



HAL
open science

Multidirectional sweeping preconditioners with non-overlapping checkerboard domain decomposition for Helmholtz problems

Ruiyang Dai, Axel Modave, Jean-François Remacle, Christophe Geuzaine

► **To cite this version:**

Ruiyang Dai, Axel Modave, Jean-François Remacle, Christophe Geuzaine. Multidirectional sweeping preconditioners with non-overlapping checkerboard domain decomposition for Helmholtz problems. *Journal of Computational Physics*, 2022, 453, pp.110887. 10.1016/j.jcp.2021.110887 . hal-03240042

HAL Id: hal-03240042

<https://hal.science/hal-03240042>

Submitted on 28 May 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.

Multidirectionnal sweeping preconditioners with non-overlapping checkerboard domain decomposition for Helmholtz problems

R. Dai^{1,3}, A. Modave², J.-F. Remacle¹, and C. Geuzaine³

¹IMMC, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium,
ruiyang.dai@uclouvain.be, jean-francois.remacle@uclouvain.be

²POEMS, CNRS, Inria, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France,
axel.modave@ensta-paris.fr

³Université de Liège, Institut Montefiore B28, 4000 Liège, Belgium,
ruiyang.dai@uclouvain.be, cgeuzaine@uliege.be

Abstract

This paper explores a family of generalized sweeping preconditioners for Helmholtz problems with non-overlapping checkerboard partition of the computational domain. The domain decomposition procedure relies on high-order transmission conditions and cross-point treatments, which cannot scale without an efficient preconditioning technique when the number of subdomains increases. With the proposed approach, existing sweeping preconditioners, such as the symmetric Gauss-Seidel and parallel double sweep preconditioners, can be applied to checkerboard partitions with different sweeping directions (*e.g.* horizontal and diagonal). Several directions can be combined thanks to the flexible version of GMRES, allowing for the rapid transfer of information in the different zones of the computational domain, then accelerating the convergence of the final iterative solution procedure. Several two-dimensional finite element results are proposed to study and to compare the sweeping preconditioners, and to illustrate the performance on cases of increasing complexity.

1 Introduction

Time-harmonic wave simulations are of interest in many scientific and engineering disciplines. For example, in radar or sonar imaging and wireless communications, the wavelength of the signal is usually several orders of magnitude smaller than the size of the domains of interest. Similarly, in seismic imaging, wave fields in complex geological media show a wide range of space-varying wavenumbers, caused by large variations in the velocity profile. Solving such time-harmonic problems numerically using finite element-type methods is notoriously difficult because it leads to (extremely) large indefinite linear systems [14], especially in the high-frequency regime. On the one hand, sparse direct solvers exhibit poor scalability w.r.t. memory and computational time for such linear systems, in particular for three-dimensional problems. On the other hand, most iterative methods that have proved successful for elliptic problems become inefficient when applied to problems with highly oscillatory solutions, and no robust and scalable preconditioner currently exists [25].

Parallel iterative solvers and parallel preconditioners, called generically “*domain decomposition methods*” (DDMs), are currently intensively studied for time-harmonic problems. These methods rely on the parallel solution of subproblems of smaller sizes, amenable to sparse direct solvers. In a finite element context, there are two types of DDMs: overlapping DDMs in which the meshes of two adjacent subdomains overlap by at least one element; and non-overlapping DDMs where adjacent subdomains only communicate through an artificial (lower dimensional) interface. The latter include *e.g.* non-overlapping Schwarz methods [11], FETI algorithms [9, 15] and the method of polarized traces [48].

In this work, we focus on non-overlapping domain decomposition solvers with optimized transmission conditions, which are well suited for time-harmonic wave problems [4, 11, 17, 18]. After non-overlapping Schwarz methods were introduced by Lions [29] for the Laplace equation and proven to converge for the Helmholtz equation by Després [11], considerable efforts have been made to develop efficient transmission

conditions to improve the rate of convergence for DDMs. The optimal convergence is obtained by using as transmission condition on each interface the Dirichlet-to-Neumann (DtN) map related to the complementary of the subdomain of interest. This DtN map is a nonlocal operator and is thus in practice very expensive to compute. Optimized Schwarz methods were introduced in [18], where the nonlocal DtN is approximated by first or second order polynomial approximations, with coefficients obtained by optimization on simple geometries. Later, quasi-optimal optimized Schwarz methods were proposed in [4] based on rational approximations related to those used in high-order absorbing boundary conditions (HABC). Conditions based on second-order operators [37], perfectly matched layers (PML) [1, 42, 46] and non-local approaches [7, 26] have also been investigated.

Even with optimal transmission conditions however, as is expected for a one-level method, the number of iterations of the DDM will grow as the number of subdomains increases. A solution for certain classes of problems is to add a component to the algorithm that is known in the DDM community as a “*coarse grid*” [16], in effect a second-level to enable longer-range information exchange than the local sharing (from one subdomain to its neighbors) of the one-level DDM. Nevertheless, the design of robust coarse grids is very challenging for wave-type problems because of the highly oscillatory behavior of the solution, and several approaches are currently investigated in the community (see *e.g.* [3, 8] and references herein). As an alternative approach, sweeping preconditioners have been proposed and studied for convection-diffusion problems in the 90’s [34, 35]. They have recently garnered a lot of interest for high-frequency Helmholtz problems [6, 13, 19, 42, 45, 46, 48], promising a number of DDM iterations that is quasi independent of the number of subdomains. However, they have two major drawbacks: they rely on intrinsically sequential operations (they are related to a LU-type factorization of the underlying iteration operator) and they are naturally only suited for layered-type domain decompositions (where the layered structure allows to explicit the LU factorization as a double sweep across the subdomains).

In this work, we explore a family of generalized sweeping preconditioners where sweeps can be done in several directions for non-overlapping domain decomposition solvers with “*checkerboard*” domain partition. This contribution relies on the availability of transmission conditions able to deal with the cross-points (*i.e.* points where more than two subdomains meet) arising in such configurations. We consider a domain decomposition solver with high-order Padé-type transmission conditions [4] and a cross-point treatment recently proposed in [33]. Sweeping preconditioners are derived in a systematic manner, based on the explicit representation of the iteration matrix in the case of checkerboard decompositions. The sweeps can be performed in Cartesian and diagonal directions, and several sweeping directions can be combined by using the flexible version of GMRES [39, 40], which allows to change the preconditioner at each iteration. For applicative cases, the resulting preconditioners provide an effective way to rapidly transfer information in the different zones of the computational domain, then accelerating the convergence of iterative solution procedure with GMRES. Our approach is related to the recent work on L-sweeps preconditioners [45] and diagonal sweeping technique [27], where directional sweeping strategies are proposed in the context of the method of polarized traces and the source transfer method, respectively. Here, the preconditioners are proposed for non-overlapping domain decomposition solvers with high-order transmission conditions, and several directions can be combined thanks to the use of flexible GMRES.

The paper is organized as follows. In Section 2, we present the non-overlapping domain decomposition algorithm with high-order transmission condition and cross-points treatment for checkerboard domain decompositions. The matrix representation of the resulting iteration operator is derived and studied in Section 3. It is used in Section 4 to derive sweeping preconditioners. In Section 5, we study and compare the resulting preconditioned domain decomposition algorithms with two-dimensional finite element benchmarks. The efficiency of these methods is demonstrated on numerical models and configurations of increasing complexity.

2 Domain decomposition algorithm for the Helmholtz equation

To describe our approach, we consider the two-dimensional scattering problem of an incident acoustic plane wave by a sound-soft obstacle of boundary Γ^{sca} . The numerical simulations are performed in a rectangular computational domain Ω of boundary $\partial\Omega = \Gamma^{\text{sca}} \cup \Gamma^\infty$, with Γ^∞ the external (artificial)




Figure 1: Configuration of the problem (*left*), illustration of the checkerboard partition (*middle*) and notation for the edges of the subdomain Ω_I (*right*).

boundary (see Figure 1, left). We seek the scattered field $u(\mathbf{x})$ that verifies

$$\begin{cases} -\Delta u - \kappa^2 u = 0, & \text{in } \Omega, \\ \partial_{\mathbf{n}} u + \mathcal{B}u = 0, & \text{on } \Gamma^\infty, \\ u = -u^{\text{inc}}, & \text{on } \Gamma^{\text{sca}}, \end{cases} \quad (1)$$

where κ is the wavenumber, u^{inc} is the incident wave, $\partial_{\mathbf{n}}$ is the exterior normal derivative and \mathcal{B} is an impedance operator to be defined. We take the convention that the time-dependence of the fields is $e^{-i\omega t}$, where ω is the angular frequency and t is the time. The impedance operator corresponding to a Sommerfeld absorbing boundary condition (ABC) on Γ^∞ is $\mathcal{B} = -i\kappa$.

2.1 Domain decomposition algorithm on a checkerboard partition

We consider a checkerboard partition of the domain Ω , that consists in a lattice of rectangular non-overlapping subdomains Ω_I ($I = 1 \dots N_{\text{dom}}$) with N_r rows and N_c columns (then, $N_{\text{dom}} = N_r \times N_c$). For each rectangular subdomain Ω_I , the four edges are denoted Γ_I^\square , with $\square \in \{x_-, y_-, x_+, y_+\}$, where the superscripts x_- , y_- , x_+ and y_+ correspond to the edges on the left, on the bottom, on the right and on the top of the subdomain, respectively (see Figure 1, right). The union of the edges then reads

$$\bigcup_{\square} \Gamma_I^\square = \Gamma_I^{x-} \cup \Gamma_I^{y-} \cup \Gamma_I^{x+} \cup \Gamma_I^{y+}.$$

To simplify the presentation, we assume that the obstacle is included in only one subdomain, the boundary of which is the union of the four edges and the boundary of the obstacle (see Figure 1, middle).

Each edge of one subdomain Ω_I is either a *boundary edge* if it belongs to the boundary of the global domain ($\Gamma_I^\square \subset \partial\Omega$) or an *interface edge* if there is a neighboring subdomain on the other side of the edge ($\Gamma_I^\square \not\subset \partial\Omega$). In this checkerboard partition, there are corners where at least two edges meet. Each corner of a subdomain is an *interior cross-point* (point that belongs to four subdomains), a *boundary cross-point* (point that belongs to two subdomains and to the exterior border $\partial\Omega$) or a corner of the main domain Ω .

With these definitions, the non-overlapping domain decomposition algorithm can be set up as follows. For each subdomain Ω_I , we seek the solution $u_I(\mathbf{x})$ of the subproblem

$$\begin{cases} -\Delta u_I - \kappa^2 u_I = 0, & \text{in } \Omega_I, \\ \partial_{\mathbf{n}_I^\square} u_I + \mathcal{B}_I^\square u_I = 0, & \text{on each } \Gamma_I^\square \subset \partial\Omega, \\ \partial_{\mathbf{n}_I^\square} u_I + \mathcal{B}_I^\square u_I = g_I^\square, & \text{on each } \Gamma_I^\square \not\subset \partial\Omega, \\ u_I = -u^{\text{inc}}, & \text{on } \partial\Omega_I \cap \Gamma^{\text{sca}}, \end{cases} \quad (2)$$

where \mathbf{n}_I^\square is the outward unit normal to the edge Γ_I^\square , \mathcal{B}_I^\square is an impedance operator and g_I^\square is a transmission variable defined on Γ_I^\square . The second and third equations in (2) are boundary and transmission conditions, respectively.

For a given boundary edge $\Gamma_I^\square \subset \partial\Omega$, we must have $\mathcal{B}_I^\square = \mathcal{B}$ to ensure the equivalence between all the subproblems and the original problem. If $\Gamma_I^\square \not\subset \partial\Omega$, there is some flexibility in the choice of \mathcal{B}_I^\square . The transmission variable is defined as

$$g_I^\square := \partial_{\mathbf{n}_I^\square} u_J + \mathcal{B}_I^\square u_J, \quad (3)$$

where u_J is the solution of the neighboring subdomain Ω_J . The transmission conditions defined on both sides of the interface enforce the continuity of the solution across the interface. Assuming that the impedance operators used on both sides of the shared interface edge $\Gamma_I^\square = \Gamma_J^{\square'} = \partial\Omega_I \cap \partial\Omega_J$ are the same (*i.e.* $\mathcal{B}_I^\square = \mathcal{B}_J^{\square'}$), the transmission variables defined on this edge verify

$$g_I^\square = -g_J^{\square'} + 2\mathcal{B}_J^{\square'} u_J, \quad (4)$$

where $g_J^{\square'}$ is the transmission variable defined on the edge $\Gamma_J^{\square'}$ of Ω_J .

The non-overlapping optimized Schwarz domain decomposition algorithm consists in solving subproblems associated to all the subdomains (equation (2)) concurrently and updating the transmissions variables using (4) in an iterative process. At each iteration $n+1$, the update formula of a transmission variable living on an interface edge Γ_I^\square of a subdomain Ω_I reads

$$g_I^{\square(n+1)} = -g_J^{\square'(n)} + 2\mathcal{B}_J^{\square'} u_J^{(n)}, \quad (5)$$

where $u_J^{(n)}$ is the solution of the neighboring subdomain Ω_J at the iteration n . The update of all the transmission variables can be recast as one application of the iteration operator \mathcal{A} defined by

$$\mathbf{g}^{(n+1)} = \mathcal{A}\mathbf{g}^{(n)} + \mathbf{b}, \quad (6)$$

where $\mathbf{g}^{(n)}$ is the set of all transmission variables defined on the interface edges and \mathbf{b} is given by the source term. It is well known that this algorithm can be seen as a Jacobi scheme applied to the linear system

$$(\mathcal{I} - \mathcal{A})\mathbf{g} = \mathbf{b}, \quad (7)$$

where \mathcal{I} is the identity operator. In order to accelerate the convergence of the procedure, this system can be solved with Krylov subspace iterative methods, such as GMRES.

2.2 Transmission operators

The convergence rate of the non-overlapping DDMs strongly depends on the impedance operator used in the transmission conditions. The optimal transmission operator corresponds to the non-local Dirichlet-to-Neumann (DtN) map related to the complementary of each subdomain. Since the cost of computing the exact DtN is prohibitive, strategies based on approximate DtN operators started to be investigated in the late 80's and early 90's (see *e.g.* [21, 36]). For Helmholtz problems, Després [2, 10] used a Robin-type operator, which is a coarse approximation of the exact DtN operator. Improved methods with optimized second-order transmission operators have next been introduced in [17, 38]. More recently, domain decomposition approaches with improved convergence rates have been proposed by using transmission conditions based on high-order absorbing boundary conditions (HABCs) [4, 5, 24, 30], perfectly matched layers (PMLs) [1, 41, 42, 46] and nonlocal operators [7, 26, 44]. As for ABCs, transmission boundary conditions related to HABCs and PMLs represent a good compromise between the basic impedance conditions (which lead to suboptimal convergence) and nonlocal approaches (which are expensive to compute).

In this work, we use the DtN operator associated to the Padé-type HABC as the impedance operator in the transmission conditions, following [4]. For an edge Γ_I^\square , the operator can be written as

$$\mathcal{B}_I^\square = -\imath\kappa\alpha \left[1 + \frac{2}{M} \sum_{i=1}^N c_i \left(1 - \alpha^2(c_i + 1) [(\alpha^2 c_i + 1) + \partial_{\tau\tau}/\kappa^2]^{-1} \right) \right], \quad \text{on } \Gamma_I^\square, \quad (8)$$

where ∂_τ is the tangential derivative and we have $\alpha = e^{\imath\phi/2}$, $c_i = \tan^2(i\pi/M)$ and $M = 2N + 1$. This Padé-type impedance operator is obtained by approximating the exact non-local DtN map associated




Figure 2: Configuration with two subdomains.

to the exterior half-plane problem (see *e.g.* [12, 31]). The symbol of the non-local operator exhibits a square-root which is replaced with the $(2N + 1)^{\text{th}}$ -order Padé approximation after a ϕ -rotation of the branch-cut. The performance of the obtained operator depends on the number of terms N and the angle of rotation ϕ . The particular parameters $N = 0$ and $\phi = 0$ leads to the basic ABC operator $\mathcal{B}_I^\square = -\imath\kappa$. See *e.g.* [23, 32] for further details.

For the effective implementation of the transmission condition, the application of the Padé-type impedance operator on a field is written in such a way that it involves only differential operators. Following an approach first used by Lindman [28] for ABCs, we introduce N auxiliary fields governed by auxiliary equations on interface edge Γ_I^\square . The application of the Padé-type impedance operator is then written as

$$\mathcal{B}_I^\square u_I = B\left(u_I, \{\varphi_{I,i}^\square\}_{i=1\dots N}\right) := -\imath\kappa\alpha \left[u_I + \frac{2}{M} \sum_{i=1}^N c_i \left(u_I + \varphi_{I,i}^\square \right) \right], \quad \text{on } \Gamma_I^\square, \quad (9)$$

with the auxiliary fields $\{\varphi_{I,i}^\square\}_{i=1\dots N}$ defined only on the edge and governed by the auxiliary equations

$$-\partial_{\tau\tau}\varphi_{I,i}^\square - \kappa^2((\alpha^2 c_i + 1)\varphi_{I,i}^\square + \alpha^2(c_i + 1)u_I) = 0, \quad \text{on } \Gamma_I^\square, \quad (10)$$

for $i = 1 \dots N$. The linear multivariate function B is introduced to simplify the expressions in the remainder of the paper. When this operator is used in a boundary condition for polygonal domains, a special treatment is required to preserve the accuracy of the solution at the corners. In the case of right-angle corners, an approach based on compatibility relations reveals to be very efficient [32].

2.3 Cross-point treatment

When the Padé-type impedance operator (9) is used in the boundary conditions and the interface conditions of the subproblem (2), a special treatment is required at the corners of the subdomain Ω_I . Indeed, the auxiliary fields governed by equation (10) on the edges require boundary conditions at the extremities of the edges, which corners of the subdomain. In this work, we use a cross-point treatment based on compatibility relations developed for the right-angle case, first proposed in [33].

To present the cross-point treatment, we consider the neighboring subdomains Ω_I and Ω_J represented on Figure 2. For the sake of shortness, we describe the methodology in the case where transmission conditions are prescribed on all the edges of both subdomains (*i.e.* they all are interface edges) and the Padé-type impedance operator is used with the same parameters on all the edges.

The subdomains share the interface edge $\Gamma_I^{x+} = \Gamma_J^{x-}$ and the interior cross-points $P_I^{x+y+} = P_J^{x-y+}$ and $P_I^{x+y-} = P_J^{x-y-}$. On the shared interface edge $\Gamma_I^{x+} = \Gamma_J^{x-}$, we have the transmissions conditions

$$\partial_{\mathbf{n}_I^{x+}} u_I + B\left(u_I, \{\varphi_{I,i}^{x+}\}_{i=1\dots N}\right) = g_I^{x+}, \quad \text{on } \Gamma_I^{x+}, \quad (11)$$

$$\partial_{\mathbf{n}_J^{x-}} u_J + B\left(u_J, \{\varphi_{J,i}^{x-}\}_{i=1\dots N}\right) = g_J^{x-}, \quad \text{on } \Gamma_J^{x-}, \quad (12)$$

where u_I and u_J are the main fields defined on Ω_I and Ω_J , respectively. The auxiliary fields $\{\varphi_{I,i}^{x+}\}_i$ and $\{\varphi_{J,i}^{x-}\}_i$ defined on the shared interface are governed by equation (10). The first set of auxiliary fields is associated to the subproblem defined on Ω_I , and the second set is associated to the one defined on Ω_J . By using the impedance operator (9) in equation (4) on both sides of the interface, we have that the transmission variables g_I^{x+} and g_J^{x-} verify

$$g_I^{x+} = -g_J^{x-} + 2B\left(u_J, \{\varphi_{J,i}^{x-}\}_{i=1\dots N}\right), \quad \text{on } \Gamma_I^{x+}, \quad (13)$$

$$g_J^{x-} = -g_I^{x+} + 2B\left(u_I, \{\varphi_{I,i}^{x+}\}_{i=1\dots N}\right), \quad \text{on } \Gamma_J^{x-}. \quad (14)$$

With these transmission variables, the transmission conditions (11) and (12) enforce weakly the continuity of the main field across the shared interface.

The cross-point treatment consists in enforcing weakly the continuity of auxiliary fields at cross-points. More precisely, only auxiliary fields defined on edges that are aligned are continuous. For instance, the auxiliary fields $\{\varphi_{I,j}^{y+}\}_j$ defined on Γ_I^{y+} and the auxiliary fields $\{\varphi_{J,j}^{y+}\}_j$ defined on Γ_J^{y+} (*i.e.* defined on the upper edges of Ω_I and Ω_J , respectively, see Figure 2) must be equal at $P_I^{x+y+} = P_J^{x-y+}$. Following the approach detailed in [33], transmission conditions with specific impedance operators are used to enforce weakly the continuity. At the cross-point, we use the transmission conditions

$$\partial_{n_I^{x+}} \varphi_{I,j}^{y+} + B\left(\varphi_{I,j}^{y+}, \{\psi_{I,ij}^{x+y+}\}_{i=1\dots N}\right) = g_{I,j}^{y+x+}, \quad \text{at } P_I^{x+y+}, \quad (15)$$

$$\partial_{n_J^{x-}} \varphi_{J,j}^{y+} + B\left(\varphi_{J,j}^{y+}, \{\psi_{J,ij}^{x-y+}\}_{i=1\dots N}\right) = g_{J,j}^{y+x-}, \quad \text{at } P_J^{x-y+}, \quad (16)$$

for $j = 1 \dots N$, with the scalar variables $\{\psi_{I,ij}^{x+y+}\}_{ij}$ and $\{\psi_{J,ij}^{x-y+}\}_{ij}$ defined as

$$\psi_{I,ij}^{x+y+} = -[\alpha^2(c_j + 1)\varphi_{I,i}^{x+} + \alpha^2(c_i + 1)\varphi_{I,j}^{y+}] / [\alpha^2c_i + \alpha^2c_j + 1], \quad \text{at } P_I^{x+y+}, \quad (17)$$

$$\psi_{J,ij}^{x-y+} = -[\alpha^2(c_j + 1)\varphi_{J,i}^{x-} + \alpha^2(c_i + 1)\varphi_{J,j}^{y+}] / [\alpha^2c_i + \alpha^2c_j + 1], \quad \text{at } P_J^{x-y+}, \quad (18)$$

for $i, j = 1 \dots N$. At the cross-point, the new transmission variables $\{g_{I,j}^{x+}\}_j$ and $\{g_{J,j}^{x-}\}_j$ verify

$$g_{I,j}^{y+x+} = -g_{J,j}^{y+x-} + 2B\left(\varphi_{J,j}^{y+}, \{\psi_{J,ij}^{x-y+}\}_{i=1\dots N}\right), \quad \text{at } P_I^{x+y+}, \quad (19)$$

$$g_{J,j}^{y+x-} = -g_{I,j}^{y+x+} + 2B\left(\varphi_{I,j}^{y+}, \{\psi_{I,ij}^{x+y+}\}_{i=1\dots N}\right), \quad \text{at } P_J^{x-y+}, \quad (20)$$

for $j = 1 \dots N$. In a nutshell, the same transmission condition with the multivariate function B is used on the interface to couple the main fields (equations (11)-(12)) and at the cross-points to couple the auxiliary fields (equations (17)-(18)). The scalar variables defined at the corners of the subdomains introduce a coupling of auxiliary fields living on adjacent edges of each subdomain. This strategy can be adapted rather straightforwardly to deal with boundary cross-points, where interface edges and boundary edges with boundary condition meet. For further details, we refer to [33].

The iterative domain decomposition algorithm is very similar to the algorithm described at the end of section 2.1. At each iteration, subproblems associated to the subdomains are solved concurrently, and transmission variables are updated. Here, the subproblem associated to Ω_I consists in finding the main field verifying system (2) and auxiliary fields verifying equations similar to (10) on each interface edge. The transmission variables are associated to interfaces edges and cross-points. They are updated with formulas similar to

$$g_I^{x+(n+1)} = -g_J^{x-(n)} + 2B\left(u_J^{(n)}, \{\varphi_{J,i}^{x-(n)}\}_{i=1\dots N}\right), \quad \text{on } \Gamma_I^{x+} \quad (21)$$

and

$$g_{I,j}^{y+x+(n+1)} = -g_{J,j}^{y+x-(n)} + 2B\left(\varphi_{J,i}^{y+(n)}, \{\psi_{J,ij}^{x-y+(n)}\}_{i=1\dots N}\right), \quad \text{on } P_I^{x+y+}, \quad (22)$$

which are obtained by rewriting equations (13) and (19) similarly to the general update formula (5). Following the approach explained in section 2.1, all the transmission variables can be merged into a

global vector $\mathbf{g}^{(n+1)}$, and the global process can be recast as one application of an iterative operator \mathcal{A} on the vector. At each iteration $n+1$, the whole process can be seen as one step of the Jacobi method to solve the linear system $(\mathcal{I} - \mathcal{A})\mathbf{g} = \mathbf{b}$, which could be solved with a Krylov subspace iterative method. Here, the main difference with most of the works is that the global vector includes transmission variables associated to both interfaces and cross-points.

3 Algebraic structure of the interface problem

In this section, we analyze the algebraic structure of the global interface problem, which can be written in an abstract form as

$$\mathcal{F}\mathbf{g} := (\mathcal{I} - \mathcal{A})\mathbf{g} = \mathbf{b}, \quad (23)$$

where \mathcal{A} is the iteration matrix, \mathbf{g} is the set of all transmission variables and \mathbf{b} is given by the source term. The global matrix $\mathcal{F} := \mathcal{I} - \mathcal{A}$ can be represented as a $N_{\text{dom}} \times N_{\text{dom}}$ sparse block matrix, which each block corresponds to the coupling between the unknowns of two subdomains. The nature of the blocks is discussed 3.1 and the sparse structure of this global block matrix is analyzed in subsection 3.2.

3.1 Identification of the blocks

Using the block representation, the abstract system (23) can be rewritten as

$$\sum_{J=1}^{N_{\text{dom}}} \mathcal{F}_I^J \mathbf{g}_J = \mathbf{b}_I, \quad I = 1 \dots N_{\text{dom}}, \quad (24)$$

where the vectors \mathbf{g}_I and \mathbf{b}_I contain all the transmission variables and the source terms, respectively, for the subdomain Ω_I . The block \mathcal{F}_I^J corresponds to a coupling between the transmission variables of the subdomains Ω_I and Ω_J . The blocks corresponding to subdomains that are not neighbours (*i.e.* which do not share any interface edge) are cancelled because there is no direct coupling between the corresponding variables. Since there are at most four neighbouring subdomains for each subdomains, there are at most four off-diagonal blocks in each line and each column of the global block matrix.

For studying the blocks, we consider a setting with transmission conditions based on the basic impedance operator for the sake of simplicity. In that case, all the transmission variables are associated to the interface edges. Since there are two transmission variables per interface edge (one for each neighboring subdomain), the total size of the vectors \mathbf{g} and \mathbf{b} is twice the number of interface edges. The size of the blocks \mathbf{g}_I and \mathbf{b}_I in these vectors corresponds to the number of interface edges for the subdomain Ω_I . The block can then be rectangular if the neighbouring subdomains Ω_I and Ω_J have different numbers of interface edges. To simplify the presentation, we assume hereafter that Ω_I and Ω_J do not touch the exterior border of the main domain (*i.e.* each of them has four interface edges and \mathcal{F}_I^J is a 4×4 matrix).

Every line of the system corresponds to a relation similar to equation (4). Let us consider the line for the transmission variable g_I^{x+} , which corresponds to the relation

$$g_I^{x+} = -g_J^{x-} + 2\mathcal{B}_J^{x-} u_J, \quad \text{on } \Gamma_I^{x+}, \quad (25)$$

where J is such that Ω_I and Ω_J are neighbouring subdomains with the shared interface edge $\Gamma_I^{x+} = \Gamma_J^{x-}$. By the linearity of the problem, the solution u_J can be split into two contributions, $u_J = v_J + w_J$. The field v_J is the solution of subproblem (2) for Ω_J where the right-hand-side term of the Dirichlet boundary condition on $\partial\Omega_J \cap \Gamma^{\text{sca}}$ is cancelled. The field w_J is the solution of subproblem (2) for Ω_J where the right-hand-side terms of the transmission conditions are cancelled. Equation (25) can then be rewritten as

$$g_I^{x+} = -g_J^{x-} + 2\mathcal{B}_J^{x-} v_J + b_J^{x-}, \quad \text{on } \Gamma_I^{x+}, \quad (26)$$

where $b_J^{x-} := 2\mathcal{B}_J^{x-} w_J$ depends only on the data of the problem, with is the incident plane wave is the present case. In order to exhibit dependences between transmission variables, we decompose the field v_J




Figure 3: Illustration of the coupling introduced by the self-coupling operator (in red) and the transfer operator (in blue) for a configuration with neighboring subdomains Ω_I and Ω_J with the shared interface edge $\Gamma_I^{x+} = \Gamma_J^{x-}$.

into several contributions. For every interface edge Γ_J^\square (with $\square \in \{x_-, y_-, x_+, y_+\}$), we introduce the field $v_J^\square(g_J^\square)$ as the solution of subproblem (2) for Ω_J with the transmission variable g_J^\square prescribed on Γ_J^\square and where all the other transmission variables and the right-hand side term of the Dirichlet condition are cancelled. By linearity, the solution of v_J can then be written as

$$v_J = \sum_{\square} v_J^\square(g_J^\square), \quad (27)$$

Using this decomposition into equation (26) gives

$$g_I^{x+} + g_J^{x-} - 2 \sum_{\square} \mathcal{B}_J^{x-} v_J^\square(g_J^\square) = b_J^{x-}, \quad \text{on } \Gamma_I^{x+}. \quad (28)$$

Defining the self-coupling and transfer operators as

$$\begin{aligned} \mathcal{E}_I^J : g_J^{x-} &\mapsto g_J^{x-} - 2\mathcal{B}_J^{x-} v_J^{x-}(g_J^{x-}) := \mathcal{E}_I^J g_J^{x-}, \\ \mathcal{T}_I^J : g_J^\square &\mapsto -2\mathcal{B}_J^{x-} v_J^\square(g_J^\square) := \mathcal{T}_I^J g_J^\square, \quad \text{for } \square \in \{y_-, x_+, y_+\}, \end{aligned} \quad (29)$$

we finally have the representation

$$g_I^{x+} + \mathcal{E}_I^J g_J^{x-} + \sum_{\square \neq x-} \mathcal{T}_I^J g_J^\square = b_J^{x-}, \quad \text{on } \Gamma_I^{x+}. \quad (30)$$

The self-coupling operator \mathcal{E}_I^J introduces a coupling between the transmission variables living on the same interface edge $\Gamma_I^{x+} = \Gamma_J^{x-}$, while the transfer operators introduces a coupling between g_I^{x+} and the transmission variables living on the other interface edges of Ω_J (*i.e.* any $\Gamma_J^\square \neq \Gamma_I^{x+}$). These couplings are illustrated in Figure 3.

Thanks to the representation in equation (30), the elements of the matrix and the right-hand side of the global system can be identified. The right-hand side of (30) is an element of \mathbf{b}_J . Looking at the first term in the left-hand side, we straightforwardly have that the blocks on the diagonal of the global matrix are identity matrices. Finally, the second and third terms correspond to elements in the off-diagonal block \mathcal{F}_I^J . The other elements of this block are equal to zero because the relation (30) corresponding to the other edges of Ω_I do not involve transmission variables of Ω_J . For instance, assuming that the shared interface edge is $\Gamma_I^{x+} = \Gamma_J^{x-}$, the block \mathcal{F}_I^J and the vector \mathbf{g}_J read

$$\mathcal{F}_I^J = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \mathcal{E}_I^J & \mathcal{T}_I^J & \mathcal{T}_I^J & \mathcal{T}_I^J \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{g}_J = \begin{bmatrix} g_J^{x-} \\ g_J^{y-} \\ g_J^{x+} \\ g_J^{y+} \end{bmatrix}. \quad (31)$$

Algorithm 1: Application of the SGS preconditioner: $\mathbf{r} \leftarrow \tilde{\mathcal{F}}_{\text{SGS}}^{-1} \mathbf{r}$.

```

// Forward sweep (application of  $\mathcal{L}^{-1}$ )
for  $S = 1 : (N_{gr} - 1)$  do
  |  $\mathbf{r}_{S+1} \leftarrow \mathbf{r}_{S+1} - \mathcal{F}_{[S+1]}^{[S]} \mathbf{r}_S$ 
end

// Backward sweep (application of  $\mathbf{U}^{-1}$ )
for  $S = N_{gr} : 2$  do
  |  $\mathbf{r}_{S-1} \leftarrow \mathbf{r}_{S-1} - \mathcal{F}_{[S-1]}^{[S]} \mathbf{r}_S$ 
end

```

Algorithm 2: Application of the SGS or DS preconditioner: $\mathbf{r} \leftarrow \tilde{\mathcal{F}}^{-1} \mathbf{r}$.

```

// Forward sweep
for  $S = 1 : (N_{gr} - 1)$  do
  parfor each  $I$  such that  $\Omega_I \subset \Omega_{[S]}$  do
    Configure boundary data for subproblem on  $\Omega_I$ :
     $u_D \leftarrow 0$  on  $\partial\Omega_I \cap \Gamma_D$ 
     $g_I^\square \leftarrow r_I^\square$  on each interface edge  $\Gamma_I^\square \not\subset \partial\Omega_{[S+1]}$ 
     $g_I^\square \leftarrow r_I^\square$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S+1]}$  (If SGS prec.)
     $g_I^\square \leftarrow 0$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S+1]}$  (If DS prec.)

    Compute  $u_I$  by solving subproblem on  $\Omega_I$ .

    Update transmission data for subproblems of the next group:
     $r_J^{\square'} \leftarrow r_J^{\square'} - r_I^\square + 2\mathcal{B}_I^\square u_I$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S+1]}$  (If SGS prec.)
     $r_J^{\square'} \leftarrow r_J^{\square'} + 2\mathcal{B}_I^\square u_I$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S+1]}$  (If DS prec.)
    with  $J$  and  $\square'$  such that  $\Omega_J \subset \Omega_{[S+1]}$  and  $\Gamma_I^\square = \Gamma_J^{\square'}$ 
  end
end

// Backward sweep
for  $S = N_{gr} : 2$  do
  parfor each  $I$  such that  $\Omega_I \subset \Omega_{[S]}$  do
    Configure boundary data for subproblem on  $\Omega_I$ :
     $u_D \leftarrow 0$  on  $\partial\Omega_I \cap \Gamma_D$ 
     $g_I^\square \leftarrow r_I^\square$  on each interface edge  $\Gamma_I^\square \not\subset \partial\Omega_{[S-1]}$ 
     $g_I^\square \leftarrow r_I^\square$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S-1]}$  (If SGS prec.)
     $g_I^\square \leftarrow 0$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S-1]}$  (If DS prec.)

    Compute  $u_I$  by solving subproblem on  $\Omega_I$ .

    Update transmission data for subproblems of the previous group:
     $r_J^{\square'} \leftarrow r_J^{\square'} - r_I^\square + 2\mathcal{B}_I^\square u_I$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S-1]}$  (If SGS prec.)
     $r_J^{\square'} \leftarrow r_J^{\square'} + 2\mathcal{B}_I^\square u_I$  on each interface edge  $\Gamma_I^\square \subset \partial\Omega_{[S-1]}$  (If DS prec.)
    with  $J$  and  $\square'$  such that  $\Omega_J \subset \Omega_{[S-1]}$  and  $\Gamma_I^\square = \Gamma_J^{\square'}$ 
  end
end

```

corresponds to the application of \mathcal{L}^{-1} , can be interpreted as a forward sweep, where information are propagated across groups of increasing number. Similarly, the loop corresponding to the application of \mathbf{U}^{-1} can be interpreted as a backward sweep, where information are propagated across groups of




Figure 5: Illustration of the SGS preconditioner with horizontal sweeps and diagonal sweeps. The numbers correspond to the subdomain numbers. Red edges correspond to transmission variables that are updated. Blue edges correspond to transmission variables that are used in the update formula.

decreasing number. The sweeps are performed in the horizontal or diagonal direction not depending on the arrangement of the subdomains. Iterations of the backward sweep are illustrated on Figure 5 for horizontal sweeps and diagonal sweeps.

In the SGS preconditioner, the forward and backward sweeps must be performed sequentially. Nevertheless, in each iteration of both loops, the subproblems of a given group of subdomains can be solved in parallel. With the horizontal sweeps, the update of the transmission variables inside each group have been avoided since the diagonal part \mathcal{D} has been replaced with an identity matrix. With the diagonal sweeps, there is no coupling between the subdomains of the same group since there is no shared edge between them.


4.2 Parallel Double Sweep (DS) preconditioner

With the parallel Double Sweep (DS) preconditioner, the \mathcal{L} and \mathcal{U} matrices are modified in such a way that applying one of them does not modify the elements of the vector used by the other. The forward and backward sweeps can then be performed in parallel, without data race, which potentially reduces the runtime per iteration by a factor two in parallel environments.

The DS preconditioner can be written as $\tilde{\mathcal{F}}_{\text{DS}} = \tilde{\mathcal{L}}\tilde{\mathcal{U}} = \tilde{\mathcal{U}}\tilde{\mathcal{L}} = \tilde{\mathcal{L}} + \tilde{\mathcal{U}} - \mathcal{I}$. To remove the dependences between \mathcal{L} and \mathcal{U} , the blocks are modified in such a way that, for a given group of subdomains $\Omega_{[S]}$, the forward sweep does not use transmission data from edges shared with a subdomain of $\Omega_{[S+1]}$ (these data are modified in the backward sweep), and the backward sweep does not use transmission data from edges shared with a subdomain of $\Omega_{[S-1]}$ (these data are modified in the forward sweep). The effective update process is illustrated in Figure 6 for both horizontal and diagonal sweeps. The application of the DS preconditioner on a vector is detailed in Algorithm 2. The main difference with the SGS preconditioner is that one or more transmission variables are cancelled when solving the subdomains. Thanks to this modification, the forward and backward sweeps can be performed in parallel.

In order to illustrate the modification of the blocks of \mathcal{L} and \mathcal{U} , we consider the 3×3 domain partition

(a) Horizontal sweeps. Resp. $\bar{u}_{(3)}^{-1}, \bar{u}_{(2)}^{-1}$



(b) Diagonal sweeps. Resp. $\bar{u}_{(5)}^{-1}, \bar{u}_{(4)}^{-1}$ and $\bar{u}_{(3)}^{-1}$




Figure 6: Illustration of the DS preconditioner with horizontal sweeps and diagonal sweeps. The numbers correspond to the subdomain numbers. Red edges correspond to transmission variables that are updated. Blue edges correspond to transmission variables that are used in the update formula.

with the diagonal-type arrangement. The blocks corresponding to the coupling of $\Omega_{[2]}$ and $\Omega_{[3]}$ reads

$$\mathcal{F}_{[3]}^{[2]} = \begin{bmatrix} \mathcal{F}_4^2 & 0 \\ \mathcal{F}_5^2 & \mathcal{F}_5^3 \\ 0 & \mathcal{F}_6^3 \end{bmatrix} = \begin{bmatrix} \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \mathcal{T}_4^2 & \mathcal{T}_4^2 & \mathcal{E}_4^2 & 0 & \emptyset & 0 & 0 \\ \emptyset & 0 & 0 & 0 & 0 & \emptyset & 0 & 0 \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \hline \emptyset & \mathcal{T}_5^2 & \mathcal{E}_5^2 & \mathcal{T}_5^2 & 0 & \emptyset & 0 & 0 \\ \emptyset & 0 & 0 & 0 & \mathcal{T}_5^3 & \emptyset & \mathcal{T}_5^3 & \mathcal{E}_5^3 \\ \emptyset & 0 & 0 & 0 & 0 & \emptyset & 0 & 0 \\ \emptyset & 0 & 0 & 0 & 0 & \emptyset & 0 & 0 \\ \hline \emptyset & 0 & 0 & 0 & \mathcal{T}_6^3 & \emptyset & \mathcal{E}_6^3 & \mathcal{T}_6^3 \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & 0 & 0 & 0 & 0 & \emptyset & 0 & 0 \end{bmatrix} \quad (40)$$

and

$$\mathcal{F}_{[2]}^{[3]} = \begin{bmatrix} \mathcal{F}_2^4 & \mathcal{F}_2^5 & 0 \\ 0 & \mathcal{F}_3^5 & \mathcal{F}_3^6 \end{bmatrix} = \begin{bmatrix} \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & 0 & 0 & \emptyset & 0 & 0 & 0 & 0 & 0 & \emptyset & \emptyset & \emptyset \\ \emptyset & 0 & 0 & \emptyset & \mathcal{E}_2^5 & \mathcal{T}_2^5 & \mathcal{T}_2^5 & \mathcal{T}_2^5 & 0 & \emptyset & \emptyset & 0 \\ \emptyset & \mathcal{E}_2^4 & \mathcal{T}_2^4 & \emptyset & 0 & 0 & 0 & 0 & 0 & \emptyset & \emptyset & 0 \\ \hline \emptyset & 0 & 0 & \emptyset & 0 & 0 & 0 & 0 & \emptyset & 0 & 0 & \emptyset \\ \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset & \emptyset \\ \emptyset & 0 & 0 & \emptyset & 0 & 0 & 0 & 0 & \mathcal{E}_3^6 & \emptyset & \emptyset & \mathcal{T}_3^6 \\ \emptyset & 0 & 0 & \emptyset & \mathcal{T}_3^5 & \mathcal{E}_3^5 & \mathcal{T}_3^5 & \mathcal{T}_3^5 & 0 & \emptyset & \emptyset & 0 \end{bmatrix}, \quad (41)$$

where rows and columns with \emptyset correspond to items on boundary edges which must be removed. These blocks belong to \mathcal{L} and \mathcal{U} , respectively. The modified blocks $\tilde{\mathcal{F}}_{[2]}^{[3]}$ and $\tilde{\mathcal{F}}_{[3]}^{[2]}$ are obtained by removing

the terms in gray. In $\tilde{\mathcal{F}}_{[3]}^{[2]}$, we cancel the terms in the 3rd, 4th, 7th and 8th columns, corresponding to transmission variables on right edge and top edge of Ω_2 and Ω_3 , shared with subdomains of the next group. In $\tilde{\mathcal{F}}_{[2]}^{[3]}$, we cancel the terms in the 2nd, 5th, 6th and 9th columns, corresponding to transmission variables on bottom edge of Ω_4 , left edge and bottom edge of Ω_5 and left edge of Ω_6 , shared with subdomains of the previous group. The modified blocks verify $\tilde{\mathcal{F}}_{[2]}^{[3]} \tilde{\mathcal{F}}_{[3]}^{[2]} = 0$.

4.3 Flexible preconditioners and parallel aspects

With both SGS and DS preconditioners, the forward and backward sweeps are performed in a direction that only depends on the blocks $-\mathcal{F}_{[I]}^{[J]}$: a horizontal direction in which the blocks in the same column are performed parallelly, and a diagonal direction in which the blocks in the same diagonal are performed parallelly. With the standard version of GMRES, only one sweeping direction must be selected. Nevertheless, in practical situations, it could be advantageous to combine different sweeping directions in order to propagate information more rapidly in different zones of the computational domain. This can be achieved thanks to the flexible version of GMRES, called F-GMRES [39, 40], where a different preconditioner can be used at each iteration. Therefore, the DS and SGS preconditioners can be used with sweeping directions that shall be modified in the course of the iterations, possibly accelerating the convergence of the iterative procedure.

The final computational procedure contains operations that can be performed simultaneously, allowing to use parallel compute architectures. The forward and backward loops are intrinsically sequential, as the different groups of subdomains have to be treated successively in a specific order. Nevertheless, parallelism can be found inside each iteration of these loops: Subproblems defined on subdomains of the same group can be solved in parallel with both preconditioners. In addition, with the DS preconditioner, the forward and backward sweeps can be performed in parallel, as discussed in the previous section. For applications requiring computations with multiple right-hand sides, strategies can also be used to accelerate the computations by solving all the problems in parallel instead of successively.

In distributed-memory parallel environments, novel questions are raised, as the placement of the subdomains on the processors shall influence the communications and the parallel efficiency. If only one subdomain is placed on each processor, all the processors will be waiting most of the time because of the sequential nature of the sweeping process. Strategies to improve the parallel efficiency have been discussed in [47] for layered-type partitions. Checkerboard partitions offer novel possibilities, as groups of subdomains can be placed on each processor, and different kind of groups can be chosen. Placement strategies have been discussed in [45] in the context of the L-sweeps preconditioners. Here, for instance, it can be advantageous to place one row of subdomains on each processor when using diagonal sweeps. This strategy could improve the parallel efficiency by reducing the waiting time of processors.

5 Computational results

In this section, the preconditioners are studied and compared by using several two-dimensional benchmarks solved with a high-order finite element method. We consider scattering benchmarks with a single source (section 5.1) and multiple sources (section 5.2), the Marmousi problem (section 5.3) and acoustic radiation from engine intake (section 5.4). The computational results presented in this article have been obtained with a single multi-core processor. Only the solution of each subproblem was performed using shared-memory parallelism.

5.1 Scattering problem with a single source

We consider the scattering of a plane wave by a sound-soft circular cylinder of radius equal to 1. The scattered field is computed on a two-dimensional square domain of size 12.5×12.5 , which is partitioned into an grid of 5×5 square subdomains. The scattering disk is placed in the middle of the subdomain that is at the left-down corner of the grid (first configuration) or the subdomain in the middle of the domain (second configuration). The Dirichlet boundary condition $u(\mathbf{x}) = -e^{ikx}$ is prescribed on the boundary

of the disk, and the Padé-type HABC is prescribed on the exterior border of the computational domain with compatibility conditions at the corners [32]. The Padé-type HABC operator is also used in the transmission conditions prescribed at the interfaces between the subdomains with a suited cross-points treatment (see section 2.3). Using a large number of auxiliary fields in the Padé-type HABC operator can improve the accuracy of the numerical solution and the efficiency of the DDM for scattering problems [4, 32]. The number of auxiliary fields $N = 8$ and the parameter $\phi = \pi/3$ are used for both exterior and transmission conditions. The wavenumber k is 2π and the wave length λ is 1.

The solution is computed using a standard high-order nodal finite element method on meshes made of triangles and generated with Gmsh [20]. Two numerical settings have been considered: P1 finite elements with 20 mesh vertices per wavelength (mesh elements of size $h = 1/20$), and P7 finite elements with 3 elements per wavelength ($h \approx 1/21$). The meshes of these two settings are made of 97868 nodes, 3966 P7 triangles and 70619 nodes, 139984 P1 triangles, respectively. The linear system resulting from the finite element discretization is solved using preconditioned versions of GMRES. Both the symmetric Gauss-Seidel (SGS) and the parallel double sweep (DS) preconditioners have been tested with three strategies:

1. Diagonal sweeps: The forward sweep goes from the left-down corner to the right-up corner of the partition, and the backward sweep does it the other way around. [SGS-D and DS-D]
2. Horizontal sweeps: The forward sweep goes from the left to the right of the partition, and the backward sweep does it the other way around. [SGS-H and DS-H]
3. Two diagonal sweeping directions are used in alternance with the flexible version of GMRES (FGMRES). The sweeps are performed between the left-down and right-up corners and between the left-up and right-down corners, alternatively. [SGS-2D and DS-2D]

First configuration: Source close to the corner of the domain

In the first configuration, the scattering disk is placed in the subdomain that is at the left-down corner of the grid. Figure 7 shows snapshots of the solutions after the first GMRES iterations with the different preconditioners. For all the preconditioners with diagonal sweeps, the solution is already good after only one iteration, which is due to two successful strategies for this specific case. First, the HABC operator used in the transmission conditions is particularly well-suited for scattering benchmarks. Then, the first sweep over the subdomains goes from the left-bottom corner to right-up corner, which follows the natural behavior of waves in this benchmark. Therefore, we get all correct information in all subdomains after the first iteration. For the preconditioners with horizontal sweeps, relying on round-trips between left boundary and right boundary, we can see that we get the complete solution only at fourth iteration.

The residual histories obtained with the different preconditioners are shown in Figure 8 for finite element schemes with P1 and P7 elements, respectively. These results confirm the visual interpretations. In both cases, the relative residual suddenly drops in residual at the first iteration when a preconditioner with diagonal sweeps is used (*i.e.* SGS-D, SGS-2D, DS-D and DS-2D), while it happens at the fourth iteration with horizontal sweeps (*i.e.* SGS-H and DS-H). Without preconditioner, eight iterations are required to reach the sudden drop in residual because the waves have to go through eight subdomains to travel from the left-down corner to the right-up corner.

By comparing the results obtained with P1 and P7 elements, we observe that the residual histories are very similar before the sudden drop in residual with both kinds of finite elements. Nevertheless, the sudden drops are much sharper and the residuals decrease more rapidly after the sudden drop with the P7 elements. After the sudden drop in residual, we observe that the decay of the residual is nearly twice slower with the DS preconditioner than with the SGS preconditioner when diagonal sweeps is used. This must be balanced with the fact that the SGS preconditioner is intrinsically a sequential procedure, while the DS preconditioner relies on two sweeps that can be done in parallel.

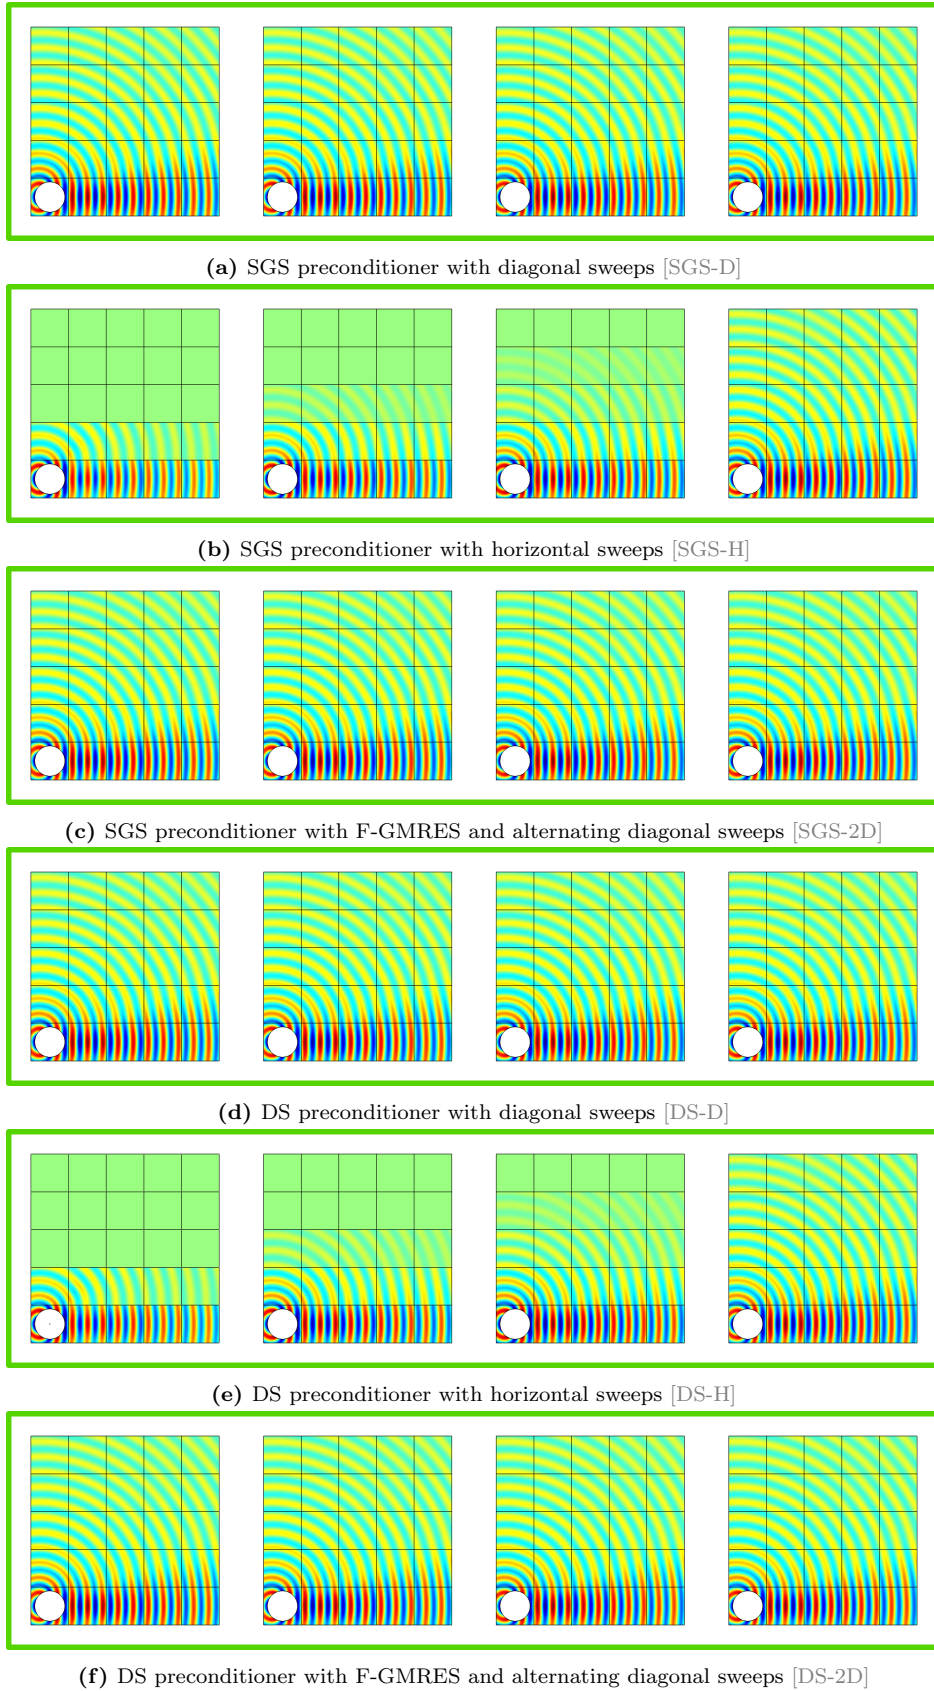


Figure 7: Scattering problem with a single source (first configuration). Snapshot of the solution after 1, 2, 3 and 4 GMRES iterations with the different preconditioners.




Figure 8: Scattering problem with a single source (first configuration). Residual history with/without preconditioner with P1 elements and 20 mesh vertices by wavelength (left) and with P7 elements and 3 elements by wavelength (right). Preconditioners with diagonal sweeps (SGS-D/DS-D), horizontal sweeps (SGS-H/DS-H) and alternating diagonal sweeps (SGS-2D/DS-2D) are considered.

Second configuration: Source in the middle of the domain

In the second configuration, the scattering disk is placed in the middle of the computational domain. Figure 9 shows snapshots of the solutions after the first GMRES iterations with the different preconditioners. First, let us focus on the solutions obtained after the very first iteration. On the snapshots, we see that the zone of influence of the source corresponds to subdomains that are mainly along the diagonal direction or along the horizontal direction starting from the center of the domain. The propagation of the source in these subdomains is obviously related to the use of preconditioners with diagonal sweeps and horizontal sweeps of the subdomains, respectively.

By contrast with the first configuration, we have different results when the SGS and DS preconditioners are used with the diagonal sweeps or the F-GMRES strategy. Because the forward sweeps and backward sweeps of DS-D do not affect each other, the subdomains on the right-up and left-down corners cannot be reached by the source after the first iteration, which is visible on Figures 9 (d) and (f). By contrast, these subdomains can be reached during the backward sweep of the SGS preconditioner, which is performed after the forward sweep (Figures 9 (a) and (c)). With diagonal sweeps, four iterations are required to get a good solution in all the subdomains with the DS preconditioner, while only two iterations are required with the SGS preconditioner. When the strategy with F-GMRES is used, the solution is very good at the second iteration, even with the DS preconditioner, because the sweeps of successive iterations are performed along both diagonal directions (Figures 9 (b) and (e)).

The residual histories obtained with the different preconditioners are shown in Figure 10 for finite element schemes with P1 and P7 elements, respectively. In all the cases, the relative residual decreases slowly until a sudden drop, which happens when the source has been propagated in all the subdomains, and when the numerical solution is close to the converged solution. The results confirm the visual observations. With the SGS preconditioner, the sudden drop occurs at the second iteration if the diagonal sweeps or the strategy with F-GMRES is used, and a third iteration is required with horizontal sweeps. With the DS preconditioner, four iterations are necessary with diagonal sweeps, while alternating diagonal sweeps realized by F-GMRES still requires two iterations.

When the source is placed in an arbitrary position in the computational domain, using flexible preconditioners, with several alternative sweeping directions, can be much more suitable than fixed preconditioners. Here, both SGS and DS preconditioners perform very well with the alternating diagonal sweeps. With P1 finite elements, the number of iterations to get the relative residual 10^{-6} is twice

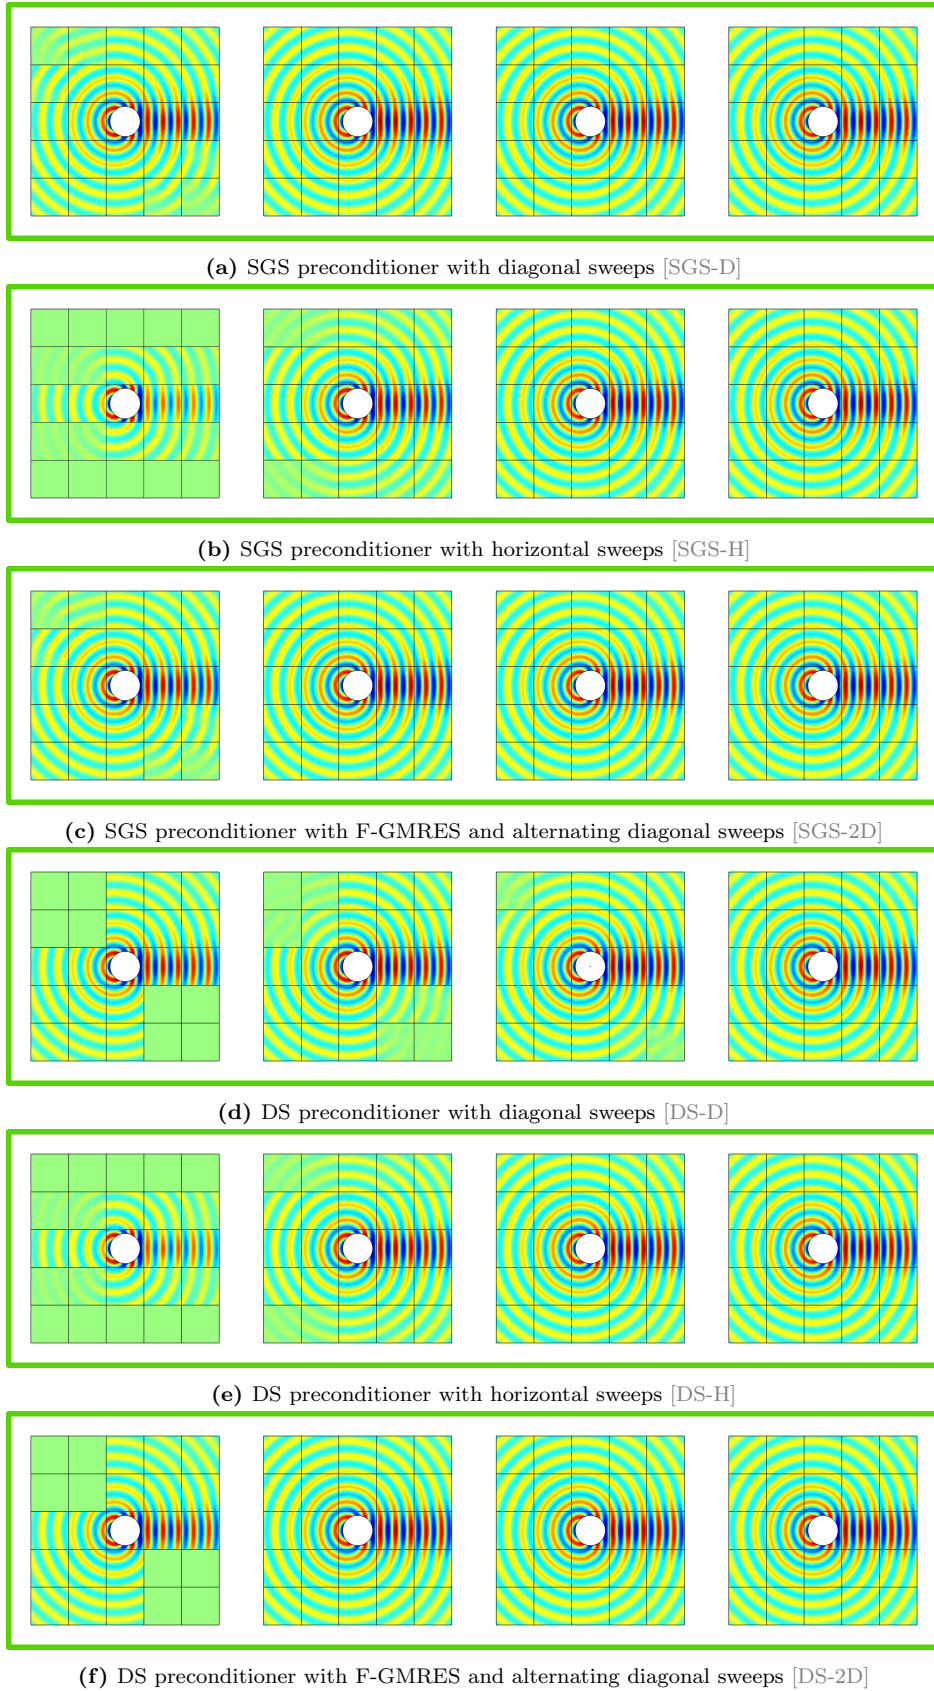


Figure 9: Scattering problem with a single source (second configuration). Snapshot of the solution after 1, 2, 3 and 4 GMRES iterations with the different preconditioners.

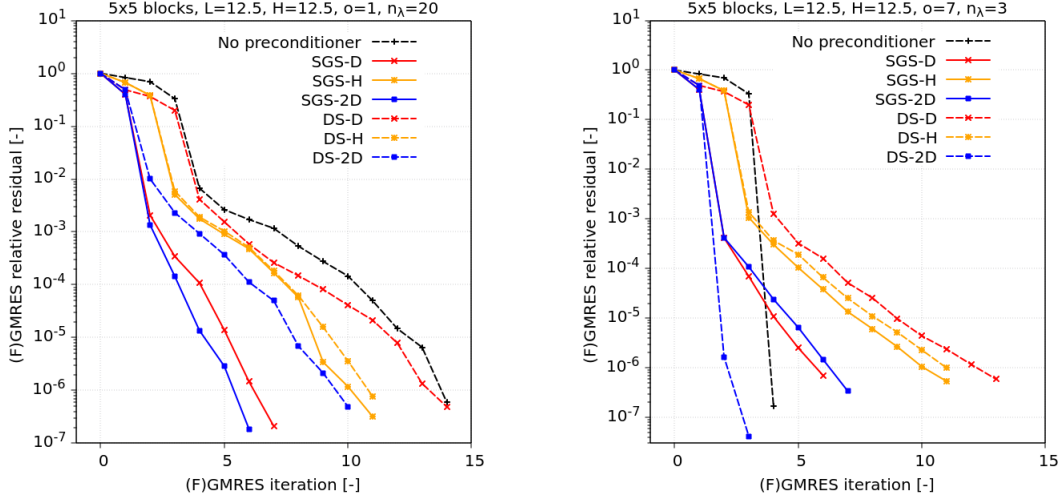


Figure 10: Scattering problem with a single source (second configuration). Residual history with/without preconditioner with P1 elements and 20 mesh vertices by wavelength (left) and with P7 elements and 3 elements by wavelength (right). Preconditioners with diagonal sweeps (SGS-D/DS-D), horizontal sweeps (SGS-H/DS-H) and alternating diagonal sweeps (SGS-2D/DS-2D) are considered.

larger with the DS preconditioner than with the SGS preconditioner. Because the sweeps of the DS preconditioner can be done in parallel, providing a speed up of 2, both SGS and DS approaches seem equivalent. By contrast, the number of iterations is similar with P7 finite elements. The parallel DS preconditioner then is more interesting for that case.

5.2 Scattering problem with multiple sources

In this section, we consider the scattering of a plane wave by two sound-soft circular cylinders of unit radius. This problem is more challenging for the DDM than the previous one because the multiple reflections between both obstacles can be complicated to capture. The simulations are performed over square grids of $N_r \times N_c$ subdomains, with $N_r = N_c = 4, 8, 12, 16$ and 20. The dimension of each subdomain is 2.5×2.5 . For each configuration, the scattering disks are placed at the left-down corner and the right-down corner of the grid. The physical and numerical parameters are the same as in the previous section. The numerical solutions for the 4×4 and 5×5 configurations are shown in Figure 11.

The numbers of GMRES iterations to reach a relative residual 10^{-6} with the different preconditioners are given in Figure 12 for different domain partitions. Contrary with the single obstacle case, the physical solution cannot be obtained in only one or two sweeps anymore because of the multiple reflections between both obstacles. We observe that the number of iterations increases with the number of subdomains when the preconditioners are used with a fixed horizontal or diagonal sweeps. By contrast, the strategy with the alternating diagonal sweeps keep the number of iterations constant, both with SGS and DS preconditioners. This is mainly due to the position of the disks: each of the two alternating diagonal sweeps is well suited to one of the sources. The number of iteration is slightly lower with the SGS preconditioner than with the DS preconditioner, but the latter will be more interesting in parallel environments considering that the forward and backward sweeps can be performed concurrently.

5.3 Marmousi benchmark

The Marmousi model is a 2D velocity model which is based on the geological structure of the Cuanza basin. This model exhibits a complex velocity profile $c(\mathbf{x})$ with realistic features (see Figure 13). It is frequently used to evaluate the performance of numerical solvers with heterogeneous media (*e.g.* [42]).




Figure 11: Scattering problem with two sources. Snapshot of the solution for configurations with 4×4 and 5×5 subdomains.




Figure 12: Scattering problem with two sources. Number of iterations with preconditioned GMRES and F-GMRES (without restart) to reach the relative residual 10^{-6} for different domain partitions.

The numerical simulations are performed over the computational domain¹ $[0, 9192] \times [0, -2904]$. The Helmholtz equation is solved over the domain, with the HABC prescribed on the boundary of the domain, and two point sources placed at coordinates $(9192/8, -10)$ and $(9192 \times 7/8, -10)$, respectively. The angular frequency is $\omega = 20\pi$, and the wavenumber is given by $k(\mathbf{x}) = \omega/c(\mathbf{x})$. Here, the maximum wavenumber is $k_{\max} = 20\pi/1500 \approx 0.042$ and the minimum wavenumber is $k_{\min} = 20\pi/4500 \approx 0.014$. In this case, the wavenumber is much smaller than 1 and the pollution term in [22] affects little. Therefore, the numerical setting is only P1 triangular elements with 20 mesh vertices per wavelength. The mesh of the computational domain is made of 137808 nodes and 274018 P1 triangles. The parameters of the HABC operator are $N = 4$ and $\phi = \pi/3$ for both the exterior boundary condition and the transmission conditions prescribed at the interfaces between the subdomains. The numerical solution is shown in Figure 14.

The numbers of GMRES iterations to reach a relative residual 10^{-6} with the different preconditioners are given in Figure 15. We have considered domain decompositions into rectangular grids of $N_r \times N_c$ subdomains, where the number of rows is $N_r = 3, 9, 15, 21, 27, 33$ and the number of columns is $N_c = 3N_r$. We observe that the number of iterations increases with the number of subdomains with

¹We assume that all the spatial dimensions are provided in the metric unit [m].




Figure 13: Marmousi benchmark. Velocity profile with values from 1500 m/s to 5500 m/s.




Figure 14: Marmousi benchmark. Numerical solution and domain partition with 3×9 subdomains.




Figure 15: Marmousi benchmark. Number of iterations with preconditioned GMRES and F-GMRES (without restart) to reach the relative residual 10^{-6} for different domain partitions. The domain is partitioned into $N_r \times N_c$ subdomains, with $N_r = 9, 15, 21, 27, 33$ and $N_c = 3N_r$.

all the preconditioners, but the increase is very slow with F-GMRES and the switching sweeping directions. Between the coarsest domain partition (3×9 subdomains) and the finest partition (33×99 subdomains), the number of iterations has increased by a factor between 5 and 9 with the preconditioners with fixed sweeping directions (SGS-D, SGS-H, DS-D, DS-H), while the factor is smaller than 2 with the switching sweeping directions (SGS-2D, DS-2D). In nearly all the cases, the SGS preconditioner with switching sweeping directions requires the smallest number of iterations, but, considering the possible parallelisation of the forward/backward sweeps with the DS preconditioner, that strategy is the best if a parallel environment is used.




Figure 16: Benchmark “Engine intake”. Representation of the computational domain and boundary conditions.

5.4 Acoustic radiation from engine intake

Description of the benchmark and domain partition

In the last benchmark, we address the computation of a time-harmonic acoustic field in a computational domain that is not rectangular. It deals with the aeroacoustics of an idealized turbofan engine intake. The geometry, shown in Figure 16, is a cylindrical duct of slowly-varying cross-section. The 2D Helmholtz equation is solved on this computational domain, which is included inside the rectangular region $[-0.3, 3.0] \times [0.0, 2.5]$. We consider a Dirichlet condition on the source, $u|_{\text{source}} = \sin((2\pi/0.50) \cdot y)$, homogeneous Neumann boundary condition on the hardwalls, and an HABC on the artificial borders where the waves must be radiated (see Figure 16). For the numerical solution, P5 finite elements are used with 5 elements per wavelength. The mesh of the computational domain is made of about 3.5×10^7 nodes and 2.8×10^6 P5 triangles for wavenumber 160π ($h \approx 1/2000$). The parameters of the HABC operators are $N = 4$ and $\phi = \pi/3$ for both interface and exterior edges. The numerical solution corresponding to wavenumber $k = 160\pi$ is shown on Figure 17.

Because the domain is not rectangular, additional steps are required to apply the proposed sweeping preconditioners, which are designed a priori only for checkerboard partitions. We have generated domain partitions of the rectangular region that contains the computational domain. This process is performed with Gmsh before the mesh generation. Then, every partition contains $N_r \times N_c$ rectangular subdomains (see Figure 18), but several “null” subdomains are fully outside the computational domain (*e.g.* subdomain Ω_J on Figure 18, right), and several subdomains are crossed by the border of the computational domain (*e.g.* subdomains Ω_I , Ω_L and Ω_K on Figure 18, right). After discretization, no unknown is associated to the null subdomains. The number of unknowns is smaller for subdomains that are crossed by the domain border than for rectangular subdomains that are fully contained inside the computational domain. Similarly, the number of discrete transmission variables is smaller on interface edges crossed by the domain boundary. In practice, the solution procedure is performed by iterating over all the subdomains, whatever they are inside or outside the computational domain. Dummy systems and dummy vectors of variables are associated to the null subdomains, and dummy transmission variables are associated to interface edges which do not belong to the computational domain (*i.e.* red dashed edges on Figure 18). Therefore, the sweeping preconditioner and our computational code can be straightforwardly used for this benchmark.

Results

Figure 19 presents the snapshot of the solution with 48×36 subdomains at the beginning of the procedure from iteration 0 to iteration 3 with preconditioners SGS-2D. The wavenumber k is 80π in Figure 19 instead of 160π considering that the visualisation is more clearly visible. At the first iteration, partial information is obtained by sweeps that go from the left-down corner to the right-up corner and the other is missed at the top of the computational domain. At the second iteration, the information at the top




Figure 17: Benchmark “*Engine intake*”. Snapshot of the numerical solution for wavenumber $k = 160\pi$.




Figure 18: Benchmark “*Engine intake*”. Mesh and example of domain partition.

of the computational domain is got added, which is contributed by alternating sweeps. At the third iteration, more details in the information is presented in the computational domain.

The number of iterations and the runtime to reach a relative residual 10^{-6} with the different preconditioners are given in Figure 20. Simulations are carried out on a Intel Xeon Phi (CPU 7210@1.30GHz) and parallelized using OpenMP interface. The runtime corresponds to the restarted (F)GMRES resolution phase (number of restart = 20). The number of threads is equal to the number of rows of subdomains.

The convergence rate with SGS-2D is the fastest. The results are comparable with DS-2D. With both preconditioners, the number of iterations is stable.

Comparing flexible preconditioners and fixed preconditioners, we can also see that switching preconditioners improve robustness of DDMs. SGS-H and DS-H are not good enough. We can see that the number of iterations with both preconditioners increases with the number of rows of subdomains. This indicates that SGS-H and DS-H might be suitable for long geometries.

Comparing all the solvers, the DDMs with switching preconditioners perform best. When the number of subdomains increases, the runtime is smaller with switching preconditioners than with the others. The reason is that the number of iterations is smaller with switching preconditioners. There isn't significant improvement on timing with fixed preconditioners SGS-D or DS-D. Although both preconditioners reduce the number of iterations, the inner steps of these preconditioners at each iteration are time-consuming.

6 Conclusion

We have proposed and compared multidirectional sweeping preconditioners for the finite element solution of Helmholtz problems with checkerboard domain decompositions. The domain decomposition algorithm relies on high-order transmission conditions and a cross-point treatment proposed in [33]. This algorithm is well suited to Helmholtz problems with checkerboard partitions of the computational domain, but it cannot scale with the number of subdomains without an efficient preconditioning technique.

While most of the sweeping techniques have been studied for layered domain partitions, we have presented generalizations for checkerboard partitions, offering flexibility in the choice of the sweeping




Figure 19: Benchmark “*Engine intake*”. Snapshot of the solution at the beginning of the procedure (iteration 0) and after 1, 2 and 3 iterations with SGS-2D. The wavenumber is 80π .

directions. Horizontal, vertical and diagonal sweeping directions can be used with symmetric Gauss-Seidel and parallel double sweep preconditioners. Several directions can be combined by using the flexible version of GMRES. For applicative cases, these preconditioners provide an efficient way to rapidly transfer information in the different zones of the computational domain, then accelerating the convergence of iterative solution procedures with GMRES. We have observed that the diagonal sweeping directions, with flipping between each iteration of the flexible GMRES, were particularly efficient in all the cases.

The multidirectional sweeping preconditioners can be straightforwardly extended to three dimensions. Thanks to the block representation of the interface system (equation (35)), every kind of sweeping

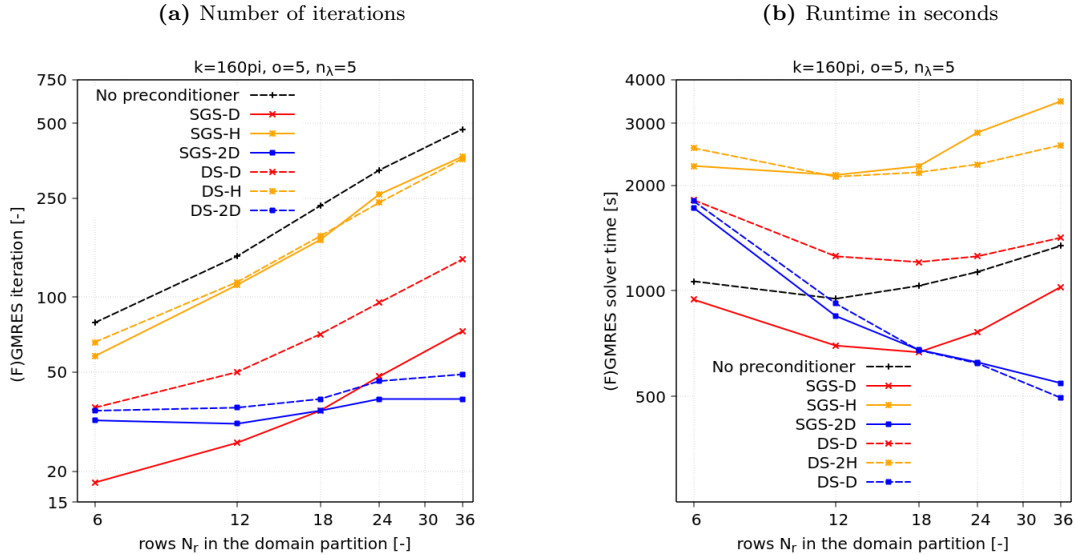


Figure 20: Benchmark “Engine intake”. Number of iterations (a) and runtime in seconds (b) with preconditioned GMRES and F-GMRES (with restart) to reach the relative residual 10^{-6} for different domain partitions. The domain is partitioned into $N_r \times N_c$ subdomains, with $N_r = 6, 12, 18, 24, 36$ and $N_c = 4N_r/3$.

preconditioner can be applied, such as the symmetric successive over-relaxation (SOR) preconditioner. A parallel quadruple sweep preconditioner, with four sweeps performing simultaneously in the $+x$, $-x$, $+y$ and $-y$ directions, can also be designed as a generalization of the parallel DS preconditioner. In future works, these preconditioners will be tested to problems with multiple right-hand sides, and in distributed-memory parallel environments, where novel questions are raised for parallel efficiency.

7 Acknowledgements

This work was carried out in the framework of the project “Large Scale Simulation of Waves in Complex Media” funded by the Communauté Française de Belgique under contract ARC WAVES 15/19-03.

References

- [1] A. V. Astaneh and M. N. Guddati. A two-level domain decomposition method with accurate interface conditions for the Helmholtz problem. *International Journal for Numerical Methods in Engineering*, 107(1):74–90, 2016.
- [2] J.-D. Benamou and B. Desprès. A domain decomposition method for the Helmholtz equation and related optimal control problems. *Journal of Computational Physics*, 136(1):68–82, 1997.
- [3] N. Bootland, V. Dolean, P. Jolivet, and P.-H. Tournier. A comparison of coarse spaces for Helmholtz problems in the high frequency regime. *arXiv preprint arXiv:2012.02678*, 2020.
- [4] Y. Boubendir, X. Antoine, and C. Geuzaine. A quasi-optimal non-overlapping domain decomposition algorithm for the Helmholtz equation. *Journal of Computational Physics*, 231(2):262–280, 2012.
- [5] Y. Boubendir and D. Midura. Non-overlapping domain decomposition algorithm based on modified transmission conditions for the Helmholtz equation. *Computers & Mathematics with Applications*, 75(6):1900–1911, 2018.
- [6] N. Bouziani, H. Calandra, and F. Nataf. An overlapping splitting double sweep method for the Helmholtz equation, 2021.
- [7] F. Collino, P. Joly, and M. Lecouvez. Exponentially convergent non overlapping domain decomposition methods for the helmholtz equation. *ESAIM: M2AN*, 54(3):775–810, 2020.
- [8] L. Conen, V. Dolean, R. Krause, and F. Nataf. A coarse space for heterogeneous Helmholtz problems based

- on the Dirichlet-to-Neumann operator. *Journal of Computational and Applied Mathematics*, 271:83–99, 2014.
- [9] A. de La Bourdonnaye, C. Farhat, A. Macedo, F. Magoules, F.-X. Roux, et al. A non-overlapping domain decomposition method for the exterior Helmholtz problem. *Contemporary Mathematics*, 218:42–66, 1998.
- [10] B. Després. Domain decomposition method and the Helmholtz problem. In G. Cohen, L. Halpern, and P. Joly, editors, *Proceedings of the First International Conference on Mathematical and Numerical Aspects of wave Propagation Phenomena (Strasbourg, France)*, pages 44–52. SIAM, 1991.
- [11] B. Després. *Méthodes de décomposition de domaine pour les problèmes de propagation d’ondes en régime Harmonique. Le théorème de Borg pour l’équation de Hill vectorielle*. PhD thesis, Paris VI University, 1991.
- [12] B. Engquist and A. Majda. Absorbing boundary conditions for numerical simulation of waves. *Proceedings of the National Academy of Sciences*, 74(5):1765–1766, 1977.
- [13] B. Engquist and L. Ying. Sweeping preconditioner for the Helmholtz equation: moving perfectly matched layers. *Multiscale Modeling & Simulation*, 9(2):686–710, 2011.
- [14] O. G. Ernst and M. J. Gander. Why it is difficult to solve Helmholtz problems with classical iterative methods. In *Numerical analysis of multiscale problems*, pages 325–363. Springer, 2012.
- [15] C. Farhat, P. Avery, R. Tezaur, and J. Li. FETI-DPH: a dual-primal domain decomposition method for acoustic scattering. *Journal of Computational Acoustics*, 13(03):499–524, 2005.
- [16] C. Farhat, P.-S. Chen, F. Risler, and F.-X. Roux. A unified framework for accelerating the convergence of iterative substructuring methods with lagrange multipliers. *International Journal for Numerical Methods in Engineering*, 42(2):257–288, 1998.
- [17] M. Gander, F. Magoules, and F. Nataf. Optimized Schwarz methods without overlap for the Helmholtz equation. *SIAM Journal on Scientific Computing*, 24(1):38–60, 2002.
- [18] M. J. Gander. Optimized Schwarz methods. *SIAM Journal on Numerical Analysis*, 44(2):699–731, 2006.
- [19] M. J. Gander and 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):3–76, 2019.
- [20] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [21] T. Hagstrom, R. P. Tewarson, and A. Jazcilevich. Numerical experiments on a domain decomposition algorithm for nonlinear elliptic boundary value problems. *Applied Mathematics Letters*, 1(3):299–302, 1988.
- [22] F. Ihlenburg and I. Babuska. Finite element solution of the helmholtz equation with high wave number part ii: The h-p version of the fem. *SIAM Journal on Numerical Analysis*, 34(1):315–358, 1997.
- [23] R. Kechroud, X. Antoine, and A. Soulaïmani. Numerical accuracy of a Padé-type non-reflecting boundary condition for the finite element solution of acoustic scattering problems at high-frequency. *International Journal for Numerical Methods in Engineering*, 64(10):1275–1302, 2005.
- [24] S. Kim and H. Zhang. Optimized Schwarz method with complete radiation transmission conditions for the Helmholtz equation in waveguides. *SIAM Journal on Numerical Analysis*, 53(3):1537–1558, 2015.
- [25] D. Lahaye, J. Tang, and K. Vuik. *Modern solvers for Helmholtz problems*. Springer, 2017.
- [26] M. Lecouvez, B. Stupfel, P. Joly, and F. Collino. Quasi-local transmission conditions for non-overlapping domain decomposition methods for the Helmholtz equation. *Comptes Rendus Physique*, 15(5):403–414, 2014.
- [27] W. Leng and L. Ju. A diagonal sweeping domain decomposition method with source transfer for the helmholtz equation, 2020. Preprint arXiv 2002.05327.
- [28] E. L. Lindman. “Free-space” boundary conditions for the time-dependent wave equation. *Journal of Computational Physics*, 18(1):66–78, 1975.
- [29] P.-L. Lions. On the Schwarz alternating method III: a variant for nonoverlapping subdomains. In *Third International Symposium on Domain Decomposition Methods for Partial Differential Equations*, 1990.
- [30] N. Marsic and H. D. Gersem. Convergence of optimized non-overlapping Schwarz method for Helmholtz problems in closed domains, 2020. Preprint arXiv 2001.01502.
- [31] F. A. Milinazzo, C. A. Zala, and G. H. Brooke. Rational square-root approximations for parabolic equation algorithms. *The Journal of the Acoustical Society of America*, 101(2):760–766, 1997.
- [32] A. Modave, C. Geuzaine, and X. Antoine. Corner treatments for high-order local absorbing boundary conditions in high-frequency acoustic scattering. *Journal of Computational Physics*, 401:109029, 2020.
- [33] A. Modave, A. Royer, X. Antoine, and C. Geuzaine. A non-overlapping domain decomposition method with high-order transmission conditions and cross-point treatment for Helmholtz problems. *Computer Methods in Applied Mechanics and Engineering*, 368:113162, 2020.
- [34] F. Nataf. On the use of open boundary conditions in block Gauss-Seidel methods for the convection-diffusion

- equations. Technical report, CMAP (Ecole Polytechnique), 1993.
- [35] F. Nataf and F. Nier. Convergence rate of some domain decomposition methods for overlapping and nonoverlapping subdomains. *Numerische mathematik*, 75(3):357–377, 1997.
 - [36] F. Nataf, F. Rogier, and E. de Sturler. Optimal interface conditions for domain decomposition methods. Technical report, CMAP (Ecole Polytechnique), 1994.
 - [37] A. Nicolopoulos. *Formulations variationnelles d'équations de Maxwell résonantes et problèmes aux coins en propagation d'ondes*. PhD thesis, Sorbonne Université, 2019.
 - [38] A. Piacentini and N. Rosa. An improved domain decomposition method for the 3D Helmholtz equation. *Computer Methods in Applied Mechanics and Engineering*, 162(1-4):113–124, 1998.
 - [39] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing*, 14(2):461–469, 1993.
 - [40] Y. Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
 - [41] A. Schädle and L. Zschiedrich. Additive Schwarz method for scattering problems using the PML method at interfaces. In *Domain Decomposition Methods in Science and Engineering XVI*, pages 205–212. Springer, 2007.
 - [42] C. C. Stolk. A rapidly converging domain decomposition method for the Helmholtz equation. *Journal of Computational Physics*, 241:240–252, 2013.
 - [43] C. C. Stolk. An improved sweeping domain decomposition preconditioner for the Helmholtz equation. *Advances in Computational Mathematics*, 43(1):45–76, 2017.
 - [44] B. Stupfel. Improved transmission conditions for a one-dimensional domain decomposition method applied to the solution of the Helmholtz equation. *Journal of Computational Physics*, 229(3):851–874, 2010.
 - [45] M. Taus, L. Zepeda-Núñez, R. J. Hewett, and L. Demanet. L-sweeps: A scalable, parallel preconditioner for the high-frequency helmholtz equation. *Journal of Computational Physics*, 420:109706, 2020.
 - [46] A. Vion and C. Geuzaine. Double sweep preconditioner for optimized Schwarz methods applied to the Helmholtz problem. *Journal of Computational Physics*, 266:171–190, 2014.
 - [47] A. Vion and C. Geuzaine. Improved sweeping preconditioners for domain decomposition algorithms applied to time-harmonic Helmholtz and Maxwell problems. *ESAIM: Proceedings and Surveys*, 61:93–111, 2018.
 - [48] L. Zepeda-Núñez and L. Demanet. The method of polarized traces for the 2D Helmholtz equation. *Journal of Computational Physics*, 308:347–388, 2016.