MULTISCALE MODEL. SIMUL.
Vol. 13, No. 2, pp. 614–631

A MULTISCALE BUTTERFLY ALGORITHM FOR MULTIDIMENSIONAL FOURIER INTEGRAL OPERATORS*

YINGZHOU LI†, HAIZHAO YANG‡, AND LEXING YING§

Abstract. This paper presents an efficient multiscale butterfly algorithm for computing Fourier integral operators (FIOs) of the form $(\mathcal{L}f)(x) = \int_{\mathbb{R}^d} a(x,\xi)e^{2\pi\imath\Phi(x,\xi)}\widehat{f}(\xi)d\xi$, where $\Phi(x,\xi)$ is a phase function, $a(x,\xi)$ is an amplitude function, and $f(x)$ is a given input. The frequency domain is hierarchically decomposed into a union of Cartesian coronas. The integral kernel $a(x,\xi)e^{2\pi\imath\Phi(x,\xi)}$ in each corona satisfies a special low-rank property that enables the application of a butterfly algorithm on the Cartesian phase-space grid. This leads to an algorithm with quasi-linear operation complexity and linear memory complexity. Different from previous butterfly methods for the FIOs, this new approach is simple and reduces the computational cost by avoiding extra coordinate transformations. Numerical examples in two and three dimensions are provided to demonstrate the practical advantages of the new algorithm.

Key words. Fourier integral operators, the butterfly algorithm, hierarchical decomposition, separated representation

AMS subject classifications. 44A55, 65R10, 65T50

DOI. 10.1137/140997658

1. Introduction.

This paper is concerned with the rapid application of Fourier integral operators (FIOs), which are defined as

$$(\mathcal{L}f)(x) = \int_{\mathbb{R}^d} a(x,\xi)e^{2\pi\imath\Phi(x,\xi)}\widehat{f}(\xi)d\xi, \tag{1}$$

where

- $a(x,\xi)$ is an amplitude function that is smooth in both x and ξ,
- $\Phi(x,\xi)$ is a phase function that is smooth in (x,ξ) for $\xi \neq 0$ and obeys the homogeneity condition of degree 1 in ξ, namely, $\Phi(x,\lambda\xi) = \lambda\Phi(x,\xi)$ for each $\lambda > 0$, and
- $\widehat{f}$ is the Fourier transform of the input f defined by

$$\widehat{f}(\xi) = \int_{\mathbb{R}^d} e^{-2\pi\imath x\cdot\xi} f(x)dx.$$

We refer the reader to Chapter IX of [26] for the analytic properties and results of the FIOs. The computation of FIOs appears quite often in the numerical solution of wave equations and related applications in computational geophysics. In a typical setting, it is often assumed that the problem is periodic (i.e., $a(x,\xi)$, $\Phi(x,\xi)$, and $f(x)$ are all periodic in x) or the function $f(x)$ decays sufficiently fast so that one can embed the

*Received by the editors November 26, 2014; accepted for publication (in revised form) April 7, 2015; published electronically June 11, 2015. This work was partially supported by the National Science Foundation under award DMS-1328230 and the U.S. Department of Energy's Advanced Scientific Computing Research program under award DE-FC02-13ER26134/DE-SC0009409.
http://www.siam.org/journals/mms/13-2/99765.html

†ICME, Stanford University, Stanford, CA 94305 (ryanli@stanford.edu).

‡Department of Mathematics, Stanford University, Stanford, CA 94305 (haizhao@math.stanford.edu).

§Department of Mathematics, Stanford University, Stanford, CA 94305, and ICME, Stanford University, Stanford, CA 94305 (lexing@stanford.edu).

problem in a sufficiently large periodic cell. A simple discretization in two dimensions considers functions f given on a Cartesian grid

$$X = \left\{ x = \left(\frac{n_1}{N}, \frac{n_2}{N} \right), 0 \le n_1, n_2 < N \text{ with } n_1, n_2 \in \mathbb{Z} \right\} \tag{2}$$

in a unit square and defines the discrete FIO by

$$(Lf)(x) = \sum_{\xi \in \Omega} a(x,\xi) e^{2\pi \imath \Phi(x,\xi)} \widehat{f}(\xi), \quad x \in X, \tag{3}$$

where

$$\Omega = \left\{ \xi = (n_1, n_2), -\frac{N}{2} \le n_1, n_2 < \frac{N}{2} \text{ with } n_1, n_2 \in \mathbb{Z} \right\}, \tag{4}$$

and $\widehat{f}$ is the discrete Fourier transform of f,

$$\widehat{f}(\xi) = \frac{1}{N^2} \sum_{x \in X} e^{-2\pi \imath x \cdot \xi} f(x).$$

The current form in (3) can be regarded as a discretization of the continuous FIO using the trapezoidal rule. In real applications, the location of grid points could be nonuniform and its quadrature rule is application dependent. In this case, the weights of the quadrature rule are absorbed in the input data $\widehat{f}$ and our algorithm applies with mild changes, as we shall discuss later.

In most examples, since $a(x,\xi)$ is a smooth symbol of order zero and type $(1,0)$ [4, 7, 16, 19, 28], $a(x,\xi)$ is numerically low-rank in the joint X and Ω domain and its numerical treatment is relatively easy. Therefore, we will simplify the problem by assuming $a(x,\xi) = 1$ in the following analysis and the algorithm description. We refer the reader to [8] for discussion on how to deal with a nonconstant amplitude function. Under this assumption, the discrete FIO discussed in this paper takes the following form:

$$(Lf)(x) = \sum_{\xi \in \Omega} e^{2\pi \imath \Phi(x,\xi)} \widehat{f}(\xi), \quad x \in X. \tag{5}$$

A direct computation of (5) takes $\mathcal{O}(N^4)$ operations, which is quadratic in the number of degrees of freedom N^2. Hence, a practical need is to design efficient and accurate algorithms to evaluate (5). This research topic is of great interest for computing wave equations, especially in geophysics [17, 20, 27, 29].

1.1. Previous work. An earlier method for the rapid computation of general FIOs is the algorithm for two-dimensional problems proposed in [7]. This method starts by partitioning the frequency domain Ω into $\mathcal{O}(\sqrt{N})$ wedges of equal angle. The integral (5) *restricted to each wedge* is then factorized into two components, both of which can be handled efficiently. The first one has a low-rank structure that leads to an $\mathcal{O}(N^2 \log N)$ fast computation, while the second one is a nonuniform Fourier transform which can be evaluated in $\mathcal{O}(N^2 \log N)$ steps with the algorithms developed in [1, 18, 25]. Summing the computational cost over all $\mathcal{O}(\sqrt{N})$ wedges gives an $\mathcal{O}(N^{2.5} \log N)$ computational cost.

Shortly after, an algorithm with quasi-linear complexity for general FIOs was proposed in [8] using the framework of the butterfly algorithms in [22, 23, 24]. This

approach introduces a polar coordinate transformation in the frequency domain to remove the singularity of $\Phi(x,\xi)$ at $\xi = 0$, proves the existence of low-rank separated approximations between certain pairs of spatial and frequency domains, and implements the low-rank approximations with oscillatory Chebyshev interpolations. The resulting algorithm evaluates (5) with $\mathcal{O}(N^2 \log N)$ operations and $\mathcal{O}(N^2)$ memory, both essentially linear in terms of the number of unknowns.

Another related research direction seeks sparse representations of the FIOs using modern basis functions from harmonic analysis. A sparse representation allows fast matrix-vector products in the transformed domain. Local Fourier transforms [3, 5, 12], wavelet-packet transforms [21], the curvelet transform [6, 9, 10, 11], the wave atom frame [14, 15], and the wave packet frame [2, 13] have been investigated for the purpose of operator sparsification. In spite of favorable asymptotic behaviors, the actual representations of the FIOs typically have a large prefactor constant in terms of both the computational time and the memory requirement. This makes them less competitive compared to the approaches in [7, 8].

1.2. Motivation. The main motivation of the current work is to improve the performance of the butterfly algorithm in [8]. As we pointed out earlier, this algorithm starts by applying a polar coordinate transformation in the frequency domain to remove the singularity of the phase function at $\xi = 0$. For this reason, we refer to this algorithm as the *polar butterfly algorithm*. More precisely, the polar butterfly algorithm introduces a polar-Cartesian coordinate transformation $T : (p_1, p_2) \to (\xi_1, \xi_2)$ such that

$$\xi = (\xi_1, \xi_2) = \frac{\sqrt{2}}{2} N p_1 e^{2\pi \imath p_2}, \quad e^{2\pi \imath p_2} = (\cos 2\pi p_2, \sin 2\pi p_2). \tag{6}$$

Let $P = T^{-1}(\Omega)$. By definition, each point $p = (p_1, p_2) \in P$ belongs to $[0,1]^2$. The new phase function $\Psi(x,p)$ in the p variable is now given by

$$\Psi(x,p) := \frac{1}{N}\Phi(x,\xi) = \frac{\sqrt{2}}{2}\Phi(x, e^{2\pi \imath p_2}) p_1, \tag{7}$$

where the last identity comes from the homogeneity of $\Phi(x,\xi)$ in ξ. Thus, computing (5) is equivalent to evaluating

$$(Lf)(x) = \sum_{\xi \in \Omega} e^{2\pi \imath \Phi(x,\xi)} \widehat{f}(\xi) = \sum_{p \in P} e^{2\pi \imath N \Psi(x,p)} \widehat{f}(T(p)). \tag{8}$$

The new phase function $\Psi(x,p)$ is smooth in the whole domain $(x,p) \subset [0,1]^2 \times [0,1]^2$ since $\Phi(x,\xi)$ is smooth in (x,ξ) for $\xi \neq 0$. This smoothness guarantees a low-rank separated approximation of $e^{2\pi \imath N \Psi(x,p)}$ when x and p are properly restricted to certain subdomains in $X \times P$ under certain geometric configuration. This low-rank property allows for the application of the butterfly algorithm in [30] and results in a fast algorithm with an $\mathcal{O}(N^2 \log N)$ computational complexity and an $\mathcal{O}(N^2)$ memory complexity.

However, the application of this polar-Cartesian transformation comes with several drawbacks, which result in a large prefactor of the computational complexity. First, due to the polar grid in the frequency domain, the points in P for the butterfly algorithm are irregularly distributed and a separate Chebyshev interpolation matrix is required for the evaluation at each point. In order to avoid the memory bottleneck from storing these interpolation matrices, the polar butterfly algorithm

generates these interpolation matrices on-the-fly during the evaluation. This turns out to be expensive in the operation count. Second, since the amplitude and phase functions are often written in the Cartesian coordinates, the polar butterfly algorithm applies the polar-Cartesian transformation for each kernel evaluation. Finally, in order to maintain a reasonable accuracy, the polar butterfly algorithm divides the frequency domain into multiple parts and applies the same butterfly algorithm to each part separately. This also increases the actual running time by a nontrivial constant factor.

1.3. Our contribution. Those drawbacks of the polar butterfly algorithm motivate us to propose a multiscale butterfly algorithm using a Cartesian grid in both the spatial and the frequency domains. To deal with the singularity of the kernel $\Phi(x,\xi)$ at $\xi = 0$, we hierarchically decompose the frequency domain into a union of nonoverlapping Cartesian coronas with a common center $\xi = 0$ (see Figure 1). More precisely, we define

$$\Omega_j = \left\{ (n_1, n_2) : \frac{N}{2^{j+1}} < \max(|n_1|, |n_2|) \le \frac{N}{2^j} \right\} \cap \Omega$$

for $j = 1, \ldots, \log N - s$, where s is just a small constant integer. The domain $\Omega_d = \Omega \setminus \cup_j \Omega_j$ is the remaining square grid at the center of constant size. Following this decomposition of the frequency domain, one can write (5) accordingly as

$$(Lf)(x) = \sum_j \left(\sum_{\xi \in \Omega_j} e^{2\pi i \Phi(x,\xi)} \widehat{f}(\xi) \right) + \sum_{\xi \in \Omega_d} e^{2\pi i \Phi(x,\xi)} \widehat{f}(\xi). \tag{9}$$

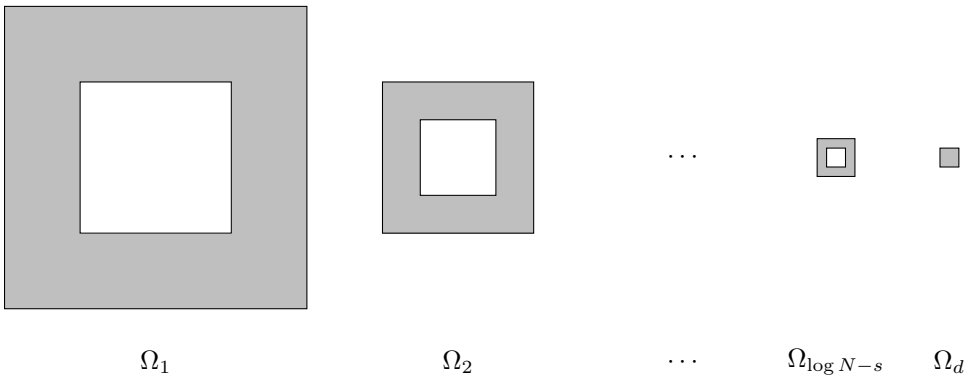

FIG. 1. *This figure shows the frequency domain decomposition of* Ω. *Each subdomain* Ω_j, $j = 1, \ldots, \log N - s$, *is a corona, and* Ω_d *is a small square domain near the origin.*

The kernel function of (9) is smooth in each subdomain Ω_j, and a Cartesian butterfly algorithm is applied to evaluate the contribution from Ω_j. For the center square Ω_d, since it contains only a constant number of points, a direct summation is used. Because of the multiscale nature of the frequency domain decompositions, we refer to this algorithm as the *multiscale butterfly algorithm*. As we shall see, the computational and memory complexities of the multiscale butterfly algorithm are still $\mathcal{O}(N^2 \log N)$ and $\mathcal{O}(N^2)$, respectively. On the other hand, the prefactors are much smaller since the multiscale butterfly is based on the Cartesian grids and requires no polar-Cartesian transformation.

1.4. Organization. The rest of this paper is organized as follows. Section 2 presents the overall structure of a butterfly algorithm. Section 3 proves a low-rank property that is essential to the multiscale butterfly algorithm. Section 4 combines the results of the previous two sections and describes the multiscale butterfly algorithm in detail. In section 5, numerical results of several examples are provided to demonstrate the efficiency of the multiscale butterfly algorithm. Finally, we conclude this paper with some discussion in section 6.

2. The butterfly algorithm. This section provides a brief description of the overall structure of the butterfly algorithm. In this section, X and Ω refer to two general sets of M points in $\mathbb{R}^2$, respectively. With no loss of generality, we assume the points in these two sets are distributed quasi-uniformly but are not necessarily the sets defined in (2) and (4).

Given an input $\{g(\xi), \xi \in \Omega\}$, the goal is to compute the potentials $\{u(x), x \in X\}$ defined by

$$u(x) = \sum_{\xi \in \Omega} K(x,\xi) g(\xi), \quad x \in X,$$

where $K(x,\xi)$ is a kernel function. Let $D_X \supset X$ and $D_\Omega \supset \Omega$ be two square domains containing X and Ω, respectively. The main data structure of the butterfly algorithm is a pair of quadtrees T_X and T_Ω. Having D_X as its root box, the tree T_X is built by recursive dyadic partitioning of D_X until each leaf box contains at most a certain number of points. The tree T_Ω is constructed by recursively partitioning in the same way. With the convention that a root node is at level 0, a leaf node is at level $L = O(\log M)$ under the quasi-uniformity condition about the point distributions, where M is the number of points in X and Ω. Throughout, we shall use A and B to denote the square boxes of T_X and T_Ω with ℓ_A and ℓ_B denoting their levels, respectively.

At the heart of the butterfly algorithm is a special low-rank property. Consider any pair of boxes $A \in T_X$ and $B \in T_\Omega$ obeying the condition $\ell_A + \ell_B = L$. In the butterfly algorithm, we assume that the submatrix $\{K(x,\xi)\}_{x \in A, \xi \in B}$ is approximately of a constant rank. More precisely, for any ϵ, there exist a constant r_ϵ independent of M and two sets of functions $\{\alpha_t^{AB}(x)\}_{1 \le t \le r_\epsilon}$ and $\{\beta_t^{AB}(\xi)\}_{1 \le t \le r_\epsilon}$ such that the following holds:

$$\left| K(x,\xi) - \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x) \beta_t^{AB}(\xi) \right| \le \epsilon \quad \forall x \in A, \forall \xi \in B. \tag{10}$$

The number r_ϵ is called the *ϵ-separation rank*. The exact form of the functions $\{\alpha_t^{AB}(x)\}_{1 \le t \le r_\epsilon}$ and $\{\beta_t^{AB}(\xi)\}_{1 \le t \le r_\epsilon}$ of course depends on the problem to which the butterfly algorithm is applied.

For a given square B in D_Ω, define $u^B(x)$ to be the *restricted potential* over the sources $\xi \in B$,

$$u^B(x) = \sum_{\xi \in B} K(x,\xi) g(\xi).$$

The low-rank property gives a compact expansion for $\{u^B(x)\}_{x \in A}$, as summing (10) over $\xi \in B$ with weights $g(\xi)$ gives

$$\left| u^B(x) - \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x) \left(\sum_{\xi \in B} \beta_t^{AB}(\xi) g(\xi) \right) \right| \le \left(\sum_{\xi \in B} |g(\xi)| \right) \epsilon \quad \forall x \in A.$$

Therefore, if one can find coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ obeying

$$\delta_t^{AB} \approx \sum_{\xi\in B} \beta_t^{AB}(\xi) g(\xi), \quad 1 \le t \le r_\epsilon, \tag{11}$$

then the restricted potential $\{u^B(x)\}_{x\in A}$ admits a compact expansion

$$\left| u^B(x) - \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\delta_t^{AB} \right| \le \left(\sum_{\xi\in B} |g(\xi)| \right) \epsilon \quad \forall x \in A.$$

A key point of the butterfly algorithm is that for each pair (A, B), the number of terms in the expansion is independent of M.

Computing $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ by means of (11) for all pairs A, B is not efficient when B is a large box because for each B there are many paired boxes A. The butterfly algorithm, however, comes with an efficient way of computing $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ recursively. The general structure of the algorithm consists of a top down traversal of T_X and a bottom up traversal of T_Ω, carried out simultaneously.

1. Construct the trees T_X and T_Ω with root nodes D_X and D_Ω.
2. Let A be the root of T_X. For each leaf box B of T_Ω, construct the expansion coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ for the potential $\{u^B(x)\}_{x\in A}$ by simply setting

$$\delta_t^{AB} = \sum_{\xi\in B} \beta_t^{AB}(\xi) g(\xi), \quad 1 \le t \le r_\epsilon. \tag{12}$$

3. For $\ell = 1, 2, \ldots, L$, visit level ℓ in T_X and level $L - \ell$ in T_Ω. For each pair (A, B) with $\ell_A = \ell$ and $\ell_B = L - \ell$, construct the expansion coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ for the potential $\{u^B(x)\}_{x\in A}$ using the low-rank representation constructed at the previous level ($\ell = 0$ is the initialization step). Let P be A's parent and C be a child of B. Throughout, we shall use the notation $C \succ B$ when C is a child of B. At level $\ell - 1$, the expansion coefficients $\{\delta_s^{PC}\}_{1\le s\le r_\epsilon}$ of $\{u^C(x)\}_{x\in P}$ are readily available and we have

$$\left| u^C(x) - \sum_{s=1}^{r_\epsilon} \alpha_s^{PC}(x)\delta_s^{PC} \right| \le \left(\sum_{\xi\in C} |g(\xi)| \right) \epsilon \quad \forall x \in P.$$

Since $u^B(x) = \sum_{C\succ B} u^C(x)$, the previous inequality implies that

$$\left| u^B(x) - \sum_{C\succ B} \sum_{s=1}^{r_\epsilon} \alpha_s^{PC}(x)\delta_s^{PC} \right| \le \left(\sum_{\xi\in B} |g(\xi)| \right) \epsilon \quad \forall x \in P.$$

Since $A \subset P$, the above approximation is of course true for any $x \in A$. However, since $\ell_A + \ell_B = L$, the sequence of restricted potentials $\{u^B(x)\}_{x\in A}$ also has a low-rank approximation of size r_ϵ, namely,

$$\left| u^B(x) - \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\delta_t^{AB} \right| \le \left(\sum_{\xi\in B} |g(\xi)| \right) \epsilon \quad \forall x \in A.$$

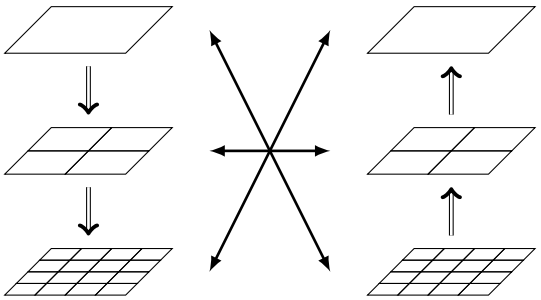

FIG. 2. *Hierarchical domain trees of the two-dimensional butterfly algorithm. Left: T_X for the spatial domain D_X. Right: T_Ω for the frequency domain D_Ω. The interactions between subdomains $A \subset D_X$ and $B \subset D_\Omega$ are represented by left-right arrow lines.*

Combining the last two approximations, we obtain that $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ should obey

$$\sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\delta_t^{AB} \approx \sum_{C\succ B}\sum_{s=1}^{r_\epsilon} \alpha_s^{PC}(x)\delta_s^{PC} \quad \forall x \in A. \tag{13}$$

This represents an overdetermined linear system for $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ in cases when $\{\delta_s^{PC}\}_{1\le s\le r_\epsilon, C\succ B}$ are available. Instead of computing $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ with a least-square method, the butterfly algorithm typically uses an efficient linear transformation approximately mapping $\{\delta_s^{PC}\}_{1\le s\le r_\epsilon, C\succ B}$ into $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$. The actual implementation of this step is very much application-dependent.

4. Finally, let $\ell = L$ and set B to be the root node of T_Ω. For each leaf box $A \in T_X$, use the constructed expansion coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ to evaluate $u(x)$ for each $x \in A$,

$$u(x) = \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\delta_t^{AB}. \tag{14}$$

A schematic illustration of this algorithm is provided in Figure 2. We would like to emphasize that the strict balance between the levels of the target boxes A and source boxes B maintained throughout this procedure is the key to obtaining the accurate low-rank separated approximations.

3. Low-rank approximations. In this section, the sets X and Ω refer to the sets defined in (2) and (4). In order to apply the algorithm in section 5, one would require the existence of the following low-rank separated representation:

$$e^{2\pi\imath\Phi(x,\xi)} \approx \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\beta_t^{AB}(\xi)$$

for any pair of boxes A and B such that $\ell_A + \ell_B = L$. However, this is not true for a general FIO kernel $e^{2\pi\imath\Phi(x,\xi)}$ due to the singularity of $\Phi(x,\xi)$ at the origin $\xi = 0$, i.e., when the square B in Ω is close to the origin of the frequency domain. However, if the frequency domain B is well separated from the origin $\xi = 0$ in a relative sense, one can prove a low-rank separated representation.

In order to make it more precise, for two given squares $A \subset X$ and $B \subset \Omega$, we introduce a new function called the residue phase function,

$$R^{AB}(x,\xi) := \Phi(x,\xi) - \Phi(c_A,\xi) - \Phi(x,c_B) + \Phi(c_A,c_B), \tag{15}$$

where c_A and c_B are the centers of A and B, respectively. Using this new definition, the kernel can be written as

$$e^{2\pi\imath\Phi(x,\xi)} = e^{2\pi\imath\Phi(c_A,\xi)} e^{2\pi\imath\Phi(x,c_B)} e^{-2\pi\imath\Phi(c_A,c_B)} e^{2\pi\imath R^{AB}(x,\xi)}. \tag{16}$$

Below is our main theorem in this paper. It provides theoretical support to the low-rank approximations we used in the multiscale butterfly algorithm. In this theorem, w_A and w_B denote the side lengths of A and B, respectively; $\mathrm{dist}(B,0)$ denotes the distance between the square B and the origin 0 in the frequency domain. The distance is given by $\mathrm{dist}(B,0) = \min_{\xi\in B} \|\xi - 0\|$. Throughout this paper, when we write $\mathcal{O}(\cdot)$, $\lesssim$ and $\gtrsim$, the implicit constant is independent of N and ϵ.

THEOREM 3.1. *Suppose $\Phi(x,\xi)$ is a phase function that is real analytic for x and ξ away from $\xi = 0$. There exist positive constants ϵ_0 and N_0 such that the following is true. Let A and B be two squares in X and Ω, respectively, obeying $w_A w_B \le 1$ and $\mathrm{dist}(B,0) \ge \frac{N}{4}$. For any positive $\epsilon \le \epsilon_0$ and $N \ge N_0$, there exists an approximation*

$$\left| e^{2\pi\imath R^{AB}(x,\xi)} - \sum_{t=1}^{r_\epsilon} \tilde{\alpha}_t^{AB}(x)\tilde{\beta}_t^{AB}(\xi) \right| \le \epsilon$$

for $x \in A$ and $\xi \in B$ with $r_\epsilon \lesssim \log^4(\frac{1}{\epsilon})$. Moreover,

- *when $w_B \le \sqrt{N}$, the functions $\{\tilde{\beta}_t^{AB}(\xi)\}_{1\le t\le r_\epsilon}$ can all be chosen as monomials in $(\xi - c_B)$ with a degree not exceeding a constant times $\log^2(1/\epsilon)$, and*
- *when $w_A \le 1/\sqrt{N}$, the functions $\{\tilde{\alpha}_t^{AB}(x)\}_{1\le t\le r_\epsilon}$ can all be chosen as monomials in $(x - c_A)$ with a degree not exceeding a constant times $\log^2(1/\epsilon)$.*

Proof. Since $w_A w_B \le 1$, we have either $w_A \le 1/\sqrt{N}$ or $w_B \le \sqrt{N}$, or we have both.

Let us first consider the case $w_B \le \sqrt{N}$. Then

$$\begin{aligned} R^{AB}(x,\xi) &= \Phi(x,\xi) - \Phi(c_A,\xi) - \Phi(x,c_B) + \Phi(c_A,c_B) \\ &= [\Phi(x,\xi) - \Phi(c_A,\xi)] - [\Phi(x,c_B) - \Phi(c_A,c_B)] \\ &= H(x,\xi) - H(x,c_B), \end{aligned}$$

where $H(x,\xi) := \Phi(x,\xi) - \Phi(c_A,\xi)$. The function $R^{AB}(x,\xi)$ inherits the smoothness from $\Phi(x,\xi)$. Applying the multivariable Taylor expansion of degree k in ξ centered at c_B gives

$$R^{AB}(x,\xi) = \sum_{1\le|i|<k} \frac{\partial_\xi^i H(x,c_B)}{i!}(\xi - c_B)^i + \sum_{|i|=k} \frac{\partial_\xi^i H(x,\xi^*)}{i!}(\xi - c_B)^i, \tag{17}$$

where ξ^* is a point in the segment between c_B and ξ. Here $i = (i_1, i_2)$ is a multi-index with $i! = i_1! i_2!$, and $|i| = i_1 + i_2$. Let us first choose the degree k so that the second sum in (17) is bounded by $\epsilon/(4\pi)$. For each i with $|i| = k$, the definition of $H(x,\xi)$ gives

$$\partial_\xi^i H(x,\xi^*) = \sum_{|j|=1} \partial_x^j \partial_\xi^i \Phi(x^*,\xi^*)(x - c_A)^j$$

for some point x^* in the segment between c_A and x. Using the fact that $\Phi(x,\xi)$ is real-analytic over $|\xi|=1$ gives that there exists a radius R such that

$$|\partial_x^j \partial_\xi^i \Phi(x,\xi)| \le Ci!j! \frac{1}{R^{|i+j|}} = Ci!j! \frac{1}{R^{k+1}}$$

for ξ with $|\xi|=1$. Here the constant C is independent of k. Since $\Phi(x,\xi)$ is homogeneous of degree 1 in ξ, a scaling argument shows that

$$|\partial_x^j \partial_\xi^i \Phi(x^*,\xi^*)| \le Ci!j! \frac{1}{R^{k+1}|\xi^*|^{k-1}}.$$

Since $\mathrm{dist}(B,0) \ge N/4$ and $w_A w_B \le 1$, we have

$$\left| \frac{\partial_\xi^i H(x,\xi^*)}{i!} (\xi - c_B)^i \right| \le \frac{2Ci!j!}{i!} \frac{1}{R^{k+1}|\xi^*|^{k-1}} w_A w_B^k \le \frac{2C}{R^{k+1}} \left(\frac{4}{\sqrt{N}} \right)^{k-1}.$$

Combining this with (17) gives

$$\begin{aligned} \left| R^{AB}(x,\xi) - \sum_{1\le|i|<k} \frac{\partial_\xi^i H(x,c_B)}{i!} (\xi - c_B)^i \right| &= \left| \sum_{|i|=k} \frac{\partial_\xi^i H(x,\xi^*)}{i!} (\xi - c_B)^i \right| \\ &\le \frac{2C(k+1)}{R^{k+1}} \left(\frac{4}{\sqrt{N}} \right)^{k-1}. \end{aligned}$$

Therefore, for a sufficient large $N_0(R)$, if $N > N_0(R)$, choosing $k = k_\epsilon = O(\log(1/\epsilon))$ ensures that the difference is bounded by $\epsilon/(4\pi)$.

The special case $k=1$ results in the following bound for $R^{AB}(x,\xi)$:

$$|R^{AB}(x,\xi)| \le \frac{4C}{R^2}.$$

To simplify the notation, we define

$$R_\epsilon^{AB}(x,\xi) := \sum_{1\le|i|<k_\epsilon} \frac{\partial_\xi^i H(x,c_B)}{i!} (\xi - c_B)^i,$$

i.e., the first sum on the right-hand side of (17) with $k = k_\epsilon$. The choice of k_ϵ together with (17) implies the bound

$$|R_\epsilon^{AB}(x,\xi)| \le \frac{4C}{R^2} + \epsilon.$$

Since $R_\epsilon^{AB}(x,\xi)$ is bounded, a direct application of Lemma 3.2 of [8] gives

$$\left| e^{2\pi\imath R_\epsilon^{AB}(x,\xi)} - \sum_{p=0}^{d_\epsilon} \frac{(2\pi\imath R_\epsilon^{AB}(x,\xi))^p}{p!} \right| \le \epsilon/2, \tag{18}$$

where $d_\epsilon = \mathcal{O}(\log(1/\epsilon))$. Since $R_\epsilon^{AB}(x,\xi)$ is a polynomial in $(\xi - c_B)$, the sum in (18) is also a polynomial in $(\xi - c_B)$ with degree bounded by $k_\epsilon d_\epsilon = \mathcal{O}(\log^2(1/\epsilon))$. Since our problem is in two dimensions, there are at most $\mathcal{O}(\log^4(1/\epsilon))$ possible monomials

in $(\xi - c_B)$ with degree bounded by $k_\epsilon d_\epsilon$. Grouping the terms with the same multi-index in ξ results in an $\mathcal{O}(\log^4(1/\epsilon))$ term ϵ-accurate separated approximation for $e^{2\pi\imath R_\epsilon^{AB}(x,\xi)}$ with the factors $\{\tilde{\beta}_t^{AB}(\xi)\}_{1\le t\le r_\epsilon}$ being monomials of $(\xi - c_B)$.

Finally, from the inequality $|e^{\imath a} - e^{\imath b}| \le |a - b|$, it is clear that a separated approximation for $e^{2\pi\imath R_\epsilon^{AB}(x,\xi)}$ with accuracy $\epsilon/2$ is also one for $e^{2\pi\imath R^{AB}(x,\xi)}$ with accuracy $\epsilon/2 + \epsilon/2 = \epsilon$. This completes the proof for the case $w_B \le \sqrt{N}$.

The proof for the case $w_A \le 1/\sqrt{N}$ is similar. The only difference is that we now group with

$$R^{AB}(x,\xi) = [\Phi(x,\xi) - \Phi(x,c_B)] - [\Phi(c_A,\xi) - \Phi(c_A,c_B)]$$

and apply the multivariable Taylor expansion in x centered at c_A instead. This results in an $\mathcal{O}(\log^4(1/\epsilon))$ term ϵ-accurate separated approximation for $e^{2\pi\imath R^{AB}(x,\xi)}$ with the factors $\{\tilde{\alpha}_t^{AB}(x)\}_{1\le t\le r_\epsilon}$ being monomials of $(x - c_A)$. ∎

Though the above proof is constructive, it is cumbersome to construct the separated approximation this way. On the other hand, the proof shows that when $w_B \le \sqrt{N}$, the ξ-dependent factors in the low-rank approximation of $e^{2\pi\imath R^{AB}(x,\xi)}$ can be monomials in $(\xi - c_B)$. Similarly, when $w_A \le 1/\sqrt{N}$, the x-dependent factors are monomials in $(x - c_A)$. This suggests using Chebyshev interpolation in x when $w_A \le 1/\sqrt{N}$ and in ξ when $w_B \le \sqrt{N}$. For this purpose, we associate with each box a Chebyshev grid as follows.

For a fixed integer q, the Chebyshev grid of order q on $[-1/2, 1/2]$ is defined by

$$\left\{ z_i = \frac{1}{2}\cos\left(\frac{i\pi}{q-1}\right) \right\}_{0\le i\le q-1}.$$

A tensor-product grid *adapted to a square* with center c and side length w is then defined via shifting and scaling as

$$\{c + w(z_i, z_j)\}_{i,j=0,1,\ldots,q-1}.$$

In what follows, M_t^B is the two-dimensional Lagrange interpolation polynomial on the Chebyshev grid adapted to the square B (i.e., using $c = c_B$ and $w = w_B$).

THEOREM 3.2. *Let A and B be as in Theorem* 3.1. *Then for any $\epsilon \le \epsilon_0$ and $N \ge N_0$ where ϵ_0 and N_0 are the constants in Theorem* 3.1, *there exists $q_\epsilon \lesssim \log^2(1/\epsilon)$ such that the following hold:*

- *when $w_B \le \sqrt{N}$, the Lagrange interpolation of $e^{2\pi\imath R^{AB}(x,\xi)}$ in ξ on a $q_\epsilon \times q_\epsilon$ Chebyshev grid $\{g_t^B\}_{1\le t\le r_\epsilon}$ adapted to B obeys*

$$\left| e^{2\pi\imath R^{AB}(x,\xi)} - \sum_{t=1}^{r_\epsilon} e^{2\pi\imath R^{AB}(x,g_t^B)} M_t^B(\xi) \right| \le \epsilon \quad \forall x \in A, \forall \xi \in B; \tag{19}$$

- *when $w_A \le 1/\sqrt{N}$, the Lagrange interpolation of $e^{2\pi\imath R^{AB}(x,\xi)}$ in x on a $q_\epsilon \times q_\epsilon$ Chebyshev grid $\{g_t^A\}_{1\le t\le r_\epsilon}$ adapted to A obeys*

$$\left| e^{2\pi\imath R^{AB}(x,\xi)} - \sum_{t=1}^{r_\epsilon} M_t^A(x) e^{2\pi\imath R^{AB}(g_t^A,\xi)} \right| \le \epsilon \quad \forall x \in A, \forall \xi \in B. \tag{20}$$

Both (19) *and* (20) *provide a low-rank approximation with $r_\epsilon = q_\epsilon^2 \lesssim \log^4(1/\epsilon)$ terms.*

The proof for this follows exactly that of Theorem 3.3 in [8].

Finally, we are ready to construct the low-rank approximation for the kernel $e^{2\pi\imath\Phi(x,\xi)}$, i.e.,

$$e^{2\pi\imath\Phi(x,\xi)} \approx \sum_{t=1}^{r_\epsilon} \alpha_t^{AB}(x)\beta_t^{AB}(\xi). \tag{21}$$

When $w_B \leq \sqrt{N}$, we multiply (19) with $e^{2\pi\imath\Phi(c_A,\xi)}e^{2\pi\imath\Phi(x,c_B)}e^{-2\pi\imath\Phi(c_A,c_B)}$, which gives that $\forall x \in A, \forall \xi \in B$

$$\left| e^{2\pi\imath\Phi(x,\xi)} - \sum_{t=1}^{r_\epsilon} e^{2\pi\imath\Phi(x,g_t^B)} \left(e^{-2\pi\imath\Phi(c_A,g_t^B)} M_t^B(\xi) e^{2\pi\imath\Phi(c_A,\xi)} \right) \right| \leq \epsilon.$$

In terms of the notation in (21), the expansion functions are given by

$$\alpha_t^{AB}(x) = e^{2\pi\imath\Phi(x,g_t^B)}, \quad \beta_t^{AB}(\xi) = e^{-2\pi\imath\Phi(c_A,g_t^B)} M_t^B(\xi) e^{2\pi\imath\Phi(c_A,\xi)}, \quad 1 \leq t \leq r_\epsilon. \tag{22}$$

This is a special interpolant of the function $e^{2\pi\imath\Phi(x,\xi)}$ in the ξ variable, which prefactors the oscillation, performs the interpolation, and then remodulates the outcome. When $w_A \leq 1/\sqrt{N}$, multiply (20) with $e^{2\pi\imath\Phi(c_A,\xi)}e^{2\pi\imath\Phi(x,c_B)}e^{-2\pi\imath\Phi(c_A,c_B)}$ and obtain that $\forall x \in A, \forall \xi \in B$

$$\left| e^{2\pi\imath\Phi(x,\xi)} - \sum_{t=1}^{r_\epsilon} \left(e^{2\pi\imath\Phi(x,c_B)} M_t^A(x) e^{-2\pi\imath\Phi(g_t^A,c_B)} \right) e^{2\pi\imath\Phi(g_t^A,\xi)} \right| \leq \epsilon.$$

The expansion functions are now

$$\alpha_t^{AB}(x) = e^{2\pi\imath\Phi(x,c_B)} M_t^A(x) e^{-2\pi\imath\Phi(g_t^A,c_B)}, \quad \beta_t^{AB}(\xi) = e^{2\pi\imath\Phi(g_t^A,\xi)}, \quad 1 \leq t \leq r_\epsilon. \tag{23}$$

Due to the presence of the demodulation and remodulation steps in the definitions (22) and (23), we refer to them as *oscillatory Chebyshev interpolations*.

4. Multiscale butterfly algorithm. In this section, we combine the low-rank approximations described in section 3 with the butterfly algorithm in section 2. Due to the restriction on the distance between B and the origin, we decompose (5) into a multiscale summation

$$(Lf)(x) = \sum_{\xi \in \Omega_d} e^{2\pi\imath\Phi(x,\xi)} \widehat{f}(\xi) + \sum_j \sum_{\xi \in \Omega_j} e^{2\pi\imath\Phi(x,\xi)} \widehat{f}(\xi), \tag{24}$$

where

$$\Omega_j = \left\{ (n_1, n_2) : \frac{N}{2^{j+1}} < \max(|n_1|, |n_2|) \leq \frac{N}{2^j} \right\} \cap \Omega$$

for $j = 1, \ldots, \log N - s$, s is a constant, and $\Omega_d = \Omega \setminus \cup_j \Omega_j$.

The term of Ω_d can be evaluated directly since $|\Omega_d| = \mathcal{O}(1)$. Let us now fix an Ω_j. Since any square B in Ω_j always stays away from the origin, the results in section 3 apply to the term for Ω_j in (24). Therefore, the butterfly algorithm as described in section 2 can be adapted to evaluate

$$\sum_{\xi \in \Omega_j} e^{2\pi\imath\Phi(x,\xi)} \widehat{f}(\xi)$$

for the Cartesian domains X and Ω_j. In contrast to the polar butterfly algorithm that works in the polar coordinates for Ω, we refer to this one as *the Cartesian butterfly algorithm*.

4.1. Cartesian butterfly algorithm. To make it more explicit, let us first consider the interaction between (X, Ω_1), with the low-rank approximation implemented using the oscillatory Chebyshev interpolation discussed in section 3.

1. *Preliminaries.* Construct two quadtrees T_X and T_{Ω_1} for X and Ω_1 by uniform hierarchical partitioning. Let b be a constant greater than or equal to 4, and define $N_1 = N$.
2. *Initialization.* For each square $A \in T_X$ of width $1/b$ and each square $B \in T_{\Omega_1}$ of width b, the low-rank approximation functions are
(25)
$$\alpha_t^{AB}(x) = e^{2\pi\imath\Phi(x,g_t^B)}, \ \ \beta_t^{AB}(\xi) = e^{-2\pi\imath\Phi(c_A,g_t^B)} M_t^B(\xi) e^{2\pi\imath\Phi(c_A,\xi)}, \ \ 1 \le t \le r_\epsilon.$$
Hence, we can define the expansion weights $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ with
$$\delta_t^{AB} := \sum_{\xi\in B} \beta_t^{AB}(\xi)\widehat{f}(\xi) = e^{-2\pi\imath\Phi(c_A,g_t^B)} \sum_{\xi\in B} \left(M_t^B(\xi) e^{2\pi\imath\Phi(c_A,\xi)} \widehat{f}(\xi) \right). \tag{26}$$
3. *Recursion.* Go up in tree T_{Ω_1} and down in tree T_X at the same time until we reach the level such that $w_B = \sqrt{N_1}$. At each level, visit all the pairs (A, B). We apply the Chebyshev interpolation in variable ξ and still define the approximation functions given in (25). Let $\{\delta_s^{PC}\}_{1\le s\le r_\epsilon}$ denote the expansion coefficients available in previous steps, where P is A's parent, C is a child of B, and s indicates the Chebyshev grid points in previous domain pairs. We define the new expansion coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ as
$$\delta_t^{AB} := e^{-2\pi\imath\Phi(c_A,g_t^B)} \sum_{C\succ B} \sum_{s=1}^{r_\epsilon} M_t^B(g_s^C) e^{2\pi\imath\Phi(c_A,g_s^C)} \delta_s^{PC}, \tag{27}$$
where we recall that the notation $C \succ B$ means that C is a child of B.
4. *Switch.* For the levels visited, the Chebyshev interpolation is applied in variable ξ, while the interpolation is applied in variable x for levels $l > \log(N_1)/2$. Hence, we are switching the interpolation method at this step. Now we are still working on level $l = \log(N_1)/2$ and the same domain pairs (A, B) in the last step. Let δ_s^{AB} denote the expansion weights obtained by Chebyshev interpolation in variable ξ in the last step. Correspondingly, $\{g_s^B\}_s$ are the grid points in B in the last step. We take advantage of the interpolation in variable x in A and generate grid points $\{g_t^A\}_{1\le t\le r_\epsilon}$ in A. Then we can define new expansion weights
$$\delta_t^{AB} := \sum_{s=1}^{r_\epsilon} e^{2\pi\imath\Phi(g_t^A,g_s^B)} \delta_s^{AB}.$$
5. *Recursion.* Go up in tree T_{Ω_1} and down in tree T_X at the same time until we reach the level such that $w_B = N_1/b$. We construct the approximation functions by Chebyshev interpolation in variable x as follows:
$$\alpha_t^{AB}(x) = e^{2\pi\imath\Phi(x,c_B)} M_t^A(x) e^{-2\pi\imath\Phi(g_t^A,c_B)}, \beta_t^{AB}(\xi) = e^{2\pi\imath\Phi(g_t^A,\xi)}. \tag{28}$$
We define the new expansion coefficients $\{\delta_t^{AB}\}_{1\le t\le r_\epsilon}$ as
$$\delta_t^{AB} := \sum_{C\succ B} e^{2\pi\imath\Phi(g_t^A,c_C)} \sum_{s=1}^{r_\epsilon} \left(M_s^P(g_t^A) e^{-2\pi\imath\Phi(g_s^P,c_C)} \delta_s^{PC} \right), \tag{29}$$
where again P is A's parent and C is a child box of B.

6. *Termination.* Finally, we reach the level that $w_B = N_1/b$. For each B on this level and for each square $A \in T_X$ of width b/N_1, we apply the approximation functions given by (28) and obtain

$$u^B(x) := e^{2\pi\imath\Phi(x,c_B)} \sum_{t=1}^{r_\epsilon} \left(M_t^A(x) e^{-2\pi\imath\Phi(g_t^A,c_B)} \delta_t^{AB} \right) \tag{30}$$

for each $x \in A$. Finally, summing over all B on this level, we have

$$u^{\Omega_1}(x) := \sum_B u^B(x) \tag{31}$$

for each $x \in A$.

We would like to emphasize that the center part of the tree T_{Ω_j} is always empty since Ω_j is a corona. Accordingly, the algorithm skips this empty part.

For a general Ω_j, the interaction between (X, Ω_j) follows a similar algorithm, except that we replace Ω_1 with Ω_j, $u^{\Omega_1}(x)$ with $u^{\Omega_j}(x)$, and N_1 with $N_j = N/2^{j-1}$ and stop at the level that $w_B = N_j/b$.

Finally, (24) is evaluated via

$$(Lf)(x) = u^{\Omega_d}(x) + \sum_j u^{\Omega_j}(x). \tag{32}$$

The multiscale butterfly algorithm can be adapted to nonuniform grid points in X and Ω. The only difference is that the oscillatory Chebyshev interpolation for low-rank approximation uses nonuniform grid points. In this case, we need to generate different interpolation matrices when we visit different leaf domain pairs $A \times B$; i.e., either A is a leaf box of T_X or B is a leaf box of T_Ω. Though this increases the operation and memory requirement, the overall operation and memory complexity remain the same.

4.2. Complexity analysis. The cost of evaluating the term of Ω_d takes at most $\mathcal{O}(N^2)$ steps since $|\Omega_d| = \mathcal{O}(1)$. Let us now consider the cost of the terms associated with $\{\Omega_j\}$.

For the interaction between X and Ω_1, the computation consists of two parts: the recursive evaluation of $\{\delta_t^{AB}\}$ and the final evaluation of $u^{\Omega_1}(x)$. The recursive part takes $\mathcal{O}(q^3 N^2 \log N)$ since there are at most $\mathcal{O}(N^2 \log N)$ pairs of squares (A, B) and the evaluation of $\{\delta_t^{AB}\}$ for each pair takes $\mathcal{O}(q^3)$ steps via dimensionwise Chebyshev interpolation. The final evaluation of $u^{\Omega_1}(x)$ clearly takes $\mathcal{O}(q^2 N^2)$ steps, as we spend $\mathcal{O}(q^2)$ on each point $x \in X$.

For the interaction between X and Ω_j, the analysis is similar. The recursive part now takes $\mathcal{O}(q^3 N_j^2 \log N_j)$ steps (with $N_j = N/2^{j-1}$), as there are at most $\mathcal{O}(N_j^2 \log N_j)$ pairs of squares involved. The final evaluation still takes $\mathcal{O}(q^2 N^2)$ steps.

Summing these contributions together results in the total computational complexity

$$\mathcal{O}(q^3 N^2 \log N) + \mathcal{O}(q^2 N^2 \log N) = \mathcal{O}(q^3 N^2 \log N) = \mathcal{O}(r_\epsilon^{3/2} N^2 \log N).$$

The multiscale butterfly algorithm is also highly efficient in terms of memory, as the Cartesian butterfly algorithm is applied sequentially to evaluate (30) for each Ω_j. Although the overall memory complexity is still $O(N^2)$, the peak memory could be significantly reduced since only $\frac{1}{b^2}$ memory of the original Cartesian butterfly algorithm is used to evaluate the FIO in Ω_j.

5. Numerical results. This section presents several numerical examples to demonstrate the effectiveness of the multiscale butterfly algorithm introduced above. In truth, FIOs usually have nonconstant amplitude functions. Nevertheless, the main computational difficulty is the oscillatory phase term. We refer the reader to [8] for algorithms that deal with nonconstant amplitude functions. The algorithms proposed in this paper are implemented in MATLAB, and the numerical results were obtained on a desktop with a 3.5 GHz CPU and 32 GB of memory. Let $\{u^d(x), x \in X\}$, $\{u^m(x), x \in X\}$, and $\{u^p(x), x \in X\}$ be the results of a discrete FIO computed by a direct matrix-vector multiplication, the multiscale butterfly algorithm, and the polar butterfly algorithm [8], respectively. To report on the accuracy, we randomly select a set S of 256 points from X and evaluate the relative errors of the multiscale butterfly algorithm and the polar butterfly algorithm by

$$\epsilon^m = \sqrt{\frac{\sum_{x\in S}|u^d(x)-u^m(x)|^2}{\sum_{x\in S}|u^d(x)|^2}} \quad \text{and} \quad \epsilon^p = \sqrt{\frac{\sum_{x\in S}|u^d(x)-u^p(x)|^2}{\sum_{x\in S}|u^d(x)|^2}}. \tag{33}$$

According to the description of the multiscale butterfly algorithm in section 4, we recursively divide Ω into $\Omega_j, j = 1, 2, \ldots, \log N - s$, where s is 5 in the following examples. This means that the center square Ω_d is of size $2^5 \times 2^5$ and the interaction from Ω_d is evaluated via a direct matrix-vector multiplication. Suppose q_ϵ is the number of Chebyshev points in each dimension. There is no sense in using butterfly algorithms to construct $\{\delta_t^{AB}\}$ when the number of points in B is fewer than q_ϵ^2. Hence, the recursion step in butterfly algorithms starts from the squares B that are a couple of levels away from the bottom of T_Ω such that each square contains at least q_ϵ^2 points. Similarly, the recursion stops at the squares in T_X that are the same number of levels away from the bottom. In the following examples, we start from level $\log N - 3$ and stop at level 3 (corresponding to $b = 2^3$ defined in section 4), which matches with q_ϵ (4 to 11).

In order to make a fair comparison, we compare the MATLAB versions of the polar butterfly algorithm and the multiscale butterfly algorithm. Hence, the running time of the polar butterfly algorithm here is slower than that in [8], which was implemented in C++.

Example 1. The first example is a generalized Radon transform whose kernel is given by

$$\begin{aligned} \Phi(x,\xi) &= x\cdot\xi + \sqrt{c_1^2(x)\xi_1^2 + c_2^2(x)\xi_2^2}, \\ c_1(x) &= (2+\sin(2\pi x_1)\sin(2\pi x_2))/3, \\ c_2(x) &= (2+\cos(2\pi x_1)\cos(2\pi x_2))/3. \end{aligned} \tag{34}$$

We assume the amplitude of this example is a constant 1. Now the FIO models an integration over ellipses where $c_1(x)$ and $c_2(x)$ are the axis lengths of the ellipse centered at the point $x \in X$. Table 1 summarize the results of this example given by the polar butterfly algorithm and the multiscale butterfly algorithm.

Example 2. Next, we provide an FIO example with a smooth amplitude function,

$$u(x) = \sum_{\xi\in\Omega} a(x,\xi) e^{2\pi\imath\Phi(x,\xi)} \widehat{f}(\xi), \tag{35}$$

TABLE 1

Comparison of the multiscale butterfly algorithm and the polar butterfly algorithm for the phase function in (34). *T_m is the running time of the multiscale butterfly algorithm; T_a is the running time of the polar butterfly algorithm; and T_m/T_p is the speedup factor.*

Multiscale butterfly			Polar butterfly			
N, q_ϵ	ϵ^m	T_m(sec)	N, q_ϵ	ϵ^p	T_p(sec)	T_p/T_m
256,5	7.89e-02	6.96e+01	256,5	4.21e-02	4.84e+02	6.96e+00
512,5	9.01e-02	3.62e+02	512,5	5.54e-02	2.34e+03	6.46e+00
1024,5	9.13e-02	1.81e+03	1024,5	4.26e-02	1.14e+04	6.31e+00
2048,5	9.47e-02	8.79e+03	2048,5	-	-	-
256,7	6.95e-03	8.20e+01	256,7	5.66e-03	5.97e+02	7.28e+00
512,7	8.43e-03	4.16e+02	512,7	5.89e-03	2.82e+03	6.79e+00
1024,7	8.45e-03	2.03e+03	1024,7	4.84e-03	1.35e+04	6.64e+00
2048,7	8.42e-03	1.04e+04	2048,7	-	-	-
256,9	3.90e-04	1.10e+02	256,9	8.25e-04	7.74e+02	7.04e+00
512,9	3.42e-04	5.39e+02	512,9	6.78e-04	3.57e+03	6.61e+00
1024,9	7.61e-04	2.74e+03	1024,9	4.18e-04	1.67e+04	6.09e+00
2048,9	4.82e-04	1.25e+04	2048,9	-	-	-
256,11	2.15e-05	1.84e+02	256,11	3.69e-05	1.15e+03	6.27e+00
512,11	1.89e-05	8.60e+02	512,11	5.53e-05	5.10e+03	5.93e+00
1024,11	1.96e-05	4.27e+03	1024,11	2.042e-05	2.30e+04	5.39e+00
2048,11	1.50e-05	1.82e+04	2048,11	-	-	-

where the amplitude and phase functions are given by

$$\begin{aligned} a(x,\xi) &= (J_0(2\pi\rho(x,\xi)) + \imath Y_0(2\pi\rho(x,\xi)))e^{-\pi\imath\rho(x,\xi)}, \\ \Phi(x,\xi) &= x\cdot\xi + \rho(x,\xi), \\ \rho(x,\xi) &= \sqrt{c_1^2(x)\xi_1^2 + c_2^2(x)\xi_2^2}, \\ c_1(x) &= (2+\sin(2\pi x_1)\sin(2\pi x_2))/3, \\ c_2(x) &= (2+\cos(2\pi x_1)\cos(2\pi x_2))/3. \end{aligned}$$

Here, J_0 and Y_0 are Bessel functions of the first and second kinds. We refer the reader to [7] for more details of the derivation of these formulas. As discussed in [8], we compute the low-rank approximation of the amplitude functions $a(x,\xi)$ first:

$$a(x,\xi) \approx \sum_{t=1}^{s_\epsilon} g_t(x)h_t(\xi).$$

In the second step, we apply the multiscale butterfly algorithm to compute

$$u_t(x) = \sum_{\xi\in\Omega} e^{2\pi\imath\Phi(x,\xi)}\widehat{f}(\xi)h_t(\xi)$$

and sum up all $g_t(x)u_t(x)$ to evaluate

$$u(x) = \sum_t g_t(x)u_t(x).$$

Table 2 summarizes the results of this example given by the direct method and the multiscale butterfly algorithm.

Note that the accuracy of the multiscale butterfly algorithm is well controlled by the number of Chebyshev points q_ϵ. This indicates that our algorithm is numerically stable in practice. Another observation is that the relative error improves on average

TABLE 2

Numerical results given by the multiscale butterfly algorithm for the FIO in (35). *T_d is the running time of the direct evaluation; T_m is the running time of the multiscale butterfly algorithm; and T_d/T_m is the speedup factor.*

N, q_ϵ	ϵ^m	T_d(sec)	T_m(sec)	T_d/T_m
256,7	5.10e-03	3.78e+03	6.07e+02	6.23e+00
512,7	7.29e-03	3.71e+04	3.50e+03	1.06e+01
1024,7	6.16e-03	6.42e+05	1.70e+04	3.77e+01
256,9	4.49e-04	2.34e+03	7.88e+02	2.97e+00
512,9	4.04e-04	3.66e+04	4.64e+03	7.90e+00
1024,9	3.88e-04	6.21e+05	2.17e+04	2.86e+01
256,11	1.86e-05	2.48e+03	1.33e+03	1.86e+00
512,11	1.80e-05	3.60e+04	6.94e+03	5.18e+00
1024,11	2.39e-05	5.96e+05	2.83e+04	2.11e+01

by a factor of 12 every time q_ϵ is incremented by a factor of 2. As we can see in those tables, for a fixed kernel and a fixed q_ϵ, the accuracy is almost independent of N. In fact, as we keep on incrementing q_ϵ, the relative error decreases until it reaches machine precision, independent of N. Hence, in practical applications, one can increase the value of q_ϵ until a desired accuracy is reached. In the comparison in Table 1, the multiscale butterfly algorithm and the polar butterfly algorithm use $q_\epsilon = \{5, 7, 9, 11\}$ and achieve comparable accuracy. Meanwhile, as we observed from Table 1, the decreasing rate of the relative error of the multiscale butterfly algorithm is larger than that of the polar butterfly algorithm. This means that if a high accuracy is desired, the multiscale butterfly algorithm requires a smaller q_ϵ compared to the polar butterfly algorithm.

The second concern about the algorithm is the asymptotic complexity. From the T_m column of Tables 1 and 2, we see that T_m almost quadrupled when the problem size doubled under the same q_ϵ. According to this, we are convinced that the empirical running time of the multiscale butterfly algorithm follows the $\mathcal{O}\left(N^2 \log N\right)$ asymptotic complexity. Note that the speedup factor over the polar butterfly algorithm is about 6 and the multiscale butterfly algorithm obtains better accuracy. This makes the multiscale butterfly algorithm quite attractive to practitioners who are interested in evaluating an FIO with a large N.

Example 3. Extending the multiscale butterfly algorithm to higher dimensions is straightforward. There are two main modifications: higher-dimensional multiscale domain decomposition and Chebyshev interpolation. In three dimensions, the frequency domain is decomposed into cubic shells instead of coronas. The kernel interpolation is applied on a three-dimensional Chebyshev grids. We apply our three-dimensional multiscale butterfly algorithm to a simple example integrating over spheres with different radii. We assume a constant amplitude function, and the kernel function is given by

$$\Phi(x,\xi) = x \cdot \xi + c(x)\sqrt{\xi_1^2 + \xi_2^2}, \quad c(x) = (3 + \sin(2\pi x_1)\sin(2\pi x_2)\sin(2\pi x_3))/4. \tag{36}$$

Table 3 summarizes the results of this example given by the direct method and the multiscale butterfly algorithm.

6. Conclusion. A simple and efficient multiscale butterfly algorithm for evaluating FIOs is introduced in this paper. This method hierarchically decomposes the frequency domain into multiscale coronas in order to avoid possible singularity of the phase function $\Phi(x, \xi)$ at $\xi = 0$. A Cartesian butterfly algorithm is applied to evaluate

TABLE 3
Numerical results given by the multiscale butterfly algorithm for the phase function in (36).

N, q_ϵ	ϵ^m	T_d(sec)	T_m(sec)	T_d/T_m
64,5	9.41e-02	1.82e+04	2.50e+03	7.31e+00
128,5	7.57e-02	6.21e+05	2.42e+04	2.57e+01
256,5	8.23e-02	3.91e+07	2.35e+05	1.66e+02
64,7	1.20e-02	1.83e+04	7.32e+03	2.50e+00
128,7	1.03e-02	6.03e+05	4.48e+04	1.35e+01
256,7	8.13e-03	4.39e+07	3.81e+05	1.15e+02

the FIO over each corona. Many drawbacks of the original butterfly algorithm based on a polar-Cartesian transform in [8] can be avoided. The new multiscale butterfly algorithm has an $\mathcal{O}(N^2 \log N)$ operation complexity with a smaller prefactor, while it keeps the same $\mathcal{O}(N^2)$ memory complexity.

REFERENCES

[1] C. ANDERSON AND M. D. DAHLEH, *Rapid computation of the discrete Fourier transform*, SIAM J. Sci. Comput., 17 (1996), pp. 913–919.

[2] F. ANDERSSON, M. V. DE HOOP, AND H. WENDT, *Multiscale discrete approximation of Fourier integral operators*, Multiscale Model. Simul., 10 (2012), pp. 111–145.

[3] A. AVERBUCH, E. BRAVERMAN, R. COIFMAN, M. ISRAELI, AND A. SIDI, *Efficient computation of oscillatory integrals via adaptive multiscale local Fourier bases*, Appl. Comput. Harmon. Anal., 9 (2000), pp. 19–53.

[4] G. BAO AND W. W. SYMES, *Computation of pseudo-differential operators*, SIAM J. Sci. Comput., 17 (1996), pp. 416–429.

[5] B. BRADIE, R. COIFMAN, AND A. GROSSMANN, *Fast numerical computations of oscillatory integrals related to acoustic scattering,* I, Appl. Comput. Harmon. Anal., 1 (1993), pp. 94–99.

[6] E. CANDÈS, L. DEMANET, D. DONOHO, AND L. YING, *Fast discrete curvelet transforms*, Multiscale Model. Simul., 5 (2006), pp. 861–899.

[7] E. CANDÈS, L. DEMANET, AND L. YING, *Fast computation of Fourier integral operators*, SIAM J. Sci. Comput., 29 (2007), pp. 2464–2493.

[8] E. CANDÈS, L. DEMANET, AND L. YING, *A fast butterfly algorithm for the computation of Fourier integral operators*, Multiscale Model. Simul., 7 (2009), pp. 1727–1750.

[9] E. CANDÈS AND D. L. DONOHO, *New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities*, Comm. Pure Appl. Math., 57 (2004), pp. 219–266.

[10] E. CANDÈS AND D. L. DONOHO, *Continuous curvelet transform.* I. *Resolution of the wavefront set*, Appl. Comput. Harmon. Anal., 19 (2005), pp. 162–197.

[11] E. CANDÈS AND D. L. DONOHO, *Continuous curvelet transform.* II. *Discretization and frames*, Appl. Comput. Harmon. Anal., 19 (2005), pp. 198–222.

[12] E. CORDERO, F. NICOLA, AND L. RODINO, *Sparsity of Gabor representation of Schrödinger propagators*, Appl. Comput. Harmon. Anal., 26 (2009), pp. 357–370.

[13] M. V. DE HOOP, G. UHLMANN, A. VASY, AND H. WENDT, *Multiscale discrete approximations of Fourier integral operators associated with canonical transformations and caustics*, Multiscale Model. Simul., 11 (2013), pp. 566–585.

[14] L. DEMANET AND L. YING, *Wave atoms and sparsity of oscillatory patterns*, Appl. Comput. Harmon. Anal., 23 (2007), pp. 368–387.

[15] L. DEMANET AND L. YING, *Scattering in flatland: Efficient representations via wave atoms*, Found. Comput. Math., 10 (2010), pp. 569–613.

[16] L. DEMANET AND L. YING, *Discrete symbol calculus*, SIAM Rev., 53 (2011), pp. 71–104.

[17] L. DEMANET AND L. YING, *Fast wave computation via Fourier integral operators.*, Math. Comp., 81 (2012), pp. 1455–1486.

[18] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM J. Sci. Comput., 14 (1993), pp. 1368–1393.

[19] L. HÖRMANDER, *Fourier integral operators.* I, Acta Math., 127 (1971), pp. 79–183.

[20] J. HU, S. FOMEL, L. DEMANET, AND L. YING, *A fast butterfly algorithm for generalized Radon transforms*, Geophysics, 78 (2013), pp. U41–U51.

[21] D. Huybrechs and S. Vandewalle, *A two-dimensional wavelet-packet transform for matrix compression of integral equations with highly oscillatory kernel*, J. Comput. Appl. Math., 197 (2006), pp. 218–232.
[22] E. Michielssen and A. Boag, *A multilevel matrix decomposition algorithm for analyzing scattering from large structures*, IEEE Trans. Antennas and Propagation, 44 (1996), pp. 1086–1093.
[23] M. O'Neil, F. Woolfe, and V. Rokhlin, *An algorithm for the rapid evaluation of special function transforms*, Appl. Comput. Harmon. Anal., 28 (2010), pp. 203–226.
[24] M. P. O'Neil, *A New Class of Analysis-based Fast Transforms*, Ph.D. thesis, Yale University, New Haven, CT, 2007.
[25] D. Potts, G. Steidl, and M. Tasche, *Fast Fourier transforms for nonequispaced data: A tutorial*, in Modern Sampling Theory, Springer, New York, 2001, pp. 247–270.
[26] E. M. Stein, *Harmonic Analysis: Real-Variable Methods, Orthogonality, and Oscillatory Integrals*, Princeton Math. Ser. 43, Princeton University Press, Princeton, NJ, 1993.
[27] D. O. Trad, T. J. Ulrych, and M. D. Sacchi, *Accurate interpolation with high-resolution time-variant radon transforms*, Geophysics, 67 (2002), pp. 644–656.
[28] H. Yang and L. Ying, *A fast algorithm for multilinear operators*, Appl. Comput. Harmon. Anal., 33 (2012), pp. 148–158.
[29] B. Yazici, L. Wang, and K. Duman, *Synthetic aperture inversion with sparsity constraints*, in Proceedings of the International Conference on Electromagnetics in Advanced Applications (ICEAA), 2011, pp. 1404–1407.
[30] L. Ying, *Sparse Fourier transform via butterfly algorithm*, SIAM J. Sci. Comput., 31 (2009), pp. 1678–1694.