

**PARALLEL SOLUTION OF HIGH-FREQUENCY HELMHOLTZ EQUATIONS
USING HIGH-ORDER FINITE DIFFERENCE SCHEMES**

Dan Gordon* **Rachel Gordon†**

*Dept. of Computer Science, University of Haifa, Haifa, Israel. Email: gordon@cs.haifa.ac.il

†Dept. of Aerospace Engineering, the Technion, Haifa, Israel. Email: rgordon@tx.technion.ac.il

Talk Abstract

We examine the solution of high-frequency Helmholtz equations using 2nd, 4th and 6th order finite difference schemes. The examples include two problems with known analytic solutions, enabling error evaluation of the different schemes on various grids (9–18 points per wavelength). We use our block-parallel CARP-CG algorithm [Parallel Computing 36, 2010] for solving the equations. The algorithm is successful at lowering the relative residual, indicating that it is a robust and reliable parallel solver of the resulting linear systems. However, lowering the error of the solution to reasonable levels is obtained only with the higher order schemes. These results corroborate the known limitations of the low order scheme at modeling the Helmholtz equation, and they indicate that CARP-CG can also be used effectively with high order finite difference schemes. The parallel evaluation uses these two problems and a more realistic third problem.

Introduction

Numerical solution of the Helmholtz equation at high frequencies (large wave numbers) is a challenging computational task. Compounding the problem is the well-known fact that 2nd order finite difference schemes do not model the problem very well and so require many grid points per wavelength. As a result, even if the linear system is solved to within a very small relative residual, the computed solution may be quite far from the true solution of the partial differential equation (PDE).

An additional problem with high wave numbers is the so-called “pollution” effect [1]. It is generally considered that the number of grid points per wavelength, N_g , should be at least 8–12. However, the pollution effect causes the relation between the wave number k and the number of intervals per domain side, N , to be non-linear: N is proportional to $k^{(p+1)/p}$, where p is the order of the accuracy of the scheme; see [8]. Hence, high order schemes have a clear advantage with large wave numbers.

The “shifted Laplacian” approach, introduced in [2], has been studied very extensively. The shift in the Laplacian is used as a preconditioner, and recent years have seen many enhancements, including the use of a com-

plex shift. For a summary and some new results, see [7]. Recently, this approach was extended to higher order schemes [8]. One potential problem with this approach is that the multigrid used for the preconditioner may be difficult to implement on unstructured grids. See also [5], [6], [16], [4], [15] for some other approaches.

This work presents numerical experiments with the block-parallel CARP-CG algorithm [12], applied to the Helmholtz equation with a large wave number, using 2nd, 4th and 6th order finite difference discretization schemes. CARP-CG is described briefly in the next section.

CARP-CG is simple to implement on structured and unstructured grids. In [13], it was shown to be a robust and highly scalable solver of high-frequency Helmholtz equations in homogeneous and heterogeneous media in 2D and 3D domains, using 2nd order finite difference schemes. CARP-CG is generally useful for linear systems with large off-diagonal elements [10], [12], and also with discontinuous coefficients [11].

Our experiments included two test cases with known analytic solutions, with $k = 300$ and $9 \leq N_g \leq 18$. CARP-CG succeeded with all schemes at lowering the residual, indicating that it is a robust and reliable parallel solver of the resulting linear systems. However, lowering the *error* of the solutions to reasonable levels was obtained only with the higher order schemes. These results corroborate the known limitations of the low order scheme at modeling the Helmholtz equation (especially at high frequencies), and they indicate that CARP-CG is also effective with high order finite difference schemes. The parallel evaluation is based on the above two problems, and also on a more realistic third problem.

The CARP-CG Algorithm – Brief Description

Consider a system of m linear equations in n variables, $Ax = b$. The Kaczmarz algorithm (KACZ) [14] is fundamental to CARP-CG. Starting from some arbitrary point, KACZ successively projects the current iterate onto a hyperplane defined by one of the equations in cyclic order. Each cycle of projections is called a KACZ sweep. The projections can be modified with a relaxation parameter $0 < \omega < 2$. If equation i has a fixed relaxation parameter

ω_i , then the projections are said to be done with cyclic relaxation.

CARP [9] is a block-parallel version of KACZ. Note that this is different from the standard block-sequential version of KACZ, which requires the equations in a block to be independent (i.e., there are no shared variables in a block). CARP divides the equations into blocks, which may overlap, and equations in a block need not be independent. In a parallel setting, every processor is in charge of a block of equations. Every processor has a copy, or “clone”, of every variable that it shares with another block. The following two steps are now repeated until convergence:

1. Every processor performs a KACZ sweep on the equations of its assigned block, updating the block’s variables. For shared variables, each processor updates its own clone of the variable.
2. The processors exchange information about the new values of the clones. Every shared variable is now updated to be the average of all its clones in the different blocks, and the new value of every shared variable is distributed among the processors which share it.

If the blocks are chosen according to spatial domains, then the exchange of data occurs only at the boundaries between domains. For a detailed parallel implementation of CARP, see [9]. An important point about CARP is that the averaging operations (between shared variables) are equivalent to certain KACZ row projections in some superspace, so CARP is equivalent to KACZ (with cyclic relaxation) in the superspace. This property provides a convenient convergence proof for CARP, and it enables the conjugate gradient (CG) acceleration of CARP.

The CG acceleration is obtained as follows. If a KACZ sweep is followed by a KACZ sweep in the opposite direction, then the resulting iteration matrix is symmetric and positive semi-definite. Hence, as shown in [3], CG can be applied to the resulting linear system, and this leads to the CGMN algorithm of [3]. CARP-CG is obtained in [12] by extending this method to KACZ with cyclic relaxation in the superspace. On one processor, CARP-CG and CGMN are identical. See [12] for details.

Results and Discussion

Tests were run on a 16-node Linux cluster, using MPICH for message passing. A fixed value of $k = 300$ was used in all cases. The 4th and 6th order schemes were done along the lines of [17], [8]; both of them result in a 9-point stencil matrix. The plots in Figures 1–4 are for one processor, and the parallel performance is examined at the end. The approximate number of variables

for different N_g are as follows (in Problems 1 and 3, the variables are complex, and in Problem 2 they are real):

$N_g =$	9	12	15	18
No. var. \approx	186,000	331,000	515,000	742,000

Problem 1. This example is taken from [8] (but with x and y interchanged for convenience). The equation $\Delta u + k^2 u = 0$ is defined on the square $[-0.5, 0.5] \times [0, 1]$. Dirichlet boundary conditions were taken on three sides: $u = 0$ for $x = -0.5$ and $x = 0.5$, and $u = \cos(\pi x)$ for $y = 0$. On the side $y = 1$, a first order absorbing boundary condition (Sommerfeld radiation condition) was taken as $u_y - i\beta u = 0$, where $\beta^2 = k^2 - \pi^2$. The analytic solution to this problem is $u(x, y) = \cos(\pi x)e^{i\beta y}$.

Fig. 1 shows the relative residual results for the three schemes, for $N_g = 18$. Here, all three convergence plots are quite similar, and those of the higher order schemes are almost identical.




Figure 1: Problem 1: relative residual for $N_g = 18$, 2nd, 4th and 6th order schemes.

Fig. 2 show the L_2 -error plots of the three schemes for $N_g = 18$. The 2nd order scheme stagnates at a value of 0.8 with this mesh size, probably due to pollution.

The L_2 -error with the 6th order scheme, for various values of N_g is shown in Fig. 3. Since higher values of N_g take more computation time, N_g should be chosen according to the desired accuracy.

Problem 2. The PDE $\Delta u + k^2 u = 0$ is defined on the unit square $[0, 1] \times [0, 1]$. The analytic solution chosen for this problem was taken as $u(x, y) = \sin(\pi x) \cos(\beta y)$, where $\beta^2 = k^2 - \pi^2$. Dirichlet boundary conditions were determined by the values of u on the boundary: $u = 0$




Figure 2: Problem 1: L_2 -error for $N_g = 18$, 2nd, 4th and 6th order schemes.

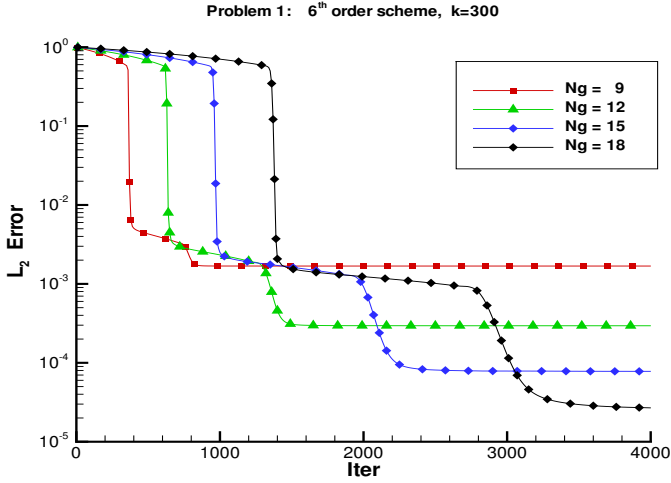


Figure 3: Problem 1: L_2 -error with 6th order scheme, for $9 \leq N_g \leq 18$.

for $x = 0$ and $x = 1$, $u = \sin(\pi x)$ for $y = 0$, and $u = \sin(\pi x) \cos \beta$ for $y = 1$.

The relative residual and L_2 -error plots are quite similar to those of Problem 1, and they are not shown. The L_2 -error plots with the 6th order scheme, for $N_g = 9, 12, 15, 18$, are shown in Fig. 4.

Parallel Performance. We now add a third problem which does not have an analytic solution, but it models a real problem with a disturbance on one side. This problem is identical to [13, Prob. 1]. The equation $\Delta u + k^2 u = 0$ is taken on $[0, 1] \times [0, 1]$. Boundary conditions on the side $y = 0$ are $u(0.5, 0) = 1$ and $u(x, 0) = 0$ for $x \neq 0.5$, i.e., there is a discontinuity at the midpoint. On the other

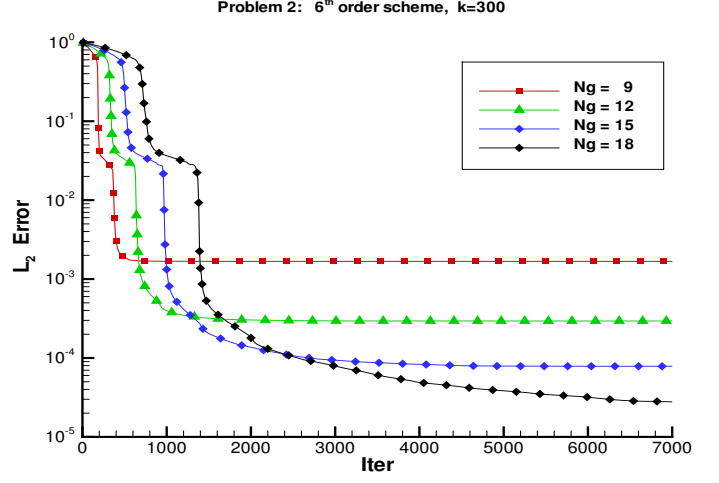


Figure 4: Problem 2: L_2 -error with 6th order scheme, for $9 \leq N_g \leq 18$.

three sides, the boundary conditions were chosen as first-order absorbing boundary condition (the Sommerfeld radiation condition): $\partial u / \partial n - iku = 0$, where n is the unit vector pointing outwards from the domain.

Table 1 shows the number of iterations required to reach a relative residual of 10^{-7} , with the 6th order scheme, for $N_g = 15$, on 1 to 16 processors. Fig. 5 shows the relative residual for Problems 2 and 3 with the 6th order scheme, for $N_g = 15$, on 1 to 16 processors.

Table 1: No. of iterations required to reach a relative residual of 10^{-7} with the 6th order scheme, for $N_g = 15$.

# Proc.	1	2	4	8	12	16
Prob. 1	2881	3516	4634	6125	4478	4983
Prob. 2	3847	3981	4328	4774	5561	5691
Prob. 3	7344	7378	7441	7572	7710	7842

Conclusions

The results show the usefulness of CARP-CG for solving the Helmholtz equation with large wave numbers, using 2nd, 4th and 6th order finite difference schemes with moderate mesh sizes. The two high order schemes are clearly preferable. Both high order schemes take the same time per iteration, so the 6th order scheme is preferable because it is more accurate.

Future research in this area will continue the study of the Helmholtz equation with high order schemes, in 2D and 3D heterogeneous domains. The 3D problems will present a greater challenge due to the large number of nonzero elements in each equation.

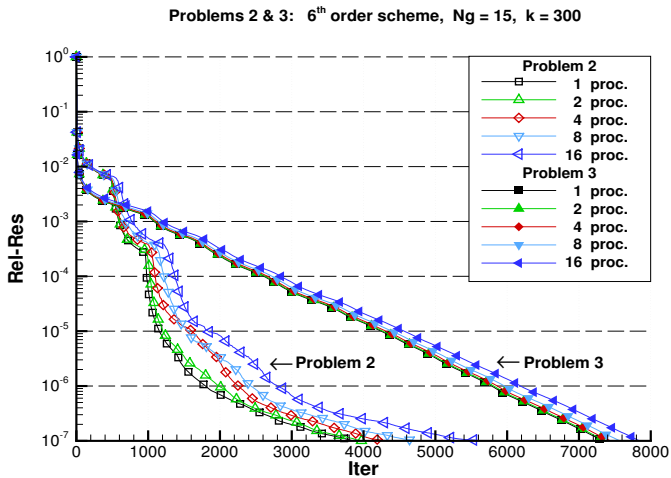


Figure 5: Problems 2 & 3: relative residual for 1–16 processors with 6th order scheme, for $N_g = 15$.

Acknowledgments. The authors wish to thank Dan Givoli and Eli Turkel for their useful comments.

References

- [1] I. M. Babuška and S. A. Sauter. Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers? *SIAM Review*, 42:451–484, 2000.
- [2] A. Bayliss, C. I. Goldstein, and E. Turkel. An iterative method for the Helmholtz equation. *J. of Computational Physics*, 49:443–457, 1983.
- [3] Å. Björck and T. Elfving. Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations. *BIT*, 19:145–163, 1979.
- [4] M. Bollhöfer, M. J. Grote, and O. Schenk. Algebraic multilevel preconditioner for the Helmholtz equation in heterogeneous media. *SIAM J. on Scientific Computing*, 31(5):3781–3805, 2009.
- [5] I. Duff, S. Gratton, X. Pinel, and X. Vasseur. Multigrid based preconditioners for the numerical solution of two-dimensional heterogeneous problems in geophysics. *International J. of Computer Mathematics*, 84(8):1167–1181, Aug. 2007.
- [6] H. C. Elman, O. G. Ernst, and D. P. O’Leary. A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations. *SIAM J. on Scientific Computing*, 23(4):1291–1315, 2001.
- [7] Y. A. Erlangga. Advances in iterative methods and preconditioners for the Helmholtz equation. *Archives of Computational Methods in Engineering*, 15:37–66, 2008.
- [8] Y. A. Erlangga and E. Turkel. Iterative schemes for high order compact discretizations to the exterior Helmholtz equation. *ESAIM: Mathematical Modelling and Numerical Analysis*, to appear, 2011.
- [9] D. Gordon and R. Gordon. Component-averaged row projections: A robust, block-parallel scheme for sparse linear systems. *SIAM J. on Scientific Computing*, 27:1092–1117, 2005.
- [10] D. Gordon and R. Gordon. CGMN revisited: robust and efficient solution of stiff linear systems derived from elliptic partial differential equations. *ACM Trans. on Mathematical Software*, 35(3): 18:1–18:27, Oct. 2008.
- [11] D. Gordon and R. Gordon. Solution methods for linear systems with large off-diagonal elements and discontinuous coefficients. *Computer Modeling in Engineering & Sciences*, 53(1):23–45, Nov. 2009.
- [12] D. Gordon and R. Gordon. CARP-CG: a robust and efficient parallel solver for linear systems, applied to strongly convection-dominated PDEs. *Parallel Computing*, 36(9):495–515, Sept. 2010.
- [13] D. Gordon and R. Gordon. Robust and highly scalable parallel solution of the Helmholtz equation with large wave numbers. Technical report, Dept. of Computer Science, University of Haifa, Israel, Oct. 2010. <http://cs.haifa.ac.il/~gordon/helm.pdf>.
- [14] S. Kaczmarz. Angenäherte Auflösung von Systemen linearer Gleichungen. *Bulletin de l’Académie Polonaise des Sciences et Lettres*, A35:355–357, 1937.
- [15] D. Osei-Kuffuor and Y. Saad. Preconditioning Helmholtz linear systems. *Applied Numerical Mathematics*, 60:420–431, 2010.
- [16] R.-E. Plessix and W. A. Mulder. Separation-of-variables as a preconditioner for an iterative Helmholtz solver. *Applied Numerical Mathematics*, 44:385–400, 2003.
- [17] I. Singer and E. Turkel. Sixth order accurate finite difference schemes for the Helmholtz equation. *J. of Computational Acoustics*, 14:339–351, 2006.