Hybrid Domain Decomposition Solvers for the Helmholtz Equation

Martin Huber[1] and Joachim Schöberl[1]

1 Introduction

When solving the Helmholtz equation with standard finite elements, the oscillatory behavior of the solution results in a large number of degrees of freedom (DoFs) required to resolve the wave, especially for high wave numbers. This together with the indefiniteness of the problem makes an iterative solution of the resulting linear system of equations difficult. Nevertheless, some advances for finding efficient preconditioners for wave type problems have been made recently. Well known is the shifted Laplace Preconditioner [5], or a sweeping preconditioner [4] based on an approximate block $LDL^\top$ factorization, which is constructed layer by layer. Especially for parallel computing platforms domain decomposition methods are very popular. Apart from optimized Schwarz methods [8], i.e. Schwarz methods which rely on optimal transmission conditions as interface condition, the FETI-H [7] and the FETI-DPH [6] method are widely used. The last two methods can be seen as further developments of the FETI and the FETI-DP methods, respectively, specialized for Helmholtz problems.

The solution strategies presented in this work are based on a mixed hybrid discontinuous galerkin formulation [13, 10] Since the hybrid formulation provides appropriate interface conditions an efficient iterative solution with Krylov space methods combined with domain decomposition preconditioners is possible. Apart from adapting a BDDC preconditioner [3, 11] to the current setting, a new Robin type domain decomposition preconditioner is constructed. This preconditioner solves in each iteration step local problems on subdomains by directly inverting the subdomain matrix. Thus, it is well suited for parallel computations. Good convergence properties of both preconditioners are demonstrated by numerical experiments. The results of this paper will be presented in more detail in [9].

2 The Mixed Hybrid Discontinuous Galerkin Formulation

Our formulation is based on the mixed form of the Helmholtz equation: Find a scalar function $u : \Omega \to \mathbb{C}$ and a vector valued function $\mathbf{p} : \Omega \to \mathbb{C}^d$

$$\operatorname{grad} u = \mathrm{i}\omega \mathbf{p} \qquad \text{and} \qquad \operatorname{div} \mathbf{p} = \mathrm{i}\omega u$$

[1]Institute for Analysis and Scientific Computing, Wiedner Hauptstrasse 8-10, A-1040 Wien, e-mail: {martin.huber}{joachim.schoeberl}@tuwien.ac.at

with an absorbing boundary condition $-\mathbf{p}\cdot\mathbf{n}+u=g$ on $\Gamma:=\partial\Omega$. As computational domain $\Omega\subset\mathbb{R}^d$ with $d=2,3$ a Lipschitz polyhedron is considered. Furthermore, $\mathbf{n}$ denotes the outer normal vector, the angular frequency ω is a positive constant and $g\in L^2(\Gamma)$. Note that [12] guarantees a unique solution.

In this paper we make use of the following notations. By $\mathscr{T}$ a triangulation with the elements T is denoted. The set of its facets F we call $\mathscr{F}$, $\mathbf{n}_F$ represents the normal vector onto a facet F, and $\mathbf{n}_T$ is the outer normal vector of an element T. Furthermore volume integrals are denoted by $(u,v)_T := \int_T u\bar{v}\,\mathrm{d}\mathbf{x}$ and surface integrals by $\langle u,v\rangle_{\partial T} := \int_{\partial T} u\bar{v}\,\mathrm{d}s$.

In order to obtain efficient solvers for the Helmholtz equation, we consider it in a mixed hybrid form. Thus, we search for $(u,\mathbf{p},u_F,p_F)\in L^2(\Omega)\times H(\mathrm{div},\mathscr{T})\times L^2(\mathscr{F})\times L^2(\mathscr{F})=:U\times V\times U_F\times V_F$ such that for all $(v,\mathbf{q},v_F,q_F)\in U\times V\times U_F\times V_F$

$$\sum_{T\in\mathscr{T}}\Big((i\omega u,v)_T-(\mathrm{div}\mathbf{p},v)_T-(u,\mathrm{div}\mathbf{q})_T-(i\omega\mathbf{p},\mathbf{q})_T+\langle u_F,\mathbf{n}_T\cdot\mathbf{q}\rangle_{\partial T} \qquad (1)$$

$$+\langle\mathbf{n}_T\cdot\mathbf{p},v_F\rangle_{\partial T}+\langle\mathbf{n}_F\cdot\mathbf{p}-p_F,\mathbf{n}_F\cdot\mathbf{q}-q_F\rangle_{\partial T}\Big)-\langle u_F,v_F\rangle_\Gamma=-\langle g,v_F\rangle_\Gamma.$$

This mixed hybrid formulation was introduced and discussed in [13]. In the formulation the space $H(\mathrm{div},\mathscr{T})$ represents an element wise $H(\mathrm{div})$ space without continuity constraints across element interfaces, and $L^2(\mathscr{F})$ is the space of L^2 functions on the facets. Consequently u_F and p_F are supported just on the facets, and they represent the values of u and $\mathbf{p}\cdot\mathbf{n}_F$ there. The problem is discretized by the finite dimensional spaces

$$U_h := \prod_{T\in\mathscr{T}} P_k(T), \qquad V_h := \prod_{T\in\mathscr{T}} RT_k(T),$$

$$U_{Fh} := \prod_{F\in\mathscr{F}} P_k(F), \qquad V_{Fh} := U_{Fh},$$

where polynomials of order k are denoted as P_k and RT_k represents a Raviart-Thomas element of order k. The discrete solutions we call u_h, $\mathbf{p}_h$, u_{Fh} and v_{Fh}, respectively.

Since there is no global coupling for the functions u_h and $\mathbf{p}_h$ across different elements, the corresponding DoFs can be eliminated cheaply on the element level via static condensation [1]. Note that this elimination corresponds on each element to the solution of a wave type problem with Robin boundary conditions, and uniqueness is guaranteed. The resulting linear system of equations needs now to be solved only for the facet DoFs.

Remark 1. The original form of equation (1) in [13] contains a penalty parameter η, which was chosen to be one in our work. For this choice the local problem on the element, which needs to be solved during static condensation, corresponds to the original problem posed on the domain with $g=\pm p_F+u_F$. The sign depends on the direction of the facet normal $\mathbf{n}_F$. Thus, g represents now the incoming impedance trace for the element.

In this work domain decomposition preconditioners will be used to solve the reduced linear system of equations for the facet unknowns u_{Fh} and p_{Fh}, which is obtained by eliminating the volume unknowns u_h and $\mathbf{p}_h$. Note that this linear system of equations is related to the skeleton of a mesh, and a domain decomposition of the skeleton is induced by a decomposition of the underlying mesh. Since impedance traces are obtained from the facet unknowns by a simple transformation of variables, transmission conditions on the interface in the sense of [2] can be enforced by guaranteeing the same value of the facet unknowns of different subdomains on the subdomain interface. Thus the mixed hybrid formulation allows in a natural way for appropriate transmission conditions for domain decomposition preconditioners.

3 The BDDC preconditioner for the mixed hybrid formulation

In this section, we adapt the BDDC preconditioner introduced by Dohrmann in [3] (compare also [11]) to wave type problems. Therefore a stabilization term has to be added to the mixed hybrid formulation, more precisely, the term

$$\sum_{T \in \mathscr{T}} \gamma \Big(\big\langle (\mathbf{n}_T \cdot \mathbf{n}_F) p_F, v_F \big\rangle_{\partial T \setminus \Gamma} + \big\langle u_F, (\mathbf{n}_T \cdot \mathbf{n}_F) q_F \big\rangle_{\partial T \setminus \Gamma} \Big), \qquad \gamma \in \mathbb{C} \qquad (2)$$

is added to (1). The parameter $\gamma \in \mathbb{C}$ is a tuning parameter, we choose based on numerical experiments. For the Helmholtz and the vector valued wave equation we made good experience with $\gamma = -0.5 - 0.1i$. These additional terms are just added for inner facets, and because of the different sign of $\mathbf{n}_T \cdot \mathbf{n}_F$ for the two neighboring elements, they cancel out when the global system of equations is assembled. Thus the problem does not change. But for domain decomposition preconditioners, which are based on submatrices assembled just for a subdomain the situation changes. These additional terms do not cancel out in the submatrices for DoFs located on the interface to other subdomains.

We use a BDDC preconditioner for this modified facet problem. The computational domain is divided into subdomains, and the DoFs on facets which just belong to one subdomain are considered to be primal, as well as the low order DoFs on interface facets. The high order DoFs on interface facets are the dual ones.

This choice leads to a large global system for the primal DoFs. Note that this system of equations consists due to the missing high order unknowns at the interface of weakly coupled subdomain blocks. Therefore it can be solved rather efficiently by direct solvers on parallel computing platforms.

4 A Robin type domain decomposition preconditioner

Like the BDDC preconditioner, the new Robin type domain decomposition (RDD) preconditioner will be applied to the skeleton problem, including the stabilizing terms (2) with the same value for γ.

Before describing the preconditioner, some notations are required. We assume, that the computational domain is divided into N subdomains Ω_i. For each subdomain a matrix A_i representing the subdomain problem is subassembled, and the global matrix A of the linear system of equations is obtained by adding these submatrices. By $\tilde{A}_i$ we denote the block of A_i which corresponds to DoFs on inner facets, i.e. facets which just belong to the domain Ω_i. The matrix $R^{(i)}$ restricts a vector to the components corresponding to these inner DoFs of the domain Ω_i. The matrix $R_D^{(i)}$ provides a weighted restriction to the domain Ω_i, i.e. when applying it, a vector entry is divided by the number of subdomains to which the corresponding DoFs belongs to. Note that an application of the prolongation matrix $R_D^{(i)\top}$ results again in a division for the interface DoFs. Thus, by summing up over all subdomains, a mean value on the interface can be created.

Using this notations, a RDD step for finding $\tilde{\mathbf{x}} := C_{RDD}^{-1}\mathbf{b}$ with $\mathbf{b}$ as right hand side of the linear system of equations and C_{RDD} as the preconditioner reads as

1) $\mathbf{y}_0 = \mathbf{0}$,
2) $\mathbf{y}_1 = \mathbf{y}_0 + \sum_{i=1}^N R^{(i)\top}\tilde{A}_i^{-1}R^{(i)}(\mathbf{b} - A\mathbf{y}_0)$,
3) $\mathbf{y}_2 = \mathbf{y}_1 + \sum_{i=1}^N R_D^{(i)\top}A_i^{-1}R_D^{(i)}(\mathbf{b} - A\mathbf{y}_1)$,
4) $\tilde{\mathbf{x}} = \mathbf{y}_2 + \sum_{i=1}^N R^{(i)\top}\tilde{A}_i^{-1}R^{(i)}(\mathbf{b} - A\mathbf{y}_2)$.

In step 2, the system of equations is solved exactly for the DoFs on the inner facets under the constraint that the solution on the interface is zero. Step 3 provides an update for the interface solution by partitioning the actual residual among the subdomains and solving the problem there exactly. A continuous interface solution is constructed by averaging the different subdomain solutions. Finally, in step 4 the solution is updated by solving the system of equations exactly for the DoFs on inner facets. Note that the interface solution remains unchanged.

The RDD-preconditioner can also be introduced in the variational projector notation. Therefore, we denote the bilinear form representing the Schur complement system, which is defined on the facet space $W := U_{Fh} \times V_{Fh}$ by a. Additionally, it is assumed that the bilinear form a can be decomposed into its subdomain contributions a_i, i.e. $a = \sum_{i=1}^N a_i$. The subspace of W containing the functions which are supported on the subdomain Ω_i is denoted by W_i, and in $\tilde{W}_i$ functions supported only on inner facets of the domain Ω_i are collected. The operator representation of the restriction matrix $R_D^{(i)}$ is called $\mathscr{R}_D^{(i)} : W \to W_i$. Thus, when applying it to any function in W, the function is restricted to the domain Ω_i, and its values on the interface facets are divided by the number of neighboring subdomains. Furthermore, $\mathscr{R}^{(i)} : W \to \tilde{W}_i$ is the restriction operator corresponding to the matrix $R^{(i)}$, and by $\mathscr{R}^{(i)\top}$ the prolongation operators are denoted.

Based on this, we define the variational projector $\mathscr{P}_D^{(i)}$ via $\mathscr{P}_D^{(i)} = \mathscr{R}_D^{(i)\top} \hat{\mathscr{P}}_D^{(i)}$ with the projector $\hat{\mathscr{P}}_D^{(i)} : W \to W_i$ and

$$a_i(\hat{\mathscr{P}}_D^{(i)} u, \phi) = a(u, \mathscr{R}_D^{(i)\top} \phi) \qquad \forall \phi \in W_i.$$

In the same way the variational projector $\mathscr{P}^{(i)}$ with $\mathscr{P}^{(i)} = \mathscr{R}^{(i)\top} \hat{\mathscr{P}}^{(i)}$ can be introduced. Here, $\hat{\mathscr{P}}^{(i)} : W \to \tilde{W}_i$ is given via

$$a_i(\hat{\mathscr{P}}^{(i)} u, \phi) = a(u, \mathscr{R}^{(i)\top} \phi) \qquad \forall \phi \in \tilde{W}_i.$$

If the operator $\mathscr{A}$ corresponds to the bilinear form a, and $\mathscr{I}$ is the identity, the error propagation operator $\mathscr{E}$ of the RDD-preconditioner reads as

$$\mathscr{E} = \mathscr{I} - \mathscr{C}_{RDD}^{-1}\mathscr{A} = \left(\mathscr{I} - \sum_{i=1}^{N} \mathscr{P}^{(i)}\right)\left(\mathscr{I} - \sum_{i=1}^{N} \mathscr{P}_D^{(i)}\right)\left(\mathscr{I} - \sum_{i=1}^{N} \mathscr{P}^{(i)}\right).$$

Remark 2. Because the system of equations is always solved exact for the DoFs on inner facets, both sets of facet DoFs u_{Fh} and p_{Fh} are not needed anymore, and the problem can be formulated just by using u_{Fh}. On the interface, both types of unknowns are still necessary in order to fix continuity conditions of the impedance traces across the interface and to guarantee convergence of the iterative solver. Nevertheless, neglecting one type of facet unknowns on inner facets saves many DoFs in an actual calculation.

5 Numerical Results

For all numerical examples which are presented in this section, we made good experience by taking a CG-solver, although, there exists no convergence theory for complex symmetric problems. We start the numerical results section by comparing the preconditioners for a simple two dimensional model problem. There, the Helmholtz equation is solved on a square $\Omega = [-1, 1]^2$ with an incoming wave from above of Gaussian amplitude, fixed by the absorbing boundary condition. The computations were done with the MPI-parallel finite element code Netgen/Ngsolve (see *http://sourceforge.net/projects/ngsolve* or [14]), which contains the software package Metis for partitioning the domain. If not said differently, a Dell R-910 Server (4 Xeon E7 CPUs with 10 cores a 2.2 GHz, 512 GB RAM) was used.

In Table 1 the iteration numbers for the BDDC and the RDD preconditioner for different wavelengths $\lambda := \frac{2\pi}{\omega}$ and mesh sizes h are given. For all computations the polynomial order was kept constant to four, and nine subdomains were used for the preconditioners.

According to the table, the BDDC preconditioner shows the highest iteration numbers close to the resolution limit at $h \approx \lambda$, which corresponds for a polynomial order of $p = 4$ to about four unknowns per wavelength. When increasing the number

Table 1 Iteration numbers of the BDDC/RDD preconditioner using 9 subdomains ($p = 4$) for different mesh sizes and wavelength λ.

λ	1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{32}$	$\frac{1}{64}$
$h = \frac{1}{4}$	45 / 78	49 / 65	60 / 61				
$h = \frac{1}{8}$	51 / 95	48 / 84	56 / 71	73 / 70			
$h = \frac{1}{16}$	56 / 123	50 / 96	49 / 83	59 / 73	80 / 72		
$h = \frac{1}{32}$	63 / 154	57 / 125	48 / 101	49 / 84	65 / 74	85 / 74	
$h = \frac{1}{64}$	66 / 202	62 / 164	56 / 127	50 / 111	50 / 89	74 / 82	101 / 89

of unknowns per wavelength, either by decreasing the mesh size or by increasing the wavelength, the number of iterations stays constant or grows slightly. For the RDD preconditioner the situation is vice versa. Although, for a large wavelength and a small mesh size the RDD preconditioner needs much more iterations than the BDDC preconditioner, it gets more and more competitive if the number of degrees of freedom per wavelength is reduced. Considering, that the RDD preconditioner is faster than the BDDC preconditioner with respect to setup-time and time per iteration, it is the method of choice for discretizations close to the resolution limit of the wave.

One reason for this behavior could be the different structure of the two solvers. While the RDD preconditioner allows just for local corrections, the BDDC solver benefits additionally from a coarse grid solution. For a decreasing wavelength, the solution gets more and more oscillatory, and the coarse grid correction, which provides communication across the whole subdomain loses its importance.

The number of iterations of the BDDC and the RDD preconditioner is also influenced by the size of the subdomains the computational domain is divided into. In Figure 1 iteration numbers of these two preconditioners are plotted for different wavelengths against the subdomain size H, both in logarithmic scale. Note that the partitioning of the domain was done by Metis, and therefore, H represents an average subdomain size. In the corresponding experiments the mesh size was kept constant to $\frac{1}{64}$ and the polynomial order to four. For the RDD preconditioner the number of iterations decreases with an increasing subdomain size. Figure 1 indicates that this decrease is proportional to $H^{-\alpha}$. According our experimental data α was estimated to be approximately 0.65. The situation is slightly different for the BDDC preconditioner. While it shows the same features for small wavelengths, i.e. for settings close to the resolution limit, the iterations stay almost constant for large wavelengths. A reason for this is, that for less oscillatory solutions the BDDC preconditioner benefits from its coarse grid correction.

Finally, we want to demonstrate the efficiency of our preconditioners with a three dimensional large scale example. The computational results presented in the following have been achieved using the Vienna Scientific Cluster 2 (VSC2). In this example, the solution of the Helmholtz equation for a grating (compare Figure 2) with period 0.14 was computed. The diameter of the computational domain was two. Thus, assuming a wave incoming from the top with Gaussian amplitude and wavelength
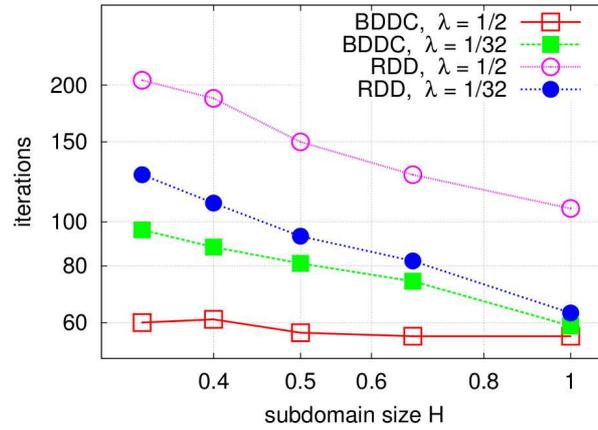
Fig. 1 Number of iterations plotted versus the size of one subdomain H for the BDDC and the RDD preconditioner. The polynomial order was 4 and $h = \frac{1}{64}$.
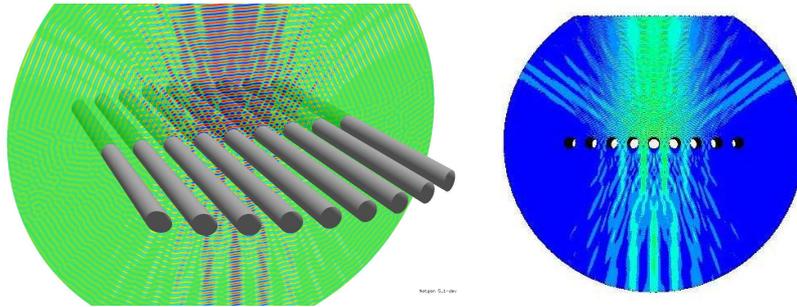
Fig. 2 Real part of the solution (left) and its absolute value (right) for a wave diffracted at a grating.

0.025 corresponds to an effective domain size of 80 wavelengths. For this setting, the left hand plot in Figure 2 shows the real part of the solution, and the absolute value is plotted on in the righthand plot. In the calculation, the underlying mesh had about 1.61 million elements with a maximal mesh size of 0.021. Selecting a polynomial order of $p = 4$ results in approximately 288.8 million volume unknowns (56.5 M. for u and 232.3 M. for $\mathbf{p}$) and 98.0 million facet unknowns (49.0 M. for u_F and p_F). Using 1200 subdomains, the assembly of the matrix took 58 seconds and the setup of the RDD preconditioner 33 seconds. The problem was solved in 12.9 min-

utes with 399 iterations on 1200 processors. Recovering the volume DoFs u_h and $\mathbf{p}_h$ from the facet DoFs u_{Fh} and p_{Fh} took 53 seconds.

References

1. Arnold, D., Brezzi, F.: Mixed and nonconforming finite element methods: Implementation, postprocessing and error estimates. RAIRO Model. Math. Anal. Numer. **19**(1), 7–32 (1985)
2. Desprès, B.: Méthodes de décomposition de domains pour les problèms de propagation d'ondes en régime harmonique. Phd thesis, Université Paris IX Dauphine (1991)
3. Dohrmann, C.: A preconditioner for substructuring based on constrained energy minimization. SIAM J. Sci. Comput. **25**(1), 246–258 (2003)
4. Engquist, B., Ying, L.: Sweeping preconditioner for the Helmholtz equation: Hierarchical matrix representation. Comm. Pure Appl. Math. **64**(5), 697–735 (2011)
5. Erlangga, Y.: Advances in iterative methods and preconditioners for the Helmholtz equation. Arch. Comput. Methods Eng. **15**(1), 37–66 (2008)
6. Farhat, C., Avery, P., Tezaur, R., Li, J.: FETI-DPH: a dual-primal domain decomposition method for acoustic scattering. J. Comput. Acoustics **13**(3), 499–524 (2005)
7. Farhat, C., Macedo, A., Lesoinne, M.: A two-level domain decomposition method for the iterative solution of high frequency exterior helmholtz problems. Numer. Math. **85**(2), 283–308 (2000)
8. Gander, M., Magoulès, F., Nataf, F.: Optimized Schwarz methods without overlap for the Helmholtz equation. SIAM J. Sci. Comput. **24**(1), 38–60 (2002)
9. Huber, M.: Hybrid discontinuous galerkin methods for the wave equation. Phd thesis, Vienna University of Technology (to appear)
10. Huber, M., Pechstein, A., Schöberl, J.: Hybrid domain decomposition solvers for the helmholtz and the time harmonic maxwell's equation. In: Domain Decomposition Methods in Science and Engineering XX, Lecture Notes in Computational Science and Engineering (to appear). Springer (2011)
11. Li, J., Widlund, O.: FETI-DP, BDDC, and block Cholesky methods. Int. J. Numer. Meth. Engrg. **66**(2), 250–271 (2006)
12. Melenk, J.: On generalized finite element methods. Phd thesis, University of Maryland (1995)
13. Monk, P., Sinwel, A., Schöberl, J.: Hybridizing Raviart-Thomas elements for the Helmholtz equation. Electromagnetics **30**(1), 149–176 (2010)
14. Schöberl, J.: NETGEN - an advanced front 2D/3D-mesh generator based on abstract rules. Comput. Vis. Sci. **1**(1), 41–52 (1997)