

Additive Schwarz Methods with Nonreflecting Boundary Conditions for the Parallel Computation of Helmholtz Problems

Lois C. McInnes, Romeo F. Susan-Resiga, David E. Keyes,
and Hafiz M. Atassi

1. Introduction

Recent advances in discretizations and preconditioners for solving the exterior Helmholtz problem are combined in a single code and their benefits evaluated on a parameterized model. Motivated by large-scale simulations, we consider iterative parallel domain decomposition algorithms of additive Schwarz type. The preconditioning action in such algorithms can be built out of nonoverlapping or overlapping subdomain solutions with homogeneous Sommerfeld-type transmission conditions on the artificially introduced subdomain interfaces. Generalizing the usual Dirichlet Schwarz interface conditions, such Sommerfeld-type conditions avoid the possibility of resonant modes and thereby assure the uniqueness of the solution in each subdomain.

The physical parameters of wavenumber and scatterer diameter and the numerical parameters of outer boundary diameter, mesh spacing, subdomain diameter, subdomain aspect ratio and orientation, subdomain overlap, subdomain solution quality (in the preconditioner), and Krylov subspace dimension interact in various ways in determining the overall convergence rate. Many of these interactions are not yet understood theoretically, thus creating interest in experimental investigation. Using the linear system solvers from the Portable Extensible Toolkit for Scientific Computation (PETSc), we begin to investigate the large parameter space and recommend certain effective algorithmic “tunings” that we believe will be valid for (at least) two-dimensional problems on distributed-memory parallel machines.

The external Helmholtz problem is the basic model of farfield propagation of waves in the frequency domain. This problem is challenging due to large discretized system sizes that arise because the computational grid must be sufficiently refined

1991 *Mathematics Subject Classification*. Primary 65N55; Secondary 65F10, 65N30, 65Y05.
Supported by U.S. Department of Energy, under Contract W-31-109-Eng-38.

Supported in part by National Science Foundation grant ECS-9527169.

Supported in part by National Science Foundation grant ECS-9527169 and by NASA Contracts NAS1-19480 and NAS1-97046.

Supported in part by National Science Foundation grant ECS-9527169.

throughout the entire problem domain to resolve monochromatic waves, which requires approximately 10–20 gridpoints per wavelength for commonly used second-order discretizations. Moreover, the conventional farfield closure, the Sommerfeld radiation condition, is accurate only at large distances from centrally located scatterers and must be replaced at the artificial outer boundary by a nonreflecting boundary condition to obtain a computational domain of practical size. To this end we employ a Dirichlet-to-Neumann (DtN) map on the outer boundary, which provides an exact boundary condition at finite distances from the scatterer. The DtN map is nonlocal but does not introduce higher order derivatives.

Although the discretized Helmholtz linear system matrix is sparse, for a large number of equations direct methods are inadequate. Moreover, the Helmholtz operator tends to be indefinite for practical values of the wavenumber and the mesh parameter, leading to ill-conditioning. As a result conventional iterative methods do not converge for all values of the wavenumber, or may converge very slowly. For example, resonances can occur when conventional Schwarz-based preconditioners are assembled from Dirichlet subdomain problems. In view of these difficulties, this work focuses on developing a family of parallel Krylov-Schwarz algorithms for Helmholtz problems based on subdomain problems with approximate local transmission boundary conditions.

It is difficult to do justice to previous work on a century-old problem that has been revisited with vigor by specialists in diverse application areas in recent years. However, we select a few references that have been of inspirational value to our own work. Keller & Givoli [13] and Harari & Hughes [10] employed the global Dirichlet-to-Neumann map (a pseudo-differential operator) as a non-reflecting BC for the truncated domain (circle or sphere) and also experimented with truncating the complexity implied by the full DtN map. Després in his doctoral dissertation [7] pioneered domain decomposition for Helmholtz problems with first-order transmission conditions on nonoverlapping interfaces between subdomains, proving convergence to a unique solution. Ghanemi [8] combined nonlocal transmission conditions with Després-style iteration. She also obtained a better rate of convergence through under-relaxation of the nonoverlapping interface conditions. Douglas & Meade [12] advocated second-order local transmission conditions for both subdomain interfaces and the outer nonreflecting boundary condition and employed underrelaxed iterations. Our colleagues in domain-decomposed Helmholtz research, Cai, Casarin, Elliott & Widlund [3] introduced Schwarz-style overlapping, used first-order transmission conditions on the overlapped interfaces, calling attention to the *wavelap* parameter, which measures the number of wavelengths in the overlap region. They have also noted the importance of a (relatively fine) coarse grid component in the Schwarz preconditioner to overcome the elliptic ill conditioning that arises asymptotically for small mesh spacing.

2. Mathematical Formulation

The scalar Helmholtz equation,

$$(1) \quad -\nabla^2 u - k^2 u = 0,$$

is derived by assuming a time-harmonic variation in the solution of the second-order, constant-coefficient wave equation. A discussion of the hierarchy of models that reduce in their purest form to (1) is given in [1]. The parameter k is the reduced wavenumber, i.e. $2\pi/\lambda$, where λ is the wavelength. In the general case

anisotropy and spatially varying coefficients may be present, although in this work we restrict attention to a homogeneous isotropic problem.

We explicitly consider only the external Helmholtz problem, in which the perturbation field u is driven by a boundary condition inhomogeneity on a nearfield boundary Γ , subdivided into Dirichlet, Γ_D , and Neumann, Γ_N , segments, one of which may be trivial. Scatterer boundary conditions of sufficient generality are

$$(2) \quad u = g_D \text{ on } \Gamma_D \quad \text{and} \quad \partial u / \partial \nu = g_N \text{ on } \Gamma_N.$$

The Sommerfeld radiation condition,

$$\lim_{r \rightarrow \infty} r^{(d-1)/2} \left(\frac{\partial u}{\partial r} + iku \right) = 0,$$

may be regarded as an expression of causality for the wave equation, in that there can be no incoming waves at infinity. Thus this condition acts as a filter that selects only the outgoing waves. (The Sommerfeld sign convention depends upon the sign convention of the exponent in the time-harmonic factor of the wave equation.)

The Dirichlet-to-Neumann Map. In finite computations the Sommerfeld boundary condition must be applied at finite distance. For B a circle (or a sphere in 3D), an integro-differential operator may be derived that maps the values of u on the artificial exterior boundary, B , to the normal derivative of u on B [9]. In two-dimensional problems this leads to an infinite series involving Hankel functions, which may be differentiated in the radial direction and evaluated at $r = R$ to yield on B :

$$(3) \quad \frac{\partial u}{\partial \nu}(R, \theta) \equiv (\mathcal{M}u)(R, \theta) = \frac{k}{\pi} \sum_{n=0}^{\infty} \frac{H_n^{(2)'}(kR)}{H_n^{(2)}(kR)} \int_0^{2\pi} \cos n(\theta - \theta') u(R, \theta') d\theta'.$$

This expression defines the Dirichlet-to-Neumann map, \mathcal{M} , where the infinite sum may be truncated to a finite approximation of $N \geq kR$ terms [10].

Finite Element Discretization. We use a Galerkin finite element formulation with isoparametric four-noded quadrilateral elements to discretize the problem specified by (1), (2), and (3) and thereby form a linear system,

$$(4) \quad Au = b.$$

Details about this system, as well as elementary properties of its pseudospectrum, are given in [14]. The Sommerfeld boundary condition has the effect of pushing the portion of the real spectrum that is close to (or at) zero in the Dirichlet case away from the origin, in the imaginary direction.

3. Schwarz-based Solution Algorithms

Brought into prominence in the current era of cache-based architectures, additive Schwarz methods have become the “workhorses” of parallel preconditioners for elliptically dominated partial differential equations in the last decade, and have recently been applied to Helmholtz problems in [3]. Although a variety of Schwarz-based techniques have been considered for Helmholtz problems by ourselves and others, this chapter focuses on accelerated overlapping iterative methods for the solution of the discrete equations (4).

Closely related to the overlapping Krylov-Schwarz method presented herein is a nonoverlapping stationary iterative scheme of Resiga and Atassi [16] for solving (1) independently in subdomains. This approach, which follows the work of Després

[7] and Ghanemi [8], uses under-relaxed impedance-type boundary conditions on subdomain interfaces and the DtN map on the exterior boundary, and features concurrent update of all subdomains for parallelism. Details of the scheme are presented in [14], where it is compared with the method featured herein.

We investigate an additive Schwarz preconditioner based on overlapping subdomains, which is accelerated by a Krylov method, such as a complex-valued version of GMRES [15]. An overlapping decomposition is defined by splitting the computational domain Ω into nonoverlapping subdomains Ω_i , with boundaries $\partial\Omega_i$, and extending each except where cut off by \cdot , and B to subdomains Ω'_i and interfaces $\partial\Omega'_i$. A (Boolean) restriction operator R'_i extracts those elements of the global vector that are local to extended subdomain i , and a prolongation operator R_i^T (without the $'$) prolongs elements local to the nonextended subdomain i back to the global vector, by extension with zeros. Let $A_i^{-1}u_i$ denote the action of solving in each extended subdomain,

$$\begin{aligned} \mathcal{L}_i u_i &= f_i \text{ in } \Omega'_i, \\ \mathcal{B}_i u_i &= \begin{cases} 0 & \text{on } \partial\Omega'_i - \partial\Omega \\ g_i & \text{on } \partial\Omega_i \cap \partial\Omega \end{cases}. \end{aligned}$$

Then a Schwarz projection M_i is defined by $M_i = R_i^T A_i^{-1} R'_i A$, and a Schwarz-preconditioned operator is then defined through $M = \sum_i M_i$. The system (4) is replaced with $Mu = \sum_i R_i^T A_i^{-1} R'_i b$. We iterate on this system with a Krylov method until convergence. This particular combination of extended restriction and unextended prolongation operators is designated ‘‘Restricted Additive Schwarz’’ (RAS) [4] and has been found by us and by others to be superior to standard Additive Schwarz, in which the prolongation is carried out to the extended subdomains. We have described the left-preconditioned form of RAS above. In practice, when comparing preconditioners as in this paper, we employ right-preconditioning, so that the residual norms available as a by-product in GMRES are not scaled by the preconditioner.


Each interior point of the original domain remains an interior point of at least one subdomain, and a standard PDE discretization is applied there. The extended interior subdomain interfaces are handled with Sommerfeld-type boundary conditions *in the preconditioner only*. Except for the use of homogeneous Sommerfeld-type boundary conditions, this method falls under the indefinite Schwarz theory of Cai & Widlund [5, 6].

4. A Model Helmholtz Problem

We use a model problem with a known exact solution to study the truncation error of this algorithm, along with the algebraic convergence rate. This model problem was employed by Givoli and Keller [9] for their demonstration of the advantages of the DtN map. The geometry and notation are defined in Figure 1. The eccentricity of the bounding circles spoils the application on B of simple boundary conditions based on normal incidence.

The explicit analytical solution u^* permits tabulation of the pointwise relative error in the numerical solution u_i^m at point i after iteration m as follows:

$$(5) \quad e_i^m = \frac{|u_i^m - u^*(x_i)|}{|u^*(x_i)|}.$$



- Inner boundary, Γ :
circle centered at $S = (-\frac{1}{2}, 0)$
with radius $R_1 = 1$
- Outer boundary, B :
circle centered at $O = (0, 0)$
with radius $R_2 = 2$
- Neumann BC on Γ , from monopole at S
- Dirichlet-to-Neumann map on B

FIGURE 1. Eccentric annulus model problem domain.

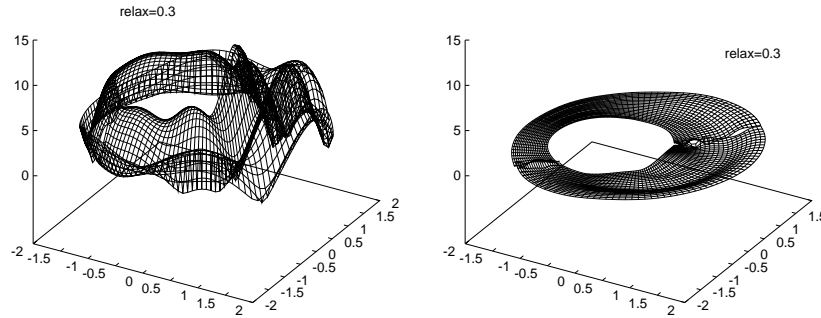


FIGURE 2. Relative error distribution of converged solution with Sommerfeld (left) and the DtN map (right) exterior boundary conditions (nonoverlapping stationary iterative method, $k = 4$).

On the basis of numerical experiments that are discussed in [14], we have determined the truncation error “floor” beneath which we need not obtain algebraic convergence of (4). We have also compared a variety of schemes for attaining that level of error. Figure 2 shows a typical plot of the relative error (5) using a nonoverlapping stationary scheme for a two-subdomain case with $k = 4$. This picture emphasizes the well known advantage of a perfectly nonreflecting DtN map over a Sommerfeld boundary condition.

5. Numerical Results

For our numerical simulations of the Helmholtz problem, we employ the Portable, Extensible Toolkit for Scientific Computing (PETSc) [2], a library that attempts to handle through a uniform interface, in a highly efficient way, the low-level details of the distributed memory hierarchy. One feature that distinguishes it from other freely available libraries of iterative methods is the capability of solving systems defined over the complex numbers. Also, PETSc’s preconditioners make it routine to vary the number of subdomains in each physical dimension into which

TABLE 1. Iteration counts and parallel execution times (in seconds) for different aspect ratio subdomains and different overlaps, 65×256 grid, 32 processors.

Subdomain Shape		Overlap $h/2$		Overlap $3h/2$		Overlap $5h/2$		Overlap $9h/2$	
Procs	Subgrid	Its	Time	Its	Time	Its	Time	Its	Time
$k = 6.5$									
1×32	8:1	131	<i>1.35s</i>	125	1.38s	125	1.46	124	1.77s
2×16	2:1	137	1.54s	128	<i>1.42s</i>	128	1.89s	126	1.45s
4×8	1:2	174	1.96s	145	<i>1.54s</i>	138	1.77s	129	1.79s
8×4	1:8	211	2.98s	169	3.54s	153	<i>2.44s</i>	134	2.57s
$k = 13.0$									
1×32	8:1	159	<i>1.92s</i>	152	1.97s	150	1.99s	146	2.13s
2×16	2:1	182	1.88s	157	<i>1.76s</i>	150	1.83s	147	1.92s
4×8	1:2	195	1.92s	164	<i>1.73s</i>	157	2.08s	153	2.27s
8×4	1:8	224	2.84s	190	<i>2.55s</i>	176	2.73s	158	3.01s

the problem will be partitioned, the amount of overlap between these subdomains, and the quality of the solution process employed on each block in the preconditioner.

5.1. Comparison of Subdomain Shape and Overlap. Like convection problems, Helmholtz problems possess “preferred” directions in that there are dominant directions of wave propagation. Unlike convection problems, these directions are not manifest in the interior equations, which are locally rotationally invariant, but enter through the boundary conditions. It is therefore of interest to study the effect of subdomain size and shape on decomposed preconditioners for Helmholtz problems. We seek to answer two questions initially for a range of wavenumbers k : how does the orientation of the cuts interact with the orientation of the waves, and how much does overlap help to “pave over” the cuts?

Our implementation permits any number of radial and circumferential cuts, provided that all subdomains consist of a rectangular subset of the radial and circumferential indices. For the purpose of playing with aspect ratio in several increments over a large ratio, we select power-of-two size discretizations. Thus, we take 64 mesh cells in the radial direction and 256 in the circumferential. To satisfy the conservative $\lambda/h \geq 20$ in all directions throughout the domain, where $\lambda = 2\pi/k$, we consider $k = 6.5$. We compare this with $k = 13.0$, in which the waves are resolved with only 10 points per wavelength in the worst-resolved part of the domain (near $(x, y) = (2, 0)$). The action of A_i^{-1} is approximated on each subdomain by application of ILU(1), with overlap as tabulated across the column sets, accelerated by restarted GMRES. Wall-clock execution times are measured on 32 nodes of an IBM SP with 120MHz quad-issue Power2 nodes with a 10^{-4} relative residual tolerance.

We readily observe in Table 1 that cuts along constant angle (which are aligned with the dominant radial direction of wave propagation) are preferable over cuts along constant radius. Convergence is very much faster with few “bad” cuts than it is with many “bad” cuts. Overlap is effective in reducing the number of iterations by about 50% in the case of many “bad” cuts, but exhibits a relatively rapid law of diminishing returns in all orientations. Since the cost per iteration rises approximately linearly in the overlap and the convergence rate benefit saturates, the

TABLE 2. Scalability for fixed global problem size, 129×512 grid, $k = 13$

Processors	Iterations	Time (Sec)	Speedup	% Efficiency
1	221	163.01	—	—
2	222	81.06	2.0	100
4	224	37.36	4.4	100
8	228	19.49	8.4	100
16	229	10.85	15.0	93
32	230	6.37	25.6	80

experimentally observed optimal overlaps (corresponding to the italicized entries in the table) are all relatively modest.

5.2. Parallel Scalability. There are several measures of parallel scalability. Two of the most important are fixed-size scalability, in which more processors are employed in solving a problem of constant size, and fixed-memory-per-node (or “Gustafson”) scalability, in which a problem’s dimension and processor granularity are scaled in proportion. For the same algorithm, we employ a finer mesh of 128 cells radially and 512 angularly and we increase the wavenumber to $k = 13$ for this fixed-size problem. As shown in Table 2, we achieve overall efficiencies of 80% or better as the number of processors increases from 1 to 32. Convergence rate suffers remarkably mildly as preconditioner granularity increases. In this fixed-size problem, the algebraic dimension of the dense matrix block corresponding to the DtN map is fixed and is equally apportioned among the processors in a sectorial decomposition.

In the Gustafson scaling, in which the overall algebraic dimension of the problem grows in proportion to the number of processors, the communication involving all exterior boundary processors that is needed to enforce the DtN map implicitly has a deleterious effect on the scaled performance. We are presently addressing this problem by means of a sparsified approximation to the DtN map. The resulting operator is still much more accurate than a purely local Sommerfeld condition, but less crippling than the full global operator. For present purposes, we present the Gustafson scaling for a problem in which the DtN map is not included in the system matrix in (4), but split off to an outer iteration.

Results are shown in Table 3, over a range of three bi-dimensional doublings. Over one million grid points are employed in the finest case. As the problem is refined, we preserve the distribution of the spectrum by scaling with hk constant (fixed number of mesh points per wavelength). It could be argued that to keep the dominating phase truncation error term uniform, we should scale with $hk^{3/2}$ constant [11]. This would make k grow less rapidly in Table 3.

We obtain a reasonable *per iteration* efficiency, but we suffer a convergence rate degradation that is Poisson-like: iteration count grows as \sqrt{P} , where P = number of subdomains. To remedy this problem, a (relatively fine) coarse grid [3] should be used in the Schwarz preconditioner.

6. Conclusions and Future Work

We have presented a parallel algorithm of Additive Schwarz type with Sommerfeld interface conditions for the wave Helmholtz problem. The benefits of DtN

TABLE 3. Scalability for fixed local problem size using an explicit implementation of the DtN map

Number of Processors	Global Dimension	k	Iters	Time per Iteration	
				Seconds	% Increase
4	129×512	13	250	.077	—
16	257×1024	26	479	.084	9
64	513×2048	52	906	.102	32

vs. Sommerfeld conditions on the exterior boundary are illustrated by comparison with analytical solution on model problem. Parallel scalability has been evaluated and is customarily good for an additive Schwarz method for a fixed-size problem. Relatively small overlaps are sufficient. The implicit DtN map, though highly accurate, intrinsically requires communication among all exterior boundary processors and hence is non-scalable, so sparsifications are under investigation. Without a global coarse grid, algorithmic scalability deteriorates in a Poisson-like manner as the mesh and processor granularity are refined.

This work is encouraging, but not definitive for parallel Helmholtz solvers. We are interested in better preconditioners to address the underlying elliptic convergence problems, and we are interested in higher-order discretizations to increase the computational work per grid point and reduce memory requirements for the same level of accuracy. Future work will include extensions to three-dimensional problems on less smooth domains, the addition of coarse grid to preconditioner, and the embedding of a Helmholtz solver in a multiphysics (Euler/Helmholtz) application in aeroacoustics.

Acknowledgements

The authors have benefitted from many discussions with X.-C. Cai, M. Casarin, F. Elliott, and O. B. Widlund.

References

1. H. M. Atassi, *Unsteady aerodynamics of vortical flows: Early and recent developments*, Aerodynamics and Aeroacoustics (K. Y. Fung, ed.), World Scientific, 1994, pp. 119–169.
2. S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, *The Portable, Extensible Toolkit for Scientific Computing, version 2.0.21*, Tech. report, Argonne National Laboratory, 1997, <http://www.mcs.anl.gov/petsc>.
3. X.-C. Cai, M. Casarin, F. Elliott, and O. B. Widlund, *Schwarz methods for the Helmholtz problem with Sommerfeld boundary conditions*, Proceedings of the 10th International Conference on Domain Decomposition Methods (J. Mandel, C. Farhat, and X.-C. Cai, eds.), AMS, 1998.
4. X.-C. Cai and M. Sarkis, *A restricted additive Schwarz preconditioner for nonsymmetric linear systems*, Tech. Report CU-CS-843-97, Computer Science Dept., Univ. of Colorado at Boulder, August 1997, http://www.cs.colorado.edu/cai/public_html/papers/ras_v0.ps.
5. X.-C. Cai and O. B. Widlund, *Domain decomposition algorithms for indefinite elliptic problems*, SIAM J. Sci. Comput. **13** (1992), 243–258.
6. ———, *Multiplicative Schwarz algorithms for nonsymmetric and indefinite elliptic problems*, SIAM J. Numer. Anal. **30** (1993), 936–952.
7. B. Després, *Méthodes de décomposition de domaines pour les problèmes de propagation d'ondes en régime harmonique*, Tech. report, University of Paris IX, 1991.
8. S. Ghanemi, *A domain decomposition method for Helmholtz scattering problems*, Proceedings of the 9th International Conference on Domain Decomposition Methods (P. E. Bjørstad, M. Espedal, and D. E. Keyes, eds.), Wiley, 1998.

9. D. Givoli, *Non-reflecting boundary conditions*, J. Comp. Phys. **94** (1991), 1–29.
10. I. Harari and T. J. R. Hughes, *Analysis of continuous formulations underlying the computation of time-harmonic acoustics in exterior domains*, Comp. Meths. Appl. Mech. Eng. **97** (1992), 103–124.
11. F. Ihlenberg and I. Babuška, *Dispersion analysis and error estimation of Galerkin finite element methods for the Helmholtz equation*, Int. J. Numer. Meths. Eng. **38** (1995), 3745–3774.
12. J. Douglas Jr. and D. B. Meade, *Second-order transmission conditions for the Helmholtz equation*, Proceedings of the 9th International Conference on Domain Decomposition Methods (P. E. Bjørstad, M. Espedal, and D. E. Keyes, eds.), Wiley, 1998.
13. J. B. Keller and D. Givoli, *Exact non-reflecting boundary conditions*, J. Comp. Phys. **82** (1989), 172–192.
14. L. C. McInnes, R. F. Susan-Resiga, D. E. Keyes, and H. M. Atassi, *A comparison of Schwarz methods for the parallel computation of exterior Helmholtz problems*, In preparation, 1998.
15. Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput. **7** (1986), 856–869.
16. R. F. Susan-Resiga and H. M. Atassi, *A domain decomposition method for the exterior Helmholtz problem*, J. Comp. Phys., To appear, 1998.

MATHEMATICS AND COMPUTER SCIENCE DIVISION, ARGONNE NATIONAL LABORATORY, ARGONNE, IL 60639-4844

E-mail address: `curfman@mcs.anl.gov`

AEROSPACE & MECHANICAL ENGINEERING DEPARTMENT, UNIVERSITY OF NOTRE DAME, NOTRE DAME, IN 46556

E-mail address: `rsusanre@light.ame.nd.edu`

COMPUTER SCIENCE DEPARTMENT, OLD DOMINION UNIVERSITY, NORFOLK, VA 23529-0162 & ICASE, NASA LANGLEY RES. CTR., HAMPTON, VA 23681-2199

E-mail address: `keyes@cs.odu.edu`

AEROSPACE & MECHANICAL ENGINEERING DEPARTMENT, UNIVERSITY OF NOTRE DAME, NOTRE DAME, IN 46556

E-mail address: `atassi@carmen.ame.nd.edu`