International Journal for Uncertainty Quantification, 1 (2): 99–117 (2011)

FAST METHOD FOR HIGH-FREQUENCY ACOUSTIC SCATTERING FROM RANDOM SCATTERERS

Paul Tsuji, [1,*] ***Dongbin Xiu,*** [2] **&** ***Lexing Ying*** [3]

[1] *ICES, University of Texas at Austin, Austin, TX 78712*
[2] *Department of Mathematics, Purdue University, West Lafayette, IN 47907*
[3] *Department of Mathematics and ICES, University of Texas at Austin, TX 78712*

Original Manuscript Submitted: 26/05/2010; Final Draft Received: 18/09/2010

This paper is concerned with the uncertainty quantification of high-frequency acoustic scattering from objects with random shape in two-dimensional space. Several new methods are introduced to efficiently estimate the mean and variance of the random radar cross section in all directions. In the physical domain, the scattering problem is solved using the boundary integral formulation and Nyström discretization; recently developed fast algorithms are adapted to accelerate the computation of the integral operator and the evaluation of the radar cross section. In the random domain, it is discovered that due to the highly oscillatory nature of the solution, the stochastic collocation method based on sparse grids does not perform well. For this particular problem, satisfactory results are obtained by using quasi–Monte Carlo methods. Numerical results are given for several test cases to illustrate the properties of the proposed approach.

KEY WORDS: *acoustic scattering, random domains, uncertainty quantification, boundary integral equations, fast algorithms, quasi–Monte Carlo methods*

1. INTRODUCTION

Acoustic and electromagnetic wave propagation in the presence of impenetrable scatterers is a commonly studied problem, with applications such as radar/sonar imaging and wireless communications. In many of these practical situations, the shape and properties of the scattering object may be slightly perturbed from the specifications of the original geometry. This may occur if a vehicle has manufacturing defects or if it has suffered damage after combat use. As a result, there is a level of uncertainty when observing physical quantities that are dependent on the characteristics of the scatterer. Quantifying this uncertainty is an important question, from an engineering point of view, and is typically done using probabilistic methods.

The work presented here deals with high-frequency acoustic scattering from an impenetrable object with a randomly perturbed surface in two dimensions. Let $D \subset \mathbb{R}^d (d = 2, 3)$ be a sound-soft scatterer with boundary ∂D sampled from a certain probability space. For a given incident field $u_I(x)$, the scattered field $u(x)$ satisfies the Helmholtz equation in the exterior of D with the following conditions:

$$\begin{aligned} &\Delta u(x) + k^2 u(x) = 0, \quad x = (x_1, x_2) \in \mathbb{R}^d \backslash D, \\ &u(x) = -u_I(x), \quad x \in \partial D, \\ &\lim_{r\to\infty} r^{(d-1)/2} \left(\partial_r u - \imath k u\right) = 0, \end{aligned} \tag{1.1}$$

where $\imath = \sqrt{-1}$, k is the wavenumber, and the wavelength $\lambda = 2\pi/k$. The last equation, known as the Sommerfeld radiation condition, enforces the scattered field to propagate from the scatterer to infinity; this ensures the uniqueness of the solution to the exterior scattering problem. Here, we are interested in the high-frequency setting, in which the

---

[*] Correspond to Paul Tsuji, E-mail: ptsuji@gmail.com

size of the domain D is much larger than the wavelength λ. The far-field pattern of the scattered field $u(x)$ is defined as

$$F(s) = \lim_{r \to \infty} \frac{u(rs)}{e^{\imath kr}/r^{(d-1)/2}}, \quad s \in \mathbb{S}, \tag{1.2}$$

where $\mathbb{S}$ is the unit circle/sphere. For many applications, the most important quantity is the radar cross section (RCS) $R(s)$, defined as

$$R(s) = |F(s)|^2, \quad s \in \mathbb{S}.$$

Because the scatterer's shape is random, there is uncertainty associated with the radar cross section $R(s)$. In practice, we are more interested in the statistical quantities of $R(s)$, such as the mean and variance.

From a numerical point of view, this problem involves two issues. The first is related to the high-frequency nature of the scattering problem. In many settings, the operating wavelength is much smaller than the radius of the scattering object in question; for example, a typical wavelength used by military communications devices ranges between several millimeters to a few centimeters, whereas the length of a fighter jet is ∼20 m. In order to accurately capture the scattering phenomena, it is commonly necessary to use a grid that resolves the oscillations of each wavelength. Hence, a large number of discretization points is necessary for such objects that are electrically or acoustically large. The standard finite element and finite difference methods for this scattering problem face several difficulties. First, the number of degrees of freedom grows as $[\mathrm{diam}(D)/\lambda]^d$. Other difficulties include artificial truncation of the unbounded computational domain, mesh generation of the scattering domain, and the large condition number of the resulting linear systems.

Because of these reasons, the most effective method for sound-soft scattering in linear homogeneous media is the boundary integral or boundary element method, where the scattered field $u(x)$ is represented as the acoustic potential generated by a layered density on ∂D that satisfies a boundary integral equation. Once this layered density is resolved, quantities such as the far-field pattern or total field can be calculated by the appropriate integrals. Compared to the aforementioned methods, the boundary integral formulation has several advantages, including the $[\mathrm{diam}(D)/\lambda]^{d-1}$ scaling of the number of unknowns, automatic treatment of the Sommerfeld radiation condition, and good conditioning properties of the resulting linear systems. The main drawback of the method is that the matrix equation which results from the discretization of a boundary integral equation is dense. In the past two decades, several efficient algorithms have been developed to speed up the iterative solution of such systems [1–7].

The randomness of the boundary surface poses the second challenge. In the traditional case, the geometry of the object in question is known and the main goal is to examine the deterministic scattered field. However, in many instances, the exact geometry of the object is not known or there is some perturbation from the geometry that would cause a notable uncertainty in the scattered field and its far-field pattern. Naturally, this problem falls into the category of stochastic modeling. The traditional approach is the Monte Carlo method [8], but it usually results in long computational times due to its slow $O(1/\sqrt{N})$ convergence with respect to the number of realizations N. More recently, a class of methods based on generalized polynomial chaos (gPC) [9, 10] have been developed and become popular in many practical applications. Most notable is the stochastic collocation method using Smolyak sparse grids [11], which may offer much better convergence properties than the Monte Carlo method while keeping the same ease of implementation. (A recent review of gPC methods can be found in [12].) The gPC methods have been applied in several cases to study random surface or roughness problems (for example, [13, 14]). For wave scattering with random shapes, the gPC method was applied in [15] and found to be effective in low-frequency scattering. However, for the high-frequency scattering problem considered here, the sparse grid collocation method does not offer a big advantage over other methods. In order to resolve the highly oscillatory solution, a higher order method is required in the random space; in addition, to properly model the rough physical domain, the random space needs to be parametrized by a larger set of random variables. Therefore, for gPC-based methods, the problem would require a high-order implementation in a large number of dimensions. This will almost certainly result in a large number of unknowns, which grows quickly for a higher-order method. This is essentially the effect of the "curse of dimensionality," though its more familiar effect is the fast growth of the number of unknowns in the physical domain. To alleviate this computational difficulty, quasi–Monte Carlo (QMC) methods based on low-discrepancy sequences are introduced. The QMC methods [16, 17] are in fact deterministic approaches based on pseudo random numbers; they have much faster convergence rates $[O(1/N)$ up to logarithmic factors] without sacrificing the generality of the Monte Carlo method, and their dependence on dimensionality is much weaker than for stochastic collocation methods.

In this paper, we combine the recent development on fast algorithms for the boundary integral solver with the QMC method to efficiently address the uncertainty quantification problem for high-frequency acoustic scattering. The rest of this paper is organized as follows. In Section 2, we derive the boundary integral formulation of the scattering problem and demonstrate how the randomness of the boundary is modeled. In Section 3, we detail the main numerical methods, including the numerical discretization, fast summation techniques, and QMC methods. In Section 4, we report the numerical experiments.

2. MATHEMATICAL FORMULATIONS

2.1 Boundary Integral Formulation

We consider the two-dimensional acoustic scattering problem with a sound-soft scatterer D. In the presence of an incident field $u_I(x)$, the scattered field $u(x)$ satisfies the following exterior boundary value problem:

$$\begin{aligned}
&\Delta u(x) + k^2 u(x) = 0, \quad x = (x_1, x_2) \in \mathbb{R}^2 \backslash D, \\
&u(x) = -u_I(x), \quad x \in \partial D, \\
&\lim_{r\to\infty} \sqrt{r}\,(\partial_r u - \imath k u) = 0.
\end{aligned} \tag{2.1}$$

It is convenient to set the wavenumber $k = 2\pi$ so that the wavelength $\lambda = 2\pi/k = 1$. We further assume that D is supported in the square $[-K/2, K/2]^2$, so that K can be considered effectively as the diameter of D. For the high-frequency problems that we are interested in, K is much larger than $\lambda = 1$. The boundary integral formulation of Eq. (2.2) utilizes the free-space fundamental solution (or the Green's function) of the 2D Helmholtz equation:

$$G(x,y) = \frac{\imath}{4} H_0^1(k|x-y|). \tag{2.2}$$

Here, H_0^1 is the zero-order Hankel function of the first kind. Using Green's third identity and the boundary condition of the sound-soft object, we can formulate the scattered field $u(x)$ as a combination of single- and double-layer potentials with surface density $\varphi(x)$ for $x \in \partial D$,

$$u(x) = \int_{\partial D} \left[\frac{\partial G(x,y)}{\partial n(y)} - \imath\eta \cdot G(x,y) \right] \varphi(y) dy, \tag{2.3}$$

where $n(y)$ is the unit normal of the scatterer surface at y and $\eta \approx k = 2\pi$. Letting x approach ∂D gives rise to the boundary integral equation

$$-u_I(x) = \frac{1}{2}\varphi(x) + \int_{\partial D} \left[\frac{\partial G(x,y)}{\partial n(y)} - \imath\eta \cdot G(x,y) \right] \varphi(y) dy. \tag{2.4}$$

Here, the extra $(1/2)\varphi(x)$ term appears because the kernels $[\partial G(x,y)]/[\partial n(y)]$ become singular as x approaches the boundary, and its limit is a combination of the δ term plus the improper integral in Eq. (2.4). The overall method to solve for $u(x)$ is as follows: one first solves for surface density $\varphi(x)$ in Eq. (2.4); after the surface density is found, it can be substituted back into Eq. (2.3) to calculate the scattered field. The total field is now found through adding the scattered field to the incident field. For more details, we refer to [18].

The far-field pattern $F(s)$ of the scatterer can also be calculated once the surface density $\varphi(x)$ is found. In the 2D case, it is given by

$$F(s) = \frac{e^{-\imath\frac{\pi}{4}}}{\sqrt{8\pi k}} \int_{\partial D} \{k\,[n(y)\cdot s] + \eta\} e^{-\imath k s\cdot y} \varphi(y) dy. \tag{2.5}$$

and the radar cross section $R(s)$ is equal to $|F(s)|^2$.

2.2 Probabilistic Modeling of Domain Uncertainty

To incorporate the uncertainty of the scatterer D, we switch to a probabilistic setting and model the surface as a random process. That is, we allow the boundary to take the form

$$\partial D_{z(\omega)} = \{x(t,\omega) = b(t) \cdot [1 + p(t,\omega)], t \in [0, 2\pi), \omega \in \Omega\},$$

where $b(t) = [b_1(t), b_2(t)]$ is the base geometry, Ω is the event space in a properly defined probability space, and $p(t,\omega)$ is the perturbation. For a fixed ω, $p(t,\omega)$ is a deterministic function representing how the base geometry $b(t)$ is scaled, while for a fixed location t, $p(t,\omega)$ is a random variable representing the uncertainty of the surface at the location associated with t. The perturbation $p(t,\omega)$ is also assumed to be sufficiently regular so that the scattering problem is well posed almost everywhere in Ω.

A critical step in modeling the random surface is to properly parametrize the random process by a finite number of independent random variables. Let $Z(\omega) = [Z_1(\omega), \ldots, Z_M(\omega)]$, $M \geq 1$, be such a set of independent random variables, whose probability distribution is $F_Z(z) = \text{Prob}(Z \leq z)$, where $z \in \mathbb{R}^M$. Without loss of generality, we focus on the continuous random variables, where a probability density function $\rho(z) = dF_Z(z)/dz$ exists. The random surface can now be expressed in terms of Z in the following manner:

$$\partial D_z = \{b(t) \cdot [1 + p(t, Z)], t \in [0, 2\pi), Z \in \mathbb{R}^M\}.$$

The requirement of the independence of $Z_1(\omega), \ldots, Z_M(\omega)$ is important for numerical purposes because most random number generators are designed to generate independent random numbers. Common tools for constructing such a finite-dimensional representation or approximation are more established for Gaussian processes. For example, spectral expansion [19] and Karhunen–Loeve expansions [20] are quite effective. For non-Gaussian processes, the parametrization procedure is still an active research topic, with many open issues. For the purpose of this paper, we simply assume that such a representation has already been established.

Now, the integral formulations given in Section 2.1 all depend on z. The density $\varphi_z(x)$ for $x \in \partial D_z$ satisfies

$$-u_I(x) = \frac{1}{2}\varphi_z(x) + \int_{\partial D_z} \left[\frac{\partial G(x,y)}{\partial n(y)} - \imath\eta \cdot G(x,y)\right] \varphi_z(y)dy. \tag{2.6}$$

The far-field pattern and the radar cross sections are equal to

$$F_z(s) = \frac{e^{-\imath\frac{\pi}{4}}}{\sqrt{8\pi k}} \int_{\partial D_z} \{kn(y)\cdot s + \eta\} e^{-\imath k s \cdot y} \varphi_z(y)dy, \quad R_z(s) = |F_z(s)^2|.$$

Finally, the mean and the variance of the observable $R(s)$ are given by

$$\begin{aligned} \mathrm{E}[R(s)] &= \int R_z(s)\rho(z)dz, \\ \mathrm{var}[R(s)] &= \int \{R_z(s) - \mathrm{E}[R(s)]\}^2 \rho(z)dz. \end{aligned} \tag{2.7}$$

It is worth noting that M, the dimensionality of the random variables z, depends on the domain uncertainty. In many realistic cases, the uncertainty presents itself with "fine" structure and as surface roughness. This implies that the random processes describing such an uncertainty should have short correlation length. Subsequently, the dimensionality M resulting from the parametrization procedure will be large. Therefore, in many practical simulations, the integrals Eq. (2.7) will be in a high-dimensional random space $\mathbb{R}^M$ with $M \gg 1$.

3. NUMERICAL METHODS

3.1 Nyström Discretization

To numerically solve for the surface density from Eq. (2.4), the Nyström method is used to discretize the integral equation. Using the periodic boundary parametrization $x(t) = [x_1(t), x_2(t)]$ for $t \in [0, 2\pi]$, the parametrized integral equation takes the following form:

$$-u_I(t) = \frac{1}{2}\varphi(t) + \int_0^{2\pi} K(t,t')\varphi(t')dt', \quad t \in [0, 2\pi], \tag{3.1}$$

with

$$K(t,t') = \left\{ \frac{\partial G[x(t), x(t')]}{\partial n[x(t')]} - \imath\eta G[x(t), x(t')] \right\} J(t'). \tag{3.2}$$

By abusing the notation slightly, we denote $u_I(t) = u_I[x(t)]$, $\varphi(t) = \varphi[x(t)]$, and the Jacobian $J(t) = J[x(t)]$. To discretize the integral equation (3.1), we create an equispaced N_t-point grid over the variable t such that $t_i = 2\pi i/N_t$ for $i = 0, 1, 2, \ldots, N_t - 1$, with $N_t = O(K)$; these points are the Nyström (or quadrature) points. The condition $N_t = O(K)$ corresponds to discretizing the boundary with $O(1)$ grid points per wavelength. The equations enforced at the Nyström points are written

$$-u_I(t_i) = \frac{1}{2}\varphi(t_i) + \int_0^{2\pi} K(t_i, t)\varphi(t)dt, \quad i = 0, 1, 2, \ldots, N_t - 1. \tag{3.3}$$

3.1.1 Quadrature Rule

The next component of the Nyström method is the quadrature rule for the integral in Eq. (3.3). More specifically, given the values of $\varphi(t_i)$ for $i = 0, 1, \ldots, N_t - 1$, one should be able to compute an accurate approximation of $\int_0^{2\pi} K(t_i, t)\varphi(t)dt$. Once the quadrature rule is determined, the resulting linear system of Eq. (3.3) is solved using iterative methods such as GMRES. If the kernel in the integrand had been smooth for all t, the standard N-point trapezoidal rule with quadrature points $\{t_j\}$ could be ideal for approximating the integral operator. Unfortunately, because the kernel $K(t_i, t)$ has a logarithmic singularity at $t = t_i$, a special quadrature rule is required. For this purpose, we utilize the modified trapezoidal rule proposed by Kapur and Rokhlin in [21]. The main idea of [21] is to build a local correction near the singularity. Taking $f(t) = K(t_i, t)\varphi(t)$ and $h = 2\pi/N_t$, the Kapur–Rokhlin quadrature rule applied to the integral in Eq. (3.3) takes the form

$$\int f(t)dt \approx \sum_{\substack{j=0 \\ j\neq i}}^{N_t-1} f(t_j)h + \sum_{\substack{j=i-m \\ j\neq i}}^{i+m} f(t_j)\beta_{|j-i|}h, \tag{3.4}$$

where the second summation is the correction term and $\beta_{|j-i|}$ are the local correction weights. One drawback that we noted about the correction weights is that they can have large negative numbers, which causes the resulting matrix to become less stable and results in an increased number of GMRES iterations.

In order to remedy this problem, we modify the approach slightly by introducing a denser grid just for the purpose of numerical integration, while keeping the original grid for the identity term $(1/2)\varphi(t)$ and incident field $u_I(t)$. In essence, this will only change the matrix-vector multiplication step in the GMRES iteration, as we will soon show; the solution $\varphi(t)$ will still be computed on the original mesh. We denote the density at these points as $\varphi(\tilde{t}_j)$ for $j = 0, 1, \ldots, \tilde{N}_t - 1$, where $\tilde{N}_t = rN_t$ for some integer refinement rate r and $\tilde{t}_j = 2\pi j/\tilde{N}_t$. In practice, r is chosen to be 4 or 8. With the more refined mesh, the quadrature formula becomes

$$\int_0^{2\pi} K(t_i, t)\varphi(t)dt \approx \sum_{\substack{j=0 \\ j\neq ri}}^{\tilde{N}_t-1} K(t_i, \tilde{t}_j)\varphi(\tilde{t}_j)\tilde{h} + \sum_{\substack{j=ri-m \\ j\neq ri}}^{ri+m} K(t_i, \tilde{t}_j)\varphi(\tilde{t}_j)\beta_{|j-ri|}\tilde{h}. \tag{3.5}$$

Here, we've run into another problem: the original grid defines φ only at points $\{t_j\}_{j=0}^{N_t-1}$, a subset of $\{\tilde{t}_j\}_{j=0}^{\tilde{N}_t-1}$; that is, we must somehow recover $\varphi(\tilde{t}_j)$ from the original grid. Because the surface of the scattering object is smooth in $\mathbb{R}^2$, we choose to use Fourier interpolation to recover the surface density on the refined grid.

We can now apply the GMRES solver to the system of equations using the refined grid in the matrix-vector multiplication within each iteration. For each iteration, we are given the density $\varphi(t_j)$ for $j = 0, 1, 2, \ldots, N_t - 1$, and we are required to calculate the right-hand side of Eq. (3.3). Based on the above discussion, we perform the following steps:

1. Given $\varphi(t_j)$ for $j = 0, 1, 2, \ldots, N_t - 1$, we use Fourier interpolation to get $\varphi(\tilde{t}_j)$ for $j = 0, 1, 2, \ldots, \tilde{N}_t - 1$.
2. For $i = 0, 1, 2, \ldots, N_t - 1$, compute
$$a_i = \sum_{\substack{j=0 \\ j \neq ri}}^{\tilde{N}_t - 1} K(t_i, \tilde{t}_j)\varphi(\tilde{t}_j)\tilde{h} \tag{3.6}$$
3. For $i = 0, 1, 2, \ldots, N_t - 1$, compute the matrix-vector product
$$b_i = \sum_{\substack{j=ri-m \\ j \neq ri}}^{ri+m} K(t_i, \tilde{t}_j)\varphi(\tilde{t}_j)\beta_{|j-ri|}\tilde{h}. \tag{3.7}$$
4. For $i = 0, 1, 2, \ldots, N_t - 1$, the right hand side of Eq. (3.3) is $\left(\frac{1}{2}\varphi(t_i) + a_i + b_i\right)$.

Step 1 of the procedure can be computed using the FFT, which takes $O(N_t \log N_t)$ operations. Because m is of $O(1)$, the amount of work necessary for step 3 is $O(N_t)$. Obviously, step 4 also takes $O(N_t)$ steps. The only step that takes $O(N_t^2)$ operations is step 2, and a fast algorithm is required to bring down this complexity.

3.1.2 Fast Pairwise Summation

Let us denote $x_i = x(t_i)$ and $\tilde{x}_j = x(\tilde{t}_j)$. Then $\tilde{x}_j = x_i$ if and only if $j = ri$. Under the new notation, Eq. (3.6) becomes
$$\sum_{\substack{j \neq ri \\ j=0}}^{\tilde{N}_t - 1} K(t_i, \tilde{t}_j)\varphi(\tilde{t}_j)\tilde{h} = \sum_{j:\tilde{x}_j \neq x_i} \left[\frac{\partial G(x_i, \tilde{x}_j)}{\partial n(\tilde{x}_j)} - \imath\eta G(x_i, \tilde{x}_j)\right] f_j \tag{3.8}$$
with $f_j = J(\tilde{t}_j)\varphi(\tilde{t}_j)\tilde{h}$.

This new formulation is close to the N-body problem of the Helmholtz kernel: Given a point set $\{p_i\}_{i=1}^N$ and sources $\{f_i\}_{i=1}^N$, one wants to evaluate at each p_i,
$$u_i = \sum_{\substack{j=0 \\ j \neq i}}^{N} G(p_i, p_j) f_j \tag{3.9}$$
Several methods [1–3] have been proposed to evaluate $\{u_i\}_{i=1}^N$ in $O(N \log N)$ steps. Here, we employ the directional multilevel method proposed in [5, 6] by one of the authors. A brief description of this algorithm is provided in Section A.1 in the Appendix for completeness.

However, our summation [Eq. (3.8)] is different from the standard N-body problem [Eq. (3.9)] in two aspects. First, we are using different source and target points; in our case, $\{x_i\}$ are the source locations and $\{\tilde{x}_j\}$ are the target locations. This fact does not change the algorithm significantly, as $\{x_i\} \subset \{\tilde{x}_j\}$; we can simply ignore potentials computed at the leaf box level for $\{\tilde{x}_j\}_{j \neq ri}$. Second, the kernel function is $[\partial G(x,y)]/[\partial n(y) - \imath\eta G(x,y)]$, a linear combination of the Green's function $G(x,y)$ with its normal derivative $[\partial G(x,y)]/[\partial n(y)]$ at the source points y. We can easily extend the directional algorithm to this kernel as well using the following argument: Because the normal derivative with respect to y is a linear operator, $[\partial G(x,y)]/[\partial n(y)] - \imath\eta G(x,y)$ as a function of x still satisfies the Helmholtz equation. Thus, the potential generated by this mixed kernel can still be reproduced by the equivalent sources of the Green's function $G(x,y)$. The only difference is that in the construction of the equivalent sources at the leaf boxes, we use the kernel $[\partial G(x,y)]/[\partial n(y)] - \imath\eta G(x,y)$ to determine the check potentials. The locations of equivalent sources and potentials stay the same.

3.1.3 Evaluation of the Far-Field Pattern

Once $\varphi(x)$ is ready, the next step is to evaluate the far-field pattern $F(s)$ numerically. Typically, one needs to compute the $F(s)$ in a finite number of directions of order $O(K)$. To that end, we discretize the unit sphere S with $N_s = O(K)$ equally spaced points s_ℓ, for $\ell = 0, 1, \ldots, N_s - 1$; as a parametrized function, it is easy to see that

$$s_\ell = (s_{\ell,1}, s_{\ell,2}) = \left[\cos\left(\frac{2\pi\ell}{N_s}\right),\ \sin\left(\frac{2\pi\ell}{N_s}\right)\right].$$

Now, for each s_ℓ, the far-field pattern $F(s_\ell)$ is given by

$$F(s_\ell) = \frac{e^{-\imath(\pi/4)}}{\sqrt{8\pi k}} \int_{\partial D} \{k[n(y)\cdot s_\ell] + \eta\} e^{-\imath k s_\ell \cdot y} \varphi(y) dy. \tag{3.10}$$

Because the integrand in the far-field operator contains no singular functions, the trapezoidal rule can be applied to approximate $F(s_\ell)$ with super algebraic convergence. Accordingly, Eq. (3.10) is approximated by

$$F(s_\ell) \approx \frac{e^{-\imath(\pi/4)}}{\sqrt{8\pi k}} \sum_{i=0}^{N_t-1} \left(k\{n[x(t_i)]\cdot s_\ell\} + \eta\right) e^{-\imath k s_\ell \cdot x(t_i)} \varphi(t_i) J(t_i) h. \tag{3.11}$$

Direct evaluation of this sum for each s_ℓ takes $O(N_s N_t) = O(K^2)$ steps, which can be very expensive when K is large. In order to speed up this calculation, we write the dot product in the brackets as the sum of two components,

$$k\{n[x(t_i)]\cdot s_\ell\} = k\left\{n_1[x(t_i)] s_{\ell,1}\right\} + k\left\{n_2[x(t_i)] s_{\ell,2}\right\}, \tag{3.12}$$

where $n = (n_1, n_2)$. Thus, the summation in Eq. (3.11) can be written as the sum of the following three sums:

$$s_{\ell,1} \frac{e^{-\imath(\pi/4)}}{\sqrt{8\pi k}} \sum_{i=0}^{N_t-1} e^{-\imath k s_\ell \cdot y(t_i)} \left\{k n_1[y(t_i)] \varphi(t_i) J(t_i) h\right\} \tag{3.13}$$

$$s_{\ell,2} \frac{e^{-\imath(\pi/4)}}{\sqrt{8\pi k}} \sum_{i=0}^{N_t-1} e^{-\imath k s_\ell \cdot y(t_i)} \left\{k n_2[y(t_i)] \varphi(t_i) J(t_i) h\right\} \tag{3.14}$$

$$\frac{e^{-\imath(\pi/4)}}{\sqrt{8\pi k}} \sum_{i=0}^{N_t-1} e^{-\imath k s_\ell \cdot y(t_i)} \left[\eta \varphi(t_i) J(t_i) h\right] \tag{3.15}$$

After appropriate rescaling, each summation becomes an instance of the sparse Fourier transform introduced in [7], where both the spatial and Fourier data are sparsely supported. More precisely, define $N = 2K$, $p_\ell = -K \cdot s_\ell$, $\xi_i = 2 \cdot x(t_i)$, $f_i = \eta\phi(t_i) J(t_i) h$, and Eq. (3.15) becomes

$$\sum_{i=0}^{N_t} e^{2\pi\imath p_\ell \cdot \xi_i / N} f_i \tag{3.16}$$

up to a constant scaling. Because we have $p_\ell \in [-N/2, N/2]^2$, $\xi_i \in [-N/2, N/2]^2$, and $N_t = O(K) = O(N)$, Eq. (3.16) fits exactly into the definition of the sparse Fourier transform. In [7], it is shown that the sparse Fourier transform can be computed in $O(K \log K)$ steps; a short description of the algorithm of [7] is outlined in Section A.2 in the Appendix for the sake of completeness. Now, since each of the three sums [Eqs. (3.13)–(3.15)] can be computed in $O(K \log K)$ steps, the total cost of computing $R(s_\ell)$ for all s_ℓ is also $O(K \log K)$.

3.2 Stochastic Algorithms: Sparse Grids and QMC Method

To efficiently evaluate the statistics defined by the integrals in Eq. (2.7), a careful approach must be taken. A popular cubature scheme used to compute these multidimensional integrals is the Smolyak sparse grid [22]; though it was shown to be effective for general purpose stochastic problems in [11], the sparse grid is found to be less effective in this situation. The main reason is that the integrals resulting from high-frequency scattering are highly oscillatory, thus requiring higher-order methods. For even moderately high dimensions, the number of points in the sparse grid grows rapidly as the accuracy level is increased. This can be seen in Table 1, where the total number of points are tabulated for moderate dimensions of $M = 8$ and $M = 10$. At the modest accuracy level of 3, the total number of points quickly exceeds 10^3, which is usually considered to be an impractical number of samples. (Details of the construction of the sparse grids can be found in numerous references such as [22].)

After extensive testing, we determined that it is more appropriate to use the QMC method for the integrals in Eq. (2.7). We follow [16] for a short description of the QMC methods. The main idea of the QMC method is the construction of low-discrepancy sequences. For any integer $b \geq 2$, we define $\mathbb{Z}_b = \{0, 1, \ldots, b-1\}$. For any integer $n \geq 1$, let us write the unique b-ary representation of n as

$$n = \sum_{j=0}^{\infty} a_j(n) b^j, \quad a_j(n) \in \mathbb{Z}_b.$$

The radical inverse function $\phi_b(n)$ is defined to be

$$\phi_b(n) = \sum_{j=0}^{\infty} a_j(n) b^{-j-1}, \quad \forall n \geq 0.$$

Clearly, we have $0 \leq \phi_b(n) \leq 1$. Two of the most commonly used low-discrepancy sequences are defined based on the radical inverse functions. Let M be an arbitrary dimension and $b_1, \ldots, b_M$ coprime to each other. The Halton sequence is defined for each integer $n > 0$ as

$$z^{(n)} = [\phi_{b_1}(n), \ldots, \phi_{b_M}(n)] \in [0,1]^M$$

The definition of the Hammersley sequence is similar. Let M be the dimension, N be the length of the sequence, and $b_1, \ldots, b_{M-1}$ coprime to each other. The Hammersley sequence is defined for $n = 1, \ldots, N$ as

$$z^{(n)} = \left[\frac{n}{N}, \phi_{b_1}(n), \ldots, \phi_{b_{M-1}}(n)\right] \in [0,1]^M.$$

For a fixed sample size N, we can generate the samples $z^{(1)}, z^{(2)}, \ldots, z^{(N)}$ using a low-discrepancy sequence (in our numerical examples, we choose the Hammersley sequence due to its lower discrepancy). For each sample $z^{(i)}$, we use the algorithms described in Section 3.1 to compute the RCS $R_{z^{(i)}}(s_\ell)$ for $\ell = 0, 1, \ldots, N_s - 1$. Once they are ready, the statistical estimations of the mean and variance are given respectively by

$$\bar{R}_N(s_\ell) = \frac{1}{N} \sum_{i=1}^{N} R_{z^{(i)}}(s_\ell)$$

$$\bar{V}_N(s_\ell) = \frac{1}{N-1} \sum_{i=1}^{N} \left[R_{z^{(i)}}(s_\ell) - \bar{R}_N(s_\ell)\right]^2.$$

TABLE 1: Number of points in Smolyak sparse grids

Dimension M	Level 1	Level 2	Level 3	Level 4
$M = 8$	17	145	849	3937
$M = 10$	21	221	1581	8801

4. NUMERICAL EXPERIMENTS

4.1 Method and Error Estimates

In this section, we present the results of some numerical experiments. Recall that the uncertainty of the scatterer is modeled by

$$\partial D_z = \{x(t) = b(t) \cdot [1 + p(t, Z)], t \in [0, 2\pi)\}$$

where $b(t) = [b_1(t), b_2(t)]$ is the base geometry and $p(t, Z)$ is the (multiplicative) random perturbation. Two base shapes on which we have tested are the cylinder and the kite (Fig. 1). These objects were chosen because they are smooth and have a simple parametrization in the two-dimensional plane:

$$\text{Cylinder: } b(t) = [b_1(t), b_2(t)] = \frac{K}{2} [\cos(t), \sin(t)] .$$

$$\text{Kite: } b(t) = [b_1(t), b_2(t)] = \frac{K}{2} \left[\frac{\cos(t) + 0.65\cos(2t) - 0.65}{1.5}, \sin(t) \right] .$$

The perturbation $p(t, z)$ is modeled as follows. First, choose a set number of frequencies or modes $\{\xi_i\}_{i=1}^{M/2}$. For simplicity, we assume that each component Z_i of the random parameter $Z = (Z_1, \ldots, Z_M)$ has a uniform probability density function over the unit interval $[0, 1]$ (this assumption can certainly be removed by performing appropriate reparametrization to each Z_i). As a result, the joint probability density function for z is the constant one function over the M-dimensional cube $[0, 1]^M$. For a given sample $Z = (Z_1, \ldots, Z_M)$, the perturbation $p(t, Z)$ is defined as

$$p(t, Z) = \frac{\mu}{K} \sum_{i=1}^{M/2} \left[\left(Z_{2i-1} - \frac{1}{2} \right) \cos(\xi_i t) + \left(Z_{2i} - \frac{1}{2} \right) \sin(\xi_i t) \right] . \tag{4.1}$$

Depending on the choice of the frequencies $\{\xi_i\}_{i=1}^{M/2}$, $p(t, Z)$ can model both low- and high-frequency perturbations.

1. Low-frequency perturbation sets $\xi_i = i$ for $i = 1, 2, \ldots, M/2$; thus, the perturbation function does not have many oscillations and the resulting boundary ∂D does not have rough edges.
2. High-frequency perturbation sets $\xi_i = iK/M$ for $i = 1, 2, \ldots, M/2$. Here, the high-frequency range extends to modes that are comparable to the size of the scattering object in terms of wavelength and the resulting boundary ∂D exhibits small-scale oscillations.

In each case, we compare the uncertainty quantification results of the Monte Carlo method and the QMC method for a fixed sample size N. In Monte Carlo method, the random parameter sample $Z^{(i)} = (Z_1^{(i)}, \ldots, Z_M^{(i)})$ for $i = 1, 2, \ldots, N$ is generated randomly for each entry $Z_j^{(i)}$ for $j = 1, 2, \ldots, M$. In QMC method, the random parameter

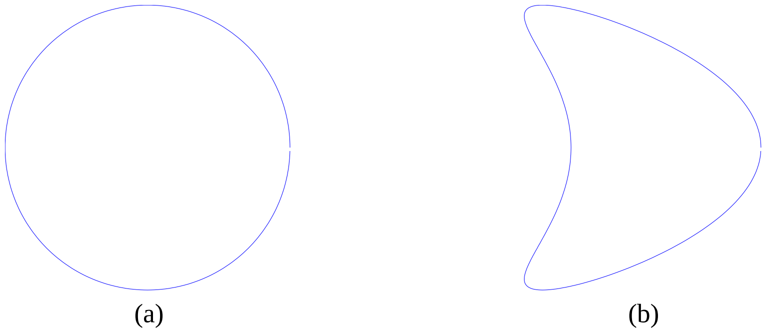
(a) (b)

FIG. 1: Base shape of the scatterers used in the test: (a) Circle and (b) kite

$Z^{(i)} = (Z_1^{(i)}, \ldots, Z_M^{(i)})$ for each i is constructed using the Hammersley sequence. For our simulations, the goal is to see how the estimations of the expected value and the variance of the radar cross section converge for the Monte Carlo and QMC methods. In each case, the statistical estimations of the mean and variance for a fixed sample size N are given respectively by

$$\bar{R}_N(s_\ell) = \frac{1}{N}\sum_{i=1}^{N} R_{z^{(i)}}(s_\ell)$$

$$\bar{V}_N(s_\ell) = \frac{1}{N-1}\sum_{i=1}^{N} \left[R_{z^{(i)}}(s_\ell) - \bar{R}_N(s_\ell)\right]^2 .$$

In order to measure the convergence rate depending on the sample size N, we estimate the error using the relative ℓ_2 norm. Suppose that $N_{\max}$ is the largest sample size used in the tests. Then for each fixed N, we define the errors $\varepsilon_{\bar{R},N}$ and $\varepsilon_{\bar{V},N}$ as

$$\varepsilon_{\bar{R},N} = \sqrt{\frac{\sum_{\ell=0}^{N_s-1} |\bar{R}_N(s_\ell) - \bar{R}_{N_{\max}}(s_\ell)|^2}{\sum_{\ell=0}^{N_s-1} |\bar{R}_{N_{\max}}(s_\ell)|^2}}$$

$$\varepsilon_{\bar{V},N} = \sqrt{\frac{\sum_{\ell=0}^{N_s-1} |\bar{V}_N(s_\ell) - \bar{V}_{N_{\max}}(s_\ell)|^2}{\sum_{\ell=0}^{N_s-1} |\bar{V}_{N_{\max}}(s_\ell)|^2}}$$

4.2 Numerical Results

In our tests, we set the diameter of the scatterer K to be 512, the number of random modes $M = 8$ and the perturbation amplitude in Eq. (4.1) $\mu = 0.1$. We choose the incident field to be a plane wave propagating in the x_1 direction, i.e.,

$$u_I(x) = e^{2\pi\imath x_1}, \tag{4.2}$$

where $x = (x_1, x_2)$; once again, the wavenumber is 2π and the wavelength λ is 1. We tested on both cylinder and kite geometries, using both low- and high-frequency perturbations mentioned earlier.

First, we show the results of stochastic collocation using Smolyak sparse grids for the random parameter, with accuracy level of 1, 2, and 3. Figure 2 shows close-ups of the variance curve for the low- and high-frequency perturbations of the kite, respectively. It is clear that stochastic collocation produces somewhat nonsensical results because it should be impossible to have a negative value for the variance of the RCS. This artifact is purely a result of utilizing negative weights in the quadrature of the random space; for this reason, stochastic collocation does not work well when the solution is highly oscillatory.

Next, we present the results of both the Monte Carlo and QMC methods. In order to measure the convergence, different sample sizes of $N = 64, 256, 1024$ are used with $N_{\max} = 1024$ for the highest-order accuracy. Figures 3 and 4 summarize the results of the cylinder for the low- and high-frequency perturbations, respectively. The errors in both cases are tabulated in Table 2. For low-frequency perturbations, the expectation and variance converge significantly faster for QMC when the sample size N increases. However, for the high-frequency perturbations, it is observed that the improvement in error for both quantities is modest at best; that is, the rougher the surface of the cylinder, the more difficult it is to accurately quantify the RCS. We have also performed tests for larger values of M and have observed similar results for both low- and high-frequency perturbations.

Figures 5 and 6 summarize the results of the kite for the high- and low-frequency perturbations, respectively, using QMC. The errors in both cases are tabulated in Table 3. The results suggests that, when the sample size N is quadrupled, the expectation for both the low frequency and high-frequency perturbations converge by a factor of 3 for the QMC method and by a factor of 2 for the standard Monte Carlo method. On the other hand, the convergence rates for the variance seem to be comparable for the two methods.

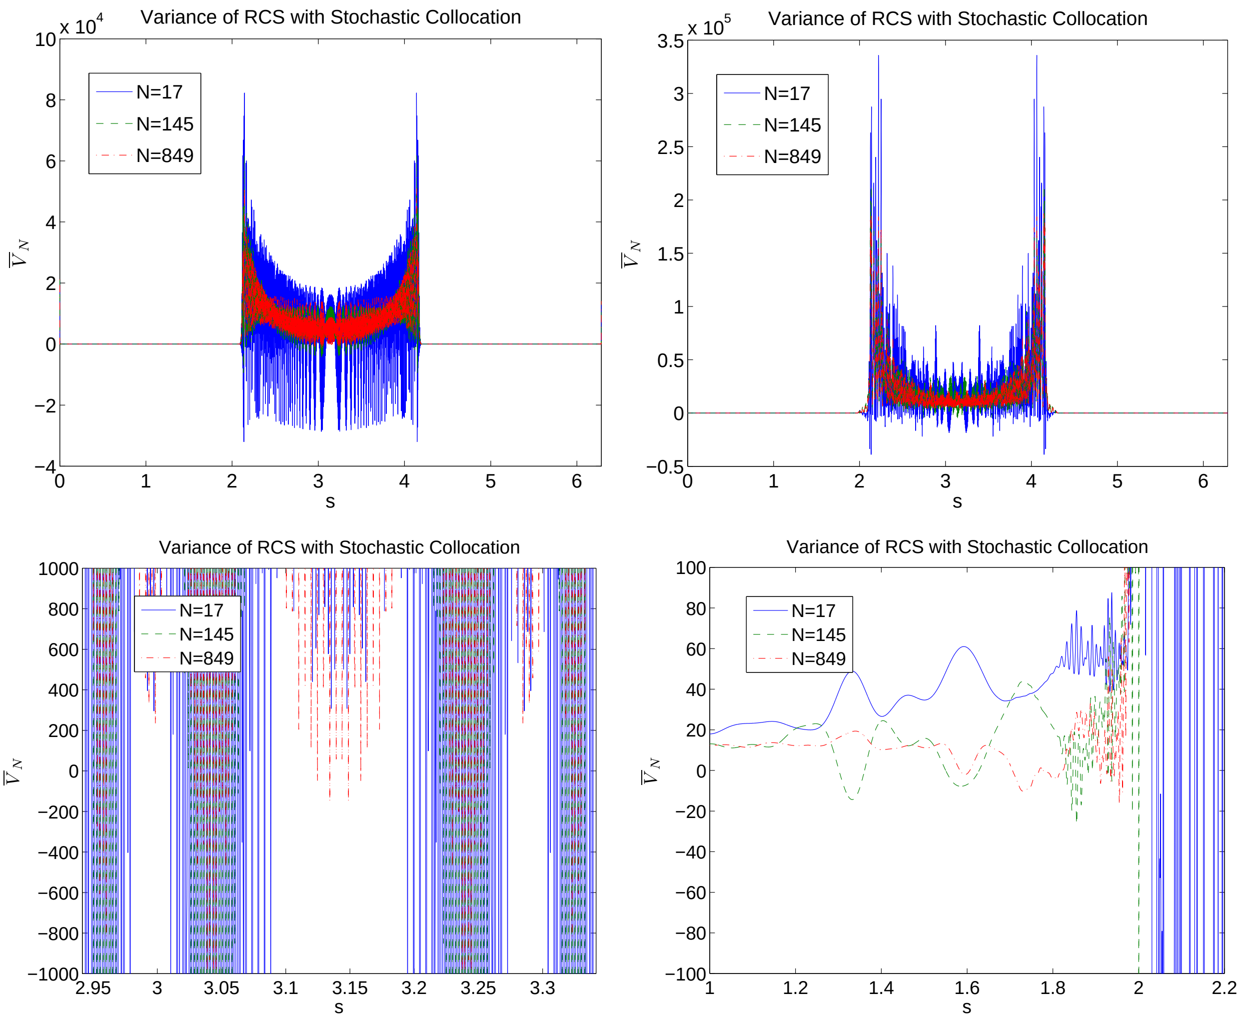

FIG. 2: Variance of the radar cross section for the kite, using stochastic collocation. (left) For low-frequency perturbations and (right) for high-frequency perturbations. For each type of perturbation, the top figure shows the full plot, while the bottom figure shows a close-up where the curves show negative variance

One difficulty we encountered in our numerical tests is the sensitivity of the RCS calculation varying with the size of the perturbation μ. For larger perturbations approaching the size of the operating wavelength, such convergence to the actual mean or variance proved to be quite difficult without having an inordinate number of samples. In order to achieve something sensible, especially for high-frequency problems, we found computationally that the perturbation size must satisfy $\mu \leq \lambda/5$.

5. CONCLUSION

In this paper, we presented a new numerical algorithm for quantifying the uncertainty of high-frequency acoustic scattering from scatterers with random shape in two-dimensional space. It allows one to effectively estimate the mean and variance of the random radar cross section in all directions. For each realization of the domain boundary, the boundary integral formulation is used with the standard Nyström discretization. The computation of the integral operator and the evaluation of the radar cross section are accelerated by the fast directional multilevel algorithm

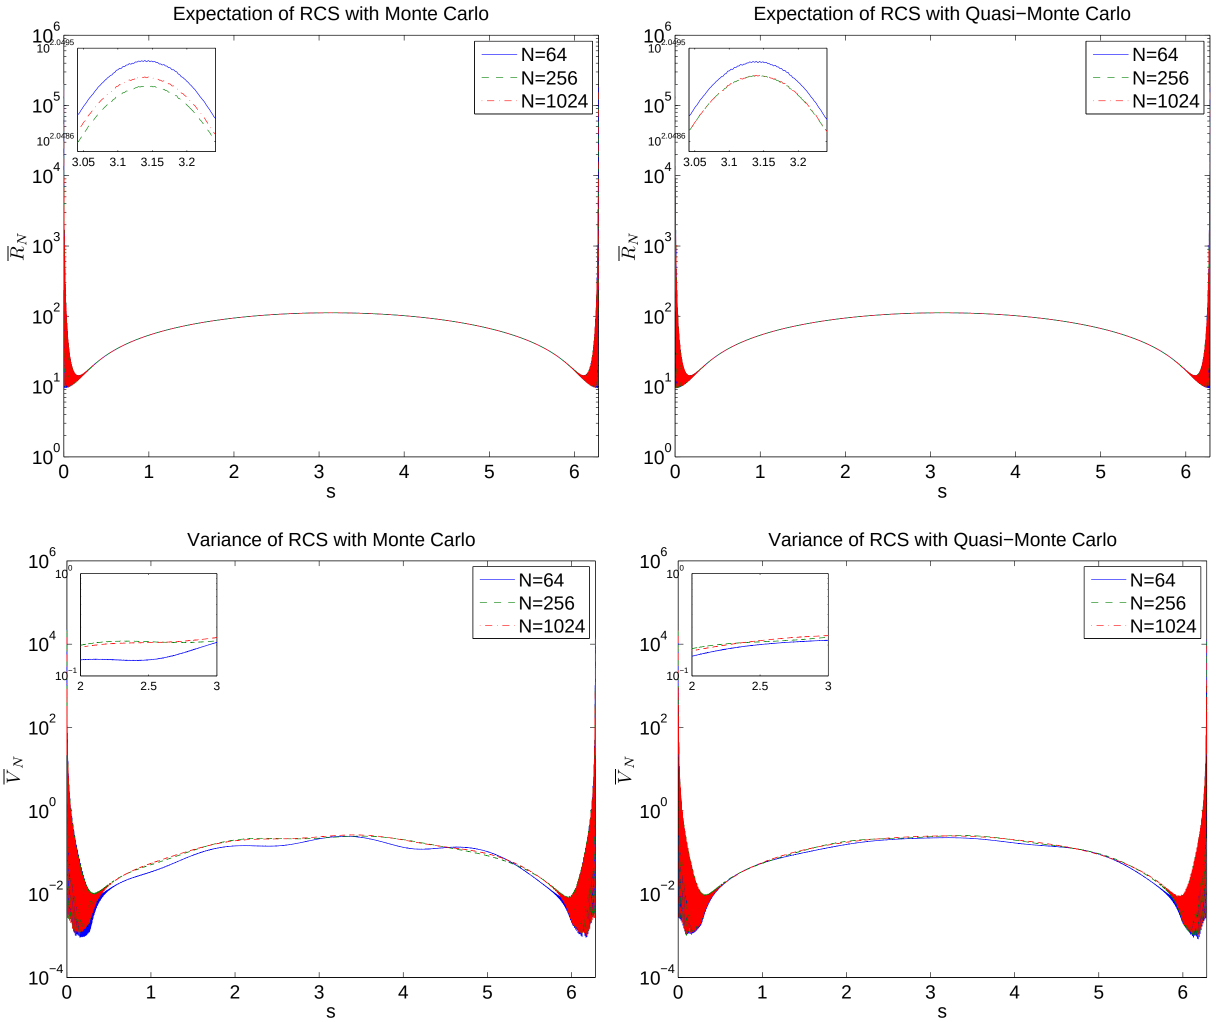

FIG. 3: Expectation and variance of the radar cross section for low-frequency perturbations on the cylinder. (left) Regular Monte Carlo and (right) use the Hammersley low-discrepancy sequence

and the butterfly algorithm (for the sparse Fourier transform). The statistical averaging is performed using the QMC method. When compared to the standard Monte Carlo method, the QMC method provides faster convergence to the mean and variance.

In our numerical tests, the random domain is modeled by a small random perturbation around a base shape. Numerical results suggest that the algorithm performs quite well when the perturbation is of low frequency compared to the wavelength of the scattering problem. More in-depth studies for high-frequency perturbations or large random perturbations are under investigation.

APPENDIX A. FAST ALGORITHMS

A.1 Fast Directional Multilevel Algorithm

Here, we briefly outline the directional method presented in [6]. Let $\{p_i\}_{i=1}^N \subset [-K/2, K/2]^2$ and $\{f_i\}_{i=1}^N$ be the sources located at $\{p_i\}_{i=1}^N$. The N-body problem of the Helmholtz kernel is to compute

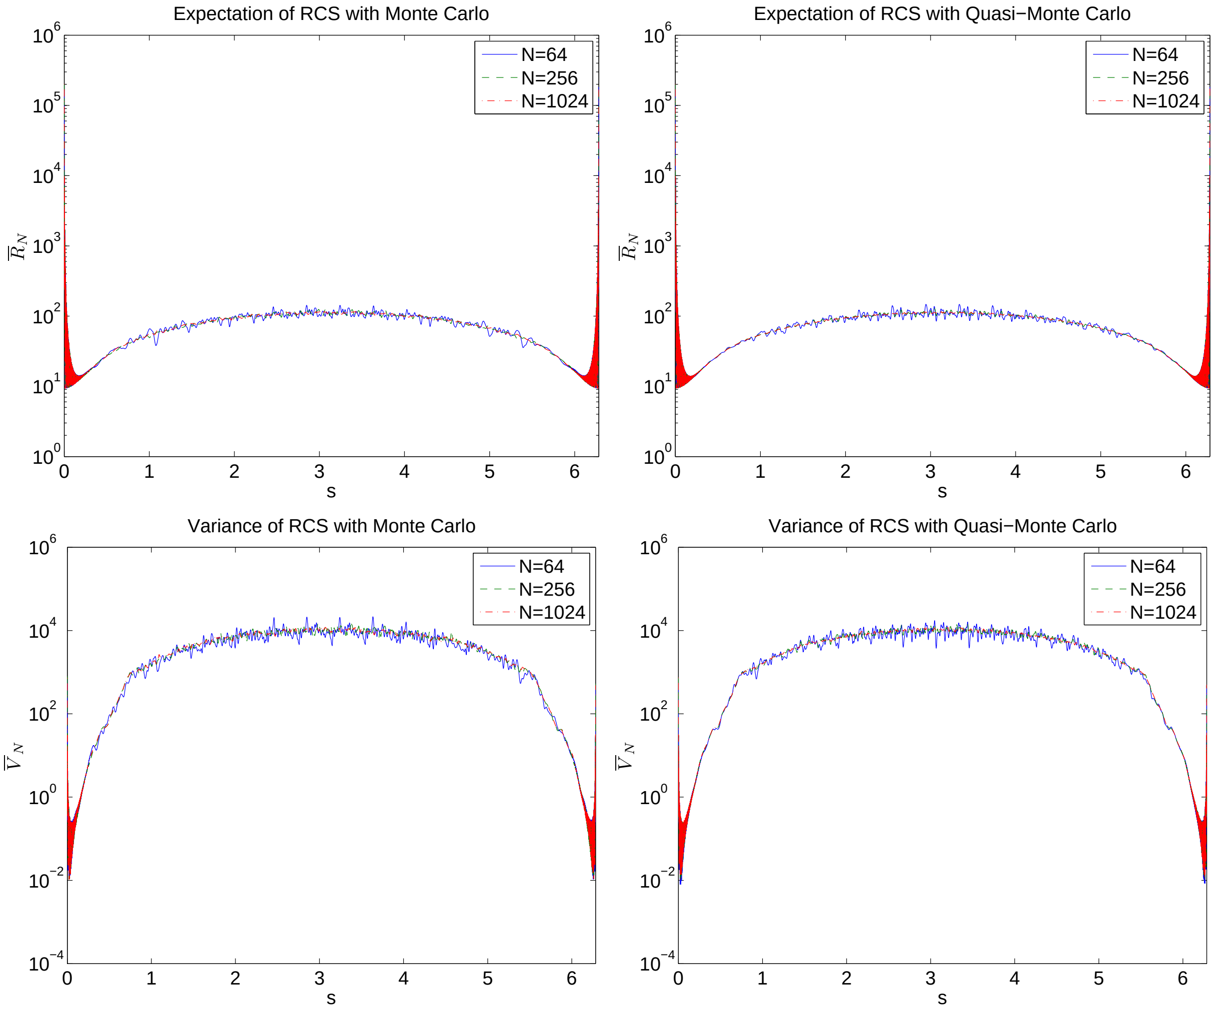

FIG. 4: Expectation and variance of the radar cross section for high-frequency perturbations on the cylinder

TABLE 2: Two-norm errors for the RCS of the cylinder geometry

	Low-frequency perturbations		High-frequency perturbations	
(Method,N)	$\varepsilon_{\bar{R},N}$	$\varepsilon_{\bar{V},N}$	$\varepsilon_{\bar{R},N}$	$\varepsilon_{\bar{V},N}$
MC, 64	6.17×10^{-5}	1.33×10^{-1}	2.35×10^{-3}	2.69×10^{-1}
MC, 256	1.50×10^{-5}	9.68×10^{-2}	1.28×10^{-3}	1.62×10^{-1}
QMC, 64	2.71×10^{-5}	8.06×10^{-2}	2.25×10^{-3}	2.27×10^{-1}
QMC, 256	4.80×10^{-6}	3.48×10^{-3}	9.32×10^{-4}	1.11×10^{-1}

$$u_i = \sum_{\substack{j=0 \\ j \neq i}}^{N} G(p_i, p_j) f_j, \tag{A.1}$$

for $i = 1, \ldots, N$, where $G(x, y)$ is the Green's function of the Helmholtz equation [Eq. (2.2)]. The main idea behind this approach is a directional low-rank property of the Helmholtz kernel. Consider a box B of width $w\lambda$ and a wedge

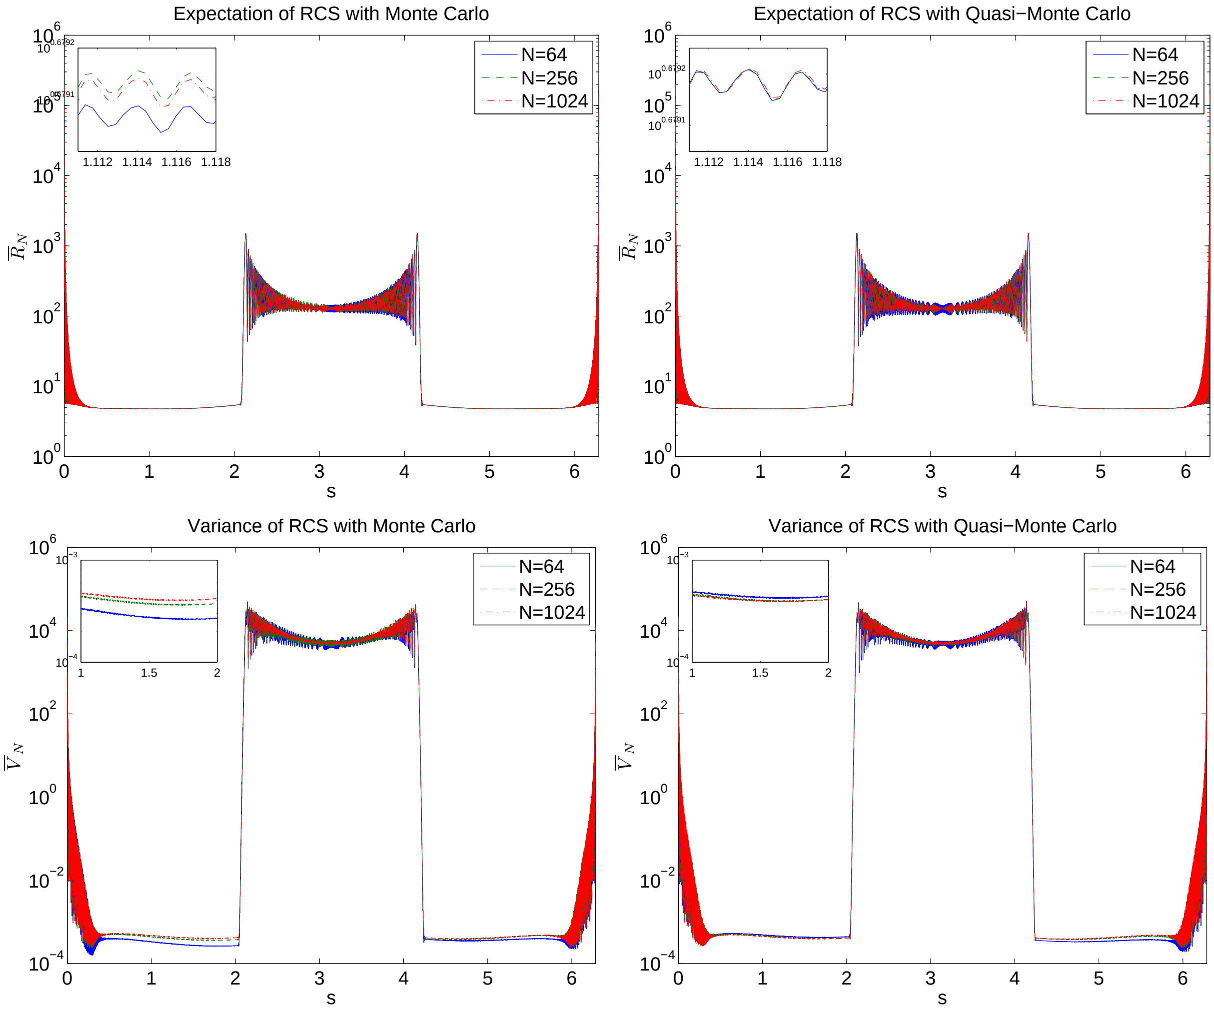

FIG. 5: Expectation and variance of the radar cross section for low-frequency perturbations on the kite scatterer

$W^{B,\ell}$ as illustrated in Fig. 7a. Because $W^{B,\ell}$ is centered at the vector ℓ with an opening angle of size $1/w$ and is a distance $w^2\lambda$ away from B, we say that $W^{B,\ell}$ and B follow a directional parabolic configuration.

We proved that for any accuracy ε, there exists a rank-r_ε separated approximation of $G(x,y)$, i.e., there are sets $\{y_q^{B,\ell}\}_{1\le q\le r_\varepsilon} \subset B$, $\{x_p^{B,\ell}\}_{1\le q\le r_\varepsilon} \subset W^{B,\ell}$, and a matrix $D=(d_{qp})_{1\le p,q\le r_\varepsilon}$ such that

$$\left| G(x,y) - \sum_{q=1}^{r_\varepsilon} G(x,y_q^{B,\ell}) \sum_{p=1}^{r_\varepsilon} d_{qp} G(x_p^{B,\ell},y) \right| \le \varepsilon \tag{A.2}$$

for $y \in B$ and $x \in W^{B,\ell}$, where the matrix $D=(d_{qp})_{1\le p,q\le r_\varepsilon}$ can be computed easily from $\{y_q^{B,\ell}\}$ and $\{x_p^{B,\ell}\}$. It is important to emphasize that the rank r_ε is independent of the size of B.

Suppose $\{f_i\}$ are the sources located at $\{y_i\}$ in B. After applying the separated approximation in Eq. (A.2) to $\{y_i\}$ and summing the approximations up with weights $\{f_i\}$, we obtain

$$\left| \sum_i G(x,y_i)f_i - \sum_{q=1}^{r_\varepsilon} G(x,y_q^{B,\ell}) \left[\sum_{p=1}^{r_\varepsilon} d_{qp} \sum_i G(x_p^{B,\ell},y_i)f_i \right] \right| = O(\varepsilon). \tag{A.3}$$

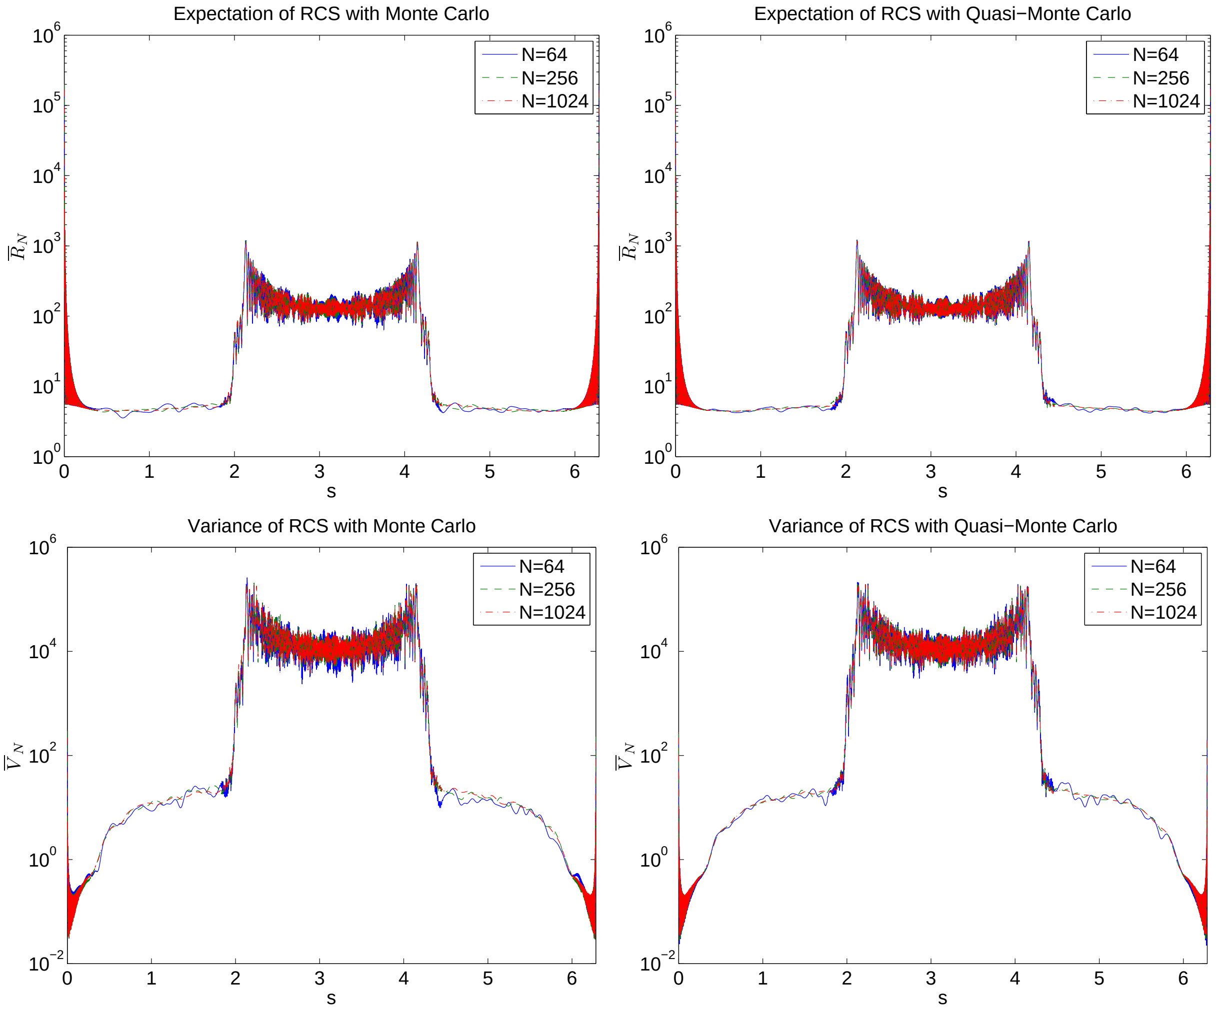

FIG. 6: The expectation and variance of the radar cross section for high-frequency perturbations on the kite scatterer

TABLE 3: Two-norm errors for the RCS of the kite geometry

	Low-frequency perturbations		High-frequency perturbations	
(Method,N)	$\varepsilon_{\bar{R},N}$	$\varepsilon_{\bar{V},N}$	$\varepsilon_{\bar{R},N}$	$\varepsilon_{\bar{V},N}$
MC, 64	1.91×10^{-3}	1.27×10^{-1}	3.39×10^{-3}	2.15×10^{-1}
MC, 256	1.07×10^{-3}	7.08×10^{-2}	1.84×10^{-3}	9.87×10^{-2}
QMC, 64	1.60×10^{-3}	9.90×10^{-2}	2.01×10^{-3}	1.29×10^{-1}
QMC, 256	5.01×10^{-4}	5.08×10^{-2}	6.45×10^{-4}	6.01×10^{-2}

This states that we can place a set of sources $\left\{f_q^{B,\ell} := \sum_p d_{qp} \sum_i G(x_p^{B,\ell}, y_i) f_i\right\}$ at points $\{y_q^{B,\ell}\}$ in order to reproduce the potential generated by the sources $\{f_i\}$ located at points $\{y_i\}$. We call these sources the directional equivalent sources of B in direction ℓ. In our algorithm, these equivalent sources play the role of the multipole expansions in the FMM algorithm [23, 24]. It is clear from Eq. (A.3) that the computation of $\{f_q^{B,\ell}\}$ utilizes only kernel evaluation and small matrix-vector multiplications.

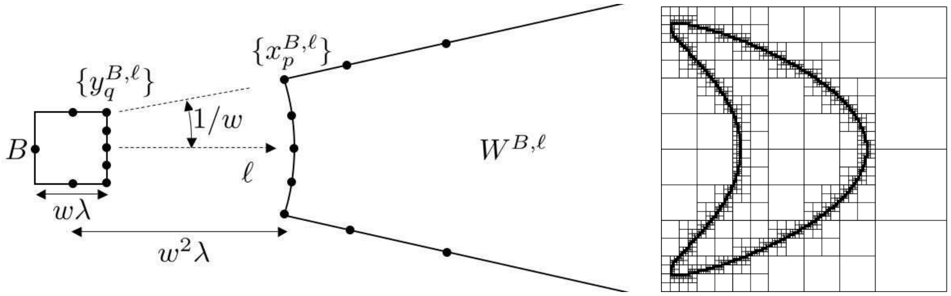

(a) (b)

Fig 7: (a) B and $W^{B,\ell}$ follows a directional parabolic configuration and (b) quadtree of a kite-shaped scatterer

Let us now reverse the role of the source and the target. Suppose that we have a set of sources $\{f_i\}$ located at points $\{x_i\}$ in $W^{B,\ell}$. Because $G(x,y) = G(y,x)$, we have

$$\left|\sum_i G(y,x_i)f_i - \sum_{p=1}^{r_\varepsilon} G(y,x_p^{B,\ell})\left[\sum_{q=1}^{r_\varepsilon} d_{qp}\sum_i G(y_q^{B,\ell},x_i)f_i\right]\right| = O(\varepsilon).$$

This means that we can reproduce the potential generated at any $y \in B$ by using the auxiliary potentials $\{u_q^{B,\ell} := \sum_i G(y_q^{B,\ell},x_i)f_i\}$. These potentials are called the directional check potentials of B in direction ℓ. In our algorithm, these potentials play the role of the local expansions of the FMM algorithm.

Our algorithm starts by constructing a quadtree that contains the whole scatterer (see Fig. 7b). A box B of width $w\lambda$ is said to be in the low-frequency regime if $w < 1$ and in the high-frequency regime if $w \geq 1$. In the high-frequency regime of the quadtree, the domain is partitioned uniformly without any adaptivity. In the low-frequency regime, a square B is partitioned as long as the number of points in B is greater than a fixed constant N_p. In order to use the low-rank separated representation in Eq. (A.2) in the high-frequency regime, we define the far-field F^B of a box B to be the region that is separated from B by a distance of $w^2\lambda$. A box A is said to be in the interaction list of B if A is in B's far-field but not in the far-field of B's parent. F^B is further partitioned into a group of directional wedges $\{W^{B,\ell}\}$, each in a cone of spanning angle $1/w$. Because the wedges of the parent box and the child box are nested, we are able to construct M2M, M2L, and L2L translations of $O(1)$ complexity as in the FMM algorithm. However, it is important to note that these translations are now directional. In the low frequency regime, the directional equivalent sources and check potentials reduce to the nondirectional equivalent sources and check potentials introduced in [25].

Putting all of these components together gives us the following directional multilevel algorithm:

1. Construct the quadtree. In the high-frequency regime, the squares are partitioned uniformly. In the low-frequency regime, a leaf square contains at most N_p points.

2. Travel up the low-frequency part of the octree. For each square B, compute the nondirectional equivalent sources following [25].

3. Travel up the high-frequency part of the quadtree. For each square B and each direction ℓ, compute $\{f_q^{B,\ell}\}$ using the directional M2M translation. We skip the squares with width greater than $\sqrt{k}\lambda$ because their interaction lists are empty.

4. Travel down the high-frequency part of the quadtree. For each square B and each direction ℓ, perform the following two steps:

    a. Transform $\{f_q^{A,\ell}\}$ of all of the squares $\{A\}$ in B's interaction list and in direction ℓ via the directional M2L translation. Next, add the result to $\{u_q^{B,\ell}\}$.

   b. Perform the directional L2L translation to transform $\{u_q^{B,\ell}\}$ into the incoming check potentials for B's children.

5. Travel down the low-frequency part of the quadtree. For each square B:

   a. Transform the nondirectional equivalent sources of all the squares $\{A\}$ in B's interaction list via the low-frequency nondirectional M2L operator. Next, add the result to the nondirectional check potentials.

   b. Perform the low-frequency directional L2L translation. Depending on whether B is a leaf square or not, add the result to the nondirectional check potentials of B's children or to the potentials at the original points inside B.

It is shown in [5, 6] that for a point set $\{p_i\}$ obtained from discretizing a scatterer boundary curve in $[-K/2, K/2]^2$, the overall cost of this algorithm is $O(K \log K)$.

A.2 Butterfly Algorithm for Sparse Fourier Transform

Recall that a sparse Fourier transform [7] is a computation of potentials in the form

$$u_i = \sum_j e^{2\pi\imath x_i \cdot k_j/N} f_j, \tag{A.4}$$

where $\{k_j\}$ is a set of $O(N)$ points sampled from a smooth curve in the Fourier domain $[-N/2, N/2]^2$, $\{f_j\}$ are the sources at $\{k_j\}$, and $\{x_i\} \subset [-N/2, N/2]^2$ is a set of $O(N)$ points sampled from another smooth curve in the spatial domain $[-N/2, N/2]^2$. The algorithm proposed in [7] first constructs adaptive quadtrees T_X and T_K for the sets $\{x_i\}$ and $\{k_j\}$, respectively. The quadtree T_X takes $[-N/2, N/2]^2$ as the top level square. Each square is partitioned recursively into four identical child squares until all leaf squares are of unit size, and only the squares that contain points in $\{x_i\}$ are kept. The quadtree T_K is constructed in the same way with $[-N/2, N/2]^2$ as the top level square and $\{k_j\}$ as the point set.

The main idea of the algorithm is based on the following geometric observation. Let A and B be two squares in T_X and T_K, respectively. If the product of their widths, $w^A w^B$, is bounded by N, then the interaction $e^{2\pi\imath x\cdot k/N}$ for $x \in A$ and $k \in K$ is numerically low rank. More precisely, for any fixed ε, there exists a number $T_\varepsilon = O(\log(1/\varepsilon))$ and two sets of functions $\{\alpha_t^{AB}(x)\}_{1\le t\le T_\varepsilon}$ and $\{\beta_t^{AB}(x)\}_{1\le t\le T_\varepsilon}$ such that

$$\left| e^{2\pi\imath x\cdot k/N} - \sum_{t=1}^{T_\varepsilon} \alpha_t^{AB}(x)\beta_t^{AB}(k) \right| \le \varepsilon.$$

In fact, the function $\alpha_t^{AB}(x)$ can be chosen to be of form $e^{2\pi\imath x\cdot k_t^B/N}$, where $\{k_t^B\}_{1\le t\le T_\varepsilon}$ belong to a two-dimensional Chebyshev grid of the square B. For a fixed accuracy ε, the size of this Chebyshev grid, T_ε, is independent of N (see [7] for details).

Let us define the partial sum $u^B(x)$ by

$$u^B(x) = \sum_{\xi_j \in B} e^{2\pi\imath x\cdot \xi_j/N} f_j \tag{A.5}$$

with the sum restricted to k inside B. The geometric observation implies that, for A and B with $w^A w^B \le N$, the restriction of $u^B(x)$ to $x \in A$ can be approximated by placing a set of equivalent sources $\{f_t^{AB}\}_{1\le t\le T_\varepsilon}$ at location $\{k_t^B\}_{1\le t\le T_\varepsilon}$. The computation of the equivalent sources $\{f_t^{AB}\}$ is done by equating the partial sum $u^B(x)$ and the potential generated by $\{f_t^{AB}\}$ at a Chebyshev grid $\{x_s^A\}$ inside A. More precisely, one solves for $\{f_t^{AB}\}$ from the following equation:

$$u_s^{AB} := u^B(x_s^A) = \sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_s^A\cdot k_t^B/N} f_t^{AB}, \quad 1 \le s, t \le T_\varepsilon.$$

Solving this system requires inverting the matrix $(e^{2\pi\imath x_s^A\cdot k_t^B/N})_{st}$. However, due to the translation-invariant property of the Fourier kernel, the matrices to be inverted for different combinations of A and B are almost identical. Furthermore, the solution of $\{f_t^{AB}\}$ can be accelerated using the tensor-product structure of the two-dimensional Fourier kernel.

Evaluating $u_s^{AB} := u^B(x_s^A)$ directly using Eq. (A.5) is computationally expensive when the square B is large. The next ingredient of the butterfly algorithm addresses how to do this efficiently. Let P be the parent of A and $B_c, c = 1, \ldots, 4$ be the children of B. From the definition [Eq. (A.5)],

$$u^B(x) = \sum_{c=1}^{4} u^{B_c}(x).$$

Now, suppose that the equivalent sources $\{f_t^{PB_c}\}$ are available already; then the quantities $\{u_s^{AB}\}$ can be approximated by

$$u_s^{AB} := u^B(x_s^A) = \sum_{c=1}^{4} u^{B_c}(x_s^A) \approx \sum_{c=1}^{4}\left(\sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_s^A k_t^{B_c}/N} f_t^{PB_c}\right), \quad 1 \le s \le T_\varepsilon$$

based on precisely the definition of the equivalent sources. This offers a much more efficient way for computing $\{u_s^{AB}\}$ as the size of the Chebyshev grid is a constant.

After putting these components together, the butterfly algorithm in [7] is in fact a systematic way to construct $\{f_t^{AB}\}_{1\le t\le T_\varepsilon}$ for all pairs of squares $A \in T_X$ and $B \in T_K$ with $w^A w^B = N$. The algorithm consists of the following steps:

1. Construct the quadtrees T_X and T_K for the point sets X and K, respectively. These trees are constructed adaptively, and all the leaf squares are of unit size.

2. Let A be the root square of T_X. For each leaf square B of T_K, compute

   $$u_s^{AB} = \sum_{k_j\in B} e^{2\pi\imath x_s^A\cdot k_j/N} f_j, \quad 1 \le s \le T_\varepsilon$$

   and solve for $\{f_t^{AB}\}_{1\le t\le T_\varepsilon}$ from

   $$u_s^{AB} = \sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_s^A\cdot k_t^B/N} f_t^{AB}, \quad 1 \le s \le T_\varepsilon.$$

3. For each $\ell = 1, 2, \ldots, \log N$, construct the equivalent sources $\{f_t^{AB}\}$ for each pair (A, B) with A at level ℓ of T_X and B at level $(\log N - \ell)$ of T_K. Let P be the parent of A, and B_c, $c = 1, \ldots, 4$ be the children of B. Compute u_s^{AB} using

   $$u_s^{AB} = \sum_{c=1}^{4}\left(\sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_s^A k_t^{B_c}/N} f_t^{PB_c}\right), \quad 1 \le s \le T_\varepsilon.$$

   Next, solve for $\{f_t^{AB}\}_{1\le t\le T_\varepsilon}$ from

   $$u_s^{AB} = \sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_s^A\cdot k_t^B/N} f_t^{AB}, \quad 1 \le s \le T_\varepsilon.$$

4. Finally, let B be the root square of T_K. For each leaf square A of T_X and for each $x_i \in A$, set

   $$u_i = \sum_{t=1}^{T_\varepsilon} e^{2\pi\imath x_i\cdot k_t^B/N} f_t^{AB}.$$

Under the assumption that the sets $\{x_i\}$ and $\{k_j\}$ are both of order $O(N)$, the overall cost of the butterfly algorithm is $O(N \log N)$, which is almost linear. We refer to [7] for the detailed complexity analysis.

REFERENCES

1. Rokhlin, V., Rapid solution of integral equations of scattering theory in two dimensions, *J. Comput. Phys.*, 86(2):414–439, 1990.
2. Rokhlin, V., Diagonal forms of translation operators for the Helmholtz equation in three dimensions, *Appl. Comput. Harmon. Anal.*, 1(1):82–93, 1993.
3. Cheng, H., Crutchfield, W. Y., Gimbutas, Z., Greengard, L. F., Ethridge, J. F., Huang, J., Rokhlin, V., Yarvin, N., and Zhao, J., A wideband fast multipole method for the Helmholtz equation in three dimensions, *J. Comput. Phys.*, 216(1):300–325, 2006.
4. Bruno, O. P. and Kunyansky, L. A., A fast, high-order algorithm for the solution of surface scattering problems: basic implementation, tests, and applications, *J. Comput. Phys.*, 169(1):80–110, 2001.
5. Engquist, B. and Ying, L., Fast directional multilevel algorithms for oscillatory kernels, *SIAM J. Sci. Comput.*, 29(4):1710–1737, 2008.
6. Engquist, B. and Ying, L., A fast directional algorithm for high frequency acoustic scattering in two dimensions, *Commun. Math. Sci.*, 7(2):327–345, 2009.
7. Ying, L., Sparse Fourier transform via butterfly algorithm, *SIAM J. Sci. Comput.*, 31(3):1678–1694, 2009.
8. Metropolis, N. and Ulam, S., The monte carlo method, *J. Am. Stat. Assoc.*, 44(247):335–341, 1949.
9. Ghanem, R. G. and Spanos, P. D., *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, 1991.
10. Xiu, D. and Karniadakis, G. E., The Wiener-Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.*, 24(2):619–644, 2002.
11. Xiu, D. and Hesthaven, J., High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.*, 27(3):1118–1139, 2005.
12. Xiu, D., Fast numerical methods for stochastic computations: A review, *Commun. Comput. Phys.*, 5(2–4):242–272, 2009.
13. Lin, G., Su, C., and Karniadakis, G., Random roughness enhances lift in supersonic flow, *Phys. Rev. Lett.*, 99:104501, 2007.
14. Lin, G., Su, C., and Karniadakis, G., Stochastic modeling of random roughness in shock scattering problems: theory and simulations, *Comput. Methods Appl. Math. Eng.*, 197(43–44):3420–3434, 2008.
15. Xiu, D. and Shen, J., An efficient spectral method for acoustic scattering from rough surfaces, *Commun. Comput. Phys.*, 2(1):54–72, 2006.
16. Niederreiter, H., *Random Number Generation and Quasi-Monte Carlo Methods*, SIAM, Philadelphia, 1992.
17. Caflisch, R. E., Monte Carlo and quasi-Monte Carlo methods., *Acta numerica, 1998*, vol. 7, Cambridge Univesity Press, Cambridge, England, pp. 1–49, 1998.
18. Colton, D. and Kress, R., *Inverse Acoustic and Electromagnetic Scattering Theory*, Springer-Verlag, 1998.
19. Shinozuka, M. and Deodatis, G., Simulation of stochastic processes by spectral representation, *Appl. Mech. Rev.*, 44(4):191–204, 1991.
20. Loève, M., *Probability Theory. I*, 4 ed., Springer-Verlag, New York, 1977.
21. Kapur, S. and Rokhlin, V., High-order corrected trapezoidal quadrature rules for singular functions, *SIAM J. Numer. Anal.*, 34(4):1331–1356, 1997.
22. Novak, E. and Ritter, K., High dimensional integration of smooth functions over cubes, *Numer. Math.*, 75:79–97, 1996.
23. Greengard, L., *The Rapid Evaluation of Potential Fields in Particle Systems*, ACM Distinguished Dissertations, MIT Press, Cambridge, MA, 1988.
24. Greengard, L. and Rokhlin, V., A fast algorithm for particle simulations, *J. Comput. Phys.*, 73(2):325–348, 1987.
25. Ying, L., Biros, G., and Zorin, D., A kernel-independent adaptive fast multipole algorithm in two and three dimensions, *J. Comput. Phys.*, 196(2):591–626, 2004.