



HAL
open science

A comparison of coarse spaces for Helmholtz problems in the high frequency regime

Niall Bootland, Victorita Dolean, Pierre Jolivet, Pierre-Henri Tournier

► **To cite this version:**

Niall Bootland, Victorita Dolean, Pierre Jolivet, Pierre-Henri Tournier. A comparison of coarse spaces for Helmholtz problems in the high frequency regime. 2021. hal-03136867

HAL Id: hal-03136867

<https://hal.science/hal-03136867>

Preprint submitted on 10 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A comparison of coarse spaces for Helmholtz problems in the high frequency regime

Niall Bootland^{a,1}, Victorita Dolean^{a,b,1}, Pierre Jolivet^c, Pierre-Henri
Tournier^{d,2}

^a*Department of Mathematics and Statistics, University of Strathclyde, Livingstone
Tower, 26 Richmond Street, Glasgow, G1 1XH, UK*

^b*Laboratoire J.A. Dieudonné, CNRS, University Côte d'Azur, Parc Valrose, 28 Avenue
Valrose, Nice, 06108, France*

^c*Institut de Recherche en Informatique de Toulouse, University of Toulouse, 2 rue
Charles Camichel, Toulouse, 31400, France*

^d*Laboratoire Jacques-Louis Lions, Sorbonne University, 4 Place
Jussieu, Paris, 75005, France*

Abstract

Solving time-harmonic wave propagation problems in the frequency domain and within heterogeneous media brings many mathematical and computational challenges, especially in the high frequency regime. We will focus here on computational challenges and try to identify the best algorithm and numerical strategy for a few well-known benchmark cases arising in applications. The aim is to cover, through numerical experimentation and consideration of the best implementation strategies, the main two-level domain decomposition methods developed in recent years for the Helmholtz equation. The theory for these methods is either out of reach with standard mathematical tools or does not cover all cases of practical interest. More precisely, we will focus on the comparison of three coarse spaces that yield two-level methods: the grid coarse space, DtN coarse space, and GenEO coarse space. We will show that they display different pros and cons, and properties depending on

Email addresses: niall.bootland@strath.ac.uk (Niall Bootland),
work@victoritadolean.com (Victorita Dolean), pierre.jolivet@enseeiht.fr (Pierre
Jolivet), tournier@ljl1.math.upmc.fr (Pierre-Henri Tournier)

¹The first two authors gratefully acknowledge support from the EPSRC grant EP/S004017/1.

²The fourth author gratefully acknowledges support from the EPSRC grant EP/R009821/1.

the problem and particular numerical setting.

Keywords: Helmholtz equations, domain decomposition methods, two-level methods, coarse spaces, high frequency

2000 MSC: 65N55, 65N35, 65F10

1. Introduction

This work is motivated by the computational challenges that arise in frequency domain simulations of wave propagation and scattering problems in heterogeneous media. Such problems appear in a broad range of engineering applications, including acoustics, electromagnetics, and seismic inversion.

The discretisation of models describing frequency domain wave problems using finite element methodology typically results in large, indefinite, and ill-conditioned linear systems. These linear systems are difficult to solve using standard methods, particularly for high frequencies and in the presence of complex heterogeneities. In order to maintain accuracy, the number of grid points must grow as a function of the frequency in such a way that, for high frequency problems, the size of the linear systems to be solved becomes prohibitive for direct methods. In such a regime, carefully designed iterative methods are required. Here, we consider a two-level domain decomposition approach for the robust parallel solution of the linear systems.

To model the wave problem, we utilise the Helmholtz equation on a domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, for the field $u(\mathbf{x}): \Omega \rightarrow \mathbb{C}$ given by

$$-\Delta u - k^2 u = f \quad \text{in } \Omega, \quad (1a)$$

$$\mathcal{C}(u) = 0 \quad \text{on } \partial\Omega, \quad (1b)$$

where \mathcal{C} incorporates some appropriate boundary conditions, $k(\mathbf{x}) > 0$ is the wave number, and $f(\mathbf{x})$ is a suitable forcing function. A key parameter is the wave number k , which relates the angular frequency ω and the wave speed c as $k = \omega/c$. The wave speed $c(\mathbf{x})$ depends on the position \mathbf{x} in the media for heterogeneous problems. Since k is proportional to the frequency, the high frequency regime constitutes the case of large k and presents particular challenges for designing effective solvers.

The difficulty in designing a good solver for the Helmholtz equation is shown very clearly in the review papers [1, 2] where one can see that there are no straightforward extensions to state-of-the-art methods for symmetric positive definite problems that tackle the indefinite or non-self adjoint problem

well. When the problems are large, however,—the case when one discretises the Helmholtz equation accurately for high wave numbers—domain decomposition methods are a natural choice [3]. Nevertheless, despite recent efforts and in view of the latest results obtained both at the theoretical [4, 5, 6] or numerical level [7, 8, 9], there is no established method outperforming all others in the case of the Helmholtz problem.

Domain decomposition methods are well suited to solve large systems of equations arising from discretisation of PDEs and are among the best-known strategies for many types of problem. However, classical versions fail to be effective and may diverge for wave propagation problems. Two key constituent parts require a more careful treatment: the transmission conditions used to transfer information between adjoining subdomains and the coarse space that allows for capturing of global behaviour and passing information between subdomains globally. In this work we consider overlapping Schwarz methods.

The use of different transmission conditions at the interfaces (artificial boundaries arising from the decomposition into subdomains) has been extensively studied over the past two decades and various works [10, 11, 12] show that these conditions can improve the convergence of Schwarz methods and preconditioners. However, good transmission conditions are not sufficient to ensure a robust behaviour with respect to heterogeneities in the problem to solve or when the number of subdomains increases. To tackle these difficulties, we need coarse information that is cheap to compute and immediately available to all subdomains.

The focus in this work is on appropriate coarse spaces. A coarse space is typically required to provide scalability with respect to the number of subdomains used. More recently, however, coarse spaces have been designed to provide robustness to model parameters, especially for large contrasts in complex heterogeneous problems. For example, the GenEO (Generalised Eigenproblems in the Overlap) coarse space has been successfully employed for the robust solution of highly heterogeneous elliptic problems [13, 14]. For the Helmholtz equation, finding a suitable coarse space is not an easy task and, being an indefinite problem, choosing a larger coarse space need not improve performance [15]. In designing coarse spaces for Helmholtz problems, we might also wish to reduce the dependence of the domain decomposition method on the wave number k . A natural idea to capture global behaviour is to use plane waves as a basis for the coarse space but it is not clear that this is suitable for heterogeneous media. Plane waves were first used within

the multigrid approach [16] before later being used to build coarse spaces for domain decomposition methods. We can cite the example of FETI(-DP)-H methods [17, 18] but they have also been used in other domain decomposition methods [19]. Nonetheless, plane waves have mainly been employed for homogeneous problems and do not have a straightforward extension to the heterogeneous case.

Even if coarse space information needs to be global and available to all domains, coarse spaces can be built locally and based on local functions. Spectral coarse spaces use basis vectors deriving from the solution of local eigenvalue problems associated with appropriate operators. Within the context of the Helmholtz equation, this is exemplified in the DtN coarse space [20]. Here, eigenproblems are formulated on subdomain interfaces based on a Dirichlet-to-Neumann (DtN) map, extending an approach for elliptic problems [21, 22]. In this work we consider two spectral coarse spaces, the DtN approach and a GenEO-type approach suited to the Helmholtz problem. We will also consider a grid coarse space approach which utilises the addition of absorption in the problem.

Our consideration of coarse spaces for Helmholtz problems in the high frequency regime provide the following main contributions of the paper:

- We bring together and outline recent work on developing coarse spaces that can be utilised to enhance domain decomposition methods for the Helmholtz problem in heterogeneous media. These approaches are then implemented in a common software, namely FreeFEM.
- We discuss implementation details and practical aspects of the methods as well as contrasting the benefits and drawbacks.
- We provide extensive numerical results on several well-known benchmark problems in 2D and 3D and compare the different approaches within a variety of settings.
- Based on the results of our numerical tests and our understanding of the implementation aspects involved, we provide an outlook on scenarios where certain methods may be more, or less, favourable.

The outline for the remainder of this work is as follows. In Section 2 we detail the boundary value problem considered, i.e., the heterogeneous Helmholtz problem and its discretisation by finite elements. In Section 3 we

introduce the basic principles of domain decomposition methods and present two versions of the one-level method, namely RAS and ORAS. A second level, or coarse space, is usually added by deflation. The three different coarse space strategies, namely the grid coarse space, DtN coarse space, and GenEO coarse space, are introduced in Section 4. Parallel implementation details of these methods are given in Section 5 and an extensive numerical study is provided in Section 6. Conclusions are then given in Section 7.

2. The heterogeneous Helmholtz problem

Our model problem consists of solving the interior Helmholtz equation (1). To be concrete, we let Ω is a bounded polygonal domain and consider specific boundary conditions on $\Gamma = \partial\Omega$. We suppose Γ is partitioned into a disjoint union $\Gamma = \Gamma_D \cup \Gamma_N \cup \Gamma_R$ where Dirichlet conditions are imposed on Γ_D , Neumann conditions on Γ_N and a Robin condition on Γ_R . Namely, we wish to solve

$$-\Delta u - k^2 u = f \quad \text{in } \Omega, \quad (2a)$$

$$u = u_{\Gamma_D} \quad \text{on } \Gamma_D, \quad (2b)$$

$$\frac{\partial u}{\partial \mathbf{n}} = 0 \quad \text{on } \Gamma_N, \quad (2c)$$

$$\frac{\partial u}{\partial \mathbf{n}} + iku = 0 \quad \text{on } \Gamma_R, \quad (2d)$$

where u_{Γ_D} is known. The Robin condition in (2d) is a standard first order approximation to the far field Sommerfeld radiation condition and, in essence, enables appropriate wave behaviour to be described in a bounded domain, allowing for incoming or outgoing waves along Ω_R . We do not require that a problem instance includes all types of boundaries but note that if $\Gamma_R = \emptyset$ then the problem will be ill-posed for certain choices of k . Furthermore, when $\Gamma_R \neq \emptyset$ the resulting linear systems, while being complex symmetric, are not Hermitian and this will be important in our choice of iterative method. However, classical iterative methods on their own are not enough to be able to solve Helmholtz problems effectively [1]. This is further amplified when applied to highly heterogeneous problems.

The heterogeneity in our model is present in the wave number $k(\mathbf{x}) > 0$, being given by ratio of the angular frequency ω and the wave speed $c(\mathbf{x})$ as $k = \omega/c$. We allow k to have jumps across different media and otherwise vary within the domain Ω such that $k \in L^\infty(\Omega)$.

To discretise (2), we use the finite element method. In order to prescribe the weak formulation, we let $V = \{u \in H^1(\Omega) : u = u_{\Gamma_D} \text{ on } \Gamma_D\}$ and, in a similar fashion, $V_0 = \{u \in H^1(\Omega) : u = 0 \text{ on } \Gamma_D\}$. The weak form of the problem is then to find $u \in V$ such that

$$a(u, v) = F(v) \quad \forall v \in V_0, \quad (3)$$

where

$$a(u, v) = \int_{\Omega} (\nabla u \cdot \nabla \bar{v} - k^2 u \bar{v}) \, d\mathbf{x} + \int_{\Gamma_R} i k u \bar{v} \, ds, \quad (4a)$$

and

$$F(v) = \int_{\Omega} f \bar{v} \, d\mathbf{x}, \quad (4b)$$

are the bilinear and linear parts, respectively. To discretise, we consider piecewise polynomial finite element approximation on a simplicial mesh \mathcal{T}^h of Ω which has a characteristic element diameter h . Denoting the associated trial space $V^h \subset V$ and test space $V_0^h \subset V_0$, the discrete problem is to find $u_h \in V^h$ such that

$$a(u_h, v_h) = F(v_h) \quad \forall v_h \in V_0^h. \quad (5)$$

Let $\{\phi_j\}_{j=1}^n$ be the nodal basis for V_0 and $\{\phi_j\}_{j=n+1}^{n+d}$ be the nodal basis for the Dirichlet boundary Γ_D , for which \mathcal{T}^h is assumed to conform. Then we can rewrite (5) as a (complex) linear system

$$A\mathbf{u} = \mathbf{f}, \quad (6)$$

where the coefficient matrix $A \in \mathbb{C}^{n \times n}$ and right-hand side vector $\mathbf{f} \in \mathbb{C}^n$ are given by $A_{i,j} = a(\phi_j, \phi_i)$ and $\mathbf{f}_i = F(\phi_i) - \sum_{l=n+1}^{n+d} a(\phi_l, \phi_i) \bar{u}_{l-n}$ respectively, for $i, j \in 1, 2, \dots, n$. Here \bar{u}_j , for $j = 1, 2, \dots, d$, are the known Dirichlet values along Γ_D corresponding to u_{Γ_D} . We then seek the solution $\mathbf{u} \in \mathbb{C}^n$ of the (in general) complex symmetric indefinite system (6) to give

$$u_h(\mathbf{x}) = \sum_{j=1}^n u_j \phi_j(\mathbf{x}) + \sum_{l=n+1}^{n+d} \bar{u}_{l-n} \phi_l(\mathbf{x}). \quad (7)$$

The wave nature of solutions to the Helmholtz equation requires a sufficiently fine mesh in order to obtain a good approximation to the true solution and this should be kept in mind when considering the choice of discretisation of the problem. In terms of increasing the wave number k , if one is to maintain the same level of accuracy of discrete solutions then the number of grid points must increase faster than k increases, due to the pollution effect [23]. This growth depends on the discretisation chosen. For instance, in the case of using a piecewise linear (P1) finite element approximation on simplicial elements of diameter h , then $k^3 h^2$ must be bounded, requiring h to shrink as $\mathcal{O}(k^{-3/2})$. For piecewise quadratic (P2) finite elements on simplicial meshes the criteria relaxes to require that h decreases as $\mathcal{O}(k^{-5/4})$. For higher order finite elements, the requirement becomes less stringent but the interpolation properties when using such approximation spaces ultimately begin to degrade.

Due to these restrictions and the desire for faster simulation times, it is common that practitioners simply consider a fixed number of points per wavelength instead, resulting in h decreasing as $\mathcal{O}(k^{-1})$. Given a fixed number, n_{ppwl} , of points per wavelength we ensure that $n_{\text{ppwl}} h \approx \lambda$, where the wavelength is given by $\lambda = 2\pi k^{-1}$. A prevalent engineering practice is to use 10 points per wavelength and in some large real-world problems of interest, such as in imaging science, it may be adequate or necessary to insist on less resolution. In light of this, we consider both 5 and 10 points per wavelength scenarios and make use of standard P2 finite element approximation throughout this work.

3. Domain decomposition

We now give details of the overlapping domain decomposition approach that we will utilise. This will be applied as a preconditioner rather than a stand-alone iterative method. Our approach is based on a two-level version of the optimised restricted additive Schwarz (ORAS) method. To provide a domain decomposition, we first partition Ω into non-overlapping subdomains $\{\Omega'_s\}_{s=1}^N$ which are resolved by the mesh \mathcal{T}^h . A layer of adjoining mesh elements is then added to provide overlapping subdomains $\{\Omega_s\}_{s=1}^N$ through the extension

$$\Omega_s = \text{Int} \left(\bigcup_{\text{supp}(\phi_j) \cap \Omega'_s \neq \emptyset} \text{supp}(\phi_j) \right), \quad (8)$$

where $\text{Int}(\cdot)$ denotes the interior of a domain and $\text{supp}(\cdot)$ the support of a function. Note that more than one layer of elements can be added in a recursive manner if subdomains with larger overlap are required.

Now that we have a domain decomposition, we can define the restriction to a given subdomain Ω_s as an operator from V^h into $V^h(\Omega_s) = \{v|_{\Omega_s} : v \in V^h\}$, namely $\mathcal{R}_s : V^h \rightarrow V^h(\Omega_s)$ where $\mathcal{R}_s v = v|_{\Omega_s}$. Let $R_s \in \mathbb{R}^{n_s \times n}$ be the matrix form of \mathcal{R}_s where n_s is the number of degrees of freedom in Ω_s . Since our subdomains overlap, we also make use of a partition of unity having matrix form $D_s \in \mathbb{R}^{n_s \times n_s}$ which is diagonal and satisfies $\sum_{s=1}^N R_s^T D_s R_s = I$; this removes “double counting” in the additive Schwarz method. Note that R_s^T acts as an extension by zero outside of Ω_s .

We can now define the restricted additive Schwarz (RAS) preconditioner

$$M_{\text{RAS}}^{-1} = \sum_{s=1}^N R_s^T D_s A_s^{-1} R_s, \quad (9)$$

where $A_s = R_s A R_s^T$ is the local Dirichlet matrix on Ω_s , that is, Dirichlet conditions are implicitly assumed on $\partial\Omega_s \setminus \partial\Omega$. This corresponds to using Dirichlet transmission conditions in the additive Schwarz method. Note that each contribution from the sum in (9) can be computed locally in parallel.

The RAS approach, using the classical choice of Dirichlet transmission conditions, applied as a stand-alone stationary iterative method need not converge since frequencies in the error smaller than the wave number k are not diminished [3]. When used as a preconditioner, the iterative solver used will typically suffer from slow convergence and may stagnate. Further, the local Dirichlet problems involving A_s are not necessarily well-posed as k^2 may be an eigenvalue of the corresponding Laplace problem, in which case A_s is singular. While methods to handle such singular systems can be applied, a different approach, which provides a convergent stand-alone method, is to change the Dirichlet transmission conditions to Robin conditions. This results in the so-called optimised restricted additive Schwarz (ORAS) method given by

$$M_{\text{ORAS}}^{-1} = \sum_{s=1}^N R_s^T D_s \hat{A}_s^{-1} R_s, \quad (10)$$

where now \widehat{A}_s is the discretisation of the local Robin problem

$$-\Delta w_s - k^2 w_s = f \quad \text{in } \Omega_s, \quad (11a)$$

$$\mathcal{C}(w_s) = 0 \quad \text{on } \partial\Omega_s \cap \partial\Omega, \quad (11b)$$

$$\frac{\partial w_s}{\partial \mathbf{n}_s} + ikw_s = 0 \quad \text{on } \partial\Omega_s \setminus \partial\Omega, \quad (11c)$$

with \mathcal{C} representing the underlying problem boundary conditions on $\partial\Omega$. The local problem (11) will always have a unique solution w_s . We note that the use of Robin conditions is not the only remedial choice. Optimal transmission conditions can also be studied and are given using a Dirichlet-to-Neumann (DtN) map [3]. However, this results in requiring pseudodifferential operators and is somewhat less practical without further approximation. We do not follow the pursuit of more advanced transmission conditions in this work.

The above approaches, RAS and ORAS, are one-level methods: they rely only on local subdomain solves and the local transmission of data. Such domain decomposition methods do not scale as we increase the number of subdomains used; that is, their convergence behaviour depends on N . To achieve robustness with respect to N , a coarse space is typically included, giving a two-level method. The coarse space is represented by a collection of column vectors Z , having full column rank. A coarse space operator $E = Z^\dagger AZ$ is constructed as well as the coarse correction operator $Q = ZE^{-1}Z^\dagger$; note the similarity to terms in (9). The inclusion of the coarse space can be done in a number of ways, the simplest being additively as

$$M_{2\text{-level}}^{-1} = M_{1\text{-level}}^{-1} + Q, \quad (12)$$

where $M_{1\text{-level}}^{-1}$ is the underlying one-level preconditioner used, such as (9) or (10). Hybrid approaches are often more effective and we shall consider the adapted deflation technique

$$M_{2\text{-level}}^{-1} = M_{1\text{-level}}^{-1}P + Q = M_{1\text{-level}}^{-1}(I - AQ) + Q, \quad (13)$$

where $P = I - AQ$ is a projection. The most crucial choice is that of Z , which provides the coarse space. This is what we shall now consider.

4. Coarse spaces

The construction of a suitable coarse space can be achieved in different ways. One natural approach is to utilise a coarse grid in order to approximate

the global behaviour. Often, the slow convergence in the one-level method can be characterised by “slow modes” which should be incorporated into the coarse space. For instance, in the Poisson problem slow modes correspond to constant functions in the kernel of the Laplace operator in each subdomain and these are used to give the Nicolaides coarse space. On the other hand, for homogeneous elasticity problems the slow modes are rigid body motions.

The notion that certain modes are responsible for slow convergence and must be incorporated into the coarse space can be made more general, in particular in the framework of spectral coarse spaces. Such coarse spaces use spectral information from an appropriate eigenproblem to identify relevant modes that should feature in the coarse space. Before considering spectral coarse spaces we first outline the grid coarse space.

Remark 1 (Notation). Following on from Section 3, for a variational problem giving rise to a system matrix B we denote by B_s the corresponding local Dirichlet matrix on Ω_s . Where Robin conditions are used on internal subdomain interfaces (the artificial boundaries) the local problem matrix is denoted by \widehat{B}_s . Meanwhile, when Neumann conditions are used on such interfaces, the local matrix is denoted by \widetilde{B}_s .

4.1. The grid coarse space method

The grid coarse space was first introduced in [24] for the absorptive Helmholtz problem and extended to incorporate impedance (or Robin) conditions in [5]. In this case the one-level method is based on the following formula

$$M_{1,\varepsilon}^{-1} = \sum_{s=1}^N R_s^T D_s \widehat{A}_{s,\varepsilon}^{-1} R_s. \quad (14)$$

where matrices $\widehat{A}_{s,\varepsilon}$ stem from the discretisation of the following local Robin problems with absorption (given by the parameter $\varepsilon \neq 0$)

$$\begin{aligned} -\Delta u_s - (k^2 + i\varepsilon)u_s &= f && \text{in } \Omega_s, \\ \mathcal{C}(u_s) &= 0 && \text{on } \partial\Omega_s \cap \partial\Omega, \\ \frac{\partial u_s}{\partial \mathbf{n}_s} + iku_s &= 0 && \text{on } \partial\Omega_s \setminus \partial\Omega. \end{aligned}$$

In order to achieve weak dependence on the wave number k and number of subdomains N , the two-level preconditioner can be written in a generic way as follows

$$M_{2,\varepsilon}^{-1} = M_{1,\varepsilon}^{-1}P + ZE^{-1}Z^\dagger, \quad (15)$$

where $M_{1,\varepsilon}^{-1}$ is the one-level preconditioner (14), Z is a rectangular matrix with full column rank, $E = Z^\dagger \widehat{A}_\varepsilon Z$ is the so-called coarse grid matrix, $Q = ZE^{-1}Z^\dagger$ is the so-called coarse grid correction matrix, and $P = I - A_\varepsilon Q$.

Perhaps the most natural coarse space is the one based on a coarser mesh, which we call the “grid coarse space”. Let us consider $\mathcal{T}^{H_{\text{coarse}}}$, a simplicial mesh of Ω with mesh diameter H_{coarse} , and $V^{H_{\text{coarse}}} \subset V$, the corresponding finite element space. Let $\mathcal{I}_0: V^{H_{\text{coarse}}} \rightarrow V^h$ be the nodal interpolation operator and define Z as the corresponding matrix. Then, in this case, $E = Z^\dagger \widehat{A}_\varepsilon Z$ is really the stiffness matrix of the problem (with absorption) discretised on the coarse mesh. Related preconditioners without absorption are used in [19].

4.2. The DtN coarse space

The Dirichlet-to-Neumann (DtN) coarse space [21, 20] is based on solving local eigenvalue problems on subdomain boundaries related to the DtN map. To define this map for the Helmholtz problem we first require the Helmholtz extension operator from the boundary of a subdomain Ω_s .

Let $\Gamma_s = \partial\Omega_s \setminus \partial\Omega$ and suppose we have Dirichlet data v_{Γ_s} on Γ_s , then the Helmholtz extension v in Ω_s is defined as the solution of

$$-\Delta v - k^2 v = 0 \quad \text{in } \Omega_s, \quad (16a)$$

$$v = v_{\Gamma_s} \quad \text{on } \Gamma_s, \quad (16b)$$

$$\mathcal{C}(v) = 0 \quad \text{on } \partial\Omega_s \cap \partial\Omega, \quad (16c)$$

where $\mathcal{C}(v) = 0$ represents the original problem boundary conditions, as in (1b). The DtN map takes in the Dirichlet data v_{Γ_s} on Γ_s and gives as output the corresponding Neumann data, that is

$$\text{DtN}_{\Omega_s}(v_{\Gamma_s}) = \left. \frac{\partial v}{\partial \mathbf{n}_s} \right|_{\Gamma_s} \quad (17)$$

where v is the Helmholtz extension defined by (16).

We now seek eigenfunctions of the DtN map locally on each subdomain Ω_s , given by solving

$$\text{DtN}_{\Omega_s}(u_{\Gamma_s}) = \lambda u_{\Gamma_s} \quad (18)$$

for eigenfunctions u_{Γ_s} and eigenvalues $\lambda \in \mathbb{C}$. To provide functions to go into the coarse space, we take the Helmholtz extension of u_{Γ_s} in Ω_s and then extend by zero into the whole domain Ω using the partition of unity.

To formulate the discrete problem, we require the coefficient matrices \tilde{A}_s corresponding to local Neumann problems on Ω_s with boundary conditions $\mathcal{C} = 0$ on $\partial\Omega_s \cap \partial\Omega$, defined analogously to that of the local Robin problems in (11). Further, we need to distinguish between degrees of freedom on the boundary Γ_s and the interior of the subdomain Ω_s , as such we let Γ_s and I_s be the set of indices on the boundary and interior respectively. We also let

$$M_{\Gamma_s} = \left(\int_{\Gamma_s} \phi_j \phi_i \right)_{i,j \in \Gamma_s} \quad (19)$$

denote the mass matrix on the subdomain interface. Using standard block notation to denote submatrices of A_s and \tilde{A}_s the discrete DtN eigenproblem is

$$\left(\tilde{A}_{\Gamma_s, \Gamma_s} - A_{\Gamma_s, I_s} A_{I_s, I_s}^{-1} A_{I_s, \Gamma_s} \right) \mathbf{u}_{\Gamma_s} = \lambda M_{\Gamma_s} \mathbf{u}_{\Gamma_s}. \quad (20)$$

The Helmholtz extension of \mathbf{u}_{Γ_s} to degrees of freedom in I_s is then given by $\mathbf{u}_{I_s} = -A_{I_s, I_s}^{-1} A_{I_s, \Gamma_s} \mathbf{u}_{\Gamma_s}$. Letting \mathbf{u}_s denote the Helmholtz extension, the corresponding vector which enters the coarse space Z is $R_s^T D_s \mathbf{u}_s$. For further details and motivation behind the DtN eigenproblems see [20].

It remains to determine which eigenfunctions of (20) should go on to be included in the coarse space. Several selection criteria were investigated in [20] and it was clear that the best choice was to select eigenvectors corresponding to eigenvalues with the smallest real part. That is, we use a threshold on the abscissa $\eta = \text{Re}(\lambda)$ given by

$$\eta < \eta_{\max}, \quad (21)$$

where η_{\max} depends on $k_s = \max_{\mathbf{x} \in \Omega_s} k(\mathbf{x})$. In particular, the choice $\eta_{\max} = k_s$ is advocated in [20]. Note that for larger η_{\max} more eigenfunctions are included in the coarse space, increasing its size and the associated computational cost. Nonetheless, it was recently showed that taking a slightly larger

threshold $\eta_{\max} = k_s^{4/3}$ can be beneficial in certain cases in order to gain robustness to the wave number [25]. However, this only occurs for the homogeneous problem with sufficiently uniform subdomains. Since it is not necessarily known in advance how many eigenvalues are below the threshold and in order to make a fair comparison in our numerical tests, we will consider using a fixed number of eigenvectors per subdomain.

4.3. The GenEO coarse space

The GenEO (Generalised Eigenproblems in the Overlap) coarse space was derived in [13] to provide a rigorously robust approach for symmetric positive definite problems even in the presence of heterogeneities. The fundamental generalised eigenproblems on Ω_s at the variational level are given by

$$a_{\Omega_s}(u, v) = \lambda a_{\Omega_s^\circ}(\Xi_s(u), \Xi_s(v)) \quad \forall v \in V(\Omega_s), \quad (22)$$

where Ξ_s represents the action of the partition of unity on Ω_s and Ω_s° is the overlapping zone, that is the part of Ω_s which overlaps with any other subdomain. Here $a_D(\cdot, \cdot)$ stems from the underlying variational problem on the domain D , in particular with problem boundary conditions on $\partial\Omega$ and natural (Neumann) conditions on parts of ∂D internal to Ω . The particular form of eigenproblem in (22) arises naturally in the analysis of [13]. We also note that (22) possesses infinite eigenvalues when there exists $u \neq 0$ such that $a_{\Omega_s^\circ}(\Xi_s(u), \Xi_s(v)) = 0 \forall v \in V(\Omega_s)$ but $a_{\Omega_s}(u, v) \neq 0$ for some $v \in V(\Omega_s)$, for example when u is supported only outside Ω_s° .

The discrete form of the eigenproblem (22) is

$$\tilde{A}_s \mathbf{u} = \lambda D_s \tilde{A}_s^\circ D_s \mathbf{u} \quad (23)$$

where \tilde{A}_s° is the (Neumann) matrix built from assembling only over elements in the overlapping zone Ω_s° . The eigenfunctions selected to enter the coarse space are the low frequency modes, that is those corresponding to the smallest eigenvalues. Typically, either a fixed number of eigenfunctions are taken per subdomain or a threshold $\lambda < \lambda_{\max}$ on the corresponding eigenvalues is used, where λ_{\max} can be chosen based on problem parameters to achieve a specified condition number of the preconditioned system; see [13].

The precise formulation of the eigenproblem and use of the overlapping zone is somewhat flexible. In particular, the overlapping zone can be replaced

with the whole subdomain, as in [14] which utilises an eigenproblem of the form

$$\tilde{A}_s \mathbf{u} = \lambda D_s A_s D_s \mathbf{u}. \quad (24)$$

It is this form of GenEO that we shall build upon shortly. We note that two separate GenEO eigenproblems can also be formulated to provide bounds on both ends of the preconditioned spectrum, as is found when using a symmetric ORAS approach in [14]. This flexibility and robustness of GenEO-type methods has yet to be fully explored, especially for problems which are not symmetric positive definite where current theory breaks down. We now consider the utility of using GenEO as basis for constructing a coarse space for the heterogeneous Helmholtz problem.

4.4. A GenEO-type coarse space for the Helmholtz equation

In pursuing a GenEO approach for the Helmholtz equation we must first note the matrices \tilde{A}_s and A_s , now stemming from the bilinear form (4a), are no longer symmetric positive definite. As such, eigenvalues λ of the eigenproblem (24) are no longer real and positive (or infinite). As a threshold criterion, as with the DtN approach, we can consider the abscissa $\eta = \text{Re}(\lambda)$ instead and seek eigenfunctions corresponding to $\eta < \eta_{\max}$. Further, since the theory breaks down without the symmetric positive definite assumption, it is no longer clear that (24) provides appropriate eigenvectors for the coarse space. Indeed, applying out of the box the GenEO method using (24) fails to provide a satisfactory method.

Since GenEO is designed for positive definite problems, a natural proposal is to use a nearby positive definite problem in the formulation of the eigenproblem, namely a Laplace problem, that is setting $k = 0$ for the purpose of constructing the coarse space. Let L_s be local Dirichlet matrix corresponding to the discrete Laplacian and \tilde{L}_s the equivalent Neumann matrix on Ω_s , then

$$\tilde{L}_s \mathbf{u} = \lambda D_s L_s D_s \mathbf{u}. \quad (25)$$

has positive real eigenvalues. While this approach can perform reasonably well when the wave number k is small, the behaviour as k grows becomes increasingly poor, as might be expected given the coarse space is independent of k .

To incorporate k , we instead link the underlying Helmholtz problem to the positive definite Laplace problem to formulate a GenEO-type eigenproblem

$$\tilde{A}_s \mathbf{u} = \lambda D_s L_s D_s \mathbf{u}. \quad (26)$$

Eigenvalues of (26) are now, in general, complex (though we note they primarily appear to cluster close to the real line) and so we suggest to threshold based on the abscissa $\eta < \eta_{\max}$. We call this GenEO-type approach for the Helmholtz problem “H-GenEO”. Some initial exploration of this method can be found in [26], where the approach is seen to perform well for a 2D waveguide problem and provide robustness to heterogeneity as well as to the wave number k , albeit requiring a comparatively large coarse space. Again, in our numerical experiments we take a fixed number of eigenvectors and vary this quantity, aiming to give a relatively fair comparison of the spectral coarse space methods.

5. Implementation details

Numerical results within this paper have been obtained using FreeFEM [27]. More precisely, `ffddm`—a light-weight layer in FreeFEM domain specific language that generates all domain decomposition data structures—was used on top of HPDDM [28], which handles the underlying computations such as matrix–vector products or preconditioner applications. When it comes to spectral two-level methods, as depicted in the previous Sections 4.2 to 4.4, the bulk of the work comes from the local eigensolves, see, e.g., eqs. (20), (23) and (26). These are performed concurrently in each subdomain, using solvers such as ARPACK [29] or SLEPc [30]. Then, the construction of the algebraic Galerkin operator Q introduced in Section 3 is performed in HPDDM. In order to deal efficiently with both large numbers of subdomains and large numbers of local eigenvectors, the assembly and exact LU factorization of Q is performed on a subset of processes from the global communicator. We use MUMPS [31] both as a subdomain solver and a coarse operator solver in our experiments. Again, we refer interested readers to [28] for more details about this subject.

For the grid coarse space method from Section 4.1, two meshes of the same domain must be considered simultaneously. First a coarse one, see an example fig. 1a, and then a fine one, see examples figs. 1b and 1c. We generate fine meshes using a uniform refinement in which edges of all triangles or tetrahedra are divided uniformly in s . In the aforementioned example, $s = 2$. Note that if one wants $n_{\text{ppw1}}h \approx \lambda$ on the fine grid, on which the original algebraic system of equations eq. (6) stems from, then the coarse mesh has to satisfy $H_{\text{coarse}} = \frac{h}{s}$. When it comes to the underlying domain decomposition preconditioners, two options will be considered next. One

where both levels of overlap on fine and coarse grids are minimal, i.e. 1, see fig. 1b, and another where the level of overlap on the coarse grid is minimal, i.e. 1, and where the level of overlap on the fine grid is s , see fig. 1c. In the latter case, note that subdomains on the fine grid are merely uniformly refined subdomains of the coarse grid. While the setup cost of such a two-level preconditioner is much lower than for a spectral method as described in the previous paragraph, applying it to a vector, for example in a Krylov method, is usually costlier. Indeed, the coarse problem is in this case often solved iteratively, instead of using an exact factorization of the coarse operator. This outer–inner strategy, which also makes the use of flexible methods such as FGMRES [32] mandatory, may not perform very well prior to some tuning of the inner (coarse) solver.

In the numerical experiments presented in this paper, the inner coarse problem in the grid coarse space method is defined with a splitting level $s = 2$ and is solved approximately with GMRES preconditioned by a one-level method, with a tolerance of 0.1. Additionally, in the spirit of “shifted Laplacian preconditioning” (see [33, 34]), the coarse problem is defined with added absorption in the equation. This improves the convergence of the one-level method, which then requires fewer iterations to reach the prescribed inner tolerance. Two-level domain decomposition preconditioners with added absorption have been proposed in the literature for the Helmholtz [35] and Maxwell [36] equations. The amount of added absorption needs to be chosen carefully: if the amount of added absorption is too small, there is no gain in the convergence of the inner coarse problem. On the other hand, if the amount is too large, the coarse problem becomes a bad approximation of the original problem (without absorption), and the number of outer iterations increases. We choose the additional imaginary term to be proportional to the wave number k as a compromise.

We conclude this section by mentioning that all the spectral preconditioners may be used in PETSc [37] through the PCHPDDM preconditioner [38]. When it comes to the grid coarse space method, we have a custom implementation in HPDDM that handles systems with multiple right-hand sides, which are encountered frequently in wave propagation phenomena [39, 40, 7]. PCMG, the PETSc machinery for geometric multigrid, does not handle such systems as of version 3.14.1.

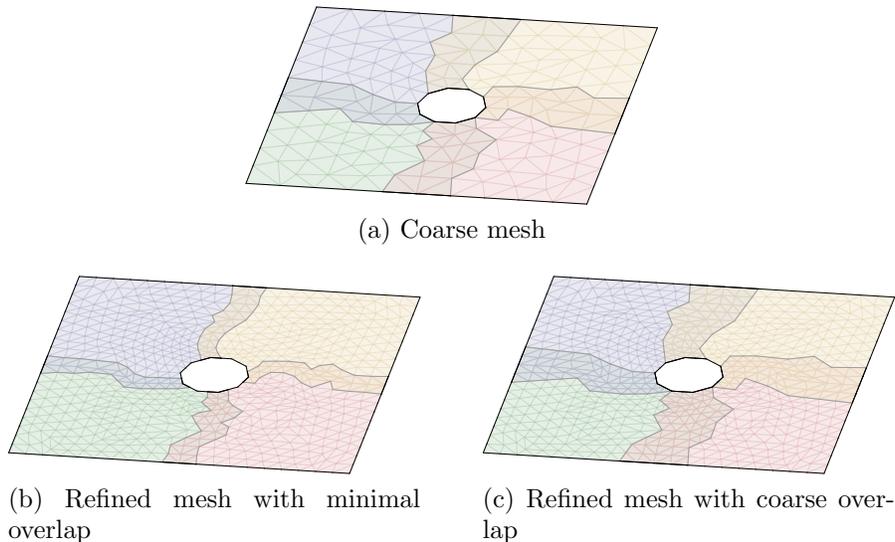


Figure 1: Different overlap definitions (bottom) for a uniformly refined simple 2D coarse mesh (top).

6. Comparative numerical studies

In this section we compare the domain decomposition preconditioners for several challenging heterogeneous test problems. We include problems in both 2D and 3D and vary different characteristics of the problem to understand how they can affect performance. In particular, we are interested in the behaviour for larger wave numbers (or equivalently larger frequencies) and for an increasing number of subdomains within the domain decomposition. We will also consider discretisations with differing number of points per wavelength and will see how this can affect performance of the preconditioners. For the approaches that use a spectral coarse space, we further investigate the choice of how many eigenvectors should be taken per subdomain.

In the results that follow, a tabulated value of $-$ signifies that the particular test instance did not run, typically this is for the smallest problem when trying to solve on too many subdomains. When the maximum number of iterations is reached before convergence, we denote this by \times . Depending on the test problem, the maximum number of iterations is set at either 500 or 1000.

6.1. The Marmousi problem

The Marmousi problem, see [41], has become a benchmark geophysics test case for the assessment of numerical methods and solvers when used in a direct or inverse problem context. It features a 2D rectangular domain with velocity data modelling a heterogeneous subsurface. We utilise this as the wave speed within our Helmholtz model with a point source placed close to the surface; see the example solution in Figure 2. The problem is posed on the domain $(0, 9.2) \times (-3, 0)$ with homogeneous Dirichlet conditions imposed on the top (surface) boundary and Robin conditions on the remaining (subsurface) boundaries, the latter being artificial boundaries set for computational purposes.

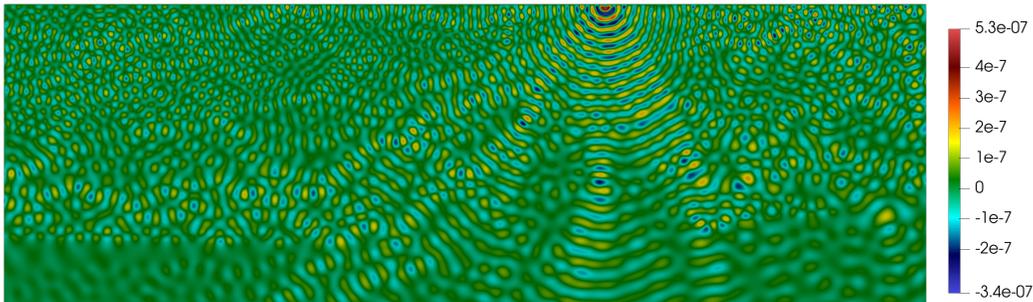


Figure 2: The real part of the solution to the Marmousi problem at a frequency of 20 Hz.

The difficulty of this test case stems both from the heterogeneous nature of the problem but also from the presence of a substantial number of wavelengths and for this reason, the problem is considered to be high-frequency. Therefore, it is important to have a sufficiently accurate discretisation, keeping in mind that for inverse problems the accuracy requirements are different (usually 5 points per wavelength) and weaker than in the case of only solving the forward problem (in common engineering practice, we use 10 points per wavelength).

We first consider discretisation with 5 points per wavelength. In Table 1 we show results for the one-level method with minimal overlap and also with minimal overlap from the coarse grid (twice the overlap—see Figure 1) as references to compare to. Further, Table 1 gives results for the two-level coarse grid approach with minimal overlap from the coarse grid. We see that both cases of the one-level method are not scalable (iteration counts increase as the number of subdomains N increases) and perform worse as

the frequency f becomes higher. On the other hand, the two-level coarse grid approach is reasonably scalable and with only mildly increasing iteration counts as the frequency increases.

Table 1: Results using the one-level and coarse grid methods for the Marmousi problem when using 5 points per wavelength, varying the frequency f and the number of subdomains N .

$f \setminus N$	One-level (min. overlap)					One-level (coarse overlap)					Coarse grid				
	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
1	30	46	78	98	–	26	39	47	64	–	15	18	19	20	–
5	57	82	117	170	236	53	76	105	154	213	26	29	28	29	31
10	75	111	173	232	330	68	102	158	212	302	32	35	41	40	42
20	89	133	193	268	373	82	125	178	248	347	34	35	42	43	44

We now turn to the two spectral coarse space methods with results in Table 2 being given for varying number of eigenvectors taken per subdomain ν . We first see that the H-GenEO method within this scenario performs poorly and, without sufficiently many eigenvectors being taken, often exhibits iteration counts worse than the one-level method. On the other hand, the DtN method generally converges faster than the one-level method and, with enough eigenvectors, can give reasonable performance. However, as the frequency increases the approach begins to struggle a little with larger iteration counts typically required for convergence even with more eigenvectors being utilised. Nonetheless, for small frequencies the approach can beat the coarse grid method with relatively few eigenvectors per subdomain and even for the larger frequency tested, when using the most subdomains and sufficient eigenvectors per subdomain the DtN approach converges faster than the coarse grid method. We note by way of caution though that taking too many eigenvectors per subdomain can eventually lead to some deterioration in the performance of the DtN method. We can also see from Table 2 that the larger coarse overlap is beneficial for the DtN method and primarily yields smaller iterations counts. On the other hand, minimal overlap appears preferable for H-GenEO.

Overall, while the DtN approach can reduce the number of iterations required for convergence compared to the coarse grid method in some cases, particularly low frequencies, the coarse grid approach exhibits greater robustness and requires relatively low iterations counts, suggesting it is the favourable approach in this scenario of using 5 points per wavelength.

We now consider the case of discretisation with 10 points per wavelength and similarly compare the different approaches. In Table 3 we show results for the one-level methods along with the coarse grid method. We notice that the iteration counts in this situation are slightly higher than when using 5 points per wavelength but otherwise the picture remains similar with coarse grid approach giving a reasonably robust method albeit with some slow growth in iteration counts as the frequency is increased.

Table 3: Results using the one-level and coarse grid methods for the Marmousi problem when using 10 points per wavelength, varying the frequency f and the number of subdomains N .

$f \setminus N$	One-level (min. overlap)					One-level (coarse overlap)					Coarse grid				
	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
1	34	49	72	143	–	30	43	63	97	–	16	18	19	21	–
5	62	94	137	191	268	58	87	126	175	246	29	29	34	34	36
10	85	136	185	272	371	78	124	172	251	346	35	41	43	46	45
20	101	152	213	299	419	92	142	198	272	389	39	47	48	49	49

Turning to the spectral coarse space approaches, we see in Table 4 that H-GenEO now becomes competitive. For larger frequencies and equivalent numbers of eigenvectors taken per subdomain, the H-GenEO method generally requires fewer iterations and this is particularly true for the large frequencies with many subdomains. We note that the minimal overlap case is preferable for H-GenEO while for DtN there is no longer a clear-cut preference between the coarse overlap and minimal overlap cases. Provided sufficiently many eigenvectors are taken per subdomain, we see that H-GenEO now outperforms the coarse grid approach in terms of iterations counts and is particularly strong for the large frequency and many subdomain situation.

Overall, in the scenario of 10 points per wavelength, we have different findings from the 5 point per wavelength scenario. Now, while the coarse grid approach is still relatively robust, the H-GenEO method can be seen as strong alternative which is able to markedly reduce the number of iterations required in the case of large frequencies and many subdomains. One possible explanation is that spectral information is not relevant or sufficiently well approximated in the case of a low number of points per wavelength.

6.2. The COBRA cavity

The COBRA cavity problem consists of a plane wave incident upon a curving cavity aperture and scattering in 3D. The geometry is shown within an example solution in Figure 3. The problem was devised by EADS Aerospaciale Matra Missiles for Workshop EM-JINA 98 and is described in [42, 43]. The COBRA cavity problem is a benchmark problem in electromagnetism (when we use the time-harmonic Maxwell’s equations as the underlying model) but here we present the Helmholtz version. The main difficulty comes from the presence of metallic curved waveguide which can cause trapping in addition to the inherent difficulties present in mid-high frequency regime wave propagation problems. This problem is naturally three-dimensional but a similar two-dimensional test can be designed making use of a cross-section of the cavity.

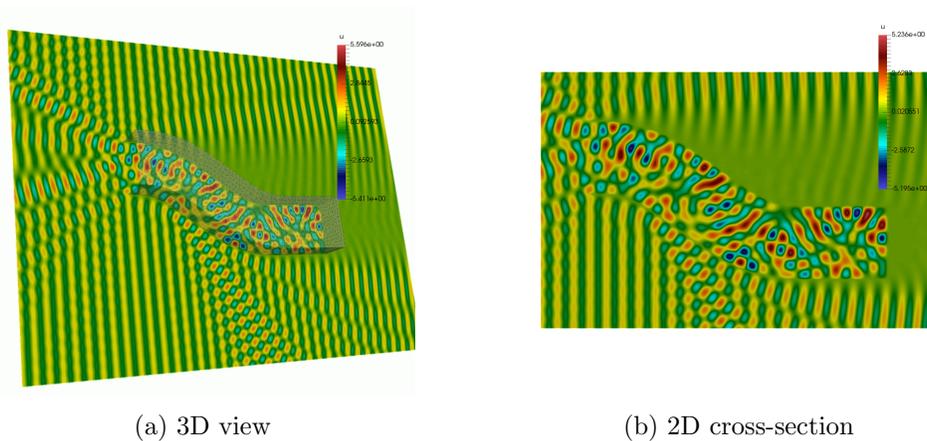


Figure 3: The COBRA cavity problem.

We first consider discretisation of the problem with 5 points per wavelength using P2 finite elements. In Table 5 we display results for the underlying one-level method and the two-level coarse grid approach. The one-level method performs poorly as expected, however, we also see that the coarse grid approach lacks robustness in this setting and performs poorly for larger wave numbers.

Table 5: Results using the one-level and coarse grid methods for the COBRA cavity problem when using 5 points per wavelength, varying the wave number k and the number of subdomains N .

$k \setminus N$	One-level				Coarse grid			
	20	40	80	160	20	40	80	160
100	71	83	104	–	21	21	22	–
200	161	203	259	303	87	90	91	91
300	326	429	574	727	277	309	330	562
400	326	426	581	705	375	450	463	471

Comparing now the spectral coarse space methods, in Table 6 we observe that, in this scenario, the H-GenEO method performs poorly and is often worse than the one-level method, even when taking fairly large numbers of eigenvectors per subdomain. The DtN method can also perform worse than the one-level method in some cases without sufficiently many eigenvectors being taken. However, the DtN approach begins to show relatively low iteration counts once enough eigenvectors are utilised, especially in the many subdomain cases, and can significantly reduce the iteration counts compared to the one-level or coarse grid approaches. Nonetheless, it appears for this 3D problem that the number of eigenvectors required in order to drive down the overall iteration count is rather high.

Table 6: Results using the H-GenEO and DtN spectral coarse space methods for the COBRA cavity problem when using 5 points per wavelength, varying the wave number k , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$k \setminus N$	H-GenEO				DtN			
		20	40	80	160	20	40	80	160
20	100	66	86	107	–	47	58	74	–
	200	165	214	284	366	154	199	250	324
	300	334	453	611	809	326	429	569	742
	400	342	440	624	810	337	425	582	727
40	100	65	85	108	–	32	34	28	–
	200	164	216	291	377	145	187	233	262
	300	348	468	638	852	319	424	567	768
	400	366	500	716	981	337	438	594	778
80	100	64	85	108	–	16	12	8	–
	200	155	214	285	382	116	132	99	63
	300	345	461	636	845	301	399	542	573
	400	408	558	771	×	335	453	650	918
160	100	63	85	108	–	7	5	5	–
	200	150	210	282	382	46	41	17	11
	300	324	441	617	830	247	263	249	144
	400	417	547	749	×	350	476	576	483
320	100	–	–	–	–	–	–	–	–
	200	145	207	281	383	10	10	50	55
	300	308	422	601	817	100	60	45	29
	400	401	527	730	994	271	233	147	44

We also consider the case of discretising the problem with 10 points per wavelength, again using P2 finite elements. Results in Table 7 show that in this case the coarse grid method performs very well and consistently leads to a small number of iterations being required for convergence; this is in stark contrast to the 5 points per wave case. For the spectral coarse spaces, as detailed in Table 8, the H-GenEO method remains poor, especially for large wave numbers, while the DtN method can perform reasonably well when enough eigenvectors are employed. In particular, the DtN method can require fewer iterations than the coarse grid method, especially for smaller frequencies, but in general struggles to compete in line with the coarse grid approach in terms of robustness.

Table 7: Results using the one-level and coarse grid methods for the COBRA cavity problem when using 10 points per wavelength, varying the wave number k and the number of subdomains N .

$k \setminus N$	One-level				Coarse grid			
	20	40	80	160	20	40	80	160
50	27	38	47	52	8	8	8	9
100	80	103	115	147	11	23	11	11
150	143	181	235	292	16	16	17	17
200	192	268	308	427	15	15	16	16

Table 8: Results using the H-GenEO and DtN spectral coarse space methods for the COBRA cavity problem when using 10 points per wavelength, varying the wave number k , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$k \setminus N$	H-GenEO				DtN			
		20	40	80	160	20	40	80	160
20	50	26	31	41	46	19	22	20	20
	100	73	94	111	130	63	77	87	112
	150	128	171	230	278	120	157	209	244
	200	186	264	315	427	178	246	283	380
40	50	25	31	40	45	13	12	10	8
	100	72	89	109	132	51	58	66	57
	150	129	168	228	279	112	130	181	211
	200	197	279	328	433	163	225	251	330
80	50	24	31	40	44	6	6	5	4
	100	68	89	106	132	28	28	19	10
	150	131	167	224	281	87	101	109	74
	200	200	270	327	428	142	193	218	255
160	50	24	29	40	44	5	4	4	4
	100	65	87	109	132	11	9	8	5
	150	126	164	220	279	50	31	21	14
	200	192	267	324	430	109	134	83	41
320	50	23	30	40	44	6	6	5	5
	100	63	85	107	133	6	6	6	5
	150	119	158	219	279	14	9	8	7
	200	182	246	311	430	44	28	15	11

6.2.1. The COBRA cavity in 2D

While the COBRA cavity is a 3D problem, we can additionally consider a 2D cross section, as illustrated in Figure 3b, and formulate the problem on such a slice. This allows us to consider a cavity problem in 2D and further consider higher wave numbers.

Results for this 2D COBRA cavity problem at 5 points per wavelength are given in Tables 9 and 10. Similarly to the full 3D problem, we see that the coarse grid method again lacks robustness in this scenario as the wave number k increases and the H-GenEO is unable to provide any benefit over the one-level method. The DtN approach is able to provide the lowest iteration counts, even with relatively few eigenvectors per subdomain, and provides a very effective method with more eigenvectors.

Table 9: Results using the one-level and coarse grid methods for the 2D COBRA cavity problem when using 5 points per wavelength, varying the wave number k and the number of subdomains N .

$k \setminus N$	One-level				Coarse grid			
	20	40	80	160	20	40	80	160
200	74	99	122	–	30	30	30	–
400	103	155	211	279	67	69	70	70
600	204	303	429	668	145	155	158	160
800	184	287	392	559	197	214	218	219
1000	188	278	400	551	242	265	269	270

Table 10: Results using the H-GenEO and DtN spectral coarse space methods for the 2D COBRA cavity problem when using 5 points per wavelength, varying the wave number k , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$k \setminus N$	H-GenEO				DtN			
		20	40	80	160	20	40	80	160
20	200	59	85	122	–	24	29	17	–
	400	101	135	202	277	78	87	77	40
	600	177	271	371	531	161	228	254	221
	800	192	293	390	517	179	274	341	371
	1000	213	302	428	597	203	288	415	521
40	200	54	83	107	–	6	7	6	–
	400	94	126	192	271	30	16	9	6
	600	164	240	348	505	84	82	38	16
	800	178	259	356	482	129	152	103	34
	1000	203	275	394	555	166	198	195	102
80	200	53	76	34	–	3	5	3	–
	400	86	120	187	261	7	5	4	3
	600	154	226	331	490	15	8	5	8
	800	164	244	338	466	37	22	8	5
	1000	184	256	370	529	74	36	14	7
160	200	53	28	1	–	2	4	1	–
	400	83	119	183	43	3	3	3	3
	600	146	218	322	480	4	3	3	6
	800	154	228	333	458	6	5	3	4
	1000	177	243	358	520	9	5	4	3
320	200	38	1	1	–	2	1	1	–
	400	81	117	43	1	3	2	3	1
	600	138	211	315	278	3	3	3	5
	800	147	221	326	448	3	3	3	3
	1000	169	233	346	512	4	3	3	3

The case of 10 points per wavelength is covered in Tables 11 and 12 and, again, we broadly see a picture akin to the full 3D problem. In this setting the coarse grid approach provides good performance but the DtN method yields lower iteration counts when sufficiently many eigenvectors are taken, again giving a very effective approach for this problem.

Table 11: Results using the one-level and coarse grid methods for the 2D COBRA cavity problem when using 10 points per wavelength, varying the wave number k and the number of subdomains N .

$k \setminus N$	One-level				Coarse grid			
	20	40	80	160	20	40	80	160
200	85	124	171	239	10	11	11	11
400	137	184	270	356	11	11	12	12
600	253	329	478	705	20	20	20	21
800	203	323	448	611	18	18	18	18
1000	217	316	470	656	19	19	19	19

Table 12: Results using the H-GenEO and DtN spectral coarse space methods for the 2D COBRA cavity problem when using 10 points per wavelength, varying the wave number k , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$k \setminus N$	H-GenEO				DtN			
		20	40	80	160	20	40	80	160
20	200	77	114	159	221	36	42	20	16
	400	136	195	273	367	111	147	172	123
	600	241	332	454	658	208	265	346	434
	800	220	338	485	692	186	282	396	533
	1000	242	363	561	787	202	293	419	608
40	200	75	108	153	220	10	7	7	5
	400	137	187	260	354	66	46	21	13
	600	236	320	439	632	152	167	128	51
	800	221	349	470	673	159	236	270	182
	1000	254	378	545	785	186	269	365	394
80	200	73	104	152	219	4	4	4	4
	400	130	174	256	347	11	8	6	5
	600	220	295	414	590	54	22	11	8
	800	213	329	454	642	88	74	25	13
	1000	256	365	528	727	141	147	80	22
160	200	71	104	148	180	2	3	4	3
	400	121	165	246	342	5	4	4	3
	600	210	282	396	577	9	5	5	3
	800	203	308	435	601	13	8	6	4
	1000	245	344	501	688	27	14	8	6
320	200	69	103	137	1	2	3	4	1
	400	116	158	239	338	3	3	4	3
	600	197	266	375	568	4	3	3	3
	800	192	295	412	579	5	4	3	3
	1000	225	324	464	653	6	5	4	4

6.3. The Overthrust problem

We will now assess our methods on the 3D $20 \times 20 \times 4.65$ km SEG/EAGE Overthrust model, as introduced and detailed in [44]. We use a homogeneous Dirichlet boundary condition at the surface and first-order absorbing boundary conditions along the other five faces of the model. The source is located at $(2.5, 2.5, 0.58)$ km and we perform a simulation with P2 finite elements on tetrahedral meshes for 0.5 Hz, 1 Hz and 2 Hz frequencies. For higher frequencies and large-scale computations on the Overthrust problem utilising

the grid coarse space see [7] and [8]. The Overthrust model can be seen as the three-dimensional counterpart of the Marmousi model and presents the difficulty of the latter (heterogeneities and high frequencies) at a larger scale since we have an extra space dimension.

We will start by testing the different approaches when using 5 points per wavelength, varying the frequency f and the number of subdomains N . Results with the one-level and coarse grid methods are shown in Table 13 while results for the H-GenEO and DtN spectral coarse space methods, varying the number of eigenvectors used per subdomain ν , are provided in Table 14. We see that in the case of low resolution the grid coarse space proves to be robust. In this scenario the H-GenEO method fails to bring improvement over the one-level method. On the other hand, the DtN approach performs well provided sufficiently many eigenvectors are included and can generally give an improvement over the coarse grid method in terms of reducing the iteration counts.

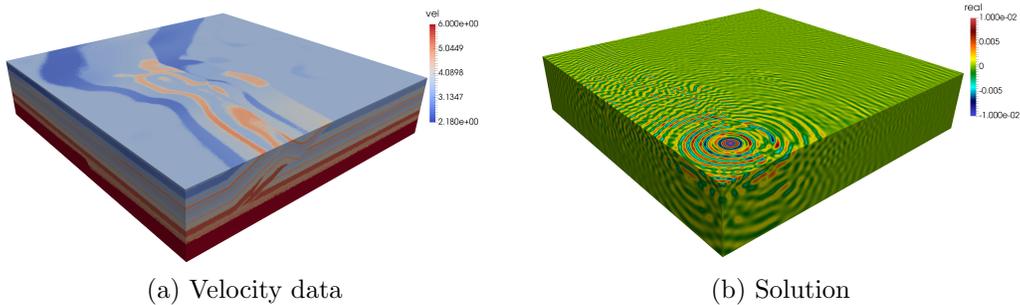


Figure 4: The Overthrust problem.

Table 13: Results using the one-level and coarse grid methods for the Overthrust problem when using 5 points per wavelength, varying the frequency f and the number of subdomains N .

$f \setminus N$	One-level (min. overlap)					One-level (coarse overlap)					Coarse grid				
	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
0.5	16	21	28	35	—	14	17	20	19	—	9	10	10	10	—
1	27	36	55	62	78	24	30	48	50	62	17	19	23	23	25
2	32	47	64	81	104	27	40	53	68	86	18	22	25	27	28

We consider now the case of higher resolution, using 10 points per wave-

length with results for smaller frequencies being given in Tables 15 and 16. Here we see that all two-level methods can perform relatively well and both H-GenEO and DtN can yield lower iteration counts than the coarse grid method without needing too many eigenvectors. If we allow for many eigenvectors then the DtN method typically requires the fewest iteration, however, we also note that for fewer eigenvectors the H-GenEO provides a greater reduction in iteration counts than DtN in the $f = 1$ Hz case and, furthermore, relatively little is gained when adding additional eigenvectors.

Table 15: Results using the one-level and coarse grid methods for the Overthrust problem when using 10 points per wavelength, varying the frequency f and the number of subdomains N .

$f \setminus N$	One-level (min. overlap)					One-level (coarse overlap)					Coarse grid				
	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
0.5	19	25	33	39	47	16	21	27	30	36	12	13	14	15	16
1	31	45	61	75	95	26	39	51	62	78	18	24	25	27	27

Finally, we examine a higher frequency case of $f = 2$ Hz where the number of degrees of freedom is now 11,334,869 and, as such, we consider utilising a larger number of subdomains N . Results for the grid coarse space are given in Table 17 and show again that the method performs well; for reference, we also report in parentheses the average number of inner iterations required for solving the coarse problem. For the spectral coarse spaces, results are given in Table 18 where we see a similar picture emerge as in the $f = 1$ Hz case. When fewer eigenvectors are used, the H-GenEO approach can converge faster than the DtN method. When more eigenvectors are used, the DtN approach becomes the method that converges fastest and can do so in fewer iterations than the grid coarse space; nonetheless, the H-GenEO method also continues to improve this time and performs similarly to the grid coarse space in terms of iteration counts. Note that the missing entries in Table 18, for the largest coarse space to be computed, are due to fact that the coarse operator (having dimension 1280×320 and dense blocks of size 320) is costly to compute and cannot be handled with current blackbox direct solvers.

Table 17: Results using the one-level and coarse grid methods for the Overthrust problem when using 10 points per wavelength for $f = 2$ Hz, varying the number of subdomains N . The number of degrees of freedom is 11,334,869.

$f \setminus N$	One-level (min. overlap)			One-level (coarse overlap)			Coarse grid		
	320	640	1280	320	640	1280	320	640	1280
2	149	185	226	125	156	189	24 (7)	24 (9)	23 (13)

Table 18: Results using the H-GenEO and DtN spectral coarse space methods for the Overthrust problem when using 10 points per wavelength for $f = 2$ Hz, varying the number of subdomains N and the number of eigenvectors used per subdomain ν . The number of degrees of freedom is 11,334,869.

$\nu \setminus N$	H-GenEO (min. overlap)			DtN (min. overlap)			H-GenEO (coarse overlap)			DtN (coarse overlap)		
	320	640	1280	320	640	1280	320	640	1280	320	640	1280
40	69	73	77	102	97	106	170	209	234	103	146	187
80	29	39	41	66	70	64	149	206	245	123	124	59
160	21	28	30	50	61	53	118	189	249	63	17	6
320	18	26	—	42	33	—	97	161	—	9	7	—

6.4. Discussion of numerical results

Numerical assessment of the three coarse spaces considered shows that, depending on the physical parameters (low or high frequency) or on the resolution considered (low resolution at 5 points per wavelength or high resolution at 10 points per wavelength), there is no systematic advantage of one method over another, as shown in the overview in Table 19. The underlying problem characteristics (spatial heterogeneities, low or high frequency) can play a strong role in the performance of each of these methods and each method has potential tuning parameters that can affect their efficacy, such as the splitting parameter s for the grid coarse space (which is maintained throughout these tests at its best value, namely 2, which gives the finest possible coarse space), or the number of local eigenvectors requested for the spectral coarse spaces. The resolution, in turn, can be dictated by the fact that the linear solver is used in the context of an inverse or a direct problem. Further, very often in geophysical computations, especially in full waveform inversion (FWI), one needs to solve repeatedly a linear system with multiple right-hand sides. In these situations, the additional precomputation required by the spectral coarse spaces becomes less of a burden compared to the inner

iterations required to solve the coarse problems when using the grid coarse space. Although apparently more expensive as they require solving local eigenvalue problems, spectral methods are potentially very promising in the context of inverse problems.

Nonetheless, we can provide some general observations from our numerical tests. The coarse grid approach is overall fairly robust in terms of the outer iteration counts required and for high wave numbers, so long as enough points per wavelength are used for discretisation. This method tended to be most favourable in the 10 points per wavelength scenarios, which is considered to be too high a level of precision in large-scale geophysical computations as it leads to a substantial, often prohibitive, number of degrees of freedom. The H-GenEO approach, while able to give some positive results for the geophysical test cases at sufficient resolution, failed to be robust for the cavity problems tested, suggesting it may only be a suitable solver method for certain types of problems. Nonetheless, in some scenarios it provides a greater reduction in iteration counts compared to the DtN method when only a relatively small number of eigenvectors are available. The DtN approach performed relatively well in our tests, arguably being the most robust method overall, so long as sufficiently many eigenvectors were utilised and, in particular, tended to be the favourable method in the 5 points per wavelength scenarios. Nonetheless, for the large 3D tests the method started to struggle somewhat unless a large number of eigenvectors were computed, but in this case the approach ultimately becomes rather memory demanding and it is challenging to deal with the coarse operator, which has large dense blocks.

Table 19: An overview of which coarse spaces perform well in the different problem scenarios tests. A ✓ indicates that the method performs well, with ✓✓ indicating this method was most favourable in a particular instance (note that for the Overthrust problem it is difficult to say that any one method is uniformly the most favourable in each case and so we omit this). A ✗ indicates that a method provided relatively little to no gain over the one-level method.

Problem	d	freq	Coarse grid		H-GenEO		DtN	
			5 ppwl	10 ppwl	5 ppwl	10 ppwl	5 ppwl	10 ppwl
Marmousi	2D	low	✓	✓	✓	✓	✓✓	✓✓
		high	✓✓	✓	✗	✓✓	✓	✓
COBRA Cavity	2D	low	✓	✓	✗	✗	✓✓	✓✓
		high	✗	✓	✗	✗	✓✓	✓✓
COBRA Cavity	3D	low	✓	✓✓	✗	✗	✓✓	✓
		high	✗	✓✓	✗	✗	✓✓	✓
Overthrust	3D	low	✓	✓	✗	✓	✓	✓
		high	✓	✓	✗	✓	✓	✓

7. Conclusions

We have presented in this work a comparison of several two-level overlapping Schwarz methods on a number of challenging problems of interest from the literature for the Helmholtz problem. Our results illustrate that each of these methods have pros and cons depending on the problem at hand and the particular numerical setting. In particular, no method establishes itself as the superior choice within a wide range of settings, though in particular cases we observe that a certain approach can be more favourable.

Note that our tests on these well-known benchmark problems can be further built upon and the conclusions refined. We have not measured timings or considered optimised implementations of the eigenvalue solvers and, in the case of the grid coarse space, the coarse grid problem is the finest possible which further requires another domain decomposition method to solve it as an inner iteration. This study is a first step in the assessment of the state of the art for two-level solvers aimed at time-harmonic wave propagation problems. Additional studies can be directed towards both the larger numerical context (taking into account higher approximation order, for example) and further applicative contexts (the solution of inverse problems in FWI, or the time-harmonic model of elastic waves which is at the same time more complex and more physically accurate). We remark that this is predominantly

a numerical study and, because of the mathematical difficulty of the models involved, a theoretical insight appears to remain out of reach, with a particular challenge being to cover the complexity of all practical situations of interest. Nonetheless, a general conclusion to be drawn is that coarse spaces are clearly a key feature in achieving scalability and robustness and, further, that the heuristics required for time-harmonic wave propagation problems are necessarily very different from the established methodology in the case of symmetric and positive definite problems.

References

- [1] O. G. Ernst, M. J. Gander, Why it is difficult to solve Helmholtz problems with classical iterative methods, in: I. G. Graham, T. Y. Hou, O. Lakkis, R. Scheichl (Eds.), *Numerical Analysis of Multiscale Problems*, Springer, 2012, pp. 325–363.
- [2] M. J. Gander, H. Zhang, A class of iterative solvers for the Helmholtz equation: Factorizations, sweeping preconditioners, source transfer, single layer potentials, polarized traces, and optimized Schwarz methods, *SIAM Review* 61 (1) (2019) 3–76.
- [3] V. Dolean, P. Jolivet, F. Nataf, *An introduction to domain decomposition methods*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2015, algorithms, theory, and parallel implementation.
URL <http://dx.doi.org/10.1137/1.9781611974065.ch1>
- [4] I. G. Graham, E. A. Spence, E. Vainikko, Recent results on domain decomposition preconditioning for the high-frequency Helmholtz equation using absorption, Lahaye D., Tang J., Vuik K. (eds) *Modern Solvers for Helmholtz Problems*. Geosystems Mathematics. Birkhäuser, Cham (2017) 3–26.
- [5] I. G. Graham, E. A. Spence, J. Zou, Domain decomposition with local impedance conditions for the Helmholtz equation, arXiv:1806.03731 (2018).
- [6] S. Gong, I. G. Graham, E. A. Spence, Domain decomposition preconditioners for high-order discretisations of the heterogeneous Helmholtz equation, arXiv:2004.03996 (2020).

- [7] V. Dolean, P. Jolivet, P.-H. Tournier, S. Operto, Iterative frequency-domain seismic wave solvers based on multi-level domain-decomposition preconditioners, in: 82th Annual EAGE Meeting (Amsterdam), Vol. arXiv:2004.06309, 2020.
- [8] V. Dolean, P. Jolivet, P.-H. Tournier, S. Operto, Large-scale frequency-domain seismic wave modeling on h -adaptive tetrahedral meshes with iterative solver and multi-level domain-decomposition preconditioners, in: SEG2020 Annual Meeting, 2020, arXiv:2004.07930.
- [9] A. Jouadé, A. Barka, Massively parallel implementation of FETI-2LM methods for the simulation of the sparse receiving array evolution of the GRAVES radar system for space surveillance and tracking, IEEE Access 7 (2019) 128968–128979.
- [10] P. Chevalier, F. Nataf, Symmetrized method with optimized second-order conditions for the Helmholtz equation, in: Domain decomposition methods, 10 (Boulder, CO, 1997), Amer. Math. Soc., Providence, RI, 1998, pp. 400–407.
- [11] F. Collino, G. Delbue, P. Joly, A. Piacentini, A new interface condition in the non-overlapping domain decomposition for the Maxwell equations Helmholtz equation and related optimal control, Comput. Methods Appl. Mech. Engrg 148 (1997) 195–207.
- [12] M. J. Gander, F. Magoulès, F. Nataf, Optimized Schwarz methods without overlap for the Helmholtz equation, SIAM J. Sci. Comput. 24 (1) (2002) 38–60.
- [13] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, R. Scheichl, Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps, Numer. Math. 126 (4) (2014) 741–770.
- [14] R. Haferssas, P. Jolivet, F. Nataf, An additive Schwarz method type theory for Lions’s algorithm and a symmetrized optimized restricted additive Schwarz method, SIAM J. Sci. Comput. 39 (4) (2017) A1345–A1365.
- [15] J. Fish, Y. Qu, Global-basis two-level method for indefinite systems. Part 1: convergence studies, Internat. J. Numer. Methods Engrg. 49 (3) (2000) 439–460.

- [16] A. Brandt, I. Livshits, Wave-ray multigrid method for standing wave equations, *Electron. Trans. Numer. Anal.* 6 (1997) 162–181.
- [17] C. Farhat, A. Macedo, M. Lesoinne, A two-level domain decomposition method for the iterative solution of high frequency exterior Helmholtz problems, *Numerische Mathematik* 85 (2) (2000) 283–308.
- [18] C. Farhat, P. Avery, R. Tezaur, J. Li, FETI-DPH: a dual-primal domain decomposition method for acoustic scattering, *Journal of Computational Acoustics* 13 (3) (2005) 499–524.
- [19] J.-H. Kimn, M. Sarkis, Restricted overlapping balancing domain decomposition methods and restricted coarse problems for the Helmholtz problem, *Computer Methods in Applied Mechanics and Engineering* 196 (8) (2007) 1507–1514, domain Decomposition Methods: recent advances and new challenges in engineering.
URL <http://www.sciencedirect.com/science/article/pii/S0045782506002696>
- [20] L. Conen, V. Dolean, R. Krause, F. Nataf, A coarse space for heterogeneous Helmholtz problems based on the Dirichlet-to-Neumann operator, *J. Comput. Appl. Math.* 271 (2014) 83–99.
- [21] F. Nataf, H. Xiang, V. Dolean, N. Spillane, A coarse space construction based on local Dirichlet-to-Neumann maps, *SIAM J. Sci. Comput.* 33 (4) (2011) 1623–1642.
- [22] V. Dolean, F. Nataf, R. Scheichl, N. Spillane, Analysis of a two-level Schwarz method with coarse spaces based on local Dirichlet-to-Neumann maps, *Comput. Methods Appl. Math.* 12 (4) (2012) 391–414.
- [23] I. M. Babuska, S. A. Sauter, Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?, *SIAM Journal on Numerical Analysis* 34 (6) (1997) 2392–2423.
- [24] I. G. Graham, E. A. Spence, E. Vainikko, Domain decomposition preconditioning for high-frequency Helmholtz problems with absorption, *Math. Comp.* 86 (307) (2017) 2089–2127.

- [25] N. Bootland, V. Dolean, On the Dirichlet-to-Neumann coarse space for solving the Helmholtz problem using domain decomposition, in: European Conference on Numerical Mathematics and Advanced Applications, Springer, 2019.
- [26] N. Bootland, Coarse spaces for Helmholtz, in: Scottish Numerical Methods Network Workshop on Iterative Methods for Partial Differential Equations, *available online at <http://personal.strath.ac.uk/jennifer.pestana/lms19/Bootland.pdf>*, 27 September 2019.
- [27] F. Hecht, New development in FreeFem++, Journal of Numerical Mathematics 20 (3-4) (2012) 251–266.
- [28] P. Jolivet, F. Hecht, F. Nataf, C. Prud’homme, Scalable domain decomposition preconditioners for heterogeneous elliptic problems, in: Proceedings of the 2013 International Conference on High Performance Computing, Networking, Storage and Analysis, SC13, ACM, 2013, pp. 80:1–80:11.
- [29] R. Lehoucq, D. Sorensen, C. Yang, ARPACK users’ guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods, Vol. 6, Society for Industrial and Applied Mathematics, 1998.
URL <http://www.caam.rice.edu/software/ARPACK>
- [30] V. Hernandez, J. E. Roman, V. Vidal, SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems, ACM Transactions on Mathematical Software 31 (3) (2005) 351–362.
URL <https://slepc.upv.es>
- [31] P. Amestoy, I. Duff, J.-Y. L’Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, SIAM Journal on Matrix Analysis and Applications 23 (1) (2001) 15–41.
URL <http://mumps.enseeiht.fr>
- [32] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM Journal on Scientific Computing 14 (2) (1993) 461–469.
- [33] Y. A. Erlangga, C. Vuik, C. W. Oosterlee, On a class of preconditioners for solving the Helmholtz equation, Appl. Numer. Math. 50 (3-4) (2004) 409–425.

- [34] D. Lahaye, J. Tang, K. Vuik, Modern solvers for Helmholtz problems, Springer, 2017.
- [35] I. G. Graham, E. A. Spence, E. Vainikko, Domain decomposition preconditioning for high-frequency Helmholtz problems with absorption, *Math. Comp.* 86 (307) (2017) 2089–2127.
- [36] M. Bonazzoli, V. Dolean, I. G. Graham, E. A. Spence, P.-H. Tournier, Domain decomposition preconditioning for the high-frequency time-harmonic maxwell equations with absorption, *Math. Comp.* 88 (2019) 2559–2604.
- [37] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc users manual, Tech. Rep. ANL-95/11 - Revision 3.14, Argonne National Laboratory (2020).
- [38] P. Jolivet, J. E. Roman, S. Zampini, KSPHPDDM and PCHPDDM: Extending PETSc with robust overlapping Schwarz preconditioners and advanced Krylov methods, *Computer & Mathematics with Applications* 84 (2021) 277–295.
URL <https://github.com/prj-/jolivet2020petsc>
- [39] P.-H. Tournier, I. Aliferis, M. Bonazzoli, M. de Buhan, M. Darbas, V. Dolean, F. Hecht, P. Jolivet, I. El Kanfoud, C. Migliaccio, F. Nataf, C. Pichot, S. Semenov, Microwave tomographic imaging of cerebrovascular accidents by using high-performance computing, *Parallel Computing* 85 (2019) 88–97.
- [40] F.-X. Roux, A. Barka, Block Krylov recycling algorithms for FETI-2LM applied to 3D electromagnetic wave scattering and radiation, *IEEE Transactions on Antennas and Propagation* 65 (4) (2017) 1886–1895.
- [41] R. J. Versteeg, G. Grau (eds.), The marmousi experience, in: Proc. EAGE workshop on Practical Aspects of Seismic Data Inversion (Copenhagen, 1990), Eur. Assoc. Explor. Geophysicists, Zeist, 1991.

- [42] J. Liu, J.-M. Jin, Scattering analysis of a large body with deep cavities, IEEE Transactions on Antennas and Propagation 51 (6) (2003) 1157–1167.
- [43] J.-M. Jin, The Finite Element Method in Electromagnetics, John Wiley & Sons, 2015.
- [44] F. Aminzadeh, J. Brac, T. Kunz, 3-D Salt and Overthrust models, SEG/EAGE 3-D Modeling Series No. 1, 1997.

Table 2: Results using the H-GenEO and DtN spectral coarse space methods for the Marmousi problem when using 5 points per wavelength, varying the frequency f , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$f \setminus N$	H-GenEO (min. overlap)					DtN (min. overlap)					H-GenEO (coarse overlap)					DtN (coarse overlap)				
		10	20	40	80	160	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
20	1	9	14	21	21	—	8	7	6	6	—	16	22	31	34	—	5	5	5	5	—
	5	46	53	67	76	100	37	46	59	74	85	58	86	120	158	201	39	54	74	83	68
	10	67	99	146	173	209	61	87	118	156	190	76	133	220	301	420	62	91	129	178	242
	20	81	120	169	248	378	75	121	157	208	275	83	125	226	370	×	75	117	158	216	300
40	1	9	13	16	21	—	10	8	6	6	—	16	22	26	32	—	7	8	6	7	—
	5	31	32	42	49	71	25	34	36	40	37	46	74	104	137	186	28	32	24	11	9
	10	66	85	105	106	126	48	65	90	102	122	100	153	233	269	373	52	80	115	131	106
	20	86	136	209	264	323	68	106	138	186	259	119	232	391	×	×	68	105	149	236	357
80	1	9	12	18	21	—	5	6	6	6	—	15	20	22	33	—	4	4	6	6	—
	5	20	20	28	41	62	15	18	13	39	67	40	62	87	127	184	10	8	8	23	19
	10	47	56	64	72	88	33	40	55	48	39	88	118	171	225	327	39	46	42	14	13
	20	87	125	171	186	217	54	86	111	152	164	153	246	404	×	×	57	104	155	188	145
160	1	9	11	17	21	—	8	7	5	6	—	16	15	21	33	—	6	7	9	6	—
	5	15	17	26	37	56	13	25	23	26	36	37	59	82	116	171	7	19	10	8	19
	10	33	40	45	56	73	21	27	30	48	80	68	94	143	196	284	18	19	17	48	29
	20	64	83	121	134	157	33	59	67	76	75	131	192	320	403	×	43	75	77	61	35

Table 4: Results using the H-GenEO and DtN spectral coarse space methods for the Marmousi problem when using 10 points per wavelength, varying the frequency f , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$f \setminus N$	H-GenEO (min. overlap)					DtN (min. overlap)					H-GenEO (coarse overlap)					DtN (coarse overlap)				
		10	20	40	80	160	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
20	1	8	8	9	15	—	8	8	5	7	—	7	7	9	14	—	5	6	4	4	—
	5	42	42	39	32	19	42	49	59	51	56	40	46	47	38	32	39	49	67	72	85
	10	78	112	125	133	99	70	105	128	157	166	72	108	133	156	138	65	99	128	169	197
	20	96	136	180	259	342	92	131	175	232	318	89	130	177	248	345	85	125	170	227	315
40	1	7	8	9	14	—	6	5	6	6	—	7	7	9	13	—	4	3	3	4	—
	5	22	19	17	13	14	25	29	31	32	33	21	20	19	17	20	25	34	43	40	38
	10	58	63	54	42	26	49	71	74	84	82	58	70	65	61	44	48	75	88	112	128
	20	101	148	169	161	148	85	120	159	190	217	94	141	183	203	214	80	116	154	194	256
80	1	7	8	9	13	—	7	7	5	5	—	6	7	9	12	—	5	4	3	4	—
	5	13	12	11	10	12	16	16	19	17	10	12	11	12	13	17	19	15	14	7	6
	10	31	28	23	17	15	28	39	40	42	41	30	31	27	23	24	29	47	55	66	34
	20	81	90	78	53	39	63	89	93	100	107	85	104	107	89	69	64	93	108	147	202
160	1	7	8	8	13	—	4	7	5	6	—	6	7	8	11	—	3	6	3	5	—
	5	10	9	10	10	12	10	11	12	17	24	8	9	10	13	16	10	9	8	21	14
	10	20	16	14	13	13	19	23	25	25	24	18	15	16	16	19	21	29	30	25	17
	20	45	40	34	25	19	35	46	47	56	59	50	52	50	39	28	36	61	77	84	76

Table 14: Results using the H-GenEO and DtN spectral coarse space methods for the Overthrust problem when using 5 points per wavelength, varying the frequency f , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$f \setminus N$	H-GenEO (min. overlap)					DtN (min. overlap)					H-GenEO (coarse overlap)					DtN (coarse overlap)				
		10	20	40	80	160	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
20	0.5	11	14	17	30	—	6	6	7	8	—	12	15	19	19	—	5	5	4	5	—
	1	25	33	47	52	60	21	28	43	45	49	25	34	44	49	56	19	28	42	48	42
	2	33	47	62	78	99	31	45	59	72	89	32	48	74	101	134	28	41	56	70	97
40	0.5	10	12	15	25	—	5	5	6	7	—	12	15	18	19	—	3	4	4	4	—
	1	24	32	44	52	59	17	25	36	38	32	24	34	45	50	56	15	20	28	14	6
	2	34	50	70	104	123	30	42	55	67	87	35	61	87	112	146	28	42	62	85	138
80	0.5	9	11	14	26	—	4	5	5	6	—	12	14	18	18	—	3	3	4	5	—
	1	22	31	44	52	60	13	19	20	14	8	23	33	45	50	56	9	9	10	5	4
	2	38	61	84	111	139	27	38	53	70	96	41	63	91	116	158	27	44	71	85	65
160	0.5	9	11	11	28	—	4	5	5	—	—	12	14	17	19	—	4	3	4	—	—
	1	19	29	44	52	60	10	10	10	7	7	23	33	46	50	56	7	6	7	5	5
	2	36	54	77	107	136	23	34	50	56	56	41	62	90	117	159	23	34	42	20	12

Table 16: Results using the H-GenEO and DtN spectral coarse space methods for the Overthrust problem when using 10 points per wavelength, varying the frequency f , the number of subdomains N , and the number of eigenvectors used per subdomain ν .

ν	$f \setminus N$	H-GenEO (min. overlap)					DtN (min. overlap)					H-GenEO (coarse overlap)					DtN (coarse overlap)				
		10	20	40	80	160	10	20	40	80	160	10	20	40	80	160	10	20	40	80	160
20	0.5	8	9	10	11	12	8	7	7	7	7	7	8	10	19	27	7	5	5	5	5
	1	24	29	34	30	17	23	31	40	44	45	22	30	42	51	64	21	30	38	45	49
40	0.5	8	9	9	11	12	6	6	6	6	6	7	8	9	18	26	4	4	4	4	4
	1	17	18	19	13	14	20	25	29	31	26	19	23	32	44	63	18	25	31	38	43
80	0.5	8	8	9	10	12	5	5	5	6	5	7	8	9	18	26	4	3	4	4	4
	1	12	11	12	11	13	15	18	22	23	15	14	14	21	37	60	15	19	26	24	13
160	0.5	7	8	9	10	11	4	5	5	5	7	7	7	8	17	26	3	4	4	4	5
	1	10	10	10	11	12	11	13	14	17	6	11	12	16	30	57	10	13	12	6	5