

A FAST BUTTERFLY ALGORITHM FOR THE COMPUTATION OF FOURIER INTEGRAL OPERATORS*

EMMANUEL CANDÈS[†], LAURENT DEMANET[‡], AND LEXING YING[§]

Abstract. This paper is concerned with the fast computation of Fourier integral operators of the general form $\int_{\mathbb{R}^d} e^{2\pi i\Phi(x,k)} f(k) dk$, where k is a frequency variable, $\Phi(x, k)$ is a phase function obeying a standard homogeneity condition, and f is a given input. This is of interest, for such fundamental computations are connected with the problem of finding numerical solutions to wave equations and also frequently arise in many applications including reflection seismology, curvilinear tomography, and others. In two dimensions, when the input and output are sampled on $N \times N$ Cartesian grids, a direct evaluation requires $O(N^4)$ operations, which is often times prohibitively expensive. This paper introduces a novel algorithm running in $O(N^2 \log N)$ time, i.e., with near-optimal computational complexity, and whose overall structure follows that of the butterfly algorithm. Underlying this algorithm is a mathematical insight concerning the restriction of the kernel $e^{2\pi i\Phi(x,k)}$ to subsets of the time and frequency domains. Whenever these subsets obey a simple geometric condition, the restricted kernel is approximately low-rank; we propose constructing such low-rank approximations using a special interpolation scheme, which prefactors the oscillatory component, interpolates the remaining nonoscillatory part, and finally remodulates the outcome. A byproduct of this scheme is that the whole algorithm is highly efficient in terms of memory requirement. Numerical results demonstrate the performance and illustrate the empirical properties of this algorithm.

Key words. Fourier integral operators, butterfly algorithm, dyadic partitioning, Lagrange interpolation, separated representation, multiscale computations

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

DOI. 10.1137/080734339

1. Introduction. This paper introduces an efficient algorithm for evaluating discrete Fourier integral operators. Let N be a positive integer, which is assumed to be an integer power of 2 with no loss of generality, and define the Cartesian grids $X = \{(i_1/N, i_2/N), 0 \leq i_1, i_2 < N\}$ and $\Omega = \{(k_1, k_2), -N/2 \leq k_1, k_2 < N/2\}$. A discrete Fourier integral operator (FIO) with constant amplitude is defined by

$$(1.1) \quad u(x) = \sum_{k \in \Omega} e^{2\pi i\Phi(x,k)} f(k), \quad x \in X,$$

where $\{f(k), k \in \Omega\}$ is a given input, $\{u(x), x \in X\}$ is the output, and, as usual, $i = \sqrt{-1}$. By an obvious analogy with problems in electrostatics, it will be convenient throughout the paper to refer to $\{f(k), k \in \Omega\}$ as *sources* and $\{u(x), x \in X\}$ as *potentials*. Here, the *phase function* $\Phi(x, k)$ is assumed to be smooth in (x, k) for $k \neq 0$ and obeys an homogeneity condition of degree 1 in k , namely, $\Phi(x, \lambda k) = \lambda\Phi(x, k)$ for each $\lambda > 0$.

*Received by the editors September 3, 2008; accepted for publication March 30, 2009; published electronically June 12, 2009.

<http://www.siam.org/journals/mms/7-4/73433.html>

[†]Applied and Computational Mathematics, Caltech, Pasadena, CA 91125 (ecandes@acm.caltech.edu). This author's research was partially supported by the Waterman Award from the National Science Foundation and by ONR grant N00014-08-1-0749.

[‡]Department of Mathematics, Stanford University, Stanford, CA 94305 (laurent@math.stanford.edu). This author's research was partially supported by National Science Foundation grant DMS-0707921.

[§]Department of Mathematics and ICES, University of Texas, Austin, TX 78712 (lexing@math.utexas.edu). This author's research was partially supported by an Alfred P. Sloan Fellowship and National Science Foundation grant DMS-0708014.

A direct numerical evaluation of (1.1) at all the points in X takes $O(N^4)$ flops, which can be very expensive for large values of N . Surveying the literature, the main obstacle to constructing fast algorithms for (1.1) is the oscillatory behavior of the kernel $e^{2\pi i\Phi(x,k)}$ when N is large, which prevents the use of the standard multiscale techniques developed in [5, 6, 26, 28]. Against this background, the contribution of this paper is to introduce a novel algorithm running in $O(N^2 \log N)$ operations, where the constant is polylogarithmic in the prescribed accuracy ε .

1.1. General strategy. Because the phase function $\Phi(x, k)$ is singular at $k = 0$, the first step consists in representing the frequency variable k in polar coordinates via the transformation

$$(1.2) \quad k = (k_1, k_2) = \frac{\sqrt{2}}{2} N p_1 e^{2\pi i p_2}, \quad e^{2\pi i p_2} = (\cos 2\pi p_2, \sin 2\pi p_2).$$

Here and below, the set of all possible points p generated from Ω is denoted by P ; see Figure 1.1(b). Note that this transformation guarantees that each point $p = (p_1, p_2)$ belongs to the unit square $[0, 1]^2$ since $-N/2 \leq k_1, k_2 < N/2$. Now it is natural to introduce a phase function $\Psi(x, p)$ in the p variable:

$$\Psi(x, p) := \frac{1}{N} \Phi(x, k) = \frac{\sqrt{2}}{2} \Phi(x, e^{2\pi i p_2} p_1),$$

where the last equality comes from the fact that $\Phi(x, k)$ is homogeneous of degree one in k . Since $\Phi(x, k)$ is smooth in (x, k) for $k \neq 0$, $\Psi(x, p)$ is a smooth function of (x, p) with x and p in $[0, 1]^2$.

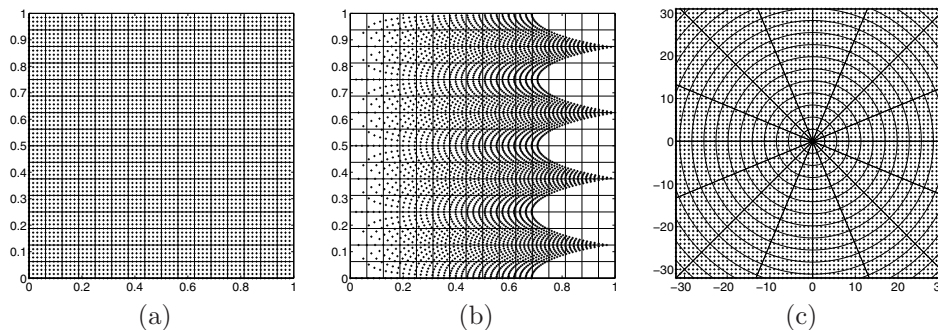


FIG. 1.1. The point distribution and hierarchical partitioning (at a fixed level) for $N = 64$. (a) The set X . (b) The set P in polar coordinates. (c) The frequency partitioning in Cartesian coordinates ($k \in \Omega$).

With these notations, we can reformulate the computational problem (1.1) as

$$u(x) = \sum_{p \in P} e^{2\pi i N \Psi(x, p)} f(p), \quad x \in X,$$

in which the sources $\{f(p)\}$ are now indexed by p instead of k . As we just mentioned, the main issue is that the kernel function $e^{2\pi i N \Psi(x, p)}$ is highly oscillatory. Our approach relies on the observation that this kernel, properly restricted to time and frequency subdomains, admits accurate and low-order separated approximations. To see why this is true, consider two square boxes A and B in $[0, 1]^2$ centered at $x_0(A)$

and $p_0(B)$, and suppose that the sidelengths $w(A)$ and $w(B)$ obey the relationship $w(A)w(B) \leq 1/N$. Introduce the new function

$$(1.3) \quad R^{AB}(x, p) := \Psi(x, p) - \Psi(x_0(A), p) - \Psi(x, p_0(B)) + \Psi(x_0(A), p_0(B)),$$

for each $x \in A$ and $p \in B$, and decompose the kernel $e^{2\pi i N \Psi(x, p)}$ as

$$(1.4) \quad e^{2\pi i N \Psi(x, p)} = e^{2\pi i N \Psi(x_0(A), p)} e^{2\pi i N \Psi(x, p_0(B))} e^{-2\pi i N \Psi(x_0(A), p_0(B))} e^{2\pi i N R^{AB}(x, p)}.$$

In (1.4), we note that each of the first three terms depends on at most one variable (x or p). Recall now the standard multi-index notation; i and j are multi-indices and for $i = (i_1, i_2)$, $i_1, i_2 \geq 0$, $|i| = i_1 + i_2$ and for $x = (x_1, x_2)$, $x^i = x_1^{i_1} x_2^{i_2}$. Applying the mean value theorem to $R^{AB}(x, p)$ successively in p and x gives

$$(1.5) \quad \begin{aligned} R^{AB}(x, p) &\leq \sup_{p^* \in B} \sum_{|j|=1} |\partial_p^j [\Psi(x, p^*) - \Psi(x_0(A), p^*)]| |(p - p_0(B))^j| \\ &\leq \sup_{x^* \in A} \sup_{p^* \in B} \sum_{|i|=1} \sum_{|j|=1} |\partial_x^i \partial_p^j \Psi(x^*, p^*)| |(x - x_0(A))^i| |(p - p_0(B))^j| \\ &= O(1/N). \end{aligned}$$

The last equation follows from the smoothness of the phase $\Psi(x, p)$ and from the assumption $w(A)w(B) \leq 1/N$. To summarize, (1.5) gives $2\pi N R^{AB}(x, p) = O(1)$ and, therefore, the complex exponential $e^{2\pi i N R^{AB}(x, p)}$ is nonoscillatory.

Under some mild smoothness condition, this observation guarantees that for any fixed accuracy ε , there exists a low-rank separated approximation of $e^{2\pi i N R^{AB}(x, p)}$, valid over $A \times B$, effectively decoupling the spatial variable x from the frequency variable p . We propose constructing this low-rank approximation using a tensor-product Chebyshev interpolation of the function $e^{2\pi i N R^{AB}(x, p)}$ in the x variable when $w(A) \leq 1/\sqrt{N}$ and in the p variable when $w(B) \leq 1/\sqrt{N}$. Since the first three terms in (1.4) depend on at most one variable, one also has a separated approximation of $e^{2\pi i N \Psi(x, p)}$ with exactly the same separation rank. Looking at (1.4), the resulting low-rank approximation of the kernel $e^{2\pi i N \Psi(x, p)}$ can be viewed as a special interpolation scheme that prefactors the oscillatory component, interpolates the remaining nonoscillatory part, and finally appends the oscillatory component. As we will see later, the separation rank providing an ε -approximation, for any fixed ε , is bounded from above by a constant independent of N . Further, if we define the partial sum generated by the sources p inside B for any fixed box B as

$$(1.6) \quad u^B(x) := \sum_{p \in B} e^{2\pi i N \Psi(x, p)} f(p),$$

then the existence of such a separated approximation implies the existence of a compact expansion for the restriction of $u^B(x)$ to A , $\{u^B(x), x \in A\}$, of the form

$$(1.7) \quad u^B(x) \approx \sum_{t=1}^r \sum_{p \in B} \alpha_t^{AB}(x) \beta_t^{AB}(p) f(p) = \sum_{t=1}^r \delta_t^{AB} \alpha_t^{AB}(x), \quad \delta_t^{AB} = \sum_{p \in B} \beta_t^{AB}(p) f(p).$$

In (1.7), the number r of expansion coefficients δ_j^{AB} is independent of N for a fixed relative error ε , as we will see later.

The problem is then to compute these compact expansions. This is where the basic structure of the butterfly algorithm [30, 31] is powerful. A brief overview is as follows. We start by building two quadrees T_X and T_P (see Figure 1.1(a) and (b)), respectively, in the spatial and frequency domains with leaf nodes at level $L = \log_2 N$. For each leaf node $B \in T_P$, we first construct the expansion coefficients for the potential $\{u^B(x), x \in A\}$, where A is the root node of T_X . This can be done efficiently because B is a very small box. Next, we go down in T_X and up in T_P simultaneously. For each pair (A, B) with A at the ℓ th level of T_X and B at the $(L - \ell)$ th level of T_P , we construct expansion coefficients for $\{u^B(x), x \in A\}$. As we shall see later, the key point is that this is done by using the expansion coefficients which have already been computed at the previous level. Finally, we arrive at level $\ell = L$, i.e., at the root node of T_P . There $u^B(x) = u(x)$, and since one has available all the compact expansions corresponding to all the leaf nodes A of T_X , one holds an approximation of the potential $u(x)$ for all $x \in X$.

1.2. Applications. The discrete equation (1.1) naturally arises as a numerical approximation of a continuous-time FIO of the general form

$$(1.8) \quad u(x) = \int_{\mathbb{R}^2} a(x, k) e^{2\pi i \Phi(x, k)} f(k) dk.$$

Note that in (1.1), the problem is simplified by setting the amplitude $a(x, k)$ to 1. The reason for making this simpler is that in most applications of interest, $a(x, k)$ is a much simpler object than the term $e^{2\pi i \Phi(x, k)}$. For instance, $a(x, k)$ often has a low-rank separated approximation, which is valid in $\mathbb{R}^2 \times \mathbb{R}^2$ and yields a fast algorithm [3, 11]. Hence, setting $a(x, k) = 1$ retains the essential computational difficulty.

A significant instance of (1.8) is the solution operator to the wave equation

$$u_{tt}(x, t) - c^2 \Delta u(x, t) = 0$$

with constant coefficients and $x \in \mathbb{R}^2$. With initial conditions of the form $u(x, 0) = u_0(x)$ and $u_t(x, 0) = 0$, say, the solution $u(x, t)$ at any time $t > 0$ is given by

$$u(x, t) = \frac{1}{2} \left(\int_{\mathbb{R}^2} e^{2\pi i(x \cdot k + c|k|t)} \hat{u}_0(k) dk + \int_{\mathbb{R}^2} e^{2\pi i(x \cdot k - c|k|t)} \hat{u}_0(k) dk \right),$$

where \hat{u}_0 is the Fourier transform of u_0 . Clearly, this is the sum of two FIOs with phase functions $\Phi_{\pm}(x, k) = x \cdot k \pm c|k|t$ and amplitudes $a_{\pm}(x, k) = 1/2$. Further, FIOs are still solution operators even in the case of inhomogeneous coefficients $c(x)$ as in

$$u_{tt}(x, t) - c^2(x) \Delta u(x, t) = 0.$$

Indeed, under very mild smoothness assumptions, the solution operator remains the sum of two FIOs, at least for sufficiently small times. The only difference is that the phases and amplitudes are a little more complicated. In particular, the phase function is the solution of a Hamilton–Jacobi equation which depends upon $c(x)$.

Another important example of FIO frequently arises in seismics. A fundamental task in reflection seismology consists in producing an image of the sharp features of an underground medium from the seismograms generated by surface explosions. In a nutshell, one builds an imaging operator which maps variations of the pressure field at the surface into variations of the sound speed of the medium (large variations indicate the presence of reflectors). This imaging operator turns out to be an FIO [4, 11].

Because FIOs are hard to compute, several algorithms with various degrees of simplification have been proposed, most notably *Kirchhoff migration*, which approximates the imaging operator as a generalized Radon transform [4, 38]. Computing this transform still has a relatively high complexity, namely, of order N^3 in two dimensions. In contrast, the algorithm proposed in this paper has an optimal $O(N^2 \log N)$ operation count, hence possibly offering a significant speedup.

1.3. Related work. Although FIOs play an important role in the analysis and computation of linear hyperbolic problems, the literature on fast computations of FIOs is surprisingly limited. The only work addressing (1.1) in this general form is the article [11] by the authors of the current paper. The operative feature in [11] is an angular partitioning of the frequency domain into \sqrt{N} wedges, each with an opening angle equal to $2\pi/\sqrt{N}$. When restricting the input to such a wedge, one can then factor the operator into a product of two simpler operators. The first operator is provably approximately low-rank (and lends itself to efficient computations), whereas the second one is a nonuniform Fourier transform which can be computed rapidly using the nonuniform fast Fourier transform (NFFT) [1, 23, 32]. The resulting algorithm has an $O(N^{2.5} \log N)$ complexity.

In a different direction, there has been a great amount of research on other types of oscillatory integral transforms. An important example is the discrete n -body problem where one wants to evaluate sums of the form

$$\sum_{1 \leq j \leq n} q_j K(|x - x_j|), \quad K(r) = e^{i\omega r}/r$$

in the high-frequency regime (ω is large). Such problems appear naturally when solving the Helmholtz equation by means of a boundary integral formulation [16, 17]. A popular approach seeks to compress the oscillatory integral operator by representing it in an appropriate basis such as a local Fourier basis or a basis extracted from the wavelet packet dictionary [2, 7, 22, 29]. This representation sparsifies the operator, thus allowing fast matrix-vector products. In spite of having good theoretical estimates, this approach has thus far been practically limited to one-dimensional boundaries. One particular issue with this approach is that the evaluation of the remaining nonnegligible coefficients sometimes requires assembling the entire matrix, which can be computationally rather expensive.

To the best of our knowledge, the most successful method for the Helmholtz kernel n -body problem in both two and three dimensions is the high-frequency fast multipole method (HF-FMM) proposed by Rokhlin and his collaborators in a series of papers [34, 35, 13]. This approach combines the analytic property of the Helmholtz kernel with an FFT-type fast algorithm to speed up the computation of the interaction between well-separated regions. If N^2 is the number of input and output points as before, the resulting algorithm has an $O(N^2 \log N)$ computational complexity. Other algorithms using similar techniques can be found in [15, 18, 19, 36].

Finally, the idea of butterfly computations has been applied to the n -body problem in several ways. The original paper of Michielssen and Boag [30] used this technique to accelerate the computation of the oscillatory interactions between well-separated regions. More recently, Engquist and Ying [24, 25] proposed a multidirectional solution to this problem, where part of the algorithm can be viewed as a butterfly computation between specially selected spatial subdomains.

1.4. Contents. The rest of this paper is organized as follows. Section 2 describes the overall structure of the butterfly algorithm. In section 3, we prove the low-rank

property of the kernel and introduce an interpolation-based method for constructing low-rank separated approximations. Section 4 develops the algorithm by incorporating our low-rank approximations into the butterfly structure. Numerical results are shown in section 5. Finally, we discuss related problems for future research in section 6.

2. The butterfly algorithm. We begin by offering a general description of the butterfly structure and then provide several concrete examples. This general structure was originally introduced in [30] and later generalized in [31].

In this section, X and P are two arbitrary point sets in \mathbb{R}^d , both of cardinality M . We are given inputs $\{f(p), p \in P\}$ and wish to compute the potentials $\{u(x), x \in X\}$ defined by

$$u(x) = \sum_{p \in P} K(x, p) f(p), \quad x \in X,$$

where $K(x, p)$ is some kernel. Let $D_X \supset X$ and $D_P \supset P$ be two square domains containing X and P , respectively. The main data structure underlying the butterfly algorithm is a pair of dyadic trees T_X and T_P . The tree T_X has D_X as its root box and is built by recursive, dyadic partitioning of D_X until each leaf box contains only a small number of points. The tree T_P recursively partitions D_P in the same way. With the convention that the root nodes are at level 0, one sees that under some uniformity condition about the point distributions, the leaf nodes are at level $L = O(\log M)$. Throughout, A and B denote the square boxes of T_X and T_P , and $\ell(A)$ and $\ell(B)$ denote their level.

The crucial property that makes the butterfly algorithm work is a special low-rank property. Consider any pair of boxes $A \in T_X$ and $B \in T_P$ obeying the condition $\ell(A) + \ell(B) = L$; we want the submatrix $\{K(x, p), x \in A, p \in B\}$ (we will sometimes loosely refer to this as the interaction between A and B) to be approximately of constant rank. More rigorously, for any ε , there must exist a constant r_ε independent of M and two sets of functions $\{\alpha_t^{AB}(x), 1 \leq t \leq r_\varepsilon\}$ and $\{\beta_t^{AB}(p), 1 \leq t \leq r_\varepsilon\}$ such that the following approximation holds:

$$(2.1) \quad \left| K(x, p) - \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \beta_t^{AB}(p) \right| \leq \varepsilon \quad \forall x \in A, \forall p \in B.$$

The number r_ε is called the ε -separation rank. The exact form of the functions $\{\alpha_t^{AB}(x)\}$ and $\{\beta_t^{AB}(p)\}$ of course depends on the problem to which the butterfly algorithm is applied, and we will give two examples at the end of this section.

Recalling the definition $u^B(x) = \sum_{p \in B} K(x, p) f(p)$, the low-rank property gives a compact expansion for $\{u^B(x), x \in A\}$ as follows: summing (2.1) over $p \in B$ with weights $f(p)$ gives

$$\left| u^B(x) - \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \left(\sum_{p \in B} \beta_t^{AB}(p) f(p) \right) \right| \leq \left(\sum_{p \in B} |f(p)| \right) \varepsilon \quad \forall x \in A.$$

Therefore, if we can find coefficients $\{\delta_t^{AB}\}_t$ obeying

$$(2.2) \quad \delta_t^{AB} \approx \sum_{p \in B} \beta_t^{AB}(p) f(p),$$

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

$$\left| u^B(x) - \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \delta_t^{AB} \right| \leq \left(\sum_{p \in B} |f(p)| \right) \varepsilon \quad \forall x \in A.$$

We would like to emphasize that for each pair (A, B) , the number of terms in the expansion is independent of M .

Computing $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ by means of (2.2) 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 for computing $\{\delta_t^{AB}\}$ recursively. The general structure of the algorithm consists of a top down traversal of T_X and a bottom up traversal of T_P , carried out simultaneously. Postponing the issue of computing the separated expansions, i.e., $\{\alpha_t^{AB}(x)\}$ and $\{\beta_t^{AB}(p)\}$, this is how the butterfly algorithm operates.

1. *Preliminaries.* Construct the trees T_X and T_P with root nodes D_X and D_P .
2. *Initialization.* Let A be the root of T_X . For each leaf box B of T_P , construct the expansion coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ for the potential $\{u^B(x), x \in A\}$ by simply setting

$$(2.3) \quad \delta_t^{AB} = \sum_{p \in B} \beta_t^{AB}(p) f(p).$$

3. *Recursion.* For $\ell = 1, 2, \dots, L$, visit level ℓ in T_X and level $L - \ell$ in T_P . For each pair (A, B) with $\ell(A) = \ell$ and $\ell(B) = L - \ell$, construct the expansion coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ for the potential $\{u^B(x), x \in A\}$. This is done by using the low-rank representation constructed at the previous level ($\ell = 0$ is the initialization step). Let A_p be A 's parent and $\{B_c\}$ be B 's children. At level $\ell - 1$, the expansion coefficients $\{\delta_{t'}^{A_p B_c}\}_{t'}$ of $\{u^{B_c}(x), x \in A_p\}$ are readily available and we have

$$\left| u^{B_c}(x) - \sum_{t'=1}^{r_\varepsilon} \alpha_{t'}^{A_p B_c}(x) \delta_{t'}^{A_p B_c} \right| \leq \left(\sum_{p \in B_c} |f(p)| \right) \varepsilon \quad \forall x \in A_p.$$

Since $u^B(x) = \sum_c u^{B_c}(x)$, the previous inequality implies that

$$\left| u^B(x) - \sum_c \sum_{t'=1}^{r_\varepsilon} \alpha_{t'}^{A_p B_c}(x) \delta_{t'}^{A_p B_c} \right| \leq \left(\sum_{p \in B} |f(p)| \right) \varepsilon \quad \forall x \in A_p.$$

Since $A \subset A_p$, 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,

$$\left| u^B(x) - \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \delta_t^{AB} \right| \leq \left(\sum_{p \in B} |f(p)| \right) \varepsilon \quad \forall x \in A.$$

Combining these last two approximations, we obtain that $\{\delta_t^{AB}\}_t$ should obey

$$(2.4) \quad \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \delta_t^{AB} \approx \sum_c \sum_{t'=1}^{r_\varepsilon} \alpha_{t'}^{A_p B_c}(x) \delta_{t'}^{A_p B_c} \quad \forall x \in A.$$

Since this is an overdetermined linear system for $\{\delta_t^{AB}\}_t$ when $\{\delta_{t'}^{A_p B_c}\}_{t',c}$ are available, one possible approach to compute $\{\delta_t^{AB}\}_t$ is to solve a least-squares problem, but this can be very costly when $|A|$ is large. Instead, the butterfly algorithm uses an approximate linear transformation mapping $\{\delta_{t'}^{A_p B_c}\}_{t',c}$ into $\{\delta_t^{AB}\}_t$, which can be computed efficiently. We will discuss how this is done in several examples at the end of this section.

4. *Termination.* Now let $\ell = L$ and set B to be the root node of T_P . For each leaf box $A \in T_X$, use the constructed expansion coefficients $\{\delta_t^{AB}\}_t$ to evaluate $u(x)$ for each $x \in A$,

$$(2.5) \quad u(x) = \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \delta_t^{AB}.$$

A schematic illustration of the algorithm is provided in Figure 2.1. We would like to emphasize that the strict balance between the levels of the target boxes A and source boxes B maintained throughout the procedure is the key to obtaining accurate low-rank separated approximations.




FIG. 2.1. Schematic illustration of the butterfly algorithm in two dimensions with four levels ($L = 3$). The tree T_X is on the left and T_P is on the right. The levels are paired as indicated so that the product of the sidelengths remains constant. The solid line pairs two square boxes A and B at level 2 (shaded in gray); low-rank approximations of the localized kernel and expansion coefficients for the localized potential are computed for each such pair. The algorithm starts at the root of T_X and at the bottom of T_P . It then traverses T_X top down and T_P bottom up and terminates when the last level (the bottom of T_X) is reached. The figure also represents the four children of any box B .

Leaving aside the computations of the separated expansion and taking for granted that constructing $\{\delta_t^{AB}\}_t$ for each pair (A, B) has, in principle, the complexity of applying a linear transform of size $O(r_\varepsilon \times r_\varepsilon)$, observe that the butterfly algorithm has low computational complexity. To be sure, the construction of T_X and T_P clearly takes at most $O(M \log M)$ operations. The initialization and termination steps take at most $O(r_\varepsilon M)$, as these steps require at most $O(r_\varepsilon)$ operations per point; see (2.3) and (2.5). The main workload is of course in the recursion step. At each fixed level ℓ , the number of pairs (A, B) under consideration is of order $O(M)$. It follows from our assumption that the number of flops required to compute all the coefficients $\{\delta_t^{AB}\}_t$ at each level ℓ is just $O(r_\varepsilon^2 M)$. Since there are only about $\log M$ levels, the number of operations in the recursion is at most of the order of $O(r_\varepsilon^2 M \log M)$. In conclusion, the overall operation count is $O(r_\varepsilon^2 M \log M)$.

The general structure of the butterfly algorithm should be clear by now, but we have left out two critical pieces, which we would need to address to apply it to specific problems.

1. What are the functions $\{\alpha_t^{AB}(x)\}$ and $\{\beta_t^{AB}(p)\}$ in the low-rank approximation (2.1) and how are they computed?
2. How do we solve for $\{\delta_t^{AB}\}_t$ from (2.4)?

The rest of this section discusses answers in two distinct examples.

Example 1. In [31], O’Neil and Rokhlin apply the butterfly algorithm to several special function transforms in one dimension. Suppose that N is a positive integer. In this setup, $D_X = D_P = [0, N]$, X and P are two sets of $M = O(N)$ points distributed uniformly or quasi-uniformly in $[0, N]$, and the kernel $K(x, p)$ parametrizes some special functions. For example, in the case of the Fourier transform, $K(x, p) = e^{2\pi i x p / N}$ so that p parametrizes a set of complex sinusoids. The trees T_X and T_P are recursive dyadic partitions of $[0, N]$ until the leaf nodes are of unit size. In this work, all the kernels under study have low-rank approximations when restricted to any pair $A \in T_X$ and $B \in T_P$ obeying $\ell(A) + \ell(B) = L = \log_2 N$.

The main tool for constructing the low-rank approximation is the interpolative decomposition proposed in [27, 14]. Given an $m \times n$ matrix Z which is approximately of rank r , the interpolative decomposition constructs an approximate factorization $Z \approx Z_C R$, where the matrix Z_C consists of a subset of r columns taken from the original matrix Z and the entries of R have values close to one. Such a decomposition requires $O(mn^2)$ operations, while storing the matrix R requires $O(rn)$ memory space. Applying this strategy to the kernel $K(x, p)$ with $x \in A$ and $p \in B$ implies that the functions $\{\alpha_t^{AB}(x), 1 \leq t \leq r\}$ are of the form $\{K(x, p_t^{AB}), 1 \leq t \leq r\}$ with $\{p_t^{AB}\} \subset B$ and the functions $\{\beta_t^{AB}(p), 1 \leq t \leq r\}$ are given by the corresponding entries in the matrix R . Due to the special form of $\{\alpha_t^{AB}(x)\}$, the coefficients $\{\delta_t^{AB}\}_t$ are often called *equivalent sources*.

Now that we have addressed the computations of $\{\alpha_t^{AB}(x)\}$ and $\{\beta_t^{AB}(p)\}$, it remains to examine how to evaluate the coefficients $\{\delta_t^{AB}\}$. In the butterfly algorithm, these coefficients are computed in the initialization step (2.3) and in the recursion step (2.4). Initially, A is the root box of T_X and B is a leaf box of T_P . To compute $\{\delta_t^{AB}, 1 \leq t \leq r\}$ in the initialization step, construct the interpolative decomposition for $K(x, p)$ with $x \in A$ and $p \in B$ obeying

$$(2.6) \quad \left| K(x, p) - \sum_{t=1}^{r_\varepsilon} K(x, p_t^{AB}) \beta_t^{AB}(p) \right| \leq \varepsilon \quad \forall x \in A, \forall p \in B.$$

Since each leaf box B contains only a constant number of points p , constructing the interpolative decomposition requires $O(N)$ operations and $O(r_\varepsilon)$ memory space for each B . Since there at most $O(N)$ of these boxes, the computational costs scales at most like $O(N^2)$. Then we simply compute $\{\delta_t^{AB}, 1 \leq t \leq r\}$ via (2.3). Once the interpolative decomposition is available, this requires $O(r_\varepsilon N)$ operation for all pairs at the 0th level.

As for (2.4), the special form of the functions $\{\alpha_t^{AB}(x), 1 \leq t \leq r_\varepsilon\}$ allows rewriting the right-hand side as

$$u^B(x) \approx \sum_c \sum_{t'=1}^{r_\varepsilon} K(x, p_{t'}^{A_p B_c}) \delta_{t'}^{A_p B_c}.$$

As a result, we can treat this quantity as the potential generated by the equivalent

sources $\{\delta_{t'}^{A_p B_c}\}_{c,t'}$ located at $\{p_{t'}^{A_p B_c}\}_{c,t'}$. In order to find $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$, construct the interpolative decomposition of $K(x, p)$ with $x \in A$ and $p \in \{p_{t'}^{A_p B_c}\}_{c,t'}$, namely,

$$(2.7) \quad \left| K(x, p) - \sum_{t=1}^{r_\varepsilon} K(x, p_t^{AB}) \beta_t^{AB}(p) \right| \leq \varepsilon \quad \forall x \in A, \forall p \in \{p_{t'}^{A_p B_c}\}_{c,t'}$$

Since the number of points in $\{p_{t'}^{A_p B_c}\}_{c,t'}$ is proportional to r_ε , this construction requires $O(r_\varepsilon^2 |A|)$ and requires $O(r_\varepsilon^2)$ memory space per pair (A, B) . Summing (2.7) over $p \in \{p_{t'}^{A_p B_c}\}_{c,t'}$ with weights $\{\delta_{t'}^{A_p B_c}\}_{c,t'}$ gives a way to compute $\{\delta_t^{AB}\}_t$. Indeed, one can set

$$\delta_t^{AB} = \sum_c \sum_{t'} \beta_t^{AB}(p_{t'}^{A_p B_c}) \delta_{t'}^{A_p B_c}, \quad 1 \leq t \leq r_\varepsilon.$$

From the above discussion, we see that the butterfly algorithm described in [31] requires a precomputation step to generate interpolative decompositions for

- $K(x, p)$ for $x \in A$, where A is the root node and $p \in B$ for each leaf node (2.6),
- and $K(x, p)$ for $x \in A$ and $p \in \{p_{t'}^{A_p B_c}\}_{c,t'}$ for each pair (A, B) with $\ell(A) = 1, 2, \dots, L$ and $\ell(A) + \ell(B) = L$ (2.7).

A simple analysis shows that these “precomputations” take $O(r_\varepsilon^2 N^2)$ operations and require $O(r_\varepsilon^2 N \log N)$ memory space. The quadratic time is very costly for problems with large N . This might be acceptable if the same FIO were applied a large number of times. However, in the situation where the operator is applied only a few times, the quadratic precomputation step becomes a huge overhead, and the computational time may even exceed that of the direct evaluation method. Moreover, the storage requirement quickly becomes a bottleneck even for problems of moderate sizes, as in practice the constant r_ε^2 is often nonnegligible.

Example 2. In [39], the butterfly algorithm is used to develop a fast algorithm for sparse Fourier transforms with both spatial and Fourier data supported on curves. Suppose N is a positive integer. In this setting, $D_X = D_P = [0, N]^2$, and X and P are two sets of $M = O(N)$ points supported on smooth curves in $[0, N]^2$. The kernel is given by $K(x, p) = e^{2\pi i x \cdot p/N}$. The quadtrees T_X and T_P are generated adaptively in order to prune branches which do not intersect with the support curves. The leaf boxes are of unit size and $L = \log_2 N$. For any pair of boxes $A \in T_X$ and $B \in T_P$ with $\ell(A) + \ell(B) = L$, it is shown that the restricted kernel $K(x, p)$ is approximately low-rank. Here, the functions $\{\alpha_t^{AB}(x), 1 \leq t \leq r\}$ take the form $\{K(x, p_t^B), 1 \leq t \leq r\}$, where $\{p_t^B\}$ is a tensor-product Chebyshev grid located inside the box B . The coefficients $\{\delta_t^{AB}\}$ —also called equivalent sources—are constructed by collocating (2.4) on a tensor-product Chebyshev grid inside the box A .

Due to the tensor-product structure of the grid $\{p_t^B, 1 \leq t \leq r\}$ and the special form of the kernel $K(x, p) = e^{2\pi i x \cdot p/N}$, one can compute $\{\delta_t^{AB}\}$ via a linear transformation which is essentially independent of the boxes A and B . As a result, one does not need the quadratic-time precomputation step and there is no need to store explicitly these linear transformations. We refer the reader to [39] for more details.

As we shall see, the algorithm introduced in this paper also makes use of tensor-product Chebyshev grids, but the low-rank approximation is constructed through interpolation rather than through collocation. Before discussing other similarities and differences, however, we first need to introduce our algorithm.

3. Low-rank approximations. Recall that our problem is to compute

$$u(x) = \sum_{p \in P} e^{2\pi i N \Psi(x,p)} f(p), \quad \Psi(x,p) = \frac{\sqrt{2}}{2} \Phi(x, e^{2\pi i p^2}) p_1,$$

for all $x \in X$, where X and P are the point sets given in Figure 1.1(a) and (b). As both X and P are contained in $[0, 1]^2$, we set $D_X = [0, 1]^2$ and likewise for D_P . Then the two quadtrees T_X and T_P recursively partition the domains D_X and D_P uniformly until the finest boxes are of sidelength $1/N$.

3.1. The low-rank property. We assume that the function $\Psi(x, p)$ is a real-analytic function in the joint variables x and p . This condition implies the existence of two constants Q and R such that

$$\sup_{x,p \in [0,1]^2} \left| \partial_x^i \partial_p^j \Psi(x,p) \right| \leq Q i! j! R^{-|i|-|j|},$$

where $i = (i_1, i_2)$ and $j = (j_1, j_2)$ are multi-indices, $i! = i_1! i_2!$, and $|i| = i_1 + i_2$. For instance, the constant R can be set as any number smaller than the uniform convergent radius of the power series of Ψ . Following [11], we term these functions (Q, R) -analytic.

The theorem below states that for each pair of boxes $(A, B) \in T_X \times T_P$ obeying $w(A)w(B) = 1/N$, the submatrix $\{e^{2\pi i N R^{AB}(x,p)}, x \in A, p \in B\}$ is approximately low-rank. Throughout, the notation $f \lesssim g$ means $f \leq Cg$ for some numerical constant C independent of N and ε .

THEOREM 3.1. *Let A and B be boxes in T_X and T_P obeying $w(A)w(B) = 1/N$. For any $\varepsilon \leq \varepsilon_0$ and $N \geq N_0$, where ε_0 and N_0 are some constants, there exists an approximation obeying*

$$\left| e^{2\pi i N R^{AB}(x,p)} - \sum_{t=1}^{r_\varepsilon} \alpha_t^{AB}(x) \beta_t^{AB}(p) \right| \leq \varepsilon$$

with $r_\varepsilon \lesssim \log^4(1/\varepsilon)$. Moreover,

- when $w(B) \leq 1/\sqrt{N}$, the functions $\{\beta_t^{AB}(p)\}_t$ can all be chosen as monomials in $(p - p_0(B))$ with a degree not exceeding a constant times $\log^2(1/\varepsilon)$,
- and when $w(A) \leq 1/\sqrt{N}$, the functions $\{\alpha_t^{AB}(x)\}_t$ can all be chosen as monomials in $(x - x_0(A))$ with a degree not exceeding a constant times $\log^2(1/\varepsilon)$.

The proof of Theorem 3.1 uses the following elementary lemma (see [11] for a proof).

LEMMA 3.2. *For each $z_0 > 0$ and $\varepsilon > 0$, set $s_\varepsilon = \lceil \max(2ez_0, \log_2(1/\varepsilon)) \rceil$. Then*

$$\left| e^{iz} - \sum_{t=0}^{s_\varepsilon-1} \frac{(iz)^t}{t!} \right| \leq \varepsilon \quad \forall |z| \leq z_0.$$

Proof of Theorem 3.1. Below, we will drop the dependence on A and B in $x_0(A)$ and $p_0(B)$, for A and B are fixed boxes. Since $w(A)w(B) = 1/N$, we either have $w(A) \leq 1/\sqrt{N}$ or $w(B) \leq 1/\sqrt{N}$ or both. Suppose for instance that $w(B) \leq 1/\sqrt{N}$. Then

$$\begin{aligned} R^{AB}(x,p) &= \Psi(x,p) - \Psi(x_0,p) - \Psi(x,p_0) + \Psi(x_0,p_0) \\ &= [\Psi(x,p) - \Psi(x_0,p)] - [\Psi(x,p_0) - \Psi(x_0,p_0)] \\ &= H_x(p) - H_x(p_0), \end{aligned}$$

where $H_x(p) := \Psi(x, p) - \Psi(x_0, p)$; the subscript indicates that we see H as a function of p and think of x as a parameter. The function $R^{AB}(x, p)$ inherits the analyticity from $\Psi(x, p)$, and its truncated Taylor expansion may be written as

$$(3.1) \quad R^{AB}(x, p) = \sum_{1 \leq |i| < K} \frac{\partial_p^i H_x(p_0)}{i!} (p - p_0)^i + \sum_{|i|=K} \frac{\partial_p^i H_x(p^*)}{i!} (p - p_0)^i,$$

where p^* is a point in the segment $[p_0, p]$. For each i with $|i| = K$, we have

$$\partial_p^i H_x(p^*) = \sum_{|j|=1} \partial_x^j \partial_p^i \Psi(x^*, p^*) (x - x_0)^j$$

for some point x^* in $[x_0, x]$ and, therefore, it follows from the (Q, R) -analycity property that

$$\left| \frac{\partial_p^i H_x(p^*)}{i!} (p - p_0)^i \right| \leq 2QR^{-(K+1)} w(A) (w(B))^K \leq 2QR^{-2} \frac{1}{N} \left(\frac{w(B)}{R} \right)^{(K-1)}.$$

Since $w(B) \leq 1/\sqrt{N}$, $1/\sqrt{N} \leq R/2 \Rightarrow w(B)/R \leq 1/2$ and, therefore, for N sufficiently large,

$$\left| \frac{\partial_p^i H_x(p^*)}{i!} (p - p_0)^i \right| \leq 2^{-(K-2)} \frac{QR^{-2}}{N}.$$

Because there are at most $K + 1$ terms with $|i| = K$, it follows that

$$(3.2) \quad 2\pi N \left| R^{AB}(x, p) - \sum_{1 \leq |i| < K} \frac{\partial_p^i H_x(p_0)}{i!} (p - p_0)^i \right| \leq \pi(K + 1) QR^{-2} 2^{-(K-3)}.$$

Set

$$(3.3) \quad K = C_0 \log(1/\varepsilon).$$

Then if C_0 is a sufficiently large numerical constant, the right-hand side of (3.2) is smaller than ε and, therefore,

$$(3.4) \quad 2\pi N |R^{AB}(x, p) - A(x, p)| \leq \varepsilon, \quad A(x, p) := \sum_{1 \leq |i| < K} \frac{\partial_p^i H_x(p_0)}{i!} (p - p_0)^i.$$

Noting that $|e^{ia} - e^{ib}| \leq |a - b|$, we see that in order to obtain an ε -accurate separated approximation for $e^{2\pi i N R^{AB}(x, p)}$, we need only construct one for $e^{2\pi i N A(x, p)}$. Our plan is to invoke Lemma 3.2. To do this, we need an estimate on $A(x, p)$. When $K = 1$, the estimate in (3.2) provides a bound of $R^{AB}(x, p)$

$$2\pi N |R^{AB}(x, p)| \leq 8\pi QR^{-2}.$$

Combining this estimate with (3.4) yields

$$2\pi N |A(x, p)| \leq 2\pi N |R^{AB}(x, p)| + \varepsilon \leq 8\pi QR^{-2} + \varepsilon.$$

By taking ε small enough, we can assume

$$2e(8\pi QR^{-2} + \varepsilon) \leq \log_2(1/\varepsilon),$$

and Lemma 3.2 gives a $\log(1/\varepsilon)$ -term ε -accurate approximation

$$\left| e^{2\pi i NA(x,p)} - \sum_{t=0}^{\log(1/\varepsilon)-1} \frac{(2\pi NA(x,p))^t}{t!} \right| \leq \varepsilon.$$

Expanding $(2\pi i A(x,p))^t$ for each t gives a sum in which each term is a function of x times a monomial $(p - p_0)^k$ of degree $|k| \lesssim \log^2(1/\varepsilon)$. Since there are at most $O(\log^4(1/\varepsilon))$ different choices for the multi-index k in the expanded formula, combining the terms with the same multi-index k yields an $O(\log^4(1/\varepsilon))$ -term 2ε -accurate separated approximation for $e^{2\pi i NR^{AB}(x,p)}$ with factors $\{\beta_t^{AB}(p)\}$ of the form $(p - p_0)^k$ as claimed.

We studied the case $w(B) \leq 1/\sqrt{N}$, but the method is identical when $w(A) \leq 1/\sqrt{N}$. Write

$$R^{AB}(x,p) = [\Psi(x,p) - \Psi(x,p_0)] - [\Psi(x_0,p) - \Psi(x_0,p_0)]$$

and follow the same procedure. The resulting approximation has $O(\log^4(1/\varepsilon))$ terms, but the factors $\{\alpha_t^{AB}(x)\}$ are now of the form $(x - x_0)^k$ with $|k| \lesssim \log^2(1/\varepsilon)$. \square

Theorem 3.1 shows that the ε -rank of $e^{2\pi i NR^{AB}(x,p)}$ is bounded by a constant multiple of $\log^4(1/\varepsilon)$ for a prescribed accuracy ε . Since

$$\Psi(x,p) = \Psi(x,p_0) + \Psi(x_0,p) - \Psi(x_0,p_0) + R^{AB}(x,p),$$

a direct consequence is that $\{e^{2\pi i N\Psi(x,p)}, x \in A, p \in B\}$ has a separated approximation of the same rank. A possible approach to compute these approximations would be to use the interpolative decomposition described in Example 1 of section 2. However, this method suffers from two main drawbacks discussed in that section limiting its applicability to relatively small problems. This is the reason why we propose below a different and faster low-rank approximation method.

3.2. Interpolation gives good low-rank approximations. The proof of Theorem 3.1 shows that when $w(B) \leq 1/\sqrt{N}$, the p -dependent factors in the low-rank approximation of $e^{2\pi i NR^{AB}(x,p)}$ are all monomials in p . Similarly, when $w(A) \leq 1/\sqrt{N}$, the x -dependent factors are monomials in x . This suggests that an alternative to obtain a low-rank separated approximation is to use polynomial interpolation in x when $w(A) \leq 1/\sqrt{N}$ and in p when $w(B) \leq 1/\sqrt{N}$.

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

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

We use this to define tensor-product grids adapted to an arbitrary squared box with center c and sidelength w as

$$\{c + w(z_{i_1}, z_{i_2}), i_1, i_2 = 0, 1, \dots, q-1\}.$$

Given a set of grid points $\{z_i \in \mathbb{R}, 0 \leq i \leq q - 1\}$, we will also consider the family of Lagrange interpolation polynomials L_i taking value 1 at z_i and 0 at the other grid points

$$L_i(z; \{z_i\}) = \prod_{0 \leq j \leq q-1, j \neq i} \frac{z - z_j}{z_i - z_j}.$$

For tensor-product grids $\{z_{1,i_1}\} \times \{z_{2,i_2}\}$, we define the two-dimensional interpolation polynomials as

$$L_i(z, \{z_i\}) = L_{i_1}(z_1, \{z_{1,i_1}\}) L_{i_2}(z_2, \{z_{2,i_2}\}), \quad i = (i_1, i_2).$$

The theorem below shows that Lagrange interpolation provides efficient low-rank approximations. In what follows, L_t^B is the two-dimensional Lagrange interpolation polynomial on the Chebyshev grid adapted to the box B .

THEOREM 3.3. *Let A and B be as in Theorem 3.1. Then for any $\varepsilon \leq \varepsilon_0$ and $N \geq N_0$, where ε_0 and N_0 are the constants in Theorem 3.1, there exists $q_\varepsilon \lesssim \log^2(1/\varepsilon)$ such that*

- when $w(B) \leq 1/\sqrt{N}$, the Lagrange interpolation of $e^{2\pi i NR^{AB}(x,p)}$ in p on a $q_\varepsilon \times q_\varepsilon$ Chebyshev grid $\{p_t^B\}$ adapted to B obeys

$$(3.5) \quad \left| e^{2\pi i NR^{AB}(x,p)} - \sum_t e^{2\pi i NR^{AB}(x,p_t^B)} L_t^B(p) \right| \leq \varepsilon \quad \forall x \in A, \forall p \in B,$$

- and when $w(A) \leq 1/\sqrt{N}$, the Lagrange interpolation of $e^{2\pi i NR^{AB}(x,p)}$ in x on a $q_\varepsilon \times q_\varepsilon$ Chebyshev grid $\{x_t^A\}$ adapted to A obeys

$$(3.6) \quad \left| e^{2\pi i NR^{AB}(x,p)} - \sum_t L_t^A(x) e^{2\pi i NR^{AB}(x_t^A,p)} \right| \leq \varepsilon \quad \forall x \in A, \forall p \in B.$$

Both (3.5) and (3.6) provide a low-rank approximation with $r_\varepsilon = q_\varepsilon^2 \lesssim \log^4(1/\varepsilon)$ terms.

The proof of the theorem depends on the following lemma.

LEMMA 3.4. *Let $f(y_1, y_2) \in C([0, 1]^2)$ and V_q be the space spanned by the monomials $y_1^{\alpha_1} y_2^{\alpha_2}$ with $0 \leq \alpha_1, \alpha_2 < q$. The projection operator Π_q mapping f into its Lagrange interpolant on the $q \times q$ tensor-product Chebyshev grid obeys*

$$\|f - \Pi_q f\| \leq (1 + C \log^2 q) \inf_{g \in V_q} \|f - g\|$$

for some numerical constant C , where $\|f\| = \sup_{y \in [0,1]^2} |f(y)|$.

The proof of this lemma is a straightforward generalization of the one-dimensional case, which can be found in [33].

Proof of Theorem 3.3. Suppose that $w(B) \leq 1/\sqrt{N}$ and pick $q_\varepsilon = K \log(1/\varepsilon)$, where $K = C_0 \log(1/\varepsilon)$ is given by (3.3) in the proof of Theorem 3.1. We fix $x \in A$ and view $e^{2\pi i NR^{AB}(x,p)}$ as a function of $p \in B$. Applying Lemma 3.4 to $e^{2\pi i NR^{AB}(x,\cdot)}$ gives

$$\left\| e^{2\pi i NR^{AB}(x,\cdot)} - \Pi_{q_\varepsilon} e^{2\pi i NR^{AB}(x,\cdot)} \right\| \leq (1 + C \log^2 q_\varepsilon) \inf_{g \in V_{q_\varepsilon}} \left\| e^{2\pi i NR^{AB}(x,\cdot)} - g \right\|.$$

Theorem 3.1 states that the functions $\{\beta_t^{AB}(p)\}_t$ are all monomials of degree less than q_ε . Therefore, for a fixed x , the low-rank approximation in that theorem belongs to V_{q_ε} and approximates $e^{2\pi iNR^{AB}(x,\cdot)}$ within ε . Combining this with the previous estimate gives

$$(3.7) \quad \left\| e^{2\pi iNR^{AB}(x,\cdot)} - \Pi_{q_\varepsilon} e^{2\pi iNR^{AB}(x,\cdot)} \right\| \leq (1 + C \log^2 q_\varepsilon) \varepsilon \leq (C_1 + C_2 \log^2(\log(1/\varepsilon))) \varepsilon,$$

where C_1 and C_2 are two constants independent of N and ε . The same analysis applies to the situation where $w(A) \leq 1/\sqrt{N}$; fix $p \in B$ and view $e^{2\pi iNR^{AB}(\cdot,p)}$ as a function of $x \in A$, repeat the same procedure, and obtain the same error bound.

The estimate (3.7) and its analogue when $w(A) \leq 1/\sqrt{N}$ are the claims (3.5) and (3.6) except for the fact that the right-hand side is of the form $(C_1 + C_2 \log^2(\log(1/\varepsilon))) \varepsilon$ rather than ε . In order to get rid of the $C_1 + C_2 \log^2(\log(1/\varepsilon))$ factor, we can repeat the proof with $\varepsilon^{(1+\delta)}$ with a small $\delta > 0$. As q_ε depends on ε only logarithmically, this increases q_ε by only a small constant factor. \square

Finally, to obtain a low-rank approximation for the real kernel $e^{2\pi iN\Psi(x,p)}$ when $w(B) \leq 1/\sqrt{N}$, multiply (3.5) with $e^{2\pi iN\Psi(x_0,p)} e^{2\pi iN\Psi(x,p_0)} e^{-2\pi iN\Psi(x_0,p_0)}$ (we again use x_0 and p_0 as shorthand for $x_0(A)$ and $p_0(B)$), which gives that for all $x \in A$, for all $p \in B$

$$\left| e^{2\pi iN\Psi(x,p)} - \sum_t e^{2\pi iN\Psi(x,p_t^B)} \left(e^{-2\pi iN\Psi(x_0,p_t^B)} L_t^B(p) e^{2\pi iN\Psi(x_0,p)} \right) \right| \leq \varepsilon.$$

In terms of the notation in (2.1), the expansion functions are given by

$$(3.8) \quad \alpha_t^{AB}(x) = e^{2\pi iN\Psi(x,p_t^B)}, \quad \beta_t^{AB}(p) = e^{-2\pi iN\Psi(x_0,p_t^B)} L_t^B(p) e^{2\pi iN\Psi(x_0,p)}.$$

This is a special interpolant of the function $e^{2\pi iN\Psi(x,p)}$ in the p -variable which (1) prefactors the oscillation, (2) performs the interpolation, and (3) remodulates the outcome. Following (2.2), the expansion coefficients $\{\delta_t^{AB}\}_t$ for the potential $\{u^B(x), x \in A\}$ should then obey the condition

$$(3.9) \quad \delta_t^{AB} \approx \sum_{p \in B} \beta_t^{AB}(p) f(p) = e^{-2\pi iN\Psi(x_0,p_t^B)} \sum_{p \in B} \left(L_t^B(p) e^{2\pi iN\Psi(x_0,p)} f(p) \right).$$

When $w(A) \leq 1/\sqrt{N}$, multiply (3.6) with $e^{2\pi iN\Psi(x_0,p)} e^{2\pi iN\Psi(x,p_0)} e^{-2\pi iN\Psi(x_0,p_0)}$ and obtain that for all $x \in A$, for all $p \in B$

$$\left| e^{2\pi iN\Psi(x,p)} - \sum_t \left(e^{2\pi iN\Psi(x,p_0)} L_t^A(x) e^{-2\pi iN\Psi(x_t^A,p_0)} \right) e^{2\pi iN\Psi(x_t^A,p)} \right| \leq \varepsilon.$$

The expansion functions are now

$$(3.10) \quad \alpha_t^{AB}(x) = e^{2\pi iN\Psi(x,p_0)} L_t^A(x) e^{-2\pi iN\Psi(x_t^A,p_0)}, \quad \beta_t^{AB}(p) = e^{2\pi iN\Psi(x_t^A,p)}.$$

The expansion coefficients $\{\delta_t^{AB}\}$ should obey

$$(3.11) \quad \delta_t^{AB} \approx \sum_{p \in B} \beta_t^{AB}(p) f(p) = \sum_{p \in B} e^{2\pi iN\Psi(x_t^A,p)} f(p) = u^B(x_t^A).$$

4. Algorithm description. This section presents our algorithm, which combines the expansions introduced in section 3 with the butterfly structure from section 2.

1. *Preliminaries.* Construct two quadrees T_X and T_P for X and P as in Figure 1.1. Each leaf node of T_X and T_P is of size $1/N \times 1/N$. Since X is a regular Cartesian grid, T_X is just a uniform hierarchical partition.
2. *Initialization.* Set A to be the root of T_X . For each leaf box $B \in T_P$, construct the expansion coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ from (3.9) by setting

$$(4.1) \quad \delta_t^{AB} = e^{-2\pi i N \Psi(x_0(A), p_t^B)} \sum_{p \in B} \left(L_t^B(p) e^{2\pi i N \Psi(x_0(A), p)} f(p) \right).$$

3. *Recursion.* For each $\ell = 1, 2, \dots, L/2$, construct the coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ for each pair (A, B) with A at level ℓ and B at the complementary level $L - \ell$ as follows: let A_p be A 's parent and $\{B_c, c = 1, 2, 3, 4\}$ be B 's children. For each child, we have available from the previous level an approximation of the form

$$u^{B_c}(x) \approx \sum_{t'} e^{2\pi i N \Psi(x, p_{t'}^{B_c})} \delta_{t'}^{A_p B_c} \quad \forall x \in A_p.$$

Summing over all children gives

$$u^B(x) \approx \sum_c \sum_{t'} e^{2\pi i N \Psi(x, p_{t'}^{B_c})} \delta_{t'}^{A_p B_c} \quad \forall x \in A_p.$$

Since $A \subset A_p$, this is also true for any $x \in A$. This means that we can treat $\{\delta_{t'}^{A_p B_c}\}$ as equivalent sources in B . As explained below, we then set the coefficients $\{\delta_t^{AB}\}_t$ as

$$(4.2) \quad \delta_t^{AB} = e^{-2\pi i N \Psi(x_0(A), p_t^B)} \sum_c \sum_{t'} L_t^B(p_{t'}^{B_c}) e^{2\pi i N \Psi(x_0(A), p_{t'}^{B_c})} \delta_{t'}^{A_p B_c}.$$

4. *Switch.* The interpolant in p may be used as the low-rank approximation as long as $\ell \leq L/2$, whereas the interpolant in x is a valid low-rank approximation as soon as $\ell \geq L/2$. Therefore, at $\ell = L/2$, we need to switch representation. Recall that for $\ell \leq L/2$ the expansion coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ may be regarded as equivalent sources, while for $\ell \geq L/2$ they approximate the values of the potential $u^B(x)$ on the Chebyshev grid $\{x_t^A, 1 \leq t \leq r_\varepsilon\}$. Hence, for any pair (A, B) with A at level $L/2$ (and likewise for B), we have $\delta_t^{AB} \approx u^B(x_t^A)$ from (3.11) so that we may set

$$(4.3) \quad \delta_t^{AB} = \sum_s e^{2\pi i N \Psi(x_t^A, p_s^B)} \delta_s^{AB}$$

(we abuse notation here since $\{\delta_t^{AB}\}$ denotes the new set of coefficients and $\{\delta_s^{AB}\}$ the older set).

5. *Recursion (end).* The rest of the recursion is analogous. For $\ell = L/2 + 1, \dots, L$, construct the coefficients $\{\delta_t^{AB}, 1 \leq t \leq r_\varepsilon\}$ as follows. With $\{\alpha_t^{AB}\}$ and $\{\beta_t^{AB}\}$ given by (3.10), we have

$$u^B(x) = \sum_c u^{B_c}(x) \approx \sum_{t', c} \alpha_{t'}^{A_p B_c}(x) \sum_{p \in B_c} \beta_{t'}^{A_p B_c}(p) f(p) \approx \sum_{t', c} \alpha_{t'}^{A_p B_c}(x) \delta_{t'}^{A_p B_c}.$$

Hence, since δ_t^{AB} should approximate $u^B(x_t^A)$ by (3.11), we simply set

$$\delta_t^{AB} = \sum_{t',c} \alpha_{t'}^{A_p B_c}(x_t^A) \delta_{t'}^{A_p B_c}.$$

Substituting α_t^{AB} with its value gives the update (4.4)

$$\delta_t^{AB} = \sum_c e^{2\pi i N \Psi(x_t^A, p_0(B_c))} \sum_{t'} \left(L_{t'}^{A_p}(x_t^A) e^{-2\pi i N \Psi(x_{t'}^A, p_0(B_c))} \delta_{t'}^{A_p B_c} \right).$$

6. *Termination.* Finally, we reach $\ell = L$ and set B to be the root box of T_P . For each leaf box A of T_X , we have

$$u^B(x) \approx \sum_t \alpha_t^{AB}(x) \delta_t^{AB}, \quad x \in A,$$

where $\{\alpha_t^{AB}\}$ is given by (3.10). Hence, for each $x \in A$, we set

$$(4.5) \quad u(x) = e^{2\pi i N \Psi(x, p_0(B))} \sum_t \left(L_t^A(x) e^{-2\pi i N \Psi(x_t^A, p_0(B))} \delta_t^{AB} \right).$$

In order to justify (4.2), recall that

$$\left| e^{2\pi i N \Psi(x, p)} - \sum_t e^{2\pi i N \Psi(x, p_t^B)} \beta_t^{AB}(p) \right| \leq \varepsilon \quad \forall x \in A, \forall p \in B,$$

where $\beta_t^{AB}(p)$ is given by (3.8). Summing the above inequality over $p \in \{p_{t'}^{B_c}\}_{t',c}$ with weights $\{\delta_{t'}^{A_p B_c}\}$ gives

$$u^B(x) \approx \sum_t e^{2\pi i N \Psi(x, p_t^B)} \sum_{c,t'} \beta_t^{AB}(p_{t'}^{B_c}) \delta_{t'}^{A_p B_c},$$

which means that we can set

$$\delta_t^{AB} = \sum_{c,t'} \beta_t^{AB}(p_{t'}^{B_c}) \delta_{t'}^{A_p B_c}.$$

Substituting β_t^{AB} with its value gives the update (4.2).

The main workload is in (4.2) and (4.4). Because of the tensor-product structures, the computations in (4.2) and (4.4) can be accelerated by performing Chebyshev interpolation one dimension at a time, reducing the number of operations from $O(r_\varepsilon^2) = O(q_\varepsilon^4)$ to $O(q_\varepsilon^3)$. As there are at most $O(N^2 \log N)$ pairs of boxes (A, B) , the recursion steps take at most $O(r_\varepsilon^{3/2} N^2 \log N)$ operations. It is not difficult to see that the remaining steps of the algorithm take at most $O(r_\varepsilon^2 N)$ operations. Hence, with $r_\varepsilon = O(\log^4(1/\varepsilon))$, this gives an overall complexity estimate of $O(\log^6(1/\varepsilon) N^2 \log N + \log^8(1/\varepsilon) N^2)$. Since the prescribed accuracy ε is a constant, our algorithm runs in $O(N^2 \log N)$ time with a constant polylogarithmic in ε . Although the dependence of this constant on $\log(1/\varepsilon)$ is quite strong, we would like to emphasize that this is only a worst case estimate. In practice, and as empirically demonstrated in section 5, this dependence is rather moderate and grows like $\log(1/\varepsilon)$.

We would like to point out that the values of $L_t^B(p_{t'}^{B_c})$ in (4.2) and of $L_{t'}^{A_p}(x_t^A)$ in (4.4) are both translation- and level-independent because of the nested structure

of the quadtree. Therefore, once these values are computed for a single pair (A, B) , they can just be reused for all pairs visited during the execution of the algorithm. In our implementation, the values of $L_t^B(p_t^{B^c})$ in (4.2) and $L_t^{A^p}(x_t^A)$ are stored in a Kronecker product form in order to facilitate the dimensionwise Chebyshev interpolation discussed in the previous paragraph.

This algorithm has two main advantages over the approach based on interpolative decomposition. First, no precomputation is required. Since the low-rank approximation uses Lagrange interpolation on fixed tensor-product Chebyshev grids, the functions $\{\alpha_t^{AB}(x)\}$ and $\{\beta_t^{AB}(p)\}$ are given explicitly by (3.8) and (3.10). In turn, this yields explicit formulas for computing the expansion coefficients $\{\delta_t^{AB}\}_t$; compare (4.1), (4.2), and (4.4). Second, this algorithm is highly efficient in terms of memory requirement. In the approach based on the interpolative decomposition method, one needs to store many linear transformations (one for each pair (A, B)), which yields a storage requirement on the order of $r_\varepsilon^2 N^2 \log N$ as observed earlier. The proposed algorithm, however, need only store the expansion coefficients $\{\delta_t^{AB}\}$. Moreover, at any point in the execution, only the expansion coefficients from two consecutive levels are actually needed. Therefore, the storage requirement is only on the order of $r_\varepsilon N^2$, which allows us to address problems with much larger sizes.

One advantage of the interpolative decomposition approach is that it often has a smaller separation rank. The reason is that the low-rank approximation is optimized for the kernel under study and, therefore, the computed rank is usually very close to the true separation rank r_ε . In contrast, our low-rank approximations are based on tensor-product Chebyshev grids and merely exploit the smoothness of the function $e^{2\pi i N R^{AB}(x,p)}$ either in x or in p . In particular, it ignores the finer structure of the kernel $e^{2\pi i N \Psi(x,p)}$ and, as a result, the computed separation rank is often significantly higher. Fortunately, this growth in the separation rank does not result in a significant increase in the computation time since the tensor-product structure and the Lagrangian interpolants dramatically decrease the computational cost.

The tensor-product Chebyshev grid is also used in the method described in Example 2 of section 2. There, the equivalent sources are supported on a Chebyshev grid in B and are constructed by collocating the potential on another Chebyshev grid in A . Because of (1) the tensor-product nature of the grids and (2) the nature of the Fourier kernel, the matrix representation of this collocation procedure has an almost (A, B) -independent Kronecker product decomposition. This offers a way of speeding up the computations and makes it unnecessary to store the matrix representation. Unfortunately, such an approach would not work for FIOs since the kernel $e^{2\pi i N \Psi(x,p)}$ does not have an (A, B) -independent tensor-product decomposition. This is why a major difference is that we use tensor-product Chebyshev grids only to interpolate the residual kernel $e^{2\pi i N R^{AB}(x,p)}$ in x or p depending on which box is smaller. The important point is that we also keep the main benefits of that approach.

Up to this point, we have been concerned only with the computation of FIOs with constant amplitudes. However, our approach can easily be extended to the general case with variable amplitudes $a(x, k)$ as in

$$(4.6) \quad u(x) = \sum_{k \in \Omega} a(x, k) e^{2\pi i \Phi(x, k)} f(k), \quad x \in X.$$

In most applications of interest, $a(x, k)$ is a simple object, i.e., much simpler than the oscillatory term $e^{2\pi i \Phi(x, k)}$. A possible approach is to follow [3, 21], where the

amplitude is assumed to have a low-rank separated approximation obeying

$$\left| a(x, k) - \sum_{t=1}^{s_\varepsilon} g_t(x)h_t(k) \right| \leq \varepsilon,$$

where the number of terms s_ε is independent of N —the size of the grids X and Ω . Such an approximation can be obtained either analytically or through the randomized procedure described in [11]. An algorithm for computing (4.6) may then operate as follows:

1. Construct the approximation $a(x, k) \approx \sum_{t=1}^{s_\varepsilon} g_t(x)h_t(k)$ with $x \in X$ and $k \in \Omega$.
2. Set $u(x) = 0$ for $x \in X$ and for each $t = 1, \dots, s_\varepsilon$,
 - (a) form the product $f_t(k) = h_t(k)f(k)$ for $k \in \Omega$,
 - (b) compute $\sum_k e^{2\pi i\Phi(x,k)} f_t(k)$ for $x \in X$ by applying the above algorithm,
 - (c) multiply the result with $g_t(x)$ for $x \in X$, and add this product to $u(x)$.

We would like to point out that the above algorithm is presented in a form that is conceptually simple. However, when applying the butterfly algorithm to the functions $\{f_t(k), t = 1, \dots, s_\varepsilon\}$ in the multiple executions of step 2(b), the following kernel evaluations are independent of $\{f_t(k)\}$ and thus performed redundantly:

- $e^{-2\pi iN\Psi(x_0(A), p_t^B)}$ and $e^{2\pi iN\Psi(x_0(A), p_t^{Bc})}$ in (4.2),
- $e^{2\pi iN\Psi(x_t^A, p_s^B)}$ in (4.3),
- $e^{2\pi iN\Psi(x_t^A, p_0(Bc))}$ and $e^{-2\pi iN\Psi(x_t^A, p_0(Bc))}$ in (4.4).

Therefore, in an efficient implementation of the above algorithm, one should “vectorize” the algorithm to operate on $\{f_t(k), t = 1, \dots, s_\varepsilon\}$ simultaneously so that redundant kernel evaluations can be avoided.

5. Numerical results. This section provides some numerical results to illustrate the empirical properties of the algorithm. The implementation is in C++ and all tests are carried out on a desktop computer with a 2.8GHz CPU.

When computing the matrix-vector product $u(x) = \sum_{k \in \Omega} e^{2\pi i\Phi(x,k)} f(k)$, we independently sample the entries of the input vector $\{f(k), k \in \Omega\}$ from the standard normal distribution so that the input vector is just white noise. Let $\{u^a(x), x \in X\}$ be the potentials computed by the algorithm. To report on the accuracy, we select a set S of 256 points from X and estimate the relative error by

$$(5.1) \quad \sqrt{\frac{\sum_{x \in S} |u(x) - u^a(x)|^2}{\sum_{x \in S} |u(x)|^2}}.$$

According to the algorithm description in section 4, the leaves of the quadtree at level $L = \log_2 N$ are of size $1/N \times 1/N$ and each contains a small number of points. However, when the number of points in a box B is much fewer than q_ε^2 , it does not make sense to construct the expansion coefficients $\{\delta_t^{AB}\}$ simply because the sources at these points would themselves provide a more compact representation. Thus in practice, the recursion starts from the boxes in T_P that are a couple of levels away from the bottom so that each box has at least q_ε^2 points in it. Similarly, the recursion stops at the boxes in T_X that are a couple of levels away from the bottom. In general, the starting and ending levels should depend on the value of q_ε . In the following examples, we start from level $\log_2 N - 3$ and stop at level 3 in T_P . This choice matches well with the values of q_ε (5 to 11) that we use here.

In our first example, we consider the computation of (1.1) with the phase function given by

$$(5.2) \quad \Phi(x, k) = x \cdot k + \sqrt{c_1^2(x)k_1^2 + c_2^2(x)k_2^2}, \quad \begin{aligned} c_1(x) &= (2 + \sin(2\pi x_1) \sin(2\pi x_2))/3, \\ c_2(x) &= (2 + \cos(2\pi x_1) \cos(2\pi x_2))/3. \end{aligned}$$

If $g(x) = \sum_{k \in \Omega} f(k)e^{2\pi i x \cdot k/N}$ is the (periodic) inverse Fourier transform of the input, this example models the integration of g over ellipses where $c_1(x)$ and $c_2(x)$ are the axis lengths of the ellipse centered at the point $x \in X$. In truth, the exact formula of this generalized Radon transform contains an amplitude term $a(x, k)$ involving Bessel functions of the first and second kinds. Nonetheless, we wish to focus on the main computational difficulty, the highly oscillatory phase in this example, and simply set the amplitude $a(x, k)$ to one. Table 5.1 summarizes the results of this example for different combinations of the grid size N (the grid is $N \times N$) and of the degree of the polynomial interpolation q .

TABLE 5.1

Computational results with the phase function given by (5.2). $N \times N$ is the size of the domain; q is the size of the Chebyshev interpolation grid in each dimension; T_a is the running time of the algorithm in seconds; T_d is the estimated running time of the direct evaluation method and T_d/T_a is the speedup factor; finally, ε_a is the accuracy estimated with (5.1).

(N, q)	$T_a(\text{sec})$	$T_d(\text{sec})$	T_d/T_a	ε_a
(256,5)	6.11e+1	3.20e+2	5.24e+0	1.26e-2
(512,5)	2.91e+2	5.59e+3	1.92e+1	1.56e-2
(1024,5)	1.48e+3	9.44e+4	6.37e+1	1.26e-2
(2048,5)	6.57e+3	1.53e+6	2.32e+2	1.75e-2
(4096,5)	3.13e+4	2.43e+7	7.74e+2	1.75e-2
(256,7)	1.18e+2	3.25e+2	2.76e+0	7.57e-4
(512,7)	5.54e+2	5.47e+3	9.87e+0	6.68e-4
(1024,7)	2.76e+3	9.48e+4	3.44e+1	6.45e-4
(2048,7)	1.23e+4	1.46e+6	1.19e+2	8.39e-4
(4096,7)	5.80e+4	2.31e+7	3.99e+2	8.18e-4
(256,9)	2.46e+2	3.10e+2	1.26e+0	3.15e-5
(512,9)	1.03e+3	5.19e+3	5.06e+0	3.14e-5
(1024,9)	4.95e+3	9.44e+4	1.91e+1	3.45e-5
(2048,9)	2.21e+4	1.48e+6	6.71e+1	4.01e-5
(4096,9)	1.02e+5	2.23e+7	2.18e+2	4.21e-5
(256,11)	4.66e+2	3.07e+2	6.59e-1	7.34e-7
(512,11)	1.69e+3	4.53e+3	2.68e+0	7.50e-7
(1024,11)	8.33e+3	9.50e+4	1.14e+1	5.23e-7
(2048,11)	3.48e+4	1.49e+6	4.27e+1	5.26e-7

Next, we use the algorithm described at the end of section 4 to study the performance in the more general setup of variable amplitudes (4.6). The second example is the exact formula for integrating over circles with radii $c(x)$ centered at the points $x \in X$

$$u(x) = \sum_{k \in \Omega} a_+(x, k)e^{2\pi i \Phi_+(x, k)} f(k) + \sum_{k \in \Omega} a_-(x, k)e^{2\pi i \Phi_-(x, k)} f(k),$$

where the amplitudes and phases are given by

$$(5.3) \quad \begin{aligned} a_{\pm}(x, k) &= (J_0(2\pi c(x)|k|) \pm iY_0(2\pi c(x)|k|)) e^{\mp 2\pi i c(x)|k|}, \\ \Phi_{\pm}(x, k) &= x \cdot k + c(x)|k|, \\ c(x) &= (3 + \sin(2\pi x_1) \sin(2\pi x_2))/4. \end{aligned}$$

(Above the functions J_0 and Y_0 are special Bessel functions. The appendix in [11] details the derivation of these formulas.) We use the randomized procedure described in [11] to construct the low-rank separated approximation for $a_{\pm}(x, k)$. For an accuracy of $1e-7$, the resulting approximation contains only three terms. Table 5.2 summarizes the results of this example for different combinations of N and q .

TABLE 5.2
Computational results with the amplitudes and phase functions given by (5.3).

(N, q)	$T_a(\text{sec})$	$T_d(\text{sec})$	T_d/T_a	ε_a
(256,5)	1.39e+2	3.20e+3	2.31e+1	1.48e-2
(512,5)	7.25e+2	5.20e+4	7.17e+1	1.62e-2
(1024,5)	3.45e+3	8.34e+5	2.42e+2	1.90e-2
(256,7)	2.69e+2	3.21e+3	1.19e+1	4.71e-4
(512,7)	1.38e+3	5.20e+4	3.78e+1	7.30e-4
(1024,7)	6.43e+3	8.35e+5	1.30e+2	6.35e-4
(256,9)	5.23e+2	3.20e+3	6.12e+0	1.59e-5
(512,9)	2.49e+3	5.17e+4	2.08e+1	2.97e-5
(1024,9)	1.15e+4	8.32e+5	7.25e+1	1.75e-5
(256,11)	1.04e+3	3.18e+3	3.06e+0	8.03e-7
(512,11)	4.10e+3	5.11e+4	1.24e+1	9.38e-7
(1024,11)	1.84e+4	8.38e+5	4.57e+1	8.01e-7

From these tables, the first observation is that the accuracy is well controlled by the size of the Chebyshev grid and that the estimated accuracy ε_a improves on average by a factor of 30 every time q is increased by a factor of 2. In practical applications, one often specifies the accuracy ε_a instead of the grid size q . To adapt to this situation, the quantity (5.1) can be used to estimate the error; whenever the error is too large, one can simply increase the value of q until the desired accuracy is reached. The second observation is that the accuracy decreases only slightly when N increases, indicating that the algorithm is numerically stable. This is due to the fact that the Lebesgue constant of the Chebyshev interpolation is almost optimal; i.e., the Chebyshev interpolation operator has almost the minimum operator norm among all Lagrange interpolants of the same order [33].

These results show that the empirical running time of our algorithm closely follows the $O(N^2 \log N)$ asymptotic complexity. Each time we double N , the size of the grid quadruples. The corresponding running time and speedup factors increase by a factor roughly equal to 4 as well. We note that for large values of N which are of interest to us and to practitioners, the numerical results show a very substantial speedup factor over direct evaluation. For instance, for $4,096 \times 4,096$ grids, we gain three orders of magnitude since one can get nearly two digits of accuracy with a speedup factor exceeding 750.

The article [11] proposed an $O(N^{2.5} \log N)$ approach based on the partitioning of the frequency domain into \sqrt{N} conical region. Though the time complexity of this former algorithm may not be optimal, we showed that it was efficient in parts because its main computational component, the nonuniform fast Fourier transform, is highly optimized. Comparing Tables 4 and 5 in [11]¹ with the numerical results presented here, we observe that both approaches take roughly the same time for $N = 256$ and 512. For $N \leq 256$, the approach based on conical partitioning is faster, as its

¹The results in Tables 4 and 5 of [11] were obtained on a desktop with a 2.6GHz CPU, which is slightly slower yet comparable to the computer used for the tests in this section. The implementation of the nonuniform fast Fourier transform in [11] was written in C++ and compiled as a MEX-function. Finally, the two examples are not exactly similar, but this slight difference is unessential.

complexity has a smaller constant. For $N \geq 512$, however, the current approach based on the butterfly algorithm clearly outperforms our former approach.

It is straightforward to generalize our algorithm to higher dimensions. In three dimensions, for example, the main modification is to use a three-dimensional Chebyshev grid to interpolate $e^{2\pi i R^{AB}(x,k)}$. Consider again a simple three-dimensional example modeling the integration over spheres with varying radii in which the now six-dimensional phase function $\Phi(x, k)$, $x, k \in \mathbb{R}^3$, is given by

$$\Phi(x, k) = x \cdot k + c(x)|k|, \quad c(x) = (3 + \sin(2\pi x_1) \sin(2\pi x_2) \sin(2\pi x_3))/4.$$

Our three-dimensional numerical results are reported in Table 5.3. In this setup, we see that our approach offers a significant speedup even for moderate values of N .

TABLE 5.3
Computational results in three dimensions with the phase function given by (5.4).

(N, q)	$T_a(\text{sec})$	$T_d(\text{sec})$	T_d/T_a	ε_a
(64,7)	1.79e+3	7.33e+3	4.10e+0	3.32e-3
(128,7)	1.58e+4	4.77e+5	3.02e+1	4.06e-3
(256,7)	1.44e+5	2.97e+7	2.06e+2	3.96e-3

6. Conclusions and discussions. This paper introduced a novel and accurate algorithm for evaluating discrete FIOs. Underlying this approach is a key mathematical property, which says that the kernel, restricted to special subdomains in time and frequency, is approximately very low-rank. Our strategy operationalizes this fact by using a multiscale partitioning of the time and frequency domain together with the butterfly structure to achieve an $O(N^2 \log N)$ asymptotic complexity.

A different way to achieve a near-optimal $O(N^2 \log N)$ complexity might be to use the curvelet transform [12, 10] of Candès and Donoho or the wave atoms [20] of Demanet and Ying. In [8, 9], Candès and Demanet proved that the curvelet representation of FIOs is optimally sparse (the wave atom representation also offers the same optimality), a property which relies on the role played by the second dyadic decomposition of Stein and his collaborators [37]. Whether one can operationalize this mathematical insight into an efficient algorithm seems an interesting direction for future research.

The geometric low-rank property together with the butterfly algorithm appears to be a very powerful combination to obtain fast algorithms for computing certain types of highly oscillatory integrals. We already discussed the work of O’Neil and Rokhlin [31], who have used the butterfly algorithm to design fast special transforms, and of Ying, who has extended this approach to develop fast algorithms for Fourier transforms with sparse data [39] and Fourier transforms with summation constraints [40]. Clearly, it would be of interest to identify wide classes of problems for which this general approach may prove powerful.

Acknowledgment. The authors thank the anonymous reviewers for their comments and suggestions.

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] 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.
- [3] G. BAO AND W. W. SYMES, *Computation of pseudo-differential operators*, SIAM J. Sci. Comput., 17 (1996), pp. 416–429.
- [4] G. BEYLKIN, *The inversion problem and applications of the generalized Radon transform*, Comm. Pure Appl. Math., 37 (1984), pp. 579–599.
- [5] G. BEYLKIN, R. COIFMAN, AND V. ROKHLIN, *Fast wavelet transforms and numerical algorithms. I*, Comm. Pure Appl. Math., 44 (1991), pp. 141–183.
- [6] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Hierarchical Matrices*, Technical report 21, Max-Planck-Institut für Mathematik in den Naturwissenschaften, Leipzig, Germany, 2003.
- [7] 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.
- [8] E. CANDÈS AND L. DEMANET, *Curvelets and Fourier integral operators*, C. R. Math. Acad. Sci. Paris, 336 (2003), pp. 395–398.
- [9] E. CANDÈS AND L. DEMANET, *The curvelet representation of wave propagators is optimally sparse*, Comm. Pure Appl. Math., 58 (2005), pp. 1472–1528.
- [10] E. CANDÈS, L. DEMANET, D. DONOHO, AND L. YING, *Fast discrete curvelet transforms*, Multiscale Model. Simul., 5 (2006), pp. 861–899.
- [11] E. CANDÈS, L. DEMANET, AND L. YING, *Fast computation of Fourier integral operators*, SIAM J. Sci. Comput., 29 (2007), pp. 2464–2493.
- [12] E. CANDÈS AND D. 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.
- [13] H. CHENG, W. Y. CRUTCHFIELD, Z. GIMBUTAS, L. F. GREENGARD, J. F. ETHRIDGE, J. HUANG, V. ROKHLIN, N. YARVIN, AND J. ZHAO, *A wideband fast multipole method for the Helmholtz equation in three dimensions*, J. Comput. Phys., 216 (2006), pp. 300–325.
- [14] H. CHENG, Z. GIMBUTAS, P. G. MARTINSSON, AND V. ROKHLIN, *On the compression of low rank matrices*, SIAM J. Sci. Comput., 26 (2005), pp. 1389–1404.
- [15] W. CHEW, E. MICHELSEN, J. M. SONG, AND J. M. JIN, EDs., *Fast and Efficient Algorithms in Computational Electromagnetics*, Artech House, Norwood, MA, 2001.
- [16] D. COLTON AND R. KRESS, *Inverse Acoustic and Electromagnetic Scattering Theory*, 2nd ed., Appl. Math Sci. 93, Springer-Verlag, Berlin, 1998.
- [17] D. L. COLTON AND R. KRESS, *Integral equation methods in scattering theory*, Pure Appl. Math. (N.Y.), John Wiley & Sons, New York, 1983.
- [18] E. DARVE, *The fast multipole method: Numerical implementation*, J. Comput. Phys., 160 (2000), pp. 195–240.
- [19] E. DARVE AND P. HAVÉ, *A fast multipole method for Maxwell equations stable at all frequencies*, Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci., 362 (2004), pp. 603–628.
- [20] L. DEMANET AND L. YING, *Wave atoms and sparsity of oscillatory patterns*, Appl. Comput. Harmon. Anal., 23 (2007), pp. 368–387.
- [21] L. DEMANET AND L. YING, *Discrete symbol calculus*, SIAM Rev., to appear.
- [22] L. DEMANET AND L. YING, *Scattering in flatland: Efficient representations via wave atoms*, Found. Comput. Math., to appear.
- [23] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM J. Sci. Comput., 14 (1993), pp. 1368–1393.
- [24] B. ENGQUIST AND L. YING, *Fast directional multilevel algorithms for oscillatory kernels*, SIAM J. Sci. Comput., 29 (2007), pp. 1710–1737.
- [25] B. ENGQUIST AND L. YING, *A fast directional algorithm for high frequency acoustic scattering in two dimensions*, Commun. Math. Sci., to appear.
- [26] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.
- [27] M. GU AND S. C. EISENSTAT, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869.
- [28] W. HACKBUSCH AND Z. P. NOWAK, *On the fast matrix multiplication in the boundary element method by panel clustering*, Numer. Math., 54 (1989), pp. 463–491.
- [29] 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.
- [30] E. MICHELSEN AND A. BOAG, *A multilevel matrix decomposition algorithm for analyzing scattering from large structures*, IEEE Trans. Antennas and Propagation, 44 (1996), pp. 1086–1093.

- [31] M. O'NEIL AND V. ROKHLIN, *A New Class of Analysis-Based Fast Transforms*, Technical report/YALE/DCS/TR1384, Yale University, New Haven, CT, 2007.
- [32] D. POTTS, G. STEIDL, AND M. TASCHE, *Fast Fourier transforms for nonequispaced data: A tutorial*, in *Modern Sampling Theory*, Appl. Numer. Harmon. Anal., Birkhäuser Boston, Boston, MA, 2001, pp. 247–270.
- [33] T. J. RIVLIN, *The Chebyshev Polynomials*, Wiley-Interscience, New York, 1974.
- [34] V. ROKHLIN, *Rapid solution of integral equations of scattering theory in two dimensions*, J. Comput. Phys., 86 (1990), pp. 414–439.
- [35] V. ROKHLIN, *Diagonal forms of translation operators for the Helmholtz equation in three dimensions*, Appl. Comput. Harmon. Anal., 1 (1993), pp. 82–93.
- [36] J. M. SONG AND W. C. CHEW, *Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering*, Microwave Opt. Tech. Lett., 10 (1995), pp. 15–19.
- [37] E. M. STEIN, *Harmonic Analysis: Real-Variable Methods, Orthogonality, and Oscillatory Integrals*, Princeton Math. Ser. 43, Princeton University Press, Princeton, NJ, 1993.
- [38] W. W. SYMES, *Mathematical Foundations of Reflection Seismology*, Technical report, Rice University, Houston, TX, 1998.
- [39] L. YING, *Sparse Fourier transform via butterfly algorithm*, SIAM J. Sci. Comput., 31 (2009), pp. 1678–1694.
- [40] L. YING AND S. FOMEL, *Fast Computation of Partial Fourier Transforms*, Technical report, University of Texas at Austin, Austin, TX, 2008.