

The Chebyshev fast Gauss and nonuniform fast Fourier transforms and their application to the evaluation of distributed heat potentials [☆]

Shravan K. Veerapaneni ^a, George Biros ^{a,b,*}

^a *Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104, USA*

^b *Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104, USA*

Received 14 March 2007; received in revised form 7 March 2008; accepted 1 May 2008

Available online 16 May 2008

Abstract

We present a method for the fast and accurate computation of distributed heat potentials in two dimensions. The distributed source is assumed to be given in terms of piecewise space–time Chebyshev polynomials. We discretize uniformly in time, whereas in space the polynomials are defined on the leaf nodes of a quadtree. The quadtree can vary at each time step. We combine a product integration rule with fast algorithms (fast heat potentials, nonuniform FFT, fast Gauss transform) to obtain a high-order accurate method with optimal complexity. If N is the number of time steps, M is the maximum number of leaf nodes over all the time steps and the input contains a q th-order polynomial representation of f , then, our method requires $\mathcal{O}(q^3 MN \log M)$ work to evaluate the heat potential at arbitrary MN space–time target locations. The overall convergence rate of the method is of order q . We present numerical experiments for $q = 4, 8, \text{ and } 16$, and we verify the theoretical convergence rate of the method. When the solution is sufficiently smooth, the 16th-order variant results in significant computational savings, even in the case in which we require only a few digits of accuracy.

© 2008 Elsevier Inc. All rights reserved.

Keywords: Heat potentials; Fast Gauss transform; Nonuniform fast Fourier transform; Chebyshev polynomials; Quadtrees; High-order quadratures

1. Introduction

In this article, we present a fast high-order algorithm for the solution of the heat equation with a distributed force $f(x, t)$:

[☆] This work was supported by the US Department of Energy under Grant DE-FG02-04ER25646, and the US National Science Foundation Grants CCF-0427985, CNS-0540372, and DMS-0612578.

* Corresponding author. Address: Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104, USA. Tel.: +1 215 898 3108; fax: +1 215 573 6334.

E-mail addresses: shravan@seas.upenn.edu (S.K. Veerapaneni), biros@seas.upenn.edu (G. Biros).

$$\frac{\partial u}{\partial t} = \Delta u + f(\mathbf{x}, t) \quad \text{in } \Omega, \quad \text{where } \Omega = [0, 1]^2 \text{ and } t \in (0, T], \tag{1}$$

$$u(\mathbf{x}, 0) = 0 \quad \text{and} \quad u(\Gamma, t) = 0, \tag{2}$$

where Γ is the boundary of Ω . Using potential theory, the solution $u(\mathbf{x}, t)$ can be written as

$$u(\mathbf{x}, t) = \int_0^t \int_{\Omega} G(\mathbf{x}, t; \mathbf{y}, \tau) f(\mathbf{y}, \tau) d\mathbf{y} d\tau, \tag{3}$$

where $G(\mathbf{x}, t; \mathbf{y}, \tau)$ is the Green’s function. In our method we assume that the distributed force f is given as a set of N quadrees $\{Q_0, Q_1, \dots, Q_{N-1}\}$. Each quadree Q_k contains a space–time piecewise Chebyshev polynomial representation of f within $\Omega \times [k\Delta t, (k + 1)\Delta t]$, where $\Delta t = \frac{T}{N}$. Our method is adaptive in space and uniform in time. The evaluation points (also referred to as “targets”) are assumed to belong to a domain $\omega \subset \Omega$.

A direct computation of (3) using generic numerical quadratures is expensive and inaccurate. It is inaccurate because the integration kernel is sharply peaked near the evaluation point. It is expensive because its evaluation requires $\mathcal{O}(NM)$ work for every \mathbf{x} , where N and M are the number of quadrature points in time and space, respectively. Therefore, the cost of evaluation at N, M temporal and spatial points would be $\mathcal{O}(N^2M^2)$.

1.1. Time domain decomposition: far and local part

Our scheme is based on an algorithm, first proposed by Greengard and Strain [13], that eliminates the history dependence of the time integral in (3) by combining two (analytically) equivalent series expansions of the heat kernel, a Fourier based and a method-of-images based:

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = 4 \sum_{\mathbf{n} \in \mathbb{Z}_+^2} \prod_{i=1}^2 e^{-|\mathbf{n}|^2 \pi^2 (t-\tau)} \sin(n_i \pi x_i) \sin(n_i \pi y_i), \quad (\mathbf{n} = \{n_1, n_2\}, |\mathbf{n}| = \sqrt{n_1^2 + n_2^2}) \tag{4}$$

$$= \frac{1}{4\pi(t-\tau)} \sum_{\mathbf{n} \in \mathbb{Z}^2} \sum_{\sigma_1 = \pm 1} \sum_{\sigma_2 = \pm 1} \sigma_1 \sigma_2 e^{-\frac{(x_1 - \sigma_1 y_1 - 2n_1)^2 + (x_2 - \sigma_2 y_2 - 2n_2)^2}{4(t-\tau)}}. \tag{5}$$

The convergence rate of the two expansions depends on $|t - \tau|$. At a time τ close to the evaluation time t , the method-of-images expansion (5) converges faster than the Fourier one. At distant times the opposite is true, the Fourier expansion (4) converges faster. This fact motivates a splitting of the volume potential (3) into a far and a local part:

$$u = u_F + u_L = \int_0^{t-\delta} \int_{\Omega} Gf + \int_{t-\delta}^t \int_{\Omega} Gf, \tag{6}$$

where δ is a small parameter. Typically, $\delta = l\Delta t$ for some positive integer l which will be defined in Section 5.1. The Fourier expansion is used to compute $u_F(\mathbf{x}, t)$ and the method of images expansion for $u_L(\mathbf{x}, t)$. For τ in $(0, t - \delta)$, the truncation error for the p -term Fourier series expansion of G is exponentially decaying with increasing p . If we assume that all target points are sufficiently far from the boundary of the domain Ω , then the terms with $|\mathbf{n}| > 0$ in the local expansion of the Green’s function are decaying exponentially within the interval $(t - \delta, t)$.¹

Far part. Upon truncation of the Fourier expansion, the far part can be approximated by

$$u_F(\mathbf{x}, t) \approx 4 \sum_{n_2=1}^p \sum_{n_1=1}^p C_{\mathbf{n}}(t, \delta) \sin(n_1 \pi x_1) \sin(n_2 \pi x_2), \tag{7}$$

$$\text{where } C_{\mathbf{n}}(t, \delta) = \int_0^{t-\delta} e^{-|\mathbf{n}|^2 \pi^2 (t-\tau)} \hat{f}_{\mathbf{n}}(\tau) d\tau, \tag{8}$$

$$\hat{f}_{\mathbf{n}}(\tau) = \int_{\Omega} f(\mathbf{y}, \tau) \sin(n_1 \pi y_1) \sin(n_2 \pi y_2) d\mathbf{y}. \tag{9}$$

¹ See Section 5 for a discussion on how to choose p and δ . The algorithm can be extended to the case in which the evaluation points are near the boundary of Ω , but the details are tedious without offering further insight.

The following recurrence is the key feature of Greengard and Strain's algorithm:

$$C_{\mathbf{n}}(t + \Delta t, \delta) = e^{-|\mathbf{n}|^2 \pi^2 \Delta t} C_{\mathbf{n}}(t, \delta) + U_{\mathbf{n}}[\hat{f}], \quad (10)$$

$$\text{where } U_{\mathbf{n}}[\hat{f}] = \int_{t-\delta}^{t+\Delta t-\delta} e^{-|\mathbf{n}|^2 \pi^2 (t+\Delta t-\tau)} \hat{f}_{\mathbf{n}}(\tau) d\tau, \quad (11)$$

Δt is the size of time step. At each time step, the far part computation involves the evaluation of the discrete sum (7) and the update of the coefficients $C_{\mathbf{n}}$ using (10).

Local part. Upon truncating the method-of-images expansion the local part can be approximated by

$$u_{\mathbf{L}}(\mathbf{x}, t) \approx \int_{t-\delta}^t \int_{\Omega} \frac{e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{4(t-\tau)}}}{4\pi(t-\tau)} f(\mathbf{y}, \tau) d\mathbf{y} d\tau. \quad (12)$$

One can show that the integrand in (12) is a smooth function in τ . Therefore, we can use any q th-order quadrature rule with nodes and weights $\{\tau_k, w_k\}_{k=1}^q$ (we discuss our choice in Section 5):

$$u_{\mathbf{L}}(\mathbf{x}, t) = \sum_{k=1}^q \frac{w_k}{4\pi(t-\tau_k)} G_{4(t-\tau_k)}[f(\cdot, \tau_k)](\mathbf{x}), \quad (13)$$

$$\text{where } G_{4(t-\tau_k)}[f(\cdot, \tau_k)](\mathbf{x}) = \int_{\Omega} e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{4(t-\tau_k)}} f(\mathbf{y}, \tau_k) d\mathbf{y}. \quad (14)$$

Therefore, the evaluation of the local part involves performing the Gauss transform of f at the quadrature nodes τ_k and then computing the discrete sum (13).

Now we can outline the algorithm. Given f , we use a time marching scheme in which the solution is time decomposed into two parts, the far and local. To compute the far part, we need to evaluate equations (9), (11), (10), and (7) – in that order. These computations involve evaluations of the sine transform of f , using the Fourier coefficients of f to update the coefficients of $u_{\mathbf{F}}$, and use of the inverse Fourier transform to get back to space coordinates. To compute the local part, we need a few fast Gauss transforms to evaluate (14) at the τ_k time points needed for the time quadrature (13). This scheme, however, suffers from several technical difficulties: it should support nonuniform discretizations, and it should be fast. Below we describe algorithmic challenges related to complexity and high-accuracy.

Complexity. A direct evaluation of the far and local parts of u becomes non-optimal because of the following operations:

- Since f is not given on a uniform grid, computing $\{\{\hat{f}_{\mathbf{n}}(\tau)\}_{n_1=1}^p\}_{n_2=1}^p$, requires $\mathcal{O}(Mp^2)$, can be very expensive as we typically expect $p^2 = \mathcal{O}(M)$. (For example, in the case in which we need to build a high-order representation of the solution, we would have that as we increase the accuracy we need more targets M and the value at each target should be increasingly accurate.)
- Also, computing the inverse sine transform of $C_{\mathbf{n}}$ can be expensive, since the evaluation points are nonuniformly distributed.
- Finally, computing the Gauss transform of f at M target locations from the M leaf nodes requires $\mathcal{O}(M^2)$ work.

Accuracy. The computation of $u_{\mathbf{F}}, u_{\mathbf{L}}$ involves the numerical evaluation of the sine transform (9), Gauss transform (14), and time integrals (11) and (12). The main challenges in constructing high-order convergent schemes for these integrals are the following:

- The kernel $e^{-|\mathbf{n}|^2 \pi^2 t}$ in the far part update (11) is sharply peaked for large values of $|\mathbf{n}|$ (see Fig. 1(a)). Special quadrature rules are needed to compute it accurately for all $|\mathbf{n}|$.
- As shown in Fig. 1(b), the kernel in the local part is sharply peaked near the evaluation point (\mathbf{x}, t) . Thus, the numerical evaluation of Gauss transform (14) will not be straightforward.
- The input $f(\mathbf{x}, t)$ is given on adaptive data structures. Hence, the numerical integration schemes should be designed to handle adaptivity.

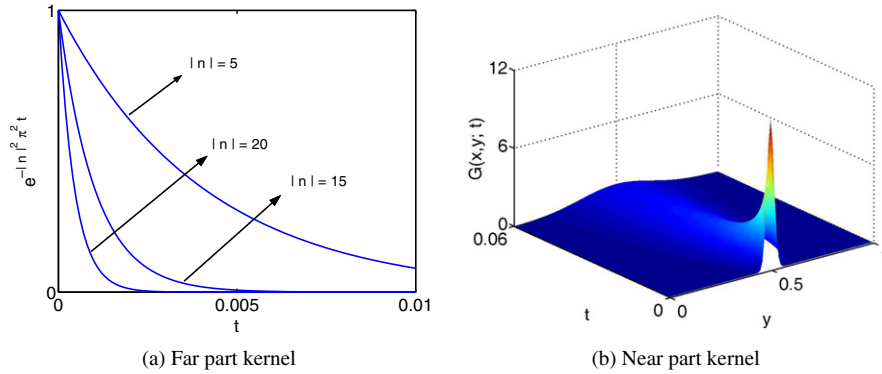


Fig. 1. Plot of kernels that appear in the far and near parts of the volume potential. The computation of far part update (11) requires convolution with the kernel $e^{-|n|^2 \pi^2 t}$ that is plotted in (a) for different values of $|n|$, with $t = 0.01$. The kernel is sharply peaked near $t = 0$ for higher values of $|n|$. In (b) we plot the one-dimensional kernel $e^{-|x-y|^2/4t}/4\pi t$ that appears in the local part evaluation, with $x = 0.5$. At $t = 0$ and as $y \rightarrow x$, the kernel is sharply peaked. Hence, a direct discretization of (11), (14) using standard quadrature rules is not computationally efficient.

1.2. Synopsis of our method

One can “easily” design schemes that are high-accurate but are too expensive to be practical. Our goal is to devise a method that can deliver both accuracy and speed. With this design goal, we have chosen a Chebyshev polynomial representation for f . As we will see in later sections, this choice allows accurate precomputation of difficult integrals using product-quadratures,² fast transformations between space and frequency domains, allows adaptive representations, and is numerically stable.

The main features of our algorithm are:

- We propose a spectral version of the fast Gauss transform to optimally compute the continuous Gauss transform of a function given by a piecewise polynomial representation (Chebyshev). We shall term this approach as the “CFGT” for Chebyshev Fast Gauss Transform. CFGT has a complexity and order of accuracy that are independent of the bandwidth the Gaussian. We compute the spatial integral in the local part (14) using CFGT.
- We propose a spectral version of the nonuniform fast Fourier transform to optimally compute the Fourier coefficients of a piecewise Chebyshev polynomial function, which we term “CNUFFT” for Chebyshev Nonuniform Fast Fourier Transform. We use CNUFFT to accelerate the computation of \hat{f}_n in (9). We use the inverse nonuniform FFT (INUFFT) for the fast computation of discrete sum (7).
- We compute the Chebyshev representation of $\hat{f}_n(\tau)$ and then compute the time integral (11) exactly using recurrence relations. Since the kernel in (11) is not approximated, the convergence of this scheme is independent of $|n|$.

In a nutshell, the main contributions of this work are the extensions of the nonuniform FFT and the fast Gauss transform for piecewise Chebyshev polynomial representations and combining the two transforms to get a fast, high-order accurate volume heat potential evaluation scheme.

² To illustrate the construction of the product integration rule, consider the computation of the transform, $F(x) = \int_{-1}^1 e^{-xy} f(y) dy$ at some particular x . Assume that $f(y) = \sum_{k=0}^{q-1} f_k T_k(y)$, where $T_k(y)$ is Chebyshev polynomial of order k and f_k is the corresponding coefficient of f . Then, we have

$$F(x) = \sum_{k=0}^{q-1} f_k I_k(x) \quad \text{and} \quad I_k = \int_{-1}^1 e^{-xy} T_k(y) dy.$$

The product integration rule is based on computing the moments I_k recursively starting from a base condition. In this example all the higher moments can be computed from I_0 .

1.3. Related work

Integral equations are widely used to model physical phenomena that involve diffusion processes, for example, see [6]. The need to integrate over the entire history has been a major hindrance for using integral equation formulations in large scale numerical simulations. The fast Gauss transform [14], reduces the complexity of computing the discrete Gauss transform

$$G_{\delta}(\mathbf{x}_i) = \sum_{j=1}^M q_j e^{-|\mathbf{x}_i - \mathbf{s}_j|^2 / \delta}, \quad i = 1, 2, \dots, N \quad (15)$$

from $\mathcal{O}(NM)$ to $\mathcal{O}(N + M)$. The FGT can be used to accelerate the computation of the solution of an initial value problem for the heat equation. Using potential theory, it can be shown that an interior boundary value problem for the heat equation can be obtained by solving integral equations that require evaluation of layered heat potentials on the boundary of domain. The algorithm proposed in [13] provides a strategy for the rapid evaluation of such potentials. Fast algorithms for the exterior domain problems were developed in [12].

Along with fast summation techniques, high-order convergent schemes are also required for the efficient evaluation of the heat potentials. To our knowledge, the only fast, high-order method for the diffusion equation based on integral equations, is that proposed by Strain in [22]: a high-order convergence was obtained by analytic integration of the monic-polynomial approximation of the input. This method is difficult to extend beyond the presented fourth-order convergent scheme, since the monic-polynomial basis is ill-conditioned. A natural question is, “do we really need a higher-order method?” The answer, provided the solution is sufficiently regular, is yes. This is because the higher-order method is significantly faster – even in the case in which we require only a few correct digits (see Section 6).

Gauss transform. Apart from applications in computational physics, FGT has been applied in diverse fields like computer vision [8], machine learning [28,27], management sciences [4,5] and others [14]. The central idea of the FGT is to use a degenerate approximation of Gaussian kernel based on Hermite polynomials. It turns out that the most expensive part of the FGT algorithm is translating the Hermite expansions from the source to the target boxes. A plane-wave expansion for the Gaussian that makes the translation operator diagonal was proposed in [15]. The complexity of translation per target box was thereby reduced from $\mathcal{O}((2n+1)^d dp^{d+1})$ to $\mathcal{O}((2n+1)^d p^d)$, where p is the number of terms retained in the truncated expansion, d is the number of dimensions and $(2n+1)^d$ are the number of cells that cover the support of the Gaussian. In [15], the authors present a modified FGT algorithm, to further reduce the complexity of the FGT to $\mathcal{O}(3dp^d)$ by transforming efficiently local expansions of the potential to neighboring boxes using plane-wave-based translation operators. Our implementation also uses such plane-wave expansions of the Gaussian.

Although much attention was devoted in the acceleration of the discrete Gauss transform, only the algorithm of Strain [22] discusses the fast computation of the continuous Gauss transform. In this paper, we extend that work to fast and higher-order computation of the continuous Gauss transform of a function defined by a piecewise Chebyshev polynomial representation.

Nonuniform FFT. We do not attempt to review all the related literature on the nonuniform FFT algorithm. The method was first introduced by [7]. A sample of articles that discuss the application of nonuniform FFT to accelerate the computation of Fourier coefficients for piecewise smooth functions include [3,9,20]. The first one uses projection of f on a spline multiresolution analysis which is then used to compute the Fourier coefficients; the second one uses Lagrange interpolation, and the last one considers piecewise constant functions. In all these papers, there is a discussion for fast computation in the case of non-smooth data. In this paper, we consider the situation in which the data is given in terms of piecewise Chebyshev polynomial coefficients rather than samples at a set of points. The method can cope with discontinuities along the edges of the quadtree nodes, but is not an efficient scheme for general discontinuous functions.

Product integration. Recursive product integration rules are not new. Integration based on recursion using Chebyshev polynomials dates back to the earlier papers by Clenshaw and Curtis [25] and Filippi [10]. The ideas were generalized by Piessens and Branders [17] to compute a variety of integral transforms. We only need recurrences for two kernels in the evaluation of (3). We construct the recurrences for these kernels in our context, which are a trivial extension of the ones presented in [17]. These kind of recurrences were also used in [24] in the context of developing high-order solvers for one dimensional diffusion equation in moving domains.

1.4. Contents

In Section 2, we discuss Chebyshev polynomials and give the construction of the product integration rules for the heat kernel. We discuss the CFGT and CNUFFT in Sections 3 and 4, respectively. In Section 5 we present the overall algorithm for the evaluation of the heat potential. We discuss numerical experiments for the transforms and the heat equation in Section 6.

2. Chebyshev representation and product integration

First, we review a few basic properties of Chebyshev polynomials (see [18] for more detailed exposition). A n th-order Chebyshev polynomial denoted by $T_n(\xi)$ defined on the interval $\xi \in [-1, 1]$ has a closed form expression given by $T_n(\xi) = \cos(n \cos^{-1} \xi)$ and $T_n(\xi)$ can be constructed from lower-order polynomials using the recurrence

$$T_n(\xi) = 2\xi T_{n-1}(\xi) - T_{n-2}(\xi), \quad T_0(\xi) = 1, \quad T_1(\xi) = \xi. \tag{16}$$

We denote the leaf nodes of a quadtree Q by $\ell_1, \ell_2, \dots, \ell_M$. In a leaf node ℓ and at a time $\tau \in [t, t + \Delta t]$, the function $f(\mathbf{x}, \tau)$ is represented using its Chebyshev coefficients

$$f(\mathbf{x}, \tau) \approx \sum_{n=0}^{q-1} \sum_{m=0}^{q-n-1} f_{n,m}^\ell(\tau) T_n(\xi_1) T_m(\xi_2), \quad (\xi_1, \xi_2) \in [-1, 1]^2, \tag{17}$$

$$\text{and } f_{n,m}^\ell(\tau) \approx \sum_{k=0}^{q-1} f_{n,m,k}^\ell T_k\left(\frac{2}{\Delta t}(\tau - t) - 1\right), \tag{18}$$

where (ξ_1, ξ_2) are local coordinates of ℓ . Since we include all the polynomials of order $q - 1$, this is a q th-order approximation of $f(\mathbf{x}, \tau)$. We use indexing scheme of [2] to represent the quadtree.

Corresponding to each time step $[k\Delta t, (k + 1)\Delta t]$, we have a quadtree Q_k . Each leaf node of Q_k contains a space–time Chebyshev polynomial representation (17, 18) of f . We note that the number of leaf nodes in each Q_k need not be the same.

2.1. Recurrence relations

To illustrate the ideas for the construction of product integration rules, we discuss the computation of $I(\alpha) = \int_0^1 e^{-\alpha x} f(x) dx$, where $f(x)$ is a locally smooth function and α is a real (or complex) number. The interval $[0, 1]$ is divided into arbitrarily-sized cells, such that in each cell, a q th-order Chebyshev polynomial approximation of $f(x)$ is within the prescribed accuracy. The integration is carried out over each cell C with center c and radius r :

$$I(\alpha) = \sum_C \sum_{n=0}^{q-1} f_n \int_{c-r}^{c+r} e^{-\alpha x} T_n\left(\frac{x-c}{r}\right) dx = \sum_C \sum_{n=0}^{q-1} f_n I_n(\alpha). \tag{19}$$

We derive recurrence relationships to compute the moments I_n in a more general setting. For any constants x_a, x_b, λ, η such that $-1 \leq \lambda x_a + \eta \leq \lambda x_b + \eta \leq 1$,

$$I_n = \int_{x_a}^{x_b} e^{-\alpha x} T_n(\lambda x + \eta) dx, \tag{20}$$

$$= \left[\frac{e^{-\alpha x}}{2\lambda} \left(\frac{T_{n+1}}{n+1} - \frac{T_{n-1}}{n-1} \right) \right]_{x_a}^{x_b} + \alpha \int_{x_a}^{x_b} \left[\frac{e^{-\alpha x}}{2\lambda} \left(\frac{T_{n+1}}{n+1} - \frac{T_{n-1}}{n-1} \right) \right], \tag{21}$$

$$\Rightarrow I_{n+1} = c_{n+1} + (n+1) \left(\frac{2\lambda}{\alpha} I_n + \frac{I_{n-1}}{n-1} \right),$$

$$\text{where } c_{n+1} = \frac{n+1}{\alpha} \left[e^{-\alpha x} \left(\frac{T_{n+1}}{n+1} - \frac{T_{n-1}}{n-1} \right) \right]_{x_a}^{x_b}. \tag{22}$$

If $\alpha \gg 1$, then $e^{-\alpha x}$ is sharply peaked and hence computing $I(\alpha)$ using numerical quadratures is quite expensive. Whereas, the convergence rate of the product integration rule discussed here depends only on the order of approximation of the input $f(x)$. Since we used polynomials of order $q - 1$, the overall rate of convergence is q . However, the recurrence (22) is unstable for smaller values of $|\alpha|$ and hence higher moments cannot be computed accurately. In this case, we resort to Olver’s algorithm [26]. (We give details in the Appendix.)

3. The Chebyshev fast Gauss transform

In this section, we present a fast, high-order algorithm for evaluating the continuous Gauss transform defined by

$$G_\delta f(\mathbf{x}) = \int_\Omega e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\delta}} f(\mathbf{y}) d\mathbf{y}, \quad \Omega = [0, 1]^2. \tag{23}$$

The input is a quadtree with leaf nodes $\{\ell_k\}_{k=1}^M$. Each leaf node contains a q th-order Chebyshev polynomial representation of f . The aim is to evaluate $\{G_\delta f(\mathbf{x}_i)\}_{i=1}^M$ at M targets within a pre-specified accuracy ϵ , for any required δ and in an optimal $\mathcal{O}(M)$ time.

For a given δ , the Gaussian centered at \mathbf{x} decays exponentially outside a fixed interval $\mathbf{x} + [-r_g\sqrt{\delta}, r_g\sqrt{\delta}]^2$, where $r_g = \sqrt{\log(1/\epsilon)}$. We call this interval the *support* of Gaussian. Below we describe two methods to compute $G_\delta f$: the direct evaluation and the kernel expansion-based algorithm. The overall algorithm combines these two methods. In this way, we obtain an algorithmic complexity that is independent of δ .

3.1. Direct evaluation

We can compute $G_\delta f$ from the piecewise polynomial representation of f using the following steps:

- We find the set of leaf nodes that cover the support of Gaussian centered at \mathbf{x} . We shall call this set the *interaction list* of \mathbf{x} and denote it by $\mathcal{I}[\mathbf{x}]$.
- We truncate the domain of integration to $\mathcal{I}[\mathbf{x}]$.
- We compute and add the contribution of each $\ell \in \mathcal{I}[\mathbf{x}]$, to obtain $G_\delta f(\mathbf{x})$.

The contribution of a leaf node ℓ , with center \mathbf{c} and side length $2r$, to $G_\delta f(\mathbf{x})$ is given by

$$\begin{aligned} G_\delta^\ell(\mathbf{x}) &= \int_{c_1-r}^{c_1+r} \int_{c_2-r}^{c_2+r} e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\delta}} f(\mathbf{y}) d\mathbf{y} = \sum_{k=0}^{q-1} \sum_{j=0}^{q-k-1} f_{k,j}^\ell \int_{-1}^1 e^{-\left(\frac{x_1-c_1}{\sqrt{\delta}} - \frac{r}{\sqrt{\delta}}\xi_1\right)^2} T_k(\xi_1) d\xi_1 \int_{-1}^1 e^{-\left(\frac{x_2-c_2}{\sqrt{\delta}} - \frac{r}{\sqrt{\delta}}\xi_2\right)^2} T_j(\xi_2) d\xi_2 \\ &= I^T[x_1 - c_1] \mathbf{f}^\ell I[x_2 - c_2], \end{aligned}$$

where \mathbf{f}^ℓ is the matrix that contains the Chebyshev coefficients of f in ℓ . Here, we performed the change of variables $\xi_k = \frac{y_k - c_k}{r}$, $k = 1, 2$. The entries of the moment vectors are given by

$$I_j[x_k - c_k] = \int_{-1}^1 e^{-\left(\frac{x_k - c_k}{\sqrt{\delta}} - \frac{r}{\sqrt{\delta}}\xi_k\right)^2} T_j(\xi_k) d\xi_k, \quad k = 1, 2 \text{ and } 0 \leq j < q. \tag{24}$$

These moments are computed exactly using the recurrence relation (67) and appropriate scaling factors λ, η . The direct evaluation is useful when the size of Gaussian support is small compared to the smallest leaf node; in this case we need $\mathcal{O}(Mq^2)$ work to evaluate $G_\delta f$ at M target points.

3.2. Degenerate kernel expansion

When the support of the Gaussian spans $\mathcal{O}(M)$ leaf nodes, direct evaluation leads to $\mathcal{O}(M^2)$ complexity. Instead, following [14,15], we use the following plane-wave expansion of the Gaussian

$$e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\delta}} = \frac{1}{4\pi} \int_{\mathbb{R}^2} e^{-\frac{|z|^2}{4}} e^{\frac{i\mathbf{z}(\mathbf{x}-\mathbf{y})}{\sqrt{\delta}}} d\mathbf{z}. \tag{25}$$

The numerical evaluation of the above equation involves the following approximations:

- The domain of integration is truncated to a finite interval $[-L, L]^2$.
- Numerical integration of the truncated plane-wave expansion.

The truncation of \mathbb{R}^2 by $[-L, L]^2$ results in a e^{-L^2} rate of convergence for the trapezoidal rule. Indeed, let $2p + 1$ be the number of trapezoidal nodes. Then, in [23] it is shown that

$$e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\delta}} \approx \frac{L^2}{4\pi p^2} \sum_{-p \leq \mathbf{k} \leq p} e^{-\left(\frac{L|\mathbf{k}|}{4p}\right)^2} e^{\frac{iL\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})}{p\sqrt{\delta}}}. \tag{26}$$

Following [14], we use the multi-index notation by which $-p \leq \mathbf{k} \leq p$ implies $-p \leq k_j \leq p$, for $j = 1, 2$. Clearly p depends on how oscillatory the integrand in the truncated plane-wave expansion is, which in turn depends on $|\mathbf{x} - \mathbf{y}|_{\max} = r_g \sqrt{\delta}$. For a required precision $\epsilon = 10^{-14}$ the parameters should be

$$r_g = 6, \quad L = 11 \quad \text{and} \quad p = 21. \tag{27}$$

The domain $\Omega = [0, 1]^2$ is divided into *uniform boxes* of length $2r_b$. In our implementation, we set $r_b = 2^{\lceil \log_2 \sqrt{\delta} \rceil - 1}$ so that a particular box is either fully embedded in some leaf node or vice-versa.

Remark. Note that we have two grids: the first one is a quadtree to represent $f(\mathbf{y})$ and the second one is a uniform grid required by FGT (see Fig. 2(a)).

Since $r_b = \mathcal{O}(\sqrt{\delta})$, the support of Gaussian centered at a particular box C spans a fixed number of boxes. These boxes belong to C 's *interaction list*, denoted by, $\mathcal{I}[C]$.³ Suppose box B belongs to $\mathcal{I}[C]$. We introduce the following terminology:

\mathbf{w}^B	the wave expansion of box B .
\mathbf{t}^{BC}	the operator that translates the wave expansion of box B to C .
\mathbf{I}^C	the local expansion of box C .

Let $\lambda_{\mathbf{k}} = \frac{iL\mathbf{k}}{p\sqrt{\delta}}$ and $\mathbf{x}_C, \mathbf{y}_B$ denote the centers of boxes C, B , respectively. Then, we define,

$$\mathbf{w}_{\mathbf{k}}^B = \int_B e^{i\lambda_{\mathbf{k}} \cdot (\mathbf{y}_B - \mathbf{y})} f(\mathbf{y}) d\mathbf{y}, \quad -p \leq \mathbf{k} \leq p, \tag{28}$$

$$\mathbf{t}_{\mathbf{k}}^{BC} = e^{i\lambda_{\mathbf{k}} \cdot (\mathbf{x}_C - \mathbf{y}_B)}, \tag{29}$$

$$\mathbf{I}^C = \sum_{B \in \mathcal{I}[C]} \mathbf{w}^B \odot \mathbf{t}^{BC}, \tag{30}$$

$$\mathbf{v}_{\mathbf{k}}^C(\mathbf{x}) = \frac{L^2}{4\pi p^2} e^{-\left(\frac{L|\mathbf{k}|}{2p}\right)^2} e^{i\lambda_{\mathbf{k}} \cdot (\mathbf{x} - \mathbf{x}_C)}, \tag{31}$$

where \odot symbolizes element-wise multiplication (Hadamard product). Following this notation, the Gauss transform at a particular point \mathbf{x} that belongs to a target box C can be written as

$$G_{\delta} f(\mathbf{x}) \approx \sum_{-p \leq \mathbf{k} \leq p} \mathbf{v}_{\mathbf{k}}^C(\mathbf{x}) \mathbf{I}_{\mathbf{k}}^C. \tag{32}$$

The main steps involved in FGT can be summarized as follows:

- At each source box B form \mathbf{w}^B .
- At each target box C gather all the wave expansions from boxes belonging to $\mathcal{I}[C]$ and form \mathbf{I}^C .
- From \mathbf{I}^C compute the Gauss transform at each target that belongs to C .

³ Notice that we already used the symbol \mathcal{I} in the direct evaluation version and we also use the same symbol in the next Section in a different context. The meaning of \mathcal{I} should be clear from the context.

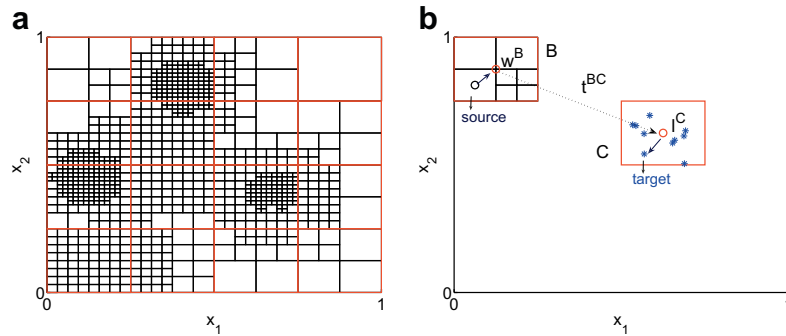


Fig. 2. Here we illustrate one of the main steps involved in FGT. (a) The leaf nodes (shown in black) of the given quadtree are assigned to regular boxes (shown in red). The size of the boxes is chosen so that each box corresponds to a node at some level of the quadtree. (b) At a source box B , the influence of all the constituent leaf nodes is encoded in the wave expansion \mathbf{w}^B . The translation operator \mathbf{t}^{BC} transfers the wave expansion of B to the target box C . At C , the influence of all the source boxes in $\mathcal{I}[C]$ is encoded in the local expansion \mathbf{I}^C . Finally, $G_{\delta f}$ at a target that belongs to C is given by $(\mathbf{I}^C, \mathbf{v}^C(\mathbf{x}))$. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

For additional details with respect to the implementation of the FGT algorithm, we refer the reader to [14,15]. The main difference of CFGT from FGT is the computation of \mathbf{w} at each source box, which we discuss next.

Computing wave expansions. The computation of \mathbf{w}^B involves finding the contribution of the leaf nodes within⁴ the box B (see Fig. 2(b))

$$\mathbf{w}^B = \sum_{\ell \in B} \mathbf{w}^\ell \odot \mathbf{t}^{\ell B}. \quad (33)$$

The definitions for \mathbf{w}^ℓ and $\mathbf{t}^{\ell B}$ are similar to (28) and (29) i.e.,

$$\mathbf{w}_{\mathbf{k}}^\ell = \int_{\ell} e^{\lambda_{\mathbf{k}} \cdot (\mathbf{c} - \mathbf{y})} f(\mathbf{y}) d\mathbf{y}, \quad \mathbf{t}_{\mathbf{k}}^{\ell B} = e^{\lambda_{\mathbf{k}} \cdot (\mathbf{y}_B - \mathbf{c})}. \quad (34)$$

Let \mathbf{c} be the center and $2r$ be the length of ℓ . Then, the wave expansion of ℓ can be computed as follows:

$$\mathbf{w}_{\mathbf{k}}^\ell = \sum_{n=0}^{q-1} \sum_{j=0}^{q-n-1} f_{n,j}^\ell \int_{c_2-r}^{c_2+r} e^{\lambda_{k_2}(c_2-y_2)} T_n\left(\frac{y_2-c_2}{r}\right) dy_2 \int_{c_1-r}^{c_1+r} e^{\lambda_{k_1}(c_1-y_1)} T_n\left(\frac{y_1-c_1}{r}\right) dy_1. \quad (35)$$

By the change of variables $\zeta = \frac{y_k - c_k}{r}$ for $k = 1, 2$ we have

$$\mathbf{w}_{\mathbf{k}}^\ell = r^2 \sum_{n=0}^{q-1} \sum_{j=0}^{q-k-1} f_{n,j}^\ell \int_{-1}^1 e^{-\lambda_{k_2} r \zeta} T_n(\zeta) d\zeta \int_{-1}^1 e^{-\lambda_{k_1} r \zeta} T_j(\zeta) d\zeta \quad \text{and} \quad \mathbf{w}^\ell = \mathbf{I}^T[r] \mathbf{f}^\ell \mathbf{I}[r], \quad (36)$$

$$\text{where} \quad \mathbf{I}_{jk}[r] = r \int_{-1}^1 e^{-\lambda_k r \zeta} T_j(\zeta) d\zeta, \quad -p \leq k \leq p \text{ and } 0 \leq j < q. \quad (37)$$

Each column of \mathbf{I} is computed using the recurrences (22).

Acceleration techniques. Although there seem to be lot of exponential evaluations, most of them can be precomputed.

1. Given the depth of the quadtree, we precompute the moments (37) for each level of the tree.
2. The translation operator $\mathbf{t}^{\ell B}$ depends on the relative position of ℓ within B . We first determine the distinct relative positions possible for the given quadtree and box size r_b . Then, we precompute the translation operator for each of these positions.

⁴ Apparently, it is possible that the box B could also be contained in a leaf node. It would be evident after we formally state CFGT that in this case, the leaf node containing B interacts with the targets directly instead of forming wave expansions.

- The operator \mathbf{t}^{BC} depends on the relative position of B and C . The number of these is bounded from above by $(2n + 1)^2$ and we precompute the corresponding translation operators.

3.3. The complete algorithm

The implementation of the original FGT [14] is not efficient for the cases in which δ is too small. In our context, δ is dictated by the quadrature nodes in time that are used to evaluate $u_L(x, t)$. These nodes can be arbitrarily close to zero. We now describe an algorithm that attains optimal complexity for any δ .

We split the given quadtree into two trees (Fig. 3): one tree in which the size of the leafs is greater than r_b and one tree that all its leaves have a size that is smaller or equal than r_b (Fig. 3(c)). The Gauss transform at the targets is the sum of contributions from these two trees. In the first tree we use the direct evaluation version, and in the second we use the kernel expansion-based algorithm.

Algorithm 1. Chebyshev fast Gauss transform

```

SPLIT THE TREE
for  $k = 1 : M$  do
  if  $r_k \leq r_b$  then
    Assign the leaf node  $Q(k)$  to  $Q_{\text{Direct}}$ .
  else
    Assign  $Q(k)$  to  $Q_{\text{Expand}}$ .
  end if
end for
 $G_\delta f(\mathbf{x}) = \text{DirectEvaluation}(Q_{\text{Direct}}, \mathbf{x}) + \text{KernelExpansion}(Q_{\text{Expand}}, \mathbf{x})$ .

```

Algorithm 2. DirectEvaluation

```

INITIALIZATION
for  $k = 1 : M$  do
   $G_\delta f(\mathbf{x}_k) = 0$ .
end for
EVALUATION
for  $k = 1 : M$  do
  Determine  $\mathcal{I}[\mathbf{x}_k]$  (the leaf nodes in the interaction list).
  for each  $\ell \in \mathcal{I}[\mathbf{x}_k]$ 
    Add contribution of each leaf node in the interaction list
    Compute the moments  $I[x_1 - c_1]$  and  $I[x_2 - c_2]$  (Eq. (24)).
     $G_\delta f(\mathbf{x}_k) += I^T[x_1 - c_1] \mathbf{f}^\ell I[x_2 - c_2]$ .
  end for
end for

```

Complexity. Let the number of leaf nodes be M and the number of targets be N . The cost of direct evaluation is $\mathcal{O}((2n + 1)^2 q^2 N \log M)$. The factor $(2n + 1)^2 \log M$ is the cost of finding the at most $(2n + 1)^2$ leafs in the interaction list of each target and q^2 is the cost of computing moments (67) at each target.

Computing the wave expansions at the leaf nodes is⁵ $\mathcal{O}((2p + 1)^2 q^2 N)$ since computing (36) for each \mathbf{k} is $\mathcal{O}(q^2)$ and we have $-p \leq \mathbf{k} \leq p$. The translation $\mathbf{w}^\ell \rightarrow \mathbf{w}^B$ requires $\mathcal{O}((2p + 1)^2)$. Therefore, computing the wave expansions at all the source boxes requires $\mathcal{O}((2p + 1)^2 q^2 N)$. The computation of the local expansion at a target box requires visiting all the $(2n + 1)^2$ source boxes in its interaction list. In [15], the authors discuss

⁵ A q th-order Chebyshev polynomial approximation has $q(q + 1)/2$ coefficients. By using symmetries, we can show that the computation of wave expansions at the leaf nodes can be reduced to $\mathcal{O}(N(p + 1)(p + 2)q(q + 1)/4)$.

a strategy to compute the local expansions by just visiting the neighbors. Using this strategy,⁶ the complexity of forming the local expansion at a target box is $\mathcal{O}((2p+1)^2)$. Computing the Gauss transform at the targets from the local expansions is $\mathcal{O}((2p+1)^2N)$. Hence, the overall complexity of the kernel expansion version of CFGT is $\mathcal{O}((2p+1)^2q^2M + (2p+1)^2N)$.

Algorithm 3. KernelExpansion

INITIALIZATION

Precompute the moments I (equation (37)), given the depth of the quadtree.

FLAGGING

Determine the source and target boxes that interact

$\mathcal{B}_0 :=$ the minimal set of source boxes that covers all the leaf nodes.

$\mathcal{C}_0 :=$ the minimal set of target boxes that covers all the targets.

Find the source boxes that are in the interaction list of target boxes

$\mathcal{B} := \mathcal{B}_0 \cap \mathcal{I}[\mathcal{C}_0]$.

Find the target boxes that are in the interaction list of source boxes

$\mathcal{C} := \mathcal{C}_0 \cap \mathcal{I}[\mathcal{B}]$.

COMPUTE WAVE EXPANSIONS

for each $B \in \mathcal{B}$ **do**

for each $\ell \in B$ **do**

Compute the wave expansion of ℓ

$\mathbf{w}^\ell = I^T[r]\mathbf{f}^\ell I[r]$.

Add the contribution of ℓ to wave expansion of B

$\mathbf{w}^B += \mathbf{w}^\ell \odot \mathbf{t}^{\ell B}$.

end for

end for

COMPUTE LOCAL EXPANSIONS

for each $C \in \mathcal{C}$ **do**

for each box $B \in \mathcal{B} \cap \mathcal{I}[C]$ **do**

Translate the wave expansion of source box B to target box C

$\mathbf{l}^C += \mathbf{w}^B \odot \mathbf{t}^{BC}$.

end for

end for

EVALUATE AT THE TARGETS

for each $C \in \mathcal{C}$ **do**

for each $\mathbf{x} \in C$ **do**

$G_\delta f(\mathbf{x}) = \langle \mathbf{v}^C(\mathbf{x}), \mathbf{l}^C \rangle$.

end for

end for

4. The Chebyshev nonuniform fast Fourier transform

We present the algorithm in one dimensional setting. Consider the computation of Fourier coefficients of $f(x)$ for $x \in [0, 2\pi]$, defined by

$$\hat{f}_k = \int_0^{2\pi} f(y) e^{-iky} dy, \quad k = -\frac{M}{2}, \dots, \frac{M}{2} - 1. \quad (38)$$

⁶ In order to use it, it is convenient to sort the target boxes using Morton ordering so that the local expansions of spatially close boxes are formed consecutively.




Fig. 3. In this figure we illustrate the tree splitting that we use in CFGT to obtain an algorithmic complexity that is independent of the value of δ . Subfigure (a) shows the quadtree, a target \mathbf{x} , and the support of the Gaussian centered at \mathbf{x} (blue box). For this δ , we can observe that a direct evaluation would be expensive because we have to loop through all the leaf nodes within the support. A kernel expansion-based approach would also be expensive because it generates too many boxes. Hence, for optimal complexity, we split the tree into two parts. Subfigure (b) shows the tree with the size of leaf nodes being greater than r_b . To evaluate the Gauss transform at a point \mathbf{x} , we just need to visit the leaf nodes in the support of the Gaussian (the shaded ones). There can be a maximum of $(2n + 1)^2$ of such leaf nodes. Subfigure (c) show the tree with the size of leaf nodes being smaller or equal than r_b . In this tree we use the kernel expansion-based algorithm. Instead of visiting all the leaf nodes within the support, we only have to gather the wave expansion of the source boxes (red boxes) that are within the support. The number of these source boxes is bounded from above by $(2n + 1)^2$. (For interpretation of the references in color in this figure legend, the reader is referred to the web version of this article.)

The input is a Chebyshev polynomial representation of $f(x)$ on a binary tree data structure imposed on $[0, 2\pi]$ with M leaf nodes. A direct computation of (38) requires $\mathcal{O}(M^2)$ work. Here, we describe an extension of non-uniform FFT that reduces the work to $\mathcal{O}(M \log M)$. (We follow the notation in [16].)

The first step in nonuniform FFT is to convolve $f(x)$ with the periodic heat kernel $g_\tau(x) = \sum_{k=-\infty}^{\infty} e^{-(x-2k\pi)^2/4\tau}$, the result of which would be a periodic function denoted by $f_\tau(x)$:

$$f_\tau(x) = f^* g_\tau(x) = \int_0^{2\pi} f(y) g_\tau(x - y) dy. \tag{39}$$

Since $f_\tau(x)$ is smooth and 2π -periodic, its Fourier coefficients can be computed by the trapezoidal rule with M_r nodes and by computing the discrete sum using FFT. So, the task at hand is to compute f_τ on an oversampled regular grid i.e., at points $x_j = jh, j = 0, 1, \dots, M_r - 1$, where $h = 1/M_r$. The oversampling ratio $R = M_r/M$ is chosen based on the desired accuracy. In the definition of g_τ , each Gaussian has a support given by $4r_g\sqrt{\tau}$, where $r_g = \sqrt{\log(1/\epsilon)}$, for desired precision ϵ . The number of oversampled grid points within the support of a Gaussian in each direction, denoted by M_{sp} , is given by $M_{sp} = 2r_g\sqrt{\tau}/h$.

Consider the computation of $f_\tau(x_j)$. The kernel $g_\tau(x_j - y)$ in (39), is the sum of the Gaussian centered at x_j and its images. Depending on the choice of τ and the location of target x_j , only a few images of g_τ are non-negligible within $[0, 2\pi]$. For illustration purposes, it is sufficient to consider just the Gaussian centered at x_j . In Fig. 4, the support of Gaussian spreads over three leaf nodes ℓ_1, ℓ_2 and ℓ_3 . The interval of integration in (39) can thus be truncated to these three leaf nodes. The contribution of ℓ_1 that has Chebyshev coefficients $\{f_k^{\ell_1}\}_{k=0}^{q-1}$ is given by

$$f_\tau^{\ell_1}(x_j) = \sum_{k=0}^{q-1} f_k^{\ell_1} I_k, \quad \text{where } I_k = \int_{c_1-r_1}^{c_1+r_1} e^{-\frac{(x_j-y)^2}{4\tau}} T_k\left(\frac{y-c_1}{r_1}\right) dy, \tag{40}$$

$$I_k = \lambda \int_{-1}^1 e^{-\left(\frac{y-\eta}{\lambda}\right)^2} T_k(y) dy, \quad \text{where } \lambda = \frac{2\sqrt{\tau}}{r_1}, \quad \eta = \frac{x_j - c_1}{r_1}. \tag{41}$$

The contributions of ℓ_2 and ℓ_3 are computed similarly. Hence, the computation of f_τ at a particular target involves visiting all the leaf nodes in which g_τ is non-negligible. Instead of finding the leaf nodes that influence a target x , we loop over the leaf nodes and determine which targets they influence. This is efficient because all the targets belong to a Cartesian grid.




Fig. 4. The support of a Gaussian centered at x_j covers the leaf nodes ℓ_1, ℓ_2 and ℓ_3 . Hence, only three leaves contribute to $f_\tau(x_j)$. The contribution of each leaf node is given by the inner product of the constituent polynomial coefficients of f and the moments of polynomials with the Gaussian. These moments are recursively computed using equation (67), which is discussed in the Appendix.

Notation. The interaction list of a particular leaf node ℓ_k is the set of all regular grid points influenced by it, denoted by $\mathcal{I}[\ell_k]$.

Notice that computing f_τ at a particular target is equivalent to computing the Gauss transform of f . However, we do not require FGT here because the targets belong to a Cartesian grid. The following lemma establishes the complexity of computing f_τ at the oversampled regular grid points.

Lemma 4.1. *The cumulative sum of length of the interaction lists of all the leaf nodes is $\mathcal{O}(M)$.*

Proof. Consider a particular leaf node ℓ_k . The number of targets within ℓ_k is bounded from above by $\frac{2r_k}{h_r} + 1$, where $h_r = 2\pi/RM$. The number of targets in $\mathcal{I}[\ell_k]$ that are to the right (and similarly to the left) of ℓ_k is given by M_{sp} . Therefore, we have

$$\sum_{k=1}^M \text{length}(\mathcal{I}[\ell_k]) \leq \sum_{k=1}^M \frac{2r_k RM}{2\pi} + 2M_{\text{sp}} + 1 = (2M_{\text{sp}} + 1)M + RM \sum_{k=1}^M \frac{r_k}{\pi}.$$

The sum of lengths of leaf nodes ($\sum_{k=1}^M 2r_k$) is the length of domain (2π), and thus

$$\sum_{k=1}^M \text{length}(\mathcal{I}[\ell_k]) \leq (2M_{\text{sp}} + 1 + R)M. \tag{42}$$

It follows that the complexity of computing f_τ , at M_r equispaced points, is $\mathcal{O}(2M(1 + R + M_{\text{sp}})q)$, since each leaf node contains q polynomial coefficients. \square

Once f_τ is computed at all the points on the oversampled grid, FFT is used to compute $\hat{f}_\tau(k)$ in $\mathcal{O}(M_r \log M_r)$ time. The Fourier coefficients of f are extracted from \hat{f}_τ using the relation [16]

$$\hat{f}_k = \sqrt{\frac{\pi}{\tau}} e^{k^2 \tau} \hat{f}_{\tau k}. \tag{43}$$

The accuracy of the transforms $f \rightarrow f_\tau \rightarrow \hat{f}_\tau$, depends on the choice of parameters τ, R, M_{sp} . The details of the analysis can be found in [7]. We follow [16] and note that for double precision $M_{\text{sp}} = 12, R = 4$ and $\tau = \pi^2/4M^2$.

Precomputation. Although the moments I in equation (40) can be computed in constant time on the fly using the recurrence (67), we can further accelerate the computations by precomputing the Chebyshev moments of the Gaussian for a set of scaling values that only depend on the depth of the tree. We set $M_r = 2^{\lceil \log_2 RM \rceil}$ to minimize the number of such distinct values (see Appendix, for λ and η).

Sine transform. To compute the far part u_F , we need to accelerate the sine transform (9) computation. To accomplish this, CNUFFT can be used. The 1D analogue of sine transform is given by

$$f_n^s = \int_0^1 f(y) \sin(n\pi y) dy, \quad n = 1, 2, \dots, M. \tag{44}$$

We cannot use the Fourier coefficients of f to compute its sine transform, since the period of underlying basis for both these transforms differ by a factor of two. Instead, we extend f to the domain $(0, 2)$ by

$$F(y) = \begin{cases} f(y) & y \in (0, 1), \\ 0 & \text{otherwise.} \end{cases} \tag{45}$$

The data structure for F is obtained by adding an extra leaf node to the binary tree representation of f with $c = \frac{3}{2}$, $r = \frac{1}{2}$ and setting the q polynomial coefficients of this leaf node to zero. The transforms on F are given by

$$F_n^s = \frac{1}{2} \int_0^2 F(y) \sin(n\pi y) dy; \quad \widehat{F}_n = \frac{1}{2} \int_0^2 F(y) e^{-n\pi y} dy. \tag{46}$$

We can compute \widehat{F} using CNUFFT and then F^s can be computed using the relation $F_n^s = -\text{Im}[\widehat{F}_n]$. The sine transform of f can then be retrieved from the relation $f_n^s = \frac{1}{2} F_n^s$. In general, the extension $f \rightarrow F$ will be discontinuous. We do not lose accuracy since CNUFFT requires only piecewise smoothness.

4.1. 2D CNUFFT

Input: A quadtree data structure \mathcal{Q} with M leaf nodes, each containing the q th-order Chebyshev coefficients of $f(\mathbf{x})$.⁷

Output: The $M_o \times M_o$ Fourier coefficients $\widehat{f}_{k_1 k_2}$, where $-\frac{M_o}{2} \leq k_1, k_2 \leq \frac{M_o}{2} - 1$.

Algorithm 4. 2D CNUFFT

INITIALIZATION

Given target precision, choose R, τ, M_{sp} .

Set $M_r = 2^{\lceil \log_2 R M_o \rceil}$.

Precompute the FGT moments I for all possible scaling values.

CONVOLUTION

for $k = 1 : M$ **do**

Determine $\mathcal{I}[\ell_k]$ (the interaction list of ℓ_k).

for each $\mathbf{x} \in \mathcal{I}[\ell_k]$ **do**

 Compute t scaling factors $\{\lambda_i, \eta_i\}_{i=1}^2$;

 For these values, extract the moments I_1, I_2 from the precomputed values.

 Add the contribution of ℓ_k to $f_\tau(\mathbf{x})$: $f_\tau(\mathbf{x}) += I_1^T \mathbf{f}^{\ell_k} I_2$.

end for

end for $\mathcal{O}(q^2 M_r^2 + (2M_{\text{sp}} + 1)^2 q^2 M)$

FFT ON OVERSAMPLED REGULAR GRID

Using FFT, compute $\widehat{f}_{\tau k_1 k_2}$, for $-\frac{M_r}{2} \leq k_1, k_2 \leq \frac{M_r}{2} - 1$ $\mathcal{O}(M_r^2 \log M_r)$.

Compute \widehat{f} from \widehat{f}_τ using analytic relations.

5. Computing heat potentials

Using CFGT and CNUFFT, we can compute the volume potential (3) in an optimal way. Here, we describe our treatment of far and local parts after providing analysis for the optimal choice of the parameters p and δ .

5.1. Optimal splitting in time

The choice of parameters $\{p, \delta\}$ affects the overall complexity and the errors due to the truncation of the kernel expansions (4) and (5). We assume that the targets are at least at a distance d away from the boundary

⁷ The linear quadtree representation of [2] stores just one index corresponding to each leaf node, the geometric information of the leaf node can be extracted from this index.

of the unit box, that is, $d = \inf\{\|\mathbf{x} - \mathbf{y}\|_2 : (\mathbf{x}, \mathbf{y}) \in \omega \times \Gamma\}$. Then, we have the following error estimates for truncation errors [19]:

$$\mathcal{E}_F(p, \delta) = \mathcal{O}\left(\frac{e^{-\pi^2 p^2 \delta}}{p\delta}\right) \quad \text{and} \quad \mathcal{E}_L(d, \delta) = \mathcal{O}\left(\frac{e^{-d^2/\delta}}{\delta}\right). \tag{47}$$

In [21,19], δ was chosen so that the truncation errors get reduced as N, M are increased. In our case, N and M are given parameters, which are decided solely based on resolving $f(\mathbf{x}, t)$ within desired accuracy. Instead, we proceed as follows. From the estimates (47), we can infer that

- Given d , we obtain $\mathcal{E}_L < \epsilon$ by requiring that $\delta \leq \delta_\epsilon$. Similarly, $\mathcal{E}_F < \epsilon$ if $p \geq p_\epsilon$.
- δ_ϵ and p_ϵ are threshold values that depend on the domains Ω and ω .

For example, if $\omega = [0.4, 0.6]^2$, then for double precision, $\delta_\epsilon = 0.002$ and $p_\epsilon = 55$.

So, for a given ϵ we must pick δ, p such that $\delta \leq \delta_\epsilon$ and $p \geq p_\epsilon$. When Δt is bigger than δ_ϵ , we set $\delta = \delta_\epsilon$ and $p = p_\epsilon$. When Δt is smaller than δ_ϵ , we have to reduce δ to preserve optimality. This in turn forces p to be more than p_ϵ to restrict \mathcal{E}_F within ϵ . We set $\delta = l\Delta t$ for some positive integer l . This means that the local part computation involves l time steps. Using CFGT, the local part requires $\mathcal{O}(Mlq^3)$ work per time step. By using CNUFFT and INUFFT, the far part requires $\mathcal{O}((M + p^2)q^3 \log p)$ work per time step. Hence, the overall complexity is $\mathcal{O}((Ml + M \log p + p^2 \log p)Nq^3)$. We chose $p = c_1\sqrt{M}$ and $l = \lceil c_2 \frac{\log M}{N\Delta t} \rceil$. Then, the truncation errors are given by

$$\mathcal{E}_F(M, N) = \mathcal{O}\left(\frac{N}{\sqrt{M} \log M} e^{-\frac{\pi^2 c_1^2 c_2 M \log M}{N}}\right), \quad \mathcal{E}_L(M, N) = \mathcal{O}\left(\frac{N}{\log M} e^{-\frac{d^2 N}{c_2 \log M}}\right). \tag{48}$$

So, by appropriately choosing the constants c_1, c_2 , we can ensure that the truncation errors are within the required accuracy. Then, the overall complexity for this choice of parameters would be $\mathcal{O}(MNq^3 \log M)$.

Next, we illustrate our construction of quadrature rules for the far and local parts for the case in which $\delta = l\Delta t$. It is straightforward to extend for the case where $\delta = \delta_\epsilon$.

5.2. Far part

The far part $u_F(\mathbf{x}, t)$ is computed by evaluating the discrete sum (7) using INUFFT. To update the coefficients C_n , we first compute the Chebyshev coefficients of $\hat{f}_n(\tau)$ and then use product integration rule to evaluate the update (11) and then the recurrence (10).

In the interval $[t - l\Delta t, t + \Delta t - l\Delta t]$, $\hat{f}_n(\tau)$ is computed at the Chebyshev nodes given by

$$\tau_i = t - l\Delta t + \frac{\Delta t}{2}(1 + \cos(\pi(i - 1)/q)), \quad i = 1, \dots, q. \tag{49}$$

At each τ_i , we use CNUFFT to compute \hat{f}_n , defined in (9). Given $\hat{f}_n(\tau_i)$, we compute its Chebyshev coefficients⁸ $\hat{f}_n(k), k = 0, \dots, q - 1$. By the change of variable $\tau = t + (1 - \theta - l)\Delta t$,

$$U_n[\hat{f}] = \Delta t e^{-|\mathbf{n}|^2 \pi^2 l \Delta t} \sum_{k=0}^{q-1} \hat{f}_n(k) \int_0^1 e^{-(|\mathbf{n}|^2 \pi^2 \Delta t)\theta} T_k(-2\theta + 1) d\theta. \tag{50}$$

For $n_1 = 1, \dots, p, n_2 = 1, \dots, p, k = 1, \dots, q$, we define

$$E_{nk} = \Delta t e^{-|\mathbf{n}|^2 \pi^2 l \Delta t} \int_0^1 e^{-(|\mathbf{n}|^2 \pi^2 \Delta t)\theta} T_k(-2\theta + 1). \tag{51}$$

Then $U_n[\hat{f}] = \sum_{k=0}^{q-1} \hat{f}_n(k) E_{nk}$. The $p^2 \times q$ entries of the matrix E can be computed in an optimal $\mathcal{O}(p^2 \times q)$ using the recurrences (22).

⁸ It is known that the q Chebyshev coefficients can be computed in $\mathcal{O}(q \log q)$ time using FFT. Also to compute \hat{f}_n , we need $f(\mathbf{y}, \tau_i)$ from the piecewise Chebyshev polynomial representation of f . This again can be accomplished in optimal time using inverse FFT.

5.3. Local part

Let us rewrite (12) as

$$u_L(\mathbf{x}, t) = \int_{t-l\Delta t}^t g(\tau) d\tau, \quad g(\tau) = \int_{\Omega} \frac{e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{4(t-\tau)}}}{4\pi(t-\tau)} f(\mathbf{y}, \tau) d\mathbf{y}. \tag{52}$$

We can easily show that g is a smooth function in time. To compute u_L , we first compute the Chebyshev coefficients of $g(\tau)$ in each of the l previous time steps:

$$u_L(\mathbf{x}, t) = \sum_{k=1}^l \int_{t-k\Delta t}^{t-(k-1)\Delta t} g(\tau) d\tau = \sum_{k=1}^l \sum_{n=0}^{q-1} g_n^k I_n, \quad \text{where } I_n = \int_{-1}^1 T_n(\xi) d\xi = \frac{1 + (-1)^n}{1 - n^2},$$

where g^k is the Chebyshev polynomial representation of g in $(t - k\Delta t, t - (k - 1)\Delta t)$. We compute g^k by evaluating $g(\tau)$ at the Chebyshev nodes

$$\tau_i = t - k\Delta t + \frac{\Delta t}{2}(1 + \cos(\pi(i - 1)/q)), \quad i = 1, \dots, q. \tag{53}$$

At each τ_i , we use CFGT to optimally compute g .

6. Numerical results

In this section, we illustrate the performance of the algorithms presented by testing them on a set of synthetic problems. The criterion we used to construct the quadtree is spectral thresholding, as discussed in [11]: a leaf node is subdivided if the tail of Chebyshev coefficients it owns has not decayed sufficiently. Upon completion of the tree construction, every leaf node contains the coefficients $f_{m,n}$ such that

$$\sum_{k=0}^{q-1} |f_{k,q-k-1}| < \epsilon^*. \tag{54}$$

We implemented the algorithms presented here in MATLAB. For each test problem, we report CPU times w.r.t. a base case.

6.1. Chebyshev FGT

We present the convergence results for CFGT on two examples. In Table 1, we present the results for the Gauss transform of a function that is uniformly smooth. In Table 2, we consider the function shown in Fig. 5.

Free space initial condition problem. Consider the initial condition problem,

$$\frac{\partial u}{\partial t} = \Delta u \quad \text{in } \mathbf{R}^d, \quad \text{s.t. } u(\mathbf{x}, 0) = u_0(\mathbf{x}). \tag{55}$$

Table 1

CPU Time and relative errors in computing $G_\delta f(\mathbf{x}) = \int_{\Omega} e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{\delta}} f(\mathbf{y}) d\mathbf{y}$ with bandwidth $\delta = 0.1$; here $f(\mathbf{y}) = \sin^2(20y_1) \sin^2(20y_2)$

$N = M$	q		8		16	
	CPU time	Error	CPU time	Error	CPU time	Error
16	1	5.56e-02	6.67e-01	1.68e-04	1.00e+00	2.41e-08
64	1.67e+00	5.66e-03	1.67e+00	7.04e-06	2.00e+00	1.22e-12
256	2.33e+00	6.83e-05	3.00e+00	4.11e-09	4.33e+00	3.00e-15
1024	7.00e+00	3.60e-06	8.00e+00	1.28e-11	1.07e+01	3.74e-15
4096	2.20e+01	2.17e-07	2.57e+01	4.65e-14	3.43e+01	4.86e-15
16384	8.03e+01	1.34e-08	9.60e+01	4.67e-15	1.23e+02	4.48e-15

N is the number of leaf nodes, M is number of target points (which are chosen randomly in the unit box $[0, 1]^2$) and q is the order of Chebyshev polynomials used to approximate $f(\mathbf{y})$ in each leaf node. CPU times are relative to the base case $N = 16$ and $q = 4$.

Table 2

Performance of 16th-order CFGT in computing $G_\delta f(\mathbf{x})$ of the function shown in Fig. 5, for different values of δ

ϵ^*	$N = M$	δ					
		10^{-2}		10^{-4}		10^{-6}	
		CPU time	Error	CPU time	Error	CPU time	Error
10^{-6}	79	1	5.50e-012	2.35	3.26e-010	0.02	8.53e-010
10^{-8}	142	1.27	2.40e-014	3.99	4.69e-012	0.04	1.22e-010
10^{-10}	265	1.95	7.91e-016	5.99	2.07e-013	0.07	8.39e-014
10^{-12}	481	2.50	7.91e-016	9.53	2.11e-013	0.15	2.70e-016

The exact solution is analytically computed. Here ϵ^* is the discretization tolerance based on which the quadtree is constructed using the criterion (54). We can see that the convergence is independent of δ . The trend in the CPU times can be explained as follows: when δ decreases, the number of source/target boxes increase and that in turn increases the total time. As δ becomes really small, the direct evaluation version of CFGT is invoked and hence CPU time is reduced. CPU times are relative to the base case $\delta = 10^{-2}$ and $N = 79$.




Fig. 5. In this figure, we give an example of a tree-based representation of the function f shown at the left most figure. The quadtrees were constructed top-down, starting from the unit square, and refining using thresholding (54) with $\epsilon^* = 10^{-5}$. The leaf nodes of the quadtree shown in the middle use cubic polynomials, while the rightmost tree uses 15th-degree polynomials. While the cubic polynomial tree uses 40,513 leaf nodes, the 15th-order polynomial tree meets the required accuracy criterion with just 49 leaf nodes. Taking into account the storage per leaf node (q^2), the high-order requires 50 times less storage. For tighter discretization tolerances, this factor becomes even greater.

The exact solution is given by

$$u(\mathbf{x}, t) = \frac{1}{4\pi t} \int_{\Omega} e^{-\frac{|\mathbf{x}-\mathbf{y}|^2}{4t}} f(\mathbf{y}) d\mathbf{y}. \quad (56)$$

Now, assuming $\Omega = [0, 1]^2$ and u_0 be given in terms of piecewise Chebyshev polynomial coefficients, we can use CFGT for fast and high-order computation of (56). Here, we solve the free space initial condition problem (56) with complicated function u_0 , as shown in Fig. 6(a).

6.2. Chebyshev NUFFT

Here, we present the convergence results for computing the Fourier coefficients using CNUFFT. We test two functions: the first is discretized using a uniform tree, while the second is discretized using a nonuniform tree (see Tables 3 and 4).

6.3. Volume potential

Here we present two examples in which we test our method for the case of the heat potential evaluation. In the first example, we use a spatial grid that is regular at each time step. In the second example, the spatial grid is nonuniform and is different at each time step.

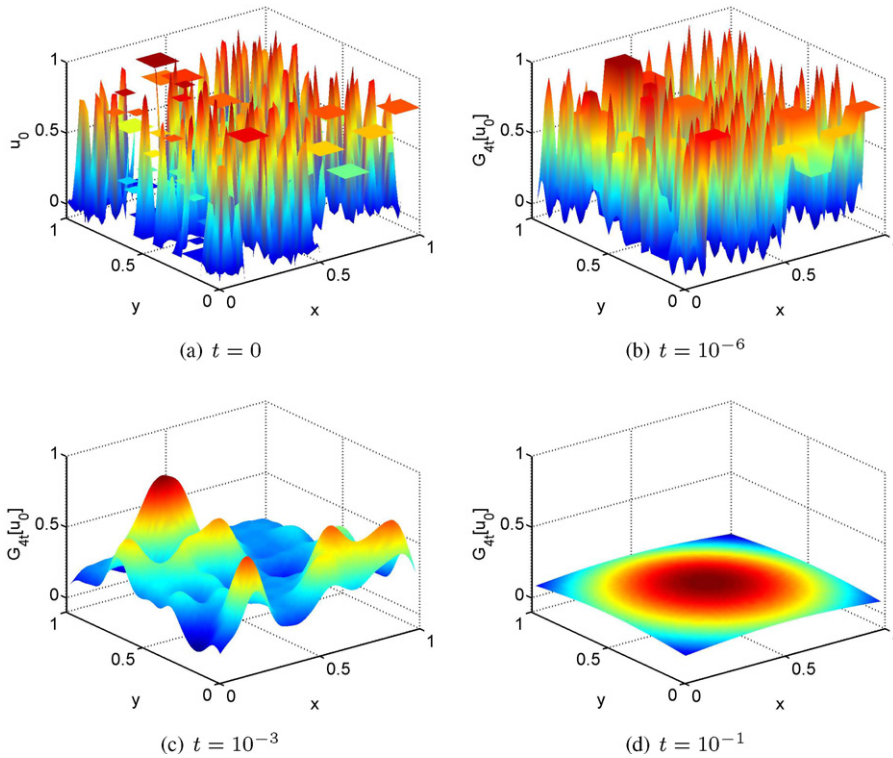


Fig. 6. Solution of the free space initial condition problem (55) at different times. The initial condition u_0 is discontinuous and oscillatory. The input is a quadtree that has 94 leaf nodes, each containing a 16th-order Chebyshev polynomial representation of u_0 .

Table 3
Results for CNUFFT of a non-periodic smooth function, $f(\mathbf{x}) = \sin(20x_1) \sin(20x_2)$ for $\mathbf{x} \in [0, 1]^2$

M	q		q		q	
	4	8	8	16	16	16
	CPU time	Error	CPU time	Error	CPU time	Error
64	1	2.91e-01	1.0	1.18e-05	2.0	1.52e-14
256	1.0	1.16e-02	2.0	2.20e-08	2.0	5.67e-14
1024	5.5	5.71e-04	7.0	2.06e-10	8.5	9.76e-14
4096	19.5	3.30e-05	21.5	9.41e-13	24.5	1.22e-13

The oversampling ratio was set to $R = 4$ and the spreading distance to $M_{sp} = 12$. CPU times are relative to $q = 4$ and $M = 64$.

Table 4
Performance of 8th- and 16th-order methods in computing the Fourier coefficients of the function given in Fig. 5

Discretization tolerance	q		q	
	8	16	16	16
	M	CPU time	M	CPU time
10^{-4}	244	5.7	37	1
10^{-6}	859	5.3	79	2.0
10^{-8}	3163	2.7e+01	142	2.3
10^{-10}	11380	1.8e+02	265	4.0
10^{-12}	42460	2.1e+03	481	4.9

Here M is the number of leaf nodes. CPU times are relative to the base case $q = 16$ and $\epsilon^* = 10^{-4}$.

Table 5
Absolute errors in computing the volume potential for Example 1

N	M	$q = 4$	$q = 8$	$q = 16$
8	4	5.83e+00	4.57e-02	7.42e-08
16	16	1.08e-01	4.03e-05	3.22e-13
32	64	1.72e-03	8.47e-08	1.89e-15
64	256	3.18e-06	3.52e-10	2.16e-15

The number of targets is M and they are randomly generated within $\omega = [0.4, 0.6]^2$. The total time is taken as $T = 1$. Here we report the discrete $L^\infty(\omega \times [0, T])$ norm of the error. In this example, we fixed the parameters to $\delta = 0.002$ and $p = 55$ so that truncation errors are within machine precision.




Fig. 7. Quadtree for $f(\mathbf{x}, t)$, defined in Example 2. The tree is generated using 8th-order Chebyshev polynomials, with discretization tolerance $\epsilon^* = 10^{-7}$. A contour plot of the solution $u(\mathbf{x}, t)$ is superimposed on the quadtrees. Evidently, the data structure changes at each step. We numerically compute $u(\mathbf{x}, t)$ at targets that are randomly located within $\omega = [0.4, 0.6]^2$ for each time step and compare with the exact solution in Table 6.

Example 1. Here, we evaluate the volume potential (3) of the function,

$$f(\mathbf{x}, t) = \sin(n\pi x_1) \sin(n\pi x_2) (n\pi \cos(n\pi t) + 2n^2 \pi^2 \sin(n\pi t)) \quad \text{with } n = 10. \quad (57)$$

The exact solution is given by

$$u(\mathbf{x}, t) = \sin(n\pi x_1) \sin(n\pi x_2) \sin(n\pi t), \quad (58)$$

which satisfies the initial and boundary conditions (2). Starting from a regular grid in space–time, we progressively refine it and the absolute errors for each discretization is reported in Table 5.

Example 2. In this example, we compute the volume potential due to a forcing function that requires time dependent spatially adaptive grid. Here $f(\mathbf{x}, t)$ is constructed by the exact solution given by⁹

$$u(\mathbf{x}, t) = \sum_{i=1}^3 t e^{-\alpha_i r_i^2} \quad r_i^2 = \|\mathbf{x} - \mathbf{c}_i(t)\|_2^2 \quad (59)$$

$$\mathbf{c}_i(t) = 0.5 + 0.05 \begin{bmatrix} \cos(2\pi(t + i/3)) \\ \sin(2\pi(t + i/3)) \end{bmatrix}, \quad \alpha_i = 400 + 100 \times 4^i.$$

The Gaussian centers $\mathbf{c}_i(t)$ move in a circular motion around the center of Ω . As the bandwidth of the Gaussians differ, the adaptive data structure is unsymmetric and hence it varies with time as shown in Fig. 7.

At each time step $[k\Delta t, (k+1)\Delta t]$, we construct the quadtree Q_k so that the polynomial coefficients of $f(\mathbf{x}, k\Delta t)$ in each leaf node obey (54).

⁹ Here $u(\mathbf{x}, t)$ satisfies the boundary conditions approximately.

7. Conclusions

We presented three transforms that can be applied to functions represented using piecewise Chebyshev polynomials: we have extended fast Gauss and nonuniform FFT transforms and, based on the fast algorithm of Greengard and Strain [13], we have presented a fast and high-order accurate algorithm for the computation of volume heat potential. We presented several numerical results in which the convergence rates of our 16th-order accurate scheme are verified.

We have set the initial condition in (2) to zero just for the ease of exposition. Our algorithm extends in an obvious way in the more general case. Given a piecewise Chebyshev polynomial representation of $u_0(\mathbf{y})$, we can use CFGT and CNUFFT for the fast and accurate evaluation of the volume potential that arises due to u_0 given by

$$\mathcal{V}[u_0](\mathbf{x}, t) = \int_{\Omega} G(\mathbf{x}, t; \mathbf{y})u_0(\mathbf{y})d\mathbf{y}.$$

Besides extending the method to the 3D case, there are two important directions that we have not explored: a posteriori error-based adaptive scheme; and the case of imposing boundary conditions on domains with complex geometries. Both directions are part of ongoing work.

Returning to the 3D case, the extension of the method to higher dimensions is straightforward using tensor products for the Gauss transform, and 3D FFTs. The complexity constants, however, increase exponentially (in the dimension) [15,16]. Optimizing and parallelizing the code becomes crucial for practical applications.

Acknowledgment

We would like to thank John Strain for bringing Ref. [26] to our attention.

Appendix A. Olver’s algorithm

Consider the computation of the moments

$$I_n(\alpha) = \int_0^1 e^{-\alpha x} T_n(2x - 1)dx, \quad 0 \leq n \leq q - 1. \tag{60}$$

We can use the recurrence (22) to compute I_n for $n > 2$ from base conditions I_0, I_1, I_2 . The recurrence (22) is an instance of second-order nonhomogeneous difference equation. Let us rewrite the recurrence as

$$I_n + a_n I_{n+1} + b_n I_{n+2} = f_n, \quad 1 \leq n \leq q - 1 \tag{61}$$

with the coefficients, a_n, b_n, f_n defined appropriately. Since this is a second-order equation, it requires two initial conditions¹⁰ I_1 and I_2 . The scheme of Olver can be written as follows:

$$\text{Solve : } I_n^N + a_n I_{n+1}^N + b_n I_{n+2}^N = f_n, \quad 1 < n \leq N. \tag{62}$$

$$\text{s.t. } I_{N+1}^N = 0, \quad I_1^N = I_1. \tag{63}$$

For the recurrence (22), we can prove that,

$$\lim_{N \rightarrow \infty} I_n^N = I_n, \quad n \geq 1. \tag{64}$$

When α is small, for which direct computation using recurrence (22) is unstable, we use Olver’s algorithm. For these values of α , the convergence (64) is quite rapid i.e., even for moderate values of N ,

$$\max_{1 < n \leq q-1} |I_n^N - I_n| < \epsilon. \tag{65}$$

¹⁰ When the recurrence (22) is written in the form (61), we see that $b_0 = 0$. Hence $n = 1$ pertains to the first initial condition, instead of $n = 0$. So, we require three base conditions for the recurrence (22).

Table 6
Performance of the 16th-order method on Example 2

ϵ^*	10^{-4}	10^{-6}	10^{-8}	10^{-10}
M_{\max}	91	136	277	415
N	8	16	32	64
Error	2.35e-06	1.00e-07	2.32e-11	1.78e-12

At each time step, the quadtree is constructed using the criterion $\sum_{k=0}^{q-1} |f_{k,q-k-1}| < \epsilon^* \|f(\mathbf{x}, k\Delta t)\|_{\infty}$ for $k = 0, 1, \dots, N - 1$. M_{\max} is the maximum number of leaf nodes over all the quadtrees at different times. Here, the targets are randomly located in $\omega = [0.4, 0.6]^2$ and the error is measured in discrete $L^{\infty}(\omega \times [0, T])$ norm.

Table 7
The threshold values $|\alpha|_{\text{th}}$ for different values of q

q	4	8	16
$ \alpha _{\text{th}}$	1/2	4	16

Whenever $\alpha \geq |\alpha|_{\text{th}}$, we compute $I_n(\alpha)$ using the recurrence (22). For $\alpha < |\alpha|_{\text{th}}$, we use Olver’s algorithm with $N = 2q + 1$.

More precisely, for double precision ($\epsilon = 10^{-12}$), we use Olver’s algorithm with $N = 2q + 1$ to compute (22), if $|\alpha| < |\alpha|_{\text{th}}$. For different values of q , we tabulate the threshold values in Table 7.

We solve (62) using LU decomposition. For smaller values of $|\alpha|$ than the threshold, we can achieve the required precision even with smaller N , but we do not carry the analysis any further.

Appendix B. Recurrences for the Gauss kernel

Defining $I_n(\lambda, \eta)$ as,

$$I_n(\lambda, \eta) = \frac{1}{\lambda} \int_{-1}^1 e^{-\left(\frac{y-\eta}{\lambda}\right)^2} T_n(y) dy, \quad 0 \leq n \leq q - 1, \tag{66}$$

we can derive the following recurrence, that is valid for $n > 3$:

$$I_n = 2\eta I_{n-1} + 2 \left((n-1)\lambda^2 + \frac{1}{n-3} \right) I_{n-2} - 2\eta \left(\frac{n-1}{n-3} \right) I_{n-3} + \frac{n-1}{n-3} I_{n-4} - C_n, \tag{67}$$

$$\text{where } C_n = \lambda(n-1) \left[e^{-\left(\frac{\eta-\eta}{\lambda}\right)^2} \left(\frac{T_{n-1}}{n-1} - \frac{T_{n-3}}{n-3} \right) \right]_{-1}^1. \tag{68}$$

The recurrence (67) is unstable for larger values of λ . Since, (67) is a fourth-order nonhomogeneous system, one can use Lozier’s algorithm [26] for stable computation for larger values of λ .

Since, however, we can precompute these moments $(I_n)^{11}$ by simply use a high-order smooth quadrature. Whenever $q > 4$ and $\lambda > \frac{1}{8}$, we use a 24th-order quadrature rule of [1], with q trapezoidal nodes to compute I_n for $4 < n < q$.

Appendix C. An extension of FGT

For some applications, one might be interested in computing the discrete Gauss transform (15) for relatively small values of δ . For instance, consider the free space initial condition problem (55) with the initial condition given by,

$$u_0(\mathbf{x}) = \sum_{j=1}^M q_j \delta(\mathbf{x} - \mathbf{y}_j). \tag{69}$$

¹¹ The direct evaluation version of CFGT does not precompute the moments. But, for the cases in which this algorithm is applicable, we can show that $\lambda \leq \frac{1}{8}$, so that the recurrence (67) can be used to compute the moments directly.

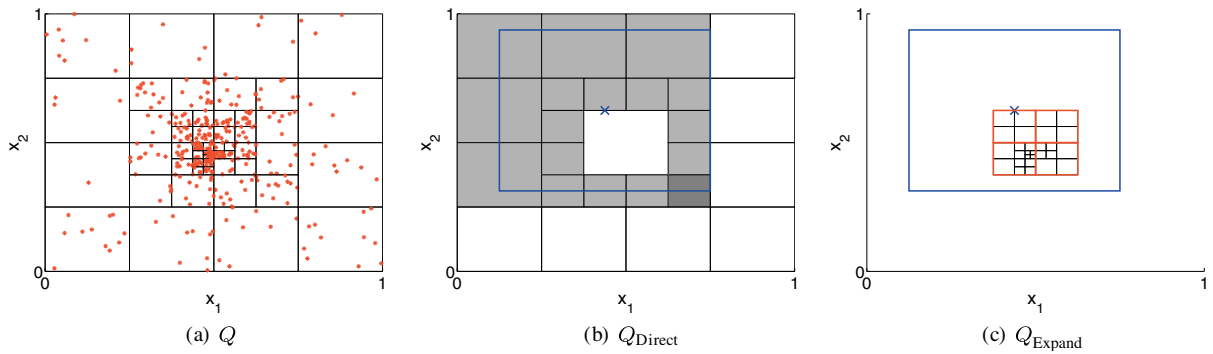


Fig. 8. Here we illustrate the extension of FGT for efficient computation of (70) at smaller times. (a) From the source distribution $\{y_j\}_{j=1}^M$ we construct a quadtree Q so that each of its leaf contains no more than s points. In this example, we have set $s = 12$. We split Q into two trees: Q_{Direct} has leaves that are larger than the source boxes and Q_{Expand} has smaller ones. (b) Here a target \mathbf{x} and the support of Gaussian centered at \mathbf{x} is shown. The leaf nodes that are shaded belong to the interaction list of \mathbf{x} . The sources that belong to these leaves interact directly. This would avoid the need to generate source boxes that are mostly empty. (c) The red boxes are the source boxes that are within the interaction list of the target box containing \mathbf{x} . The sources that belong to these boxes interact with \mathbf{x} through Hermite/Wave expansions. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

As this function cannot be represented using piecewise polynomials, CFGT is not applicable. Substituting u_0 in (56), we get

$$u(\mathbf{x}, t) = \frac{1}{4\pi t} \sum_{j=1}^M q_j e^{-\frac{|\mathbf{x}-y_j|^2}{4t}}. \quad (70)$$

The aim is to compute u at $\{\mathbf{x}_k\}_{k=1}^M$ and at any given time t . A direct application of FGT on (70) would be expensive at small values of t . This is because the number of boxes that FGT requires is inversely proportional to \sqrt{t} . We now describe an extension of the fast Gauss transform [14] that would enable the fast computation of (70) for any required t . The strategy is an extension of the idea used in CFGT to the discrete case. We explain this in Fig. 8.

References

- [1] Bradley K. Alpert, Hybrid Gauss-trapezoidal quadrature rules, *SIAM Journal on Scientific Computing* 20 (1999) 1551–1584.
- [2] Laurent Balmelli, J. Kovacevic, M. Vetterli, Quadtrees for embedded surface visualization: constraints and efficient data structures, pp. 487–491.
- [3] Gregory Beylkin, On the fast Fourier transform of a function with singularities, *Applied and Computational Harmonic Analysis* 2 (1995) 363–381.
- [4] Mark Broadie, Y. Yamamoto, Application of the fast Gauss transform to option pricing, *Management Science* 49 (2003) 1071–1088.
- [5] Mark Broadie, Y. Yamamoto, A double-exponential fast Gauss transform algorithm for pricing discrete path-dependent options, *Operations Research* 53 (2005) 764–779.
- [6] J. Crank, *Free and Moving Boundary Problems*, Oxford Science Publications, 1984.
- [7] Alok Dutt, V. Rokhlin, Fast Fourier transforms for nonequispaced data, *SIAM Journal on Scientific Computing* 14 (1993) 1368–1393.
- [8] Ahmed Elgammal, R. Duraiswami, L.S. Davis, Efficient kernel density estimation using the fast Gauss transform with applications to color modeling and tracking, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (2003) 1499–1504.
- [9] Guo-Xin Fan, Q.H. Liu, Fast Fourier transform for discontinuous functions, *IEEE Transactions on Antennas and Propagation* 52 (2004) 461–465.
- [10] S. Filippi, Angenaherte Tschebyscheff-approximation einer stammfunktion – eine modifikation des verfahrens von clenshaw and curtis, *Numerische Mathematik* 6 (1964).
- [11] Leslie Greengard, June-Yub Lee, A direct adaptive Poisson solver of arbitrary order accuracy, *Journal of Computational Physics* 125 (1996) 415–424.
- [12] Leslie Greengard, Patrick Lin, Spectral approximation of the free-space heat kernel, *Applied and Computational Harmonic Analysis* 9 (2000) 83–97.
- [13] Leslie Greengard, John Strain, A fast algorithm for the evaluation of heat potentials, *Communications on Pure and Applied Mathematics* XLIII (1990) 949–963.

- [14] Leslie Greengard, John Strain, The fast Gauss transform, *SIAM Journal on Scientific and Statistical Computing* 12 (1991) 79–94.
- [15] Leslie Greengard, Xiaobai Sun, A new version of the fast Gauss transform, *Documenta Mathematica* III (1998) 575–584.
- [16] June-Yub Lee, Leslie Greengard, Accelerating the nonuniform fast Fourier transform, *SIAM Review* 46 (2004) 443–454.
- [17] R. Piessens, M. Branders, Numerical-solution of integral-equations of mathematical physics using Tschebyscheff polynomials, *Journal of Computational Physics* 21 (1976) 178–196.
- [18] Theodore J. Rivlin, *Chebyshev Polynomials*, Wiley-Interscience, 1990.
- [19] J.A. Sethian, J. Strain, Crystal growth and dendritic solidification, *Journal of Computational Physics* 98 (1992).
- [20] Eugene Sorets, Fast Fourier transforms of piecewise constant functions, *Journal of Computational Physics* 116 (1995) 369–379.
- [21] John Strain, Fast potential theory. II. Layer potentials and discrete sums, *Journal of Computational Physics* 99 (1992) 251–270.
- [22] John Strain, Fast adaptive methods for the free-space heat equation, *SIAM Journal on Scientific Computing* 15 (1994) 185–206.
- [23] Xiaobai Sun, Yajuan Bao, A Kronecker product representation of the fast Gauss transform, *SIAM Journal on Matrix Analysis and Applications* 24 (2002) 768–786.
- [24] Shrawan K. Veerapaneni, George Biros, A high-order solver for the heat equation in 1D domains with moving boundaries, *SIAM Journal on Scientific Computing* 29 (2007) 2581–2606.
- [25] C.W. Clenshaw, A.R. Curtis, A method for numerical integration on an automatic computer, *Numerische Mathematik* 2 (1960) 197–205.
- [26] Jet Wimp, *Computation with Recurrence Relations*, Pitman Advanced Pub. Program, 1984.
- [27] Changjiang Yang, Ramani Duraiswami, Larry Davis, Efficient kernel machines using the improved fast Gauss transform, *Advances in Neural Information Processing Systems*, vol. 17, MIT Press, Cambridge, MA, 2005, pp. 1561–1568.
- [28] Changjiang Yang, Ramani Duraiswami, Nail A. Gumerov, Larry Davis, Improved fast Gauss transform and efficient kernel density estimation, in: *ICCV'03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, IEEE Computer Society, 2003, p. 464.