

A Fast, High-Order Algorithm for the Solution of Surface Scattering Problems: Basic Implementation, Tests, and Applications¹

Oscar P. Bruno and Leonid A. Kunyansky

*Applied and Computational Mathematics, California Institute of Technology,
217-50, Pasadena, California 91125*

E-mail: bruno@acm.caltech.edu; leonk@acm.caltech.edu

Received September 21, 2000; revised December 18, 2000

We present a new algorithm for the numerical solution of problems of acoustic scattering by surfaces in three-dimensional space. This algorithm evaluates scattered fields through fast, high-order solution of the corresponding boundary integral equation. The high-order accuracy of our solver is achieved through use of *partitions of unity* together with *analytical* resolution of kernel singularities. The acceleration, in turn, results from use of a novel approach which, based on high-order “*two-face*” *equivalent source* approximations, reduces the evaluation of far interactions to evaluation of 3-D fast Fourier transforms (FFTs). This approach is faster and substantially more accurate, and it runs on dramatically smaller memories than other FFT and k -space methods. The present algorithm computes one matrix-vector multiplication in $\mathcal{O}(N^{6/5} \log N)$ to $\mathcal{O}(N^{4/3} \log N)$ operations, where N is the number of surface discretization points. The latter estimate applies to smooth surfaces, for which our high-order algorithm provides accurate solutions with small values of N ; the former, more favorable count is valid for highly complex surfaces requiring significant amounts of subwavelength sampling. Further, our approach exhibits super-algebraic convergence; it can be applied to smooth and nonsmooth scatterers, and it does not suffer from accuracy breakdowns of any kind. In this paper we introduce the main algorithmic components in our approach, and we demonstrate its performance with a variety of numerical results. In particular, we show that the present algorithm can evaluate accurately in a personal computer scattering from bodies of acoustical sizes of several hundreds. © 2001 Academic Press

¹ The U.S. Government’s right to retain a nonexclusive royalty-free license in and to the copyright covering this paper, for governmental purposes, is acknowledged.

INTRODUCTION

The calculation of electromagnetic scattering from large two-dimensional surfaces remains one of the most important and challenging problems in computational science. Roughly, these problems present difficulties, as they require accurate descriptions and manipulation of highly oscillatory functions. Scattering problems involving one-dimensional integrals have been efficiently treated by means of high-order integrators (including the exponentially accurate trapezoidal rule and other high-order schemes [13, 18, 22, 24]), which reduce dramatically the complexity necessary to meet a given accuracy requirement. Fast, high-order treatment of problems of scattering by two-dimensional surfaces requires much more delicate treatments, however [2, 3, 10, 11, 23–26, 28].

In this context, use of high-order integrators is necessary to guarantee accurate results but is generally not sufficient. In scattering problems requiring a large number N of discretization points, use of a simple-minded integration scheme would usually lead to inordinately long computing times. A number of fast algorithms for 3-D scattering have been introduced in the past two decades [2, 3, 11, 12, 23, 26, 28, 29]. These methods are considerably faster than classical nonaccelerated algorithms—they run in $\mathcal{O}(N \log N)$ to $\mathcal{O}(N^{3/2} \log N)$ operations in contrast with the $\mathcal{O}(N^2)$ operations required by nonaccelerated schemes—and thus, they allow for computations involving rather large scattering surfaces. None of the existing fast implementations exhibits high-order convergence, however. As a result, the error of such fast computations turns out to be on the order of a fraction of a decibel (or several percent) even for the simplest test scatterers (see [10, 28]).

In this paper we present a fast, high-order algorithm for the solution of problems of acoustic scattering from smooth surfaces in three dimensions (see also [5]). The present algorithm computes scattered fields in $\mathcal{O}(N^{6/5} \log N)$ to $\mathcal{O}(N^{4/3} \log N)$ operations. The latter estimate applies to smooth surfaces, for which our high-order algorithm provides accurate solutions with small values of N ; the former, more favorable count is valid for highly complex surfaces requiring significant amounts of subwavelength sampling. A variety of numerical experiments indicate that this algorithm performs exceptionally well and, in fact, that it yields, in competitive running times, accuracies considerably higher than those rendered by other methods. Explicit comparisons with a number of well-known solvers are provided in Section 6.

The present algorithm is the result of our attempts to generalize the methods of [8, 9] to problems of scattering by surfaces in space. In that work problems of scattering by two-dimensional heterogeneous bodies and the associated *planar surface integrals* were treated via a combination of spherical wave expansions for the free space Green function (the addition theorem) and the fast Fourier transform (FFT). Further, high-order integration was obtained by analytical resolution of singularities. The resulting fast high-order method is very accurate, and it can handle easily a wide variety of complex electrically large scatterers.

For a variety of reasons the required generalizations of these ideas to problems of surface scattering are not direct. On one hand, for general curved *geometries* a numerical curved-surface version of the addition theorem would be difficult to obtain. Further, the methods for analytical resolution of singularities used in [8, 9] cannot be applied directly here due, in part, to the topological and geometric characteristics of surfaces in three-dimensional space. Finally, use of in-surface FFTs, which can be incorporated without difficulty for the geometries treated in [8, 9], does not seem easy to implement in the cases considered

presently. All of these difficulties can be surmounted, however: the fast, high-order algorithm presented in this paper is based on analytical resolution of singularities and FFTs.

Our basic high-order integrator involves use of partitions of unity—to deal with topological characteristics of closed surfaces—and analytical resolution of singularities—to avoid costly refinement strategies. Use of this algorithm without acceleration would lead to the customary $\mathcal{O}(N^2)$ operation count (where N is the size of the surface grid). The constant of proportionality in this complexity estimate is rather small, however, so that even without acceleration the present high-order integrator is an efficient solver for small- to medium-sized problems; see Section 6.2 for comparisons with other high-order integrators. For large problems, however, use of acceleration is imperative.

Two well-known approaches to acceleration have been available for a number of years: the fast multipole method (FMM) [12, 25, 26, 28, 29] and a broad class of FFT accelerated techniques and k -space methods [2, 3, 11, 23]. FMM-based algorithms provide considerable acceleration: they run in as little as $\mathcal{O}(N \log N)$ operations per iteration. However, to the best of our knowledge, high-order accuracy has not been demonstrated in FMM computations of wave scattering. A possible explanation for this fact is that the FMM approach [12, 26] depends critically on certain mappings which contain multiplication by Hankel functions of high order. These operations are associated with a substantial amount of ill conditioning, which leads to accuracy limitations known as the “subwavelength breakdown problem” (see [14, p. 51; 15; 19; 21]). These instabilities may prevail and mask the asymptotic high-order convergence of any underlying high-order integrator, however accurate. The search for stable FMM solvers continues to this day [19], and the feasibility of such designs is yet to be demonstrated. In contrast, the FFT acceleration techniques are stable.

The accelerator we introduce is closely related to two of the most advanced FFT methods developed recently [2, 23]. An important element common to these two methods and our technique is a concept of equivalent (or auxiliary) sources, located on a subset of a 3-D Cartesian grid. In all three cases, the intensities of these sources are chosen to approximate the field radiated by the scatterer, which allows fast computation of the numerous “non-adjacent interactions” through the use of the 3-D FFT. Surface problems such as the ones we consider are treated in [2, 23] by means of equivalent sources located in a *volumetric* grid—in such a way that equivalent sources with nonzero intensities occupy *all Cartesian nodes adjacent to the scatterer*. Since the spacing of this Cartesian grid cannot be coarsened beyond some threshold, and, further, since the 3-D convolution should be performed throughout the whole volume occupied by the body, for surface problems such a scheme requires an $\mathcal{O}(N^{3/2})$ FFT. Therefore, previous FFT surface scattering solvers require $\mathcal{O}(N^{3/2})$ units of RAM and they run in $\mathcal{O}(N^{3/2} \log N)$ operations.

Our algorithm, in contrast, subdivides the volume occupied by the scatterer into a number of (relatively large) cubic cells, and it places equivalent sources *on the faces* of those cells. As we will show, such a design significantly reduces the sizes of the required FFTs—to as few as $\mathcal{O}(N^{6/5})$ to $\mathcal{O}(N^{4/3})$ points—with proportional improvement in storage requirements and operation count. Further, it results in super-algebraic convergence of the equivalent source approximations *as the size of the scatterer is increased*. In view of its high-order character and its improved acceleration technique, the present algorithm can evaluate solutions to large scattering problems in short computing times, very accurately and with very small memory requirements; see Section 6.

Interestingly, the two main components of the present approach can be used independently by of each other. The acceleration method, on one hand, can be used in combination with

any discretization strategy, including Galerkin approaches such as the method of moments, or Nystrom approaches of any kind. This acceleration algorithm is itself extremely accurate, so that the overall accuracy of a given implementation is determined by the corresponding accuracy of the integration rule. Our high-order integrator, on the other hand, which leads to the customary $O(N^2)$ operation count when used without acceleration, can be combined with other acceleration techniques if desired. The fast high-order combination we use is very competitive, as is clearly demonstrated by the numerical results of Section 6, and by the corresponding error analysis presented in the companion paper [20], which, in what follows, is referred to as Part II.

1. MATHEMATICAL FORMULATION

As we have mentioned, the present implementation of our methods applies to acoustic scattering by a sound-soft obstacle: outside the obstacle D , the scattered field $\psi(\mathbf{r})$ arising from the incoming wave $\psi^i(\mathbf{r})$ is the unique radiating solution of the Helmholtz equation [13]

$$\Delta\psi(\mathbf{r}) + k^2\psi(\mathbf{r}) = 0, \quad \mathbf{r} \in \mathbb{R}^3 \setminus \bar{D}, \quad (1)$$

with the boundary condition

$$\psi(\mathbf{r}) = -\psi^i(\mathbf{r}), \quad \mathbf{r} \in \partial D. \quad (2)$$

Here k denotes the wavenumber (so that $\lambda = 2\pi/k$ is the wavelength), and $r = |\mathbf{r}|$. Extension of our methods to the corresponding electromagnetic Maxwell problem will be presented elsewhere.

A solution to the Helmholtz problem above can be obtained through an integral equation formulation including the acoustic single- and double-layer potentials,

$$(S\varphi)(\mathbf{r}) = \int_{\partial D} \Phi(\mathbf{r}, \mathbf{r}')\varphi(\mathbf{r}') ds(\mathbf{r}') \quad (3)$$

and

$$(K\varphi)(\mathbf{r}) = \int_{\partial D} \frac{\partial\Phi(\mathbf{r}, \mathbf{r}')}{\partial\nu(\mathbf{r}')} \varphi(\mathbf{r}') ds(\mathbf{r}'). \quad (4)$$

Here $\Phi(\mathbf{r}, \mathbf{r}') = e^{ik|\mathbf{r}-\mathbf{r}'|}/4\pi|\mathbf{r}-\mathbf{r}'|$ is the Green function for the Helmholtz equation, and $\nu(\mathbf{r}')$ is the external normal to the surface at point \mathbf{r}' . Explicitly, the scattered field can be obtained easily once the integral equation for the unknown density $\varphi(\mathbf{r})$,

$$\frac{1}{2}\varphi(\mathbf{r}) + (K\varphi)(\mathbf{r}) - i\gamma(S\varphi)(\mathbf{r}) = \psi^i(\mathbf{r}), \quad \mathbf{r} \in \partial D, \quad (5)$$

has been solved; see [13]. Here γ is an arbitrary positive constant; appropriate choices of this parameter can be very advantageous in practice—see Section 6.

Naturally, the possibility of producing fast and accurate solutions for our problems hinges on our ability to evaluate the integrals (3) and (4) accurately and efficiently. In attempting to develop such accurate and efficient integrators one faces two main difficulties, namely,

accurate evaluation of the singular *adjacent interactions*—without undue compromise of speed—and fast evaluation of the voluminous number of *nonadjacent interactions*—without compromise in accuracy.

The difficulties connected with the high-order evaluation of *adjacent interactions* are caused, mainly, by the singular nature of the integral kernels $\Phi(\mathbf{r}, \mathbf{r}')$ and $\partial\Phi(\mathbf{r}, \mathbf{r}')/\partial v(\mathbf{r}')$ at $\mathbf{r}' = \mathbf{r}$. While, certainly, the well-known strategy of “singularity subtraction” gives rise to bounded integrands, integration of such bounded functions by means of classical high-order methods does not exhibit high-order accuracy, since the subsequent derivatives of the integrand are themselves unbounded. Thus, specialized quadrature rules must be developed and used to achieve high-order-accurate integration. Our high-order adjacency integrator, which is described in detail in Section 3, is based on the use of partitions of unity and analytical resolution of singularities. This approach compares favorably with previous high-order methods: it does not require costly setup manipulations and it leads to substantially more accurate and faster numerics. Comparisons of the accuracies and timings produced by our method with those given by other high-order methods are presented in Section 6.

2. PARTITIONS OF UNITY AND DISCRETIZATION

In order to deal with topological characteristics of closed surfaces and the singular character of integrands we utilize a strategy based on local parametrizations together with *fixed* and *floating* partitions of unity (POU), as explained in what follows.

To describe the scattering surface ∂D we cover it by a number K of overlapping coordinate patches \mathcal{P}^k , $k = 1, \dots, K$ (local charts, in the language of differential geometry), such that

1. Each patch \mathcal{P}^k is an open set within ∂D for $k = 1, \dots, K$, and the sets \mathcal{P}^k , $k = 1, \dots, K$, cover ∂D (Fig. 1).
2. The set \mathcal{P}^k is the image of a coordinate open set \mathcal{H}^k , contained in the plane, via a smooth invertible parameterization

$$\mathbf{r}^k = \mathbf{r}^k(u^k, v^k) \text{ defined for } (u^k, v^k) \in \mathcal{H}^k, k = 1, \dots, K,$$

which admits a smooth inverse, and such that the vector product

$$\mathbf{V}^k = \mathbf{V}^k(u^k, v^k) = \frac{\partial \mathbf{r}^k}{\partial u^k} \times \frac{\partial \mathbf{r}^k}{\partial v^k} \quad (6)$$

is bounded away from zero in \mathcal{H}^k . We assume, as we may, that each one of the vectors \mathbf{V}^k is an outward normal, so that the outward unit normal on \mathcal{P}^k is given by

$$\mathbf{v}^k = \frac{\mathbf{V}^k}{|\mathbf{V}^k|}.$$

In practice, we find it advantageous to utilize as large patches as the geometry permits. This approach thus provides a compromise between a desirable global discretization and the impossibility of describing a complicated surface by a single equation.

We will utilize a partition of unity subordinated to this covering of ∂D , that is, a set of functions $\{w^k(\mathbf{r}), k = 1, \dots, K\}$, such that

1. w^k is defined, smooth, and nonnegative in ∂D , and it vanishes outside \mathcal{P}^k , and
2. $\sum_{k=1}^K w^k = 1$ throughout ∂D .

It is not difficult to construct numerically such a POU for a given surface. For efficiency it is preferable to use a POU with small derivatives, which can be arranged by allowing for substantial overlap of the patches \mathcal{P}^k . Use of this partition of unity, which will be referred to as the “fixed” POU, allows us to reduce integration over the surface to evaluation of integrals of smooth functions $\varphi^k(u^k, v^k)$ compactly supported in the planar sets \mathcal{H}^k multiplied by the singular kernels. Indeed, defining

$$\varphi^k(u^k, v^k) = \varphi(\mathbf{r}^k(u^k, v^k)) \cdot w^k(\mathbf{r}^k(u^k, v^k)) \quad (7)$$

we see that an integral over the entire surface ∂D evaluated at the point \mathbf{r} may be obtained as a sum of integrals over \mathcal{H}^k for $k = 1, \dots, K$:

$$(S\varphi)(\mathbf{r}) = \sum_{k=1}^K \int_{\mathcal{H}^k} \Phi(\mathbf{r}, \mathbf{r}^k(u^k, v^k)) \varphi^k(u^k, v^k) J_k(u^k, v^k) du^k dv^k, \quad (8)$$

$$(K\varphi)(\mathbf{r}) = \sum_{k=1}^K \int_{\mathcal{H}^k} \frac{\partial \Phi(\mathbf{r}, \mathbf{r}^k(u^k, v^k))}{\partial v(\mathbf{r}^k(u^k, v^k))} \varphi^k(u^k, v^k) J_k(u^k, v^k) du^k dv^k. \quad (9)$$

At this stage we may describe our discretized problem, which we obtain through discretization of the densities $\varphi^k(u^k, v^k)$ for $k = 1, \dots, K$. To do this we use a *Cartesian* set of nodes (u_ℓ^k, v_m^k) within \mathcal{H}^k , and we denote by $\varphi_{\ell,m}^k$ the corresponding nodal values of our approximation to the density: these values are the unknowns we seek to obtain. To compute approximations to the integrals (3) and (4) we need to prescribe a quadrature rule which, utilizing the values $\varphi_{\ell,m}^k$, evaluates the patch integrals in Eqs. (8) and (9) to high order.

To design our quadrature rule we first note that the high-order integration problem depends significantly on whether the integrand is singular within the integration domain. In detail, if \mathbf{r} lies outside \mathcal{P}^k then $\mathbf{r}' = \mathbf{r}^k(u^k, v^k)$ does not coincide with \mathbf{r} for $(u^k, v^k) \in \mathcal{H}^k$, and the kernels $\Phi(\mathbf{r}, \mathbf{r}')$ and $\partial \Phi(\mathbf{r}, \mathbf{r}')/\partial v(\mathbf{r}')$ remain nonsingular within the patch. In this case we see, further, that, because the densities $\varphi^k(u^k, v^k)$ vanish to high order at the boundary of the relevant integration domains, the functions to be integrated can actually be viewed as smooth periodic functions. However, for such functions the trapezoidal rule yields quadratures with super-algebraic convergence, and our integration problem is solved in this case. Interestingly, we will utilize trapezoidal rules throughout the paper to construct an efficient higher order integrator of the singular integrals as well—as we explain in what follows.

Let us thus consider the problem of evaluation of an integral over the patch \mathcal{P}^k for a point \mathbf{r} within \mathcal{P}^k . (Note that such a point can correspond to either a node (u_ℓ^k, v_m^k) of the coordinate grid on the patch \mathcal{P}^k or to some node (u_q^j, v_r^j) associated with a different patch \mathcal{P}^j which overlaps \mathcal{P}^k .) In this case we will use an infinitely smooth finitely supported function $\eta_{\mathbf{r}}(\mathbf{r}')$ to split our integration problem in the form

$$\int \dots ds(\mathbf{r}') = \int \dots (1 - \eta_{\mathbf{r}}(\mathbf{r}')) ds(\mathbf{r}') + \int \dots \eta_{\mathbf{r}}(\mathbf{r}') ds(\mathbf{r}'), \quad (10)$$

where $\eta_{\mathbf{r}}(\mathbf{r}') = 1$ for $|\mathbf{r} - \mathbf{r}'| < r_0$ and $\eta_{\mathbf{r}}(\mathbf{r}')$ vanishes for $|\mathbf{r} - \mathbf{r}'| \geq r_1$. The pair $(\eta_{\mathbf{r}}, 1 - \eta_{\mathbf{r}})$ is, of course, a partition of unity. One such POU needs to be constructed for each

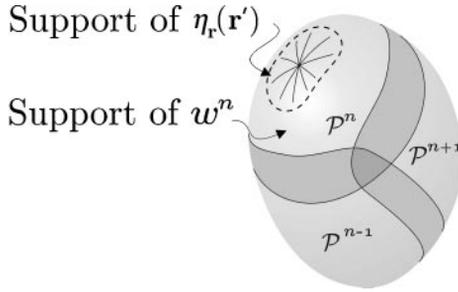


FIG. 1. POU covering and region of singular integration.

target point \mathbf{r}' : the collection of all such partitions of unity will be referred to as the *floating* POU. We see that the integrand containing the factor $(1 - \eta_r(\mathbf{r}'))$ is, again, a smooth periodic function, and the corresponding integral may be accurately evaluated by means of the trapezoidal rule; the integral containing the factor $\eta_r(\mathbf{r}')$, in turn, will be evaluated by means of a specialized high-order singular integrator constructed in Section 3.

(The floating POU allows us to restrict use of a specialized high-order singular integrator (see Section 3) to a neighborhood of each singular point—thus limiting the computational cost of such singular integrations and allowing for a fast overall integrator. The support of the floating POU will be chosen in such a way as to optimize the overall complexity of the algorithm while maintaining super-algebraic convergence; see Section 4. The choice $\eta = 1$ (large r_0 , no localization) gives rise to an $\mathcal{O}(N^2)$ integrator—which, as we will see, can itself be quite efficient for small- to medium-size problems.)

The POU scheme described in this section is depicted in Fig. 1. We thus see the surface is covered by large patches which have associated “fixed” partitions of unity. Varying the target point \mathbf{r} , in addition, we have a two-function “floating” partition of unity. The radial lines in Fig. 1 relate to the details of our singular high-order integrator, which we describe in the following section. The nonsingular integrals will be computed by means of trapezoidal rules, appropriately accelerated by means of the methods of Section 4.

3. SINGULAR INTEGRATION: VARYING POLAR COORDINATE SYSTEM

In this section we describe the high-order adjacency-integrator that we use for evaluation of the last integral in Eq. (10). In detail, on each patch \mathcal{P}^k we need to integrate numerically products of a smooth function $f_k(u', v') = \phi^k(u', v') J_k(u', v') \eta_r(\mathbf{r}'(u', v'))$ with kernels of each one of the following types:

$$\begin{aligned}
 G_1(\mathbf{R}) &= \frac{\cos k|\mathbf{R}|}{|\mathbf{R}|}, \\
 G_2(\mathbf{R}) &= \frac{\sin k|\mathbf{R}|}{|\mathbf{R}|}, \\
 G_3(\mathbf{R}) &= \left(k \cos k|\mathbf{R}| - \frac{\sin k|\mathbf{R}|}{|\mathbf{R}|} \right) \frac{\mathbf{R} \cdot \mathbf{v}(r)}{\mathbf{R}^2}, \\
 G_4(\mathbf{R}) &= \frac{\cos k|\mathbf{R}|}{|\mathbf{R}|} \frac{\mathbf{R} \cdot \mathbf{v}(r)}{\mathbf{R}^2}, \\
 G_5(\mathbf{R}) &= \sin k|\mathbf{R}| \frac{\mathbf{R} \cdot \mathbf{v}(r)}{2\mathbf{R}^2}.
 \end{aligned} \tag{11}$$

Here, for an arbitrary evaluation point $\mathbf{r} \in \partial D$ we have set $\mathbf{R} = \mathbf{r} - \mathbf{r}^k(u', v')$; the kernels $G_1(\mathbf{R})$ and $G_2(\mathbf{R})$ correspond to the single-layer potential, while $G_3(\mathbf{R})$ through $G_5(\mathbf{R})$ occur in the double-layer integral.

It is easy to design a high-order integration algorithm for the products of $f_k(u', v')$ with $G_2(\mathbf{r})$ and $G_3(\mathbf{r})$: these kernels are real analytic functions, and the function $f_k(u', v')$ is smooth and vanishes with all its derivatives for $|\mathbf{R}| \geq r_1$ due to the properties of $\eta_{\mathbf{r}}(\mathbf{r}'(u', v'))$. Thus, the trapezoidal rule provides a high-order quadrature rule in this case.

This approach is not appropriate, however, for any of the remaining kernels. Indeed, the kernels $G_1(\mathbf{r})$ and $G_4(\mathbf{r})$ are singular, while $G_5(\mathbf{r})$ is discontinuous at $|\mathbf{R}| = 0$. As we have mentioned, our approach to high-order quadrature of such kernels is based on *analytical resolution* of the associated singularities—that is, reduction of the singular integral to quadrature of smooth functions.

To obtain such analytical resolutions for integrals evaluated at a discretization point $(u, v) \in \mathcal{H}^k$ (where $\mathbf{r}^k(u, v) = \mathbf{r}$) we use a system of polar coordinates centered at (u, v) : $u' - u = \rho \cos \theta$, $v' - v = \rho \sin \theta$. (Again, \mathbf{r} may correspond to either a node $\mathbf{r}^k(u_\ell^k, v_m^k)$ of the coordinate grid on the current integration patch or to some node $\mathbf{r}^j(u_q^j, v_j^j)$ associated with a different patch \mathcal{P}^j which overlaps \mathcal{P}^k .)

In this system of coordinates the relevant integrals can be made to read

$$I_i(u, v) = \frac{1}{2} \int_0^{2\pi} L_i(u, v, \theta) d\theta, \quad i = 1, 4, 5, \quad (12)$$

where we have set

$$\begin{aligned} L_1(u, v, \theta) &= \int_{-r_1}^{r_1} f_k^*(\rho, \theta) \frac{|\rho|}{|\mathbf{R}|} \cos k|\mathbf{R}| d\rho, \\ L_4(u, v, \theta) &= \int_{-r_1}^{r_1} f_k^*(\rho, \theta) \frac{|\rho|}{|\mathbf{R}|} \cos k|\mathbf{R}| \frac{\mathbf{R} \cdot \mathbf{v}(r)}{\mathbf{R}^2} d\rho, \\ L_5(u, v, \theta) &= \int_{-r_1}^{r_1} f_k^*(\rho, \theta) \frac{|\rho|}{|\mathbf{R}|} |\mathbf{R}| \sin k|\mathbf{R}| \frac{\mathbf{R} \cdot \mathbf{v}(r)}{\mathbf{R}^2} d\rho, \end{aligned} \quad (13)$$

with

$$\begin{aligned} \mathbf{R} &= \mathbf{R}(\rho, \theta) = \mathbf{r}^*(\rho, \theta) - \mathbf{r}(u', v'), \\ \mathbf{r}^*(\rho, \theta) &= \mathbf{r}(u + \rho \cos \theta, v + \rho \sin \theta), \\ f_k^*(\rho, \theta) &= f_k(u + \rho \cos \theta, v + \rho \sin \theta). \end{aligned}$$

For an infinitely smooth surface the expressions

$$\frac{|\rho|}{|\mathbf{R}|}, \quad (14)$$

$$\frac{\mathbf{R} \cdot \mathbf{v}(r)}{\mathbf{R}^2} \quad (15)$$

can be easily shown to be infinitely smooth functions of ρ for any fixed direction θ . In the limit $\rho \rightarrow 0$ the values of these expressions can be found explicitly as

$$A \equiv \lim_{\rho \rightarrow 0} \frac{|\rho|}{|\mathbf{R}|} = |\mathbf{r}_u(u, v) \cos \theta + \mathbf{r}_v(u, v) \sin \theta|^{-1},$$

$$\lim_{\rho \rightarrow 0} \frac{\mathbf{R} \cdot \nu(r)}{\mathbf{R}^2} = -\frac{1}{2} \kappa_{\mathbf{n}}(u, v, \theta) = -\frac{A^2}{2} \left(\frac{d^2 \mathbf{r}^*(\rho, \theta)}{d\rho^2} \cdot \nu(r) \right),$$

where $\kappa_{\mathbf{n}}(u, v, \theta)$ is the curvature of the surface at the point $\mathbf{r}(u, v)$ in the direction $\mathbf{r}_u(u, v) \cos \theta + \mathbf{r}_v(u, v) \sin \theta$.

Since all the other factors in the integrands of (13) are smooth functions of ρ , and since the integrands vanish at the ends of the integration interval together with all their derivatives, the use of the trapezoidal rule provides radial quadratures of high-order accuracy. The trapezoidal rule also provides an appropriate high-order integrator for the angular integration of Eq. (12), since the corresponding integrands $L_i(u, v, \theta)$, $i = 1, 4, 5$, are themselves smooth periodic functions of θ . By symmetry, the range of the angular integration can be reduced from 2π to π .

(An additional aspect that needs to be considered in this context concerns cancellation errors. Indeed, since we utilize several parametrizations, it is possible for the quantity $|\mathbf{R}|$ to become very small while remaining nonzero. In this case, special care should be exercised to avoid cancellation errors when evaluating Eqs. (14) and (15); indeed, it is not hard to compute these expressions by means of appropriate polynomial interpolations and/or Taylor expansions—thus avoiding explicit calculation of differences.)

We have thus shown that use of polar coordinates provides an effective analytical resolution of the singular terms and that it allows high-order integration by means of the trapezoidal rule. The corresponding radial quadrature points, however, do not lie on the Cartesian grid associated with the given coordinate patch; see Fig. 2. Thus, use of an appropriate (fast) interpolation strategy is necessary for evaluation of the necessary function values at the radial integrations points. Efficiency is of utmost importance here, since we use one such polar coordinate transformation *at each target point* (u, v) . An efficient interpolation method can in fact be obtained—as described in what follows.

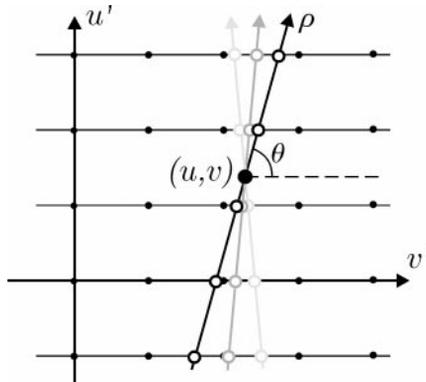


FIG. 2. Integration in a polar system; empty circles indicate discretization points for the integration with respect to ρ .

3.1. Efficient Interpolation Scheme

To obtain accurately interpolated values of the integrand $\varphi^k(u', v')$ at points lying on the lines $u' = u + \rho \cos \theta_\ell$, $v' = v + \rho \sin \theta_\ell$, $\theta_\ell = \pi \ell / n$, $\ell = 0, \dots, n-1$ from the corresponding values at the Cartesian nodes we proceed as follows. The integrand is given by its values at the nodes (u_i, v_j) , $u_i = i \Delta u$, $v_j = j \Delta v$, $i = -n_u, \dots, n_u$, $j = -n_v, \dots, n_v$. To evaluate an integral along the line $u' = u + \rho \cos \theta_\ell$, $v' = v + \rho \sin \theta_\ell$, for $\theta_\ell \in [\pi/4, 3\pi/4]$, we utilize a trapezoidal quadrature rule with the discretization step $\Delta \rho_\theta$ equal to $\Delta v / \sin \theta_\ell$, so that the integrand has to be evaluated at the points $\rho_k = k \Delta \rho_\theta - v / \sin \theta_\ell$. Note that the quadrature rule is chosen so that all of these points are situated on the straight lines $v_j = j \Delta v$, as shown in Fig. 2.

Since for each fixed value $v' = v_j = j \Delta v$ the function is known at equidistant points $u_i = i \Delta u$, a one-dimensional interpolation in the u direction suffices to provide all the required values. To speed up calculations while maintaining high-order accuracy, we use the following interpolation and approximation algorithm:

1. Obtain the Fourier coefficients of the given function for each one of the lines $v' = v_j = j \Delta v$ by means of the fast Fourier transform.
2. Evaluate the resulting Fourier series and their first derivatives on a much finer equispaced grid. These evaluations can be obtained efficiently, again, by means of the fast Fourier transform; in all numerical examples provided in the present text the spacing of the refined grid was 16 times smaller than that of the original grid.
3. On each interval of the refined grid construct an interpolating polynomial of degree 3 such that its values and the values of its first derivatives coincide with those of the trigonometric polynomial at the endpoints of the interval.

As a result of this procedure we obtain polynomial splines that closely approximate the interpolating Fourier series. The use of the fast Fourier transform makes the interpolation times negligible compared to that required by the other stages of the algorithm; evaluation of the splines at the required points requires few multiplications and additions per point and it is therefore very fast as well. (Clearly our interpolating algorithm, which is based on use of cubic interpolations, is fourth order accurate in the subgrid spacing. One could certainly use increasingly larger subgrids and Chebyshev interpolation to produce an interpolation technique of super-algebraic convergence. This is a matter of limited interest in practice, however. Indeed, in the cases we have considered, the cubic interpolation method with a 16-fold refinement described above matches the accuracy of the underlying trigonometric approximation to $\mathcal{O}(10^{-9})$, in computing times on the order of 1% of the time required by the overall computation. These accuracies are higher than those of interest in all of the problems we have treated, so that the use of more sophisticated interpolating techniques does not seem necessary.)

The radial integrations have thus far been restricted to lines determined by angles θ_ℓ in the interval $[\pi/4, 3\pi/4]$. Integration over the lines corresponding to the complementary set of angles $\theta_\ell \in [0, \pi/4] \cap [3\pi/4, \pi]$ can be performed similarly; in this case interpolations along the lines $u_i = i \Delta u$ should be used.

4. EVALUATION OF NONSINGULAR INTERACTIONS

Unlike the singular quadrature problem, the evaluation of nonsingular interactions does not present challenges from the point of view of accuracy, since in this case all integrands

are smooth and, in view of our smooth patching strategy, the trapezoidal rule yields high-order accuracy (see [20] for a detailed error analysis). The main task here, however, relates to acceleration. Indeed, most interactions are of nonadjacency type, and a direct use of an integrator such as the trapezoidal rule would lead to a complete $\mathcal{O}(N^2)$ algorithm.

In the following section we introduce our precise concept of adjacency, which is central to our approach. According to this concept (which bears some similarities to the ones used in other accelerated methods such as the FMM [25] and the k -space method [23]), the nonsingular interactions are further classified into nonadjacent nonsingular interactions and adjacent nonsingular interactions. The evaluation of nonadjacent interactions is discussed in Sections 4.1–4.3 below. We note that the combination of the integrator of Section 3 with that of Sections 4.1–4.3 accounts for most of the surface interactions but not all of them. Our treatment of certain remaining “adjacent nonsingular” interactions is discussed in Section 4.4—where, in addition, we make specific choices with regard to the parameters defining the floating partition of unity of Sections 2 and 3 and thus complete the description of our integration algorithm.

4.1. Two-Face Equivalent Source Representations

Our acceleration strategy is based on certain distributions of “equivalent sources,” which we describe in what follows. We begin by considering a cube C of side A containing the given obstacle, which we then partition into a number L^3 of identical, nonoverlapping cubic cells c_i of side $H = A/L$, so that there are L cells along each edge of the cube. (For elongated obstacles a 3-D slab is preferable; for simplicity of presentation, however, we will limit our discussion to a covering by a cube.) We note that each one of the surface discretization points (also called “true sources” in what follows) is contained in one of the cells c_i ; typically, however, most cells contain no true sources in their interior. As we shall see in Section 4.3, it is necessary for our method to use cells c_i , which do not admit inner resonances—eigenfunctions of the Dirichlet Laplacian—for the given wavenumber k . This can be ensured easily by adjusting slightly the cube size A and the cell size H if necessary.

We seek to substitute the surface sources contained in a cubic cell c_i by certain “equivalent sources” on the faces of c_i , in such a way that the field produced by the c_i -equivalent sources coincides, to high-order accuracy, with the field generated by the c_i -true sources at all points in space that are “not adjacent” to c_i . More precisely, the approximation corresponding to a cell c_i will be valid outside the concentric cube \mathcal{S}_i of side $3H$, with exponentially small errors. As we show below, for computational efficiency it is favorable to use a sequence of three independent sets of equivalent sources, located on three corresponding sets of points Π^ℓ , $\ell = 1, 2, 3$. To define Π^ℓ , note that the union of all cell faces parallel to the coordinate plane $x_\ell = 0$ consists of a number L of parallel squares of side A . Upon choice of an equivalent-source step size h (see Remark 4 and Section 5), we place identical two-dimensional Cartesian grids of points on each one of these squares; Π^ℓ is then defined as the union of all these grids.

(We note that such distributions of equivalent sources, which by design are contained in three sets of “sparsely spaced” parallel planes, lead to significant benefits: On one hand they give rise to a high-order accelerator; see also Table VII and Part II [20]. On the other hand, sparsity leads to FFTs which are significantly smaller than those arising from fully volumetric equivalent-source distributions [2, 23]—and, thus, to substantially reduced memory requirements and faster numerics.)

The high-order equivalent-source representations we seek for nonadjacent fields result from consideration of spherical harmonics, Bessel and Hankel functions $Y_n^m(\cdot)$, $j_n(\cdot)$, and $h_n^{(1)}(\cdot)$, and the spherical wave series

$$u(\mathbf{r}) = ik \sum_{n=0}^{\infty} \sum_{m=-n}^n b_{n,m} h_n^{(1)}(k|\mathbf{r}|) Y_n^m(\mathbf{r}/|\mathbf{r}|), \quad (16)$$

which, for a field $u(\mathbf{r})$ generated by sources inside a sphere B_a of radius a , provides a convergent expansion of the field everywhere outside B_a . Roughly speaking, at a distance $D = |\mathbf{r}|$ from the center of B_a , the error in a truncated spherical wave expansion (16) containing the $n \leq n_t$ terms only,

$$u(\mathbf{r}) = ik \sum_{n=0}^{n_t} \sum_{m=-n}^n b_{n,m} h_n^{(1)}(k|\mathbf{r}|) Y_n^m(\mathbf{r}/|\mathbf{r}|), \quad (17)$$

is of order

$$\mathcal{O}\left(\left(\frac{a}{D}\right)^{n_t}\right) \quad (18)$$

for sufficiently large n_t ; see Remark 4. This estimate, which follows from the asymptotic properties of the special functions and consideration of the addition theorem [13],

$$\frac{e^{ik|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|} = ik \sum_{n=0}^{\infty} \sum_{m=-n}^n h_n^{(1)}(k|\mathbf{r}|) Y_n^m(\mathbf{r}/|\mathbf{r}|) j_n(k|\mathbf{r}'|) \overline{Y_n^m(\mathbf{r}'/|\mathbf{r}'|)}, \quad (19)$$

is discussed in detail in [20].

It follows from these considerations that, at a distance D from a given cell c_i , the fields generated by the c_i -surface and c_i -equivalent sources can be approximated with a prescribed accuracy by truncated spherical wave expansions of the form (17) with an appropriate value of n_t . However, these expansions converge slowly in at least some of the 26 cells surrounding c_i . This leads to our precise concept of adjacency: two points \mathbf{r} and \mathbf{r}' in space are considered to be adjacent if they both lie within the same cell c_i or if one lies within a cell c_i and the other lies in any of the 26 neighboring cells. Equivalently, a point $\mathbf{r} \in c_i$ is adjacent to every point $\mathbf{r}' \in \mathcal{S}_i$ and nonadjacent to every point $\mathbf{r}' \notin \mathcal{S}_i$.

We can now describe the specifics of our implementation. For a fixed value of $\ell = 1, 2, 3$, we associate to each cell c_i -equivalent sources (monopoles $\xi_{i,j}^{(m)\ell} \Phi(\mathbf{r}, \mathbf{r}_{i,j}^\ell)$ and dipoles $\xi_{i,j}^{(d)\ell} \partial \Phi(\mathbf{r}, \mathbf{r}_{i,j}^\ell) / \partial x_\ell$ of intensity $\xi_{i,j}^{(m)\ell}$ and $\xi_{i,j}^{(d)\ell}$, respectively), which are placed at points $\mathbf{r}_{i,j}^\ell$, $j = 1, \dots, M^{equiv}$, contained within certain subsets Π_i^ℓ of Π^ℓ . In detail, the set Π_i^ℓ consists of all the points in Π^ℓ which lie within the union of two circular domains concentric with (and containing) the faces of c_i in Π^ℓ , as shown in Fig. 3. The radius of these domains is chosen to be equal to—or slightly larger than—the length of the diagonals of the faces; experimentally it was found that increase of the radius up to 40% leads to somewhat higher accuracy and relatively small increases in computing times.

Remark 1. As shown in Part II [20], use of a combined set of monopoles and dipoles does indeed suffice to provide convergent approximations from the types of two-face representations we use. Two-face distributions of one type, either monopole or dipole type, do not give rise to such convergent approximations, however.

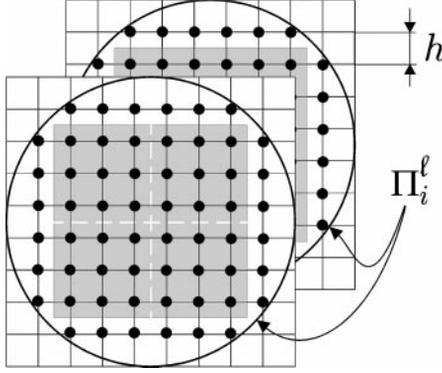


FIG. 3. Locations of the equivalent sources (black circles); gray squares indicate faces of a cell c_i .

Like the field $\psi^{c_i, true}$ radiated by the c_i -true sources, the field $\psi^{c_i, eq}$ radiated by the c_i -equivalent sources

$$\psi^{c_i, eq}(\mathbf{r}) = \sum_{j=1}^{\frac{1}{2}M^{equiv}} \left(\xi_{i,j}^{(m)\ell} \Phi(\mathbf{r}, \mathbf{r}_{i,j}^\ell) + \xi_{i,j}^{(d)\ell} \frac{\partial \Phi(\mathbf{r}, \mathbf{r}')}{\partial x'_\ell} \Big|_{\mathbf{r}'=\mathbf{r}_{i,j}^\ell} \right) \quad (20)$$

can be expressed through an expansion of the form (16). Let us take a value n_t of the truncation parameter such that the truncated expansion (17) corresponding to $\psi^{c_i, true}$ approximates this quantity to within the prescribed error tolerance. We thus see that $\psi^{c_i, true}$ can be approximated by $\psi^{c_i, eq}$ at points nonadjacent to c_i and to within a given tolerance $\mathcal{O}(\varepsilon)$ if and only if the intensities of the equivalent sources can be selected in such a way that:

1. The coefficients of the spherical wave expansion (16) for $\psi^{c_i, true}$ are approximated “well” by the corresponding expansion coefficients for $\psi^{c_i, eq}$ for all $n \leq n_t$. More precisely, the truncated spherical wave expansion of order n_t for these two fields should differ in no more than $\mathcal{O}(\varepsilon)$ everywhere outside \mathcal{S}_i

2. The n_t -tail (i.e., the sums for $n > n_t$) of the spherical wave expansion for $\psi^{c_i, eq}$ should be of $\mathcal{O}(\varepsilon)$ or less everywhere outside \mathcal{S}_i .

In practice, it was found that these two conditions can be guaranteed to hold provided

$$M^{equiv} \gtrsim n_t^2 \quad (21)$$

equivalent sources are used, and provided the intensities are chosen so that the vector formed by the differences $(\psi^{c_i, eq}(\mathbf{r}) - \psi^{c_i, true}(\mathbf{r}))$ (as \mathbf{r} varies over a number $n^{coll} \approx 2M^{equiv}$ collocation points on $\partial\mathcal{S}_i$) is minimized in the mean-square norm. Thus, the intensities are obtained as the least-squares solution of an overdetermined linear system of the form

$$\mathbf{A}\xi = \mathbf{b}, \quad (22)$$

where \mathbf{A} is an $n^{coll} \times M^{equiv}$ matrix. Note that this prescription does not require explicit use of the spherical wave expansions (16), and thus it completely bypasses costly evaluations of special functions.

Remark 2. Since all the cells c_i are identical, the matrix \mathbf{A} above is one and the same for all the c_i 's. Thus, the QR decomposition of \mathbf{A} (which we use to solve the least-squares problem under consideration) need only be computed once and stored for repeated use.

Remark 3. Our choice for the positions of the equivalent sources and similar choices for the corresponding positions of the n^{coll} collocation points can lead to substantial reductions in computation time. Our prescription in these regards takes advantage of certain symmetries which can be exploited to induce a block diagonal structure in the matrix \mathbf{A} , and, consequently, a significant eightfold reduction in the computational expense required to produce the intensities of the equivalent sources; see Appendix A.

Remark 4. The asymptotic regime of exponential decay (18) for the error in the approximation (17) is achieved for values of $n_t \geq 2kA/L$; that is, n_t must exceed 4π times the acoustical size of the cells (measured in λ); see Part II [20]. Thus, the parameter $M^{equiv} \geq n_t^2$ must exceed the value $(2kA/L)^2$ for the substitution to exhibit exponentially accurate approximations. As shown in Part II, the prescription $M^{equiv} = n_t^2$ with n_t given by

$$n_t = \max \left\{ \frac{2kA}{L}, -2 \frac{\log(\varepsilon)}{\log(3)} \right\} \quad (23)$$

gives rise to errors of $\mathcal{O}(\varepsilon)$ in the equivalent source approximations; the associated maximal step size $h = h^{max}$ in the Cartesian grids of Π^ℓ is then, approximately, $h^{max} \approx (2A)/(Ln_t)$.

Remark 5. The surface distributions of equivalent sources we use are not unrelated to those occurring in the Laplace solver of [1]. Indeed, that FMM algorithm uses distributions of equivalent sources on spherical surfaces to represent the fields generated by portions of the scattering surface. Such spherical arrangements are not suitable for use in conjunction with FFTs. Our equivalent sources, in contrast, like those arising in other k -space methods, are designed to work as part of FFT-based algorithms. Our two-face distributions of equivalent sources are useful in that they give rise to reduced FFTs and thus to substantially reduced computing times and memory requirements.

4.2. FFTs: Nonadjacent Field Values on a Sparse 3-D Grid

The methods introduced in the previous section allow us to identify the field produced by the portion of the scattering surface contained within a cell c_i with the corresponding field produced by the equivalent sources on Π_i^ℓ everywhere outside \mathcal{S}_i . In other words, recalling the definition of adjacency of Section 4.1, we define $\psi^{a,true}(\mathbf{r})$ and $\psi^{na,true}(\mathbf{r})$ as the fields induced at point \mathbf{r} by the adjacent and nonadjacent true sources, respectively, so that

$$\psi(\mathbf{r}) = \psi^{a,true}(\mathbf{r}) + \psi^{(na,true)}(\mathbf{r}). \quad (24)$$

Further, denoting by $\psi^{(na,eq)\ell}(\mathbf{r})$ the field induced at point \mathbf{r} by all the equivalent sources in Π^ℓ nonadjacent to \mathbf{r} ($\ell = 1, 2, 3$), we may write

$$\psi(\mathbf{r}) = \psi^{a,true}(\mathbf{r}) + \psi^{(na,eq)\ell}(\mathbf{r}) + \mathcal{O}(\varepsilon), \quad (25)$$

where the $\mathcal{O}(\varepsilon)$ term is the prescribed numerical tolerance in the equivalent source approximation; see Section 4.1.

The acceleration algorithm utilizes a quantity related to those occurring in Eq. (25), namely, the field $\psi^{(*)\ell}(\mathbf{r})$ produced at a grid point \mathbf{r} by all of the equivalent sources in Π^ℓ except the one at point \mathbf{r} —where the Green function is infinite. Using the notation

$$\Phi^*(\mathbf{r}, \mathbf{r}') = \begin{cases} \Phi(\mathbf{r}, \mathbf{r}') & \text{for } \mathbf{r} \neq \mathbf{r}' \\ 0 & \text{for } \mathbf{r} = \mathbf{r}' \end{cases}, \quad \Gamma^{(*)\ell}(\mathbf{r}, \mathbf{r}') = \begin{cases} \partial\Phi(\mathbf{r}, \mathbf{r}')/\partial x'_\ell & \text{for } \mathbf{r} \neq \mathbf{r}' \\ 0 & \text{for } \mathbf{r} = \mathbf{r}' \end{cases} \quad (26)$$

we have

$$\psi^{(*)\ell}(\mathbf{r}) = \sum_{\mathbf{r}' \in \Pi^\ell} (\xi_{\mathbf{r}'}^{(m)\ell} \Phi^*(\mathbf{r}, \mathbf{r}') + \xi_{\mathbf{r}'}^{(d)\ell} \Gamma^{(*)\ell}(\mathbf{r}, \mathbf{r}')) \quad (27)$$

(compare with (20)). Here the intensities $\xi_{\mathbf{r}'}^{(m)\ell}$ and $\xi_{\mathbf{r}'}^{(d)\ell}$ denote the sum of all intensities of equivalent sources located at point \mathbf{r}' :

$$\xi_{\mathbf{r}'}^{(m)\ell} = \sum_{\mathbf{r}''_{i,j}=\mathbf{r}'} \xi_{\mathbf{r}''_{i,j}}^{(m)\ell}, \quad \xi_{\mathbf{r}'}^{(d)\ell} = \sum_{\mathbf{r}''_{i,j}=\mathbf{r}'} \xi_{\mathbf{r}''_{i,j}}^{(d)\ell}. \quad (28)$$

We note that the functions $\psi^{(*)\ell}$ ($\ell = 1, 2, 3$) do not approximate any of the physical quantities under consideration since, at any given $\mathbf{r} \in \partial D$, the quantity $\psi^{(*)\ell}(\mathbf{r})$ contains only poor approximations of contributions from sites adjacent to \mathbf{r} . Subtraction of these poor approximations would then complete the evaluation of $\psi^{(na,eq)\ell}(\mathbf{r})$.

The importance of the quantities $\psi^{(*)\ell}$ lies, of course, in that, being exact convolutions in a Cartesian grid, they can be evaluated accurately and efficiently by means of the fast Fourier transform. Having obtained $\psi^{(*)\ell}$, one can proceed to compute $\psi^{(na,eq)\ell}$ by subtracting from $\psi^{(*)\ell}$ contributions of adjacent cells. Since the equivalent sources are located in the nodes of a Cartesian grid, the contributions from adjacent cells are given by (small) three-dimensional convolutions—which can be evaluated efficiently by means of FFTs. Upon subtraction, the values of $\psi^{(na,eq)\ell}$ result. These values provide good approximations for the field $\psi^{(na,true)\ell}$; using $\psi^{(na,eq)\ell}$ for $\ell = 1, 2, 3$ we thus obtain approximations for $\psi^{(na,true)\ell}$ throughout the boundary of each cell c_i . The evaluation of the surface values of such nonadjacent true fields is the subject of the following section.

4.3. Evaluation of the Surface Values of $\psi^{na,eq}$

Once $\psi^{(na,true)\ell}(\mathbf{r})$ is known for \mathbf{r} on the faces of a cell c_i , this function can be evaluated at points inside c_i as the solution of a Dirichlet problem. Such Dirichlet problems can be solved uniquely and in a stable manner since the size of the cells has been chosen so that internal resonances do not occur (Section 4.1); see also [20] for an error analysis in these regards.

In order to obtain the field inside a cell, we use a discretized plane wave expansion [16],

$$\psi^{(na,true)\ell}(\mathbf{r}) \approx \sum_{j=1}^{n^{wave}} \zeta_j \exp i k \mathbf{u}_j \mathbf{r}, \quad (29)$$

where \mathbf{u}_j are unit vectors defining directions of wave propagation, and $\zeta = (\zeta_1, \zeta_2, \dots, \zeta_{n^{wave}})$ is a vector of expansion coefficients. Since $\psi^{(na,true)\ell}(\mathbf{r})$ does not contain contributions

from cells adjacent to \mathbf{r} , $\psi^{na,eq}(\mathbf{r})$ this expansion converges exponentially fast with increasing number of wavevectors u_j ; see [20] for details. A number of other representations could be used in the solution of our Dirichlet problems, including an inner-field spherical wave expansion; the expansion (29) seems advantageous in that it does not require evaluation of special functions.

The choice of the unit vectors in the representation (29) is rather arbitrary: it is only necessary for the vectors \mathbf{u}_j to sample the surface of the unit sphere with a sufficient degree of uniformity. The prescription we use for the distribution of the \mathbf{u}_j 's leads to significant reductions in the operation count; see Appendix A for details. The coefficients ζ_j , on the other hand, are chosen so that the expansion (29) matches the field values on the boundary of c_i . As in Section 4.1, stability considerations require discretizations and plane waves to be set up in such a way that the associated linear algebra problem is overdetermined, in the present case by a factor of 1.5 or less—depending on the acoustical size of the cells c_i .

The matching procedure thus requires solution of an overdetermined system of linear equations

$$\mathbf{B}\zeta = \mathbf{d}, \quad (30)$$

where $\mathbf{d} = (d_1, d_2, \dots, d_{n^{face}})$ is a vector containing known values of $\psi^{(na,true)}$ at all the equivalent-source points $\mathbf{r}_{i,j}^\ell$ contained in the boundary of c_i , and \mathbf{B} is an $n^{face} \times n^{wave}$ matrix ($n^{face} > n^{wave}$). Note that not all of the $\frac{1}{2}M^{equiv}$ equivalent source locations in Π_i^ℓ are used here as matching points; see Fig. 3. This least-squares problem is solved by means of a QR decomposition; as in Section 4.1, since the geometry of all cells is identical the QR factorization of \mathbf{B} needs to be computed only once, as factors can be stored for repeated use.

4.4. Adjacent Nonsingular Interactions

The floating POU $\eta_{\mathbf{r}}$ of Section 2 was introduced as a means to reduce the domain in which the polar coordinate integration of Section 3 is operative—thus avoiding the substantial $\mathcal{O}(N^2)$ overall operation count which would otherwise be required—but the size of the support of $\eta_{\mathbf{r}}$ was left unspecified. The FFT acceleration technique of Sections 4.1–4.3, in turn, effectively accelerates the evaluation of nonadjacent contributions, that is, the contribution to a cell c_i from sources lying outside the cube \mathcal{S}_i . It is therefore reasonable to restrict the support of $\eta_{\mathbf{r}}$ to lie within \mathcal{S}_i for every $\mathbf{r} \in c_i$. Thus our prescription is: Taking r_1 to equal the cell size H , $\eta_{\mathbf{r}}(\mathbf{r}')$ is a C^∞ function of \mathbf{r}' , $\eta_{\mathbf{r}}(\mathbf{r}') = 1$ for $|\mathbf{r} - \mathbf{r}'| \leq r_0$ and $\eta_{\mathbf{r}}(\mathbf{r}')$ vanishes for $|\mathbf{r} - \mathbf{r}'| \geq r_1$. The specific value of r_0 and the particular function to be used have limited impact as long as the derivatives of $\eta_{\mathbf{r}}$ are not excessively large. In fact, it is advantageous to set up this floating POU in such a way that its spatial derivatives are as small as possible: a specific form we have used frequently is

$$\eta_{\mathbf{r}}(\mathbf{r}') = e^{\frac{2e^{-1/t}}{t-1}}, \quad t = |\mathbf{r} - \mathbf{r}'|/r_1.$$

To complete our integration algorithm we need to add the integral containing the term $(1 - \eta_{\mathbf{r}})$ in Eq. (10) over the portion of the scattering surface which lies within \mathcal{S}_i . To do this we simply add the corresponding point sources at the Cartesian discretization points in parameter space whose images lie inside \mathcal{S}_i . The overall procedure results in spectral accuracy, since together with the integral over the nonadjacent sources (given by the acceleration

scheme), this makes up the trapezoidal rule sum for the first integral on the right-hand side of Eq. (10), whose integrand is a smooth periodic function in (u, v) parameter space.

Sections 2–4 provide a full description of our integration algorithm for any given choice of the discretization parameters N , L , and M^{equiv} . Various rationales leading to optimal choices of these parameters are discussed in the following section.

5. DISCRETIZATION PARAMETERS AND OPERATION COUNT

The explicit algorithmic prescriptions introduced in the previous sections depend on three main discretization parameters, namely, the number N of surface discretization points, the number L of planes defining the sets Π^ℓ , and the number M^{equiv} of equivalent sources in each one of the sets Π_i^ℓ . Alternatively, we may characterize our discretizations by means of the three numbers (N, L, M) , where M is the total number of points within each one of the planes defining Π^ℓ . (With reference to the equivalent-source spacing h introduced in Section 4.1 we have $h = A/\sqrt{M}$.) Note that

$$2\pi M/L^2 \lesssim M^{equiv} \lesssim 4\pi M/L^2. \quad (31)$$

The precise relation between M^{equiv} and M depends on the radius used for the domains Π_i^ℓ ; see Section 4.1.

Specific choices of the parameters (N, L, M) are to be made to account for the character of each problem under consideration. Thus, for example, the number N of discretization points must be selected in order to correctly sample acoustic wavelengths and scattering surfaces. The parameters L and M , however, ought to be chosen to optimize the operation count for a prescribed accuracy ε .

To analyze and optimize the complexity of our algorithm we need to consider some basic associated quantities; these parameters, all of which can be expressed in terms of (N, L, M) , are listed in Table I. With these quantities it is not difficult to determine the number of operations required by the various stages of our algorithm.

Let us consider, for example, the cost of the most important part in our acceleration algorithm, namely, the convolution step. There are in fact six convolutions in this step, two for each of the three sets of equivalent source planes, one for the monopole distribution, and the other for the dipole distribution. Each one of these convolutions is computed through an FFT on the three-dimensional grid of equivalent sources. Since, as stated in Table I,

TABLE I
Number of Elements in Various Discretization Sets

Set	Number of elements
Surface discretization points	N
Planes defining the sets Π_ℓ	L
Equivalent sources on each of the Π_ℓ planes	M
Equivalent sources along the edge of a cell c_i	$M^{1/2}/L$
Π_i^ℓ	$M^{equiv} \approx 2\pi M/L^2$ to $4\pi M/L^2$
Nonempty cells	$\mathcal{O}(L^2)$
Surface points in each nonempty cell, average	$\mathcal{O}(N/L^2)$

TABLE II
Number of Operations Required by Various Stages of the Algorithm

Operation	Flops (per iteration)
Trigonometric interpolation (Section 3.1)	$\mathcal{O}(N \log N)$
Evaluation of adjacent interactions (Sections 3 and 4.4)	$\mathcal{O}(N^2/L^2)$
Evaluation of the vectors \mathbf{b} and \mathbf{d} for all cells (Sections 4.1 and 4.3)	$\mathcal{O}(MN/L^2)$
Solution of all least-squares problems (Sections 4.1 and 4.3)	$\mathcal{O}(M^2/L^2)$
Convolution (Section 4.2)	$\mathcal{O}(ML \log ML)$
Evaluation of surface values of the equivalent-source field (Section 4.3)	$\mathcal{O}(MN/L^2)$
Correction of the surface values of the equivalent-source field (Section 4.2)	$\mathcal{O}(M \log(M/L^2))$

a set of L planes each containing M equivalent sources is used, there is a total of LM equivalent sources in each convolution. Thus, the convolution step reduces to evaluation of a fixed number of FFTs of size $\mathcal{O}(LM)$, and thus, the total cost of the convolution step is $\mathcal{O}(LM \log(LM))$. All other necessary counts result similarly from the parameters listed in Table I. The resulting complexities are listed in Table II.

In order to proceed with the parameter optimization we may safely assume that the relation

$$M \leq \mathcal{O}(N) \quad (32)$$

holds. Indeed, in view of Eq. (23), M^{equiv} is roughly of the same order of magnitude as the number of points in the face of a cell c_i as dictated by the Nyquist frequency, so that M is approximately equal to the number of points required to Nyquist-sample a surface of area A^2 —while N must be at least that and possibly larger if subwavelength surface features need to be resolved. (Generally, a significant amount of subwavelength sampling is required to accurately resolve geometric features of a given surface, although simple, slowly varying geometries such as spheres can be sampled at rates determined by the radiation wavelength only.)

Under the condition (32), the total number of operations \mathcal{T} —which equals the sum of the quantities in the right column of Table II—is given by

$$\mathcal{T} = \mathcal{O}(N^2/L^2) + \mathcal{O}(M^3/L^6) + \mathcal{O}(ML \log ML). \quad (33)$$

In view of the relations

$$M^{equiv} = n_t^2, \quad n_t = \max \left\{ \frac{2kA}{L}, 2 \frac{-\log(\varepsilon)}{\log(3)} \right\}, \quad \text{and} \quad M^{equiv} = \mathcal{O}(M/L^2) \quad (34)$$

(see (21), (23), and (31)), it follows that for a given *fixed* value of N , \mathcal{T} actually depends on a single parameter, say, $\mathcal{T} = \mathcal{T}(L)$, and thus, that once the surface sampling rate has been chosen, the optimal discretization parameters L and M can be obtained as the solution of a one-dimensional minimization problem.

To minimize \mathcal{T} we define

$$c(\varepsilon) = -\frac{\log(\varepsilon)}{\log(3)}, \quad (35)$$

and we consider the two alternatives that arise as the maximum in Eq. (34) is realized either by $2kA/L$ or by $2c(\varepsilon)$.

Case A. $\frac{kA}{L} \geq c(\varepsilon)$. In this case we have $M = \mathcal{O}(k^2 A^2)$; to simplify the optimization calculation we define an exponent β by

$$\beta(N, kA) = \log(k^2 A^2) / \log N, \quad (36)$$

so that the total number \mathcal{T} of operations is given by

$$\mathcal{T} = \mathcal{O}(N^2/L^2) + \mathcal{O}(N^{3\beta}/L^6) + \mathcal{O}(N^\beta L \log(N^\beta L)). \quad (37)$$

An exact minimization of this expression with respect to L requires solution of the transcendental equation

$$\mathcal{O}(L^3) = \frac{\mathcal{O}(N^{2-\beta})}{\mathcal{O}(1) + \mathcal{O}(\log(N^\beta L)) + \mathcal{O}(N^{2\beta}/L^7)}. \quad (38)$$

We will find an approximate solution to this equation satisfying the condition

$$\frac{N^{2\beta}}{L^7} \ll 1. \quad (39)$$

Neglecting this term in the denominator of Eq. (38) and approximating the lower order term $1 + \log(N^\beta L)$ by 1 we obtain the approximate minimizer

$$L_0 = N^{\frac{2-\beta}{3}}. \quad (40)$$

We see that, indeed, condition (39) is satisfied as long as $\beta < 14/13$ —a restriction which can safely be assumed; see Eq. (42) below.

The value $L = L_0$ leads to an overall operation count \mathcal{T} given by

$$\mathcal{T} = \mathcal{O}(N^{(2+2\beta)/3} \log(N^{(2+2\beta)/3})). \quad (41)$$

An exact solution of Eq. (38) could in principle lead to additional reductions in the overall complexity of the algorithm, but such reductions are small—of the order of the constants involved in the “order of magnitude” calculation or logarithmic at best—and they will not be pursued further here.

The algorithmic prescriptions implicit in the complexity count of Eq. (41) depend on the parameter β —which gives a measure of the size N of the discretization used for the scattering surface relative to the size $\mathcal{O}(k^2 A^2)$ of a grid needed to correctly sample the wavelength of the incoming radiation alone in a planar surface of size A . The parameter β equals 1 when the surface discretization step is of order of a wavelength ($N = \mathcal{O}(k^2 A^2)$); for more refined surface discretization β is less than 1. Further, the condition $kA \geq Lc(\varepsilon)$ assumed in the present Case A leads to the inequality $N^{\beta/2} \geq N^{(2-\beta)/3}$ or $\beta \geq \frac{4}{5}$, so that, in the present case we have

$$\frac{4}{5} \leq \beta \leq 1. \quad (42)$$

The corresponding values of \mathcal{T} vary from $\mathcal{O}(N^{6/5} \log N^{6/5})$ to $\mathcal{O}(N^{4/3} \log N^{4/3})$.

We now consider the second alternative, which arises as the maximum in Eq. (34) is realized by $2c(\varepsilon)$ —which corresponds to the case $\beta < 4/5$.

Case B. $kA < Lc(\varepsilon) = L \frac{-\log(\varepsilon)}{\log(3)}$. In this case we have $M = \mathcal{O}(L^2)$ and Eq. (33) gives

$$\mathcal{T} = \mathcal{O}(N^2/L^2) + \mathcal{O}(L^3 \log L^3).$$

The corresponding minimizer satisfies the transcendental equation

$$\mathcal{O}(L^5 \log L^3) = \mathcal{O}(N^2).$$

An approximate solution to this equation is given by

$$L_0 = N^{2/5}, \quad (43)$$

which leads to the operation count

$$\mathcal{T} = \mathcal{O}(N^{6/5} \log N).$$

The restriction $\beta < 4/5$ on the values of β in the present case $kA < Lc(\varepsilon)$ follows from Eqs. (43).

Cases A and B provide the optimal operation counts attainable by our algorithm under all possible scenarios—that is, for all possible values of the parameter β . We thus see that for simple surfaces, for which $\beta = 1$, the overall operation count is $\mathcal{O}(N^{4/3} \log N)$. Complex surfaces, containing subwavelength features, however, need to be discretized much more finely than required by the radiation wavelength only—thus leading to values of $\beta < 1$. In sum: (a) For smooth surfaces, for which our high-order algorithm provides accurate solutions for small values of N , the overall complexity count is $\mathcal{O}(N^{4/3} \log N)$; (b) for highly complex surfaces, on the other hand, we have $\beta < 4/5$, N is large, and the overall complexity count is $\mathcal{O}(N^{6/5} \log N)$.

6. NUMERICAL RESULTS AND PERFORMANCE COMPARISONS

In the following sections we compare our results with those provided by some of the most competitive algorithms in existence today. Thus, in Section 6.1 we compare the overall performance of our method with that of FISC [28], in Section 6.2 we compare our Nystrom high-order integrator to that of [10], and in Section 6.3 we delineate the distinctions between our approach and the FFT-based algorithms of [2, 23]—with highlights on the advantages offered by the present approach.

(The following caveat should be taken into account when considering the data presented in the following sections: Our results correspond to solutions of three-dimensional acoustic scattering problems—solutions of the Helmholtz equation—whereas the FastScat and FISC data of [10, 28] correspond to solutions of the Maxwell equations. There are, of course, some differences between the Helmholtz and Maxwell problems; in particular, the unknowns in the Maxwell integral equations are two-dimensional vectors, as opposed to the single scalar unknown arising in the Helmholtz integral equation. However, our methods apply to the full Maxwell problem, and their performance in that case is expected to be similar to the one presented here.)

Solutions of the linear systems arising from discretization of Eq. (5) were obtained in all cases by means of a version of the iterative solver GMRES [27] in fully complex arithmetic—which leads to a lower number of iterations than the corresponding double-dimension real

TABLE III
Scattering by Spheres of Radii 12λ and 24λ as Computed
by FISC and the Present Algorithm

Algorithm	Radius	Time	RAM	Unknowns	$\varepsilon_{\%}$	Computer
FISC	12λ	12 h	1.8 Gb	602112	6.9%	SGI Power Challenge R8000
Present (NA)	12λ	6.5 h	24 Mb	26214	0.18%	Pentium II 400 MHz
Present	12λ	16 h	120 Mb	87318	0.0014%	Pentium II 400 MHz
FISC	24λ	8×5 h	5 Gb	2408448	7.6%	SGI Origin 2000 (8 proc.)
Present	24λ	33 h	<600 Mb	349830	0.025%	Pentium II 400 MHz

Note. The RCS error $\varepsilon_{\%}$ is defined in Appendix B.

problem; see [7]. Finally, in all cases we used the value $\gamma = \max\{3, A/\lambda\}$ for the coupling constant in Eq. (5), where A is the diameter of the scatterer. Indeed, we have found that this value of γ leads to a substantially reduced number of GMRES iterations.

6.1. Comparison with the FMM Approach of [28]

FMM-based algorithms provide considerable acceleration: they run in as little as $\mathcal{O}(N \log N)$ operations per iteration. As mentioned in the introduction, to the best of our knowledge, high-order accuracy has not been demonstrated in FMM computations of wave scattering. A possible explanation for this fact is that the FMM approach [12, 26] depends critically on certain mappings which contain multiplication by Hankel functions of high order. These operations are associated with a substantial amount of ill conditioning, which leads to accuracy limitations known as the “subwavelength breakdown problem”; see [14, p. 51; 21, p. 576]. These instabilities may prevail and mask the asymptotic high-order convergence of any underlying high-order integrator, however accurate. The search for stable FMM solvers continues to this day [19], and the feasibility of such designs is yet to be demonstrated. In contrast, the FFT acceleration techniques are stable.

Table III compares the performance of our algorithms to the FMM implementation of FISC [28]. We see that the present algorithm achieves considerably higher accuracy than those of [28] with lesser computational resources.

6.2. Nystrom Local Integrator

In Section 3 we have described a quadrature algorithm which evaluates singular integrals to high order. Although this algorithm is mainly intended for evaluation of adjacent interactions, it can be applied to evaluation of nonadjacent interactions as well. Indeed, letting η in equation (10) be a function with a large support, perhaps even

$$\eta \equiv 1, \tag{44}$$

leads to a scheme in which a large fraction of the interactions are handled by the adjacency integrator, resulting in a (high order) $\mathcal{O}(N^2)$ method.

(The speedup advantages provided by the fast nonadjacency integrator of Section 4 can, of course, be very substantial, as we demonstrate in Section 6 below. As we will see, however, the direct quadrature rule defined by (44) is so efficient that, for small- to medium-size problems, it can perform better than the accelerated scheme.)

TABLE IV
Convergence Study: Scattering by Sphere of Radius Equal to 2.7λ
Radius—Nonaccelerated Computations

Patches	Unknowns	Discretization		
		density	ε_∞	ε_2
$6 \times 17 \times 17$	1350	3 per 1λ	0.1	2.9×10^{-2}
$6 \times 33 \times 33$	5766	6 per 1λ	9.0×10^{-4}	1.8×10^{-4}
$6 \times 65 \times 65$	23790	12 per 1λ	3.6×10^{-6}	1.4×10^{-6}
$6 \times 129 \times 129$	93726	24 per 1λ	1.6×10^{-8}	5.6×10^{-9}

Note. The error norms ε_2 and ε_∞ are defined in Appendix B.

In Tables IV and V we present computations of scattering by a small sphere obtained from our basic high-order integrator (with a large-support η and without use of FFT acceleration). The convergence study of Table IV demonstrates clearly the high-order nature of the present local integrator—the corresponding high-order convergence of the accelerator is demonstrated in the following section. Table IV shows that our integrator can produce meaningful results from use of as few as three points per wavelength; subsequent doubling of the discretization density consistently result in accuracy improvements of two orders of magnitude.

Table V, on the other hand, provides a comparison of our results with those produced by the high-order Nystrom and Galerkin discretization techniques of [10] (program FastScat). In the computing time portion of Table V we only show *the setup time reported in [10]*, since in that work a slow LU decomposition was used to solve the resulting linear system. In the entries corresponding to our algorithm we show the *full computing time* required for the solution of the boundary integral equation. The notation NA in our tables indicates results obtained by means of the nonaccelerated version of our algorithm.

In examining Table V it should be borne in mind that different computers were used (a Sparc 10 in [10] and a 400-MHz PC in our work) and different problems were solved (a Maxwell system in [10] and the Helmholtz equation in our work). It should also be emphasized that, as mentioned above, only the setup time of the high-order integrator of [10] is shown. We see that our method produces substantially more accurate results than those of [10], in total computing times which are comparable or smaller than the setup portions reported in that work.

TABLE V
Performance of Three High-Order Methods: The High-Order Nystrom
and Galerkin Techniques of [10] and the Present Algorithm

Algorithm	Radius	Time	Unknowns	$\varepsilon_\%$
Nystrom [10]	2.7λ	1953 s (setup)	5400	2.2%
Galerkin [10]	2.7λ	38803 s (setup)	5400	0.48%
Present (NA)	2.7λ	294 s	2526	0.1%
Present (NA)	2.7λ	1430 s	5430	0.0045%

Note. The RCS error $\varepsilon_\%$ is defined in Appendix B.

6.3. Comparison with AIM

As discussed in the introduction, the FFT acceleration technique described in Section 4 differs substantially from all previous FFT-based techniques, including the AIM [2] and the precorrected technique of [23]. Indeed, the present technique uses surface rather than volumetric distributions of equivalent sources, and it therefore leads to (1) substantially reduced memory requirements, (2) spectrally convergent approximations, and, most importantly, (3) an improved operation count.

With regard to point (1) we note that, indeed, for an N point discretization, the traditional FFT surface scattering solvers such as the AIM require an $\mathcal{O}(N^{3/2})$ FFT—and a corresponding $\mathcal{O}(N^{3/2})$ amount of RAM. The present version of our method requires six FFTs of size $\mathcal{O}(N^q)$, $6/5 \leq q \leq 4/3$. This implies a significantly lower memory requirement, specially for the most complex scattering problems ($q = 6/5$). Even for a sphere, for which $q = 4/3$, the application of our technique to the 350,000 unknown problem of Table III requires only one-eighth of the memory that would be required by the volumetric FFT techniques. Together with the spectral convergence properties, such reductions have allowed us to compute very accurately, and on a personal computer, scattering from the bodies of sizes close to the largest reported up to now [2, 28]. (Forty IBM SP2 nodes were used in the latter work to treat scatterers of diameters up to $70\lambda \times 40\lambda \times 15\lambda$; no error estimates were given in that work. The largest bodies we have treated on a single-processor 400-MHz Pentium II (1 Gb of RAM) are an ellipsoid of diameters $100\lambda \times 25\lambda \times 25\lambda$ and a sphere of diameter 48λ —in both cases with highly accurate results; see Tables III and VI.)

(To facilitate evaluation of errors in the case of ellipsoidal scatterers (Table VI), we used an off-center source of radiation inside the ellipsoid. The Helmholtz equation was then solved for the boundary conditions induced on the surface of the ellipsoid by this point source. The exact solution to this problem outside the ellipsoid equals, quite simply, the field created by the source itself, so that errors in a numerical solution can be computed exactly. Table VI presents the corresponding error values for the numerical far field in two different error norms.)

Two additional advantages arising from the use of surface distributions of equivalent sources concern the high-order character of the method and its operation count. Indeed, to increase the approximation order of FFT-based methods (which has not been done before), one must represent larger groups of true sources by correspondingly larger groups of equivalent sources in the Cartesian grid. It can be shown that the minimal required number K of equivalent sources in a group surrounding a portion P of the scattering surface is proportional to the area of P . We thus have a choice of how to distribute equivalent sources around P . To do this we point out that the radiation from the true sources in P can be represented *exactly* by a source distribution on a *surface*(!) which encloses P . In other words, it is *sufficient* to use $\mathcal{O}(K)$ equivalent sources on the boundaries of cubic cells, and, in fact, *it*

TABLE VI
Scattering from Large Ellipsoids: Point Source inside the Body

Size	No. of iterations	Time/iteration	RAM	Unknowns	ε_∞	ε_2
$80\lambda \times 20\lambda \times 20\lambda$	15	5 h 22 min	600 M	691206	1.4×10^{-4}	2.9×10^{-5}
$100\lambda \times 25\lambda \times 25\lambda$	15	5 h 29 min	600 M	691206	1.1×10^{-3}	2.4×10^{-4}

Note. The error norms ε_2 and ε_∞ are defined in Appendix B.

TABLE VII
Accuracy of Two-Face Approximation

kH	Body size	Equivalent sources/cell	Maximum error
8	512	832	2.8×10^{-6}
12	1728	1792	1.0×10^{-7}
16	4096	3152	1.8×10^{-9}
24	13824	7120	4.5×10^{-10}

Note. The number of sources (monopoles plus dipoles) used corresponds to a fixed density of equivalent sources (2π per wavelength).

is optimal to do so. Indeed, use of a volumetric distribution of equivalent sources, which is certainly possible, would lead either to (a) a substantially larger local linear algebra problem for the determination of the intensities of the equivalent sources—if the volume is filled with $K^{3/2}$ sources—and thus, to an increased operation count, or to (b) a reduced resolution if only K sources are used volumetrically, since this leads to a coarser spacing of equivalent sources. (We note that these alternatives do not arise when low-order approximations, such as the ones in [2, 23], are used.) The two-face equivalent source distributions introduced in Section 4 constitute further extensions and improvements on the approach based on surface equivalent-source distributions.

The crude discussion presented above has been substantiated with a complete mathematical theory, which is presented in Part II. Here we present the numerical results of Table VII, which serve as an experimental validation of the proposed acceleration technique: Using our prescriptions, we will seek to approximate a field of a unit source located inside a cubic cell c_i of size H —two faces of which are depicted in gray in Fig. 3. The test source is placed at the middle of an edge of c_i , halfway between the two planes containing equivalent sources. (Our experiments show that this is the most challenging location for a test source.) The values of the original and the approximating fields are then compared on the faces of a concentric cube S_i of side $3H$. The approximation errors for different cell sizes are shown in Table VII. The column “Sources” in this table shows the number of equivalent sources used in one cell. The column “Body size” shows the acoustical size kA of

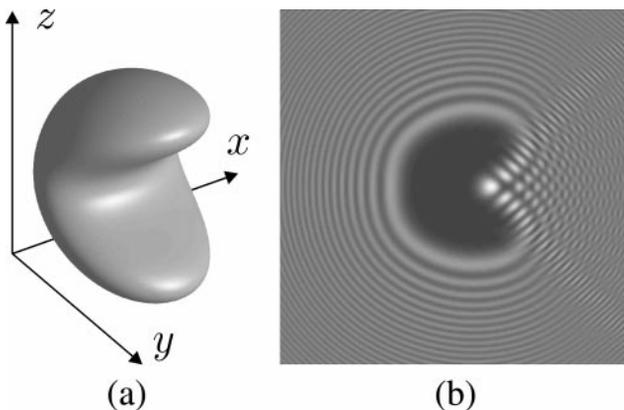


FIG. 4. Scattering by a bean-shaped obstacle; (a) the obstacle; (b) intensity of the scattered field behind the obstacle (shadow) from a plane wave propagating in the x direction.

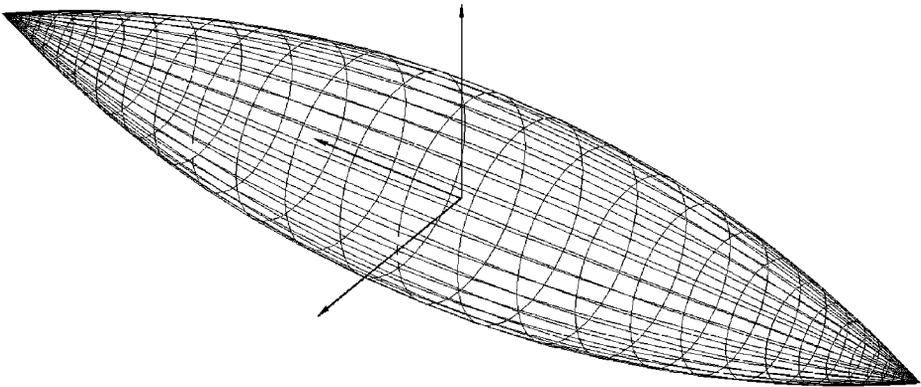


FIG. 5. Ogive geometry presented in Ref. [30].

a hypothetical scatterer which, according to our prescriptions (assuming a $\mathcal{O}(N^{4/3} \log(N))$ algorithm), would optimally use the quoted numbers of equivalent sources (roughly $(kH)^3$). We note the high accuracy of the approximation for relatively small cells, as well as the increased accuracy that results as the size of the obstacle is increased.

6.4. Numerical Results for Other Geometries

Of course, no aspect of our algorithm is restricted to consideration of simple shapes, such as spheres or ellipsoids. To demonstrate this we consider two additional geometries, namely, the bean-shaped scatterer depicted in Fig. 4a and the ogive surface shown in Fig. 5. The bean-shaped surface is defined by the equation

$$\frac{x^2}{a^2(1 - \alpha_3 \cos \frac{\pi z}{R})} + \frac{(\alpha_1 R \cos \frac{\pi z}{R} + y)^2}{b^2(1 - \alpha_2 \cos \frac{\pi z}{R})} + \frac{z^2}{c^2} = R^2,$$

with $a = 0.8$, $b = 0.8$, $c = 1$, $\alpha_1 = 0.3$, $\alpha_2 = 0.4$, and $\alpha_3 = 0.1$. The largest dimension A of this body (in the z direction) is equal to $2R$.

In order to investigate the accuracy of the algorithm in this case we first calculated solutions corresponding to the boundary conditions given by a point source inside the body, at a distance 0.5 from the origin in the direction $\theta = 10^\circ$, $\varphi = 0^\circ$. As explained in the previous section, the exact solution for such boundary conditions is known in closed form. The results of this experiment are given in Table VIII. We see that the performance of the algorithm is not altered and that none of the good qualities demonstrated above are a result of the simplicity of the geometries used. Figure 4b presents the forward scattering from the bean-shaped object of size $A = 30\lambda$ under plane-wave incidence. The near field behind the

TABLE VIII
Scattering by the Bean-Shaped Obstacle: Point Source

Size A	Unknowns	No. of iterations	Time/iteration	ε_∞	ε_2
30λ	271190	18	1 h 50 min	1.3×10^{-3}	3.0×10^{-4}
30λ	617910	18	4 h 28 min	1.5×10^{-4}	3.2×10^{-5}

Note. The error norms ε_2 and ε_∞ are defined in Appendix B.

TABLE IX
Scattering by an Ogive

Type	Size	Unknowns	Iterations	Time/iteration	ε_∞	ε_2
Nonaccelerated	1λ	1568	20	69 s	2.5×10^{-3}	1.4×10^{-3}
Nonaccelerated	1λ	6336	17	12 min 45 s	3.8×10^{-5}	2.2×10^{-5}
Nonaccelerated	1λ	25472	17	3 h 27 min	9.8×10^{-7}	4.8×10^{-7}
Accelerated	10λ	34112	13	26 min	3.8×10^{-4}	2.1×10^{-4}
Accelerated	20λ	34112	14	14 min	6.0×10^{-3}	2.4×10^{-3}
Accelerated	20λ	72320	19	67 min	5.4×10^{-5}	2.1×10^{-5}

Note. The error norms ε_2 and ε_∞ are defined in Appendix B.

object—on a plane parallel to the coordinate plane Oyz and located on a distance A from the origin—is shown in Fig. 4b.

Table IX displays a set of preliminary results obtained for scattering from a singular surface, the ogive depicted in Fig. 5, for acoustical sizes (distances between tips) equal to 1λ , 10λ , and 20λ . Here we again used boundary conditions as given by a unit source located inside the ogive for which, as mentioned above, the exact solution is known. (A convergence study for plane-wave boundary conditions showed errors consistent with those displayed in Table IX.) For the larger sizes we used the accelerator described in Section 4; note the substantial improvements in computing times resulting from the acceleration algorithm.

(In this paper we have not described our general approach to resolution of geometric singularities. In the particular case of the ogive’s conical singularities, the resolution of the singular integrands was achieved through a combination of two changes of variables: a polar change of variables similar to that described in Section 3 followed by a polynomial change of variables which regularizes the Hölder-type singularity of the underlying density; complete details are given in [6].)

A brief comment is in order with regard to adaptivity: When refined discretizations are required to resolve singularities of the scattering surface—as they are indeed implicit in our treatment of the ogive tips discussed above—the following acceleration scheme can be used. A global (coarser) Cartesian grid is utilized to compute long-range interactions between distant portions of the scattering surface, together with as many local (fine) grids as necessary to appropriately cover the singularity regions. The fine grids are then used to accelerate short-range interactions within each one of the singular regions, and thus a multiscale FFT acceleration scheme results.

APPENDIX A: ACCELERATION OF THE LEAST-SQUARES SOLVERS

We show how the solution of the linear systems (22) and (30) can be reduced to solution of eight linear systems of eight times smaller size. For definiteness we present full details for the former system

$$\mathbf{A}\xi = \mathbf{b}; \tag{45}$$

as mentioned below, the latter system can be handled similarly.

The vector ξ in (45) contains the intensities of the equivalent sources $\xi = (\xi_\ell^m, \xi_\ell^d)$; the right-hand side $\mathbf{b} = (b_1, b_2, \dots, b_{n_{coll}})$, on the other hand, is a vector containing values of

the field generated at n^{coll} collocation points on ∂S_i by the true surface sources, and \mathbf{A} is a $(n^{coll} \times n^{source})$ matrix each entry $a_{i,j}$ of which is equal to the field of the j th equivalent source at the i th collocation point.

As shown in Fig. 3, the array of equivalent sources is designed in such a way that an even number of points in each row and each column of this array lies within each face of each one of the cells c_i . This arrangement induces a useful set of symmetries in the set Π_i^ℓ . Indeed, we see that, in coordinates centered at the center of c_i , if a point $\mathbf{r} = (x, y, z)$ is in the set then so are each one of the following eight points:

$$\begin{aligned}
 \mathbf{r}_1 &= (x, y, z), \\
 \mathbf{r}_2 &= (x, -y, z), \\
 \mathbf{r}_3 &= (-x, -y, z), \\
 \mathbf{r}_4 &= (-x, y, z), \\
 \mathbf{r}_5 &= (x, y, -z), \\
 \mathbf{r}_6 &= (x, -y, -z), \\
 \mathbf{r}_7 &= (-x, -y, -z), \\
 \mathbf{r}_8 &= (-x, y, -z).
 \end{aligned} \tag{46}$$

A set Q of points which, like Π_i^ℓ , contains all of the points (46) every time it contains (x, y, z) , will be referred to as an invariant set. Clearly, an invariant set of cardinality n splits into $n/8$ invariant subsets of eight points, each of which is “minimal,” in the sense that it contains no nontrivial invariant subsets. (We note that a minimal invariant set is generated by any one of its elements by means of the symmetries (46).)

To take advantage of the symmetries underlying Π_i^ℓ we also utilize an invariant set Q_i of collocation points on ∂S_i . Further, we reorder the vectors ξ , \mathbf{b} and the matrix \mathbf{A} so that the entries corresponding to each minimal invariant group of points are numbered consecutively. (We place all monopoles first followed by all dipoles, so that the left half of the matrix \mathbf{A} corresponds to the intensities of monopoles, while the right part is related to the intensities of dipoles.) Under such numbering the matrix \mathbf{A} in (45) consists of $(n^{source}/8) \times (n^{coll}/8)$ blocks $\mathbf{P}_{r,s}$ of size 8×8 ; note that the entries of the block $\mathbf{P}_{r,s}$ relating the r th group of sources to the s th group of collocation points contains the values of the field generated by the source group r (with intensities equal to 1) at s th collocation group.

At this stage, appropriate changes of basis can be used to reduce each one of the blocks $\mathbf{P}_{r,s}$ to diagonal form. We first treat the case in which the r th group of sources is a group of monopoles; in this case we introduce an orthonormal basis in \mathbb{R}^8 consisting of the following vectors $\mathbf{v}_j \in \mathbb{R}^8$, $j = 1, \dots, 8$:

$$\begin{aligned}
 \mathbf{v}_1 &= \frac{1}{\sqrt{8}}(1, 1, 1, 1, 1, 1, 1, 1)^T, \\
 \mathbf{v}_2 &= \frac{1}{\sqrt{8}}(1, -1, 1, -1, 1, -1, 1, -1)^T, \\
 \mathbf{v}_3 &= \frac{1}{\sqrt{8}}(1, 1, -1, -1, 1, 1, -1, -1)^T, \\
 \mathbf{v}_4 &= \frac{1}{\sqrt{8}}(1, -1, -1, 1, 1, -1, -1, 1)^T, \\
 \mathbf{v}_5 &= \frac{1}{\sqrt{8}}(1, 1, 1, 1, -1, -1, -1, -1)^T,
 \end{aligned} \tag{47}$$

$$\begin{aligned}\mathbf{v}_6 &= \frac{1}{\sqrt{8}}(1, -1, 1, -1, -1, 1, -1, 1)^T, \\ \mathbf{v}_7 &= \frac{1}{\sqrt{8}}(1, 1, -1, -1, -1, -1, 1, 1)^T, \\ \mathbf{v}_8 &= \frac{1}{\sqrt{8}}(1, -1, -1, 1, -1, 1, 1, -1)^T.\end{aligned}$$

It is easy to check that, indeed, in the basis (47), the linear operator $\mathbf{P}_{r,s}$ is diagonal; that is, calling \mathbf{V} the (orthogonal and symmetric) matrix (V_1, \dots, V_8) , we have that the matrix

$$\tilde{\mathbf{P}}_{r,s} = \mathbf{V}^T \mathbf{P}_{r,s} \mathbf{V} = \mathbf{V} \mathbf{P}_{r,s} \mathbf{V} \quad (48)$$

is diagonal. To show this we consider the field generated by a set of eight monopoles located at points making up a minimal invariant subset Q of Π_i^ℓ , and with strengths given by the coordinates of the vector \mathbf{v}_j for some j . Let R , in turn, be one of the minimal invariant subsets of the set of collocation points. We note that an interchange of two points $\mathbf{r}_{k_1} \leftrightarrow \mathbf{r}_{k_2}$ in Q has the same effect on the field as a corresponding permutation of the coordinates of \mathbf{v}_j . In particular, each of the maps

$$x_k \mapsto -x_k, \quad k = 1, 2, 3, \quad (49)$$

induces a permutation of the coordinates of each of the vectors \mathbf{v}_j —which has the same effect as that caused by the corresponding permutation of monopole locations. By inspection, we see that the former permutations map a vector \mathbf{v}_j either into \mathbf{v}_j itself or into $-\mathbf{v}_j$. Calling \mathbf{w}_j the vector of values of the field at the points R , taken with an ordering as induced by Eq. (46), we see that, again, the induced permutations map a vector \mathbf{w}_j into \mathbf{w}_j if $\mathbf{v}_j \mapsto \mathbf{v}_j$, and they map \mathbf{w}_j into $-\mathbf{w}_j$ if $\mathbf{v}_j \mapsto -\mathbf{v}_j$. From this fact it is easy to check that \mathbf{w}_j is proportional to \mathbf{v}_j and, thus, that the matrix (48) is diagonal as claimed—thus concluding our treatment of 8×8 blocks in the monopole case.

Dipoles can be treated in a similar manner: the only difference in this case relates to the fact that, for an intensity vector \mathbf{v}_j as in (47), the transformation (49) for $k = \ell$ maps \mathbf{w}_j into $-\mathbf{w}_j$ if $\mathbf{v}_j \mapsto \mathbf{v}_j$, and it maps \mathbf{w}_j into \mathbf{w}_j if $\mathbf{v}_j \mapsto -\mathbf{v}_j$ —that is, the sign changes are opposite for monopoles and dipoles in the $k = \ell$ case; in the cases $k \neq \ell$ they are identical. This situation results from the fact that all dipoles used in Eq. (20) are given by derivatives in the positive x_ℓ direction. Use of dipoles defined by means of normals exterior to the cell faces gives rise to symmetries identical to those considered earlier, and, thus, the diagonalization in new variables (which may differ from the actual dipole intensities by a sign) can be performed by means of the changes of basis given for the monopole case. We have thus reduced to diagonal form each one of the 8×8 blocks which make up the matrix \mathbf{A} , and we have thus produced an equivalent sparse system. An additional permutation of the basis transforms the overall system (45) into a block-diagonal form, with blocks of size $(n^{\text{source}}/8) \times (n^{\text{coll}}/8)$, as claimed.

The block diagonalization described above thus reduces the problem of computing the QR factorization of A to that of obtaining the QR factorization of eight matrices of size $(n^{\text{source}}/8) \times (n^{\text{coll}}/8)$. This leads to a 64-fold reduction in the cost of computation of the QR factorization, and an 8-fold reduction in the cost of solution of each least-squares

problem. (Note that the latter operation is required for each nonempty cell and for every iteration of the linear-system solver.)

The symmetries in the linear system (30) can be exploited in an entirely analogous manner to produce similar cost reductions in the solution of the in-cell Dirichlet problem described in Section 4.3.

APPENDIX B: ERROR ESTIMATION

The errors mentioned in Sections 6.1 and 6.2 for the codes FISC and FastScat were originally presented in Ref. [10, 28] as L^2 errors in the bistatic radar cross section (RCS) $4\pi |u_\infty(\hat{\mathbf{x}})|^2$ measured in decibels; i.e.,

$$\varepsilon_{\text{dB}} = \left\{ \frac{1}{4\pi} \int_{\mathbb{S}^2} (10 \log_{10} |u_\infty^{\text{calc}}(\hat{\mathbf{x}})|^2 - 10 \log_{10} |u_\infty^{\text{exact}}(\hat{\mathbf{x}})|^2)^2 d\hat{\mathbf{x}} \right\}^{1/2},$$

where $u_\infty^{\text{exact}}(\hat{\mathbf{x}})$ and $u_\infty^{\text{calc}}(\hat{\mathbf{x}})$ represent the exact and calculated far-field values in the direction $\hat{\mathbf{x}}$. In Sections 6.1 and 6.2 we preferred to utilize the closely related L^2 norm of the relative pointwise error in $|u_\infty(\hat{\mathbf{x}})|^2$ given in percent,

$$\varepsilon_{\%} = 100 \left\{ \frac{1}{4\pi} \int_{\mathbb{S}^2} [(|u_\infty^{\text{calc}}(\hat{\mathbf{x}})|^2 - |u_\infty^{\text{exact}}(\hat{\mathbf{x}})|^2) / |u_\infty^{\text{exact}}(\hat{\mathbf{x}})|^2]^2 d\hat{\mathbf{x}} \right\}^{1/2},$$

which gives a direct indication of the number of digits of accuracy in the RCS. It is easy to see that, for errors of the order of, say, 10% or lower, $\varepsilon_{\%}$ and ε_{dB} are related by a constant factor:

$$\varepsilon_{\%} \approx 10 \ln 10 \varepsilon_{\text{dB}}. \quad (50)$$

The values of the error $\varepsilon_{\%}$ produced by FastScat and FISC as presented in Sections 6.1 and 6.2 were computed from the values of ε_{dB} given in Ref. [10, 28] by means of Eq. (50).

In Sections 6.3 and 6.4 we study the far-field values $u_\infty(\hat{\mathbf{x}})$ and, to gain a precise insight into the accuracies obtained for this quantity, we use two different error measures, namely, the relative L^2 norm

$$\varepsilon_2 = \left\{ \int_{\mathbb{S}^2} |u_\infty^{\text{calc}}(\hat{\mathbf{x}}) - u_\infty^{\text{exact}}(\hat{\mathbf{x}})|^2 d\hat{\mathbf{x}} \right\}^{1/2} / \left\{ \int_{\mathbb{S}^2} |u_\infty^{\text{exact}}(\hat{\mathbf{x}})|^2 d\hat{\mathbf{x}} \right\}^{1/2} \quad (51)$$

and the maximum norm

$$\varepsilon_\infty = \max_{\hat{\mathbf{x}}} |u_\infty^{\text{calc}}(\hat{\mathbf{x}}) - u_\infty^{\text{exact}}(\hat{\mathbf{x}})|. \quad (52)$$

ACKNOWLEDGMENTS

This effort is sponsored by the Air Force Office of Scientific Research, Air Force Materials Command, USAF, under Grants F49620-96-1-0008 and F49620-99-1-0010. O.B. gratefully acknowledges support from NSF (through an NYI award and through Contracts DMS-9523292 and DMS-9816802) and from the Powell Research Foundation. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

REFERENCES

1. C. R. Anderson, An implementation of the fast multipole method without multipoles, *SIAM J. Sci. Stat. Comput.* **13**(4), 923 (1992).
2. E. Bleszynski, M. Bleszynski, and T. Jaroszewicz, AIM: Adaptive integral method for solving large-scale electromagnetic scattering and radiation problems, *Radio Sci.* **31**, 1225 (1996).
3. N. Bojarski, The k -space formulation of the scattering problem in the time domain, *J. Acoust. Soc. Am.* **72**, 570 (1982).
4. Deleted in proof.
5. O. P. Bruno and L. A. Kunyansky, Surface scattering in 3-D: An accelerated high-order solver, submitted for publication.
6. O. P. Bruno and L. A. Kunyansky, High-order solution of wave scattering problems: Corners, edges and other geometric singularities, in preparation.
7. O. Bruno and M. Hyde, High order solution of scattering by penetrable bodies, in preparation.
8. O. Bruno and A. Sei, A high order solver for problems of scattering by heterogeneous bodies, in *Proceedings of the 13th Annual Review of Progress in Applied Computational Electromagnetism* (Applied Computational Electromagnetics Society, 1997), pp. 1296–1302.
9. O. Bruno and A. Sei, A fast high-order solver for EM scattering from complex penetrable bodies: TE case, *IEEE Trans. Antennas Propag.* **48**, 1862 (2000).
10. L. Canino, J. Ottusch, M. Stalzer, J. Visher, and S. Wandzura, Numerical solution of the Helmholtz equation using a high-order Nystrom discretization, *J. Comput. Phys.* **146**, 627 (1998).
11. M. F. Catedra, E. Cago, and L. Nuno, A numerical scheme to obtain the RCS of three-dimensional bodies of resonant size using the conjugate gradient method and the fast fourier transform, *IEEE Trans. Antennas Propag.* **37**, 528 (1989).
12. R. Coifman, V. Rokhlin, and S. Wandzura, The fast multipole method for the wave equation: A pedestrian prescription, *IEEE Antennas Propag. Mag.* **35**, 7 (1993).
13. D. Colton and R. Kress, *Inverse Acoustic and Electromagnetic Scattering Theory* (Springer-Verlag, Berlin/Heidelberg, 1998).
14. B. Dembart and E. Yip, The accuracy of fast multipole methods for Maxwell's equations, *IEEE Comput. Sci. Eng.* **4**, 48 (1998).
15. M. Epton and B. Dembart, Multipole translation theory for the three-dimensional Laplace and Helmholtz equations, *SIAM J. Sci. Comput.* **16**, 865 (1995).
16. L. B. Felsen and N. Marcuvitz, *Radiation and Scattering of Waves* (Prentice-Hall, New York, 1973).
17. Deleted in proof.
18. A. Greenbaum, L. Greengard and G. McFadden, Laplace equation and the Dirichlet–Neumann map in multiply connected domains, *J. Comput. Phys.* **105**, 267 (1993).
19. L. Greengard, J. F. Huang, V. Rokhlin, and S. Wandzura, Accelerating fast multipole methods for the Helmholtz equation at low frequencies, *IEEE Comput. Sci. Eng.* **5**, 32 (1998).
20. L. A. Kunyansky and O. P. Bruno, A fast, high-order algorithm for the solution of surface scattering problems. II. Theoretical considerations, submitted for publication.
21. C. Labreuche, A convergence theorem for the fast multipole method for 2-dimensional scattering problems, *Math. Comput.* **67**, 553 (1998).
22. E. Martensen, Über eine methode zum räumlichen Neumannschen problem mit einer anwendung für torusartige berandungen, *Acta Math.* **109**, 75 (1963).
23. J. R. Phillips and J. K. White, A precorrected-FFT method for electrostatic analysis of complicated 3-D structures, *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems* **16**, 1059 (1997).
24. V. Rokhlin, Rapid solution of integral equations of classical potential theory, *J. Comput. Phys.* **60**, 187 (1985).
25. V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, *J. Comput. Phys.* **86**, 414 (1990).

26. V. Rokhlin, Diagonal form of translation operators for the Helmholtz equation in three dimensions, *Appl. Comput. Harmonic Anal.* **1**, 82 (1993).
27. Y. Saad and M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems, *SIAM J. Sci. Stat. Comput.* **7**, 856 (1986).
28. J. M. Song, C. C. Lu, W. C. Chew, and S. W. Lee, Fast Illinois solver code (FISC), *IEEE Antennas Propag. Mag.* **40**, 27 (1998).
29. J. M. Song, C. C. Lu, and W. C. Chew, Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects, *IEEE Trans. Antennas Propag.* **45**, 1488 (1997).
30. A. C. Woo, H. T. G. Wang, M. J. Schuh, and M. L. Sanders, Benchmark radar targets for the validation of computational electromagnetics programs, *IEEE Antennas Propag. Mag.* **35**, 84 (1993).