

Fast Algorithms for the Multi-dimensional Jacobi Polynomial Transform

James Bremer
Department of Mathematics
University of California, Davis, CA, USA
bremer@math.ucdavis.edu

Qiyuan Pang
Department of Mathematics
Purdue University, West Lafayette, IN, USA
qpang@purdue.edu

Haizhao Yang
Department of Mathematics
Purdue University, West Lafayette, IN, USA*
National University of Singapore, Singapore†
haizhao@nus.edu.sg

September 11, 2019

Abstract

We use the well-known observation that the solutions of Jacobi's differential equation can be represented via the non-oscillatory phase and amplitude functions to develop a fast algorithm for computing multi-dimensional Jacobi polynomial transforms. More explicitly, it follows from this observation that the matrix corresponding to the discrete Jacobi transform is the Hadamard product of a numerically low-rank matrix and a multi-dimensional discrete Fourier transform (DFT) matrix. The application of the Hadamard product can be carried out via r^d fast Fourier transforms (FFTs), where $r = O(\frac{\log n}{\log \log n})$ and d is the dimension, resulting in a nearly optimal algorithm to compute the multidimensional Jacobi polynomial transform.

Keywords. Multi-dimensional Jacobi polynomial transform, nonuniform transforms, non-oscillatory phase function, randomized low-rank approximation, fast Fourier transform.

1 Introduction

Jacobi polynomials contain a wide class of orthogonal polynomials, e.g., the Gegenbauer polynomials, and thus also the Legendre, Zernike, and Chebyshev polynomials, which have been studied and applied extensively in mathematical analysis and practical applications [1, 2, 3]. Previously, the Legendre and Chebyshev polynomials, not only in the univariate cases but also in the multivariate cases, have been essential tools in solving partial differential equations with spectral methods

*Current address.

†On leave from the National University of Singapore, where part of the work was done.

[4, 5, 6, 7, 8]. Fast Chebyshev polynomial transforms in 1D to 3D have been proposed in [9, 10, 11] to increase the numerical efficiency of these transforms when the problem size is large. Recently, Jacobi polynomials in spectral approximations were revisited and shown to be an efficient tool for problems with degenerated or singular coefficients [12], for singular differential equations [13, 14, 15], and for optimal error estimates for p -version of finite element methods [16, 17, 18]. The Jacobi polynomials were also generalized to a wider class of functions to simplify the approximation analysis with more precise error estimates and well-conditioned algorithms in [19]. Therefore, it is important to develop efficient computational tools for multivariate Jacobi polynomial transforms.

The one-dimensional forward discrete Jacobi transform consists of evaluating an expansion of the form

$$f(x) = \sum_{\nu=1}^n \alpha_{\nu} P_{\nu-1}^{(a,b)}(x), \quad (1)$$

where $P_{\nu}^{(a,b)}$ denotes the order- ν Jacobi polynomial of the first kind corresponding to parameters a and b , at a collection of points $X = \{x_i\}_{i=1,\dots,n} \subset (-1, 1)$. The one-dimensional inverse discrete Jacobi transform is the process of computing the coefficients $\{\alpha_{\nu}\}_{\nu=1,\dots,n}$ in an expansion of the form (1) given its values at a collection of distinct points X . For the sake of brevity, we will generally drop the adjective “discrete” and simply use the terms “forward Jacobi transform” and “inverse Jacobi transform” when referring to these operations. Often, the points X are the nodes of the n -point Gauss-Jacobi quadrature rule

$$\int_{-1}^1 p(x)(1+x)^a(1-x)^b dx \approx \sum_{\nu=1}^n p(x_{\nu})\omega_{\nu} \quad (2)$$

that is exact when p is a polynomial of degree less than or equal to $2n - 1$. However, it is useful to consider more general sets of points as well. In the case in which the points X are the nodes of the Gauss-Jacobi quadrature rule, we say the corresponding transform and inverse transform are uniform; otherwise, we describe them as nonuniform. To form an orthonormal matrix to represent the transform for numerical purpose, we will rescale the transform in (1) with quadrature weights later.

Many methods for rapidly applying the Jacobi transform and various special cases of the Jacobi transform have been developed. Examples include the algorithms of [20, 21, 22, 23] for applying the Legendre transform, and those of [24] for forming and evaluating expansions for Jacobi polynomials in general. See also, [9] and its references for extensive information on numerical algorithms for forming and manipulating Chebyshev expansions. Almost all such algorithms can be placed into one of two categories. Algorithms in the first category, such as [25, 26, 27, 28, 2, 29, 30, 31, 39], make use of the structure of the connection matrices which take the coefficients in the expansion of a function in terms of one set of Jacobi polynomials to the coefficients in the expansion of the same function for a different set of Jacobi polynomials. They typically operate by computing the Chebyshev coefficients of expansion and then applying a connection matrix or a series of connection matrices to obtain the coefficients in the desired expansion. The computation of the Chebyshev expansion can be carried out efficiently in $O(n \log n)$ operations via the nonuniform FFT (NUFFT) [32, 33] and the application of each connection matrix can be performed in a number of operations which grows linearly or quasi-linearly via the fast multipole method (FMM) [22, 34, 35, 24] or the fast eigendecomposition of semiseparable matrices [36, 37, 24].

The second class of algorithms, of which [38] is a major example, make use of butterfly methods and similar techniques for the application of oscillatory matrices. Most algorithms in this category, including that of in [38], require precomputation with $O(n^2)$ running times. However, [40]

describes an algorithm based on this approach whose total running time is $O(n \log^2(n))$, when both parameters a and b are in the interval $(-1/2, 1/2)$. It makes use of the observation that the solutions of Jacobi's differential equation can be accurately represented via non-oscillatory phase and amplitude functions to apply the Jacobi transform rapidly. The fact that certain differential equations admit non-oscillatory phase functions has long been known [41]; a general theorem was established in [42] and a numerical method for the computation of non-oscillatory phase functions was given in [43]. In the case of Jacobi's differential equation with $t \in [0, \pi]$ and $\nu > 0$, there exists a smooth amplitude function $M^{(a,b)}(t, \nu)$ and a smooth phase function $\psi^{(a,b)}(t, \nu)$ such that

$$\tilde{P}_\nu^{(a,b)}(t) = M^{(a,b)}(t, \nu) \cos(\psi^{(a,b)}(t, \nu)) \quad (3)$$

and

$$\tilde{Q}_\nu^{(a,b)}(t) = M^{(a,b)}(t, \nu) \sin(\psi^{(a,b)}(t, \nu)), \quad (4)$$

where $\tilde{P}_\nu^{(a,b)}$ and $\tilde{Q}_\nu^{(a,b)}$ are referred to as the modified Jacobi functions of the first and second kind, respectively. They are defined in terms of the Jacobi functions of the first and second kinds $P_\nu^{(a,b)}$ and $Q_\nu^{(a,b)}$ (see [41] for definitions) via the formulas

$$\tilde{P}_\nu^{(a,b)}(t) = C_\nu^{(a,b)} P_\nu^{(a,b)}(\cos(t)) \sin\left(\frac{t}{2}\right)^{a+\frac{1}{2}} \cos\left(\frac{t}{2}\right)^{b+\frac{1}{2}} \quad (5)$$

and

$$\tilde{Q}_\nu^{(a,b)}(t) = C_\nu^{(a,b)} Q_\nu^{(a,b)}(\cos(t)) \sin\left(\frac{t}{2}\right)^{a+\frac{1}{2}} \cos\left(\frac{t}{2}\right)^{b+\frac{1}{2}}, \quad (6)$$

where

$$C_\nu^{(a,b)} = \sqrt{(2\nu + a + b + 1) \frac{\Gamma(1 + \nu)\Gamma(1 + \nu + a + b)}{\Gamma(1 + \nu + a)\Gamma(1 + \nu + b)}}. \quad (7)$$

The normalization constant is chosen to ensure the $L^2(0, \pi)$ norm of $\tilde{P}_\nu^{(a,b)}$ is 1 when ν is an integer. Indeed, the set $\{\tilde{P}_j^{(a,b)}\}_{j=0}^\infty$ is an orthonormal basis for $L^2(0, \pi)$. The change of variables $x = \cos(t)$ makes the singularities in phase and amplitude functions for Jacobi's differential equation more tractable. Obviously, there is no substantial difference in treating expansions of the type (1) and those of the form

$$f(t) = \sum_{\nu=1}^n \alpha_\nu \tilde{P}_{\nu-1}^{(a,b)}(t) = \sum_{\nu=1}^n \alpha_\nu M^{(a,b)}(t, \nu) \cos(\psi^{(a,b)}(t, \nu)), \quad (8)$$

and we will consider the latter form here, more specially, the latter form with scaling by quadrature weights. Moreover, we will denote by $\{t_k\}_{k=1, \dots, n}$ and $\{w_k\}_{k=1, \dots, n}$ the nodes and weights of the trigonometric Gauss-Jacobi quadrature rule

$$\int_0^\pi f(\cos(t)) \cos^{2a+1}\left(\frac{t}{2}\right) \sin^{2b+1}\left(\frac{t}{2}\right) dt \approx \sum_{k=1}^n f(\cos(t_k)) \cos^{2a+1}\left(\frac{t_k}{2}\right) \sin^{2b+1}\left(\frac{t_k}{2}\right) w_k \quad (9)$$

obtained by applying the change of variables $x = \cos(t)$ to (2). Again, we refer to transforms which use the set of points $\{t_k\}_{k=1, \dots, n}$ as uniform and those which use some other point sets as nonuniform.

In [40], it was observed that the second summation in (8) could be interpreted as the application of a structured matrix that is the real part of a Hadamard product of a NUFFT matrix and a

numerically low-rank matrix. This leads to a method for its computation which takes quasi-linear time [33, 44, 45]. Expansions in terms of the functions of the second kind $\{\tilde{Q}_j^{(a,b)}\}_{j=0}^{\infty}$ can be handled similarly. This non-oscillatory representation method will be described in more detail in Section 2.2.

In this paper, we generalize the one-dimensional Jacobi polynomial transform of [40] to the multi-dimensional case. The transformation matrix in the multi-dimensional case is still the real part of a Hadamard product of a multi-dimensional NUFFT matrix and a numerically low-rank matrix in the form of tensor products. However, the numerical rank of the low-rank matrix might increase quickly in the dimension following the algorithm in [40] and hence the multi-dimensional NUFFT in [32, 33] might not be sufficiently efficient. To obtain a viable multi-dimensional Jacobi polynomial transform, we reformulate the Hadamard product into a new one, which is the Hadamard product of an over-sampled Fourier transform matrix and a numerically low-rank matrix, and propose a fast randomized SVD to achieve near optimality in the low-rank approximation. This leads to an efficient implementation of two-dimensional and three-dimensional Jacobi polynomial transforms. Besides, as will be demonstrated numerically, the new method proposed in this paper is faster and more robust than the method in [40] and works for a and b in the regime $(-1, 1)$, which is a larger interval than the regime $(-\frac{1}{2}, \frac{1}{2})$ in [40].

The remainder of the paper is organized as follows. In Section 2, we will first briefly introduce the fast SVD via randomized sampling and the theory of non-oscillatory phase functions. In Section 3, a variant of the fast one-dimensional Jacobi polynomial transform in [40] is given; it is followed by a description of our fast multi-dimensional Jacobi polynomial transform. In Section 4, numerical results which demonstrate the efficiency of our algorithm are described. Finally, we will conclude our discussion in Section 5.

2 Preliminaries

In this section, we will revisit the linear scaling randomized SVD introduced in [46] and the non-oscillatory phase and amplitude functions in [40] to make the presentation self-contained.

2.1 Approximate SVD via Randomized Sampling

For a numerically low-rank matrix $Z \in \mathbb{C}^{n \times n}$ with an $O(1)$ algorithm to evaluate an arbitrary entry, [46] introduced an $O(nr^2)$ algorithm to construct a rank- r approximate SVD, $Z \approx U_0 \Sigma_0 V_0^*$, from $O(r)$ randomly selected rows and columns of Z , where $*$ denotes the conjugate transpose, U_0 and $V_0 \in \mathbb{C}^{n \times r}$, $\Sigma_0 \in \mathbb{C}^{r \times r}$ is a diagonal matrix with approximate singular values in the diagonal part in a descent order.

Here, we adopt the standard notation for a submatrix in MATLAB: given a row index set I and a column index set J , $Z_{I,J} = Z(I, J)$ is the submatrix with entries from rows in I and columns in J ; we also use “:” to denote the entire columns or rows of the matrix, i.e., $Z_{I,:} = Z(I, :)$ and $Z_{:,J} = Z(:, J)$. With these handy notations, we briefly introduce the randomized SVD as follows.

Algorithm 2.1. *Approximate SVD via randomized sampling*

1. Let Π_{col} and Π_{row} denote the important columns and rows of Z that are used to form the column and row bases. Initially $\Pi_{col} = \emptyset$ and $\Pi_{row} = \emptyset$.
2. Randomly sample rq rows and denote their indices by S_{row} . Let $I = S_{row} \cup \Pi_{row}$. Here $q = O(1)$ is a multiplicative oversampling parameter. Perform a pivoted QR decomposition

of $Z_{I,:}$ to get

$$Z_{I,:}P = QR, \quad (10)$$

where P is the resulting permutation matrix and $R = (r_{ij})$ is an $O(r) \times n$ upper triangular matrix. Define the important column index set Π_{col} to be the first r columns picked within the pivoted QR decomposition.

3. Randomly sample rq columns and denote their indices by S_{col} . Let $J = S_{col} \cup \Pi_{col}$. Perform a pivoted LQ decomposition of $Z_{:,J}$ to get

$$PZ_{:,J} = LQ, \quad (11)$$

where P is the resulting permutation matrix and $L = (l_{ij})$ is an $m \times O(r)$ lower triangular matrix. Define the important row index set Π_{row} to be the first r rows picked within the pivoted LQ decomposition.

4. Repeat steps 2 and 3 a few times to ensure Π_{col} and Π_{row} sufficiently sample the important columns and rows of Z .
5. Apply the pivoted QR factorization to $Z_{:\Pi_{col}}$ and let Q_{col} be the matrix of the first r columns of the Q matrix. Similarly, apply the pivoted QR factorization to $Z_{\Pi_{row},:}^*$ and let Q_{row} be the matrix of the first r columns of the Q matrix.
6. We seek a middle matrix M such that $Z \approx Q_{col}MQ_{row}^*$. To solve this problem efficiently, we approximately reduce it to a least-squares problem of a smaller size. Let S_{col} and S_{row} be the index sets of a few extra randomly sampled columns and rows. Let $J = \Pi_{col} \cup S_{col}$ and $I = \Pi_{row} \cup S_{row}$. A simple least-squares solution to the problem

$$\min_M \| Z_{I,J} - (Q_{col})_{I,:}M(Q_{row}^*)_{:,J} \|, \quad (12)$$

gives $M = (Q_{col})_{I,:}^\dagger Z_{I,J} (Q_{row}^*)_{:,J}^\dagger$, where $(\cdot)^\dagger$ stands for the pseudo-inverse.

7. Compute an SVD $M \approx U_M \Sigma_M V_M^*$. Then the low-rank approximation of $Z \approx U_0 S_0 V_0^*$ is given by

$$U_0 = Q_{col}U_M; \Sigma_0 = \Sigma_M; V_0^* = V_M^*Q_{row}^*. \quad (13)$$

In our numerical implementation, iterating Steps 2 and 3 twice is empirically sufficient to achieve accurate low-rank approximations via Algorithm 2.1. Similar arguments as in [47] for a randomized CUR factorization can be applied to quantify the error and success probability rigorously for Algorithm 2.1. But at this point, we only focus on the application of Algorithm 2.1 to the fast Jacobi polynomial transform without theoretical analysis.

Note that our goal in the Jacobi polynomial transform is to construct a low-rank approximation of the low-rank matrix, i.e., $Z \approx UV^*$ with U and $V \in \mathbb{C}^{n \times r}$, up to a fixed relative error ε , rather than a fixed rank. Algorithm 2.1 can also be embedded into an iterative process that gradually increases the rank parameter r to achieve the desired accuracy. We denote the rank parameter as r_ε when it is large enough to obtain the accuracy ε .

When Z originates from the discretization of a smooth function $Z(x, y)$ at the grid points $\{x_i\}_{1 \leq i \leq n}$ and $\{y_i\}_{1 \leq i \leq n}$, i.e., $Z_{ij} = Z(x_i, y_j)$, another standard method for constructing a low-rank factorization of $Z \approx UV^*$ is Lagrange interpolation at Chebyshev grids in x or y . For example, let

$C_\mu := \{\mu_i\}_{1 \leq i \leq r}$ denote the set of r Chebyshev grids in the domain of x , $U \in \mathbb{R}^{n \times r}$ be the matrix consisting of the i -th Lagrange polynomial in x corresponding to μ_i as its i -th column, and $V \in \mathbb{C}^{n \times r}$ be the matrix such that V_{ij} is equal to $\bar{Z}(\mu_j, y_i)$, where $\bar{\cdot}$ denotes the conjugate operator, then $Z \approx UV^*$ by the Lagrange interpolation. Usually, an oversampling parameter q is used via setting $r = qr_\epsilon$. Then a rank- r_ϵ truncated SVD of $U \approx U_0 \Sigma_0 V_0^*$ gives a compressed rank- r_ϵ approximation of $Z \approx U_0 (VV_0 \Sigma_0)^*$, where $U_0 \in \mathbb{C}^{n \times r_\epsilon}$ and $VV_0 \Sigma_0 \in \mathbb{C}^{n \times r_\epsilon}$.

The fast nonuniform FFT in [33] and the one-dimensional Jacobi polynomial transform in [40] adopted low-rank approximation via Lagrange interpolation to deal with the low-rank term in their Hadamard products of low-rank and (nonuniform) FFT matrices without an extra truncated SVD. In this paper, we propose to use the randomized SVD via random sampling to obtain nearly optimal rank in the low-rank approximation.

2.2 Non-oscillatory phase and amplitude functions

Given a pair of parameters a and b in $(-\frac{1}{2}, \frac{1}{2})$, and a maximum degree of interest $N_{\max} > 27$, we revisit the fast algorithms in [40] for constructing non-oscillatory phase and amplitude functions $\psi^{(a,b)}(t, \nu)$ and $M^{(a,b)}(t, \nu)$ such that

$$\tilde{P}_\nu^{(a,b)}(t) = M^{(a,b)}(t, \nu) \cos\left(\psi^{(a,b)}(t, \nu)\right) \quad (14)$$

and

$$\tilde{Q}_\nu^{(a,b)}(t) = M^{(a,b)}(t, \nu) \sin\left(\psi^{(a,b)}(t, \nu)\right) \quad (15)$$

for $t \in [\frac{1}{N_{\max}}, \pi - \frac{1}{N_{\max}}]$ and $\nu \in (27, N_{\max}]$. The polynomials with t and ν out of these ranges can be evaluated by the well-known three-term recurrence relations or various asymptotic expansions; the lower bound of 27 was chosen to obtain optimal numerical performance.

Next, we present a few facts regarding the phase and amplitude functions related to Jacobi's differential equations. It is well known that the functions $\tilde{P}_\nu^{(a,b)}$ and $\tilde{Q}_\nu^{(a,b)}$ satisfy the second order differential equation

$$y''(t) + q_\nu^{(a,b)}(t)y(t) = 0, \quad (16)$$

where

$$q_\nu^{(\alpha,\beta)}(t) = \left(\nu + \frac{\alpha + \beta + 1}{2}\right)^2 + \frac{\frac{1}{4} - \alpha^2}{4 \sin^2(\frac{t}{2})} + \frac{\frac{1}{4} - \beta^2}{4 \cos^2(\frac{t}{2})}. \quad (17)$$

We refer to Equation (16) as Jacobi differential equation. Following the derivation in [40], we can show that the pair $\{\tilde{P}_\nu^{(a,b)}, \tilde{Q}_\nu^{(a,b)}\}$ of real-valued, linearly independent solutions of (16) satisfies

$$\tilde{P}_\nu^{(a,b)}(t) = \sqrt{W} \frac{\cos(\psi^{(a,b)}(t, \nu))}{\sqrt{|\partial_t \psi^{(a,b)}(t, \nu)|}}, \quad (18)$$

and

$$\tilde{Q}_\nu^{(a,b)}(t) = \sqrt{W} \frac{\sin(\psi^{(a,b)}(t, \nu))}{\sqrt{|\partial_t \psi^{(a,b)}(t, \nu)|}}, \quad (19)$$

where W is the necessarily positive constant Wronskian of the pair $\{\tilde{P}_\nu^{(a,b)}, \tilde{Q}_\nu^{(a,b)}\}$, and $\psi^{(a,b)}(t, \nu)$ is the non-oscillatory phase function of the pair $\{\tilde{P}_\nu^{(a,b)}, \tilde{Q}_\nu^{(a,b)}\}$ satisfying

$$\psi^{(a,b)}(t, \nu) = C + \int_{\sigma_1}^t \frac{W}{(\tilde{P}_\nu^{(a,b)}(s))^2 + (\tilde{Q}_\nu^{(a,b)}(s))^2} ds \quad (20)$$

with C an appropriately chosen constant related to σ_1 . Note that $\partial_t \psi^{(a,b)}(t, \nu) > 0$ since $W > 0$. Hence, the non-oscillatory amplitude function of $\{\tilde{P}_\nu^{(a,b)}, \tilde{Q}_\nu^{(a,b)}\}$ can be defined as

$$M^{(a,b)}(t, \nu) = \sqrt{\frac{W}{|\partial_t \psi^{(a,b)}(t, \nu)|}} = \sqrt{\frac{W}{\partial_t \psi^{(a,b)}(t, \nu)}}. \quad (21)$$

Through straightforward computation, it can be verified that the square $N^{(a,b)}(t, \nu) = (M^{(a,b)}(t, \nu))^2$ of the amplitude function satisfies the third order linear ordinary differential equation (ODE)

$$\partial_{ttt} N^{(a,b)}(t, \nu) + 4q_\nu^{(a,b)}(t) \partial_t N^{(a,b)}(t, \nu) + 2(q_\nu^{(a,b)}(t))' N^{(a,b)}(t, \nu) = 0 \quad \text{for all } \sigma_1 < t < \sigma_2. \quad (22)$$

Obviously,

$$\left(M^{(a,b)}(t, \nu)\right)^2 = \left(P_\nu^{(a,b)}(t)\right)^2 + \left(Q_\nu^{(a,b)}(t)\right)^2. \quad (23)$$

Hence, to set up the initial condition for the ODE in (23), we can specify the values of $(M^{(a,b)}(t, \nu))^2$ and its first two derivatives in t at a point on the interval $[\frac{1}{N_{\max}}, \pi - \frac{1}{N_{\max}}]$ for any ν using various asymptotic expansion for $P_\gamma^{(a,b)}$ and $Q_\gamma^{(a,b)}$. Afterwards, we uniquely determine $M^{(a,b)}(t, \nu)$ via solving the ODE (22) using a variant of the integral equation method of [32] (or any standard method for stiff problems).

To obtain the values of $\psi^{(a,b)}(t, \nu)$, we first calculate the values of

$$\frac{d}{dt} \psi^{(a,b)}(t, \nu) \quad (24)$$

via (21). Next, we obtain the values of the function $\tilde{\psi}$ defined via

$$\tilde{\psi}(t) = \int_{\alpha_1}^t \frac{d}{ds} \psi^{(a,b)}(s, \nu) ds \quad (25)$$

at any point in $[\frac{1}{N_{\max}}, \pi - \frac{1}{N_{\max}}]$ via spectral integration. There is an unknown constant connecting $\tilde{\psi}$ with $\psi^{(a,b)}(t, \nu)$; that is,

$$\psi^{(a,b)}(t, \nu) = \tilde{\psi}(t) + C. \quad (26)$$

To evaluate C , we first use a combination of asymptotic and series expansions to calculate $\tilde{P}_\nu^{(a,b)}$ at the point α_1 . Since $\tilde{\psi}(\alpha_1) = 0$, it follows that

$$\tilde{P}_\gamma^{(a,b)}(\alpha_1) = M^{(a,b)}(\alpha_1, \gamma) \cos(C), \quad (27)$$

and C can be calculated in the obvious fashion.

The above discussion has introduced an algorithm to evaluate $\psi^{(a,b)}(t, \nu)$ and $M^{(a,b)}(t, \nu)$ in the whole domain $(t, \nu) \in [\frac{1}{N_{\max}}, \pi - \frac{1}{N_{\max}}] \times (27, N_{\max}]$. To achieve a fast algorithm, we note that it is enough to conduct calculation on selected important grid points of $[\frac{1}{N_{\max}}, \pi - \frac{1}{N_{\max}}] \times (27, N_{\max}]$ through the above calculation, since we can evaluate $\psi^{(a,b)}(t, \nu)$ and $M^{(a,b)}(t, \nu)$ via interpolation with their values on the important grid points.

It is well known that a polynomial $p(t)$ (here $p(t)$ is either $\tilde{P}_\nu^{(a,b)}(t)$ or $\tilde{Q}_\nu^{(a,b)}(t)$) can be evaluated in a numerically stable fashion using the barycentric Chebyshev interpolation formula [48]

$$p(t) = \sum_{j=1}^k \frac{w_j}{t - x_j} p(x_j) \Big/ \sum_{j=1}^k \frac{w_j}{t - x_j}, \quad (28)$$

where x_1, \dots, x_k are the nodes of the k -point Chebyshev grid on a sufficiently small interval (σ_1, σ_2) (such that $p(t)$ is a polynomial of degree at most $k - 1$) and

$$w_j = \begin{cases} (-1)^j, & 1 < j < k; \\ (-1)^j \frac{1}{2}, & \text{otherwise.} \end{cases} \quad (29)$$

Hence, it is sufficient to evaluate $\psi^{(a,b)}(t, \nu)$ and $M^{(a,b)}(t, \nu)$, which lead to $\tilde{P}_\nu^{(a,b)}(t)$ and $\tilde{Q}_\nu^{(a,b)}(t)$, at a tensor product of Chebyshev grid points in t and ν , and calculate them at an arbitrary location (t, ν) via a bivariate piecewise barycentric Chebyshev interpolation.

More explicitly, let m_ν be the least integer such that $3^{m_\nu} \geq N_{\max}$, and let m_t be equal to twice the least integer l such that $\frac{\pi}{2} 2^{-l+1} \leq \frac{1}{N_{\max}}$. Next, we define $\beta_j = \max\{3^{j+2}, N_{\max}\}$ for $j = 1, \dots, m_\nu$, $\alpha_i = \frac{\pi}{2} 2^{i-m_t/2}$ for $i = 1, \dots, m_t/2$, and $\alpha_i = \pi - \frac{\pi}{2} 2^{m_t/2+1-i}$ for $i = m_t/2 + 1, \dots, m_t$. Now we let

$$\tau_1, \dots, \tau_{M_t} \quad (30)$$

denote the 16-point piecewise Chebyshev grid on the intervals

$$(\alpha_1, \alpha_2), (\alpha_2, \alpha_3), \dots, (\alpha_{m_t-1}, \alpha_{m_t}) \quad (31)$$

with $M_t = 15(m_t - 1) + 1$ points in total; let

$$\gamma_1, \dots, \gamma_{M_\nu} \quad (32)$$

denote the nodes of the 24-point piecewise Chebyshev grid on the intervals

$$(\beta_1, \beta_2), (\beta_2, \beta_3), \dots, (\beta_{m_\nu-1}, \beta_{m_\nu}) \quad (33)$$

with $M_\nu = 23(m_\nu - 1) + 1$ points in total. The piecewise Chebyshev grids (30) and (32) form a tensor product

$$\{(\tau_i, \gamma_j) : i = 1, \dots, M_t, j = 1, \dots, M_\nu\}. \quad (34)$$

Then $\psi^{(a,b)}(t, \nu)$ and $M^{(a,b)}(t, \nu)$ can be obtained via barycentric Chebyshev interpolations in t and ν via their values at the tensor product restricted to the piece $(\alpha_i, \alpha_{i+1}) \times (\beta_j, \beta_{j+1})$ that (t, ν) belongs to for some i and j .

Here we summarize the operation complexity of the whole algorithm above. A detailed discussion can be found in [40]. ODE solvers are applied for $\mathcal{O}(\log(N_{\max}))$ values of ν , and the solution of the solver is evaluated at $\mathcal{O}(\log(N_{\max}))$ values of t , making the total running time of the procedure just described $\mathcal{O}(\log^2(N_{\max}))$.

Once the values of $\psi_\nu^{(a,b)}$ and $M_\nu^{(a,b)}$ are ready at the tensor product of the piecewise Chebyshev grids (30) and (32), they can be evaluated for any t and ν via repeated application of the barycentric Chebyshev interpolation formula in the same number of operations which is independent of ν and t .

3 Fast Jacobi polynomial transforms

Without loss of generality, we assume that $\psi^{(a,b)}(t, \nu)$, $M^{(a,b)}(t, \nu)$, $P_\nu^{(a,b)}$, and $Q_\nu^{(a,b)}$ can be evaluated in $O(1)$ operations for any $t \in (0, \pi)$ and $\nu \in [0, N_{\max}]$. A fast algorithm for rapidly computing Gauss-Jacobi quadrature nodes and weights in the Jacobi polynomial transform in (8) has also been introduced in [40]. We refer the reader to [40] for details and assume that quadrature nodes and weights are available in this section.

3.1 One-dimensional transform and its inverse

Here we propose a new variant of the one-dimensional Jacobi polynomial transform for $\tilde{P}_j^{(a,b)}(t)$ in [40]. The new algorithm simplifies the discussion of the fast algorithm. The transform for $\tilde{Q}_j^{(a,b)}(t)$ is similar. For our purpose, we consider the n^{th} order uniform forward Jacobi polynomial transform with a scaling, that is, calculate the vector of values

$$\begin{pmatrix} f(t_1)\sqrt{w_1} \\ f(t_2)\sqrt{w_2} \\ \vdots \\ f(t_n)\sqrt{w_n} \end{pmatrix} \quad (35)$$

given the vector

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{pmatrix} \quad (36)$$

of the coefficients in the expansion

$$f(t) = \sum_{j=0}^{n-1} \alpha_j \tilde{P}_j^{(a,b)}(t); \quad (37)$$

here, $t_1, \dots, t_n, w_1, \dots, w_n$ are the nodes and weights of the n -point trigonometric Gauss-Jacobi quadrature rule corresponding to the parameters a and b . The “nonuniform” forward Jacobi transform does not require trigonometric Gauss-Jacobi quadrature nodes and weights.

In the uniform transform, the properties of the trigonometric Gauss-Jacobi quadrature rule and the weighting by square roots in (35) ensure that the $n \times n$ matrix $W\mathcal{J}_n^{(a,b)}$ taking (36) to (35) is orthogonal, where W is the $n \times n$ matrix

$$W = \begin{pmatrix} \sqrt{w_1} & 0 & 0 & 0 \\ 0 & \sqrt{w_2} & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \sqrt{w_n} \end{pmatrix}, \quad (38)$$

and $\mathcal{J}_n^{(a,b)}$ is the $n \times n$ matrix representing the sum in (37). Hence, the inverse transform is a simple matvec by the transpose of $W\mathcal{J}_n^{(a,b)}$. In the nonuniform case, $\mathcal{J}_n^{(a,b)}$ is usually a very ill-conditioned matrix and the inverse transform requires solving a challenging system of linear equations, which will be discussed in a separate paper in the future.

It follows from (14) that the (j, k) -th entry of $\mathcal{J}_n^{(a,b)}$ is

$$\left(\mathcal{J}_n^{(a,b)}\right)_{jk} = \begin{cases} M^{(a,b)}(t_j, k-1) \cos(\psi^{(a,b)}(t_j, k-1)), & \text{if } k-1 > 27, \\ \tilde{P}_{k-1}^{(a,b)}(t_j), & \text{otherwise,} \end{cases} \quad (39)$$

since the part consisting of low-degree polynomials is computed via three-term recurrence relations or various asymptotic expansions. Let $\mathcal{V}_n^{(a,b)} \in \mathbb{R}^{n \times 27}$ be the submatrix of $\mathcal{J}_n^{(a,b)}$ corresponding to the part of polynomials with degree less than 27, and $\mathcal{G}_n^{(a,b)} \in \mathbb{R}^{n \times n-27}$ be the rest of $\mathcal{J}_n^{(a,b)}$, then

$$\mathcal{J}_n^{(a,b)} = \left[\mathcal{V}_n^{(a,b)}, \mathcal{G}_n^{(a,b)} \right].$$

Note that $\mathcal{G}_n^{(a,b)}$ is the real part of a discrete Fourier integral transform that can be evaluated with a quasi-linear complexity by the algorithm of [44] for applying Fourier integral transforms and a special case of that in [45] for general oscillatory integral transforms. Motivated by this observation, the algorithm in [40] constructs a low-rank matrix $\mathcal{A}_n^{(a,b)} \in \mathbb{C}^{n \times (n-27)}$ whose (j, k) -th entry is

$$M^{(a,b)}(t_j, k+27) \exp\left(\mathbf{i}\left(\psi^{(a,b)}(t_j, k+27) - (k+27)t_j\right)\right). \quad (40)$$

Denote the $n \times (n-27)$ nonuniform FFT matrix, whose (j, k) -th entry is

$$\exp(i(k+27)t_j), \quad (41)$$

as \mathcal{N}_n . Then

$$\mathcal{J}_n^{(a,b)} = \left[\mathcal{V}_n^{(a,b)}, \Re(\mathcal{A}_n^{(a,b)} \otimes \mathcal{N}_n) \right], \quad (42)$$

where “ \otimes ” denotes the Hadamard product. It can be immediately seen from (42) that $\mathcal{J}_n^{(a,b)}$ can be applied through a small number of NUFFTs [32, 33] and a simple direct summation for $\mathcal{V}_n^{(a,b)}$ in a quasi-linear scaling of n .

More specifically, let a low-rank approximation of $\mathcal{A}_n^{(a,b)}$ be

$$\mathcal{A}_n^{(a,b)} \approx \sum_{j=1}^r u_j v_j^T, \quad (43)$$

where u_j and v_j are column vectors for each j . Then the transform from (36) to (35) can be approximated by the following sum,

$$W \mathcal{J}_n^{(a,b)} \alpha \approx W \mathcal{V}_n^{(a,b)} \alpha_{\mathcal{V}} + W \Re\left(\sum_{j=1}^r D_{u_j} \mathcal{N}_n D_{v_j} \alpha_{\mathcal{G}}\right), \quad (44)$$

where $\alpha_{\mathcal{V}}$ is the subvector of α with the first 27 entries, $\alpha_{\mathcal{G}}$ is the subvector of α with the rest entries, D_u denotes a diagonal matrix with a column vector u on its diagonal. The formula (44) can be carried out by r NUFFTs in $O(rn \log n)$ arithmetic operations. Generally speaking, r is related to n and it was conjectured in [40] that

$$r = O\left(\frac{\log n}{\log \log n}\right)$$

via excessive numerical experiments.

In practice, inspired by the NUFFT in [33], we can replace the Hadamard product of a low-rank matrix and a NUFFT matrix in (42) with a Hadamard product of a low-rank matrix and a DFT matrix, the matvec of which can be carried out more efficiently with a few numbers of FFTs.

Consider another matrix $\mathcal{B}_n^{(a,b)} \in \mathbb{C}^{n \times (n-27)}$ whose (j, k) -th entry is defined via

$$M^{(a,b)}(t_j, k+27) \exp\left(\mathbf{i}\left(\psi^{(a,b)}(t_j, k+27) - 2\pi \frac{[t_j n / 2\pi]}{n} (k+27)\right)\right), \quad (45)$$

where $[x]$ denotes the integer nearest to x . Then we have

$$\mathcal{J}_n^{(a,b)} = \left[\mathcal{V}_n^{(a,b)}, \Re(\mathcal{B}_n^{(a,b)} \otimes \mathcal{F}_n) \right], \quad (46)$$

where \mathcal{F}_n is an $n \times (n - 27)$ matrix whose (j, k) -th entry is

$$\exp\left(i(2\pi \frac{\lfloor t_j n / 2\pi \rfloor}{n} (k + 27))\right). \quad (47)$$

It's obvious that \mathcal{F}_n in (46) is a row permutation of an inverse DFT matrix. Note that the difference between the phase functions of $\mathcal{A}_n^{(a,b)}$ and $\mathcal{B}_n^{(a,b)}$ is less than π . Hence, $\mathcal{B}_n^{(a,b)}$ is also a low-rank matrix and we can apply the randomized SVD in Algorithm 2.1 to construct a low-rank approximation of $\mathcal{B}_n^{(a,b)}$ in $O(r^2n)$ operations.

Suppose we have constructed an approximate rank- r SVD of $\mathcal{B}_n^{(a,b)}$ up to a desired accuracy using Algorithm 2.1; that is, suppose that we have computed the factorization

$$\mathcal{B}_n^{(a,b)} \approx USV, \quad (48)$$

where $U \in \mathbb{C}^{n \times r}$, $V \in \mathbb{C}^{r \times (n-27)}$, and $S \in \mathbb{R}^{r \times r}$ is a positive definite diagonal matrix. By rearranging the factors above, we have

$$\mathcal{B}_n^{(a,b)} \approx (US^{\frac{1}{2}})(S^{\frac{1}{2}}V) = u_1v_1^T + \dots + u_rv_r^T, \quad (49)$$

where u_i and v_i^T denote the i^{th} column vector of $US^{\frac{1}{2}}$ and the i^{th} row vector of $S^{\frac{1}{2}}V$, respectively, and T denotes the matrix transpose. Once (49) is ready, we have

$$W\mathcal{J}_n^{(a,b)}\alpha \approx W\mathcal{V}_n^{(a,b)}\alpha_{\mathcal{V}} + W\Re\left(\left(\sum_{j=1}^r u_jv_j^T\right) \otimes \mathcal{F}_n\right)\alpha_{\mathcal{G}} = W\mathcal{V}_n^{(a,b)}\alpha_{\mathcal{V}} + W\Re\left(\sum_{j=1}^r D_{u_j}\mathcal{F}_nD_{v_j}\alpha_{\mathcal{G}}\right), \quad (50)$$

where D_u denotes a diagonal matrix with u on its diagonal. Formula (50) indicates that the Jacobi polynomial transform can be evaluated efficiently via r inverse FFTs, which requires $O(rn \log n)$ arithmetic operations and $O(rn)$ memory. SVD gives an optimal numerical rank of a matrix given a fixed accuracy in the low-rank approximation. Hence, the rank in our new method should be smaller than or equal to the rank in [40]. Numerical performance in this paper and [40] leads to the conjecture that the rank

$$r = O\left(\frac{\log n}{\log \log n}\right).$$

Compared to (44) used in the original fast Jacobi transform in [40], the number of inverse FFTs in (50) has been optimized. Thus Formula (50) would take fewer operations to compute a multi-dimensional transform. Note that, compared to the original method in [40], though our new method using (50) would cost more time to construct a low-rank approximation, the optimized rank of the low-rank approximation would accelerate the application of the multi-dimensional Jacobi transform. Besides, as we mentioned previously, the new method works in a larger range of a and b . In Section 4, we will provide numerical comparisons to demonstrate the superiority of the new method over the method in [40].

The fast inverse Jacobi transform in the uniform case can be carried out in a similar manner, since $W\mathcal{J}_n^{(a,b)}$ is an orthonormal matrix. In fact, the inverse transform can be computed via

$$\alpha_{\mathcal{V}} = \left(\mathcal{V}_n^{(a,b)}\right)^T W^T W f,$$

and

$$\alpha_{\mathcal{G}} \approx \Re\left(\left(\sum_{j=1}^r u_jv_j^T\right) \otimes \mathcal{F}_n\right)^T W^T W f = \Re\left(\sum_{j=1}^r D_{v_j}\mathcal{F}_n^T D_{u_j}W^T W f\right), \quad (51)$$

where \mathcal{F}_n^T is a permutation of the DFT matrix. Therefore, the inverse Jacobi polynomial transform can also be computed via r FFTs.

It is worth emphasizing that the fast algorithm introduced in this section also works for non-uniform Jacobi polynomial transforms since the low-rankness of $\mathcal{B}_n^{(a,b)}$ is independent of the samples in t and the quadrature weights in the uniform Jacobi polynomial transform. There will be numerical examples later to verify this. The inverse transform in the non-uniform case requires solving a highly ill-conditioned linear system, which will be reported in a separate paper in the future.

3.2 Two-dimensional transform and its inverse

In this section, we extend our algorithm to the two-dimensional case, using the new method developed in the last section. We only focus on the transform for $\tilde{P}_\nu^{(a,b)}(t)$; the one for $\tilde{Q}_{k-1}^{(a,b)}(t)$ is similar. Furthermore, we concentrate on transforms admitting a tensor-product structure which is usually the case in real applications. We will adopt the MATLAB notation $A(\cdot)$ or $vec(A)$ as a vector resulting from reshaping the matrix A into a vector. We also use similar notations as in Section 3.1, e.g., $\mathcal{J}_{n,\cdot}^{(a,b)}$, $\mathcal{V}_{n,\cdot}^{(a,b)}$, $\mathcal{G}_{n,\cdot}^{(a,b)}$, $\mathcal{B}_{n,\cdot}^{(a,b)}$, $\mathcal{F}_{n,\cdot}$, and W . denote corresponding matrices analogous to their counterparts in Section 3.1, respectively, with “.” specifying a variable x or y in the spatial domain. In the rest of this section, we always assume a low-rank approximation

$$\mathcal{B}_{n,\cdot}^{(a,b)} = \sum_{i=1}^r u_{\cdot,i} v_{\cdot,i}^T \quad (52)$$

has been obtained by Algorithm 2.1 up to a desired accuracy for “.” as x or y .

Given locations $\{x_i\}_{i=1,\dots,n} \subset (0, \pi)$ and $\{y_i\}_{i=1,\dots,n} \subset (0, \pi)$, with no substantial difference, the forward and inverse two-dimensional Jacobi polynomial transforms arise in the following Jacobi expansion

$$f(x_i, y_j) = \sum_{k=1}^n \sum_{\ell=1}^n \alpha(k, \ell) \tilde{P}_{k-1}^{(a,b)}(x_i) \tilde{P}_{\ell-1}^{(a,b)}(y_j), \quad \text{for } i, j = 1, \dots, n, \quad (53)$$

where α denotes an expansion coefficients matrix. The forward and inverse transform can be defined analogously as in the one-dimensional case. When both $\{x_i\}_{i=1,\dots,n}$ and $\{y_i\}_{i=1,\dots,n}$ are exactly the nodes of the trigonometric Gauss-Jacobi quadrature rule, the corresponding transform is referred to as the uniform transform. With additional weight matrices W_x and W_y , the forward transformation matrix $(W_x \mathcal{J}_{n,x}^{(a,b)}) \odot (W_y \mathcal{J}_{n,y}^{(a,b)})$ taking $\alpha(\cdot)$ to $(W_x \odot W_y) f(\cdot)$ is orthogonal, where “ \odot ” denotes the Kronecker product. When $\{x_i\}_{i=1,\dots,n}$, $\{y_i\}_{i=1,\dots,n}$, and weights are not given by the trigonometric Gauss-Jacobi quadrature rule, we refer the corresponding transform as a non-uniform transform.

With the low-rank approximations in (52) available, we have

$$\left[(W_x \mathcal{J}_{n,x}^{(a,b)}) \odot (W_y \mathcal{J}_{n,y}^{(a,b)}) \right] \alpha(\cdot) = vec \left[((W_x \mathcal{J}_{n,x}^{(a,b)}) ((W_y \mathcal{J}_{n,y}^{(a,b)}) \alpha)^T)^T \right]. \quad (54)$$

This formula indicates that 2D uniform Jacobi polynomial transforms can be evaluated via $2n$ 1D forward transforms, thus $O((r_x + r_y)n)$ 1D inverse FFTs, which results in $O((r_x + r_y)n^2 \log(n))$ arithmetic operations and $O((r_x + r_y)n^2)$ memory. We would like to emphasize that this algorithm also works for a 2D non-uniform forward transform since it is a tensor form of the one-dimensional transform. There will be numerical examples later to verify this. Note that although the left-hand side and the right-hand side of Formula (54) are mathematically equivalent, we adopt the right-hand side in implementation to compute the 2D transform because Level 3 BLAS can be applied

to accelerate matrix-matrix multiplications. Analogous techniques can be used to compute other multidimensional transforms in this paper.

The fast 2D inverse Jacobi transform in the uniform case can be carried out in a similar manner, since both $W_x \mathcal{J}_{n,x}^{(a,b)}$ and $W_y \mathcal{J}_{n,y}^{(a,b)}$ are orthonormal matrices. In fact, the 2D inverse transform can be computed via

$$\alpha(\cdot) = \left[(\mathcal{J}_{n,x}^{(a,b)})^T \odot (\mathcal{J}_{n,y}^{(a,b)})^T \right] \tilde{f}(\cdot) = \text{vec}((\mathcal{J}_{n,y}^{(a,b)})^T \tilde{f} \mathcal{J}_{n,x}^{(a,b)}), \quad (55)$$

where $\tilde{f} := W_y^2 f W_x^2$. Again, Level 3 BLAS and $2n$ 1D inverse transforms can be applied to compute this 2D transform. Therefore, 2D inverse Jacobi polynomial transforms can also be evaluated via $O((r_x + r_y)n)$ 1D FFTs.

3.3 Three-dimensional transform and its inverse

In this section, we continue to extend our algorithm for 3D Jacobi polynomial transform and its inverse. Analogously, we just discuss the transform for $\tilde{P}_\nu^{(a,b)}(t)$ and the transform for $\tilde{Q}_{k-1}^{(a,b)}(t)$ is similar. Just like the 2D case, we mainly focus on those transforms with a tensor-product structure. The notations in Section 3.2 will be inherited here.

Given locations $\{x_i\}_{i=1,\dots,n}$, $\{y_i\}_{i=1,\dots,n}$, and $\{z_i\}_{i=1,\dots,n}$ in $(0, \pi)$, with no substantial difference, the forward and inverse three-dimensional Jacobi polynomial transforms arise in the following Jacobi expansion

$$f(x_i, y_j, z_k) = \sum_{h=1}^n \sum_{\ell=1}^n \sum_{m=1}^n \alpha(k, \ell, m) \tilde{P}_{h-1}^{(a,b)}(x_i) \tilde{P}_{\ell-1}^{(a,b)}(y_j) \tilde{P}_{m-1}^{(a,b)}(z_k), \quad \text{for } i, j, k = 1, \dots, n, \quad (56)$$

where α denotes a three-dimensional tensor containing expansion coefficients. The forward and inverse transforms can be defined analogously to the three-dimensional case. When $\{x_i\}_{i=1,\dots,n}$, $\{y_i\}_{i=1,\dots,n}$, and $\{z_i\}_{i=1,\dots,n}$ are all exactly the nodes of the trigonometric Gauss-Jacobi quadrature rule, the corresponding transform is referred to as the uniform transform. Otherwise, we refer the corresponding transform as a non-uniform transform.

In the uniform transform, in order to take advantage of orthogonality, we consider the tensor product $(W_x \mathcal{J}_{n,x}^{(a,b)}) \odot (W_y \mathcal{J}_{n,y}^{(a,b)}) \odot (W_z \mathcal{J}_{n,z}^{(a,b)})$ that takes $\alpha(\cdot)$ to $(W_x \odot W_y \odot W_z) f(\cdot)$. Once the related low-rank approximations,

$$\mathcal{B}_{n,\cdot}^{(a,b)} = \sum_{i=1}^r u_{\cdot,i} v_{\cdot,i}^T, \quad (57)$$

have been obtained, we have

$$(W_x \odot W_y \odot W_z) f(\cdot) = (W_x \odot W_y \odot W_z) (\mathcal{J}_{n,x}^{(a,b)} \odot \mathcal{J}_{n,y}^{(a,b)} \odot \mathcal{J}_{n,z}^{(a,b)}) \alpha(\cdot), \quad (58)$$

where

$$(\mathcal{J}_{n,x}^{(a,b)} \odot \mathcal{J}_{n,y}^{(a,b)} \odot \mathcal{J}_{n,z}^{(a,b)}) \alpha(\cdot) = \text{vec} \left[((\mathcal{J}_{n,y}^{(a,b)} \odot \mathcal{J}_{n,z}^{(a,b)}) \tilde{\alpha}) (\mathcal{J}_{n,x}^{(a,b)})^T \right] \quad (59)$$

and $\tilde{\alpha}$ is a $n^2 \times n$ matrix reshaped from α .

Formula (59) implies that a 3D uniform Jacobi polynomial transform can be computed by evaluating the right-hand side using Level 3 BLAS routines with n 2D forward transforms and n^2 1D forward transforms, which are dominated by totally $O((r_x + r_y + r_z)n^2)$ inverse FFTs and result in $O((r_x + r_y + r_z)n^3 \log(n))$ arithmetic operations and $O((r_x + r_y + r_z)n^3)$ memory. Again, we




Figure 1: Numerical results for the comparison of the CHEB and RS method for 3D forward uniform transform. Top: the factorization and application time (in second) from left to right, respectively. Bottom left and right: the numerical ranks of the low-rank matrices provided by CHEB and RS, and the relative errors of the fast matvec compared to the exact summation.

would like to emphasize that this algorithm also works for 3D non-uniform forward transforms. There will be numerical examples later to verify this.

The fast 3D inverse (uniform) Jacobi transform in the uniform case can be carried out in a similar manner, since $W_x \mathcal{J}_{n,x}^{(a,b)}$, $W_y \mathcal{J}_{n,y}^{(a,b)}$ and $W_z \mathcal{J}_{n,z}^{(a,b)}$ are all orthonormal matrices. In fact, the 3D inverse transform can be computed via

$$(\mathcal{J}_{n,x}^{(a,b)} \odot \mathcal{J}_{n,y}^{(a,b)} \odot \mathcal{J}_{n,z}^{(a,b)})^T \tilde{f}(\cdot) = \text{vec} \left[(((\mathcal{J}_{n,y}^{(a,b)})^T \odot (\mathcal{J}_{n,z}^{(a,b)})^T) \bar{f}) \mathcal{J}_{n,x}^{(a,b)} \right] \quad (60)$$

where $\tilde{f}(\cdot) = (W_x \odot W_y \odot W_z)^2 f(\cdot)$ and \bar{f} is an $n^2 \times n$ matrix reshaped from \tilde{f} .

Analogously, the right-hand side allows Level 3 BLAS with n inverse 2D transforms and n^2 1D inverse transforms. Therefore, the 3D inverse Jacobi polynomial transform can also be evaluated via $O((r_x + r_y + r_z)n^2)$ 1D FFTs, resulting in a nearly linear scaling algorithm.

4 Numerical results

This section presents several numerical examples to demonstrate the effectiveness of the algorithms proposed above. Section 4.1 demonstrates the stability of our algorithm. Section 4.2 provides a comparison of the original method in [40] named as method **CHEB** using (44) in Section 3.1, our new method named as **RS** using (50) in Section 3.1 to demonstrate the superiority of our algorithm. In Section 4.3, we apply our new method to 2D and 3D Jacobi polynomial transforms with parameters $a = b = 0.40$ to show the complexity of our algorithm. All implementations are in MATLAB[®] on a server computer with 28 processors and 2.6 GHz CPU. But only one processor was used in every single experiment. We have made our code, including that for all of the experiments described here, available on GitLab at the following address:

a=b \ log ₂ n	1D			2D			3D		
	10	15	20	6	9	12	6	7	8
-0.75	1.00	13.7	42.3	2.50	0.10	0.005	3.05	0.89	2.36
-0.50	0.001	0.23	4.25	1E-5	1E-4	0.001	1E-4	7E-5	7E-5
-0.25	0.33	4.65	44.3	1.57	0.07	0.005	1.87	1.04	3.56
0.00	0.69	8.10	60.0	3.16	0.24	0.009	3.81	0.99	1.22
0.25	0.71	1.95	23.7	4.74	0.36	0.009	5.90	1.52	1.66
0.50	0.30	0.34	4.34	0.03	1E-4	0.001	0.97	0.23	0.35
0.75	2.06	4.80	60.1	5.81	0.37	0.02	7.13	2.13	1.41

Table 1: Numerical results of Equation (61) in the cases of 1D, 2D and 3D to measure stability. All numbers in this table are measured in the unit of $1E - 8$. Hence, the quantity in Equation (61) is of order $1E - 8$ in all test cases.

rank \ log ₂ n	a=b					
	14	15	16	17	18	19
0.0	18, 20	19, 21	19, 22	20, 23	21, 24	20, 26
0.1	19, 20	19, 21	19, 23	20, 23	21, 24	20, 26
0.2	18, 20	18, 21	19, 22	20, 23	21, 24	20, 26
0.3	18, 19	19, 21	20, 22	19, 23	21, 24	20, 25
0.4	18, 19	19, 19	19, 21	19, 22	21, 24	19, 25
0.5	9, 9	9, 9	9, 9	9, 9	9, 9	9, 9
0.6	18, 19	18, 20	19, 21	19, 23	20, 23	20, 24
0.7	18, 19	18, 20	20, 21	20, 22	21, 23	20, 24
0.8	18, 19	19, 20	20, 21	19, 21	20, 23	20, 24
0.9	18, 18	19, 20	20, 21	19, 21	20, 22	19, 23
-0.1	18, 20	19, 21	19, 23	19, 23	20, 25	20, 26
-0.2	17, 20	18, 21	18, 22	19, 23	20, 24	18, 27
-0.3	18, 19	18, 21	18, 22	18, 23	20, 24	18, 25
-0.4	17, 19	17, 19	17, 21	18, 22	19, 24	18, 25
-0.5	2, 9	2, 9	3, 9	3, 9	4, 9	4, 9
-0.6	16, 19	18, 20	17, 21	18, 22	18, 23	18, 24
-0.7	17, 19	18, 20	18, 21	18, 22	19, 23	18, 24
-0.8	17, 21	18, 21	18, 22	18, 23	19, 25	19, 25
-0.9	18, 20	18, 21	18, 22	18, 22	19, 23	19, 24

Table 2: Numerical ranks of the CHEB and RS algorithms for 1D forward uniform transform up to a fixed accuracy $1E - 8$ when $a = b$. There are two numbers in each item: the left one denotes the numerical rank by the RS algorithm and the right one is for the CHEB method in [40].

error a=b	$\log_2 n$	14	15	16	17	18	19
0.0		0.17, 0.66	0.28, 0.41	0.13, 0.77	0.13, 0.67	0.95, 0.57	0.80, 0.28
0.1		0.10, 0.60	0.45, 0.46	1.58, 0.20	0.73, 0.51	1.04, 1.06	5.47, 0.45
0.2		0.48, 0.25	1.01, 0.73	1.51, 0.77	0.96, 0.41	0.73, 1.08	8.97, 0.22
0.3		0.91, 0.69	0.32, 0.47	0.37, 0.25	0.15, 0.43	0.92, 0.31	2.41, 0.80
0.4		0.30, 1.19	0.34, 1.99	0.71, 1.21	0.62, 1.57	0.52, 0.53	4.38, 0.24
0.5		0.31, 0.32	0.31, 0.27	0.17, 0.35	0.11, 0.14	0.32, 0.45	0.17, 0.27
0.6		0.32, 0.40	0.92, 0.41	1.51, 0.33	0.94, 0.18	3.98, 0.76	2.53, 1.72
0.7		0.57, 0.40	1.03, 1.40	0.61, 1.94	1.72, 1.17	2.42, 1.58	4.07, 2.30
0.8		0.79, 0.93	0.73, 0.53	0.83, 1.43	4.85, 3.09	7.15, 1.45	12.5, 1.58
0.9		1.03, 2.92	1.01, 0.84	0.80, 2.36	3.57, 4.84	4.40, 2.62	26.9, 5.43
-0.1		0.37, 0.58	0.69, 0.33	1.09, 0.22	1.80, 0.59	1.46, 0.37	7.44, 0.39
-0.2		0.45, 0.40	0.33, 0.41	1.53, 0.97	1.65, 0.48	1.39, 0.81	9.30, 0.21
-0.3		0.08, 0.72	0.36, 0.33	0.87, 0.18	1.54, 0.40	1.88, 0.17	16.4, 0.95
-0.4		0.12, 0.98	0.64, 0.93	1.57, 0.85	2.03, 0.97	1.27, 0.30	4.22, 0.14
-0.5		0.03, 0.47	0.05, 0.40	0.05, 0.51	0.31, 0.31	0.03, 0.48	0.08, 0.47
-0.6		0.57, 0.15	0.17, 0.47	1.44, 0.15	0.99, 0.15	1.33, 0.27	3.97, 1.79
-0.7		0.40, 0.19	0.45, 1.72	1.61, 1.18	3.13, 1.68	1.56, 0.83	5.56, 3.62
-0.8		0.68, 3E06	1.43, 1E06	2.55, 8E05	3.34, 7E05	4.19, 5E05	4.87, 3E05
-0.9		0.62, 3E06	1.25, 1E06	2.97, 8E05	3.21, 6E05	5.31, 5E05	8.73, 3E05

Table 3: Numerical relative accuracy of the CHEB and RS algorithms for 1D forward uniform transform up to a fixed accuracy $1E - 8$ when $a = b$. The units of all data in this table are $1E - 8$. There are two floats in each item: the left one denotes the relative accuracy of the low-rank factorization by the RS algorithm and the right one is for the CHEB algorithm in [40].




Figure 2: Numerical results for the 2D uniform forward transform (the first row), the 2D uniform inverse transform (the second row), and the 2D nonuniform forward transform. The running time (in second), the numerical rank of the low-rank matrix in (52), and the relative error of the fast algorithm compared to the exact summation are visualized from left to right columns.

<https://gitlab.com/FastOrthPolynomial/Jacobi.git>

In our numerical results, n denotes the number of grid points per dimension; “fac” and “app” stand for the factorization time for setting up the algorithm and the application time for applying the algorithm, respectively; “releerr” means the relative error of the fast matvec. The rank parameter r (the upper bound of the rank for each dimension of a low rank factorization) and the accuracy parameter ϵ in all low-rank factorization algorithms are set to $\lceil 2 \log_2(N) \rceil$ and $1e-8$, respectively, where $\lceil \cdot \rceil$ is the ceiling function. To make fair comparisons, the oversampling parameter q in Algorithm 2.1 for the RS method is set such that $q \times r$ is approximately equal to the number of piecewise Chebyshev grids in (30) used in the CHEB method. We perform the same experiment 10 times and summarize the average statistics in the following figures. All timings are measured in seconds.

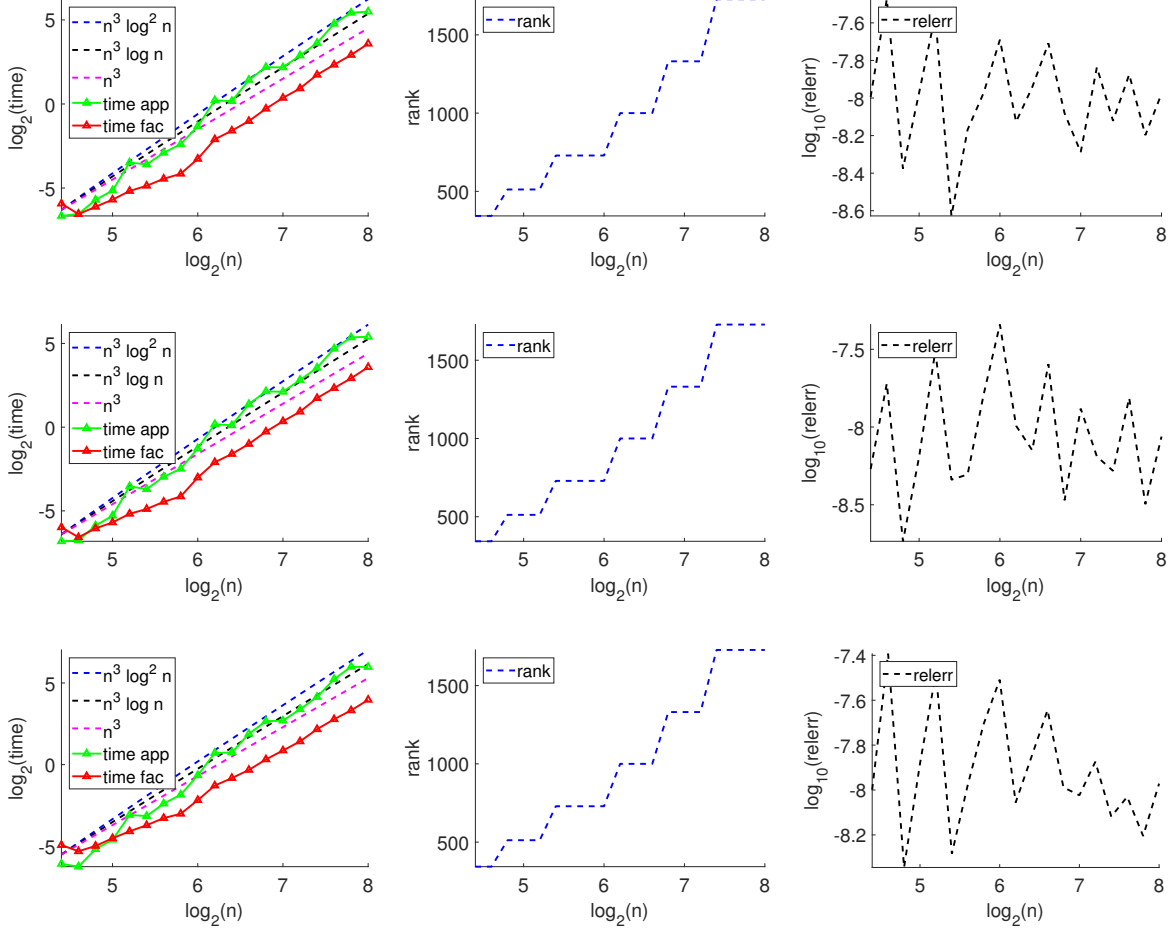


Figure 3: Numerical results for the 3D uniform forward transform (the first row), the 3D uniform inverse transform (the second row), and the 3D nonuniform forward transform. The running time (in second), the numerical rank of the low-rank matrix in (57), and the relative error of the fast algorithm compared to the exact summation are visualized from left to right columns.

4.1 Stability

In the case of “uniform” transform, the transformation matrix is orthonormal after weighting (Eq. 50) and hence the condition number is 1. Hence, the forward and inverse transforms are well-conditioned. The proposed fast algorithm explores low-rank approximation to special structures in the Jacobi polynomial matrix $\mathcal{J}_{n,x}^{(a,b)}$ and may result in perturbation error. To ease the concern of numerical stability, we evaluate the following quantity for random vectors v to measure the stability of the propose fast forward and inverse transforms:

$$\|(W \mathcal{J}_{n,x}^{(a,b)})^{-1} W \mathcal{J}_{n,x}^{(a,b)} v - v\| / \|v\|, \quad (61)$$

where $(W \mathcal{J}_{n,x}^{(a,b)})^{-1} = (W \mathcal{J}_{n,x}^{(a,b)})^T$ due to the orthonormal property. Table 1 summarizes the above quantity for different a 's and b 's in different dimensions and shows that the proposed fast algorithm is stable.

4.2 Comparison of CHEB and RS methods

4.2.1 Comparisons in the one-dimensional case for a and b in $(-1, 1)$

First, we provide comparisons between CHEB and RS algorithms when a and b are set to be different values in the range $(-1, 1)$. The numerical results are summarized in Table 2 and Table 3. The numerical results show that both CHEB and RS achieve desired accuracy when a and b are not close to -1 ; RS works but CHEB fails when a and b are close to -1 . Hence, the newly proposed RS algorithm works in a larger range of a and b than the CHEB method in [40].

4.2.2 Comparisons in the three-dimensional case for a and b in $(-1, 1)$

Second, we provide performance comparisons of CHEB algorithm and RS algorithm in the three-dimensional case with $a = b = 0.40$ to demonstrate that the RS is more efficient than the CHEB algorithm. In the one and two-dimensional cases, RS is also faster than CHEB but the speed-up is less obvious.

As we can see from the results summarized in Figure 1, the RS method provides a more compact matrix compression while keeping the compression accuracy competitive to that of the CHEB algorithm. The more compact compression results in faster set-up and application time for the Jacobi transform for all problem sizes.

4.3 Performance of fast multi-dimensional transforms

Finally, we will present examples for two and three-dimensional Jacobi polynomial transforms using the RS algorithm. In either the two or three-dimensional case, there are numerical results for one uniform forward transform, one uniform inverse transform, and one nonuniform forward transform. a and b are all set to 0.40 in these examples.

Figure 2 and 3 summarize the numerical results in 2D and 3D, respectively. In cases, no matter forward or inverse transform, uniform or nonuniform transform, 2D or 3D, the factorization and application time of our algorithm scales like $O(n^d \log n)$ or $O(n^d \log^2 n)$, where d is the dimension. The numerical rank gradually increases like $O(\log n)$ as the problem size increases. The relative error of the fast algorithm is in line with the desired accuracy in the low-rank factorization.

5 Conclusion

This paper proposed a fast algorithm for multi-dimensional Jacobi polynomial transforms based on the observation that the solution of Jacobi's differential equation can be represented via non-oscillatory phase and amplitude functions. In particular, the transformation matrix corresponding to the discrete transform is a Hadamard product of a numerically low-rank matrix and a multi-dimensional discrete Fourier transform matrix. After constructing a low-rank approximation to the numerical low-rank matrix, the application of the Hadamard product can be carried out via $O(r^d)$ fast Fourier transforms, where $r = O(\frac{\log n}{\log \log n})$, n is the grid size per dimension, and d is the dimension, resulting in a nearly optimal algorithm to compute the multi-dimensional Jacobi polynomial transform.

We proposed to apply the randomized SVD to construct low-rank factorizations, resulting in a faster Jacobi transform in high-dimensional spaces. Moreover, numerical experiments show that the new fast transform works for a larger class of Jacobi polynomials with parameter a and b in the interval $(-1, 1)$ than the one-dimensional algorithm in [40].

For other values of a and b outside $(-1, 1)$, the Jacobi transformation matrix is no longer purely oscillatory and hence the proposed method is not applicable. We will tackle this issue in the future via hierarchically dividing the transformation matrix into sub-matrices that are either purely non-oscillatory, which can be handled by fast low-rank factorization, or purely oscillatory, which can be applied via the fast algorithm in this paper.

The fast inverse in the nonuniform case is still an open problem. It involves a highly ill-conditioned linear system of equations. We are working on an efficient preconditioner for this linear system and will summarize our work in a separate paper.

Acknowledgments. H. Y. was partially supported by Grant R-146-000-251-133 in the Department of Mathematics at the National University of Singapore, by the Ministry of Education in Singapore under the grant MOE2018-T2-2-147, and the start-up grant of the Department of Mathematics at Purdue University.

References

- [1] Richard Askey. *Orthogonal Polynomials and Special Functions*. Society for Industrial and Applied Mathematics, 1975.
- [2] Gabor Szeg. *Orthogonal polynomials*, volume 23. American Mathematical Soc., 1939.
- [3] Aleksandr Filippovich Timan. *Theory of Approximation of Functions of a Real Variable*. Dover books on advanced mathematics. Dover Publications, 1994.
- [4] John P. Boyd. *Chebyshev and Fourier Spectral Methods: Second Revised Edition*. Dover Books on Mathematics. Dover Publications, 2001.
- [5] Claudio Canuto, M. Yousuff Hussaini, Alfio Quarteroni, and Thomas A. Zang. *Spectral Methods: Fundamentals in Single Domains*, pages 3–37. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- [6] Daniele Funaro. *Polynomial Approximation of Differential Equations*. Number v. 8 in Lecture notes in physics: Monographs. Springer-Verlag, 1992.
- [7] David Gottlieb and Steven A. Orszag. *Numerical Analysis of Spectral Methods*. Society for Industrial and Applied Mathematics, 1977.
- [8] Guo Ben-Yu. *Spectral Methods and Their Applications*. WORLD SCIENTIFIC, 1998.
- [9] Zachary Battles and Lloyd N. Trefethen. An extension of matlab to continuous functions and operators. *SIAM Journal on Scientific Computing*, 25(5):1743–1770, 2004.
- [10] Alex Townsend and Lloyd N. Trefethen. An Extension of Chebfun to Two Dimensions. *SIAM Journal on Scientific Computing*, 35(6):C495–C518, 2013.
- [11] Behnam Hashemi and Lloyd N. Trefethen. Chebfun in Three Dimensions. *SIAM Journal on Scientific Computing*, 39(5):C341–C363, 2017.
- [12] Christine Bernardi and Yvon Maday. *Spectral methods*, volume 5 of *Handbook of Numerical Analysis*. Elsevier, 1997.
- [13] Guo Ben-Yu. Gegenbauer Approximation and Its Applications to Differential Equations on the Whole Line. *Journal of Mathematical Analysis and Applications*, 226(1):180 – 206, 1998.

- [14] Ben-Yu Guo. Jacobi Approximations in Certain Hilbert Spaces and Their Applications to Singular Differential Equations. *Journal of Mathematical Analysis and Applications*, 243(2):373 – 408, 2000.
- [15] Ben-Yu Guo. Jacobi spectral approximations to differential equations on the half line. *Journal of Computational Mathematics*, 18(1):95–112, 2000.
- [16] Ivo Babuška and Benqi Guo. Optimal estimates for lower and upper bounds of approximation errors in the p-version of the finite element method in two dimensions. *Numerische Mathematik*, 85(2):219–255, Apr 2000.
- [17] Ivo Babuška and Benqi Guo. Direct and Inverse Approximation Theorems for the p-Version of the Finite Element Method in the Framework of Weighted Besov Spaces. Part I: Approximability of Functions in the Weighted Besov Spaces. *SIAM Journal on Numerical Analysis*, 39(5):1512–1538, 2002.
- [18] Benqi Guo and Ivo Babuška. Direct and inverse approximation theorems for the p-version of the finite element method in the framework of weighted Besov spaces, Part III: Inverse approximation theorems. *Journal of Approximation Theory*, 173:122 – 157, 2013.
- [19] Ben-Yu Guo, Jie Shen, and Li-Lian Wang. Generalized jacobi polynomials/functions and their applications. *Applied Numerical Mathematics*, 59(5):1011 – 1028, 2009.
- [20] Arieh Iserles. A fast and simple algorithm for the computation of Legendre coefficients. *Numerische Mathematik*, 117(3):529–553, Mar 2011.
- [21] Nicholas Hale and Alex Townsend. A fast, simple, and stable Chebyshev–Legendre transform using an asymptotic formula. *SIAM Journal on Scientific Computing*, 36(1):A148–A167, 2014.
- [22] Bradley K Alpert and Vladimir Rokhlin. A fast algorithm for the evaluation of Legendre expansions. *SIAM Journal on Scientific and Statistical Computing*, 12(1):158–179, 1991.
- [23] Vladimir Rokhlin. A fast algorithm for the evaluation of Legendre expansions. *SIAM Journal on Sci. and Stat. Computing*, 12:158–179, 1991.
- [24] Jens Keiner. *Fast Polynomial Transforms*. Logos Verlag, 2011.
- [25] Richard Askey. *Orthogonal polynomials and special functions*, volume 21. Siam, 1975.
- [26] George E Andrews, Richard Askey, and Ranjan Roy. Special functions, volume 71 of encyclopedia of mathematics and its applications, 1999.
- [27] P. Maroni and Z. da Rocha. Connection coefficients between orthogonal polynomials and the canonical sequence: an approach based on symbolic computation. *Numerical Algorithms*, 47(3):291–314, Mar 2008.
- [28] Luogeng Hua. *Harmonic analysis of functions of several complex variables in the classical domains*. Number 6. American Mathematical Soc., 1963.
- [29] Alex Townsend, Marcus Webb, and Sheehan Olver. Fast polynomial transforms based on toeplitz and hankel matrices. *Mathematics of Computation*, 87(312):1913–1934, 2017.

- [30] Richard Mikal Slevinsky. On the use of hahns asymptotic formula and stabilized recurrence for a fast, simple and stable chebyshevjacobi transform. *IMA Journal of Numerical Analysis*, 38(1):102–124, 2018.
- [31] Jie Shen, Yingwei Wang, and Jianlin Xia. Fast structured Jacobi-Jacobi transforms. *Mathematics of Computation*, 88(318):1743–1772, 2018.
- [32] Leslie Greengard and June-Yub Lee. Accelerating the nonuniform fast Fourier transform. *SIAM Review*, 46(3):443–454, 2004.
- [33] Diego Ruiz-Antolín and Alex Townsend. A nonuniform fast Fourier transform based on low rank approximation. *SIAM Journal on Scientific Computing*, 40(1):A529–A547, 2018.
- [34] Leslie Greengard and Vladimir Rokhlin. A fast algorithm for particle simulations. *Journal of computational physics*, 73(2):325–348, 1987.
- [35] Jens Keiner. Computing with expansions in Gegenbauer polynomials. *SIAM Journal on Scientific Computing*, 31(3):2151–2171, 2009.
- [36] Raf Vandebril, Marc Van Barel, and Nicola Mastronardi. *Matrix Computations and Semiseparable Matrices*. The Johns Hopkins University Press, Baltimore, MD, 2008.
- [37] Jens Keiner. Gegenbauer polynomials and semiseparable matrices. *Electronic Transactions on Numerical Analysis*, 30:26–53, 2008.
- [38] Michael O’Neil, Franco Woolfe, and Vladimir Rokhlin. An algorithm for the rapid evaluation of special function transforms. *Applied and Computational Harmonic Analysis*, 28:203–226, 2010.
- [39] D. Potts, G. Steidl, and M. Tasche. Fast algorithms for discrete polynomial transforms. *Mathematics of Computation*, 67:1577–1590, October 1998.
- [40] James Bremer and Haizhao Yang. Fast algorithms for Jacobi expansions via nonoscillatory phase functions. *IMA Journal of Numerical Analysis*, 04 2019.
- [41] *NIST Digital Library of Mathematical Functions*. <http://dlmf.nist.gov/>, Release 1.0.13 of 2016-09-16. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller and B. V. Saunders, eds.
- [42] Zhu Heitman, James Bremer, and Vladimir Rokhlin. On the existence of nonoscillatory phase functions for second order ordinary differential equations in the high-frequency regime. *Journal of Computational Physics*, 290:1 – 27, 2015.
- [43] James Bremer. On the numerical solution of second order differential equations in the high-frequency regime. *Applied and Computational Harmonic Analysis*, 44:312–349, 2018.
- [44] Emmanuel Candès, Laurent Demanet, and Lexing Ying. Fast computation of Fourier integral operators. *SIAM Journal on Scientific Computing*, 29(6):2464–2493, 2007.
- [45] Haizhao Yang. A unified framework for oscillatory integral transforms: When to use nufft or butterfly factorization? *J. Comput. Physics*, 388:103–122, 2018.
- [46] Björn Engquist and Lexing Ying. A fast directional algorithm for high frequency acoustic scattering in two dimensions. *Commun. Math. Sci.*, 7(2):327–345, 06 2009.

- [47] Jiawei Chiu and Laurent Demanet. Sublinear randomized algorithms for skeleton decompositions. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1361–1383, 2013.
- [48] Lloyd N. Trefethen. *Approximation Theory and Approximation Practice*. Society for Industrial and Applied Mathematics, 2013.