Sparse preconditioners for dense complex linear systems arising in some radar cross section computations

M. Ganesh[1] S. C. Hawkins[2]

(Received 28 July 2006; revised 10 April 2007)

Abstract

We present a sparse preconditioner for efficient iterative solution of large dense linear systems that arise in radar cross section computations for a perfectly conducting scatterer using a high order surface integral equation algorithm. The algorithm allows the linear systems to be assembled efficiently but overall efficiency of the method can only be achieved using iterative solvers with an appropriate preconditioner. We demonstrate the effectiveness and efficiency of our preconditioner for electromagnetic scattering linear systems with tens of thousands of unknowns arising in radar cross section computations for small to medium electromagnetic-sized scatterers.

See http://anziamj.austms.org.au/ojs/index.php/ANZIAMJ/article/view/61 for this article, Published July 11, 2007. ISSN 1446-8735

Contents

1 Introduction

The radar cross section (RCS) of an object indicates how visible (or invisible) the object is when illuminated by an electromagnetic wave. Numerical approximations to the RCS are an essential tool for designing objects that must not be visible, for example stealth aircraft. An efficient high order algorithm for computing the RCS of a perfectly conducting obstacle D (with surface ∂D) is described by Ganesh and Hawkins [6], a variant of an earlier method [7]. The RCS is induced by an incident plane wave

$$\boldsymbol{E}^{\mathrm{i}}(\boldsymbol{x}) = ik\widehat{\boldsymbol{p}}_0 e^{ik\boldsymbol{x}\cdot\widehat{\boldsymbol{d}}_0}\,, \quad \boldsymbol{H}^{\mathrm{i}}(\boldsymbol{x}) = ik(\widehat{\boldsymbol{d}}_0 \times \widehat{\boldsymbol{p}}_0)e^{ik\boldsymbol{x}\cdot\widehat{\boldsymbol{d}}_0}\,,$$

impinging on the obstacle. Here $\widehat{\boldsymbol{d}}_0$ is the direction of the incident wave, $\widehat{\boldsymbol{p}}_0$ is its polarization (perpendicular to $\widehat{\boldsymbol{d}}_0$), and k is the wavenumber of the incident wave. The resulting scattered electromagnetic field $[\boldsymbol{E}, \boldsymbol{H}]$ satisfies the time-harmonic Maxwell equations

$$\operatorname{curl}\boldsymbol{E}(\boldsymbol{x}) - ik\boldsymbol{H}(\boldsymbol{x}) = \boldsymbol{0}\,, \quad \operatorname{curl}\boldsymbol{H}(\boldsymbol{x}) + ik\boldsymbol{E}(\boldsymbol{x}) = \boldsymbol{0}\,, \quad \boldsymbol{x} \in \mathbb{R}^3 \setminus \bar{D}\,, \quad (1)$$

the Silver–Müller radiation condition

$$\lim_{|\boldsymbol{x}|\to\infty} [\boldsymbol{H}(\boldsymbol{x}) \times \boldsymbol{x} - |\boldsymbol{x}|\boldsymbol{E}(\boldsymbol{x})] = \boldsymbol{0}\,, \quad (2)$$

and the perfect conductor boundary condition

$$\boldsymbol{n}(\boldsymbol{x}) \times \boldsymbol{E}(\boldsymbol{x}) = -\boldsymbol{n}(\boldsymbol{x}) \times \boldsymbol{E}^{\mathrm{i}}(\boldsymbol{x}) =: \boldsymbol{f}(\boldsymbol{x}), \quad \boldsymbol{x} \in \partial D, \tag{3}$$

where $\boldsymbol{n}(\boldsymbol{x})$ denotes the unit outward normal at $\boldsymbol{x} \in \partial D$. The RCS is computed from the far field of $\boldsymbol{E}$. Two types of RCS are of particular interest, depending on whether the transmitter and receiver are apart or co-located:

1. the RCS for all directions $\widehat{\boldsymbol{x}}$, with a fixed incident direction $\widehat{\boldsymbol{d}}_0$;
2. the RCS for all directions $\widehat{\boldsymbol{x}}$, with varying incident directions $\widehat{\boldsymbol{d}}_0 = -\widehat{\boldsymbol{x}}$.

These are the bistatic and monostatic RCS respectively [8]. Usually the bistatic and monostatic RCS are represented by their values at a discrete set of measurement directions. In the monostatic case the Maxwell equations (1) must be solved with a different boundary condition for each evaluation direction, because the boundary condition (3) depends on the incident wave direction. Hundreds of measurement directions are required to resolve the monostatic RCS, whose complexity depends on the shape of the perfect conductor and its diameter compared to the incident wavelength.

The preconditioning approach in this work is focused on matrices arising from a particular high order algorithm [6]. Implementation of that fully discrete Galerkin algorithm [6] requires solution of linear systems involving a large, dense, and non-Hermitian matrix with right hand sides depending on the incident wave. In the case of monostatic RCS computations, the discretized electromagnetic scattering matrix is fixed and the right hand side changes for each evaluation direction. Using the high order algorithm [6] one computes accurate approximations to the RCS of medium electromagnetic-sized obstacles using a few tens of thousands of unknowns [6]. Linear systems of this size are solved using the LU factorization of the matrix. This is advantageous in monostatic RCS computations because once the matrix has been factorized, the linear system can be solved for many right hand sides at little extra cost.

TABLE 1: Comparison of the wall time required to setup and factorize the electromagnetic scattering matrix on a cluster of eight nodes, with each node having two dual-core 2 GHz Opteron processors.

Unknowns	matrix setup time	LU factorization time
17326	1030 s	380 s
22186	1670 s	882 s
27646	2920 s	1660 s
33706	4390 s	2950 s
40266	7060 s	4980 s

It is well known that (for three dimensional problems) the matrix setup time in standard Galerkin boundary integral algorithms is usually higher than the matrix solver time, due to evaluation of a quadruple integral (or double surface integral) for each entry of the Galerkin matrix. Efficiently setting up the Galerkin boundary integral matrix is crucial for fast algorithms. Using our high order algorithm and efficient matrix setup techniques [6], we compute the discretized $N \times N$ electromagnetic scattering matrix in $O(N^{2.5})$ operations. (Also, the algorithm [6] requires fewer unknowns than many industrial standard algorithms.) However, the LU factorization of a dense matrix typically requires $O(N^3)$ operations. Table 1 compares the wall time required to set up and factorize various size matrices arising in our algorithm on a distributed memory parallel cluster architecture. Both matrix assembly and factorization are performed in parallel. The matrix is stored by columns in a wrapped interleaved storage. The LU factorization is computed by a column oriented routine with blocking. For small N, the LU factorization is fast compared with the matrix setup, but for large N, the LU factorization cost becomes significant. Thus we are motivated to solve the linear system more quickly using an iterative solver.

Since the surface integral based electromagnetic scattering matrices are complex, dense, and non-Hermitian, the Generalized Minimal Residual algorithm (GMRES) [11] is a natural choice for the iterative solver, and we use it

in this work combined with efficient sparse preconditioners. Fast solution of linear systems using iterative solvers requires a well conditioned iteration matrix. The matrix produced by our algorithm is generally not well-conditioned, and becomes more ill-conditioned with increasing frequency of the incident wave, and with increasing complexity of the scattering object. This article presents a preconditioner that improves the conditioning of the iteration matrix and allows us to solve our linear systems iteratively with just a few tens of iterations. The preconditioner is also cheap to compute.

Preconditioners are well established for linear systems involving sparse matrices. In the case of linear systems involving dense matrices, we require that the preconditioner be sparse to reduce memory use and minimize CPU time. Such preconditioners can be computed directly from the discretized surface integral equation matrix using a sparse approximate inverse technique [1, 9]. Computation of a sparse approximate inverse preconditioner can be accelerated if a sparse approximation to the scattering matrix is used, without degrading significantly the quality of the preconditioner [1, 2]. Another approach is to precondition using the inverse of some easily inverted sparse approximation to the scattering matrix [4]. Alternatively, standard sparse matrix preconditioning techniques can be applied to a sparse approximation of the matrix [10, 13]. Techniques based on incomplete factorization of a sparse approximation to the matrix have been tested for matrices arising from boundary element schemes for electromagnetic scattering [2, 3] but were found to be poor in that setting. However, we show that for the linear systems arising in our high order algorithm, such techniques produce very effective preconditioners.

Section 2 outlines the fully discrete Galerkin scheme that induces the electromagnetic scattering linear systems considered. Section 3 describes our preconditioner and presents numerical results for several test objects at a range of frequencies demonstrating that our preconditioner reduces the number of required GMRES iterations by several hundreds.

2 Boundary integral formulation

Our algorithm [6] is based on the boundary integral formulation for (1)–(3). The RCS is computed from the tangential density $\boldsymbol{w}$, which solves the surface integral equation

$$\boldsymbol{w}(\boldsymbol{x}) + (\mathcal{M}\boldsymbol{w})(\boldsymbol{x}) = 2\boldsymbol{f}(\boldsymbol{x}), \tag{4}$$

with the modified magnetic dipole operator

$$(\mathcal{M}\boldsymbol{a})(\boldsymbol{x}) = 2\int_{\partial D} \boldsymbol{n}(\boldsymbol{x}) \times \operatorname{curl}_{\boldsymbol{x}} \left\{\Phi(\boldsymbol{x},\boldsymbol{y})[I - \boldsymbol{n}(\boldsymbol{y})\boldsymbol{n}(\boldsymbol{y})^T]\boldsymbol{a}(\boldsymbol{y})\right\} ds(\boldsymbol{y}), \tag{5}$$

for $\boldsymbol{x} \in \partial D$, where Φ is the fundamental solution of the Helmholtz equation [5, Equation (2.1)]. The tangential projection operator $I - \boldsymbol{n}(\boldsymbol{y})\boldsymbol{n}(\boldsymbol{y})^T$ is introduced for technical reasons; it allows the use of a non-tangential basis for the solution of (4) by the Galerkin method.

We assume that the perfect conductor surface ∂D is parametrized by a bijective map $\boldsymbol{q} : \partial B \to \partial D$, where ∂B denotes the unit sphere. Under this assumption, the surface integral equation (4) can be rewritten as a surface integral equation on the unit sphere,

$$\boldsymbol{W}(\widehat{\boldsymbol{x}}) + \mathcal{M}\boldsymbol{W}(\widehat{\boldsymbol{x}}) = \boldsymbol{F}(\widehat{\boldsymbol{x}}), \quad \widehat{\boldsymbol{x}} \in \partial B, \quad \boldsymbol{W} = \boldsymbol{w} \circ \boldsymbol{q}, \quad \boldsymbol{F} = \boldsymbol{f} \circ \boldsymbol{q}, \tag{6}$$

where the surface integral operator $\mathcal{M}$ is derived from $\mathcal{M}$ following Ganesh and Hawkins [7].

Our algorithm [6] solves (6) using a fully discrete Galerkin scheme. The ansatz space for the Galerkin scheme is based on an $N(= 3(n+1)^2 - 2)$ dimensional space spanned by the vector spherical polynomials

$$\boldsymbol{Y}^{(1)}_{l,j} = \frac{1}{\sqrt{l(l+1)}} \operatorname{Grad} Y_{l,j}, \quad \boldsymbol{Y}^{(2)}_{l,j} = \boldsymbol{\nu} \times \boldsymbol{Y}^{(1)}_{l,j}, \quad \boldsymbol{Y}^{(3)}_{l,j} = \boldsymbol{\nu} Y_{l,j}, \tag{7}$$

for $0 \le l \le n$, $|j| \le l$ (for convenience we define $\boldsymbol{Y}^{(1)}_{0,0} = \boldsymbol{0} = \boldsymbol{Y}^{(2)}_{0,0}$), where the $Y_{l,j}$ are spherical harmonics [6] and $\boldsymbol{\nu}$ is the unit outward normal to ∂B.

Our Galerkin approximation to the solution $\boldsymbol{W}$ of (6) is

$$\boldsymbol{W}_n = \sum_{l=0}^{n} \sum_{|j|\le l} \sum_{\widetilde{k}=1}^{3} x_{lj\widetilde{k}} \boldsymbol{Y}_{l,j}^{(\widetilde{k})} ,$$

where the coefficients $x_{lj\widetilde{k}}$ are computed by solving the linear system

$$Ax = b , \tag{8}$$

where $A = I + M$,

$$M_{l'j'k',lj\widetilde{k}} = \langle \mathcal{M}' \boldsymbol{Y}_{l,j}^{(\widetilde{k})} , \boldsymbol{Y}_{l',j'}^{(k')} \rangle , \quad b_{l'j'k'} = \langle \boldsymbol{f} , \boldsymbol{Y}_{l',j'}^{(k')} \rangle ,$$

and $\mathcal{M}'$ is a fully discrete approximation to $\mathcal{M}$ [6]. The inner product $\langle \cdot, \cdot \rangle$ on ∂B is a quadrature approximation to the L_2 inner product [6]. The $N \times N$ matrix A is dense, complex, and non-Hermitian. The matrix is assembled in $O(N^{2.5})$ operations using an efficient assembly scheme [6].

3 Preconditioner and numerical experiments

To accelerate the convergence of GMRES we introduce preconditioning, and solve the right preconditioned system

$$AS^{-1}y = b , \quad x = S^{-1}y , \tag{9}$$

where the matrix S is chosen so that AS^{-1} is well conditioned, and vector products with S^{-1} are cheap to compute.

Efficient preconditioning requires that S be a sparse approximation to A. Several strategies for choosing S are described by Carpentieri et al. [2]. Because our Galerkin basis is globally supported, we use an algebraic method to obtain a sparse approximation $S = (s_{ij})$ to $A = (a_{ij})$,

$$s_{ij} = \begin{cases} a_{ij} , & \text{if } |a_{ij}| > \tau \max_{\hat{i}} |a_{\hat{i}j}|, \\ 0 , & \text{otherwise.} \end{cases}$$

The thresholding parameter τ is chosen to give sufficiently good approximation, but not make S too dense. This column-wise thresholding strategy can be performed in parallel without communication between CPUs, because A is stored by columns.

The matrix S can be inverted by computing a sparse LU factorization, but LU factorization generally produces lots of fill-in, which increases the required storage and the cost of the factorization. In this work we reduce fill-in by approximately inverting S using an incomplete LU factorization, where small entries in the LU factors are discarded. We use the Sparskit routine ILUT [12] modified for complex matrices, with drop tolerance ϵ for L and U. Although incomplete LU factorization based preconditioners were found to be poor in some studies [2, 3], results reported in these references used ILU(0), which does not allow any fill-in. We show below that our preconditioner, which uses ILUT(ϵ), is very effective for our algorithm.

We demonstrate the accelerated convergence obtained with our preconditioner for RCS computations. The high order algorithm yields very accurate approximations to the RCS [6], and so we require very accurate solutions to the resulting linear system. In all experiments in this section, we terminate the GMRES iterations when the residual norm has been reduced by a factor of 10^{-10}. In our tables, a $\times$ indicates that GMRES has not converged to this accuracy within a total of 1000 iterations. We allow this relatively large number of iterations so that we can demonstrate convergence of GMRES (without preconditioning) for several of our test problems.

All CPU times presented in this section are the wall time measured on a cluster with five nodes, with each node having two dual-core 2 GHz Opteron processors. We use all five nodes for $N = 27646$ unknowns, but for $N < 27646$ we use only one node. In our experiments the sparse preconditioner is computed in serial on a single node. Parallel computation of the preconditioner is possible—parallel incomplete LU factorization codes are available [11] but are not used in this work. To develop parallel LU factorization codes for general sparse matrices is non-trivial. In almost all of our experiments the

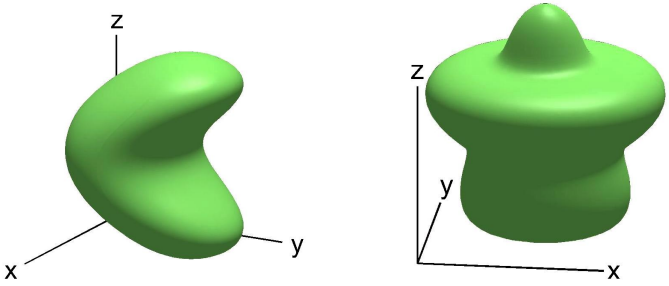

FIGURE 1: Bean and fountain shaped obstacles.

time required to compute the preconditioner is very small compared to the matrix assembly time.

We remark that the preconditioner can be easily stored and used to accelerate solutions for monostatic RCS computations involving many right hand sides. We use compressed sparse column storage for the preconditioner. In our tables preconditioner density is the density of the combined ILU factors of the preconditioner, computed as the total number of nonzeros, expressed as a percentage of N^2.

The linear systems that we solve arise in RCS computations for four different perfectly conducting scatterers: a sphere, sph(siz_obs); an ellipsoid with aspect ratio $4:3:2$, ell(siz_obs); a bean shaped obstacle, bean(siz_obs); and a fountain shaped obstacle, fount(siz_obs). The bean and fountain shaped obstacles are as in Figure 1. Here siz_obs is the diameter of the obstacle. In general, the difficulty of the problem depends on the shape of the obstacle and its electromagnetic-size siz_obs/λ, where $\lambda = 2\pi/k$ is the wavelength of the incident wave. If the frequency of the incident wave is ω, then $\lambda = c/\omega$ where c is the speed of light.

The effectiveness of the sparse preconditioner is demonstrated in Table 2 where we give the number of iterations and CPU time required to solve the linear system using GMRES and preconditioned GMRES. The incident wave

has horizontal incident direction and azimuthal angle 90.0°. In the tables, PGMRES denotes GMRES preconditioned with our preconditioner with thresholding parameters $\tau = \epsilon = 10^{-3}$ as described above. GMRES without preconditioning converges only in four of eight test problems, and even in the cases where it converges the number of iterations are much higher than PGMRES, leading to substantially higher CPU time. In general the number of iterations required to solve the linear system varies with the incident wave direction. We have performed similar comparisons for a range of incident directions with qualitatively similar results to those reported in Table 2.

In Tables 3 and 4 we demonstrate that the number of iterations required with our preconditioner undergoes little growth as the incident wave frequency increases. Similar behaviour was obtained for all objects tested.

In Table 5 we present results for a sparser variant of our preconditioner obtained with parameter $\tau = 10^{-2}$. Using the sparser preconditioner substantially reduces the time required to assemble the preconditioner. For example, the assembly time for the bean(16λ) case with $\tau = 10^{-2}$ is almost half that for $\tau = 10^{-3}$. On the other hand, the preconditioner obtained with $\tau = 10^{-2}$ requires more GMRES iterations than the preconditioner with $\tau = 10^{-3}$. The best choice of preconditioning parameter τ will depend on the number of linear systems to be solved. In particular, for monostatic RCS computations it is desirable to have fewer iterations.

4 Conclusions

We have demonstrated an efficient and effective preconditioner for linear systems arising in a high order Galerkin scheme for RCS computations. Using this preconditioner and GMRES, we are able to solve the linear systems efficiently to simulate three-dimensional scattered fields.

TABLE 2: Comparison of CPU time and number of iterations for PGMRES and GMRES with incident wave direction $\phi = 90.0°$.

Obstacle	unknowns	GMRES itns t_{solve}	PGMRES preconditioner density	PGMRES preconditioner assembly	PGMRES itns t_{solve}
bean(8λ)	13066	536 304.3 s	4.973%	101.6 s	9 7.9 s
fount(8λ)	13066	641 370.7 s	0.458%	17.3 s	6 4.7 s
sph(16λ)	13066	211 112.3 s	0.008%	16.9 s	2 2.3 s
ell(16λ)	13066	×	7.715%	131.0 s	14 12.9 s
bean(16λ)	27646	×	10.552%	2729.8 s	15 28.8 s
fount(16λ)	27646	×	0.382%	168.8 s	7 6.6 s
sph(24λ)	27646	297 246.8 s	0.004%	163.8 s	2 2.4 s
ell(24λ)	27646	×	7.386%	1174.4 s	15 25.0 s

TABLE 3: Comparison of CPU time and number of iterations for PGMRES and GMRES for incident wave direction $\phi = 90.0°$.

Obstacle	unknowns	GMRES	PGMRES		
		itns	preconditioner		itns
		t_{solve}	density	assembly	t_{solve}
sph(λ)	766	16 0.0 s	0.131%	0.1 s	2 0.0 s
sph(8λ)	3886	100 5.0 s	0.026%	1.4 s	2 0.2 s
sph(16λ)	13066	211 112.3 s	0.008%	16.9 s	2 2.3 s
sph(24λ)	27646	297 246.8 s	0.004%	163.8 s	2 2.4 s

TABLE 4: Comparison of CPU time and number of iterations for PGMRES and GMRES for incident wave direction $\phi = 90.0°$

Obstacle	unknowns	GMRES	PGMRES		
		itns	preconditioner		itns
		t_{solve}	density	assembly	t_{solve}
fount(λ)	3886	31 1.6 s	1.106%	1.5 s	5 0.4 s
fount(8λ)	13066	641 370.7 s	0.458%	17.3 s	6 4.7 s
fount(16λ)	27646	×	0.382%	168.8 s	7 6.6 s

TABLE 5: Comparison of CPU time and number of iterations for PGMRES with $\tau = 10^{-2}$ and GMRES for incident wave direction $\phi = 90.0°$.

Obstacle	unknowns	GMRES	PGMRES		
		itns	preconditioner		itns
		t_{solve}	density	assembly	t_{solve}
bean(8λ)	13066	536 304.3 s	2.913%	57.5 s	30 22.8 s
fount(8λ)	13066	641 370.7 s	0.232%	17.0 s	11 8.3 s
sph(16λ)	13066	211 112.3 s	0.008%	16.9 s	2 2.1 s
ell(16λ)	13066	×	4.937%	82.1 s	105 83.9 s
bean(16λ)	27646	×	6.727%	1522.8 s	197 285.1 s
fount(16λ)	27646	×	0.257%	165.1 s	26 20.5 s
sph(24λ)	27646	297 246.8 s	0.004%	164.4 s	2 2.6 s
ell(24λ)	27646	×	4.579%	647.6 s	214 277.2 s

Acknowledgements Support of the Australian Research Council is gratefully acknowledged.

References

[1] G. Alleon, Michele Benzi, and L. Giraud. Sparse approximate inverse preconditioning for dense linear systems arising in computational electromagnetics. *Numerical Algorithms*, 16:1–15, 1997. doi:10.1023/A:1019170609950. C237

[2] B. Carpentieri, I. S. Duff, and L. Giraud. Sparse pattern selection strategies for robust frobenius-norm minimization preconditioners in electromagnetism. *Numer. Linear Algebra Appl.*, 7:667–685, 2000. http://dx.doi.org/10.1002/1099-1506(200010/12)7:7/8<667::AID-NLA218>3.0.CO;2-X. C237, C239, C240

[3] B. Carpentieri, I. S. Duff, L. Giraud, and M. Magolu monga Made. Sparse symmetric preconditioners for dense linear systems in electromagnetism. *Numer. Linear Algebra Appl.*, 11:753–771, 2004. doi:10.1002/nla.345. C237, C240

[4] K. Chen. Efficient iterative solution of linear systems from discretizing singular integral equations. *Elec. Tran. Numer. Anal.*, 2:76–91, 1994. http://etna.mcs.kent.edu/vol.2.1994/pp76-91.dir/pp76-91.pdf. C237

[5] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer, 1998. C238

[6] M. Ganesh and S. C. Hawkins. An efficient surface integral method for the time-harmonic Maxwell equations. *ANZIAM*, 48:C17–C33, 2007.

http://anziamj.austms.org.au/ojs/index.php/ANZIAMJ/article/view/60/43. C234, C235, C236, C238, C239, C240

[7] M. Ganesh and S. C. Hawkins. A spectrally accurate algorithm for electromagnetic scattering in three dimensions. *Numerical Algorithms*, 43:25–60, 2006. doi:10.1007/s11075-006-9033-7. C234, C238

[8] E. F. Knott, J. F. Shaeffer, and M. T. Tuley. *Radar Cross Section*. SciTech Publishing Inc, 2004. C235

[9] L. Y. Kolotolina. Explicit preconditioning of systems of linear algebraic equations. *J. Sov. Math.*, 43:2566–2573, 1988. C237

[10] J. Lee, J. Zhang, and C. Lu. Incomplete LU preconditioning for large scale dense complex linear systems from electromagnetic wave scattering problems. *J. Comp. Phys.*, 185:158–175, 2003. doi:10.1016/S0021-9991(02)00052-9. C237

[11] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, second edition, 2003. C236, C240

[12] Y. Saad. SPARSKIT. http://www-users.cs.umn.edu/~saad/software/home.html, 2005. C240

[13] K. Sertel and J. L. Volakis. Incomplete LU preconditioner for FMM implementation. *Micro. Opt. Tech. Lett.*, 26:265–267, 2000. http://dx.doi.org/10.1002/1098-2760(20000820)26:4<265::AID-MOP18>3.0.CO;2-O. C237

Author addresses

1. **M. Ganesh**, Department of Mathematical and Computer Sciences, Colorado School of Mines, Golden, CO 80401.
mailto:mganesh@mines.edu

2. **S. C. Hawkins**, School of Mathematics and Statistics, University of New South Wales, Sydney, NSW 2052, Australia.
mailto:stuart@maths.unsw.edu.au