

Fast Preconditioned Krylov Methods for Boundary Integral Equations in Electromagnetic Scattering

Bruno Carpentieri

*University of Groningen, Institute of Mathematics and Computing Science,
Groningen
The Netherlands*

1. Introduction

An accurate numerical solution of Electromagnetic scattering problems is critically demanded in the simulation of many real-life applications, such as in the design of industrial processes and in the study of wave propagation phenomena. Electromagnetic (EM) scattering problems address the physical issue of computing the diffraction pattern of the EM radiation that is propagated by a complex body, illuminated by an incident wave. An explicit solution is possible only for simple targets, e.g. for spherical bodies; complicated geometries impose to use approximate numerical techniques.

Until the emergence of high-performance computing in the early eighties, the analysis of scattering problems was afforded by using approximate high frequency techniques such as the shooting and bouncing ray method (SBR) (Lee et al. (1988)). Ray-based asymptotic methods like SBR and the uniform theory of diffraction are based on the idea that EM scattering becomes a localized phenomenon as the size of the scatterer increases with respect to the wavelength. In the last decades, due to impressive advances in computer technology and the introduction of innovative algorithms with limited computational and memory requirement, a more rigorous numerical solution has become possible for many practical applications.

Finite-difference (FD) (Kunz & Luebbers (1993); Taflove (1995)), finite-element (FE) (Silvester & Ferrari (1990); Volakis et al. (1998)) and finite-volume (FV) methods (Bonnet et al. (1998); Botha (2006)) can be used to discretize the Maxwell's equations into a finite volume surrounding the scatterer, giving rise to sparse systems of linear equations. Upon inversion of the system, a solution is computed for all excitations. More recently, alternative approaches based on integral equations are becoming increasingly popular for solving high-frequency EM scattering problems. They reformulate the Maxwell's equations in the frequency domain and solve for the electric and the magnetic currents induced on the surface of the object. Thus integral methods require only a simple description of the surface of the target by means of triangular facets (see an example of discretization in Figure 1). This means that a 3D problem is reduced to solving a 2D surface problem, simplifying considerably the mesh generation especially in the case of moving objects. No artificial boundaries need to be imposed and boundary conditions are automatically satisfied in the case of perfectly conducting objects.

Another interest to use surface discretizations is that they noticeably reduce the effect of grid dispersion errors. Grid dispersion errors occur when a wave has a different phase velocity on

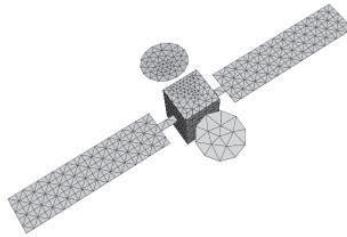


Fig. 1. Example of surface discretization in an integral equation context. Each unknown of the problem is associated to an edge in the mesh. Courtesy of the EMC-CERFACS Group in Toulouse.

the grid compared to the exact solution; they tend to accumulate in space and may introduce spurious solutions over large 3D simulation regions (Bayliss et al. (1985); Jr. (1994); Lee & Cangellaris (1992)). For second-order accurate differential schemes, to alleviate this problem the grid density may grow up to $\mathcal{O}((kd)^3)$ unknowns in 2D and of $\mathcal{O}((kd)^{4.5})$ in 3D, where k is the wavenumber and d is the approximate diameter of the simulation region. Therefore, the overall solution cost may increase considerably also for practical (*i.e.* finite) values of wavenumber (Chew et al. (1997)).

Boundary element discretizations are applied in many scientific and engineering areas beside electromagnetics and acoustics, e.g. in biomagnetic and bioelectric inverse modeling, magnetostatic and biomolecular problems, and many other applications (Forsman, Gropp, Kettunen & Levine (1995); Yokota, Bardhan, Knepley, Barba & Hamada (2011)). The potential drawback is that they lead, upon discretization, to large and dense linear systems to invert. Hence fast numerical linear algebra methods and efficient parallelization techniques are urged for solving large-scale boundary element equations efficiently on modern computers. In this chapter we overview some relevant techniques. In Section 2 we introduce the boundary integral formulation for EM scattering from perfectly conducting objects. In Section 4 we discuss fast iterative solution strategies based on preconditioned Krylov methods for solving the dense linear system arising from the discretization. In Section 5 we focus our attention on the design of the preconditioner, that is a critical component of Krylov methods in this context. We conclude our study in Section 5 with some final remarks.

2. The integral equation context

In an integral equation context, the standard EM scattering problem may be formulated in variational form as follows:

Find the surface current \vec{j} such that for all tangential test functions \vec{j}^t , we have

$$\int_{\Gamma} \int_{\Gamma} G(|y-x|) \left(\vec{j}(x) \cdot \vec{j}^t(y) - \frac{1}{k^2} \operatorname{div}_{\Gamma} \vec{j}(x) \cdot \operatorname{div}_{\Gamma} \vec{j}^t(y) \right) dx dy = \frac{i}{kZ_0} \int_{\Gamma} \vec{E}_{inc}(x) \cdot \vec{j}^t(x) dx. \quad (1)$$

Eqn. (1) is called Electric Field Integral Equation (EFIE) (see Bilotti & Vegni (2003); Li et al. (2005)); we denote by $G(|y - x|) = \frac{e^{ik|y-x|}}{4\pi|y-x|}$ the Green's function of Helmholtz equation, Γ is the boundary of the object, k the wave number and $Z_0 = \sqrt{\mu_0/\epsilon_0}$ the characteristic impedance of vacuum (ϵ is the electric permittivity and μ the magnetic permeability). Given a continuously differentiable vector field $\vec{j}(x)$ represented in Cartesian coordinates on a 3D Euclidean space as $\vec{j}(x_1, x_2, x_3) = j_{x_1}(x_1, x_2, x_3)\vec{e}_{x_1} + j_{x_2}(x_1, x_2, x_3)\vec{e}_{x_2} + j_{x_3}(x_1, x_2, x_3)\vec{e}_{x_3}$, where $\vec{e}_{x_1}, \vec{e}_{x_2}, \vec{e}_{x_3}$ are the unit basis vectors of the Euclidean space, we denote by $\text{div}\vec{j}(x)$ the divergence operator defined as

$$\text{div}\vec{j}(x) = \frac{\partial j_{x_1}}{\partial x_1} + \frac{\partial j_{x_2}}{\partial x_2} + \frac{\partial j_{x_3}}{\partial x_3}.$$

The EFIE formulation can be applied to arbitrary geometries such as those with cavities, disconnected parts, breaks on the surface; hence, it is very popular in industry.

For closed targets, the Magnetic Field Integral Equation (MFIE) can be used, which reads

$$\int_{\Gamma} (\vec{R}_{ext} j \wedge \vec{v}) \cdot \vec{j}^{\vec{t}} + \frac{1}{2} \int_{\Gamma} \vec{j} \cdot \vec{j}^{\vec{t}} = - \int_{\Gamma} (\vec{H}_{inc} \wedge \vec{v}) \cdot \vec{j}^{\vec{t}}.$$

The operator $\vec{R}_{ext} j$ is defined as

$$\vec{R}_{ext} j(y) = \int_{\Gamma} \text{grad}_y G(|y - x|) \wedge \vec{j}(x) dx,$$

and is evaluated in the domain exterior to the object.

Both formulations suffer from interior resonances which make the numerical solution more problematic at some frequencies known as resonant frequencies, especially for large objects. The problem can be solved by combining linearly EFIE and MFIE. The resulting integral equation, known as Combined Field Integral Equation (CFIE), is the formulation of choice for closed targets. We point the reader to Gibson (2008) for a thorough presentation of integral equations in electromagnetism.

On discretizing Eqn. (1) in space by the Method of Moments (MoM) over a mesh containing n edges, the surface current \vec{j} is expanded into a set of basis functions $\{\vec{\varphi}_i\}_{1 \leq i \leq n}$ with compact support (the Rao-Wilton-Glisson basis, Rao et al. (1982), is a popular choice), then the integral equation is applied to a set of tangential test functions $\vec{j}^{\vec{t}}$. Selecting $\vec{j}^{\vec{t}} = \vec{\varphi}_j$, we are led to compute the set of coefficients $\{\lambda_i\}_{1 \leq i \leq n}$ such that

$$\begin{aligned} \sum_{1 \leq i \leq n} \lambda_i \left[\int_{\Gamma} \int_{\Gamma} G(|y - x|) \left(\vec{\varphi}_i(x) \cdot \vec{\varphi}_j(x) - \frac{1}{k^2} \text{div}_{\Gamma} \vec{\varphi}_i(x) \cdot \text{div}_{\Gamma} \vec{\varphi}_j(y) \right) dx dy \right] = \\ = \frac{i}{kZ_0} \int_{\Gamma} \vec{E}_{inc}(x) \cdot \vec{\varphi}_j(x) dx, \end{aligned} \quad (2)$$

for each $1 \leq i \leq n$. The set of equations (2) can be recast in matrix form as

$$A\lambda = b, \quad (3)$$

where $A = [A_{ij}]$ and $b = [b_i]$ have elements, respectively,

$$A_{ij} = \int_{\Gamma} \int_{\Gamma} G(|y-x|) \left(\vec{\varphi}_i(x) \cdot \vec{\varphi}_j(y) - \frac{1}{k^2} \operatorname{div}_{\Gamma} \vec{\varphi}_i(x) \cdot \operatorname{div}_{\Gamma} \vec{\varphi}_j(y) \right) dx dy,$$

$$b_j = \frac{i}{kZ_0} \int_{\Gamma} \vec{E}_{inc}(x) \cdot \vec{\varphi}_j(y) dx.$$

The set of unknowns are associated with the vectorial flux across an edge in the mesh. The right-hand side varies with the frequency and the direction of the illuminating wave.

3. Fast matrix solvers for boundary element equations

Linear systems issued from boundary element discretizations may be very large in applications, although their size is typically much smaller compared to those arising from FE or FV formulations of the same problem. The number of unknowns grows linearly with the size of the scatterer and quadratically with the frequency of the incoming radiation (Bendali (1984)). A target with size of a few tens of wavelength, illuminated at $\mathcal{O}(1)$ GHz of frequency, may lead to meshes with several million points (Sylvand (2002)). Some efficient out-of-core dense direct solvers based on variants of Gaussian elimination have been proposed for solving blocks of right-hand sides, see e.g. Alléon, Amram, Durante, Homs, Pogarieloff & Farhat (1997); Chew & Wang (1993). However, the memory requirements of direct methods are not affordable for solving such systems in realistic applications, even on modern parallel computers. Iterative methods can solve the problems of space of direct methods because they are based on matrix-vector (M-V) multiplications. In general terms, a modern integral equation solver is the mix of a robust iterative method, a fast algorithm for computing cheap approximate M-V products, and an efficient preconditioner to speed-up the convergence.

3.1 The choice of the iterative method

Krylov methods are among the most popular accelerators because of their ability to deliver good rates of convergence and to handle very large problems efficiently. They look for the solution of the system $Ax = b$ in the Krylov space $K_k(A, b) = \operatorname{span}\{b, Ab, A^2b, \dots, A^{k-1}b\}$. This is a good space from which to construct approximate solutions for a nonsingular linear system because it is intimately related to A^{-1} . In fact the inverse of any nonsingular matrix A can be written in terms of powers of A with the help of the minimal polynomial $q(t)$ of A , which is the unique monic polynomial of minimal degree such that $q(A) = 0$. If $\lambda_1, \dots, \lambda_d$ are the distinct eigenvalues of A , m_j is the index of λ_j , and we define m as

$$m = \sum_{j=1}^d m_j,$$

then

$$q(t) = \prod_{j=1}^d (t - \lambda_j)^{m_j}. \quad (4)$$

Writing $q(t)$ in the form

$$q(t) = \sum_{j=0}^m \alpha_j t^j,$$

Solver	Iterations (CPU time)
CORS	601 (253*)
BiCOR	785 (334)
GMRES(50)	2191 (469)
QMR	878 (548)
BiCGSTAB	1065 (444)

Table 1. Number of iterations and CPU time (in seconds) required by Krylov methods to reduce the initial residual to $\mathcal{O}(10^{-8})$. An asterisk "*" indicates the fastest run. The problem is shown in Figure 2

we have

$$A^{-1} = -\frac{1}{\alpha_0} \sum_{j=0}^m \alpha_{j+1} A^j \quad , \quad \alpha_0 = \prod_{j=1}^d (-\lambda_j)^{m_j} \neq 0.$$

This shows that, if the minimal polynomial of A has degree m , then the solution of $Ax = b$ lies in the space $K_m(A, b)$. The smaller the degree of the minimal polynomial, the faster the expected rate of convergence of a Krylov method (see Ipsen & Meyer (1998)).

One issue is the choice of the suitable Krylov algorithm. Most integral formulations for surface and hybrid surface/volume scattering give rise to indefinite linear systems that cannot be solved using the Conjugate Gradient method (see discussions in Section 2). The GMRES method by Saad & Schultz (1986) is virtually always used for solving dense non-Hermitian linear systems as it is an optimal iterative solver, in the sense that it minimizes the 2-norm of the residual over the corresponding Krylov space. It generally requires the least number of iterations to converge. However, the optimality of GMRES comes at a price. The cost of applying the method increases with the iterations, and it may sometimes become prohibitively expensive for solving practical applications. As an attempt to limit the costs of GMRES, the algorithm is often restarted. After a given number of steps k , the approximate solution is computed from the generated Krylov subspace. Then the Krylov subspace is destroyed, and a new space is reconstructed using the latest residual.

On the other hand, non-optimal methods attempt to limit the costs of GMRES while preserving its favourable convergence properties. In Table 1, we show the number of iterations required by Krylov methods to reduce the initial residual to $\mathcal{O}(10^{-8})$ starting from the zero vector on the problem shown in Figure 2. For simplicity, the right-hand side of the linear system is set up so that the initial solution is the vector of all ones. We do not use preconditioning. In addition to restarted GMRES, we consider complex versions of iterative algorithms based on Lanczos biorthogonalization, such as BiCGSTAB (van der Vorst (1992)) and QMR (Freund & Nachtigal (1994)) and on the recently developed Lanczos biconjugate A -orthonormalization, such as BiCOR and CORS (Carpentieri et al. (2011); Jing, Huang, Zhang, Li, Cheng, Ren, Duan, Sogabe & Carpentieri (2009)). We clearly observe the importance of the choice of the iterative method. In our experiments, the CORS method is the fastest non-Hermitian solver with respect to CPU time on most selected examples except GMRES with large restart. Indeed, unrestarted GMRES may outperform all other Krylov methods and should be used when memory is not a concern. However, reorthogonalization costs may penalize the GMRES convergence in large-scale applications, so using high values of restart may not be convenient (or even not affordable for the memory) as shown in Carpentieri et al. (2005). In Table 1 we select a value of 50 for the restart parameter.

The BiCOR and CORS methods are introduced in Carpentieri et al. (2011); Jing, Huang, Zhang, Li, Cheng, Ren, Duan, Sogabe & Carpentieri (2009). They search for the approximate

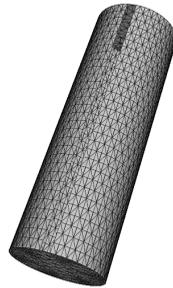


Fig. 2. Test problem: an open cylindrical surface. Characteristics of the associated linear system: size=6268, frequency=362 MHz, $\kappa_1(A) = \mathcal{O}(10^5)$. Courtesy of the EMC-CERFACS Group in Toulouse.

solution in the Krylov subspace $K_m(A, r_0)$ by applying a Petrov-Galerkin approach and imposing the residual be orthogonal to the constraints subspace $A^H K_m(A^H, r_0^*)$; the shadow residual r_0^* is chosen to be equal to $r_0^* = Ar_0$. The basis vector representations for the subspaces $K_m(A, r_0)$ and $A^H K_m(A^H, r_0^*)$ are computed by means of the biconjugate A -Orthonormalization procedure. Starting from two vectors v_1 and w_1 chosen to satisfy the condition $\omega_1^H A v_1$, the method ideally builds up a pair of biconjugate A -orthonormal bases $v_j, j = 1, 2, \dots, m$ and $w_i, i = 1, 2, \dots, m$, respectively for the dual Krylov subspaces $K_m(A; v_1)$ and $K_m(A^H; w_1)$, satisfying the condition $\omega_i^H A v_j = \delta_{ij}, 1 \leq i, j \leq m$. We point the reader to Jing, Carpentieri & Huang (2009) for further experiments with iterative Krylov methods for surface integral equations.

A significant amount of work has been devoted in the last years to design fast algorithms that can reduce the $\mathcal{O}(n^2)$ computational complexity for the M-V product operation required at each step of a Krylov method, such as the Fast Multipole Method (FMM) (Greengard & Rokhlin (1987); Rokhlin (1990)), the panel clustering method (Hackbush & Nowak (1989)), the \mathcal{H} -matrix approach (Hackbush (1999)), wavelet techniques (Alpert et al. (1993); Bond & Vavasis (1994)), the adaptive cross approximation method (Bebendorf (2000)), the impedance matrix localization method (Canning (1990)), the multilevel matrix decomposition algorithm (Michielssen & Boag (1996)) and others. In particular, the combination of iterative Krylov subspace solvers and FMM is a popular approach for solving integral equations. For Helmholtz and Maxwell problems, FMM algorithms enable to speedup M-V multiplications with boundary element matrices down to $\mathcal{O}(n \log n)$ algorithmic and memory complexity depending on the problem and on the specific implementation, see e.g. Cheng et al. (2006); Darrigrand (2002); Darve & Havé (2004); Dembart & Epton (1994); Engheta et al. (1992); Song & Chew (1995); Tausch (2004). Two-level implementations of FMM can reduce the cost of the M-V product operation from $\mathcal{O}(n^2)$ to $\mathcal{O}(n^{3/2})$, a three level algorithm down to $\mathcal{O}(n^{4/3})$ and the Multilevel Fast Multipole Algorithm (MLFMA) to $\mathcal{O}(n \log n)$.

Multipole techniques exploit the rapid decay of the Green's function and compute interactions amongst degrees of freedom in the mesh at different levels of accuracy depending on their physical distance. The 3D mesh of the object is partitioned recursively into boxes of roughly equal size until the size becomes small compared with the wavelength. The hierarchical partitioning of the object is typically represented using a tree-structured data called *oct-tree*

(see Figure 3). Multipole coefficients are computed for all boxes starting from the smallest

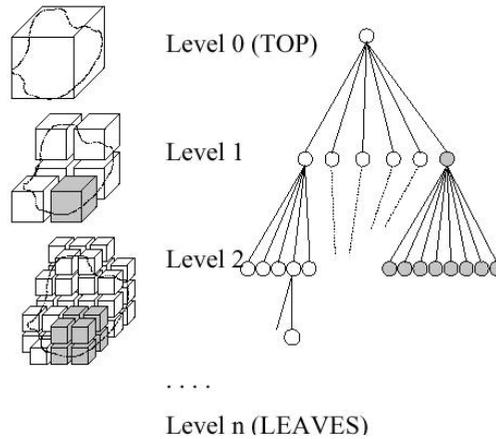


Fig. 3. The *oct-tree* data structure representation in the FMM algorithm. Each cube has up to eight children and one parent box except for the largest cube which encloses the whole domain.

ones, that are the leaves, and recursively for each parent cube in the tree by summing together multipole coefficients of its children. Interactions of degrees of freedom within one observation box and its close neighboring boxes are computed exactly using MoM; depending on the frequency, they generate between 1% and 2% of the entries of A . Interactions with boxes that are not neighbors of the observation box but whose parent in the oct-tree is a neighbor of the box parent are computed using FMM (see Figure 4). All other interactions are computed

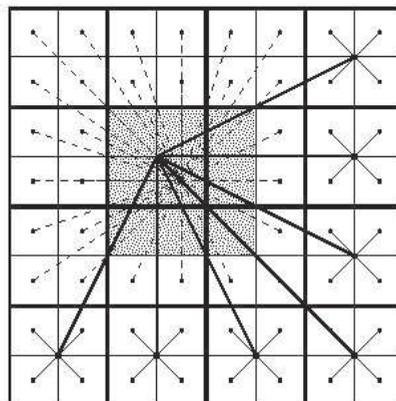


Fig. 4. Interactions in the multilevel FMM algorithm. Interactions for the gray boxes are computed directly. We denote by dashed lines cubes that are not neighbors of the cube itself but whose parent is a neighbor of the cube's parent. These interactions are computed using the FMM. All other interactions are computed hierarchically on a coarser level.

hierarchically on a coarser level by traversing the oct-tree. Multiple techniques have been efficiently implemented on distributed memory parallel computers proving to be scalable to several million discretization points, see for instance the FISC code developed at University of Illinois by Song & Chew (1998); Song et al. (1997; 1998), the INRIA/EADS integral equation code AS_ELFIP by Sylvand (2002; 2003), the Bilkent University code by Ergül & Gürel (2007; 2008) and others.

4. Algebraic preconditioning for boundary integral equations

Krylov methods may converge very slowly in practice, mainly due to bad spectral conditioning of the linear system. Relation (4) implies that the dimension of the solution space, and therefore the convergence properties, are mostly dictated by the eigenvalue distribution of A . The spectral properties may vary noticeably depending on the integral operator as well as on object shape and material. Problems with cavities or open surfaces typically require more iterations to converge than closed objects of the same physical size, and nonuniform meshes often produce ill-conditioned MoM matrices. On EFIE, the iteration count of Krylov solvers may increase as $\mathcal{O}(n^{0.5})$ when the number of unknowns n is related to the wavenumber, see for instance experiments reported in Song & Chew (1998), whereas on CFIE the number of iterations typically increases as $\mathcal{O}(n^{0.25})$.

On the other hand, if preconditioning A by a nonsingular matrix M the eigenvalues of $M^{-1}A$ fall into a few clusters, say t of them, whose diameters are small enough, then $M^{-1}A$ behaves numerically like a matrix with t distinct eigenvalues. As a result, we would expect t iterations of a Krylov method to produce reasonably accurate approximations. The matrix M is called the *preconditioner* matrix; preconditioning can be applied from the left as $M^{-1}Ax = M^{-1}b$ as well as from the right as $AM^{-1}y = b$ with $x = M^{-1}y$.

Optimal analytic preconditioners have been proposed for surface integral equations, see e.g. Antoine et al. (2004); Christiansen & Nédélec (2002); Steinbach & Wendland (1998). But they are problem-dependent. In this study, we consider purely algebraic techniques which compute the preconditioner only using information contained in the coefficient matrix of the linear system. Although far from optimal for any specific problem, algebraic methods can be applied to different operators and to changes in the geometry only by tuning a few parameters, and may often be developed from existent public-domain software implementations.

We are interested to develop techniques that have $\mathcal{O}(n \log n)$ algorithmic and memory complexity in the construction and in the application phase like FMM, and may be implemented efficiently within multipole codes. For memory concerns, we compute the preconditioner by initially decomposing the linear system in the form

$$(S + B)x = b \tag{5}$$

where S is a sparse matrix retaining the most relevant contributions to the singular integrals and is easy to invert, while B can be dense. If the continuous operator \mathcal{S} underlying S is bounded and the operator \mathcal{B} underlying B is compact, then $\mathcal{S}^{-1}\mathcal{B}$ is compact and

$$\mathcal{S}^{-1}(S + B) = \mathcal{I} + \mathcal{S}^{-1}B.$$

We may expect that the preconditioned system $(I + S^{-1}B)x = S^{-1}b$ has a good clusterization of eigenvalues close to one, see e.g. Chen (1994) and (Chen, 2005, pp. 182-185).

The simplest approach to compute the local matrix S is to drop the small entries of A below a threshold (Cosnau (1996); Kolotilina (1988); Vavasis (1992)). When all the entries of A are not explicitly available, it may be necessary to use information extracted from the physical mesh of the problem. In an integral equation context, the surface of the object is discretized using a triangular mesh; each degree of freedom (DOF), or equivalently each unknown of the linear system, is associated to an edge of the mesh. Therefore, the sparsity pattern of S can be defined according to the concept of level k neighbours (see Figure 5(a)). Level 1 neighbours of a DOF are the DOF plus the four DOFs belonging to the two triangles that share the edge corresponding to the DOF itself. Level 2 neighbours are all the level 1 neighbours plus the DOFs in the triangles that are neighbours of the two triangles considered at level 1, and so forth. Due to the very localized nature of the Green's function, by retaining a few (two or three) levels of neighbours for each DOF an effective approximation may be constructed.

Comparative experiments show that there is little to choose. Both matrix- and mesh-based approaches can provide very good approximations S to the dense coefficient matrix for low sparsity ratio between 1% and 2% (Carpentieri et al. (2000)). The mesh-based approach is straightforward to implement in FMM codes as the object is typically partitioned using geometric information (see Figure 5(b)). Multipole algorithms yield a matrix decomposition

$$A = A_{diag} + A_{near} + A_{far}, \tag{6}$$

where A_{diag} is the block diagonal part of A associated with interactions of basis functions belonging to the same box, A_{near} is the block near-diagonal part of A associated with interactions within one level of neighboring boxes (they are 8 in 2D and 26 in 3D), and A_{far} is the far-field part of A . Therefore, in a multipole setting a suitable choice for the local matrix may be $S = A_{diag} + A_{near}$.

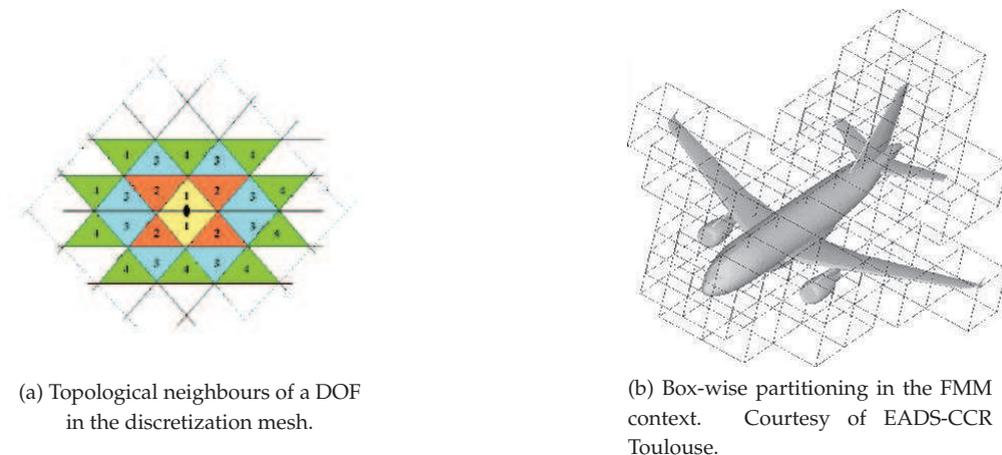


Fig. 5. Mesh-based pattern selection strategies to compute local interactions in an integral equation context.

4.1 Comparison of standard preconditioners

To illustrate the difficulty of finding a good preconditioner for this problem class, in Table 2 we report one experiments with the GMRES solver and various algebraic preconditioners applied to a scattering problem from an open cylindric surface illuminated at 200 MHz of frequency and modeled using EFIE. The system has $n = 1299$ unknowns and is a low resolution testcase than the problem in Figure 2. In connection with GMRES, we consider preconditioners M of either implicit type (which approximately factorize S) or of explicit type (which approximately invert S) at roughly the same number of nonzero entries in M . We adopt the following acronyms:

- *None*, means that no preconditioner is used;
- *Diag*, a simple diagonal scaling, *i.e.* M is the diagonal of S ;
- *SSOR*, the symmetric successive overrelaxation method $M = \frac{(D+\omega E)D^{-1}(D+\omega E^T)}{\omega(2-\omega)}$, where we denote by D the diagonal of S and E is the strict lower triangular part of S ;
- *ILU(0)* by Saad (1996), the lower/upper incomplete LU factorization $M = \tilde{L}\tilde{U}$, $\tilde{L} \approx L$, $\tilde{U} \approx U$, $S = LU$, where the sparsity pattern of \tilde{L} (resp. \tilde{U}) is equal to that of the lower (resp. upper) triangular part of S ;
- *SPAI* by Grote & Huckle (1997), an approximate inverse preconditioner $M \approx S^{-1}$ computed by minimizing $\|I - SM\|_F$. The same pattern of S is imposed to M .
- *AINV* by Benzi et al. (1996), a sparse approximate inverse computed in factorized form by applying an incomplete biconjugation process to S , and dropping small entries below a threshold in the inverse factors.

Density of $S = 3.18\%$ - Density of $M = 1.99\%$

Precond.	GMRES(30)	GMRES(80)	GMRES(∞)
<i>None</i>	-	-	302
<i>Diag</i>	-	-	272
<i>SSOR</i>	-	717	184
<i>ILU(0)</i>	-	454	135
<i>SPAI</i>	308	70	70
<i>AINV</i>	-	-	-

Table 2. Number of iterations using GMRES and various preconditioners on a test problem, a cylinder with an open surface, discretized with $n = 1299$ edges. The tolerance is set to 10^{-8} . The symbol '-' means that no convergence was achieved after 1000 iterations. The results are for right preconditioning.

We see that many standard methods fail. Simple preconditioners, like the diagonal of A , diagonal blocks, or a band, may be effective when the coefficient matrix has some degree of diagonal dominance (Song et al. (1997)). For ill-conditioned and indefinite matrices, more robust methods are needed. Techniques that are successful for solving partial differential equations may be successfully adopted for integral equations; in the next section, we analyse some of these methods.

4.2 Sparse approximate inverses preconditioner

Approximate inverse methods are very attractive for parallelism. They explicitly compute and store an approximation of the inverse of the coefficient matrix $M \approx S^{-1}$, which may be used as preconditioner by performing one or more sparse M-V products operations at each step of an iterative solver. As shown in Figure 6, due to the rapid decay of the Green’s function the entries of A^{-1} may have a very similar structure to those of A , so that a very sparse preconditioner M may effectively capture the large contributions to the inverse.

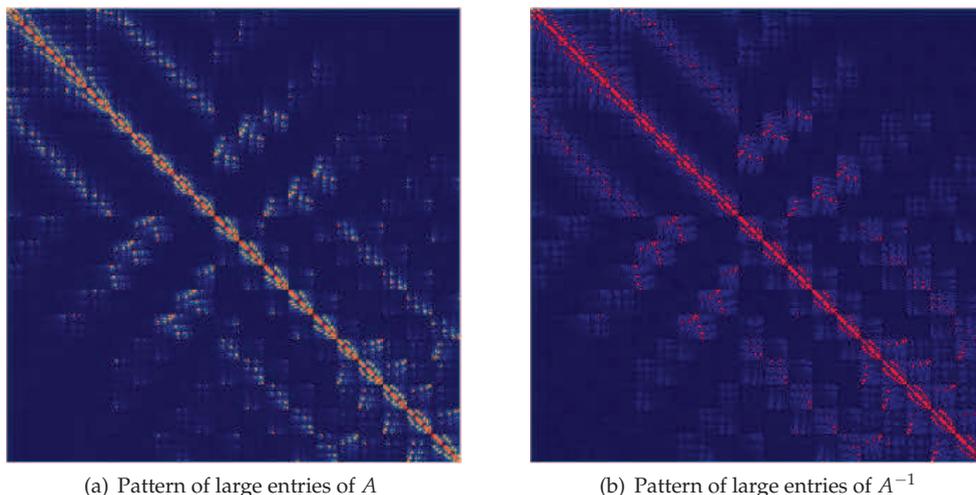


Fig. 6. Structure of the large entries of A (on the left) and of A^{-1} (on the right). Large to small entries are depicted in different colors, from red to green, yellow and blue. The test problem is a small sphere.

The actual entries of M may be computed by minimizing the error matrix $\|I - SM\|_F$ for right preconditioning ($\|I - MS\|_F$ resp. left preconditioning). The Frobenius-norm allows to decouple the constrained minimization problem into n independent linear least-squares problems, one for each column (resp. row) of M when preconditioning from the right (resp. from the left). The independence of the least-squares problems can be immediately seen from the identity

$$\|I - SM\|_F^2 = \sum_{j=1}^n \|e_j - Sm_{\bullet j}\|_2^2, \tag{7}$$

where e_j is the j th canonical unit vector and $m_{\bullet j}$ is the column vector representing the j th column of M . In the case of right preconditioning, the analogous relation

$$\|I - MS\|_F^2 = \|I - S^T M^T\|_F^2 = \sum_{j=1}^n \|e_j - S^T m_{j\bullet}\|_2^2 \tag{8}$$

holds, where $m_{j\bullet}$ is the column vector representing the j th row of M . The preconditioner is not guaranteed to be nonsingular in general, and additionally it does not preserve any possible symmetry of A . The condition to ensure non-singularity of M may be derived from the following estimates of the accuracy of the approximate inverse (Grote & Huckle (1997)):

THEOREM 1. Let $r_j = Sm_j - e_j$ be the residual associated with column m_j for $j = 1, 2, \dots, n$, and

$$q = \max_{1 \leq j \leq n} \left\{ \text{nnz}(r_j) \right\} \ll n. \text{ Suppose that } \|r_j\|_2 < t \text{ for } j = 1, 2, \dots, n, \text{ then we have}$$

$$\begin{aligned} \|SM - I\|_F &\leq \sqrt{nt}, & \|M - S^{-1}\|_F &\leq \|S^{-1}\|_2 \sqrt{nt}, \\ \|SM - I\|_2 &\leq \sqrt{nt}, & \|M - S^{-1}\|_2 &\leq \|S^{-1}\|_2 \sqrt{nt}, \\ \|SM - I\|_1 &\leq \sqrt{qt}, & \|M - S^{-1}\|_1 &\leq \|S^{-1}\|_1 \sqrt{qt}. \end{aligned}$$

■

Owing to this result, all the eigenvalues of SM lie in the disk centered in 1 and of radius \sqrt{qt} ; the value of q is not known a priori, though, so that one might enforce the condition $\sqrt{nt} < 1$ to prevent singularity or near-singularity of the preconditioned matrix. In practice it may be too costly to compute M with such a small t . For some problems, it may be observed a lack of robustness of the approximate inverse due to the clustering of small eigenvalues in the spectrum of the preconditioned matrix. Stabilization techniques based on eigenvalue deflation may be used to enhance the robustness of M , see e.g. Carpentieri et al. (2003).

The most critical component is the computation of the nonzero structure of M . From Figure 6, we see that the sparse pattern of S may be a suitable pattern for M . Denoting by

$$\mathcal{P} = \{ (i, j) \in [1, n]^2 \text{ s.t. } m_{ij} \neq 0 \}$$

the nonzero structure of the approximate inverse, we may automatically determine the pattern of the nonzero entries of the j th column of M as

$$\mathcal{C}_j = \{ i \in [1, n] \text{ s.t. } (i, j) \in \mathcal{P} \}.$$

and compute the associated entries by solving a small size dense least-squares problem. The least-squares solution involves only those columns of S indexed by \mathcal{C}_j ; we indicate this subset by $S(:, \mathcal{C}_j)$. Because S is sparse, many rows in $S(:, \mathcal{C}_j)$ are usually null, not affecting the solution of the least-squares problems (7). Thus denoting by \mathcal{R}_j the set of indices corresponding to the nonzero rows in $S(:, \mathcal{C}_j)$, by $\hat{S} = S(\mathcal{R}_j, \mathcal{C}_j)$, by $\hat{m}_j = m_j(\mathcal{C}_j)$, and by $\hat{e}_j = e_j(\mathcal{C}_j)$, the actual “reduced” least-squares problems to solve are

$$\min \|\hat{e}_j - \hat{S}\hat{m}_j\|_2, \quad j = 1, \dots, n. \quad (9)$$

Usually problems (9) have much smaller size than problems (7) and can be efficiently solved by dense QR factorization. The parallel implementation of the approximate inverse is highly scalable as shown in Table 3, while the numerical performance typically tend to deteriorate for increasing matrix size as can be seen in Table 4.

Approximate inverses may be also computed in factorized form as $M = \tilde{G}\tilde{Z}$, where $\tilde{G} \approx U^{-1}$ and $\tilde{Z} \approx U^{-1}$ are approximation of the inverse triangular factors of S , see for instance Alléon, Benzi & Giraud (1997); Chen (1998); Rahola (1998); Samant et al. (1996). One example of such preconditioner is the *AINV* method by Benzi et al. (1996), a sparse approximate inverse computed in factorized form by applying an incomplete biconjugation process to S and dropping small entries below a threshold in the inverse factors. However, disappointing results with factorized approximate inverses have been reported on this problem class, see e.g. Carpentieri et al. (2004). The reason of failure is that for many integral formulations like EFIE and CFIE, the inverse factors may be totally unstructured. In this case, selecting

n (procs)	Construction	Elapsed time
	time (sec)	precond (sec)
112908 (8)	513	0.39
221952 (16)	497	0.43
451632 (32)	509	0.48
900912 (64)	514	0.60

Table 3. Parallel scalability of the approximate inverse for solving large-scale boundary integral equations on a model problem.

$dof / freq$	FROB		GMRES(∞)		GMRES(120)	
	Density	Time	Iter	Time	Iter	Time
23676 / 1.3 Ghz	0.94	2m	438	20m	+2000	55m
104793 / 2.6 "	0.19	6m	234	20m	253	17m
419172 / 5.2 "	0.05	21m	413	2h 44m	571	2h 26m
943137 / 7.8 "	0.02	49m	454	3h 35m•	589	5h 55m

Table 4. Numerical scalability of the approximate inverse for solving large-scale boundary integral equations. The symbol • means run on 32 processors. Notation: m means minutes, h hours.

a priori the sparse pattern for the factors can be extremely hard and dynamic pattern selection strategies, that drop small entries below a user-defined threshold during the computation, may be very difficult to tune as they can easily discard relevant information and lead to a very poor preconditioner. For those problems, finding the appropriate threshold to enable a good trade-off between sparsity and numerical efficiency is challenging and very problem-dependent.

4.3 Incomplete LU factorization preconditioner

ILU -type methods compute an approximate triangular decomposition of S by means of an incomplete Gaussian elimination process. The ILU preconditioner writes as $M = \tilde{L}\tilde{U}$, $\tilde{L} \approx L$, $\tilde{U} \approx U$ where L and U denote respectively the lower and upper triangular factors of the standard LU factorization of S . This class of methods is virtually always used for solving sparse linear systems. However, mixed success is reported on dense matrix problems, due to the indefiniteness of the systems arising from the discretization. The root of the problem is that small pivots often appear during the factorization, leading to highly ill-conditioned triangular factors and unstable triangular solves (Carpentieri et al. (2004)).

In Table 5 we show an experiment with an ILU preconditioner computed from the sparse approximation S to A , using different values of density for S . The test case is a sphere of 1 meter length illuminated at 300 MHz; the problem is modeled using EFIE and the mesh is discretized with 2430 edges. The set \mathcal{F} of fill-in entries to be kept for the approximate lower triangular factor \tilde{L} is defined by

$$\mathcal{F} = \{ (k, i) \mid lev(l_{k,i}) \leq \ell \} ,$$

where the integer ℓ denotes a user specified maximal fill-in level. The level $lev(l_{k,i})$ of the coefficient $l_{k,i}$ of L is computed as follows:

Initialization

$$lev(l_{k,i}) = \begin{cases} 0 & \text{if } l_{k,i} \neq 0 \text{ or } k = i \\ \infty & \text{otherwise} \end{cases}$$

Factorization

$$lev(l_{k,i}) = \min \left\{ lev(l_{k,i}), lev(l_{i,j}) + lev(l_{k,j}) + 1 \right\}.$$

Observe that the larger ℓ , the higher the density of the preconditioner. We denote the resulting preconditioner by $ILU(\ell)$ Saad (1996).

In our results, increasing the fill-in parameter may produce much more robust preconditioners than $ILU(0)$ applied to a denser sparse approximation of the original matrix; $ILU(1)$ may deliver a good rate of convergence provided the coefficient matrix is not too sparse. However, the factorization of a very sparse approximation (up to 2%) of the coefficient matrix can be stable and accelerate significantly the convergence, especially if at least one level of fill-in is retained. Then, for higher values of the density of S the factors may become progressively ill-conditioned, the triangular solves unstable and consequently the preconditioner is useless. The table also shows that ill-conditioning of the factors is not related to ill-conditioning of A .

Density of $S = 2\%$				
$IC(\text{level})$	Density of L	$\kappa_\infty(L)$	GMRES(30)	GMRES(50)
$IC(0)$	2.0%	$2 \cdot 10^3$	378	245
$IC(1)$	5.1%	$1 \cdot 10^3$	79	68
$IC(2)$	9.1%	$9 \cdot 10^2$	58	48
Density of $S = 4\%$				
$IC(\text{level})$	Density of L	$\kappa_\infty(L)$	GMRES(30)	GMRES(50)
$IC(0)$	4.0%	$6 \cdot 10^9$	–	–
$IC(1)$	11.7%	$2 \cdot 10^5$	–	–
$IC(2)$	19.0%	$7 \cdot 10^3$	40	38
Density of $S = 6\%$				
$IC(\text{level})$	Density of L	$\kappa_\infty(L)$	GMRES(30)	GMRES(50)
$IC(0)$	6.0%	$8 \cdot 10^{11}$	–	–
$IC(1)$	18.8%	$5 \cdot 10^{11}$	–	–
$IC(2)$	29.6%	$7 \cdot 10^4$	–	–

Table 5. Number of iterations of GMRES varying the sparsity level of S and the level of fill-in of the approximate factor L on a spherical model problem ($n = 2430, \kappa_\infty(A) = \|A\|_\infty \|A^{-1}\|_\infty \approx \mathcal{O}(10^2)$). The symbol ‘-’ means that convergence was not obtained after 500 iterations.

A complex diagonal compensation can help to compute a more stable preconditioner, by shifting along the imaginary axis the eigenvalues close to zero in the spectrum of the coefficient matrix. However, the value of the shift is not easy to tune *a priori* and its effect

on the convergence is difficult to predict (Carpentieri et al. (2004)). Pivoting may be a more robust approach to overcome the problem according to reported experiment by Malas & Gürel (2007); in this case, the i th row of the factor is computed as soon as $perm_{tol} \times |s_{ij}| > |s_{ii}|$, where $perm_{tol}$ is the permutation tolerance and s_{ij} are the entries of S .

We follow a different approach. We report on experiments with multilevel inverse-based ILU factorization methods to possibly remedy numerical instabilities. Following Bollhöfer & Saad (2006), we initially rescale and reorder the initial matrix A as

$$P^T D_l A D_r Q = \hat{A}, \tag{10}$$

which yields $\hat{A}\hat{x} = \hat{b}$ for appropriate \hat{x}, \hat{b} . The initial step may consist of an optional maximum weight matching (Duff & Koster (1999)). By rescaling and a one-sided permutation, it attempts to improve the diagonal dominance. After that, a symmetric reordering is applied to reduce the fill-/bandwidth. The latter can also be used without an a priori matching step, only rescaling the entries and symmetrically permuting the rows and the columns. This is of particular interest for (almost) symmetrically structured problems. Next, an inverse-based ILU with static diagonal pivoting is computed. I.e., during the approximate incomplete factorization $\hat{A} \approx LDU$ such that L, U^H are unit lower triangular factors and D is block diagonal, the norms $\|L^{-1}\|, \|U^{-1}\|$ are estimated. If at factorization step l a prescribed bound κ is exceeded, the current row l and column l are permuted to the lower right end of the matrix. Otherwise the approximate factorization is continued. One single pass leads to an approximate partial factorization

$$\Pi^T \hat{A} \Pi = \begin{pmatrix} B & F \\ E & C \end{pmatrix} \approx \begin{pmatrix} L_B & 0 \\ L_E & I \end{pmatrix} \begin{pmatrix} D_B & 0 \\ 0 & S_C \end{pmatrix} \begin{pmatrix} U_B & U_F \\ 0 & I \end{pmatrix} \equiv L_1 D_1 U_1, \tag{11}$$

with a suitable leading block B and a suitable permutation matrix Π , where $\|L_1^{-1}\| \leq \kappa, \|U_1^{-1}\| \leq \kappa$. The remaining system S_C approximates $C - EB^{-1}F$. From the relations

$$\begin{cases} B\hat{x}_1 + F\hat{x}_2 = \hat{b}_1 \\ E\hat{x}_1 + C\hat{x}_2 = \hat{b}_2 \end{cases} \Rightarrow \begin{cases} \hat{x}_1 = B^{-1}(\hat{b}_1 - F\hat{x}_2) \\ (C - EB^{-1}F)\hat{x}_2 = \hat{b}_2 - EB^{-1}\hat{b}_1 \end{cases} ,$$

at each step of an iterative solver we need to store and invert only blocks with B and $S_C \approx C - EB^{-1}F$ while for reasons of memory efficiency, L_E, U_F are discarded and implicitly represented via $L_E \approx EU_B^{-1}D_B^{-1}$ (resp. $U_F \approx D_B^{-1}L_B^{-1}F$). When the scaling, reordering and the factorization is successively applied to S_C , a multilevel variant of (10) is computed. E.g., after a one additional level we obtain

$$\tilde{P}\tilde{D}_l A \tilde{D}_r \tilde{Q} = \left[\begin{array}{c|c|c} B & F_1 & F_2 \\ \hline E_1 & C_{11} & C_{12} \\ \hline E_2 & C_{21} & C_{22} \end{array} \right] \approx \left[\begin{array}{c|c|c} L_B & 0 & 0 \\ \hline L_{E_1} & I & 0 \\ \hline L_{E_2} & L_{C_{21}} & I \end{array} \right] \left[\begin{array}{c|c|c} D_B & 0 & 0 \\ \hline 0 & D_{C_{11}} & 0 \\ \hline 0 & 0 & S_{22} \end{array} \right] \left[\begin{array}{c|c|c} U_B & U_{F_1} & U_{F_2} \\ \hline 0 & I & U_{C_{12}} \\ \hline 0 & 0 & I \end{array} \right].$$

The multilevel algorithm ends at some step m when either S_C is factored completely or it becomes considerably dense and switches to a dense LAPACK solver. After computing an m -step ILU decomposition, for preconditioning we have to apply $L_m^{-1}AU_m^{-1}$. From the error equation $E_m = A - L_m D_m U_m$, we see that $\|L_m^{-1}\|$ and $\|U_m^{-1}\|$ contribute to the inverse error $L_m^{-1}E_m U_m^{-1}$. Monitoring the growth of these two quantities during the partial factorization

is essential to preserve the numerical stability of the solver, as can be observed comparing results in Table 5 and Table 6.

Density of $S = 2\%$			
threshold	Density of L	GMRES(30)	GMRES(50)
1.0e-3	0.30	29	29
Density of $S = 4\%$			
MILU	Density of L	GMRES(30)	GMRES(50)
1.0e-3	0.39	26	26
Density of $S = 6\%$			
MILU	Density of L	GMRES(30)	GMRES(50)
1.0e-3	0.46	24	24

Table 6. Number of iterations of GMRES using a multilevel inverse-based ILU factorization as preconditioner. The model problem is the same as in Table 5.

5. Concluding remarks

We have discussed some fast iterative solution techniques for solving surface boundary integral equations. High-frequency simulations of large structures are extremely demanding for scalable solvers and large computing resources. We have reviewed recent advances for the class of Krylov subspace methods, sparse approximate inverses, incomplete LU factorizations.

Other approach have been applied in this area of research. Multigrid methods are provably optimal algorithms for solving various classes of partial differential equations. Attempts to apply these techniques to dense linear systems have obtained mixed success. Early experiments on boundary element equations are reported with geometric versions on simple model problems, typically the hypersingular and single-layer potential integral operators arising from the Laplace equation (Bramble et al. (1994); Petersdorff & Stephan (1992); Rjasanow (1987)). Multigrids require a hierarchy of nested meshes to setup the principal components of the algorithm, *i.e.* a coarsening strategy to decrease the number of unknowns, grid transfer operators to move from a grid to another one, coarse grid operators and smoothing procedure, see e.g. Hackbusch (1985). Thus they are difficult to implement. On the other hand, algebraic multigrid algorithms use only single grid information extracted from either the graph or the entries of the coefficient matrix and are nearly as effective as geometric algorithms in reducing the number of iterations, see e.g Braess (1995); Brandt (1999); Ruge & Stüben (1987); Vanek et al. (1996). Langer et al. propose to apply an auxiliary sparse matrix reflecting the local topology of the mesh on the fine grid to setup all the components of the multigrid algorithm in a purely algebraic setting (Langer et al. (2003)). This *gray-box* approach is fairly robust on model problems and maintains the algorithmic and memory complexity of the M-V product operation (Langer & Pusch (2005)), thus it is well suited to be combined with MLFMA. See also Carpentieri et al. (2007) for another multigrid-type solver.

Preconditioners based on wavelet techniques are also receiving interest. The wavelet compression of integral operators with smooth kernels yields nearly sparse matrices with at most $\mathcal{O}(n \log^a n)$ nonzero entries, where a is a small constant that depends on the operator and the wavelet used, see e.g. earlier work by Beylkin et al. (1991); Dahmen et al. (1993); Harbrecht & Schneider (2004); Hawkins et al. (2007); Lage & Schwab (1999). The compressed

matrix is spectrally equivalent to the original matrix and preconditioning is often needed (Chan & Chen (2000; 2002); Chan et al. (1997); Chen (1999); Ford & Chen (2001); Hawkins & Chen (2005); Hawkins et al. (2005)). Some efficient wavelet preconditioning algorithms have been proposed, based on bordered block structure (Ford & Chen (2001); Hawkins et al. (2005)), multi-level preconditioners (Chan & Chen (2002)), and sparse approximate inverses. However, most experiments with wavelet preconditioners are reported for model problems, e.g. Calderon-Zygmund type matrix, single and double layer potentials, the hyper-singular operator. For oscillatory kernels the compressed matrix may be fairly dense and wavelet techniques are less useful. For Helmholtz problems, wavelet Galerkin schemes yield matrices with approximately $\mathcal{O}(kn)$ (k is the wavenumber) which becomes $\mathcal{O}(n^2)$ when the number of unknowns is related to k .

Further investigations are necessary to identify the best class of methods for the given problem and the selected computer hardware. The use of more powerful (but also more complex) computing facilities should help in the search for additional speed, but it will also mean that there will be even more factors that need to be considered when attempting to identify the optimal approach in the future.

6. References

- Alléon, G., Amram, S., Durante, N., Homsı, P., Pogarieloff, D. & Farhat, C. (1997). Massively parallel processing boosts the solution of industrial electromagnetic problems: High performance out-of-core solution of complex dense systems, in M. Heath, V. Torczon, G. Astfalk, P. E. Björkstam, A. H. Karp, C. H. Koebel, V. Kumar, R. F. Lucas, L. T. Watson & E. D. E. Womble (eds), *Proceedings of the Eighth SIAM Conference on Parallel Computing*, SIAM Book, Philadelphia. Conference held in Minneapolis, Minnesota, USA.
- Alléon, G., Benzi, M. & Giraud, L. (1997). Sparse approximate inverse preconditioning for dense linear systems arising in computational electromagnetics, *Numerical Algorithms* 16: 1–15.
- Alpert, B., Beylkin, G., Coifman, R. & Rokhlin, V. (1993). Wavelet-like bases for the fast solution of second-kind integral equations, *SIAM J. Scientific and Statistical Computing* 14: 159–184.
- Antoine, X., Bendali, A. & Darbas, M. (2004). Analytic preconditioners for the electric field integral equation, *International Journal for Numerical Methods in Engineering* 61(8): 1310–1331.
- Bayliss, A., Goldstein, C. I. & Turkel, E. (1985). On accuracy conditions for the numerical computation of waves, *J. Comp. Phys.* 59: 396–404.
- Bebendorf, M. (2000). Approximation of boundary element matrices, *Numerische Mathematik* 86(4): 565–589.
- Bendali, A. (1984). *Approximation par elements finis de surface de problemes de diffraction des ondes electro-magnetiques*, PhD thesis, Université Paris VI.
- Benzi, M., Meyer, C. & Tuma, M. (1996). A sparse approximate inverse preconditioner for the conjugate gradient method., *SIAM J. Scientific Computing* 17: 1135–1149.
- Beylkin, G., Coifman, R. & Rokhlin, V. (1991). Fast wavelet transforms and numerical algorithms, *Comm. Pure Appl. Math.* 44: 141–183.
- Bilotti, F. & Vegni, C. (2003). MoM entire domain basis functions for convex polygonal patches, *Journal of Electromagnetic Waves and Applications* 17(11): 1519–1538.
- Bollhöfer, M. & Saad, Y. (2006). Multilevel preconditioners constructed from inverse-based ILUs, *SIAM J. Scientific Computing* 27(5): 1627–1650.

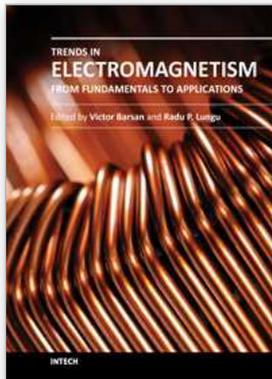
- Bond, D. & Vavasis, S. (1994). Fast wavelet transforms for matrices arising from boundary element methods, *Technical Report 94-174*, Cornell University Ithaca, NY, USA.
- Bonnet, P., Ferrieres, X., Grando, J., Alliot, J. & Fontaine, J. (1998). Frequency-domain finite volume method for electromagnetic scattering, *Antennas and Propagation Society International Symposium, 1998. IEEE* 1(21-26): 252–255.
- Botha, M. (2006). Solving the volume integral equations of electromagnetic scattering, *J. Comput. Phys.* 218(1): 141–158.
- Braess, D. (1995). Towards algebraic multigrid for elliptic problems of second order, *Computing* 55(4): 379–393.
- Bramble, J., Leyk, Z. & Pasciak, J. (1994). The analysis of multigrid algorithms for pseudodifferential operators of order minus one, *Math. Comput.* 63(208): 461–478.
- Brandt, A. (1999). General highly accurate algebraic coarsening, *Electronic Trans. Num. Anal* 10: 1–20.
- Canning, F. (1990). The impedance matrix localization (IML) method for moment-method calculations, *IEEE Antennas and Propagation Magazine* .
- Carpentieri, B., Duff, I. & Giraud, L. (2003). A class of spectral two-level preconditioners, *SIAM J. Scientific Computing* 25(2): 749–765.
- Carpentieri, B., Duff, I., Giraud, L. & Monga Made, M. M. (2004). Sparse symmetric preconditioners for dense linear systems in electromagnetism., *Numerical Linear Algebra with Applications* 11(8-9): 753–771.
- Carpentieri, B., Duff, I., Giraud, L. & Sylvand, G. (2005). Combining fast multipole techniques and an approximate inverse preconditioner for large electromagnetism calculations, *SIAM J. Scientific Computing* 27(3): 774–792.
- Carpentieri, B., Duff, I. S. & Giraud, L. (2000). Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism, *Numerical Linear Algebra with Applications* 7(7-8): 667–685.
- Carpentieri, B., Giraud, L. & Gratton, S. (2007). Additive and multiplicative two-level spectral preconditioning for general linear systems., *SIAM J. Scientific Computing* 29(4): 1593–1612.
- Carpentieri, B., Jing, Y.-F. & Huang, T.-Z. (2011). The BiCOR and CORS algorithms for solving nonsymmetric linear systems, *SIAM J. Scientific Computing* . In press.
- Chan, T. & Chen, K. (2000). Two-stage preconditioners using wavelet band splitting and sparse approximation, *UCLA CAM report CAM00-26*, Dept of Mathematics, UCLA, USA.
- Chan, T. & Chen, K. (2002). On two variants of an algebraic wavelet preconditioner, *SIAM Journal on Scientific Computing* 24(1): 260–283.
- Chan, T., Tang, W.-P. & Wan, W. (1997). Fast wavelet based sparse approximate inverse preconditioner, *BIT* 37(3): 644–660.
- Chen, K. (1994). Efficient iterative solution of linear systems from discretizing singular integral equations, *Elec. Trans. Numer. Anal.* 2: 76–91.
- Chen, K. (1998). On a class of preconditioning methods for dense linear systems from boundary elements., *SIAM J. Scientific Computing* 20(2): 684–698.
- Chen, K. (1999). Discrete wavelet transforms accelerated sparse preconditioners for dense boundary element systems, *Electronic Transactions on Numerical Analysis* 8: 138–153.
- Chen, K. (2005). *Matrix Preconditioning Techniques and Applications*, Cambridge University Press.
- Cheng, H., Crutchfield, W., Gimbutas, Z., Greengard, L., Ethridge, J., Huang, J., Rokhlin, V., Yarvin, N. & Zhao, J. (2006). A wideband fast multipole method for the Helmholtz equation in three dimensions, *J. Comput. Phys.* 216(1): 300–325.

- Chew, W. C., Jin, J. M., Lu, C. C., Michielssen, E. & Song, J. M. (1997). Fast solution methods in electromagnetics, *IEEE Transactions on Antennas and Propagation* 45(3): 533–543.
- Chew, W. & Wang, Y. (1993). A recursive T-matrix approach for the solution of electromagnetic scattering by many spheres, *IEEE Transactions on Antennas and Propagation* 41(12): 1633–1639.
- Christiansen, S. & Nédélec, J.-C. (2002). A preconditioner for the electric field integral equation based on Calderon formulas, *SIAM J. Numer. Anal.* 40(3): 1100–1135.
- Cosnau, A. (1996). Etude d'un préconditionneur pour les matrices complexes dense symmetric issues des équations de Maxwell en formulation intégrale., *Note technique ONERA, ONERA*. 142328.96/DI/MT.
- Dahmen, W., Prössdorf, S. & Schneider, R. (1993). Wavelet approximation methods for pseudodifferential equations ii: Matrix compression and fast solution, *Advances in Computational Mathematics* 1(3): 259–335.
- Darrigrand, E. (2002). Coupling of fast multipole method and microlocal discretization for the 3D Helmholtz equation, *J. Comput. Phys.* 181(1): 126–154.
- Darve, E. & Havé, P. (2004). Efficient fast multipole method for low-frequency scattering, *J. Comput. Phys.* 197(1): 341–363.
- Dembart, B. & Epton, M. (1994). A 3D fast multipole method for electromagnetics with multiple levels, *Tech. Rep. ISSTECH-97-004*, The Boeing Company, Seattle, WA.
- Duff, I. S. & Koster, J. (1999). The design and use of algorithms for permuting large entries to the diagonal of sparse matrices, *SIAM J. Matrix Analysis and Applications* 20(4): 889–901.
- Engheta, N., Murphy, W., Rokhlin, V. & Vassiliou, M. (1992). The fast multipole method (FMM) for electromagnetic scattering problems, *IEEE Transactions on Antennas and Propagation* 40(6): 634–641.
- Ergül, Ö. & Gürel, L. (2007). Fast and accurate solutions of extremely large integral-equation problems discretized with tens of millions of unknowns, *Electron. Lett.* 43(9): 499–500.
- Ergül, Ö. & Gürel, L. (2008). Efficient parallelization of the multilevel fast multipole algorithm for the solution of large-scale scattering problems, *IEEE Transactions on Antennas and Propagation* 56(8): 2335–2345.
- Ford, J. & Chen, K. (2001). Wavelet-based preconditioners for dense matrices with non-smooth local features, *BIT* 41(2): 282–307.
- Forsman, K., Gropp, W., Kettunen, L. & Levine, D. (1995). Volume integral equations in nonlinear 3D magnetostatics, *International Journal of Numerical Methods in Engineering* 38: 2655–2675.
- Yokota, R., Bardhan, J. P., Knepley, M. G., Barba, L.A. & Hamada, T. (2011). Biomolecular electrostatics using a fast multipole BEM on up to 512 gpus and a billion unknowns, *Computer Physics Communications* 182(6): 1272–1283.
- Freund, R. W. & Nachtigal, N. M. (1994). An implementation of the QMR method based on coupled two-term recurrences., *SIAM J. Scientific Computing* 15(2): 313–337.
- Gibson, W. (2008). *The method of moments in electromagnetics.*, Boca Raton, FL: Chapman and Hall/CRC.
- Greengard, L. & Rokhlin, V. (1987). A fast algorithm for particle simulations, *Journal of Computational Physics* 73: 325–348.
- Grote, M. & Huckle, T. (1997). Parallel preconditionings with sparse approximate inverses., *SIAM J. Scientific Computing* 18: 838–853.
- Hackbusch, W. (1985). *Multigrid Methods and Applications*, Springer-Verlag, Berlin.
- Hackbush, W. (1999). A sparse matrix arithmetic based on \mathcal{H} -matrices, *Computing* 62(2): 89–108.

- Hackbush, W. & Nowak, Z. (1989). On the fast matrix multiplication in the boundary element method by panel clustering, *Numerische Mathematik* 54(4): 463–491.
- Harbrecht, H. & Schneider, R. (2004). Wavelet based fast solution of boundary integral equations, in N. C. et al. (ed.), *Abstract and Applied Analysis*, World Scientific Publishing Company, pp. 139–162. Proceedings of the International Conference in Hanoi, 2002.
- Hawkins, S. & Chen, K. (2005). An implicit wavelet sparse approximate inverse preconditioner, *SIAM J. Scientific Computing* 27(2): 667–686.
- Hawkins, S., Chen, K. & Harris, P. (2005). An operator splitting preconditioner for matrices arising from a wavelet boundary element method for the Helmholtz equation, *International Journal of Wavelets, Multiresolution and Information Processing* 3 (4): 601–620.
- Hawkins, S., Chen, K. & Harris, P. (2007). On the influence of the wavenumber on compression in a wavelet boundary element method for the Helmholtz equation, *International Journal of Numerical Analysis and Modeling* 4(1): 48–63.
- Ipsen, I. C. F. & Meyer, C. D. (1998). The idea behind Krylov methods, *The American Mathematical Monthly* 105(10): 889–899.
- Jing, Y.-F., Carpentieri, B. & Huang, T.-Z. (2009). Experiments with Lanczos biconjugate A-orthonormalization methods for MoM discretizations of Maxwell's equations, *Progress In Electromagnetics Research, PIER* 99 pp. 427–451.
- Jing, Y.-F., Huang, T.-Z., Zhang, Y., Li, L., Cheng, G.-H., Ren, Z.-G., Duan, Y., Sogabe, T. & Carpentieri, B. (2009). Lanczos-type variants of the COCR method for complex nonsymmetric linear systems, *Journal of Computational Physics* 228(17): 6376–6394.
- Jr., W. S. (1994). Errors due to spatial discretization and numerical precision in the finite-element method, *IEEE Trans. Ant. Prop.* 42(11): 1565–1569.
- Kolotilina, L. Y. (1988). Explicit preconditioning of systems of linear algebraic equations with dense matrices., *J. Sov. Math.* 43: 2566–2573. English translation of a paper first published in *Zapiski Nauchnykh Seminarov Leningradskogo Otdeleniya Matematicheskogo im. V.A. Steklova AN SSSR* 154 (1986) 90-100.
- Kunz, K. & Luebbers, R. (1993). *The Finite Difference Time Domain Method for Electromagnetics*, CRC Press, Boca Raton.
- Lage, C. & Schwab, C. (1999). Wavelet Galerkin algorithms for boundary integral equations, *SIAM J. Scientific Computing* 20(6): 2195–2222.
- Langer, U. & Pusch, D. (2005). Data-sparse algebraic multigrid methods for large scale boundary element equations, *Applied Numerical Mathematics* 54(3-4): 406–424.
- Langer, U., Pusch, D. & Reitzinger, S. (2003). Efficient preconditioners for boundary element matrices based on grey-box algebraic multigrid methods, *International Journal for Numerical Methods in Engineering* 58(13): 1937–1953.
- Lee, R. & Cangellaris, A. (1992). A study of discretization error in the finite element approximation of wave solution, *IEEE Trans. Ant. Prop.* 40(5): 542–549.
- Lee, S., Ling, H. & Chou, R. (1988). Ray tube integration in shooting and bouncing ray method, *Micro. Opt. Tech. Lett.* 1: 285–289.
- Li, J., Li, L. & Gan, Y. (2005). Method of moments analysis of waveguide slot antennas using the EFIE, *Journal of Electromagnetic Waves and Applications* 19(13): 1729–1748.
- Malas, T. & Gürel, L. (2007). Incomplete LU preconditioning with multilevel fast multipole algorithm for electromagnetic scattering, *SIAM J. Scientific Computing* 29(4): 1476–1494.

- Michielssen, E. & Boag, A. (1996). A multilevel matrix decomposition algorithm for analyzing scattering from large structures, *IEEE Transactions on Antennas and Propagation* 44(8): 1086–1093.
- Petersdorff, T. v. & Stephan, E. (1992). Multigrid solvers and preconditioners for first kind integral equations, *Numer. Meth. for PDE* 8: 443–450.
- Rahola, J. (1998). Experiments on iterative methods and the fast multipole method in electromagnetic scattering calculations, *Technical Report TR/PA/98/49*, CERFACS, Toulouse, France.
- Rao, S., Wilton, D. & Glisson, A. (1982). Electromagnetic scattering by surfaces of arbitrary shape, *IEEE Trans. Antennas Propagat.* AP-30: 409–418.
- Rjasanow, S. (1987). Zweigittermethode für eine modellaufgabe bei bem-diskretisierung, *Wiss. Z. Tech. Univ. Karl-Marx-Stadt* 29(2): 230–235.
- Rokhlin, V. (1990). Rapid solution of integral equations of scattering theory in two dimensions, *J. Comp. Phys.* 86(2): 414–439.
- Ruge, J. & Stüben, K. (1987). *Multigrid Methods, Frontiers in Applied Mathematics 3*, S.F. McCormick ed., SIAM, Philadelphia, PA, chapter Algebraic multigrid (AMG), pp. 73–130.
- Saad, Y. (1996). *Iterative Methods for Sparse Linear Systems.*, PWS Publishing, New York.
- Saad, Y. & Schultz, M. (1986). GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems., *SIAM J. Scientific and Statistical Computing* 7: 856–869.
- Samant, A., Michielssen, E. & Saylor, P. (1996). Approximate inverse based preconditioners for 2D dense matrix problems, *Technical Report CCEM-11-96*, University of Illinois.
- Silvester, P. & Ferrari, R. (1990). *Finite Elements for Electrical Engineers*, Cambridge University Press, Cambridge.
- Song, J. & Chew, W. (1995). Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering, *Mico. Opt. Tech. Lett.* 10(1).
- Song, J. & Chew, W. (1998). The Fast Illinois Solver Code: Requirements and scaling properties, *IEEE Computational Science and Engineering* 5(3): 19–23.
- Song, J., Lu, C.-C. & Chew, W. (1997). Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects, *IEEE Transactions on Antennas and Propagation* 45(10): 1488–1493.
- Song, J., Lu, C., Chew, W. & Lee, S. (1998). Fast illinois solver code (FISC), *IEEE Antennas and Propagation Magazine* 40(3): 27–34.
- Steinbach, O. & Wendland, W. (1998). The construction of some efficient preconditioners in the boundary element method., *Adv. Comput. Math.* 9(1-2): 191–216.
- Sylvand, G. (2002). *La Méthode Multipôle Rapide en Electromagnétisme : Performances, Parallélisation, Applications*, PhD thesis, Ecole Nationale des Ponts et Chaussées.
- Sylvand, G. (2003). Complex industrial computations in electromagnetism using the fast multipole method, in G. Cohen, E. Heikkola, P. Joly & P. Neittaanmaki (eds), *Mathematical and Numerical Aspects of Wave Propagation*, Springer, pp. 657–662. Proceedings of Waves 2003.
- Taflove, A. (1995). *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, Artech House, Boston.
- Tausch, J. (2004). The variable order fast multipole method for boundary integral equations of the second kind, *Computing* 72(3-4): 267–291.
- van der Vorst, H. (1992). Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems., *SIAM J. Scientific and Statistical Computing* 13: 631–644.

- Vanek, P., Mandel, J. & Brezina, M. (1996). Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, *Computing* 56: 179–196.
- Vavasis, S. (1992). Preconditioning for boundary integral equations, *SIAM J. Matrix Analysis and Applications* 13: 905–925.
- Volakis, J., Chatterjee, A. & Kempel, L. (1998). *Finite element methods for electromagnetics*, IEEE Press, Piscataway, NJ.



Trends in Electromagnetism - From Fundamentals to Applications

Edited by Dr. Victor Barsan

ISBN 978-953-51-0267-0

Hard cover, 290 pages

Publisher InTech

Published online 23, March, 2012

Published in print edition March, 2012

Among the branches of classical physics, electromagnetism is the domain which experiences the most spectacular development, both in its fundamental and practical aspects. The quantum corrections which generate non-linear terms of the standard Maxwell equations, their specific form in curved spaces, whose predictions can be confronted with the cosmic polarization rotation, or the topological model of electromagnetism, constructed with electromagnetic knots, are significant examples of recent theoretical developments. The similarities of the Sturm-Liouville problems in electromagnetism and quantum mechanics make possible deep analogies between the wave propagation in waveguides, ballistic electron movement in mesoscopic conductors and light propagation on optical fibers, facilitating a better understanding of these topics and fostering the transfer of techniques and results from one domain to another. Industrial applications, like magnetic refrigeration at room temperature or use of metamaterials for antenna couplers and covers, are of utmost practical interest. So, this book offers an interesting and useful reading for a broad category of specialists.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Bruno Carpentieri (2012). Fast Preconditioned Krylov Methods for Boundary Integral Equations in Electromagnetic Scattering, Trends in Electromagnetism - From Fundamentals to Applications, Dr. Victor Barsan (Ed.), ISBN: 978-953-51-0267-0, InTech, Available from: <http://www.intechopen.com/books/trends-in-electromagnetism-from-fundamentals-to-applications/fast-preconditioned-krylov-methods-for-solving-large-scale-boundary-integral-equations-in-electromag>

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821