Efficient Simulation of High-Frequency Wave Propagation Using Domain Decomposition and High-Order FEM

Mahadevan Ganesh and Charles Morgenstern

Abstract

Large scale scientific computing models, requiring iterative algebraic solvers, are needed to simulate high-frequency wave propagation. This is because large degrees of freedom are needed to avoid the celebrated Helmholtz computer model pollution effects. Using low-order finite difference or finite element methods (FDM/FEM), such issues have been well investigated for low and medium frequency models (typically at most 50 wavelengths per diameter of the wave propagation domain). Standard FDM/FEM based discretizations of the time-harmonic Helmholtz wave propagation model lead to sign-indefinite systems with eigenvalues in the left half of the complex plane. Hence standard iterative methods (such as GMRES/BiCGstab) perform poorly, and additional techniques such as multigrid (MG) or decomposition of the domain are required for efficient and practical simulation of high-frequency FDM/FEM Helmholtz models. In this work, we investigate the use of multiple additive Schwarz type domain decomposition (DD) approximations to efficiently simulate high-frequency wave propagation with high-order FEM. We compare our DD based results with those obtained using a standard geometric MG approach for over 100 wavelength models.

1 Introduction

Various domain decomposition (DD) methods have been shown to be effective for numerous applications [7]. Recently there has been substantial interest in understanding applications of DD for simulation of sign-indefinite wave propagation models (see [5] and references therein). Most recent works have focused on simple cases with low-order methods on rectangular domains, and consider only low- to medium-frequency problems in numerical experiments.

For high-frequency models, standard low-order FDM/FEM models require a computationally prohibitive large number of degrees of freedom (DoF) for accurate solutions. Thus low-order computer models are not practical for high-frequency problems. In this article, we consider acoustic wave propagation in non-convex media with reentrant corners which are features that occur in practical applications. We implement a high-order FEM model to facilitate efficient simulation of high-frequency problems with 100 or more wavelengths per diameter of the computational domain.

The main contribution of this short note from a preliminary version of our article [3] is to develop a high-order FEM for a two dimensional homogeneous media wave propagation model in conjunction with several variants of the classical additive Schwarz

(AS) method for high-frequency problems. The full heterogeneous media article [3] will include high-order FEM models for efficient simulation of high-frequency wave propagation in both two and three dimensions. We demonstrate the efficiency of the algorithms by comparing with a similar geometric multigrid (GMG) approach [4].

We consider the following model for time-harmonic acoustic wave propagation. Let $k > 0$ be the constant wavenumber, and $\Omega \subset \mathbb{R}^2$ be a bounded Lipschitz domain with outward pointing unit normal vector ν and diameter L. Let $g \in L^2(\partial\Omega)$ be the absorbing boundary data, typically induced by the impinging incident wave (plane or point-source) from $\mathbb{R}^2 \setminus \overline{\Omega}$. The continuous model with an absorbing boundary for the acoustic wave propagation is to find $u \in H^1(\Omega)$ such that

$$\Delta u(\mathbf{x}) + k^2 u(\mathbf{x}) = 0, \quad \text{in } \Omega, \quad \text{and} \quad \frac{\partial u}{\partial \nu} - \mathrm{i}ku = g, \quad \text{on } \partial\Omega. \tag{1}$$

We define the wavelength as $\lambda = 2\pi/k$, refer the model (1) to be an X wavelength problem if $L = X\lambda$, and define the frequency of the model to be high if $X \geq 50$.

Standard finite element method (FEM) computer models for this problem require large degrees of freedom (DoF) for accurate solutions, and the DoF required increase as k increases. With the standard low-order FEM with spline degree $p = 1$ we require the mesh size $h \sim k^{-3/2}$ for pollution free solutions [5]. With this choice, the resulting DoF becomes prohibitive for direct methods as k becomes large. While the DoF are decreased, high-order FEM models also require iterative methods as the frequency of the problem becomes very large. The resulting system of equations are poorly conditioned and sign-indefinite, and traditional iterative methods are not possible or perform poorly. Thus efficient iterative methods for FEM computer models of Helmholtz problems is an active area of research. We decrease the required DoF in this paper by using a high-order FEM with $p = 4$ and develop an efficient DD iterative method framework.

The articles [1, 2, 4] (and references therein) consider preconditioning the system with a discrete version of the following complex shifted Laplacian wave propagation model with added absorption. As a preconditioner, we use the discrete system resulting from the Galerkin discretization of the standard variational formulation to the following problem. Find $u \in H^1(\Omega)$ such that

$$\Delta u(\mathbf{x}) + (k^2 + \mathrm{i}\,\mathcal{E})u(\mathbf{x}) = 0 \quad \text{in } \Omega, \quad \text{and} \quad \frac{\partial v}{\partial \nu} - \mathrm{i}kv = g \quad \text{on } \partial\Omega, \tag{2}$$

where $\mathcal{E} > 0$ is a free shift parameter.

Most recent work regarding the shifted Laplacian considers using a multigrid (MG) approximation of a discrete version of (2). With MG approximations of the preconditioner, it has been shown that the key to efficient simulation is a proper choice of the shift parameter $\mathcal{E}$. The recent work [2] provided analysis for the proper choice of the shift $\mathcal{E} \sim k$ in order for the shifted Laplacian to be a good preconditioner, requiring wavenumber independent GMRES iterations, for the original problem. However, this work does not apply when DD or MG are used to efficiently approximate the preconditioner. In the case of MG approximations, based on numerical experience, it is well known that a shift $\mathcal{E} \sim k^2$ provides optimal convergence of the preconditioned iterative method (see [1, 4]). However, analysis of this choice for MG is an open problem.

The recent preprint [5] suggests that additive Schwarz (AS) approximations of the shifted Laplacian problem in (2) with shift $\mathcal{E} \sim k^\alpha$ for $\alpha \in [0, 2]$ provide efficient preconditioners for the original model. Further using a standard low-order method, it has been shown in [5] that the choice $\alpha \in [0, 1.2]$ is optimal with a slight degradation as α approaches 2. This is surprising since $\alpha = 2$ is the common choice with MG. However, to achieve these results in [5], two-level AS methods are implemented with relatively fine coarse grid selection which makes their choice of DD implementation more computationally expensive than common MG approaches [4]. Further, the authors of [5] require the use of hybrid AS methods which are less parallelizable than classical AS methods.

In this work, we consider the proper choice of shift $\mathcal{E}$ when implementing hybrid and classical AS methods with parameter selections which provide both high and low computational costs for a high-order h-p computer model. We consider the use of overlapping and non-overlapping AS and hybrid AS type algorithms to approximate the preconditioner.

2 A High-Order FEM Model

Our FEM model is based on a high-order discretization of the standard variational formulations for the problem in (1) and the corresponding preconditioner problem from (2). For our model problems we choose the boundary data g in (1) and (2) such that a point source at $(3, 0)$ is the exact solution to the problem. In Figure 1 we provide simulations for the non-convex model problem.

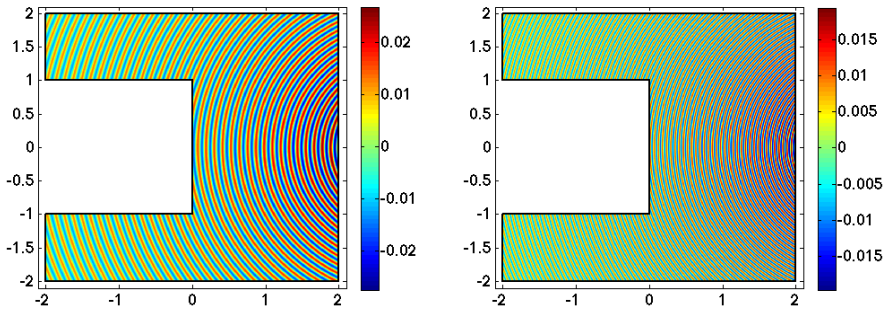

Figure 1: Simulations of the real part of the total field for the model problem with 50 and 100 wavelengths per diameter of Ω (left and right)

The FEM model is based on the following variational formulation of the continuous models in $V = H^1(\Omega)$. The variational formulation for (2) using the L^2 inner product notation is to find $u \in V$ such that

$$a_{\mathcal{E}}(u, v) = F(v), \quad \text{for all } v \in V, \quad \text{where} \tag{3}$$

$$a_{\mathcal{E}}(u,v) = \langle \nabla u, \nabla v \rangle_{L^2(\Omega)} - (k^2 + \mathrm{i}\,\mathcal{E}) \langle u, v \rangle_{L^2(\Omega)} - ik \langle u, v \rangle_{L^2(\partial\Omega)}, \quad F(v) = \langle g, v \rangle_{L^2(\partial\Omega)}. \tag{4}$$

For a chosen mesh parameter h, the discrete Ritz-Galerkin system for solving the problem is posed on a finite dimensional subspace $V_p^h \subset V$, spanned by WEB-spline [6] basis functions B_j of degree p for $j \in I^h$, where I^h is an index set corresponding to node points used to construct the basis. The FEM approximate solution u_h is then

$$u_h = \sum_{j \in I_h} c_j B_j \ , \qquad \text{where} \qquad a(u_h, v) = F(v), \qquad \text{for all} \qquad v \in V_h. \tag{5}$$

The second equation in (5) yields a system of linear equations for the unknown coefficients in first equation in (5). We denote the resulting system of equations in matrix form as $A_{\mathcal{E}}$. If $\mathcal{E} = 0$, the problem in (2) reduces to the original problem in (1), and we denote the sesquilinear form $a(\cdot, \cdot)$ and system of equations as A.

We use a non-standard high-order FEM approach based on multivariate basis splines on uniform Cartesian grids which facilitates simple implementation of geometric multigrid (GMG) [4]. This choice is natural for the model problem considered here, but our work can be extended to general complex shaped domains including curved geometries (see [4]). We omit details of the method in this short article and refer interested readers to [3, 4, 6] for two and three dimensional models.

With spline-degree $p = 4$ fixed, we consider mesh-grid width $h = (1/2)^m$ for $m \in \mathbb{N}$, and we desire to achieve relative $L^2(\Omega)$-norm error of approximately 10^{-3} or less (i.e. relative error of approximately 0.1% or less). We begin by demonstrating the required choice of m to achieve our desired tolerance in Tables 1 and 2. In Table 1 we consider $m = 6$, and we only achieve our desired tolerance with up to $L = 80\lambda$. However, in Table 2 we implement $m = 7$ and achieve our desired tolerance for problems with up to $L = 150\lambda$. With $h = (1/2)^7$ there are 4.83 and 7.24 elements per wavelength when $L = 150\lambda$ and 100λ respectively. This is much fewer elements per wavelength than is necessary for accurate results with standard low-order FEM.

Table 1: $L^2(\Omega)$-norm error with $p = 4$, and $h = (1/2)^6$ varying frequency

L	50λ	60λ	70λ	80λ	90λ	100λ
Error	8.6457e-05	2.4574e-04	6.5773e-04	1.8285e-03	5.4235e-03	1.6285e-02

Table 2: $L^2(\Omega)$-norm error with $p = 4$, and $h = (1/2)^7$ varying frequency

L	100λ	110λ	120λ	130λ	140λ	150λ
Error	8.8281e-05	1.5471e-04	2.6905e-04	4.7323e-04	8.5096e-04	1.5618e-03

If we implement the same grid width $h = (1/2)^7$ in Table 3 with $p = 1, 2$, and 3 for the 150 wavelength problem, we do not achieve our desired accuracy. However, in Table 3 we see that with the FEM implemented we have similar DoF in the FEM system for all choices of p, and this implies solving the FEM system has a similar computational cost. Thus our high-order FEM provides a significant decrease in the required DoF and the computational cost of solving the FEM system for accurate solutions to high-frequency problems.

Table 3: $L^2(\Omega)$-norm error with $L = 150\lambda$, and $h = (1/2)^7$, varying p.

p	1	2	3	4
Error	1.3591e+00	4.2992e-01	2.2751e-02	1.5618e-03
DoF	197,889	199,172	200,457	201,744

3 Domain Decomposition Algorithms

Consider the system of equations given by the matrix $A_{\mathcal{E}}$ which results from the high-order Galerkin FEM discretization of the shifted Laplacian problem (2) in the FEM subspace V_p^h corresponding to a fine grid width of h and WEB-spline basis degree p. Here we define the notation and method used for decomposing the FEM space V_p^h into N subspaces for the AS type DD algorithms.

Assume that we have a collection of open subsets $\{\tilde{\Omega}_\ell : \ell = 1, \cdots, N\}$ such that $\tilde{\Omega}_\ell \subset \mathbb{R}^d$, and the collection forms an overlapping cover of $\overline{\Omega}$. Details on the choice of $\tilde{\Omega}_\ell$ will be provide later. We are interested in the portion of the cover interior to $\overline{\Omega}$, so we define the collection $\{\Omega_\ell : \ell = 1, \cdots, N\}$ where $\Omega_\ell = \tilde{\Omega}_\ell \cap \overline{\Omega}$. For $\ell = 1, \cdots, N$ we assume that $\overline{\Omega}_\ell$ is non-empty and a union of elements from the fine grid intersected with the closure of the domain. For each Ω_ℓ the associated subspace of V_p^h is

$$V^\ell = \{v_h \in V_p^h : \text{supp}(v_h) \subset \overline{\Omega}_\ell\}.$$

The number of DoF for V^ℓ is the cardinality of the set $\{x_j : j \in I^h(\Omega_\ell)\}$, where $I^h(\Omega_\ell)$ is a suitable index set, and x_j corresponds to the lower left corner of support for the associated basis function.

In practice, for $\ell = 1, \cdots, N$, we choose the index set $I^h(\Omega_\ell)$ first, and this determines the collection $\{\Omega_\ell : \ell = 1, \cdots, N\}$. Further, we begin by choosing a non-overlapping collection $\{\hat{\Omega}_\ell : \ell = 1, \cdots, N\}$, and we carefully choose the overlap to maximize the amount of coupling retained from the original FEM system $A_{\mathcal{E}}$ while minimizing the added DoF in the subdomains. We choose $I^h(\hat{\Omega}_\ell)$ such that for all $j \in I^h(\hat{\Omega}_\ell)$ the lower left corner x_j of associated basis functions ϕ_j lies within a rectangle. We keep all the rectangles approximately the same size and denote the approximate maximum diameter of the subdomains as H_s. The parameter H_s determines the number of subdomains and the DoF in each subdomain. We plot an example choice of $I^h(\hat{\Omega}_\ell)$ for the model problem by plotting the corresponding x_j in different colors for $\ell = 1, \cdots, N$ in Figure 2.

After creating the non-overlapping decomposition, we choose the overlap as follows. To simplify notation we denote $I := I^h$ in the following to avoid excessive superscripts. Define a graph $G = (I, E)$ where $I = \{1, \cdots, n\}$ are the n indices corresponding to the DoF in the FEM system and E is the edge set such that $E = \{(i, j) | (A_{\mathcal{E}})_{i,j} \neq 0\}$. Using the standard variational formulation of the problem the non-zero pattern is symmetric, and thus G is undirected.

We denote $I_i^0 := I^h(\hat{\Omega}_i)$ to simplify the presentation. Our work above partitions the graph into N non-overlapping subsets I_ℓ^0 for $\ell = 1, \cdots, N$ such that $I = \bigcup_{\ell=1}^{N} I_\ell^0$.

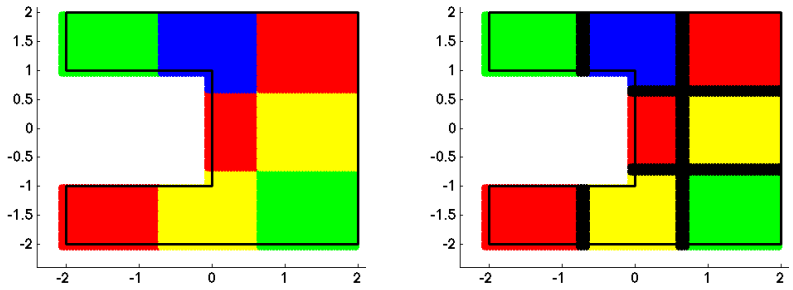

Figure 2: Subdomain choices for the model problem with $h = 0.015625$, $p = 4$, $H_s = 1$ and $\delta = 0$, and 1 (left and right). Points corresponding to DoF in different subdomains are plotted in different colors. Points corresponding to DoF included in multiple subdomains are plotted in black. The same color had to be used for multiple subdomains.

We now define I_i^1 for $i = 1, \cdots, N$ to be the one-overlap partition of I where for each $i = 1, \cdots, N$, $I_i^0 \subset I_i^1$, and I_i^1 is created by including all the immediate neighboring vertices in the graph to the vertices in I_i^0. We use this same method to further define the δ-overlap partition of I

$$I = \bigcup_{i=1}^{N} I_i^\delta,$$

where I_i^δ has δ levels of overlap with neighboring subdomains for integer $\delta \geq 0$.

Commonly the overlap is instead chosen geometrically by including in a non-overlapping subdomain neighboring DoF within a certain distance of the boundary of the subdomain (see [5]). However, our method chooses additional DoF for each non-overlapping subdomain based on the coupling of the system. Both methods do achieve similar results, but our method is well suited for high-order FEM. Since high-order FEM increases the density and coupling of the system $A_\mathcal{E}$ as p increases, the actual geometric width of the overlap created increases as p increases for fixed δ. If we were to use a geometric method to choose our overlap, we would have to manually increase the width of the overlap as p increased. We provide an example of an overlapping decomposition for the model problem in Figure 2.

Using the described subspaces and subdomains, we implement the following two-level algorithms: non-overlapping AS, and hybrid additive Schwarz (HAS), and overlapping AS, restricted additive Schwarz (RAS), averaged additive Schwarz (AVE), and hybrid restricted additive Schwarz (HRAS). These algorithms are well known, and we omit the details in this short article. We refer interested readers to [5] for full definitions of all the methods and to [7] for an overview of general DD methods.

Because we implement two-level methods, we require an additional choice of a coarse grid subspace and corresponding restriction and prolongation operators. Details of these choices are omitted here, interested readers are referred to [3, 4, 6]. We additionally use this same choice of coarse grid and grid transfer operators for our GMG implementation used for comparison in the numerical results.

3.1 Computational Cost

The key reason for using a DD or GMG [4] approximation of the preconditioner is to decrease the computational cost of calculating the preconditioner. Here we demonstrate the computational cost for the GMG and DD methods implemented in this article explicitly through tables, and we use this to justify our choices of the coarse grid width H_c and subdomain diameter H_s used in our numerical experiments.

We present the DoF in the coarse grid systems in Table 4. This same choice of coarse grid is used for both the GMG and DD approximations. The DoF in the fine grid with $h = (1/2)^7$ is $201,744$, as shown in Table 3. For our comparison of the GMG and DD methods we consider a two-level GMG with $H_c = (1/2)^6$ compared with DD methods using the coarse grid width $H_c = (1/2)^6$. We additionally consider a three-level GMG with coarsest grid width $H_c = (1/2)^5$ compared to DD methods with coarse grid width $H_c = (1/2)^5$. This provides equivalent DoF in the coarse grid where we implement a direct solve for both the GMG and DD methods.

Table 4: DoF in the FEM system with varying H_c for the model problem with $p = 4$.

H_c	$(1/2)^3$	$(1/2)^4$	$(1/2)^5$	$(1/2)^6$
DoF	1,104	3,728	13,584	51,728

Table 5: Maximum DoF in a subdomain and the number of subdomains N for the non-convex domain with $p = 4$, $h = (1/2)^7$, and varying H_s, and δ.

H_s	$(1/2)^0$	$(1/2)^1$	$(1/2)^2$	$(1/2)^3$
N	9	25	81	289
DoF $\delta = 0$	29,584	10,816	3,364	961
DoF $\delta = 1$	31,680	12,321	4,356	1,521

For the DD methods, the choice of subdomain diameter H_s provides an additional parameter to control the DoF in the subdomain systems. We consider the DoF required for the subdomain systems in Table 5. We choose H_s so that the resulting DoF in the subdomain system are comparable to the DoF in the coarse grid system, and this results in good load balancing since the subdomain and coarse grid systems can be solved in parallel for the non-hybrid DD algorithms. Thus for the cases where $H_c = (1/2)^6$ and $(1/2)^5$ we choose $H_s = (1/2)^0$ and $(1/2)^1$ respectively. This results in slightly fewer DoF in the subdomain systems compared with the coarse grid system as can be seen from comparing Tables 4 and 5.

When $H_c = (1/2)^6$ and $H_s = (1/2)^0$, for our classical AS preconditioners we must solve a system of dimension $51,728$ and nine systems of dimension $29,584$ when $\delta = 0$ or $31,680$ when $\delta = 1$, but these solves can be done completely in parallel. For the GMG algorithm we solve one system of dimension $51,728$ and require some additional work for smoothing on the fine grids. Similar analysis can be done for $H_c = (1/2)^5$ and $H_s = (1/2)^1$ using Tables 4–5. Thus with parallel implementation the cost of the DD and GMG methods are similar.

4 Numerical Results

We consider the performance of the DD approximations for the preconditioner through numerical evidence. For comparison we consider the use of a GMG approximation [4] of the preconditioner. For all results we fix the fine grid width $h = (1/2)^7$. For the one-level GMG algorithm with coarse grid width $H_c = (1/2)^6$, we use a $V(1,1)$-cycle. For the two-level MG algorithm with coarsest grid width $H_c = (1/2)^5$ we implement an $F(1,1)$-cycle motivated by the previous work [4]. For both algorithms we use weighted Jacobi smoothing with weight $\omega = 0.5$ and implement a direct solve on the coarsest grid. These choices of parameters and cycles were shown to be effective preconditioners for similar problems with $p = 3$ and $L \leq 100\lambda$ in [4].

Table 6: Number of BiCGstab iterations required when preconditioned with GMG, AS and HAS with $\delta = 0$ with coarse grid $H_c = (1/2)^6$ and subdomain diameter $H_s = (1/2)^0$ for fixed $h = (1/2)^7$ and $p = 4$.

$\mathcal{E} = k^2$						
L	100λ	110λ	120λ	130λ	140λ	150λ
GMG	267	310	339	375	398	486
AS	244	278	319	329	365	409
HAS	294	329	363	377	410	449
$\mathcal{E} = k^{3/2}$						
L	100λ	110λ	120λ	130λ	140λ	150λ
GMG	65	69	71	79	81	169
AS	31	31	32	34	36	41
HAS	30	30	30	32	33	45
$\mathcal{E} = k$						
L	100λ	110λ	120λ	130λ	140λ	150λ
GMG	31	35	36	44	62	>500
AS	28	27	28	29	38	50
HAS	8	8	10	14	42	84

In Table 6 we consider the results with fixed coarse grid $h = (1/2)^6$ for both the GMG and DD algorithms. For this initial comparison we consider only the non-overlapping DD algorithms, AS and HAS. Because our coarse grid is very fine for this case, we achieve optimal results for all algorithms with complex shift $\mathcal{E} < k^2$ even for the GMG preconditioner where $\mathcal{E} \sim k^2$ is the standard choice. The choice of $\mathcal{E} = k^{3/2}$ or k depends on the size of $H_c k$ and the choice of algorithm. We note that the fewest required iterations are obtained with the DD algorithms when compared to the GMG. The choice of AS or the hybrid HAS depends on the size of $H_c k$. However, the non-hybrid algorithms have similar computational cost to GMG when all subdomain system solves are done in parallel, and this is not possible with the hybrid HAS. Thus we prefer to use non-hybrid methods when the corresponding hybrid methods provide only slightly fewer iterations as with non-overlapping AS and HAS.

Table 7: Number of BiCGstab iterations required preconditioned with RAS, AVE, and HRAS with $\delta = 1$, coarse grid $H_c = (1/2)^6$ and subdomain diameter $H_s = (1/2)^0$ for fixed $h = (1/2)^7$ and $p = 4$.

$\mathcal{E} = k^{3/2}$						
L	100λ	110λ	120λ	130λ	140λ	150λ
RAS	29	33	30	31	32	35
AVE	29	30	29	31	31	35
HRAS	25	26	28	28	30	41
$\mathcal{E} = k$						
L	100λ	110λ	120λ	130λ	140λ	150λ
RAS	18	16	18	19	34	73
AVE	18	17	18	20	33	73
HRAS	5	5	6	11	25	70

The efficiency of the preconditioner can be further improved for fixed coarse grid $h = (1/2)^6$ by considering overlapping DD algorithms (see Table 7). We omit the results with shift $\mathcal{E} = k^2$ because similar to Table 6 the overlapping algorithms perform poorly with this choice. The most noticeable improvements compared to the non-overlapping algorithms in Table 6 are seen for the highest frequency problems when $H_c k$ becomes large for the RAS and AVE, and for the lowest frequency problems when $H_c k$ is small for HRAS. For problems up to 120 wavelengths HRAS is a competitive choice, and this is expected based on the recent work [5], but for a robust algorithm the AVE and RAS preconditioners provide efficient preconditioners for all frequencies considered. With all three preconditioners the optimal choice of $\mathcal{E}$ depends on the size of $H_c k$.

Table 8: Number of BiCGstab iterations required preconditioned with GMG, RAS, AVE, and HRAS with $\delta = 1$, coarse grid $H_c = (1/2)^5$, and subdomain diameter $H_s = (1/2)^1$ for fixed $h = (1/2)^7$, and $p = 4$.

$\mathcal{E} = k^2$						
L	100λ	110λ	120λ	130λ	140λ	150λ
GMG	263	275	328	385	380	432
RAS	452	451	467	432	463	496
AVE	433	456	445	441	460	>500
HRAS	298	281	320	327	345	387
$\mathcal{E} = k^{3/2}$						
L	100λ	110λ	120λ	130λ	140λ	150λ
GMG	>500	>500	>500	>500	>500	>500
RAS	78	109	98	71	73	91
AVE	76	91	101	78	74	81
HRAS	83	90	99	63	66	77

We conclude this article in Table 8 by comparing our preferred overlapping DD algorithms with GMG with coarse grid $H_c = (1/2)^5$. This choice provides less computationally expensive approximation of the preconditioner compared to Tables 6–7, and in practice we would like to choose H_c as large as possible and H_s as small as possible.

We see in Table 8 that with less DoF in the coarsest grid the GMG now requires the standard choice of $\mathcal{E} \sim k^2$ in order for BiCGstab to converge in less than 500 iterations. However, the DD methods perform best when $\mathcal{E} = k^{3/2}$. Because $H_c k$ is increased all algorithms perform poorly when $\mathcal{E} = k$, and thus we omit these results. We see a significant decrease in the required iterations using DD algorithms compared to GMG. We emphasize again that with parallel processing the computational cost of the GMG, AVE, and RAS methods is similar, but HRAS is more expensive because the coarse and subdomain solves cannot be done in parallel. Thus simulating high-frequency problems with our high-order FEM algorithm, we prefer to use DD methods to traditional GMG to approximate the shifted Laplacian preconditioner. In particular we see the most efficient results with the overlapping RAS and AVE preconditioners.

References

[1] Y.A. Erlangga, C.W. Oosterlee, and C. Vuik. A novel multigrid based preconditioner for heterogeneous Helmholtz problems. *SIAM J. Sci. Comput.*, 27:1471–1492, 2006.

[2] M.J. Gander, I.G. Graham, and E.A. Spence. Applying GMRES to the Helmholtz equation with shifted Laplacian preconditioning: what is the largest shift for which wavenumber-independent convergence is guaranteed? *Numer. Math.*, 131:567–614, 2015.

[3] M. Ganesh and C. Morgenstern. High-order FEM domain decomposition models for high-frequency wave propagation in heterogeneous media. *Under Preparation*, 2016.

[4] M. Ganesh and C. Morgenstern. An efficient multigrid algorithm for heterogeneous acoustic media sign-indefinite high-order FEM models. *Numer. Linear Algebr (Revised version under review)*, 2016
(`http://inside.mines.edu/~mganesh/final_mg_revised.pdf`).

[5] I.G. Graham, E.A. Spence, and E. Vainikko. Domain decomposition preconditioning for high-frequency Helmholtz problems using absorption. *preprint*, 2015.

[6] K. Höllig. *Finite Element Methods with B-Splines*. Number 26 in Frontiers in Applied Mathematics. SIAM, Philadelphia, 2003.

[7] A. Toselli and O. Widlund. *Domain Decomposition Methods – Algorithms and Theory*. Springer, 2005.