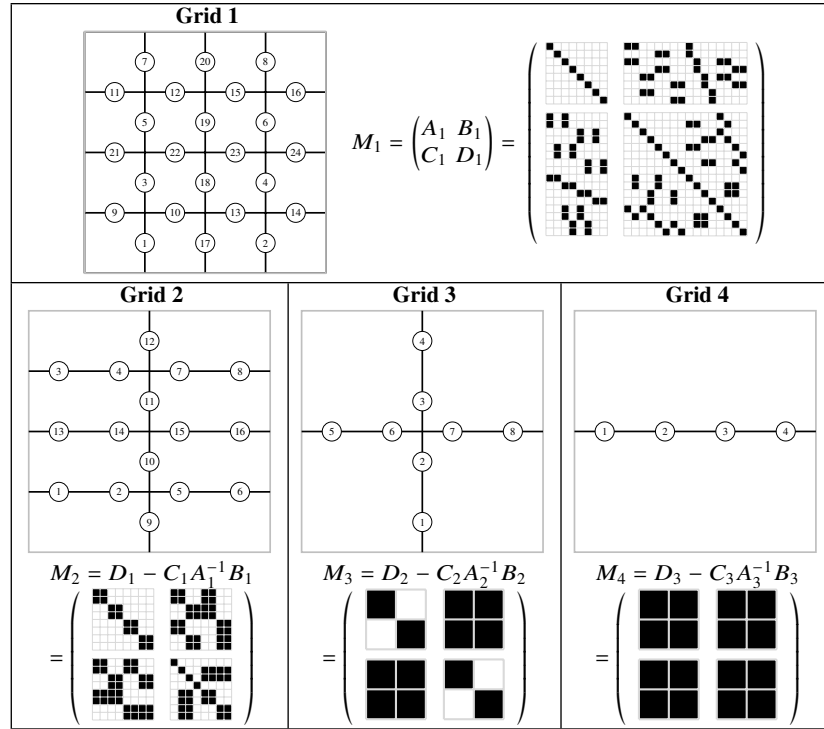
Eliminating  $x$  gives  $x = A^{-1}By + A^{-1}h$ , substituting into the outgoing relations (5) produces the fused pair operator

$$T_{\text{pair}} = \underbrace{D}_{\begin{bmatrix} T_{E_1 E_1}^{(e_1)} & 0 \\ 0 & T_{E_2 E_2}^{(e_2)} \end{bmatrix}} - \underbrace{C}_{\begin{bmatrix} T_{E_1 \alpha}^{(e_1)} & 0 \\ 0 & T_{E_2 \beta}^{(e_2)} \end{bmatrix}} A^{-1} B, \quad H_{\text{pair}} = \begin{bmatrix} H_{E_1}^{(e_1)} \\ H_{E_2}^{(e_2)} \end{bmatrix} + CA^{-1} \begin{bmatrix} H_\alpha^{(e_1)} \\ H_\beta^{(e_2)} \end{bmatrix}. \tag{11}$$

where  $T_{\text{pair}}$  is clearly a Schur complement.

#### 4.2 Nested dissection: Global scheme

Figure 1 illustrates the face-merging procedure for the skeleton system  $M_1$  on a  $4 \times 4$



**Fig. 1** Face merging and sparsity patterns for a  $4 \times 4$  element mesh **Grid 1**: Faces 1 to 8 are eliminated, merging pairs of elements. These faces' dofs form the top left  $1 \times 1$ -face-block diagonal part of  $M_1$  since they are not linked *directly* between each other, but through another face, e.g. face 1 is related to face 2 through face 17. **Grid 2**: Faces 1 to 8 are eliminated by pairs, merging  $1 \times 2$  subdomains by one of their largest sides. These faces' dofs form the top left  $2 \times 2$ -face-block diagonal matrix. **Grid 3**: Faces 1 to 4 are eliminated by pairs, merging  $2 \times 2$  subdomains. These faces' dofs form the top left  $2 \times 2$ -face-block diagonal matrix. **Grid 4**: Faces 1 to 4 are now fully coupled,  $M_4$  is dense.

element mesh. The sequence Grid 1–Grid 4 is nested dissection in reverse. It lists, from fine to coarse, the elimination sets used by the solver. The key consequence of using this hierarchy is that at every level  $\ell$  the eliminated-face block  $A_\ell$  is block diagonal, with blocks growing as subdomains are merged. Applying  $A_\ell^{-1}$  to form the next Schur complement has controlled cost, while the fill is pushed to coarser levels.

Let  $f_i$  denote the face labeled  $i$  in Grid 1. The first elimination set is  $\mathcal{F}_1 = \{f_1, \dots, f_8\}$ . It is maximal among the nested dissection ordered subsets of Grid 1 faces that are pairwise element-disjoint. Thus the faces in  $\mathcal{F}_1$  do not couple directly and can be eliminated independently. Eliminating  $\mathcal{F}_1$  merges the element pairs adjacent to these faces, producing the  $(1 \times 2)$ -element subdomains shown in Grid 2. Given the nested dissection face ordering,  $M_1$  is partitioned as  $M_1 = \begin{pmatrix} A_1 & B_1 \\ C_1 & D_1 \end{pmatrix}$  and eliminat-

ing  $\mathcal{F}_1$  yields  $M_2 = D_1 - C_1 A_1^{-1} B_1$ , the reduced face system shown in Grid 2. In this example the diagonal blocks of  $M_2$  are the fused ItI operators for the merged element pairs. The remaining nonzeros in  $M_2$  encode couplings between these merged pairs across the faces of Grid 2. The coarse grids in Grid 2–Grid 4 are tensor-product in the geometric sense that the merges produce axis-aligned rectangular subdomains arranged in a Cartesian tiling.

At level  $\ell \geq 1$  we analogously eliminate a maximal element-disjoint set  $\mathcal{F}_\ell$  in the level- $\ell$  face graph and obtain  $M_{\ell+1} = D_\ell - C_\ell A_\ell^{-1} B_\ell$ . The eliminated unknowns at level  $\ell$  decompose into disjoint merged-subdomain groups. Accordingly  $A_\ell$  is block diagonal, with one block per merged subdomain at that level. Grid 3 and Grid 4 visualize the next two levels. The blocks in  $A_\ell$  grow, and the Schur complements become progressively denser as the hierarchy coarsens. Repeating this construction and using the associativity of Schur complements [5] yields the nested-dissection factorization of the spectral-element face system.

### 4.3 Solver recast as a multigrid relaxation scheme

The block-inverse relation introduced in [15] takes the form

$$M^{-1} = \begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \underbrace{\begin{bmatrix} A^{-1} & 0 \\ 0 & 0 \end{bmatrix}}_F + \underbrace{\begin{bmatrix} -A^{-1}B \\ I \end{bmatrix}}_P \underbrace{\left( D - CA^{-1}B \right)^{-1}}_{S^{-1}} \underbrace{\begin{bmatrix} 0 & I \end{bmatrix}}_R \left( I - M \underbrace{\begin{bmatrix} A^{-1} & 0 \\ 0 & 0 \end{bmatrix}}_F \right). \quad (12)$$

This identity motivates the definition of a recursive multigrid algorithm without post-smoothing rather than a single relaxation step: the local inversion  $A^{-1}$  acts as a smoother, and the reduced system  $S$  defines the next level. The recursive iteration  $\text{MG}(M) = F + PS^{-1}R(I - MF)$ , where  $S^{-1}$  is obtained by applying the same procedure to  $S$ . A single coarse call yields a V-cycle; multiple ones define a  $\gamma$ -cycle—both fully consistent with the hierarchical merging in the HPS method. We employ MG as a preconditioner for flexible GMRES, with the coarse solve performed by a fixed number of unpreconditioned GMRES iterations.

## 5 Numerical experiments

We consider one of the problems from [14], with  $b(\mathbf{x}) = 1.5e^{-160[(x-0.5)^2+(y-0.5)^2]}$  and  $s(\mathbf{x}) = -\kappa^2 b(\mathbf{x})e^{i\kappa x}$ , representing scattering by a Gaussian bump. We use polynomial degree 16, a residual tolerance of  $10^{-8}$ , and a frequency giving 9.6 points per wavelength, yielding about  $10^{-7}$  accuracy (verifying the estimate in [14]) for roughly one million degrees of freedom before skeletonization.

Figure 2 shows the solution, and Table 1 reports results obtained in MATLAB, varying the number of levels. The table lists memory footprint, build time, total

iterations, and solve time for different fixed coarse iteration counts and  $\gamma$  values. The face sets used to build the multilevel grids are those described in Section 4.2. The problem was run on a laptop with 32 GB RAM and a hybrid processor (6 hyper-threading cores @ 4.7 GHz and 8 cores @ 3.5 GHz). Although cache effects favor certain configurations, an overall timing trend can be observed. The method demonstrates that performance can be tuned to available memory and the number of solves required, while being faster than the unpreconditioned case in many configurations. For context, we compare against the classic direct HPS method, which uses an exact coarse space. It can be observed that accepting a few iterations can save a significant amount of memory footprint.

## 6 Conclusion

We provide a flexible iterative variant of the otherwise direct HPS method for variable-coefficient Helmholtz problems arising, for instance, in wave propagation and geological prospection: on different coarse levels we replace the exact solve by a small, fixed number of FGMRES iterations, explicitly trading a few Krylov steps for a reduced coarse-space memory footprint. The large-scale (including 3D) assessments in [14, 18] indicate that the coarse-level operators can dominate memory, even when factorized/compressed via SVD-type techniques; our goal is to mitigate this bottleneck. Since current architectures tend to be more compute-rich than memory-rich, we expect this tradeoff to enable larger problem sizes at fixed accuracy by paying only a few additional FGMRES iterations.

## References


1. Babb, T., Gillman, A., Hao, S., Martinsson, P.G.: An accelerated poisson solver based on a multidomain spectral discretization. *BIT* **58**(4), 851–879 (2018)
2. Beams, N.N., Gillman, A., Hewett, R.J.: A parallel shared-memory implementation of a high-order accurate solution technique for variable coefficient Helmholtz problems. *Comput. Math. Appl.* **79**(4), 996–1011 (2020)
3. Cai, X.C., Goldstein, C., Pasciak, J.E.: Multilevel iteration for mixed finite element systems with penalty. *SIAM J. Sci. Comput.* **14**(4), 1072–1088 (1993)
4. Chandler-Wilde, S.N.: The impedance boundary value problem for the Helmholtz equation in a half-plane. *Math. Methods Appl. Sci.* **20**(10), 813–840 (1997)
5. Crabtree, D.E., Haynsworth, E.V.: An identity for the Schur complement of a matrix. *Proc. Am. Math. Soc.* **22**(2), 364–366 (1969)
6. Després, B.: Méthodes de décomposition de domaine pour les problèmes de propagation d’ondes en régime harmonique. Ph.D. thesis, Université Paris IX Dauphine (1991)
7. Gander, M.J., Zhang, H.: A class of iterative solvers for the Helmholtz equation: Factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods. *SIAM Rev.* **61**(1), 3–76 (2019)
8. Gillman, A., Barnett, A.H., Martinsson, P.G.: A spectrally accurate direct solution technique for frequency-domain scattering problems with variable media. *BIT* **55**(1), 141–170 (2015)

**Table 1 PMem:** Preconditioner Memory Footprint [MB], **It:** Flexible GMRES iterations with restart at 60, **Bt:** Build time [s], **St:** Solve time [s], **c.i.:** coarse GMRES iterations. Results for  $10^6$  dofs at 9.6 points per wavelength.

Case		PMem	It	Bt	St
Unpreconditioned		0	669	0	85
Exact coarse space		3108	1	75	4

#levels	PMem	Bt	$\gamma = 1$						$\gamma = 2$					
			4 c.i.		5 c.i.		6 c.i.		2 c.i.		3 c.i.		4 c.i.	
			It	St	It	St	It	St	It	St	It	St	It	St
2	46	6	37	53	22	44	16	45	83	71	32	44	18	39
3	460	15	23	42	15	40	11	40	24	55	11	42	7	42
4	805	20	18	30	12	27	9	27	11	48	6	41	4	43
5	1202	27	13	36	9	34	7	36	5	55	<b>3</b>	<b>71</b>	1	52
6	1527	31	11	22	7	18	5	16	2	47	1	46		
7	1897	38	9	28	6	26	4	23			<b>1</b>	<b>90</b>		
8	2185	43	8	19	5	15	4	15						
9	2502	45	8	28	5	23	4	24						
10	2724	52	7	19	5	17	4	15						
11	2946	63	<b>7</b>	<b>26</b>	<b>5</b>	<b>24</b>	<b>4</b>	<b>24</b>						
12	3051	67	3	11	2	8	1	6						



**Fig. 2** Solution of the variable-coefficient Helmholtz problem.

#### Notes.

$\eta$  is chosen equal to  $\kappa$  to obtain  $10^6$  dofs while yielding  $\approx 10^{-7}$  accuracy; the estimate in [14] was confirmed in these runs. FGMRES is initialized with the zero vector. PMem includes the full memory usage of the program. Restart 60 was selected to avoid exhausting laptop memory. Bold entries satisfy  $B_t + S_t > 85$ .

- Gillman, A., Martinsson, P.G.: A direct solver with  $O(n)$  complexity for variable coefficient elliptic PDEs discretized via a high-order composite spectral collocation method. *SIAM J. Sci. Comput.* **36**(4), A2023–A2046 (2014)
- Graham, I.G., Sauter, S.A.: Stability and finite element error analysis for the Helmholtz equation with variable coefficients. *Math. Comp.* **89**, 105–138 (2020)
- Kinsler, L.E., Frey, A.R., Coppens, A.B., Sanders, J.V.: *Fundamentals of acoustics*. Wiley (2000)
- Kyurkchan, A.G., Smirnova, N.I.: *Mathematical Modeling in Diffraction Theory*. Elsevier, Amsterdam (2016)
- Lighthill, M.J.: *Waves in Fluids*. Cambridge University Press, Cambridge, England (1978)
- Lucero Lorca, J.P., Beams, N., Beecroft, D., Gillman, A.: An iterative solver for the HPS discretization applied to three dimensional Helmholtz problems. *SIAM J. Sci. Comput.* **46**(1), A80–A104 (2024)
- Lucero Lorca, J.P., Rosenberg, D., Jankov, I., McCoid, C., Gander, M.: On an efficient line smoother for the  $p$ -multigrid  $\gamma$ -cycle. *arXiv:2504.10710* (2025)
- Martinsson, P.: A direct solver for variable coefficient elliptic PDEs discretized via a composite spectral collocation method. *J. Comput. Phys.* **242**, 460–479 (2013)
- Martinsson, P.G.: The Hierarchical Poincaré–Steklov (HPS) solver for elliptic PDEs: A tutorial. *arXiv:1506.01308* (2015)
- Melia, O., Fortunato, D., Hoskins, J., Willett, R.: Hardware acceleration for HPS algorithms in two and three dimensions. *J. Comput. Phys.* **547**, 114549 (2026)
- Outrata, M., Lucero Lorca, J.P.: Towards modular Hierarchical Poincaré–Steklov solvers. Submitted in the proceedings of DD29 (2025)
- Pfeiffer, H., Kidder, L., Scheel, M., Teukolsky, S.: A multidomain spectral method for solving elliptic equations. *Comput. Phys. Commun.* **152**(3), 253–273 (2003)