

Efficient Iterative Solution of the Three-Dimensional Helmholtz Equation

Howard C. Elman¹ and Dianne P. O’Leary²

*Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland,
College Park, Maryland 20742*

E-mail: elman@cs.umd.edu and oleary@cs.umd.edu

Received August 19, 1997; revised January 27, 1998

We examine preconditioners for the discrete indefinite Helmholtz equation on a three-dimensional box-shaped domain with Sommerfeld-like boundary conditions. The preconditioners are of two types. The first is derived by discretization of a related continuous operator that differs from the original only in its boundary conditions. The second is derived by a block Toeplitz approximation to the discretized problem. The resulting preconditioning matrices allow the use of fast transform methods and differ from the discrete Helmholtz operator by an operator of low rank. We present experimental results demonstrating that when these methods are combined with Krylov subspace iteration, convergence rates depend only mildly on both the wave number and discretization mesh size. In addition, the methods display high efficiencies in an implementation on an IBM SP-2 parallel computer. © 1998 Academic Press

1. INTRODUCTION

The problem considered in this paper is to compute the numerical solution of the Helmholtz equation

$$-\Delta u - k^2 u = f. \tag{1}$$

This equation arises in numerous physical applications [9, pp. 640ff]. Here we consider a three-dimensional box-shaped domain $\Omega = (a_1, b_1) \times (a_2, b_2) \times (a_3, b_3) \subset \mathbb{R}^3$, with

¹ This work was supported by the U.S. Office of Naval Research under Grant N000149510338 and by the U.S. National Science Foundation under Grant DMS-9423133.

² This work was supported by the U.S. Office of Naval Research under Grant N000149510338 and by a General Research Board Grant from the Office of Graduate Studies and Research of the University of Maryland.

Sommerfeld-like boundary conditions

$$u_n - iku = 0 \tag{2}$$

on $\partial\Omega$, which constitute an approximation to the Sommerfeld radiation condition

$$\lim_{r \rightarrow \infty} r(u_n - iku) = 0, \tag{3}$$

used in models of acoustic scattering [17].

Discretization of the problem (1)–(2) results in a linear system of equations

$$Au = f. \tag{4}$$

Since the problem is fully three-dimensional, any reasonable discretization will contain a large number of unknowns and require considerable storage. Direct methods based on Gaussian elimination with partial pivoting require a prohibitive amount of additional storage and thus have limited use. Multilevel methods suffer from the requirement that coarse spaces used must be fine enough to accurately represent the solution; see e.g., [5, 7, 19, 20, 21]. In addition, the complex symmetric coefficient matrix A typically has eigenvalues with both positive and negative real parts. This can cause difficulties for iterative solution methods, and preconditioning of the matrix is essential in order to attain efficiency.

In this paper, we propose solving the discrete Helmholtz equation using Krylov subspace iterative methods with a preconditioning methodology derived from fast direct methods. The basic principle behind fast direct solvers is to apply an inexpensive transformation to break a problem into a number of lower-dimensional but independent problems. Many solvers use fast Fourier transforms (FFTs) to achieve separation of variables and then solve the resulting set of decoupled problems using sparse matrix methods. Fast direct methods are standard tools for solving the Poisson equation on regular domains with Dirichlet, Neumann, or periodic boundary conditions [6]; they can be adapted to other domains via a capacitance matrix or embedding methods [14, 24]. They have been used for the three-dimensional Helmholtz equation with Dirichlet or Neumann boundary conditions on an irregular domain [23], and for the two-dimensional problem in polar coordinates with nonreflecting boundary conditions [18] (derived from a Dirichlet-to-Neumann mapping) in [11, 15]. In this work, we develop efficient solvers for problems with Sommerfeld-like boundary conditions on box-shaped domains. Combining our techniques with capacitance matrix methods would produce solvers for general geometries in Cartesian coordinates, including exterior problems.

Our idea generalizes some results developed for two-dimensional Helmholtz problems by Ernst and Golub [12]. (See also [8, Sec. 4] for variants applied to definite elliptic problems.) We approximate the discrete operator A with a matrix Q that can be treated with fast direct methods. For finite difference discretizations, we derive Q by defining and discretizing the differential operator in the same way as for A except that the boundary conditions on either two or four faces of Ω are replaced by more convenient ones (Dirichlet or Neumann). The resulting matrix Q differs from A by a (relatively) low-rank operator and can be used as a preconditioner for A , to accelerate the convergence of iterative solvers

based on Krylov subspaces. We also develop variants of these ideas for finite element discretizations (on uniform grids), focusing on trilinear elements. Here, rather than explicitly modifying the boundary conditions to construct Q , we use the fact that the discrete operator A is close to a block Toeplitz matrix and replace certain sub-blocks of A by Toeplitz approximations that are amenable to fast transforms. For both types of discretizations, we will demonstrate empirically that Q meets the requirements for an effective preconditioner:

- Applying the action of Q^{-1} to a vector is not too expensive. For our preconditioners, using Q^{-1} entails a set of FFTs together with solution of smaller dimensional problems (see Section 2).
- Q greatly reduces the number of iterations needed by Krylov subspace methods to solve (4).

In particular, we will show that for several choices of Q , the experimental convergence behavior of preconditioned restarted GMRES [25] depends only mildly on both the wave number k and the discretization mesh size. In addition, we will demonstrate how the methods can be implemented on a parallel computer with high efficiency.

The paper is organized as follows. In Section 2, we show how fast transform methods can be used to generate preconditioning operators for finite difference discretizations of the Helmholtz equation and we develop variants applicable to low-order finite element discretizations, using trilinear elements as a specific example. In Section 3, we present the results of a series of numerical experiments demonstrating the performance of the preconditioners. In most cases, there is virtually no increase in iteration counts as the mesh size is refined for fixed wave numbers; especially for finite differences, there is only slight dependence on the wave numbers. In addition, we show that the new methods are more effective than a standard algebraic preconditioner based on symmetric successive overrelaxation (SSOR) [30]. In Section 4, we show how the methods can be implemented on parallel computers, and we demonstrate their parallel efficiency using experiments on a sixteen processor IBM SP-2. Finally, in Section 5, we make some concluding remarks.

2. THE PRECONDITIONERS

Good preconditioners for Krylov subspace iterations can be determined in two ways:

- preconditioners derived from operators related to the desired operator.
- preconditioners derived from matrices related to the desired matrix.

We will use both approaches in our work. For simplicity, we restrict our attention to $\Omega = (0, 1)^3$, the unit box.

2.1. Preconditioners for Finite Difference Discretizations. Given positive integers m_x , m_y , and m_z , let problem (1)–(2) be discretized by the seven-point (second order accurate) finite difference operator on a uniform mesh with cells of size $h_x \times h_y \times h_z$, where $h_x = 1/(m_x + 1)$, $h_y = 1/(m_y + 1)$, $h_z = 1/(m_z + 1)$. Assume that the normal derivatives in the boundary conditions are approximated by one-sided differences and that the discrete boundary conditions are used to eliminate all unknowns on $\partial\Omega$. The resulting “stencil” at interior grid points is

where

$$P = I_{m_z} \otimes T_y^{(m_y)} + T_z^{(m_z)} \otimes I_{m_y} - k^2 I_{m_z m_y}.$$

The eigenvectors of T_{m_x} are the columns of the orthogonal matrix U_s corresponding to a discrete sine series representation. Therefore, the product $v = U_s^T w$ can be formed by computing the discrete Fourier sine transform of the vector w , and $w = U_s v$ is the inverse transform of the vector v . If D_d is the diagonal matrix of eigenvalues of T_{m_x} , then Q_d can be represented as

$$Q_d = (I_{m_z m_y} \otimes U_s) (I_{m_z m_y} \otimes D_d + P \otimes I_{m_x}) (I_{m_z m_y} \otimes U_s^T). \quad (6)$$

Reordering the rows and columns of the matrix $I_{m_z m_y} \otimes D_d + P \otimes I_{m_x}$ in (6) by lines within each y - z plane yields a matrix \mathcal{P} containing m_x diagonal blocks, each with the same nonzero structure as a five-point finite difference discretization of a two-dimensional problem.

Using Q_d as a preconditioner entails applying the action of Q_d^{-1} to a vector v at each step of an iterative algorithm. The discussion above shows that this computation can be performed in the following sequence of steps:

- (1) Perform $m_x m_y$ sine transform operations in the x coordinate directions to compute $v_1 = (I_{m_z m_y} \otimes U_s^T) v$.
- (2) Solve m_x two-dimensional problems, one in each y - z plane, to compute v_2 .
- (3) Perform $m_x m_y$ inverse sine transform operations in the x coordinate directions to compute $w = (I_{m_z m_y} \otimes U_s) v_2$.

The solution of the two-dimensional problems required in Step (2) can be done using a variety of techniques, including general sparse direct methods [10], band solvers, and domain decomposition.

Consider a variant of this preconditioner, derived using Dirichlet boundary conditions at two pairs of opposite faces of $\partial\Omega$: $x = 0, 1$ and $y = 0, 1$. The resulting preconditioning matrix is

$$Q_{dd} = I_{m_z} \otimes I_{m_y} \otimes T_{m_x} + I_{m_z} \otimes T_{m_y} \otimes I_{m_x} + T_{m_z}^{(z)} \otimes I_{m_y} \otimes I_{m_x} - k^2 I_{m_x m_y m_z}. \quad (7)$$

A set of sine transforms in the x -direction still decouples the problem into m_x two-dimensional subproblems as in (6), but now if this is followed by a set of sine transforms in the y -direction, the result is $m_x m_y$ independent one-dimensional (tridiagonal) problems. Step (2) of the computation of the action of Q_{dd}^{-1} then has the following form:

- (2a) Perform $m_z m_x$ sine transform operations in the y coordinate directions.
- (2b) Solve $m_x m_y$ tridiagonal systems in the z coordinate directions.
- (2c) Perform $m_z m_x$ inverse sine transform operations in the y coordinate directions.

We will consider four preconditioning operators of this type:

- Q_d derived from Dirichlet boundary conditions at $x = 0, 1$.
- Q_{dd} derived from Dirichlet boundary conditions at $x = 0, 1$ $y = 0, 1$.
- Q_n derived from homogeneous Neumann boundary conditions at $x = 0, 1$.

The fast solver here involves discrete cosine transforms and solution of m_x two-dimensional problems.

- Q_{nn} derived from homogeneous Neumann boundary conditions at $x = 0, 1, y = 0, 1$. The fast solver involves discrete cosine transforms in two directions plus solution of $m_x m_y$ one-dimensional (tridiagonal) problems.

If m_x and m_y are powers of two, then the time required to compute $Q_{dd}^{-1}v$ or $Q_{nn}^{-1}v$ is proportional to $m_x m_y m_z (\log m_x + \log m_y + 1)$. For $Q_d^{-1}v$ or $Q_n^{-1}v$, the time is proportional to $m_x m_y m_z \log m_x$ plus the time to solve the two-dimensional problems.

Clearly, it is possible to derive other variants of these ideas based on other boundary conditions. As long as the Sommerfeld-like boundary conditions are retained in at least one coordinate direction, the preconditioning operators that use any combination of Dirichlet and Neumann boundary conditions are all nonsingular for any value of k .

2.2. Preconditioners for a Finite Element Discretization. The weak formulation of the Helmholtz equation (1)–(2) is to find $u \in \mathcal{S}$ such that

$$a(u, v) = (f, v) \quad \text{for all } v \in \mathcal{S},$$

where

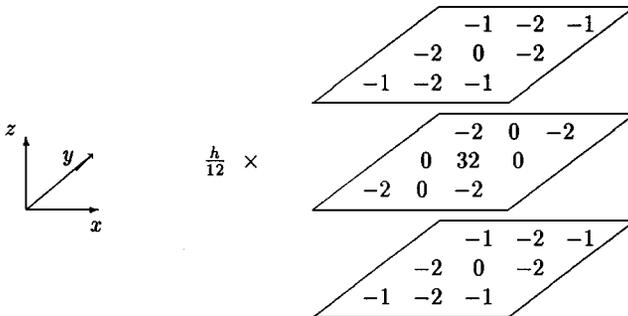
$$\begin{aligned}
 a(u, v) &= \int_{\Omega} (\nabla u \cdot \nabla \bar{v} - k^2 u \bar{v}) - ik \int_{\partial\Omega} u \bar{v} \\
 (f, v) &= \int_{\Omega} f \bar{v}
 \end{aligned}
 \tag{8}$$

and \mathcal{S} is an appropriate Sobolev space (often, $H^1(\Omega)$), depending on f . Note that the boundary conditions (2) are explicitly incorporated into the weak form.

Let \mathcal{S}_h denote the finite dimensional subspace of \mathcal{S} determined from continuous piecewise trilinear basis functions on rectangular elements. The discrete weak form is to find $u_h \in \mathcal{S}_h$ such that

$$a(u_h, v) = (f, v) \quad \text{for all } v \in \mathcal{S}_h.$$

Assuming uniform cells of size $h \times h \times h$ and using a natural ordering of unknowns, the resulting coefficient matrix A again has block tridiagonal structure. (Different values of h can be used in each coordinate direction; we restrict our attention to cubic cells only to simplify the notation.) We describe it using its stencil at interior mesh points, which consists of the contribution from the stiffness matrix (from $(\nabla u \cdot \nabla v)$),



together with the contribution from the mass matrix (from $k^2(u, v)$),

TABLE 1

Iteration Counts for GMRES (20) with Preconditioners That Combine One-Dimensional Transforms with Two-Dimensional Sparse Direct Solves: Finite Differences with Non-Smooth Solution

k	Q_d (sine + 2D solves)				k	Q_n (cosine + 2D solves)			
	m					m			
	20	40	60	80	20	40	60	80	
1	11	11	12	—	1	4	3	2	—
5	9	10	11	—	5	7	6	5	—
10	9	11	11	—	10	10	8	8	—
20	12	12	12	—	20	15	14	12	—
30	13	14	14	—	30	20	18	18	—
40	12	18	17	—	40	31	21	20	—
50	16	19	24	—	50	44	26	27	—

This ensures that the norm minimized by GMRES is independent of the choice of the preconditioner. The iteration is stopped when

$$\frac{\|r_j\|_2}{\|f\|_2} < 10^{-5},$$

where $r_j = f - Au_j$ is the residual for the iterate u_j , $u_0 \equiv 0$, and the norm is the usual vector Euclidean norm. All computations with these preconditioners were performed on multiprocessors of an IBM SP-2 computer in double precision. (See Section 4 for details.)

We use a set of uniform three-dimensional grids of size $m \times m \times m$ for $20 \leq m \leq 80$, together with a variety of wave numbers k . For any k , accurate discrete solutions of (1) will be obtained only if the mesh is fine enough to resolve the features of the problem. A commonly used criterion is for the mesh to include at least ten grid points per wave, i.e., to require $k \leq 2\pi/(10h)$. See [16] for rigorous justification of this criterion for one-dimensional problems, instead of the more stringent requirement that $k^3 h^2$ be bounded. In the tabulated data shown below, results for problems with at least ten grid points per wave lie above the jagged line; data lying below these lines are included only to show trends and do not correspond to physically meaningful computations. Dashed lines (—) correspond to problems that are too large for practical computation on this configuration.

Tables 1 and 2 show results for non-smooth problems and finite difference discretizations. Table 1 shows that iteration counts required by GMRES (20) with the preconditioners that entail a set of trigonometric transforms together with direct solution of two-dimensional subproblems. The entries on the left are for the preconditioner Q_d , which uses sine transforms, and the entries on the right are for Q_n , which uses cosine transforms. We will refer to these here as the “two-dimensional” solvers. Table 2 shows analogous results for Q_{dd} and Q_{nn} , where the two-dimensional subproblems are also treated using fast transforms; we will refer to these (less costly) preconditioners as the “one-dimensional” solvers. Table 3 shows analogous results for smooth problems. For the sake of brevity, here we only discuss the one-dimensional preconditioners Q_{dd} and Q_{nn} .

The following trends are evident in these tables:

TABLE 2

Iteration Counts for GMRES (20) with Preconditioners That Combine Two Sets of One-Dimensional Transforms with One-Dimensional Tridiagonal Solves: Finite Difference with Non-Smooth Solution

Q_{dd} (sine + 1D solves)					Q_{nn} (cosine + 1D solves)				
m					m				
k	20	40	60	80	k	20	40	60	80
1	14	15	13	12	1	5	4	3	3
5	12	13	14	14	5	10	9	8	8
10	13	14	15	15	10	19	17	16	15
20	20	25	21	21	20	50	41	38	35
30	36	30	29	28	30	87	76	77	69
40	34	53	47	45	40	133	96	106	95
50	46	76	69	69	50	264	142	174	165

(1) For non-smooth problems, iteration counts are insensitive to mesh size. Indeed, the counts often decrease as h decreases. Performance is much better than would be expected from the rank ($m^2 - m$) of the difference between A and the preconditioner Q .

(2) Iteration counts for the methods based on the sine transform increase very modestly with the wave number k ; counts for the cosine transform are more sensitive to k but they are smaller when k is small.

(3) Fewer iterations are required with two-dimensional solvers (for which the boundary conditions determining Q are more like those determining A) than with the one-dimensional solvers. As we will see in Section 4, however, this is at the expense of significant extra work.

(4) The counts for smooth problems are somewhat higher than in the non-smooth case. The qualitative dependence of performance on wave number is largely the same, as is the

TABLE 3

Iteration Counts for GMRES (20) with Preconditioners That Combine Two Sets of One-Dimensional Transforms with One-Dimensional Tridiagonal Solves: Finite Differences with Smooth Solution

Q_{dd} (sine + 1D solves)					Q_{nn} (cosine + 1D solves)				
m					m				
k	20	40	60	80	k	20	40	60	80
1	16	23	36	42	1	6	6	6	6
5	13	18	23	30	5	11	11	11	11
10	14	20	26	30	10	17	18	19	18
20	18	22	28	32	20	51	57	57	58
30	26	37	36	40	30	78	94	101	101
40	29	55	55	53	40	101	118	135	137
50	24	75	69	81	50	155	171	196	206

TABLE 4

Iteration Counts for GMRES (20) with Preconditioners That Combine One-Dimensional Transforms with Two-Dimensional Sparse Direct Solves: Trilinear Finite Elements with Non-Smooth Solution

Q_d (sine + 2D solves)					Q_n (cosine + 2D solves)				
m					m				
k	20	40	60	80	k	20	40	60	80
1	24	25	26	—	1	23	22	22	—
5	32	29	28	—	5	32	30	29	—
10	48	37	34	—	10	51	41	39	—
20	87	65	58	—	20	85	67	59	—
30	300*	85	79	—	30	300*	104	89	—
40	300*	146	89	—	40	300*	245	127	—
50	300*	300*	158	—	50	300*	300*	300*	—

dependence on mesh size, except in the case of small wave numbers with sine transforms; here the iteration counts appear to grow roughly linearly with h^{-1} .

Tables 4–6 show the iteration counts for the trilinear finite element discretization. Here we use \hat{Q}_{dd} and \hat{Q}_{nn} for the one-dimensional solvers; these were more effective than the variants Q_{dd} and Q_{nn} . Numbers with an asterisk correspond to the maximum number of iterations permitted. Many of the trends are the same as for finite differences: iteration counts for non-smooth problems are essentially independent of the mesh size; the sine-based methods are generally more effective than the cosine-based methods; and, for smooth problems, costs and dependence of the sine-based method on mesh size are somewhat greater. The sensitivity to wave number is somewhat more pronounced than for finite differences, although for the larger values of k considered, the iteration counts decrease as the mesh is refined. This

TABLE 5

Iteration Counts for GMRES (20) with Preconditioners That Combine Two Sets of One-Dimensional Transforms with One-Dimensional Tridiagonal Solves: Trilinear Finite Elements with Non-Smooth Solution

\hat{Q}_{dd} (sine + 1D solves)					\hat{Q}_{nn} (cosine + 1D solves)				
m					m				
k	20	40	60	80	k	20	40	60	80
1	14	15	15	14	1	12	12	11	11
5	14	14	14	14	5	21	19	18	18
10	19	17	17	17	10	43	40	39	38
20	51	35	31	30	20	156	138	123	112
30	164	73	51	41	30	300*	253	214	198
40	300*	135	97	77	40	300*	300*	300*	300*
50	300*	231	146	140	50	300*	300*	300*	300*

TABLE 6

Iteration Counts for GMRES (20) with Preconditioners That Combine Two Sets of One-Dimensional Transforms with One-Dimensional Tridiagonal Solves: Trilinear Finite Elements with Smooth Solution

\hat{Q}_{dd} (sine + 1D solves)					\hat{Q}_{nm} (cosine + 1D solves)				
m					m				
k	20	40	60	80	k	20	40	60	80
1	15	19	23	32	1	12	12	12	12
5	14	17	19	23	5	18	20	20	19
10	17	20	24	27	10	35	35	34	36
20	34	32	31	33	20	115	112	101	97
30	108	56	46	46	30	206	180	167	168
40	189	92	71	73	40	246	300*	300*	300*
50	272	150	115	123	50	300*	300*	300*	300*

suggests that for these wave numbers the asymptotic behavior (as $h \rightarrow 0$) of the solvers is being approached only for the finest meshes considered here. In these tests, performance is less dependent on the smoothness of the solution than in the experiments with finite differences.

In a few tests without preconditioning, GMRES (20) required an average of eight times more steps than with the Q_{dd} preconditioner for non-smooth problems, and at least fifteen times more steps for smooth problems. (These tests were for finite differences, $n = 20$ and 40 and $k \leq 20$; a maximum of 300 steps was used, and for the smooth problems the stopping criterion was not satisfied in a majority of the runs.)

To compare the performance of the new preconditioners with a "standard" algebraic preconditioner, we show in Tables 7 and 8 iteration counts obtained using the SSOR preconditioner [30] (with $\omega = 1$) with GMRES (20). This method has been used to good effect for solving the Helmholtz equation with multiple right hand sides (derived, e.g., from incident waves at different angles) in [13], cf. also [3]. These tests were run on a uniprocessor Sun SPARCstation 20 in Matlab, and we considered only mesh sizes less than or equal to $m = 60$. (Storage limitations permitted only $m \leq 50$ for trilinear elements.) In almost all

TABLE 7

Iteration Counts of GMRES (20) with SSOR ($\omega = 1$) Preconditioning: Non-Smooth Solutions

Finite differences				Trilinear elements			
m				m			
k	20	40	60	k	20	40	50
1	26	19	20	1	20	19	19
5	26	28	24	5	17	27	26
10	37	49	50	10	27	38	42
20	76	78	87	20	80	61	66
30	200	103	103	30	300*	90	85

TABLE 8

Iteration Counts of GMRES (20) with SSOR ($\omega = 1$) Preconditioning: Smooth Solutions

Finite differences				Trilinear elements			
m				m			
k	20	40	60	k	20	40	50
1	52	200	300*	1	26	99	122
5	40	122	209	5	24	62	97
10	57	170	249	10	37	78	113
20	110	159	219	20	96	99	117
30	237	175	222	30	300*	122	127

cases, these iteration counts are larger than the analogous entries from Tables 1–6, and the growth in iteration counts with increasing wave numbers is considerably more pronounced. The differences are especially dramatic for smooth problems. (We also remark, however, that the SSOR method was surprisingly insensitive to mesh size in the non-smooth case. An analysis of behavior like this for a simple elliptic model problem is given in [22].) In addition, these tests were run with a “natural” ordering of the unknowns, which does not have efficient parallel implementation. Tests with point red-black ordering for the finite difference example (for $n \leq 40$, $k \leq 20$, and non-smooth problems) required on average 65% more steps.

4. PARALLEL IMPLEMENTATION AND PERFORMANCE

The algorithms developed in Section 2 are well adapted for parallel computation in both shared memory and message-passing environments. We have implemented them on an IBM SP-2 computer assuming that the number of processors does not exceed m_x or m_z . In this section, we describe the implementation and present the results of experiments showing parallel performance.

4.1. *Parallel implementation.* For ease of illustration, we describe our implementation under the assumption that there are $p = 4$ processors. The mesh is partitioned among processors as in Fig. 1; if m_z is divisible by p then each processor contains m_z/p blocks of

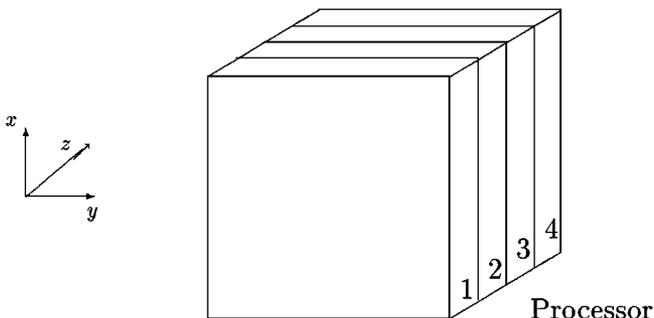


FIG. 1. Partitioning of the three-dimensional mesh among 4 processors.

$m_x m_y$ grid points oriented along the x - y plane. If m_z is not divisible by p , then the number of x - y planes assigned to each processor differs by at most one.

4.1.1. *Parallel implementation of the iterative method.* We use the original form of the restarted GMRES algorithm, based on the Arnoldi basis, as presented in [25]. With the partitioning described above, the GMRES algorithm is easy to parallelize: each processor is responsible for storing and updating at most

$$\hat{m}_z \equiv \left\lceil \frac{m_z}{p} \right\rceil m_x m_y$$

unknowns and for maintaining the factored form of a Hessenberg matrix of size $r \times r$ (using r steps of GMRES between restarts). Communication is necessary only for inner products and matrix-vector products. The cost per iteration is $O(\hat{m}_z)$ floating point operation plus accumulation of (on average) $(r + 3)/2$ sums of scalars across processors, plus the cost of matrix-vector products and preconditioning.

4.1.2. *Parallel implementation of matrix-vector products.* Computation of products of the finite difference or finite element matrix with a vector requires that each processor send its highest-numbered x - y plane to the processor numbered one greater, and its lowest-numbered x - y plane to the processor numbered one less (if these processors exist). The stencil can then be applied to the local data. The cost per matrix multiply is $O(\hat{m}_z)$ floating point operations plus 2 sends and receives of $m_x m_y$ numbers per processor.

4.1.3. *Parallel implementation of the preconditioning.* The preconditioning computation is a somewhat more complex operation. Let

$$\hat{m}_x \equiv \left\lceil \frac{m_x}{p} \right\rceil m_y m_z.$$

We arrange the work in the chart below, estimating the cost assuming that m_x and m_y are powers of 2. Note that the costs of \hat{Q}_{dd} and \hat{Q}_{nm} are identical to those of Q_{dd} .

Operation	Cost
Each processor computes sine or cosine transforms in the x direction on its local data.	$O(\hat{m}_z \log m_x)$ operations.
The data are rearranged so that each processor has an approximately equal number of y - z planes. This requires a nontrivial amount of communication: the data movement from the perspective of Processor 2 is shown in Fig. 2.	Each processor sends at most $\lceil \frac{m_z}{p} \rceil \lceil \frac{m_x}{p} \rceil m_y$ numbers to every other processor.
Each processor then solves its assigned two-dimensional problems. For preconditioners Q_{dd} and Q_{nm} , this entails sine or cosine transforms in the y direction followed by solution of tridiagonal systems, followed by inverse sine or cosine transforms. For Q_d and Q_n , a two-dimensional problem is solved in each of the y - z planes.	For Q_d and Q_n , the cost is the solution of $\lceil \frac{m_x}{p} \rceil$ problems of size $m_y m_z$. For Q_{dd} and Q_{nm} the cost is $O(\hat{m}_x (\log m_y + 1))$ operations.
The data are rearranged to the original configuration.	Each processor sends at most $\lceil \frac{m_z}{p} \rceil \lceil \frac{m_x}{p} \rceil m_y$ numbers to every other processor.
Each processor computes inverse sine or cosine transforms in the x direction.	$O(\hat{m}_z \log m_x)$ operations.

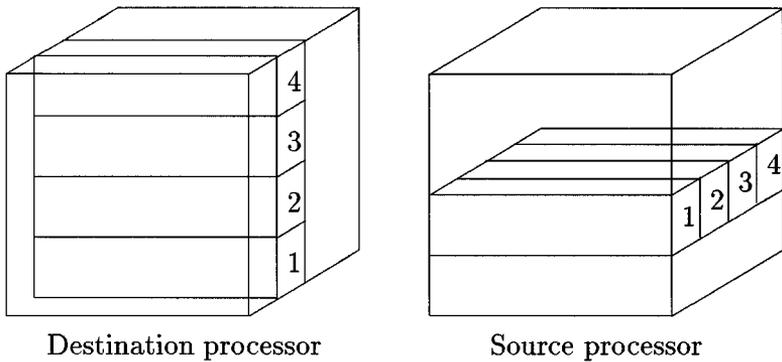


FIG. 2. Destination and source processors for data movement, from the perspective of Processor 2.

Remark 4.1. For $m_x = m_y = m_z = m$, the total arithmetic cost for Q_{dd} or Q_{nn} is proportional to $\frac{m^3 \log m}{p}$, and the communication cost (assuming no contention for messages sent simultaneously) is proportional to $\frac{m^3}{p}$. The communication cost is the same for the two-dimensional direct solvers Q_d or Q_n , and the smallest possible arithmetic cost is $O(\frac{m^4}{p})$ if a nested dissection method is used [10]; this strategy would neglect any need for pivoting. For the tests described below, we used a bandsolver (which allows pivoting), resulting in cost proportional to $\frac{m^5}{p}$.

Remark 4.2. For Q_{dd} and Q_{nn} , data movement can be completely masked by computation of the sine or cosine transforms in the y direction, provided communication speed is not too slow and the hardware supports overlapping of communication and computation. This option is not currently supported on the SP-2.

4.2. Parallel performance. We now describe the performance of the solvers using the transform-based preconditioners on a sixteen processor IBM SP-2 computer, a distributed memory machine with explicit message passing. The system contains sixteen RS6000/390 processing nodes running AIX, interconnected via a proprietary interprocessor communications switch. The computational component of the program was written in Fortran90 and compiled using the mpxlf90 compiler with the optimization (-O) switch. All modules were taken from off-the-shelf software and modified to conform to Fortran90: GMRES is derived from the TEMPLATES package [2]; the two-dimensional direct solvers use the LAPACK [1] bandsolver `cgbsv` and `cgbtrs`, as do the tridiagonal solvers used for the one-dimensional preconditioners; and the sine-transform and cosine-transform routines are from FFTPACK [28]. All tests used double precision complex floating point computations. Communication was performed using MPI [26] with nonblocking sends (MPI_ISEND) and blocking receives (MPI_RECV). Inner products were performed and broadcast using MPI_ALLREDUCE.

The results on parallel performance are summarized in Table 9 for finite differences and Table 10 for trilinear finite elements. Again, dashed lines correspond to problems that were too large. The entries show CPU times for solving the discrete problem with non-smooth solution on various grid sizes with the sine-based preconditioners. For an understanding of parallel performance it suffices to look at just one wave number, which in these cases was $k = 5$; results for the cosine-based preconditioners also lead to the same conclusions. We

TABLE 9

CPU Times for Solving the Finite Difference Discretization with $k = 5$, for Various Grid Sizes

Number of processors	Q_d (sine + 2D solves)			Number of processors	Q_{dd} (sine + 1D solves)		
	m				m		
	16	32	64		16	32	64
1	1.89	25.37	—	1	1.62	13.12	—
2	.98	12.75	—	2	.85	6.91	—
4	.57	6.72	—	4	.47	3.74	28.24
8	.43	3.56	—	8	.40	2.15	14.55
16	.52	2.14	27.72	16	.52	1.37	8.14

used grid sizes that are powers of two only for convenience so that the grid parameter m_z divides the number of processors; results when this is not the case are similar. The timings reflect the averages over three runs; these runs were made in time sharing mode and other users had access to the machine.

It is evident from these data that all the methods display a large amount of parallelism. The average speedups in going from p to $2p$ processors for $m = 32$ is 1.77 for the one-dimensional solvers and 1.90 for the two-dimensional solvers; similar efficiencies are observed for $m = 64$ and $m = 16$ when p is small. The higher parallel efficiencies for the two-dimensional solvers reflect the larger ratio of arithmetic to communication for these preconditioners. Table 11 shows the total speedup, i.e., ratio of CPU time on one processor to CPU time on p processors, for $m = 32$. The typically larger speedups observed for the finite element problems also stem from the larger amount of arithmetic (caused by denser coefficient matrices) for these problems. There is some degradation of performance, especially for the one-dimensional solvers, when the number of processors increases. We attribute this to the relatively large number of synchronization points required by the modified Gram–Schmidt computation in each iteration of the GMRES algorithm, and to the decreased amount of computation performed on each processor. On a very large number of processors, performance can be improved using a Krylov subspace method that avoids such synchronizations, e.g., GMRES with the classical Gram–Schmidt

TABLE 10

CPU Times for Solving the Trilinear Finite Element Discretization with $k = 5$, for Various Grid Sizes

Number of processors	Q_d (sine + 2D solves)			Number of processors	Q_{dd} (sine + 1D solves)		
	m				m		
	16	32	64		16	32	64
1	9.20	91.56	—	1	3.41	24.82	—
2	4.38	45.81	—	2	1.71	12.70	—
4	2.09	24.41	—	4	.80	7.03	49.34
8	1.39	12.23	—	8	.65	3.71	25.10
16	1.20	6.67	70.59	16	.80	2.59	15.21

TABLE 11
Speedups for $\hat{m} = 32$

Number of processors	Fin. diff. Q_d	Fin. diff. Q_{dd}	Fin. elem. Q_d	Fin. elem. \hat{Q}_{dd}
2	1.99	1.90	2.00	1.95
4	3.78	3.51	3.91	3.53
8	7.13	6.10	7.49	6.69
16	11.86	9.58	13.73	9.58

orthogonalization, or methods such as quasi-minum residual or BiCGSTAB that do not require orthogonalization against a growing collection of old vectors.³ For the more compute-intensive two-dimensional solvers, there is little performance degradation. Despite this, overall costs of the one-dimensional solvers were significantly lower, even when these required more iterations (i.e., for finite differences).

Remark 4.3. In repeated examples of groups of three runs, we found some variation in the average CPU times, on the order of 10%. This explains some instances (e.g., finite elements, $Q_d, m = 16$) where doubling the processors led to speedups greater than two. These variations may derive from contention for the communication switch, although we also encountered some nontrivial variation on sixteen processors, when we occupied the whole machine.

5. CONCLUDING REMARKS

The preconditioners presented here enable efficient parallel solution of the three-dimensional Helmholtz equation on large uniform grids. Performance of restarted GMRES with these preconditioners is relatively insensitive to discretization mesh size and wave number, and the algorithms are highly parallelizable.

We have some limited computational experience using QMR, a lower-storage alternative to GMRES, on this set of test problems and preconditioners and find that the convergence rate is similar.

We have presented results of these methods on simple box-shaped domains. A problem with Sommerfeld boundary conditions on only a portion of the boundary (arising, e.g., from a nozzle configuration) would be somewhat easier to handle. The methods can also be adapted for use on exterior Helmholtz problems by applying the capacitance method of [23], perhaps using a nonuniform grid in the neighborhood of the scatterer. In forthcoming work we will present the results of applying the algorithms to problems with varying k , corresponding to an inhomogeneous medium. Our results indicate that the methods are practical for such problems, too.

The techniques are potentially applicable in non-Cartesian coordinate systems, although the “fast” solvers in such settings are not nearly as fast, relying on generation and solution of general block tridiagonal systems [27]. We also expect them to perform well for more accurate local approximations to the radiation boundary conditions (3), of the type considered in [4].

³ For example, in a few tests with the classical Gram–Schmidt orthogonalization and finite differences, we found improved speedups of 13.7 and 10.3 for Q_d and Q_{dd} , respectively, on sixteen processors; cf. Table 11.

ACKNOWLEDGMENTS

We thank Alan Sussman for a great deal of advice on using the IBM-SP2, Ilya Zavorine for help with programming, and Olof Widlund for some helpful discussions.

REFERENCES

1. E. Anderson *et al.*, *LAPACK Users' Guide* (SIAM, Philadelphia, 1995), 2nd ed.
2. R. Barrett *et al.*, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods* (SIAM, Philadelphia, 1993).
3. A. Bayliss, C. I. Goldstein, and E. Turkel, An iterative method for the Helmholtz equation, *J. Comput. Phys.* **49**, 443 (1983).
4. A. Bayliss, M. Gunzburger, and E. Turkel, Boundary conditions for the numerical solution of elliptic equations in exterior domains, *SIAM J. Appl. Math.* **42**, 430 (1982).
5. J. H. Bramble, J. E. Pasciak, and J. Xu, The analysis of multigrid algorithms for nonsymmetric and indefinite problems, *Math. Comp.* **51**, 389 (1988).
6. B. L. Buzbee, F. W. Dorr, J. A. George, and G. H. Golub, The direct solution of the discrete Poisson equation on irregular regions, *SIAM J. Numer. Anal.* **8**, 722 (1971).
7. X.-C. Cai and O. B. Widlund, Domain decomposition algorithms for indefinite elliptic problems, *SIAM J. Sci. Stat. Comput.* **13**, 243 (1992).
8. R. H. Chan and M. K. Ng, Conjugate gradient methods for Toeplitz systems, *SIAM Rev.* **38**, 427 (1996).
9. R. Dautray and J.-L. Lions, *Mathematical Analysis and Numerical Methods for Science and Technology* (Springer-Verlag, New York, 1990), Vol. 1.
10. I. S. Duff, A. M. Erisman, and J. K. Reid, *Direct Methods for Sparse Matrices* (Clarendon, Oxford, 1986).
11. O. Ernst, *Fast Numerical Solution of Exterior Helmholtz Problems with Radiation Boundary Condition by Imbedding*, Ph.D. thesis, Stanford University, Program in Scientific Computing and Computational Mathematics, 1994.
12. O. Ernst and G. H. Golub, A domain decomposition approach to solving the Helmholtz equation with a radiation boundary condition, in *Domain Decomposition in Science and Engineering*, edited by A. Quarteroni, J. Periaux, Y. Kuznetsov, and O. Widlund (Amer. Math. Soc. Providence, 1994), pp. 177–192.
13. R. W. Freund and M. Malhotra, A block-QMR Algorithm for non-Hermitian linear systems with multiple right hand sides, *Linear Algebra Appl.* **254**, 197 (1997).
14. J. A. George, *The Use of Direct Methods for the Solution of the Discrete Poisson Equation on Non-rectangular Regions*, Tech. Rep. STAN-CS-70-159, Computer Science Department, Stanford University, Stanford, CA 1970.
15. E. Heikkola, T. Rossi, P. Tarvainen, and Y. Kuznetsov, *Efficient Preconditioners Based on Fictitious Domains for Elliptic fe-Problem with Lagrange Multipliers*, Tech. Rep. 11/1996, University of Jyväskylä, 1996.
16. F. Ihlenburg and I. Babuška, Finite element solution of the Helmholtz equation with high wave number. Part I. The h-version of the FEM, *Comput. Math. Appl.* **30**, 9 (1995).
17. M. C. Junger and D. Feit, *Sound, Structures and Their Interaction* (MIT Press, Cambridge, MA, 1986), 2nd ed.
18. J. B. Keller and D. Givoli, Exact nonreflecting boundary conditions, *J. Comput. Phys.* **82**, 172 (1989).
19. S. Kim, A parallizable iterative procedure for the Helmholtz problem, *Appl. Numer. Math.* **14**, 435 (1994).
20. S. Kim, Parallel multidomain iterative algorithms for the Helmholtz wave equation, *Appl. Numer. Math.* **17**, 411 (1995).
21. M. Malhotra and P. M. Pinsky, Parallel preconditioning based on h-hierarchical finite elements with applications to acoustics, *Comput. Methods Appl. Mech. Engrg.* **155**, 97 (1998).
22. A. E. Naiman, I. M. Babuška, and H. C. Elman, A note on conjugate gradient convergence, *Numer. Math.* **76**, 209 (1997).
23. D. P. O'Leary and O. Widlund, Capacitance matrix methods for the Helmholtz equation on general three dimensional regions, *Math. Comput.* **33**, 849 (1979).

24. W. Proskurowski and O. Widlund, On the numerical solution of Helmholtz's equation by the capacitance matrix method, *Math. Comput.* **30**, 433 (1976).
25. Y. Saad and M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* **7**, 856 (1986).
26. M. Snir, S. W. Otto, S. Huss-Lederman, D. W. Walker, and J. Dongarra, *MPI: The Complete Reference* (MIT Press, Cambridge, MA, 1996).
27. P. Swarztrauber and R. Sweet, *Efficient Fortran Subprograms for the Solution of Elliptic Equations*, Tech. Rep. NCAR TN/IA-109, National Center for Atmospheric Research, Boulder, CO, 1975; code available through <http://www.netlib.org>.
28. P. N. Swarztrauber, Vectorizing the fft's, in *Parallel Computations*, edited by G. Rodrigue (Academic Press, New York, 1982), pp. 51–83.
29. C. Van Loan, *Computational Frameworks for the Fast Fourier Transform* (SIAM, Philadelphia, 1992).
30. D. M. Young, *Iterative Solution of Large Linear Systems* (Academic Press, New York, 1970).