Extension of Two-level Schwarz Preconditioners to Symmetric Indefinite Problems
TR2008-914

Alan Leong

New York University, Courant Institute of Mathematical Sciences
251 Mercer Street
New York, NY 10012
September 2008

A thesis submitted in partial fulfillment
of the requirements for the degree of
master's of science
Department of Mathematics
New York University
September 2008

Advisor: Olof Widlund

Abstract

Two-level overlapping Schwarz preconditioners are extended for use for a class of large, symmetric, indefinite systems of linear algebraic equations. The focus is on an enriched coarse space with additional basis functions built from free space solutions of the underlying partial differential equation. GMRES is used to accelerate the convergence of preconditioned systems. Both additive and hybrid Schwarz methods are considered and reports are given on extensive numerical experiments.

1 Introduction

Iterative solvers can provide a significant advantage in the computational and storage requirements of solving sparse linear systems. They are distinct from direct solvers in that iterative solvers repeatedly apply an operator until a convergence criterion is met rather than fully factoring the matrix as would be done for a direct solver. Such sparse systems commonly arise from finite element solutions of partial differential equations (PDEs), see [1]. Krylov space methods comprise a significant class of iterative solvers. These solvers function by calculating an optimal approximation in a Krylov space, with the dimension of the Krylov space being enlarged at each iteration, see [12, §6.2]. Two prominent examples are the conjugate gradient (CG) algorithm and the generalized minimal residual (GMRES) algorithm. These methods are applicable to symmetric positive definite and general systems respectively. A significant drawback, compared to direct solvers, is a much greater sensitivity to the conditioning of the problem. To overcome this difficulty, preconditioning techniques can be used to construct a well conditioned system from an ill-conditioned system, without altering the solution. Two goals of preconditioning are increasing the speed and robustness of an iterative solver. Preconditioning aids in attaining uniform performance regardless of the conditioning of the non-transformed system, making the preconditioned solver more robust [12, §10]. Domain decomposition methods are a class of preconditioning techniques especially applicable to linear systems arising from PDEs approximated by finite elements. They provide a division of the problem into many subproblems, each significantly smaller than the full problem. This makes domain decomposition a good option for use with parallel computing [3, §1].

The domain decomposition methods that will be discussed involve two levels of problems. First, a coarse level problem is used to capture global information over the full domain. Then the domain is partitioned into subdomains and a fine triangulation is used to associate a local problem to each subdomain. The finite element tearing and interconnecting (FETI), see [6] [13], and balancing domain decomposition by constraints (BDDC), see [8] [13], methods utilize a partition of the domain into non-overlapping subdomains. Two-level Schwarz methods involve decomposing the domain into overlapping subdomains [12][13].

Many well developed and effective domain decomposition methods require that the linear system be symmetric, positive definite. Most of the theoretical results pertaining to classical two-level Schwarz methods require that the

problem be positive definite. Some physical systems are modeled using PDEs that give rise to indefinite systems. Additionally, inverse iteration is a common method of approximating eigenvectors. Inverse iteration requires solving systems of the form $(A-\mu I)x^{(k)} = x^{(k-1)}$ with μ near an eigenvalue of A, see [14, §27]. Such a problem clearly becomes indefinite once the eigenvalue selected is sufficiently deep into the spectrum. As some important problems involve solving indefinite systems, extending these iterative methods to indefinite problems is a useful endeavor. A study of the dual primal variant of FETI (FETI-DP) by Farhat and Li resulted in the FETI-DPH algorithm which is applicable to some indefinite problems [5]. The work of Li and Tu has extended BDDC to some classes of indefinite problems [8]. Both the FETI-DPH algorithm and the extension of the BDDC algorithm rely on enlarging the coarse problem to overcome the difficulties resulting from an indefinite problem [5][8][13]. Two-level Schwarz methods have previously been extended to some indefinite problems by Cai and Widlund, see [3][4]. Their approach focused on carefully controlling the diameter of the elements in the coarse triangulation, see [3] [4]. The present discussion will focus on attempting to apply two-level Schwarz methods to some indefinite problems by enlarging the coarse problem in the spirit of the aforementioned work of Farhat, Li, and Tu, see [5] [8].

1.1 Overview

The present discussion will proceed in four stages. First, some properties of two-level Schwarz methods are reviewed for the case of positive definite problems. This is accompanied by numerical experiments that serve to establish familiarity with the behavior of Schwarz preconditioners and vet the implementations of the Schwarz preconditioners. Second, some Schwarz preconditioners with expanded coarse spaces are proposed to handle indefinite problems. The proposed preconditioners closely parallel the work by Farhat and Li in [5] on FETI-DPH. FETI is replaced by two-level Schwarz, and the analogous steps are taken to construct additional coarse basis functions based on plane waves. The effectiveness of the resulting preconditioners is explored through the use of some numerical experiments. Third, spectral techniques are used to compare the proposed preconditioners. The various preconditioned systems are examined using spectral techniques in an attempt to verify that preconditioners with faster convergence correspond to preconditioned systems that are expected to be conducive to fast GMRES convergence. Lastly, inverse iteration is presented as an application of the preconditioning techniques developed.

In the numerical experiments throughout, Ω is a square domain. The coarse and fine triangulations are based on square grids. The stiffness matrices of the problems are obtained by finite elements. In particular, each square is divided into two triangles, and piecewise linear, nodal, conforming elements are used. These basis functions are linear on each triangle. Note that this gives a stiffness matrix identical to the classical five point finite difference scheme, see [1, §4]. The number of coarse mesh points along one edge of the domain is denoted by N. Similarly, H is the coarse mesh size and h is the fine mesh size. The

amount of overlap is denoted by δ which refers to the amount of overlap added to each edge of a typical, non-overlapping subdomain. Thus, the overlapping subdomains have side length $H+2\delta$. Condition number estimates are denoted by κ. Residuals are measured in the ℓ_2 norm. The iteration count is the number of iterations of conjugate gradients that are necessary to reduce the ℓ_2 norm of the initial residual by at least 10^8. Lastly, "relative residual" is the magnitude of the residual once the iteration halts, relative to the initial residual. Halting of the iteration occurs either by reaching the convergence conditions, or the maximal number of iterations allowed. The magnitude of the residual is measured relative to the initial residual.

2 Two-Level Overlapping Schwarz Methods

The domain decomposition used for two-level overlapping Schwarz methods differs from that used for the FETI-DP and BDDC algorithms in that the subdomains used for Schwarz methods overlap. That is, for adjacent, open subdomains Ω_i, Ω_j,

$$\Omega_i \cap \Omega_j \neq \emptyset$$

for the case of a Schwarz method, while for FETI-DP and BDDC,

$$\Omega_i \cap \Omega_j = \emptyset.$$

Such an overlapping partition for a two-level Schwarz method can be formed by first taking a partition of Ω into non-overlapping, open subdomains Ω_i. These are then enlarged to give a set of overlapping subdomains Ω_i'. The construction of the decomposition of Ω begins with two triangulations of Ω. One triangulation is coarse, while the other is a refinement of the first. Associating a finite element space with the coarse triangulation gives the coarse space V_0. The coarse triangulation also gives a natural partition into non-overlapping Ω_i. The Ω_i' are constructed by adding layers of the fine triangulation adjacent to $\partial\Omega_i$. For $i = 1, 2, \ldots, N$, a local finite element space V_i is associated to each Ω_i' using the fine triangulation. Lastly, let V be a finite element space on Ω utilizing the fine triangulation on the full domain. Proceeding with the domain decomposition, restriction maps

$$R_i : V \to V_i$$

are defined as well as interpolation maps

$$R_i^T : V_i \to V$$

which are extensions by zero. These lead to local problems $\widetilde{A_i}$ for $i = 1, 2, \ldots, n$, given by (1)

$$\widetilde{A_i} = R_i A R_i^T \tag{1}$$

along with a coarse (global) problem, $\widetilde{A_0}$ given by (2)

$$\widetilde{A_0} = R_0 A R_0^T. \tag{2}$$

From these problems, Schwarz projectors P_i are constructed, as in (3)

$$P_i = R_i^T \widetilde{A_i}^{-1} R_i A. \tag{3}$$

It can be shown that the P_i are projections, see [13, Lemma 2.1]. A preconditioned problem is obtained by solving the coarse problem (2) and local problems (1), then assembling the solutions in some manner, see [13, §3]. The specific procedure for solving and combining these problems can be formulated by combining the projectors in different manners, see [13, §1.4, §2.2, §2.5.2].

Recall that, for the numerical experiments, Ω is a square. The Ω_i are taken to be a set of squares with side length H and vertices coinciding with coarse mesh points. Furthermore, the Ω_i' are obtained by adding a number of additional fine mesh layers. By maintaining a uniform size for the local problems, the matrix corresponding to the local problems can be factored once and used to solve all of the local problems. To maintain this uniformity, subdomains Ω_i adjacent to $\partial\Omega$ are enlarged to overlapping subdomains Ω_i' by adding a width of 2δ along the edge opposite the edge contained in $\partial\Omega$. Interior subdomains Ω_i are enlarged to overlapping subdomains Ω_i by adding a fixed width of δ to each of the four sides. This gives square overlapping subdomains Ω_i' of side length $H + 2\delta$.

2.1 Additive Schwarz Preconditioner

The additive variant of the Schwarz preconditioner is constructed by simply summing all of the Schwarz projectors [13, §2.2]. That is,

$$P_{ad} = \sum_{i=0}^{N} P_i. \tag{4}$$

This variant has the advantage of all local problems as well as the coarse problem being independent of one another, allowing for parallel execution. However, convergence rates tend to be slower than for multiplicative variants.

2.2 Hybrid Schwarz Preconditioner

A hybrid Schwarz Preconditioner can be constructed by

$$P_{hy} = P_0 + (I - P_0) \left(\sum_{i=1}^{N} P_i \right) (I - P_0) \tag{5}$$

see [13, §2.2]. This variant treats the coarse problem multiplicatively, and all of the local problems additively.

The discussion for this paper will be limited to the additive and hybrid variants.

2.3 Positive Definite Problems

The theoretical results give the following bound on the condition number for the additive Schwarz preconditioner, which is quoted without proof.

Theorem 1. *If exact solvers are utilized on all subspaces, then*

$$\kappa \leq C_1 + C_2 \left(\frac{H}{\delta} \right)$$

where κ is the condition number of the additive Schwarz operator and the constants C_1, C_2 depend on N^c but are independent of h, H, and δ [13, Theorem 3.13].

N^c is the number of colors necessary to color the undirected graph formed by assigning a node to each subdomain, and placing an edge between nodes corresponding to overlapping subdomains [13, §2.5.1].

Thus, as $H\delta^{-1}$ decreases, κ is expected to steadily decrease. However, when $H\delta^{-1}$ approaches 2, a sudden change in κ can be expected as this corresponds to three-fold overlap, a setup in which each overlapping subdomain overlaps two additional subdomains in each direction, increasing N^c, resulting in different values of C_1, C_2.

Experiments for the case of positive definite problems begins by considering the Poisson equation (6) defined on a square domain Ω with the Dirichlet condition $u = 0$ on $\partial\Omega$ as a model problem,

$$-\Delta u = f. \tag{6}$$

3 Numerical Results

Tables 2 through 4 focus on verifying that the implementation generates numerical results consistent with this theoretical bound on κ. This is achieved in three steps. First, $H\delta^{-1}$ is fixed while the fine and coarse mesh sizes, h and H respectively, are allowed to vary. As C_1 and C_2 are independent of h, H, and δ, the condition number κ is expected to remain fixed as $H\delta^{-1}$ is fixed. Then this is carried out with a few different right hand sides, as indicated in Table 1, to verify that this bound is also independent of the choice of right hand side. In all four cases, zero Dirichlet conditions are used. The following domains are used. For (7), $\Omega = [-3\pi/2, 3\pi/2] \times [-3\pi/2, 3\pi/2]$. In the case of (8) and (9), $\Omega = [-1, 1] \times [-1, 1]$. Lastly, for (10), $\Omega = [-\pi, \pi] \times [-\pi, \pi]$. Finally, H and h are fixed while increasing δ. The behavior of κ is expected to change near $H\delta^{-1} = 2$ as three-fold overlap occurs everywhere when $H\delta^{-1} = 2$.

The estimate of the condition number is obtained by estimating the maximal eigenvalue $\lambda_{\max}$ and the minimal eigenvalue $\lambda_{\min}$ and taking

$$\kappa = \frac{\lambda_{\max}}{\lambda_{\min}}$$

Table 1: Functions used for right hand sides

$$f(x,y) = 8\cos(2x)\cos(2y) \tag{7}$$
$$f(x,y) = (4x^2-2)e^{1-x^2} + (4y^2-2)e^{1-y^2} + (4-4y^2-4x^2)e^{2-x^2-y^2} \tag{8}$$
$$f(x,y) = 4 - 2y^2 - 2x^2 \tag{9}$$
$$f(x,y) = 32\sin(4x)\sin(4y) \tag{10}$$

Table 2: Condition number and convergence data for (9) $f = 4 - 2y^2 - 2x^2$

Unknowns	h	H	δ	κ	Iteration Count	Relative Residual
7921	2/91	2/5	2/15	5.7462813	19	8.73835e-09
22201	2/151	2/5	2/15	5.6421541	19	5.0427e-09
108241	2/331	2/5	2/15	5.4990406	18	5.396e-09
395641	2/631	2/5	2/15	5.2893106	17	9.63763e-09

as the condition number estimate. The relevant eigenvalues, $\lambda_{\max}, \lambda_{\min}$ are estimated by calculating the eigenvalues of the tridiagonal matrix generated in the course of the conjugate gradients iteration, see [10, §4.4], [12, §6.7.3].

Numerical experiments are carried out using MATLAB. The conjugate gradient algorithm is implemented based on the description in [12, Algorithm 9.1]. The Schwarz preconditioners are implemented following the relevant descriptions in [13]. The associated coarse and local problems are solved using the Cholesky factorization implementation provided in MATLAB. Residuals are measured in the ℓ_2 norm.

3.1 Condition Number Results

The data contained in Table 2 is generated using (9) for the right hand side. The condition number, κ, decreases a small amount as h is decreased. However, this change is not significant relative to the size of κ, which is consistent with the theoretical results.

To examine the effect of changing the right hand side, all four proposed right hand sides are considered and the condition number (κ) is plotted against the fine mesh size (h) in Figure 1. For each different right hand side, a fixed value for H and δ was used. Again, κ remains stable in the presence of different right hand sides, in line with the results expected by the theory.

Returning to the case of (9) as the right hand side, H and δ are allowed to vary to arrive at the values in Table 3. Performing a least squares fit of the condition number to a quadratic polynomial in $H\delta^{-1}$, the following quadratic

Figure 1: Condition number for fixed $H/\delta = 3$ varying mesh size

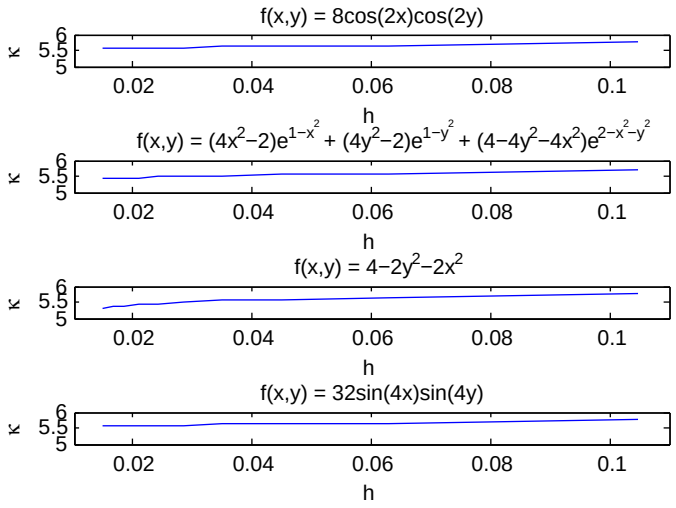

Table 3: Condition number with varying H/δ

H	δ	H/δ	κ
1/4	5/64	16/5	5.7926800
1/4	1/16	4	6.0885131
1/4	3/64	16/3	6.4235317
1/4	1/32	8	7.7075691
1/4	1/64	16	12.970830

model is obtained.

$$\kappa = 0.021232431 \left(\frac{H}{\delta}\right)^2 + 0.15331958 \left(\frac{H}{\delta}\right) + 5.0852774$$

Additionally, the following linear model of the data is computed, again by a least squares fit.

$$\kappa = 0.57083990 \left(\frac{H}{\delta}\right) + 3.6256879$$

In Figure 2 both of these models are plotted along with the five calculated data points indicated in Table 3. Note that the coefficient for the quadratic term is roughly one tenth of the coefficient for the linear term. This indicates that the implementation of the additive Schwarz method is behaving in nearly the linear fashion with respect to $H\delta^{-1}$, in agreement with the theoretical result.

Lastly, the effect of varying the amount of overlap (δ) is examined. For these experiments, (9) is once again used to generate the right hand side. A fixed fine

Figure 2: Linear and quadratic models of κ

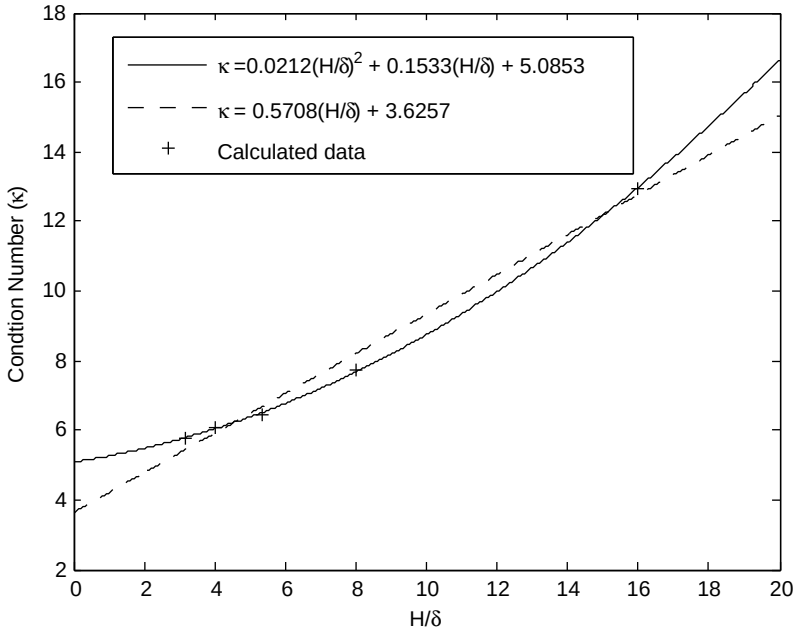

Table 4: Condition number with respect to varying δ

Layers	H/δ	κ	$\lambda_{\min}$	$\lambda_{\max}$
28	8/7	1.6093010e+01	1.0000401e+00	1.6093656e+01
26	16/13	1.6042061e+01	1.0000054e+00	1.6042147e+01
24	4/3	1.6011506e+01	9.9996022e-01	1.6010869e+01
22	16/11	1.6002531e+01	9.9984187e-01	1.6000000e+01
20	8/5	9.2426883e+00	9.9939168e-01	9.2370658e+00
18	16/9	9.1689795e+00	9.9445863e-01	9.1181708e+00
16	2	9.7939680e+00	9.2491105e-01	9.0585493e+00
14	16/7	1.0299989e+01	8.7574232e-01	9.0201362e+00
12	8/3	1.1141623e+01	8.0786779e-01	9.0009580e+00
10	16/5	5.7926800e+00	7.6431925e-01	4.4274568e+00
8	4	6.0885131e+00	7.0375532e-01	4.2848235e+00
6	16/3	6.4235317e+00	6.4787062e-01	4.1616174e+00
4	8	7.7075691e+00	5.2807739e-01	4.0701930e+00
2	16	1.2970830e+01	3.0988126e-01	4.0194171e+00

mesh size of $h^{-1} = 128$ is used. The coarse mesh size is fixed at $H^{-1} = 4$, and δ is allowed to vary from $\delta = 1/64$ to $\delta = 7/32$. Selected values are shown in Table 4.

3.2 General comments

Numerical results for the positive definite problem are consistent with the established theoretical results. The condition number κ demonstrates independence from the selection of right hand side and h. Additionally, nearly linear behavior of κ with respect to $H\delta^{-1}$ is observed. As expected, a steady decrease of κ is observed as $H\delta^{-1}$ is decreased. Additionally, an increase in κ near $H\delta^{-1} = 2$ is observed. This is expected as the subdomains intersecting the boundary $\partial\Omega$ attain three-fold overlap just prior to $H\delta^{-1} = 2$ and all other subdomains attain three-fold overlap when $H\delta^{-1} = 2$. Three-fold overlap increases N^C which Theorem 1 indicates will change C_1, C_2, and hence the bound on κ. In Figure 3, the lower bound on the spectrum steadily increases as $H\delta^{-1}$ decreases. The upper bound on the spectrum remains stable until $H\delta^{-1}$ approaches 2, at which point it rapidly increases, degrading the conditioning of the problem. Three-fold overlap is achieved by the domains intersecting $\partial\Omega$ due to specific choices made in the implementations of the Schwarz preconditioner related to maintaining a uniform size for the local problems.

Figure 3: Condition number versus varying overlap

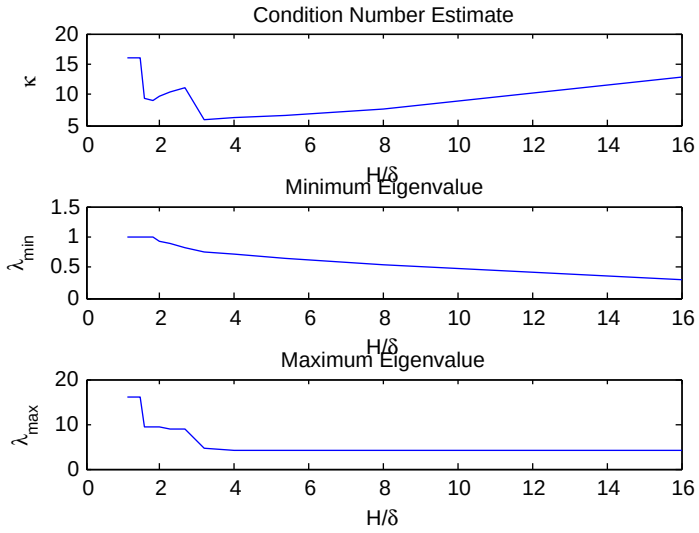

4 Transition to indefinite problems

The model indefinite problem to be investigated is a shifted variant (11) of the Poisson equation (6) with zero Dirichlet conditions,

$$(-\Delta - \gamma^2)u = f. \tag{11}$$

Note that for sufficiently large γ^2, the shifted problem (11) is no longer positive definite. Thus, two changes are made to the preconditioned solver. First, the Cholesky factorization used to solve the coarse problem in the Schwarz preconditioner is replaced by an LU factorization. Additionally, the Cholesky factorization used to solve the local problems should be replaced by an LU factorization if the local problems are allowed to be indefinite. Alternatively, by restricting the size of H, the local problems can be made positive definite, in which case, Cholesky factorizations are preferred for performance reasons. For all of the numerical experiments presented, the latter approach is taken; a sufficiently small H is selected to ensure positive definite local problems. Second, the preconditioned CG algorithm is replaced by the left preconditioned generalized minimal residual method (GMRES).

4.1 GMRES

CG is a Krylov space method and the best known iterative solver for positive definite systems, see [12, §6.7]. GMRES is an iterative, Krylov space method suitable for nonsymmetric and indefinite problems, see [11] [12, §6.5]. Krylov

methods seek an approximate solution of a system $Ax = b$ in the associated Krylov subspace. The Krylov subspace is defined by (12),

$$\mathcal{K}_m(A, v) = \text{span}\{v, Av, \ldots, A^{(m-1)}v\}. \tag{12}$$

A difficulty in using GMRES is the lack of an explicit calculation of the residual at each step, see [12, §6.5.3]. The solution selected for this implementation is to explicitly calculate the residual every 5 iterations. Additionally, if a large number of iterations is necessary, storage and computational requirements can become quite high. To counteract this, the GMRES algorithm can be restarted by running the algorithm for a fixed number of iterations, then taking the initial guess x_0 to be the current approximation and restarting the iteration [12, §6.5.5]. GMRES also possesses several desirable qualities. The algorithm will only break down if it is fully converged, that is, the solution is exact. Otherwise, the iteration will continue until the convergence criteria are met [11, §3.4]. Additionally, the restarted variant of GMRES tends to require less storage space and less computation than other similar methods [11].

GMRES calculates an approximate solution by constructing an element y of the Krylov subspace related to the initial guess x_0. At the m^{th} iteration y is selected from the m^{th} Krylov subspace $\mathcal{K}_m$ to minimize the residual $\|A(x_0+y)-b\|_2$ [12, §6.5]. The mechanics of GMRES are related to the Arnoldi iteration, see [14, §35] [12, §6]. The Arnoldi iteration constructs an orthonormal basis of the Krylov space $\mathcal{K}_m(A, v)$ [12, §6.3]. Additionally, it forms an upper Hessenberg matrix. The eigenvalues of this upper Hessenberg matrix converge to the eigenvalues of A. In particular, the eigenvalues of the Hessenberg matrix will tend to converge to the extremal eigenvalues of A quickly, see [14, §34]. This suggests a rough means for evaluating the efficacy of various preconditioning techniques. The eigenvalues of the Hessenberg matrix formed in the GMRES iteration are calculated. This gives an approximation of the spectrum of A, denoted by $\sigma(A)$. Comparing different techniques, a preconditioner can be expected to be more effective if the eigenvalues of the Hessenberg matrix are conducive to rapid convergence of GMRES.

The n^{th} GMRES approximation can also be thought of in terms of a polynomial approximation problem. The quantity

$$\|p_n(A)b\|$$

is minimized by finding an optimal p_n in P_n, the space of monic polynomials of degree at most n with $p_n(0) = 1$ [14, §35].

The convergence rate of the GMRES iteration can be estimated by the following result, quoted without proof.

Theorem 2. *Let A be diagonalizable and V be a matrix of eigenvectors. Let $\sigma(A)$ be the spectrum of A. Then after n iterations, the residual of the GMRES iteration satisfies*

$$\frac{\|r_n\|}{\|b\|} \leq \kappa(V) \inf_{p_n \in P_n} \sup_{z \in \sigma(A)} |p_n(z)|.$$

[14, Theorem 35.2].

If V is assumed to be well conditioned, then this reduces to bounding $p_n(z)$ on the spectrum of A. This leads to an analysis of the convergence of GMRES based on the eigenvalues of A. The worst case is when the eigenvalues of A surround the origin [14, §35]. As spectral information about A is typically not known, one approach is to bound the spectrum by an elliptical region disjoint from the origin, and minimize the polynomial over this region [9]. This approach is developed by Manteuffel for analyzing the Tchebychev iteration by confining the eigenvalues to a closed elliptical region contained in the right half plane, see [9]. Specifically, the following result regarding the polynomial minimization problem is quoted without proof.

Theorem 3. *Assume A is $M \times M$ and diagonalizable with eigenvalues denoted $\lambda_1, \ldots, \lambda_M$ and let*

$$\varepsilon^{(m)} = \min_{p_n \in P_n} \max_{\lambda \in \sigma(A)} |p_n(\lambda)|$$

where $\sigma(A)$ denotes the spectrum of A. Also, assume the eigenvalues $\lambda_1, \ldots, \lambda_\nu$ are contained in the left half plane for some $0 < \nu < M$. Lastly, assume the eigenvalues $\lambda_{\nu+1}, \ldots, \lambda_M$ are confined to a closed disk with center $C > 0$ and radius $R < C$. Then

$$\varepsilon^{(m)} \leq \left(\frac{D}{d}\right)^{\mu} \left(\frac{R}{C}\right)^{m-\nu}$$

where

$$D = \max_{\substack{i=1,\ldots\nu \\ j=\nu+1,\ldots M}} |\lambda_i - \lambda_j| \text{ and } d = \min_{i=1,\ldots,\nu} |\lambda_i|$$

[11, Theorem 5].

This result outlines several conditions that will aid in rapid convergence of GMRES. First, keeping all of the eigenvalues within a bounded region of small, finite diameter. This will control D. Bounding the eigenvalues in the left half plane away from the origin prevents d from becoming very small. Together, these two conditions control the term Dd^{-1}. If the eigenvalues in the right half plane are not near the imaginary axis, have small imaginary part, and real parts all of comparable size, then it will typically be possible to construct a circle such that the term RC^{-1} is small. Controlling Dd^{-1} and RC^{-1} in this manner promotes rapid convergence of GMRES.

4.2 Comparison of CG and GMRES

The aforementioned modifications allow for the naïve application of the additive and hybrid Schwarz preconditioners to indefinite problems. Before proceeding with such experiments, some previous experiments on the Poisson equation (6) are rerun to verify that these modifications have not significantly impacted performance. Returning to positive definite problems, Table 5 compares the number of iteration counts for identical problems, when CG is replaced by GMRES. The numerical experiments are conducted restarting the GMRES algorithm every 70 iterations. Additionally, this implementation is set to halt after a maximum of

Table 5: Comparison of CG and GMRES performance on positive definite problems

h	H/δ	CG Iterations	GMRES Iterations	CG Residual	GMRES Residual
$f(x,y) = 8\cos(2x)\cos(2y)$					
$\pi/30$	3	20	20	3.59874e-09	2.18897e-09
$\pi/70$	3	19	20	6.63909e-09	2.06143e-09
$f(x,y) = (4x^2-2)e^{1-x^2} + (4y^2-2)e^{1-y^2} + (4-4y^2-4x^2)e^{2-x^2-y^2}$					
$1/45$	3	20	20	4.19206e-09	9.35552e-09
$1/105$	3	19	20	6.41544e-09	1.59943e-09
$f(x,y) = 4 - 2y^2 - 2x^2$					
$1/45$	3	19	20	8.73835e-09	2.4265e-09
$1/105$	3	18	20	4.7383e-09	1.50827e-09
$f(x,y) = 32\sin(4x)\sin(4y)$					
$\pi/45$	3	20	20	3.07881e-09	2.35214e-09
$\pi/105$	3	19	20	3.64368e-09	1.84288e-09

490 iterations if the convergence criterion, a reduction of the ℓ_2 residual by a factor of 10^8, is not met.

The results in Table 5 indicate that the convergence rate of the two algorithms remain comparable. In some cases, the iteration counts are identical, in others the GMRES iteration count is slightly higher, but this can be expected as the GMRES iteration counts only have a resolution down to 5 iterations. That is, if GMRES converges in any of the sixteenth through twentieth iterations, the calculated iteration count will be 20.

4.3 Indefinite Problems

Solving (11) for sufficiently large values of γ^2 gives rise to an indefinite linear system. As shown in Tables 7 and 8 such problems significantly degrade the performance of the additive and hybrid Schwarz preconditioners. Previous work on applying Schwarz preconditioners to indefinite problems has focused on using the coarse problem to approximate the eigenvalues in the left half plane and confine the spectrum of the preconditioned system to a bounded set in the right half plane see [3, §1] [4]. Theorem 3 indicates that such goals are in line with achieving rapid convergence of GMRES. A major result finds that this can be achieved, for some additive variants, by restricting the size of H, see [3, Theorem 1]. This approach can also be applied to multiplicative and multilevel variants of the Schwarz method, see [4]. Numerical results pertaining to this previous work can be found in [2].

One general framework for the construction of preconditioners for indefinite systems is based on combining a coarse problem, to control the global behavior, with a preconditioner for symmetric positive definite systems, to handle the local problems, which can typically be made positive definite [2]. This approach

is developed by Xu and Cai in [15]. They found that obtaining a good solution to the coarse problem is essential to making this approach work [15]. Such an approach was taken by Farhat and Li in the construction of FETI-DPH. The coarse problem construction consists of constructing additional coarse basis functions based on elastic waves [5].

Presently, this approach is attempted with two-level Schwarz methods. The coarse space of the classical Schwarz preconditioners, used in the previous experiments, consists of conforming, linear, nodal basis functions, see [1, §4]. To this, additional plane wave based basis functions are added. These additional basis functions are constructed by picking a plane wave $\psi(x, y)$ and considering an edge $\mathcal{E}_{ij}$ bordering two non-overlapping subdomains Ω_i, Ω_j. Zero Dirichlet conditions are imposed along $(\partial\Omega_i \cup \partial\Omega_j) \setminus \mathcal{E}_{ij}$. Along $\mathcal{E}_{ij}$, Dirichlet conditions are imposed corresponding to evaluating ψ at the nodes that lie on $\mathcal{E}_{ij}$. At the nodes at the corners of Ω_i, Ω_j, zero Dirichlet conditions are used. By solving the shifted problem (13) on Ω_i and Ω_j using the above boundary conditions, a basis function supported on $\Omega_i \cup \Omega_j$ that is equal to ψ on $\mathcal{E}_{ij}$ is obtained,

$$-(\Delta - \gamma^2)u = 0. \tag{13}$$

The harmonic extension is also considered. The approach is the same, but (13) is replaced by (14),

$$-\Delta u = 0. \tag{14}$$

Lastly, H is always selected to ensure that the local problems are positive definite. This places a restriction on the maximum size of H, but is a less restrictive condition than is necessary to control the behavior of indefinite problems by only restricting H.

5 Numerical Experiments

Numerical calculations are carried out in MATLAB. The preconditioned GMRES algorithm is implemented as described in [12, Algorithm 9.4]. The residual is checked every 5 iterations, giving the iteration counts displayed a resolution of 5 iterations. The GMRES iteration is restarted every 70 iterations and is halted at 490 iterations regardless of convergence conditions. Otherwise, the iteration is halted when the magnitude of the residual, relative to the initial residual, is less than 10^{-8}. The Schwarz preconditioners are implemented following the descriptions in [13]. The coarse and local subproblems are solved directly using MATLAB primitives. When auxiliary basis functions are used, the extension is calculated once on a pair of adjacent subdomains and the same extension is used for all edges in that orientation, horizontal or vertical. This approach makes for significant savings in computational cost. In spite of the fact that the values of the plane wave are expected to be different on different edges, preliminary numerical experiments showed no significant difference in using this approach rather than calculating the extension separately for each edge.

Table 6: Values of γ^2

γ^2	Number of negative eigenvalues	$\min\|\gamma^2 - \lambda\|$
100	26	1.1361676
130.69743296	37	2.4897986e-09
182.49	50	4.0004826e-03
182.494	50	4.8264084e-07
182.49400048	50	2.6408316e-09
200	54	2.1237771
241.65664799	67	5.3987037e-10

The shifts (γ^2) are selected to investigate the effect of having additional negative eigenvalues. Additionally, some shifts are selected to be very close to eigenvalues of (6) with an eye towards applicability to inverse iteration among other applications that require solving nearly singular, indefinite problems. The system resulting from such γ^2 is expected to be nearly singular, and hence it is worthwhile to check how this extension of a Schwarz method reacts to such an ill-conditioned problem.

With the above goals in mind, values of γ^2 are selected and both their index and distance to the closest eigenvalue of the discrete Laplacian is listed in Table 6. Two values of the shift $\gamma^2 = 100, 200$ are chosen as large round numbers. The remaining five are chosen to be near eigenvalues of $-\Delta$. Three, $\gamma^2 = 130.69743296, 182.49400048, 241.65664799$ are very close to eigenvalues of $-\Delta$. As it will be seen that performance significantly suffers at $\gamma^2 = 182.49400048$, two additional shifts $\gamma^2 = 182.49$ and $\gamma^2 = 182.494$ are selected to further investigate the sensitivity to ill conditioned problems.

All of the following experiments are conducted using the right hand side corresponding to (9) and domain $\Omega = [-1, 1] \times [-1, 1]$. First, the performance of the additive and hybrid Schwarz preconditioners, with no additional basis functions, is investigated. For this a fixed fine mesh size of $h^{-1} = 128$ is used. The coarse mesh size is either $H^{-1} = 4$ or $H^{-1} = 8$. The coarse mesh is refined when δ is allowed to increase. This ensures positive definite local problems. For $H^{-1} = 4$, 2 layers of the fine mesh are added to generate an overlapping partition, giving $\delta = 1/64$. For $H^{-1} = 8$, the overlap (δ) is allowed to vary, investigating the impact of various values of $H\delta^{-1}$.

As expected, Tables 7 and 8 demonstrate a significant loss of performance when the Schwarz methods are applied to indefinite problems without modification. Both the additive and hybrid variants show similar performance. Only the shifts of $\gamma^2 = 100$ and $\gamma^2 = 130.69743296$ converge with any consistency. The shift of $\gamma^2 = 200$ converges for $H\delta^{-1} = 8$ and $H\delta^{-1} = 4$, when considering the additive variant, and for $H\delta^{-1} = 4$ when considering the hybrid variants. In both cases, after a large number of iterations. The problems corresponding to $\gamma^2 = 182.49400048$ and $\gamma^2 = 241.65664799$ do not converge in any of the attempted experiments. In the case of $\gamma^2 = 100$, performance, of both the additive and hybrid Schwarz methods, can be improved by decreasing $H\delta^{-1}$.

Table 7: Naïve application of the additive Schwarz preconditioner to indefinite problems

Shift (γ^2)	H/δ	δ	Iteration Count	Relative Residual
100	16	1/64	160	9.2022082e-09
100	8	1/64	175	3.7035524e-09
100	4	1/32	110	3.1862903e-09
100	8/3	3/64	105	5.1990863e-09
100	2	1/16	115	6.9153993e-09
130.69743296	16	1/64	280	9.7193482e-09
130.69743296	8	1/64	90	3.4528576e-09
130.69743296	4	1/32	105	4.2315848e-09
130.69743296	8/3	3/64	85	9.2248977e-09
130.69743296	2	1/16	490	3.2667106e-07
182.49400048	16	1/64	490	7.0102286e-07
182.49400048	8	1/64	490	6.3981694e-06
182.49400048	4	1/32	490	1.1882915e-05
182.49400048	8/3	3/64	490	1.8578846e-05
182.49400048	2	1/16	490	7.1418732e-05
200	16	1/64	490	3.9363058e-04
200	8	1/64	460	9.7757165e-09
200	4	1/32	345	5.6168722e-09
200	8/3	3/64	490	1.0678323e-05
200	2	1/16	490	3.8780113e-05
241.65664799	16	1/64	490	2.4490145e-03
241.65664799	8	1/64	490	3.5289017e-06
241.65664799	4	1/32	490	5.0083013e-06
241.65664799	8/3	3/64	490	3.6150524e-05
241.65664799	2	1/16	490	1.3841250e-04

Table 8: Naïve application of a hybrid Schwarz preconditioner to indefinite problems

Shift (γ^2)	H/δ	δ	Iteration Count	Relative Residual
100	16	1/64	160	8.5188128e-09
100	8	1/64	115	1.6310423e-09
100	4	1/32	110	3.3416916e-09
100	8/3	3/64	110	1.8559358e-09
100	2	1/16	115	4.9628975e-09
130.69743296	16	1/64	280	9.7193482e-09
130.69743296	8	1/64	90	3.4528576e-09
130.69743296	4	1/32	105	4.2315848e-09
130.69743296	8/3	3/64	85	9.2248977e-09
130.69743296	2	1/16	490	3.2667106e-07
182.49400048	16	1/64	490	7.0102286e-07
182.49400048	8	1/64	490	6.3981694e-06
182.49400048	4	1/32	490	1.1882915e-05
182.49400048	8/3	3/64	490	1.8578846e-05
182.49400048	2	1/16	490	7.1418732e-05
200	16	1/64	490	3.8425803e-04
200	8	1/64	490	1.3190229e-08
200	4	1/32	345	7.5851212e-09
200	8/3	3/64	490	1.3166859e-05
200	2	1/16	490	4.9391982e-05
241.65664799	16	1/64	490	2.4490145e-03
241.65664799	8	1/64	490	3.5289017e-06
241.65664799	4	1/32	490	5.0083013e-06
241.65664799	8/3	3/64	490	3.6150524e-05
241.65664799	2	1/16	490	1.3841250e-04

Table 9: Angles used to generate waves

2 Waves	0, $\pi/2$
3 Waves	0, $\pi/2$, $5\pi/3$
4 Waves	0, $\pi/2$, $4\pi/7$, $5\pi/3$
5 Waves	0, $\pi/2$, $4\pi/9$, $4\pi/7$, $5\pi/3$
6 Waves	0, $\pi/2$, $4\pi/9$, $4\pi/7$, $5\pi/3$, $16\pi/9$

The other values of γ^2 exhibit at best erratic behavior. There is a general lack of robustness when the unmodified Schwarz preconditioners are applied to indefinite problems. The solver remains sensitive to the conditioning of the non-transformed system. The smallest shifts ($\gamma^2 = 100, 130.69743296$) can be handled without modification to the method, but all of the other selected values for γ^2 result in problems for which convergence is obtained in, at best, a few isolated cases and with a high iteration count.

5.1 Extended Coarse Spaces

In Tables 10, 11, 12, and 13 the coarse space is extended by adding additional basis functions based on waves to the nodal basis functions. For Tables 10 and 11, the basis functions are constructed by solving the shifted problem (13) on a pair of adjacent subdomains. In Tables 12 and 13 the harmonic extension is taken, in which the basis functions are obtained by solving (14) on a pair of adjacent subdomains. The waves are defined by selecting a collection of angels θ_i and then defining

$$\gamma_i^1 = \gamma \cos(\theta_i)$$
$$\gamma_i^2 = \gamma \sin(\theta_i)$$

and letting

$$\psi_i = \cos(\gamma_i^1)\cos(\gamma_i^2) - \sin(\gamma_i^1)\sin(\gamma_i^2). \tag{15}$$

Note that (15) is equal to the real part of (16),

$$e^{i(\gamma^1 x + \gamma^2 y)}. \tag{16}$$

The collection of angles used for each number of waves is indicated in Table 9.

The results of constructing additional basis functions by using the shifted problem (13) to extend boundary conditions imposed by a plane wave are shown in Tables 10 and 11, for the additive and hybrid Schwarz preconditioners respectively. In the case of both the additive and hybrid preconditioner, significant gains are made for both $\gamma^2 = 100$ and $\gamma^2 = 200$. Results are much more mixed for the values of γ^2 near eigenvalues of the discrete operator. In the case of $\gamma^2 = 130.69743296$, convergence is consistently achieved, and very good performance is observed for both the additive and hybrid Schwarz preconditioners using two and six plane waves. In the case of three, four, and five plane waves,

Table 10: Plane waves extended by shifted problem with additive Schwarz

Shift	H/δ	$\dim(V_0)$	Iteration Count	Relative Residual
2 Waves				
100	16	273	220	9.0020667e-09
130.69743296	16	273	115	5.7344353e-09
182.49	16	273	490	9.5943094e-06
182.494	16	273	490	1.0496348e-05
182.49400048	16	273	490	1.0861869e-05
200	16	273	165	9.0781428e-09
241.65664799	16	273	490	4.4347124e-04
3 Waves				
100	16	385	445	3.2862458e-09
130.69743296	16	385	150	8.6464041e-09
182.49	16	385	190	7.9089172e-09
182.494	16	385	490	9.9117405e-06
182.49400048	16	385	490	1.0269273e-05
200	16	385	140	5.6488709e-09
241.65664799	16	385	490	2.7035221e-04
4 Waves				
100	16	497	100	8.8998653e-09
130.69743296	16	497	170	2.2681643e-09
182.49	16	497	125	8.1372466e-09
182.494	16	497	430	4.4020150e-09
182.49400048	16	497	490	1.2790079e-05
200	16	497	140	3.9716990e-09
241.65664799	16	497	490	6.6143717e-06
5 Waves				
100	16	609	105	7.1042291e-09
130.69743296	16	609	150	3.5704010e-09
182.49	16	609	125	8.7572122e-09
182.494	16	609	490	1.2046810e-06
182.49400048	16	609	490	1.3130895e-05
200	16	609	120	7.9501368e-09
241.65664799	16	609	490	3.7911369e-06
6 Waves				
100	16	721	115	6.6731707e-09
130.69743296	16	721	105	8.8475400e-09
182.49	16	721	130	2.4251270e-09
182.494	16	721	435	8.7781901e-09
182.49400048	16	721	490	1.3644642e-05
200	16	721	115	8.6174278e-09
241.65664799	16	721	490	3.2385933e-06

Table 11: Plane waves extended by shifted problem with hybrid Schwarz

Shift	H/δ	$\dim(V_0)$	Iteration Count	Relative Residual
2 Waves				
100	16	273	195	5.7538532e-09
130.69743296	16	273	115	5.7344353e-09
182.49	16	273	490	9.5943094e-06
182.494	16	273	490	1.0496348e-05
182.49400048	16	273	490	1.0861869e-05
200	16	273	180	6.3094789e-09
241.65664799	16	273	490	4.4347124e-04
3 Waves				
100	16	385	250	7.4647556e-09
130.69743296	16	385	150	8.6464041e-09
182.49	16	385	190	7.9089172e-09
182.494	16	385	490	9.9117405e-06
182.49400048	16	385	490	1.0269273e-05
200	16	385	140	5.2669508e-09
241.65664799	16	385	490	2.7035221e-04
4 Waves				
100	16	497	160	9.2948254e-09
130.69743296	16	497	170	2.2681643e-09
182.49	16	497	125	8.1372466e-09
182.494	16	497	430	4.4020150e-09
182.49400048	16	497	490	1.2790079e-05
200	16	497	130	5.7535688e-09
241.65664799	16	497	490	6.6143717e-06
5 Waves				
100	16	609	105	4.7627806e-09
130.69743296	16	609	150	3.5704010e-09
182.49	16	609	125	8.7572122e-09
182.494	16	609	490	1.2046810e-06
182.49400048	16	609	490	1.3130895e-05
200	16	609	95	9.6563359e-09
241.65664799	16	609	490	3.7911369e-06
6 Waves				
100	16	721	115	4.9701544e-09
130.69743296	16	721	105	8.8475400e-09
182.49	16	721	130	2.4251270e-09
182.494	16	721	435	8.7781901e-09
182.49400048	16	721	490	1.3644642e-05
200	16	721	100	7.8886135e-09
241.65664799	16	721	490	3.2385933e-06

performance remains acceptable, but is worse than the cases of two and six plane waves. The problems resulting from $\gamma^2 = 182.49400048$ and $\gamma^2 = 241.65663799$ do not converge, within the allotted iterations, for any of the proposed preconditioners. Selecting $\gamma^2 = 182.49$, the distance to the nearest eigenvalue is increased to roughly 10^{-3}. In this case, convergence is attained when using three or more waves with either the additive or hybrid Schwarz preconditioner. Furthermore, as the number of waves used is increased, performance is further improved. Using $\gamma^2 = 182.494$ reduces the distance to the nearest eigenvalue to roughly 10^{-7}. In this case, convergence is only attained when using six waves, and at a high iteration count.

For choices of γ^2 not near eigenvalues, the addition of plane wave based basis functions reduces the iteration count. Additionally, when using six plane waves, the iteration counts for the three smallest shifts as well as for $\gamma^2 = 200$ are comparable to one another, ranging from 100 to 130 iterations. This indicates that the sensitivity of the preconditioners to nearly singular problems increases as the number of negative eigenvalues, for the non-transformed system, increases. Furthermore, no extended coarse space tested is able to bring about convergence of the preconditioned GMRES algorithm for $\gamma^2 = 182.49400048$ or $\gamma^2 = 241.65664799$, in the allotted number of iterations. In the case of $\gamma^2 = 182.49400048$, the relative residual remains on the order of 10^{-5} for the various number of plane waves used. The proposed preconditioners are unable to improve the performance for this nearly singular problem, relative to the case of no additional basis functions. For $\gamma^2 = 241.65664799$, the relative residual at the time of termination is reduced by a factor of 10^2 in the case of both the additive and hybrid Schwarz preconditioner, as the number of plane waves used is increased. In spite of the fact that convergence is not achieved, performance is improved in the case of $\gamma^2 = 241.65664799$. In cases where the problem both possess many negative eigenvalues, and is nearly singular, success proves limited.

In addition to the overall positive results, there are also some instances of sudden and significant jumps in the iteration counts. In particular, when moving from 2 to 3 plane waves to generate auxiliary basis functions, there is a significant increase in iteration count for $\gamma^2 = 100$, for both the additive and hybrid variant. For the additive variant, the increase is from 220 to 445 iterations; the hybrid case increases from 195 to 250 iterations. This suggests that the problem is not sufficiently controlled by the preconditioner using 2 plane waves. Furthermore, this suggests that some instances of low iteration counts may be the result of serendipitous circumstances of the problem rather than the preconditioner. An attempt will be made to investigate this behavior further, using spectral techniques, in Section 6.

Using the shifted problem (13) to construct auxiliary basis functions from plane waves increases the set of tractable, indefinite problems. Comparing the results in Tables 7 and 8 to those in Tables 10 and 11, more problems converge when using an extended coarse space, than without. Problems that are both indefinite and nearly singular present the most difficulty. Considering $\gamma^2 = 182.49, 182.494, 182, 49400048$, the distance to the nearest eigenvalue is roughly

10^{-3}, 10^{-7} and 10^{-9} respectively. The first is tractable for some of the proposed preconditioners. The second is only tractable when using six plane waves, and converges at a high iteration count. The third, which is closest to an eigenvalue, is not tractable for any of the extended coarse spaces considered. When γ^2 is not near an eigenvalue, a larger value of $\gamma^2 = 200$ results in a tractable problem, in spite of the fact that it results in more negative eigenvalues.

In contrast with the results in Tables 10 and 11, Tables 12 and 13 show that constructing additional basis functions by taking the harmonic extension of plane waves gives little to no enhancement of the performance when $H^{-1} = 4$. Only the problems corresponding to $\gamma^2 = 100$ and $\gamma^2 = 130.69743296$ consistently converge within the allotted number of iterations. However, these cases are also tractable without using an extended coarse space. Additionally, no significant gains in performance are seen in moving from two plane waves to six plane waves. With few exceptions, those cases that do not converge in the allotted iterations also show little change in the relative residual when the maximum iteration count is reached. This changes somewhat in Section 5.2, when H is reduced from $H^{-1} = 4$ to $H^{-1} = 8$. This is consistent with existing theory as controls on the size of H are often necessary to control the behavior of indefinite problems. For $H^{-1} = 8$, the problem corresponding to $\gamma^2 = 200$ becomes tractable. However, success using the extended coarse space constructed by harmonic extension remains limited compared to constructing the auxiliary basis function by the shifted problem (13). Additionally, a smaller value of H results in a larger coarse space. Constructing the auxiliary basis functions by harmonic extension is less effective than using the shifted problem (13).

5.2 Variable Overlap

The results in Tables 10 through 13 use a domain decomposition with fixed $H^{-1} = 4$ and $\delta = 1/64$, while the number of plane wave based auxiliary basis functions is allowed to vary. In Tables 14 through 17 the focus turns to varying the overlap parameter δ. To this end, the number of plane waves used to generate additional basis functions is fixed at four, with the relevant angles indicated in Table 9. The fine mesh remains fixed at $h^{-1} = 128$. The coarse mesh H is reduced to a fixed value of $H^{-1} = 8$, while δ is allowed to vary. A smaller choice of H is used to ensure that the local problems remain positive definite as δ is increased. For the construction of auxiliary basis functions, extension by the shifted problem (13) is used for Tables 14 and 15; harmonic extension is used in Tables 16 and 17 for the additive and hybrid variants respectively. The right hand side is generated by (9), and $\Omega = [-1, 1] \times [-1, 1]$.

Table 14 and Table 15 show similar results for the problems based on values of γ^2 far from eigenvalues. The cases of $\gamma^2 = 182.49400048$ and $\gamma^2 = 241.65664799$ remain intractable. With regard to $\gamma^2 = 100, 200$, both the additive and hybrid Schwarz methods are tolerant of an increase of overlap up to $H\delta^{-1} = 4$, but for more generous overlap (smaller $H\delta^{-1}$) performance is significantly degraded. Additionally, for the smallest shift of $\gamma^2 = 100$, overlap has no significant impact on performance for either the additive or hybrid variants of the Schwarz method.

Table 12: Plane waves using harmonic extension with additive Schwarz

Shift	H/δ	$\dim(V_0)$	Iteration Count	Relative Residual
2 Waves				
100	16	273	155	9.5599874e-09
130.69743296	16	273	210	5.5983885e-09
182.49	16	273	490	1.3694516e-05
182.494	16	273	490	1.3425646e-05
182.49400048	16	273	490	1.3362950e-05
200	16	273	490	7.5903591e-04
241.65664799	16	273	490	3.2330384e-03
3 Waves				
100	16	385	190	3.1540690e-09
130.69743296	16	385	195	8.3096089e-09
182.49	16	385	490	2.0739665e-05
182.494	16	385	490	2.1880907e-05
182.49400048	16	385	490	2.1881024e-05
200	16	385	490	1.5680353e-03
241.65664799	16	385	490	6.4959085e-03
4 Waves				
100	16	497	185	5.3416358e-09
130.69743296	16	497	160	9.9489685e-09
182.49	16	497	490	1.9858520e-05
182.494	16	497	490	2.0834160e-05
182.49400048	16	497	490	2.0834273e-05
200	16	497	490	1.5134567e-03
241.65664799	16	497	490	5.0067063e-03
5 Waves				
100	16	609	135	9.6812984e-09
130.69743296	16	609	180	4.8219924e-09
182.49	16	609	490	2.3016904e-05
182.494	16	609	490	2.4123890e-05
182.49400048	16	609	490	2.4124024e-05
200	16	609	490	1.1325171e-03
241.65664799	16	609	490	4.5240462e-03
6 Waves				
100	16	721	140	7.5493674e-09
130.69743296	16	721	170	9.8077059e-09
182.49	16	721	490	2.3134788e-05
182.494	16	721	490	2.4143136e-05
182.49400048	16	721	490	2.4143252e-05
200	16	721	490	1.5565399e-03
241.65664799	16	721	490	4.3871924e-03

Table 13: Plane waves using harmonic extension with hybrid Schwarz

Shift	H/δ	$\dim(V_0)$	Iteration Count	Relative Residual
2 Waves				
100	16	273	155	5.4350322e-09
130.69743296	16	273	210	5.5983885e-09
182.49	16	273	490	1.3694516e-05
182.494	16	273	490	1.3425646e-05
182.49400048	16	273	490	1.3362950e-05
200	16	273	490	1.1798718e-03
241.65664799	16	273	490	3.2330384e-03
3 Waves				
100	16	385	190	3.3046439e-09
130.69743296	16	385	195	8.3096089e-09
182.49	16	385	490	2.0739665e-05
182.494	16	385	490	2.1880907e-05
182.49400048	16	385	490	2.1881024e-05
200	16	385	490	1.6457573e-03
241.65664799	16	385	490	6.4959085e-03
4 Waves				
100	16	497	185	5.7190104e-09
130.69743296	16	497	160	9.9489685e-09
182.49	16	497	490	1.9858520e-05
182.494	16	497	490	2.0834160e-05
182.49400048	16	497	490	2.0834273e-05
200	16	497	490	1.5639719e-03
241.65664799	16	497	490	5.0067063e-03
5 Waves				
100	16	609	140	8.2622400e-09
130.69743296	16	609	180	4.8219924e-09
182.49	16	609	490	2.3016904e-05
182.494	16	609	490	2.4123890e-05
182.49400048	16	609	490	2.4124024e-05
200	16	609	490	1.4449118e-03
241.65664799	16	609	490	4.5240462e-03
6 Waves				
100	16	721	140	8.9535096e-09
130.69743296	16	721	170	9.8077059e-09
182.49	16	721	490	2.3134788e-05
182.494	16	721	490	2.4143136e-05
182.49400048	16	721	490	2.4143252e-05
200	16	721	490	1.4677963e-03
241.65664799	16	721	490	4.3871924e-03

Table 14: Varying Overlap Additive Schwarz

Shift	H/δ	δ	$\dim(V_0)$	Iteration Count	Relative Residual
100	8	1/64	2145	100	3.0436788e-09
100	4	1/32	2145	105	1.6813723e-09
100	8/3	3/64	2145	100	2.3063731e-09
100	2	1/16	2145	105	6.3473931e-09
130.69743296	8	1/64	2145	100	2.0669144e-09
130.69743296	4	1/32	2145	100	2.4692700e-09
130.69743296	8/3	3/64	2145	70	6.1354683e-09
130.69743296	2	1/16	2145	490	2.1269374e-06
182.49	8	1/64	2145	450	7.0643878e-09
182.49	4	1/32	2145	490	8.2930068e-08
182.49	8/3	3/64	2145	490	3.4256439e-05
182.49	2	1/16	2145	490	2.7476827e-05
182.494	8	1/64	2145	490	1.0614526e-05
182.494	4	1/32	2145	490	3.4189226e-05
182.494	8/3	3/64	2145	490	3.5193717e-05
182.494	2	1/16	2145	490	2.7412401e-05
182.49400048	8	1/64	2145	490	1.4089799e-05
182.49400048	4	1/32	2145	490	3.7245845e-05
182.49400048	8/3	3/64	2145	490	3.5014129e-05
182.49400048	2	1/16	2145	490	2.7416902e-05
200	8	1/64	2145	200	4.9055077e-09
200	4	1/32	2145	185	4.0763359e-09
200	8/3	3/64	2145	490	9.4330359e-07
200	2	1/16	2145	490	8.6148982e-06
241.65664799	8	1/64	2145	490	5.6337035e-06
241.65664799	4	1/32	2145	490	1.9988713e-05
241.65664799	8/3	3/64	2145	490	1.2148124e-05
241.65664799	2	1/16	2145	490	9.4844144e-05

Table 15: Varying Overlap Hybrid Schwarz

Shift	H/δ	δ	$\dim(V_0)$	Iteration Count	Relative Residual
100	8	1/64	2145	100	9.9830696e-09
100	4	1/32	2145	80	8.2644625e-09
100	8/3	3/64	2145	100	7.9927872e-09
100	2	1/16	2145	110	9.4885554e-09
130.69743296	8	1/64	2145	95	8.3296453e-09
130.69743296	4	1/32	2145	95	4.1112115e-09
130.69743296	8/3	3/64	2145	70	4.8243226e-09
130.69743296	2	1/16	2145	345	8.1940962e-09
182.49	8	1/64	2145	450	7.0643878e-09
182.49	4	1/32	2145	490	8.2930068e-08
182.49	8/3	3/64	2145	490	3.4256439e-05
182.49	2	1/16	2145	490	2.7476827e-05
182.494	8	1/64	2145	490	1.0614526e-05
182.494	4	1/32	2145	490	3.4189226e-05
182.494	8/3	3/64	2145	490	3.5193717e-05
182.494	2	1/16	2145	490	2.7412401e-05
182.49400048	8	1/64	2145	490	2.5168852e-05
182.49400048	4	1/32	2145	490	4.9951271e-05
182.49400048	8/3	3/64	2145	490	3.5081303e-05
182.49400048	2	1/16	2145	490	2.7871986e-05
200	8	1/64	2145	240	6.8829858e-09
200	4	1/32	2145	255	7.8889987e-09
200	8/3	3/64	2145	490	3.0408008e-06
200	2	1/16	2145	490	9.8470379e-06
241.65664799	8	1/64	2145	490	5.7060866e-06
241.65664799	4	1/32	2145	490	1.0440383e-05
241.65664799	8/3	3/64	2145	490	2.5973163e-05
241.65664799	2	1/16	2145	490	1.2204557e-04

Table 16: Varying Overlap Additive Schwarz with Harmonic Extension

Shift	H/δ	δ	$\dim(V_0)$	Iteration Count	Relative Residual
100	8	1/64	2145	125	7.3654342e-09
100	4	1/32	2145	105	7.8852340e-09
100	8/3	3/64	2145	110	8.3193554e-09
100	2	1/16	2145	115	7.5720233e-09
130.69743296	8	1/64	2145	120	3.5076260e-09
130.69743296	4	1/32	2145	115	5.1788017e-09
130.69743296	8/3	3/64	2145	105	9.6312223e-09
130.69743296	2	1/16	2145	490	6.1314029e-06
182.49	8	1/64	2145	490	1.0038766e-05
182.49	4	1/32	2145	490	1.6420108e-05
182.49	8/3	3/64	2145	490	8.0804478e-06
182.49	2	1/16	2145	490	1.1027626e-05
182.494	8	1/64	2145	490	8.5703894e-06
182.494	4	1/32	2145	490	1.5053449e-05
182.494	8/3	3/64	2145	490	7.9481144e-06
182.494	2	1/16	2145	490	1.0905083e-05
182.49400048	8	1/64	2145	490	1.6047363e-05
182.49400048	4	1/32	2145	490	8.6262014e-06
182.49400048	8/3	3/64	2145	490	8.4163102e-06
182.49400048	2	1/16	2145	490	1.1413697e-05
200	8	1/64	2145	260	5.9892001e-09
200	4	1/32	2145	230	9.7942780e-09
200	8/3	3/64	2145	490	1.1043534e-05
200	2	1/16	2145	490	3.5100252e-05
241.65664799	8	1/64	2145	490	1.2608351e-04
241.65664799	4	1/32	2145	490	5.4577373e-04
241.65664799	8/3	3/64	2145	490	3.5170903e-05
241.65664799	2	1/16	2145	490	4.8865863e-04

For the shift $\gamma^2 = 130.69743296$, overlap can be increased to $H\delta^{-1} = 8/3$. Indeed, a performance improvement is seen in this case for both the additive and hybrid variants. In the cases of $\gamma^2 = 182.49$ and $\gamma^2 = 182.494$, a decrease of $H\delta^{-1}$ from $H\delta^{-1} = 16$ to $H\delta^{-1} = 8$ proves catastrophic to the performance of the preconditioner. The cases in which greater overlap improves performance are very limited. In general, performance is either unchanged or significantly degraded with increased overlap. This suggests that for an extended coarse space, with auxiliary basis functions constructed from plane waves by the shifted problem (13), small overlap is to be preferred.

In the case of harmonic extension, the reduction of H from $H^{-1} = 4$ to $H^{-1} = 8$ makes the case of $\gamma^2 = 200$ tractable. This is likely due more to the reduction of H than the increase in δ as other results, not shown, indicate that

Table 17: Varying Overlap Hybrid Schwarz with Harmonic Extension

Shift	H/δ	δ	$\dim(V_0)$	Iteration Count	Relative Residual
100	8	1/64	2145	125	4.5624092e-09
100	4	1/32	2145	105	4.4728500e-09
100	8/3	3/64	2145	105	8.8771967e-09
100	2	1/16	2145	105	6.2445180e-09
130.69743296	8	1/64	2145	95	8.3296453e-09
130.69743296	4	1/32	2145	95	4.1112115e-09
130.69743296	8/3	3/64	2145	70	4.8243226e-09
130.69743296	2	1/16	2145	345	8.1940962e-09
182.49	8	1/64	2145	490	1.0038766e-05
182.49	4	1/32	2145	490	1.6420108e-05
182.49	8/3	3/64	2145	490	8.0804478e-06
182.49	2	1/16	2145	490	1.1027626e-05
182.494	8	1/64	2145	490	8.5703894e-06
182.494	4	1/32	2145	490	1.5053449e-05
182.494	8/3	3/64	2145	490	7.9481144e-06
182.494	2	1/16	2145	490	1.0905083e-05
182.49400048	8	1/64	2145	490	2.5168852e-05
182.49400048	4	1/32	2145	490	4.9951271e-05
182.49400048	8/3	3/64	2145	490	3.5081303e-05
182.49400048	2	1/16	2145	490	2.7871986e-05
200	8	1/64	2145	260	7.0694182e-09
200	4	1/32	2145	325	6.0351320e-09
200	8/3	3/64	2145	490	1.7271959e-05
200	2	1/16	2145	490	4.3629406e-05
241.65664799	8	1/64	2145	490	5.7060866e-06
241.65664799	4	1/32	2145	490	1.0440383e-05
241.65664799	8/3	3/64	2145	490	2.5973163e-05
241.65664799	2	1/16	2145	490	1.2204557e-04

for $H^{-1} = 4$ and $H\delta^{-1} = 8$, convergence does not occur within the allotted iterations. This is reasonable as the prior work on using Schwarz methods to solve indefinite problems has shown using a sufficiently small value of H is essential [3][4]. Otherwise, convergence behavior is similar to that observed in the previous results displayed in Tables 12 and 13. The problem corresponding to $\gamma^2 = 100$ is insensitive to decreasing $H\delta^{-1}$. Performance gains are achieved in the case of $\gamma^2 = 130.69743296$ as $H\delta^{-1}$ is reduced, but reduction to $H\delta^{-1} = 8/3$ or less results in significant degradation of performance. In the case of $\gamma^2 = 200$, increasing overlap beyond $H\delta^{-1} = 4$ results in degraded performance. Overall, small overlap proves to be preferable to generous overlap.

6 Spectral Analysis

In Section 4.1, it was noted that the efficacy of a preconditioner is related to minimizing polynomials over the spectrum of the transformed system. Additionally, in Section 4.1 some conditions conducive to rapid convergence of GMRES are outlined. Configurations of eigenvalues where the absolute values of all eigenvalues are similar, and the eigenvalues in the right half plane can be confined to a bounded circular, or elliptical, region with small radius are desirable. The worst case is for the eigenvalues to surround the origin. Furthermore, the relevant eigenvalues can be approximated in a rough fashion by the eigenvalues of the upper Hessenberg matrix formed in the GMRES iteration. The proposed preconditioners, with extended coarse spaces, are now analyzed by spectral means in an attempt to verify that those displaying better performance are indicated by spectra conducive to fast GMRES convergence. The relevant eigenvalues are displayed in Figures 4 through 8 for the additive and hybrid preconditioners, with auxiliary basis functions generated by the shifted problem (13). Figures 9 through 12 display the spectral information for some selected cases of the additive and hybrid preconditioners without an extended coarse space. Lastly, Figures 13 through 18 display spectral information for some cases of the additive and hybrid preconditioners with auxiliary basis functions generated by harmonic extension.

The preconditioners being considered in Figures 4 through 8 use a coarse space enlarged by constructing plane wave based basis functions using the shifted problem (13). The cases of $\gamma^2 = 100$ and $\gamma^2 = 130.69743296$ correspond to problems in which good acceleration is observed for some of the proposed preconditioners. The case of $\gamma^2 = 241.65664799$ illustrates the failure of the proposed preconditioners to provide significant acceleration. In all cases, there are real eigenvalues in the left half plane. This is expected as the original systems are indefinite and no specific steps are taken to ensure the transformed system is positive definite. Beyond this, there are significant differences in the location of the complex eigenvalues, as well as the magnitudes of the eigenvalues.

The two successful cases (Figures 4 through 7) show encouraging results for the application of GMRES. In the case of $\gamma^2 = 100$, both the additive and hybrid Schwarz methods, at 4 and 6 plane waves, result in the negative eigenvalues

Figure 4: Additive Schwarz $\gamma^2 = 100$ – waves extended by shifted problem

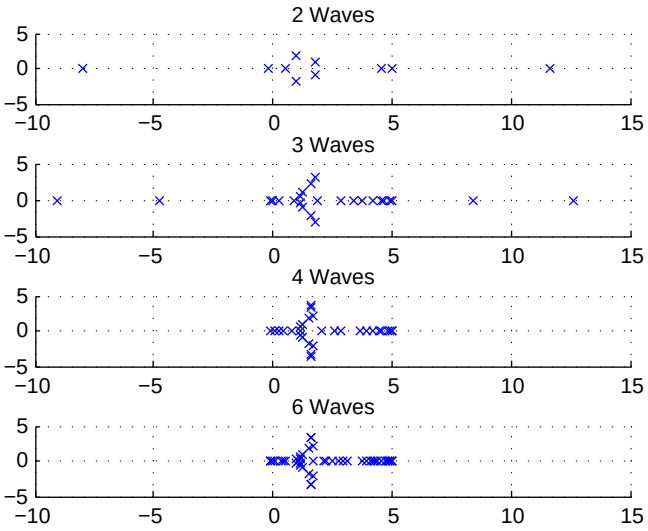

Figure 5: Hybrid Schwarz $\gamma^2 = 100$ – waves extended by shifted problem

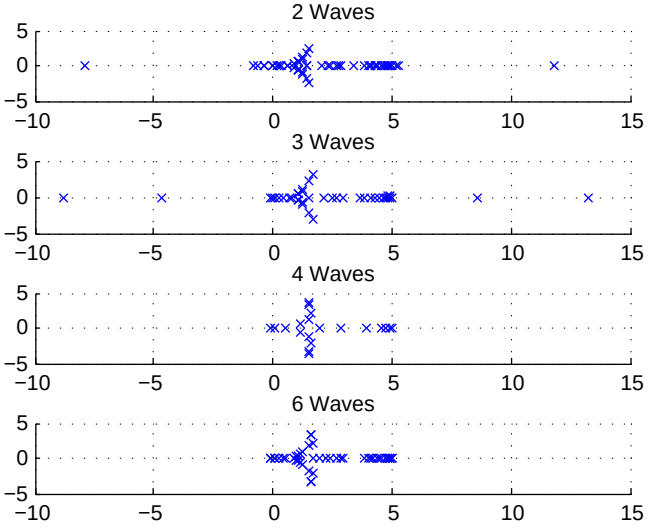

Figure 6: Additive Schwarz $\gamma^2 = 130.69743296$ – waves extended by shifted problem

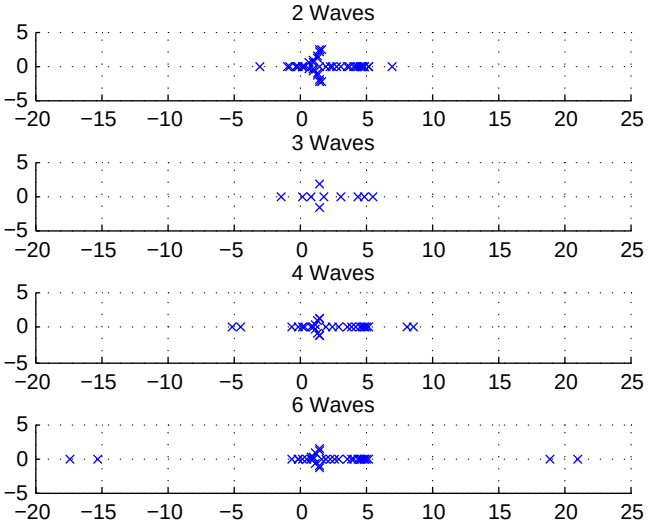

Figure 7: Hybrid Schwarz $\gamma^2 = 130.69743296$ – waves extended by shifted problem

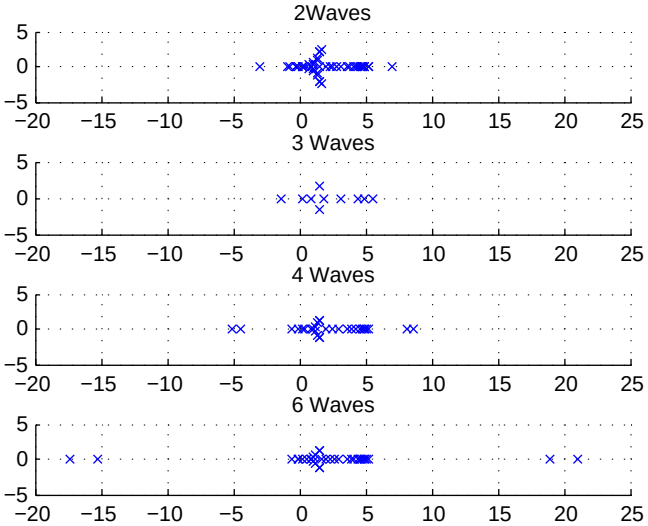

Figure 8: Additive Schwarz $\gamma^2 = 241.65664799$ – waves extended by shifted problem

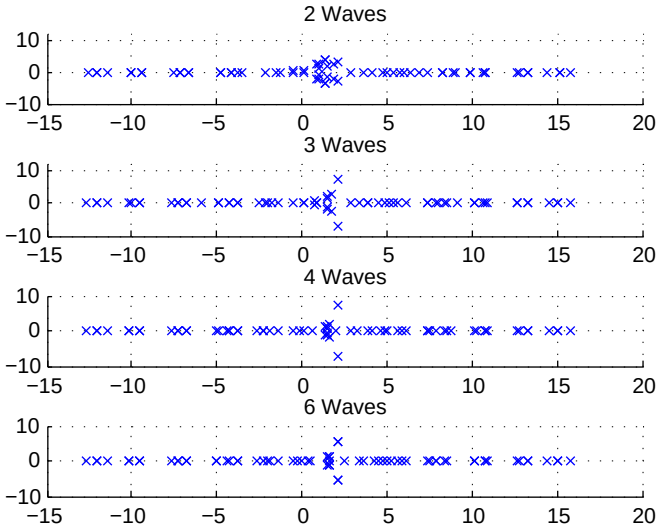

becoming small in absolute value. The positive real eigenvalues also decrease in absolute value, making the real eigenvalues, both positive and negative, all of more comparable size. Furthermore, the complex eigenvalues pull away from the imaginary axis and can be contained within a bounded circular subset of the right half plane. This configuration of the spectrum suggests that the preconditioned system will be more amenable to the convergence of GMRES at an accelerated rate. In terms of the previous comments regarding Theorem 3, the movement of the eigenvalues towards the origin aids in controlling the Dd^{-1} term. The fact that the negative eigenvalues lie near the origin is not a significant step backwards as the non-transformed system also possessed negative eigenvalues near the origin. The movement of the negative eigenvalues towards the eigenvalues near the origin reduces D. The movement of the complex eigenvalues, in the right half plane, away from the imaginary axis combined with the general reduction of the real part of these eigenvalues allows them to be enclosed in a circular region that gives a smaller value to the RC^{-1} term. This is supported by the iteration counts in Tables 10 and 11, as the cases of four, five, and six waves outperform the preconditioners using zero, two, or three waves. In the case of $\gamma^2 = 100$ (Figures 4 and 5), there are two discouraging developments when the transition is made from two plane waves to three. First, there are more complex eigenvalues with large imaginary part, lying near the imaginary axis. Secondly, there is an additional negative real eigenvalue, with large magnitude. These conditions provide a partial explanation for the jump

in the iteration count seen in Tables 10 and 11. Moving to a higher number of waves, the complex eigenvalues pull back from the imaginary axis, and the real eigenvalues become very close to one another in absolute value. In the case of $\gamma^2 = 130.69743296$, the complex eigenvalues are in good positions for all preconditioners shown. With regard to the real eigenvalues, additional outliers appear in the cases of three, four, and six waves. This is consistent with the mildly degraded performance versus two plane waves as seen in Tables 10 and 11. Curiously, performance is improved moving from four to six plane waves in spite of the larger outliers. This indicates that more sophisticated techniques are necessary for a full analysis of the preconditioner. The spectral data pertaining to the cases of $\gamma^2 = 100$ and $\gamma^2 = 130.69743296$ (Figures 4 through 7), correlate with the results presented in Tables 10 and 11. Configurations of eigenvalues that are expected to promoted accelerated GMRES convergence correspond to cases in which good performance is observed. Figure 8 illustrates the spectral analysis of a case that does not converge ($\gamma^2 = 241.65664799$). In the cases of two and three plane waves the location of the complex eigenvalues is less encouraging. In the case of two plane waves, a few complex eigenvalues are in the left half plane. When using three plane waves, some of the complex eigenvalues have small real part. Complex eigenvalues in both the left and right half planes results in the origin being surrounded. Having complex eigenvalues in the right half plane, with small real part, makes it more difficult to bound the eigenvalues in the right half plane in a favorable, closed, circular region. These issues with the complex eigenvalues are resolved in the cases of using four and six plane waves. This also coincides with the reduction in relative residual observed in Tables 10 and 11. In all of the results displayed related to $\gamma^2 = 241.65664799$, the real eigenvalues remain distributed across a significant subset of the real axis. Controlling the size of the real eigenvalues would likely provide significant performance gains for the case of $\gamma^2 = 241.65664799$.

The focus now turns to analysis of the approaches with more limited success. The results listed in Tables 7 and 8 indicate $\gamma^2 = 100$ as a case in which an unmodified Schwarz preconditioner has significant success while the case of $\gamma^2 = 182.49400048$ did not see significant acceleration. These two cases are selected for the spectral analysis, to illustrate conditions when the problem can be sufficiently controlled by controlling H and δ, as well as a configuration of the spectrum when this approach fails. Figures 9 and 11 contain results for the additive variant, while Figures 10 and 12 address the hybrid variant.

The eigenvalues shown in Figures 9 and 10 show similar behavior to the previous successful cases shown in Figures 4 through 7. The complex eigenvalues move away from the imaginary axis as $H\delta^{-1}$ decreases. The negative eigenvalues also move towards the origin, bringing them closer to one another in absolute value. This behavior is similar to what is observed in Figure 4 through 7, which represent successful preconditioning attempts. This is consistent with the results in Tables 7 and 8, which show good performance, in terms of iteration count, for moderate values of $H\delta^{-1}$. Tables 7 and 8 also indicate that attempts to solve the systems for the case of $\gamma^2 = 182.49400048$ are unsuccessful. Examining the spectrum for the case of $\gamma^2 = 182.49400048$ in Figures 11 and 12, conditions

Figure 9: Additive Schwarz $\gamma^2 = 100$ – no waves

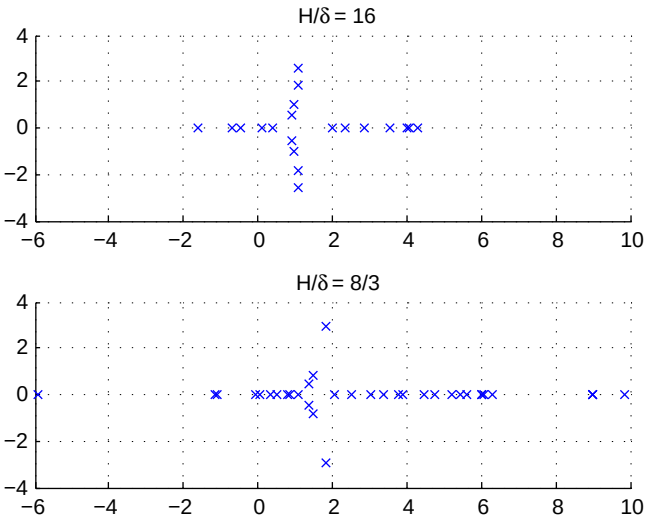

Figure 10: Hybrid Schwarz $\gamma^2 = 100$ – no waves

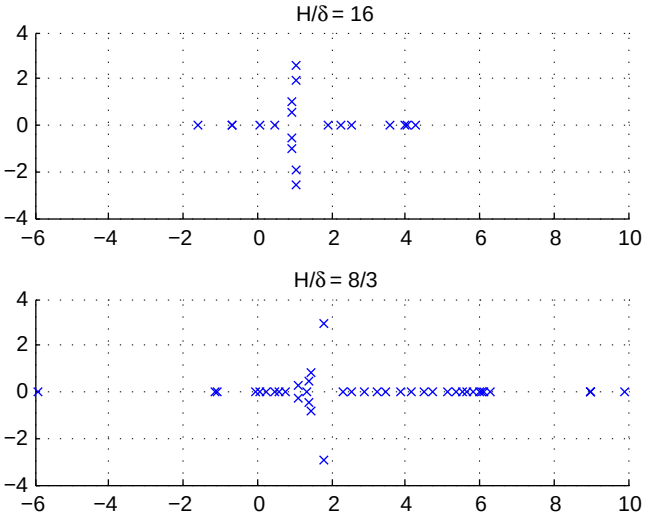

Figure 11: Additive Schwarz $\gamma^2 = 182.49400048$ – no waves

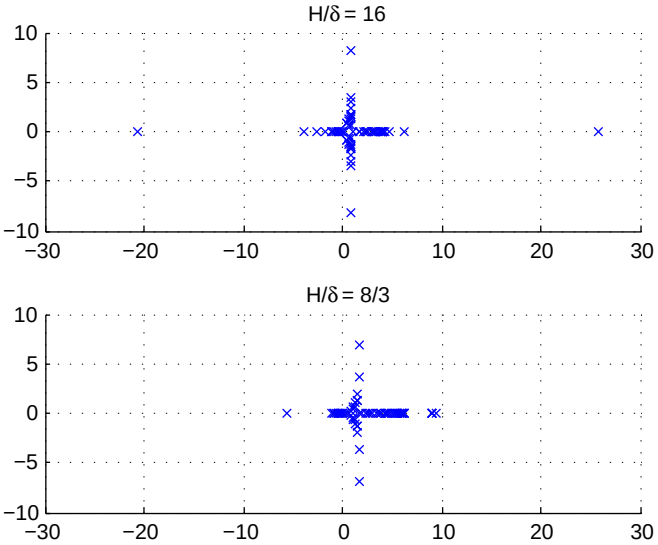

Figure 12: Hybrid Schwarz $\gamma^2 = 182.49400048$ – no waves

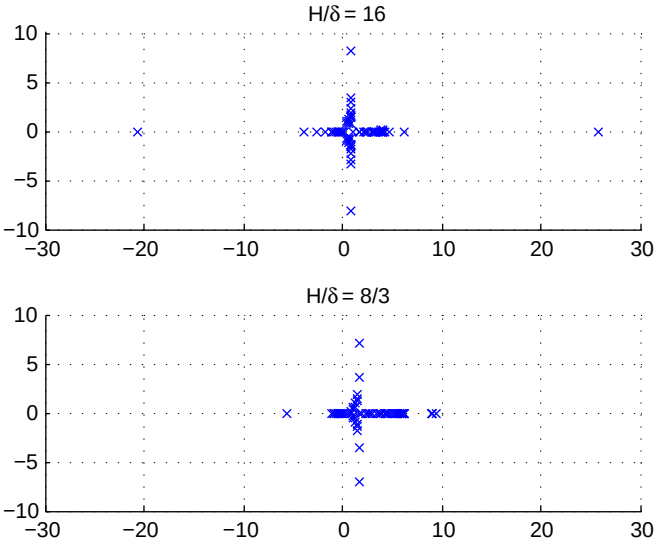

are found to be less favorable for GMRES. For the case of $H\delta^{-1} = 16$, there are complex eigenvalues near the origin that lie close to the imaginary axis. Decreasing $H\delta^{-1}$ to $H\delta^{-1} = 8/3$ is successful in pulling the complex eigenvalues further into the right half plane, but complex eigenvalues with large imaginary part remain. Additionally, outliers along the real axis are reduced, but remain large relative to the other real eigenvalues. The configuration of the spectrum for $\gamma^2 = 182.49400048$ is not the worst case scenario for GMRES, with the origin surrounded, but neither is it sufficiently conducive to rapid convergence of GMRES to allow convergence in the allotted iterations. In terms of Theorem 3, the outlier eigenvalues result in a large value for Dd^{-1}. Eigenvalues located near the imaginary axis, and to a lesser extent other complex eigenvalues with large imaginary part, require that a circle centered about a point on the positive, real axis have large radius to enclose them. Thus, the RC^{-1} term will be less favorable. This is borne out in the results in Tables 7 and 8, where reducing $H\delta^{-1}$ is unable to bring about convergence in the allotted iterations.

Turning to the proposed preconditioner using an extended coarse space in which the plane wave conditions are extended by harmonic extension, two values of γ^2 are selected to analyze the results, for $H^{-1} = 4$ and $\delta = 1/64$, with a variable number of plane waves. Following the results in Tables 12 and 13, $\gamma^2 = 100$ is selected as a convergent case, and $\gamma^2 = 182.49400048$ as an example of a case where convergence is not achieved. In Figures 17 and 18, H is reduced to $H^{-1} = 8$, and $H\delta^{-1}$ is fixed at 4. The spectra for four values of γ^2 are examined for the additive and hybrid variants respectively. This data is drawn from the experiments with variable overlap in Tables 16 and 17 as these experiments utilized a coarse mesh of $H^{-1} = 8$, which demonstrated a greater degree of success when using harmonic extension.

The spectra resulting from applying the additive and hybrid variants, with harmonic extension, to the case of $\gamma^2 = 100$ are displayed in Figure 13 and Figure 14 respectively. Good results, very similar to Figures 4 and 5, are observed. The imaginary parts of the complex eigenvalues are reduced. Also, the absolute values of the real eigenvalues are reduced. These results are consistent with the results in Tables 12 and 13 which show iteration counts similar to those in the case of constructing auxiliary basis functions by extending the plane wave conditions with the shifted problem (13). The situation illustrated for $\gamma^2 = 182.49400048$ in Figures 15 and 16 corresponds to a much less successful problem. Complex eigenvalues, albeit with small imaginary part, are found close to the origin. In the case of $\gamma^2 = 182.48400048$, the complex eigenvalues are confined to the right half plane, but lie much closer to the origin than in the case of either extension by the shifted problem (13), or no auxiliary basis functions. When attempting to contain the eigenvalues in a closed circular region, these eigenvalues in the right half plane near the imaginary axis require that the circular region have a radius (R) that is large relative to the positive real number (C) about which it is centered. This results in a larger value of RC^{-1} as in Theorem 3, which suggests GMRES will converge at a slower rate. This configuration of eigenvalues shown in Figures 15 and 16 is expected to be worse for GMRES convergence than those seen in Figures 11 and 12 or Figures

Figure 13: Additive Schwarz $\gamma^2 = 100$ – harmonic extension

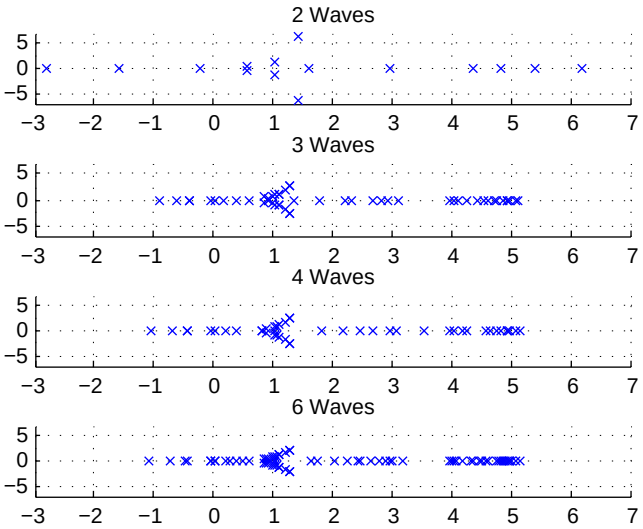

Figure 14: Hybrid Schwarz $\gamma^2 = 100$ – harmonic extension

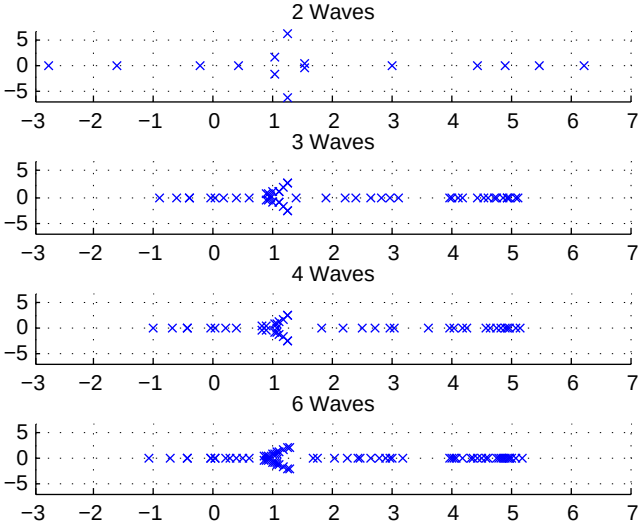

Figure 15: Additive Schwarz $\gamma^2 = 182.49400048$ – harmonic extension

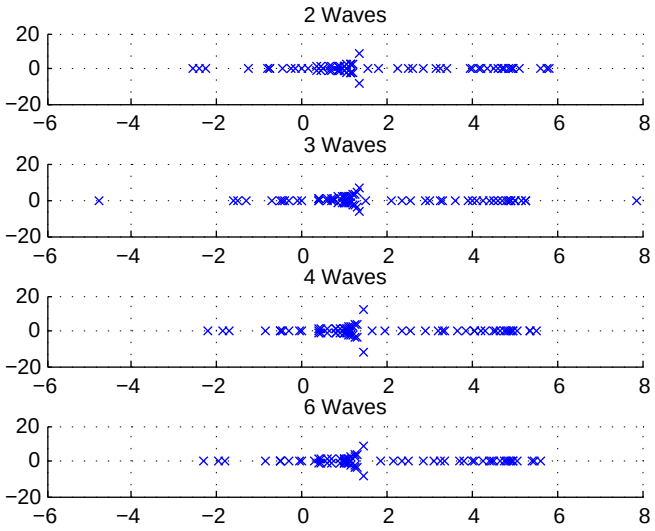

Figure 16: Hybrid Schwarz $\gamma^2 = 182.49400048$ – harmonic extension

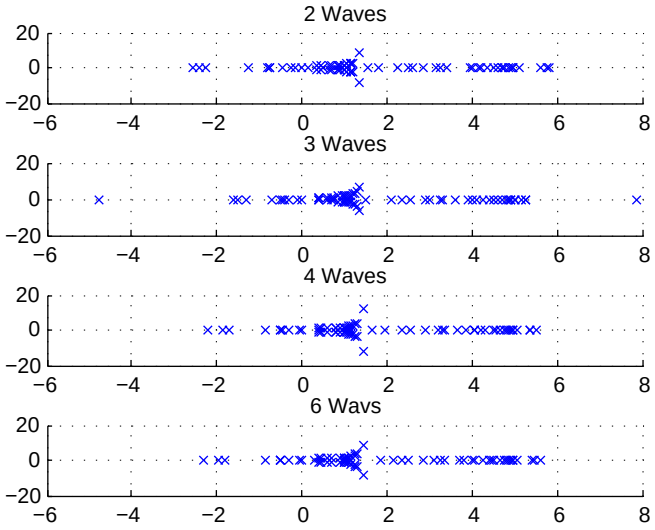

Figure 17: Additive Schwarz $H^{-1} = 8$, $H/\delta = 4$ – harmonic extension

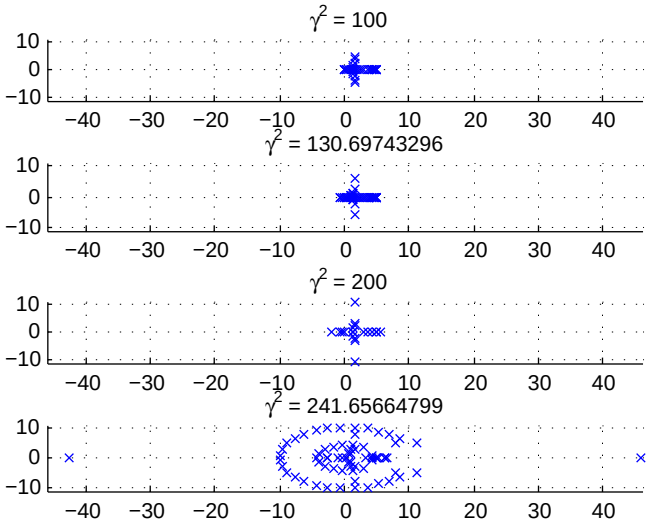

Figure 18: Hybrid Schwarz $H^{-1} = 8$, $H/\delta = 4$ – harmonic extension

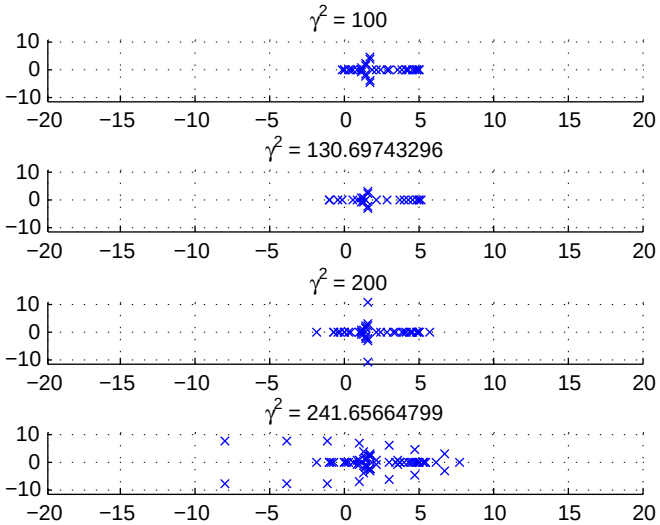

6 and 7. Indeed, this is borne out in the iteration counts in Tables 12 and 13. In the case of using harmonic extension with $H^{-1} = 8$ and $H\delta^{-1} = 4$, Figures 17 and 18 indicate favorable configurations of the spectrum for the same three values of γ^2 that converge in the case of constructing the extended coarse space by the shifted problem (13) with $H^{-1} = 4$. In the case of $\gamma^2 = 241.65664799$, the origin is surrounded in a severe manner, which is consistent with the lack of convergence seen in Tables 16 and 17. The spectral analysis pertaining to harmonic extension indicates that a tighter control on H is necessary for harmonic extension, than extension by the shifted problem (13). Convergence is slower in the case of $\gamma^2 = 200$, but attainable, provided H is small enough. Additionally, the spectral results pertaining to $\gamma^2 = 241.65664799$ indicate that the situation in the failure of harmonic extension is worse than that for extension by the shifted problem (13).

7 Inverse Iteration

Lastly, inverse iteration is presented as an application of the preconditioner using an extended coarse space constructed by the shifted problem (13). The domain $\Omega = [0,1] \times [0,1]$ is selected to reduce the size of H without requiring a larger number of unknowns. A fixed set of four plane waves, as described in Table 9, is used with a fine mesh of $h^{-1} = 256$, a coarse mesh of $H^{-1} = 16$, and an overlap of $\delta = 1/128$. The eigenvalues of the discrete operator $-\Delta$ are known, see [7]. The eigenvalue

$$\lambda = (256^2)\left(4\left(\sin^2\left(\frac{3\pi}{512}\right) + \sin^2\left(\frac{3\pi}{512}\right)\right)\right) \tag{17}$$

is selected as it is a simple eigenvalue, giving the iteration only one eigenvector to converge to. The eigenvector is known and corresponds to a discretization of the continuous function (18),

$$u(x,y) = \sin(3\pi x)\sin(3\pi y). \tag{18}$$

Inverse iteration proceeds by solving the equation

$$(A - \mu I)\, x^{(k)} = x^{(k-1)} \tag{19}$$

for μ near an eigenvalue and $x^{(0)}$ is an initial guess. Ideally, $x^{(k)}$ converges to an eigenvector rapidly. In this experiment, A is taken to be the discrete Laplacian, and $\mu = 177.63281$ is selected to be near λ. This gives

$$|\lambda - \mu| = 4.4717084\text{e-}06.$$

Additionally, (19) becomes (11) with $\gamma^2 = \mu$ and $f = x^{(k-1)}$, $u = x^{(k)}$. The initial guess is taken to be 1 on all points interior to Ω and 0 on $\partial\Omega$. Initial results pertaining to the convergence of inverse iteration are displayed in Table

Table 18: Inverse Iteration Convergence $\mu = 177.63281$

	$\cos(\theta)$	θ	GMRES Iterations
$x^{(1)}$	9.9999999999993028e-01	3.7342202920144023e-07	150
$x^{(2)}$	9.9999999999999600e-01	8.9406967163085964e-08	320
$x^{(3)}$	1.0000000000000078e+00	(1.2467205919833109e-07)$\imath$	245

18. To measure how close the iterate is to being an eigenvector, the angle between the iterate and the exact eigenvector is measured by

$$\cos(\theta) = \frac{\langle x^{(k)}, u \rangle}{\|x^{(k)}\| \, \|u\|}$$

where u is the exact eigenvector.

Inverse iteration proves successful. After one iteration, the angle between the iterate and the exact solution is on the order of 10^{-7}. Further iterations are able to reduce this to 10^{-8}, although it stagnates at that point. Additionally, due to roundoff error, $\cos(\theta)$ exceeds 1, precluding such an analysis. Since no significant progress, in terms of θ, is made after the first iteration, the convergence behavior of the GMRES iteration for the first iteration is examined more closely in Table 19. The usual convergence tests are overridden and results for the first 450 GMRES iterations are displayed. Every five iterations, the angle θ is checked, as well as the relative residual.

In Table 19, rapid reduction of both θ and the relative residual is seen through the first 25 iterations. The solution after the first 25 iterations is shown in Figure 19. For comparison, the exact eigenvector is displayed in Figure 20. Additional iterations further reduce both θ and the relative residual, but at a much slower rate. After 105 GMRES iterations, θ attains its terminal value of roughly 10^{-7}. The residual reaches the convergence criterion of a size less than 10^{-8} at 150 iterations. Continuing further, the relative residual is further reduced to roughly 10^{-13}, after 450 GMRES iterations, but θ remains on the order of 10^{-7}. Inverse iteration attains a solution near an eigenvector, and stagnates. Convergence to an approximate eigenvector occurs prior to the convergence of GMRES, measured by residuals.

The proposed Schwarz preconditioner proves suitable for inverse iteration. The results indicate that a shift μ that results in rapid convergence of inverse iteration results in a problem that is tractable for the proposed preconditioner. The GMRES iteration count, for the first iteration, is good. The results in Table 18 indicate that there is little reason to proceed beyond the first iteration. The results in Table 19 show that the point at which θ stagnates is reached prior to the convergence of GMRES, with respect to the size of the residual.

Table 19: Inverse Iteration Convergence $\mu = 177.63281$, first iteration detail

GMRES Iterations	θ	Relative Residual
5	1.1718342e+00	6.4557306e-03
10	2.5117419e-01	5.2623644e-04
15	3.4354446e-03	1.8665505e-05
20	4.7137513e-04	1.3510048e-06
25	2.8809031e-04	5.7036202e-07
30	2.5456238e-04	5.4100448e-07
35	2.9087708e-04	7.3798994e-07
40	3.0472594e-04	8.3836569e-07
45	3.1226001e-04	8.9177995e-07
50	3.1696442e-04	9.2460320e-07
60	3.2249828e-04	9.6273025e-07
70	3.2564147e-04	9.8416711e-07
80	2.3348607e-04	6.3013000e-07
90	2.1708927e-04	5.6388845e-07
100	2.2535759e-06	4.5326443e-07
105	4.3774935e-07	5.3411952e-08
110	4.1803003e-07	4.4842162e-08
120	4.1268416e-07	4.4164077e-08
130	4.0998508e-07	4.4017933e-08
140	4.3084718e-07	4.3954080e-08
150	3.7342203e-07	4.5940920e-09
200	3.5855798e-07	1.3669559e-09
250	3.7073658e-07	1.5497012e-10
300	3.7163388e-07	1.5223128e-11
350	3.6833307e-07	2.1273620e-11
400	3.9028528e-07	3.5113924e-13
450	4.1052631e-07	2.6650111e-13

Figure 19: Solution after 25 GMRES iterations

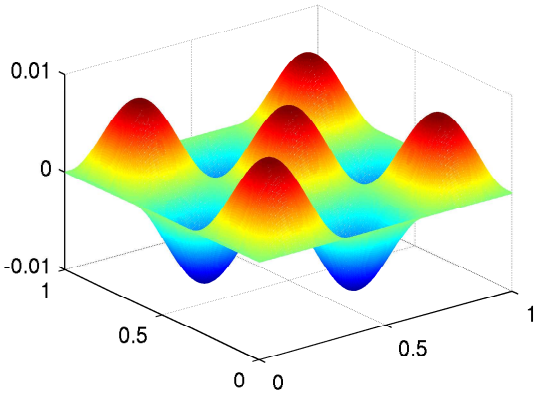

Figure 20: Exact eigenvector

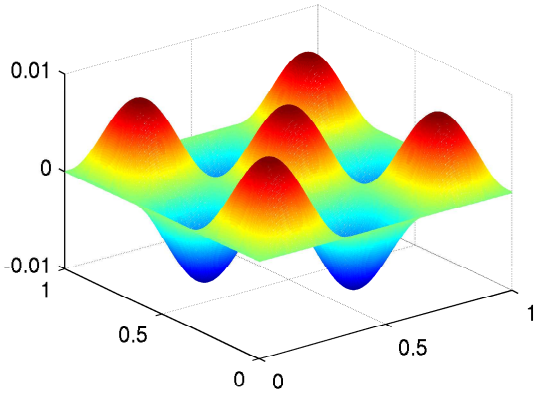

8 Concluding Remarks

The preconditioners considered enjoy varying degrees of success. The most successful approach considered uses an extended coarse space in which auxiliary basis functions are constructed via the shifted problem (13). The approach of controlling the size of H, with no auxiliary basis functions, is effective, but requires increasingly smaller H as the parameter γ^2 increases as seen in the previous work by Cai and Widlund, see [3] [4]. Results displayed in Tables 7 and 8 show very limited success for attempting to control the behavior of the system using only H and δ as γ^2 increases. Above $\gamma^2 = 130.69743296$, only one attempt at $\gamma^2 = 200$ converges, but does so at over 300 GMRES iterations. The results in Tables 10 and 11 indicate that, in the case of basis functions constructed by the shifted problem, (13) allows a very large value of $H^{-1} = 4$ to be used while still attaining convergence in under 200 GMRES iterations. However, for larger values of γ^2, the proposed preconditioners become sensitive to the conditioning of the non-transformed system. This shift is also found to be suitable for rapid convergence of inverse iteration. Tables 12 and 13 indicate that this is not the case for basis functions constructed by harmonic extension. Only the two smallest shifts result in tractable problems. It is only in Tables 16 and 17, when the coarse mesh is reduced to $H^{-1} = 8$, that a few more of the test cases converge. In these additional cases, convergence typically requires more than 200 GMRES iterations. A combination of controlling the size of H and auxiliary basis functions constructed from the shifted problem (13) enjoys the greatest success in terms of accelerating convergence.

The results of adding additional plane wave based basis functions point to a progression of larger classes of indefinite problems that become tractable. The classical Schwarz preconditioners exhibit poor performance, as the number of negative eigenvalues is increased, as well as situations where the problem is close to singular. When the coarse space is extended by adding basis functions based on a small number of plane waves, problems with more negative eigenvalues, such as in the case of $\gamma^2 = 200$, become tractable, but problems that are nearly singular continue to present difficulties. As the number of plane waves used is further increased, some of the problems with γ^2 near an eigenvalue of the operator become tractable, but problems with both γ^2 near an eigenvalue and a large number of negative eigenvalues remain intractable.

In some cases, sudden jumps in the iteration count are observed as additional plane waves are used to extend the coarse space. Justification of these jumps can be found in the relevant spectral analysis. The spectral analysis of the proposed preconditioners reveals that the preconditioning techniques are indeed successful in constructing systems conducive to accelerated convergence of GMRES. The successful cases share similar spectra indicating that preconditioning, rather than fortuitous circumstances, is responsible for improved performance. Additionally, in the cases where a jump in the iteration count is observed, a behavior is observed in the spectral analysis that provides an explanation. New eigenvalues with either large real part or large complex part appear, but come back under control with the addition of more plane wave based basis functions. In all

successful cases considered, the eigenvalues in the left half plane are restricted to the negative real axis. Furthermore, the more effective preconditioners gather the negative eigenvalues closer to the origin, reducing the absolute value of the negative eigenvalues relative to the negative eigenvalue with smallest absolute value. In the right half plane, the real eigenvalues gather towards the smallest positive eigenvalue. Also, the complex eigenvalues pull away from the imaginary axis. This allows the eigenvalues in the right half plane to be confined to a bounded subset of the right half plane. The situation in the right half plane combined with the negative eigenvalues gathering into a small subset of the negative real axis give a situation conducive to rapid GMRES convergence.

In addition to indicating the successful cases, the spectral techniques also indicate the unsuccessful attempts. In the cases where convergence is not achieved, the spectrum of the upper Hessenberg matrix varies, but consistently assumes configurations less supportive of rapid GMRES convergence. In most cases, complex eigenvalues remain near the imaginary axis, at times even appearing on both sides of the imaginary axis. Additionally, the real eigenvalues tend to remain distributed across a much larger subset of the real axis. In the case of Figures 15 and 16, the locations of the real eigenvalues are good, but the locations of the complex eigenvalues present problems. While the necessary spectral analysis indicated by the theory is too involved for a careful comparison of the successful cases, it is an effective tool for distinguishing preconditioners that achieve reasonable acceleration from those that provide little or no benefit.

Among the Schwarz preconditioners considered, extending the coarse space with basis functions constructed from plane waves, using the shifted problem (13) along with small overlap, proves to be the best choice. In the case of $\gamma^2 = 182.49$, an increase in overlap from $H\delta^{-1} = 16$ to $H\delta^{-1} = 8$ proves catastrophic for performance. This indicates that small overlap should be preferred for performance reasons. The approach of extending the coarse space is able to handle the various choices of γ^2 that are intractable for the Schwarz preconditioners without an extended coarse space. This makes a larger class of problems tractable without requiring a very small value of H. Consistently reducing the size of H results in an increased number of local problems. Constructing the auxiliary basis functions by harmonic extension enjoys some success with smaller H. The need to control the size of H is a common theme in applying domain decomposition methods to indefinite systems, but a small H also significantly increases the size of the global problem, both by increasing the size of the unextended coarse space, and increasing the number of auxiliary basis functions added. In the experiments considered, extension by the shifted problem (13) achieves very good performance with five or six wave equations and $H^{-1} = 4$, which corresponds to a coarse space of dimension 609 or 721 respectively. The smaller value of $H^{-1} = 8$ necessary to use harmonic extension leads to a coarse space of dimension 2145 when using four plane waves. Moreover, the use of a harmonic extension results in slower convergence. Tables 14 and 15 indicate that for the case of extension by (13), increased overlap does little to improve performance. Additionally, small overlap helps limit the size of the local problems. Given the lack of significant performance gains from allowing generous

overlap, this further indicates that small overlap should be preferred. Additionally, reduction of $H\delta^{-1}$, at times, requires a reduction of H as increasing δ can lead to indefinite local problems if H is large.

The experiments in Section 7 indicate a successful application of the proposed preconditioner to inverse iteration. The use of auxiliary basis functions constructed by the shifted problem (13) and a smaller value of H allow for the calculation of an eigenvector corresponding to the eigenvalue λ given by (17). Using $H^{-1} = 16$, the case of $\gamma^2 = 177.63281$ is tractable. Inverse iteration yields a solution that is close to the exact eigenvector after a single iteration, with an acceptable GMRES iteration count. The results also indicate that there is little point in continuing beyond this first iteration.

The most effective Schwarz preconditioner considered is a Schwarz method with a coarse space expanded by plane wave based basis functions constructed by using the shifted problem (13), and local problems with small overlap. The use of an extended coarse space gives the best performance increases seen in the results presented. Additionally, this approach enjoys sufficient robustness to make a combination of GMRES with such a preconditioner suitable for eigenvector calculation via inverse iteration. This approach also allows for larger values of γ^2 without requiring as significant a reduction of H as is necessary for approaches without an extended coarse space. Additionally, reductions of H further enhance the robustness of the preconditioner. Increasing the overlap proves to not be of significant use in accelerating the convergence. Keeping the overlap small also helps to maintain positive definite local problems with large H. Using large values of H maintains a smaller coarse space dimension, and limits the number of local problems. Augmenting the coarse space with basis functions, constructed using the shifted problem (13) in conjunction with small overlap, yields a robust preconditioner that provides for good acceleration of convergence while also keeping down the size of the global problem and maintaining a small number of local problems.

References

[1] Dietrich Braess. *Finite Elements*. Cambridge University Press, second edition, 2001. Translated by Larry L. Schumaker.

[2] Xiao-Chuan Cai, William D. Gropp, and David E. Keyes. A comparison of some domain decomposition and ILU preconditioned iterative methods for nonsymmetric elliptic problems. *Numer. Linear Algebra Appl.*, 1(5):477–504, 1994.

[3] Xiao-Chuan Cai and Olof B. Widlund. Domain decomposition algorithms for indefinite elliptic problems. *SIAM Journal on Scientific and Statistical Computing*, 13(1):243–258, 1992.

[4] Xiao-Chuan Cai and Olof B. Widlund. Multiplicative Schwarz algorithms for some nonsymmetric and indefinite problems. *SIAM Journal on Numerical Analysis*, 30(4):936–952, 1993.

[5] Charbel Farhat and Jing Li. An iterative domain decomposition method for the solution of a class of indefinite problems in computational structural dynamics. *Applied Numerical Mathematics*, 54(2):150–166, 2005.

[6] Charbel Farhat and Francois-Xavier Roux. A method of finite element tearing and interconnecting and its parallel solution algorithm. *International Journal for Numerical Methods in Engineering*, 32(6):1205–1227, 1991.

[7] Arieh Iserles. *A First Course in the Numerical Analysis of Differential Equations*. Cambridge University Press, New York, 2004.

[8] Jing Li and Xuemin Tu. Convergence analysis of a balancing domain decomposition method for solving interior Helmholtz equations. submitted to Numer. Linear Algebra Appl., 2008.

[9] Thomas A. Manteuffel. The Tchebychev iteration for nonsymmetric linear systems. *Numerische Mathematik*, 28(3):307–327, 1977.

[10] Dianne P. O'Leary and Olof Widlund. Capacitance matrix methods for the Helmholtz equation on general three-dimensional regions. *Mathematics of Computation*, 33(147):849–879, 1979.

[11] Youcef Saad and Martin H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.

[12] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, second edition, 2003.

[13] Andrea Toselli and Olof Widlund. *Domain Decomposition Methods – Algorithms and Theory*. Springer – Verlag, New York, 2005.

[14] Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra*. SIAM, 1997.

[15] Jinchao Xu and Xiao-Chuan Cai. A preconditioned GMRES method for nonsymmetric or indefinite problems. *Mathematics of Computation*, 59(200):311–319, 1992.