

Short Note

Fast algorithms for spherical harmonic expansions, III

Mark Tygert

Courant Institute of Mathematical Sciences, NYU, 251 Mercer St., New York, NY 10012, United States

ARTICLE INFO

Article history:

Received 13 December 2009
 Received in revised form 5 April 2010
 Accepted 10 May 2010
 Available online 23 May 2010

Keywords:

Butterfly
 Algorithm
 Spherical harmonic
 Transform
 Interpolative decomposition

ABSTRACT

We accelerate the computation of spherical harmonic transforms, using what is known as the butterfly scheme. This provides a convenient alternative to the approach taken in the second paper from this series on “Fast algorithms for spherical harmonic expansions”. The requisite precomputations become manageable when organized as a “depth-first traversal” of the program’s control-flow graph, rather than as the perhaps more natural “breadth-first traversal” that processes one-by-one each level of the multilevel procedure. We illustrate the results via several numerical examples.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

The butterfly algorithm, introduced in [10,11], is a procedure for rapidly applying certain matrices to arbitrary vectors. (Section 3 below provides a brief introduction to the butterfly.) The present paper uses the butterfly method in order to accelerate spherical harmonic transforms. The butterfly procedure does not require the use of extended-precision arithmetic in order to attain accuracy very close to the machine precision, not even in its precomputations – unlike the alternative approach taken in the predecessor [15] of the present paper.

Unlike some previous works on the butterfly, the present article does not use on-the-fly evaluation of individual entries of the matrices whose applications to vectors are being accelerated. Instead, we require only efficient evaluation of full columns of the matrices, in order to make the precomputations affordable. Furthermore, efficient evaluation of full columns enables the acceleration of the application to vectors of both the matrices and their transposes. On-the-fly evaluation of columns of the matrices associated with spherical harmonic transforms is available via the three-term recurrence relations satisfied by associated Legendre functions (see, for example, Section 5 below).

The precomputations for the butterfly become affordable when organized as a “depth-first traversal” of the program’s control-flow graph, rather than as the perhaps more natural “breadth-first traversal” that processes one-by-one each level of the multilevel butterfly procedure (see Section 4 below).

The present article is supposed to complement [11,15], combining ideas from both. Although the present paper is self-contained in principle, we strongly encourage the reader to begin with [11,15]. The original is [10]. Major recent developments are in [4,17]. The introduction in [15] summarizes most prior work on computing fast spherical harmonic transforms; a new application appears in [12]. These articles and their references highlight the computational use of spherical harmonic transforms in meteorology and quantum chemistry. The structure of the remainder of the present article is as follows: Section 2 reviews elementary facts about spherical harmonic transforms. Section 3 describes basic tools from previous works.

E-mail address: tygert@aya.yale.edu

Section 4 organizes the preprocessing for the butterfly to make memory requirements affordable. Section 5 outlines the application of the butterfly scheme to the computation of spherical harmonic transforms. Section 6 describes the results of several numerical tests. Section 7 draws some conclusions.

Throughout, we abbreviate “interpolative decomposition” to “ID” (see Section 3.1 for a description of the ID). The butterfly procedures formulated in [10,11] and the present paper all use the ID for efficiency.

2. An overview of spherical harmonic transforms

The spherical harmonic expansion of a bandlimited function f on the surface of the sphere has the form

$$f(\theta, \varphi) = \sum_{k=0}^{2l-1} \sum_{m=-k}^k \beta_k^m \bar{P}_k^{|m|}(\cos(\theta)) e^{im\varphi}, \quad (1)$$

where (θ, φ) are the standard spherical coordinates on the two-dimensional surface of the unit sphere in \mathbb{R}^3 , $\theta \in (0, \pi)$ and $\varphi \in (0, 2\pi)$, and $\bar{P}_k^{|m|}$ is the normalized associated Legendre function of degree k and order $|m|$ (see, for example, Section 3.3 for the definition of normalized associated Legendre functions). Please note that the superscript m in β_k^m denotes an index, rather than a power. “Normalized” refers to the fact that the normalized associated Legendre functions of a fixed order are orthonormal on $(-1, 1)$ with respect to the standard inner product. Obviously, the expansion (1) contains $4l^2$ terms. The complexity of the function f determines l .

In many areas of scientific computing, particularly those using spectral methods for the numerical solution of partial differential equations, we need to evaluate the coefficients β_k^m in an expansion of the form (1) for a function f given by a table of its values at a collection of appropriately chosen nodes on the two-dimensional surface of the unit sphere. Conversely, given the coefficients β_k^m in (1), we often need to evaluate f at a collection of points on the surface of the sphere. The former is known as the forward spherical harmonic transform, and the latter is known as the inverse spherical harmonic transform. A standard discretization of the surface of the sphere is the “tensor product”, consisting of all pairs of the form (θ_k, φ_j) , with $\cos(\theta_0), \cos(\theta_1), \dots, \cos(\theta_{2l-2}), \cos(\theta_{2l-1})$ being the Gauss–Legendre quadrature nodes of degree $2l$, that is,

$$-1 < \cos(\theta_0) < \cos(\theta_1) < \dots < \cos(\theta_{2l-2}) < \cos(\theta_{2l-1}) < 1 \quad (2)$$

and

$$\bar{P}_{2l}^0(\cos(\theta_k)) = 0 \quad (3)$$

for $k = 0, 1, \dots, 2l-2, 2l-1$, and with $\varphi_0, \varphi_1, \dots, \varphi_{4l-3}, \varphi_{4l-2}$ being equispaced on the interval $(0, 2\pi)$, that is,

$$\varphi_j = \frac{2\pi(j + \frac{1}{2})}{4l-1} \quad (4)$$

for $j = 0, 1, \dots, 4l-3, 4l-2$. This leads immediately to numerical schemes for both the forward and inverse spherical harmonic transforms whose costs are proportional to l^3 .

Indeed, given a function f defined on the two-dimensional surface of the unit sphere by (1), we can rewrite (1) in the form

$$f(\theta, \varphi) = \sum_{m=-2l+1}^{2l-1} e^{im\varphi} \sum_{k=|m|}^{2l-1} \beta_k^m \bar{P}_k^{|m|}(\cos(\theta)). \quad (5)$$

For a fixed value of θ , each of the sums over k in (5) contains no more than $2l$ terms, and there are $4l-1$ such sums (one for each value of m); since the inverse spherical harmonic transform involves $2l$ values $\theta_0, \theta_1, \dots, \theta_{2l-2}, \theta_{2l-1}$, the cost of evaluating all sums over k in (5) is proportional to l^3 . Once all sums over k have been evaluated, each sum over m may be evaluated for a cost proportional to l (since each of them contains $4l-1$ terms), and there are $(2l)(4l-1)$ such sums to be evaluated (one for each pair (θ_k, φ_j)), leading to costs proportional to l^3 for the evaluation of all sums over m in (5). The cost of the evaluation of the whole inverse spherical harmonic transform (in the form (5)) is the sum of the costs for the sums over k and the sums over m , and is also proportional to l^3 ; a virtually identical calculation shows that the cost of evaluating the forward spherical harmonic transform is also proportional to l^3 .

A trivial modification of the scheme described in the preceding paragraph uses the fast Fourier transform (FFT) to evaluate the sums over m in (5), approximately halving the operation count of the entire procedure. Several other careful considerations (see, for example, [2,13]) are able to reduce the costs by 50% or so, but there is no simple trick for reducing the costs of the whole spherical harmonic transform (either forward or inverse) below l^3 . The present paper presents faster (albeit more complicated) algorithms for both forward and inverse spherical harmonic transforms. Specifically, the present article provides a fast algorithm for evaluating a sum over k in (5) at $\theta = \theta_0, \theta_1, \dots, \theta_{2l-2}, \theta_{2l-1}$, given the coefficients $\beta_{|m|}^m, \beta_{|m|+1}^m, \dots, \beta_{2l-2}^m, \beta_{2l-1}^m$, for a fixed m . Moreover, the present paper provides a fast algorithm for the inverse procedure of determining the coefficients $\beta_{|m|}^m, \beta_{|m|+1}^m, \dots, \beta_{2l-2}^m, \beta_{2l-1}^m$ from the values of a sum over k in (5) at $\theta = \theta_0, \theta_1, \dots, \theta_{2l-2}, \theta_{2l-1}$. FFTs or fast discrete sine and cosine transforms can be used to handle the sums over m in (5) efficiently. See [12] for a detailed summary and novel application of the overall method. The present article modifies portions of the method of [12,15], focusing exclusively on the modifications.

3. Preliminaries

In this section, we summarize certain facts from mathematical and numerical analysis, used in Sections 4 and 5. Section 3.1 describes interpolative decompositions (IDs). Section 3.2 outlines the butterfly algorithm. Section 3.3 summarizes basic properties of normalized associated Legendre functions.

3.1. Interpolative decompositions

In this subsection, we define interpolative decompositions (IDs) and summarize their properties.

The following lemma states that, for any $m \times n$ matrix A of rank k , there exist an $m \times k$ matrix $A^{(k)}$ whose columns constitute a subset of the columns of A , and a $k \times n$ matrix \tilde{A} , such that

1. some subset of the columns of \tilde{A} makes up the $k \times k$ identity matrix,
2. \tilde{A} is not too large, and
3. $A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} = A_{m \times n}$.

Moreover, the lemma provides an approximation

$$A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} \approx A_{m \times n} \tag{6}$$

when the exact rank of A is greater than k , yet the $(k + 1)$ st greatest singular value of A is still small. The lemma is a reformulation of Theorem 3.2 in [9] and Theorem 3 in [5]; its proof is based on techniques described in [6,8,16]. We will refer to the approximation in (6) of A as an interpolative decomposition (ID). We call \tilde{A} the “interpolation matrix” of the ID.

Lemma 3.1. *Suppose that m and n are positive integers, and A is a real $m \times n$ matrix.*

Then, for any positive integer k with $k \leq m$ and $k \leq n$, there exist a real $k \times n$ matrix \tilde{A} , and a real $m \times k$ matrix $A^{(k)}$ whose columns constitute a subset of the columns of A , such that

1. some subset of the columns of \tilde{A} makes up the $k \times k$ identity matrix,
2. no entry of \tilde{A} has an absolute value greater than 1,
3. the spectral norm (that is, the l^2 -operator norm) of \tilde{A} satisfies $\|\tilde{A}_{k \times n}\|_2 \leq \sqrt{k(n - k) + 1}$,
4. the least (that is, the k th greatest) singular value of \tilde{A} is at least 1,
5. $A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} = A_{m \times n}$ when $k = m$ or $k = n$, and
6. when $k < m$ and $k < n$, the spectral norm (that is, the l^2 -operator norm) of $A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} - A_{m \times n}$ satisfies

$$\|A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} - A_{m \times n}\|_2 \leq \sqrt{k(n - k) + 1} \sigma_{k+1}, \tag{7}$$

where σ_{k+1} is the $(k + 1)$ st greatest singular value of A .

Properties 1, 2, 3, and 4 in Lemma 3.1 ensure that the ID $A^{(k)} \cdot \tilde{A}$ of A is a numerically stable representation. Also, Property 3 follows directly from Properties 1 and 2, and Property 4 follows directly from Property 1.

Remark 3.2. Existing algorithms for the computation of the matrices $A^{(k)}$ and \tilde{A} in Lemma 3.1 are computationally expensive. We use instead the algorithm of [5,8] to produce matrices $A^{(k)}$ and \tilde{A} which satisfy slightly weaker conditions than those in Lemma 3.1. We compute $A^{(k)}$ and \tilde{A} such that

1. some subset of the columns of \tilde{A} makes up the $k \times k$ identity matrix,
2. no entry of \tilde{A} has an absolute value greater than 2,
3. the spectral norm (that is, the l^2 -operator norm) of \tilde{A} satisfies $\|\tilde{A}_{k \times n}\|_2 \leq \sqrt{4k(n - k) + 1}$,
4. the least (that is, the k th greatest) singular value of \tilde{A} is at least 1,
5. $A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} = A_{m \times n}$ when $k = m$ or $k = n$, and
6. when $k < m$ and $k < n$, the spectral norm (that is, the l^2 -operator norm) of $A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} - A_{m \times n}$ satisfies

$$\|A_{m \times k}^{(k)} \cdot \tilde{A}_{k \times n} - A_{m \times n}\|_2 \leq \sqrt{4k(n - k) + 1} \sigma_{k+1}, \tag{8}$$

where σ_{k+1} is the $(k + 1)$ st greatest singular value of A .

For any positive real number ε , the algorithm can identify the least k such that $\|A^{(k)} \cdot \tilde{A} - A\|_2 \approx \varepsilon$. Furthermore, the algorithm computes both $A^{(k)}$ and \tilde{A} using at most

$$C_{ID} = \mathcal{O}(kmn \log(n)) \tag{9}$$

floating-point operations, typically requiring only

$$C'_{\text{ID}} = \mathcal{O}(kmn). \quad (10)$$

3.2. The butterfly algorithm

In this subsection, we outline a simple case of the butterfly algorithm from [10,11]; see [11] for a detailed description.

Suppose that n is a positive integer, and A is an $n \times n$ matrix. Suppose further that ε and C are positive real numbers, and k is a positive integer, such that any contiguous rectangular subblock of A containing at most Cn entries can be approximated to precision ε by a matrix whose rank is k (using the Frobenius/Hilbert–Schmidt norm to measure the accuracy of the approximation); we will refer to this hypothesis as “the rank property”. The running-time of the algorithm will be proportional to k^2/C ; taking C to be roughly proportional to k suffices for many matrices of interest (including nonequispaced and discrete Fourier transforms), so ideally k should be small. We will say that two matrices G and H are equal to precision ε , denoted $G \approx H$, to mean that the spectral norm (that is, the l^2 -operator norm) of $G - H$ is $\mathcal{O}(\varepsilon)$.

We now explicitly use the rank property for subblocks of multiple heights, to illustrate the basic structure of the butterfly scheme.

Consider any two adjacent contiguous rectangular subblocks L and R of A , each containing at most Cn entries and having the same numbers of rows, with L on the left and R on the right. Due to the rank property, there exist IDs

$$L \approx L^{(k)} \cdot \tilde{L} \quad (11)$$

and

$$R \approx R^{(k)} \cdot \tilde{R}, \quad (12)$$

where $L^{(k)}$ is a matrix having k columns, which constitute a subset of the columns of L , $R^{(k)}$ is a matrix having k columns, which constitute a subset of the columns of R , \tilde{L} and \tilde{R} are matrices each having k rows, and all entries of \tilde{L} and \tilde{R} have absolute values of at most 2.

To set notation, we concatenate the matrices L and R , and split the columns of the result in half (or approximately in half), obtaining T on top and B on the bottom:

$$(L|R) = \begin{pmatrix} T \\ B \end{pmatrix}. \quad (13)$$

Observe that the matrices T and B each have at most Cn entries (since L and R each have at most Cn entries). Similarly, we concatenate the matrices $L^{(k)}$ and $R^{(k)}$, and split the columns of the result in half (or approximately in half), obtaining $T^{(2k)}$ and $B^{(2k)}$:

$$(L^{(k)}|R^{(k)}) = \begin{pmatrix} T^{(2k)} \\ B^{(2k)} \end{pmatrix}. \quad (14)$$

Observe that the $2k$ columns of $T^{(2k)}$ are also columns of T , and that the $2k$ columns of $B^{(2k)}$ are also columns of B . Due to the rank property, there exist IDs

$$T^{(2k)} \approx T^{(k)} \cdot \widetilde{T^{(2k)}} \quad (15)$$

and

$$B^{(2k)} \approx B^{(k)} \cdot \widetilde{B^{(2k)}}, \quad (16)$$

where $T^{(k)}$ is a matrix having k columns, which constitute a subset of the columns of $T^{(2k)}$, $B^{(k)}$ is a matrix having k columns, which constitute a subset of the columns of $B^{(2k)}$, $\widetilde{T^{(2k)}}$ and $\widetilde{B^{(2k)}}$ are matrices each having k rows, and all entries of $\widetilde{T^{(2k)}}$ and $\widetilde{B^{(2k)}}$ have absolute values of at most 2.

Combining (11)–(16) yields that

$$T \approx T^{(k)} \cdot \widetilde{T^{(2k)}} \cdot \begin{pmatrix} \tilde{L} & \mathbf{0} \\ \mathbf{0} & \tilde{R} \end{pmatrix} \quad (17)$$

and

$$B \approx B^{(k)} \cdot \widetilde{B^{(2k)}} \cdot \begin{pmatrix} \tilde{L} & \mathbf{0} \\ \mathbf{0} & \tilde{R} \end{pmatrix}. \quad (18)$$

If we use m to denote the number of rows in L (which is the same as the number of rows in R), then the number of columns in L (or R) is at most Cn/m , and so the total number of entries in the matrices in the right-hand sides of (11) and (12) can be as large as $2mk + 2k(Cn/m)$, whereas the total number of nonzero entries in the matrices in the right-hand sides of (17) and (18) is at most $mk + 4k^2 + 2k(Cn/m)$. If m is nearly as large as possible — nearly n — and k and C are much smaller than n ,

then $mk + 4k^2 + 2k(Cn/m)$ is about half $(2mk + 2k(Cn/m))$. Thus, the representation provided in (17) and (18) of the merged matrix from (13) is more efficient than that provided in (11) and (12), both in terms of the memory required for storage, and in terms of the number of operations required for applications to vectors. Notice the advantage of using the rank property for blocks of multiple heights.

Naturally, we may repeat this process of merging adjacent blocks and splitting in half the columns of the result, updating the compressed representations after every split. We start by partitioning A into blocks each dimensioned $n \times \lfloor C \rfloor$ (except possibly for the rightmost block, which may have fewer than $\lfloor C \rfloor$ columns), and then repeatedly group unprocessed blocks (of whatever dimensions) into disjoint pairs, processing these pairs by merging and splitting them into new, unprocessed blocks having fewer rows. The resulting multilevel representation of A allows us to apply A with precision ε from the left to any column vector, or from the right to any row vector, using just $\mathcal{O}((k^2/C)n \log(n))$ floating-point operations (there will be $\mathcal{O}(\log(n))$ levels in the scheme, and each level except for the last will only involve $\mathcal{O}(n/C)$ interpolation matrices of dimensions $k \times (2k)$, such as $\widetilde{T}^{(2k)}$ and $\widetilde{B}^{(2k)}$). Fig. 1 illustrates the resulting partitionings of A into blocks of various dimensions (but with every block having the same number of entries), when $n = 8$ and $C = 1$. For further details, see [11].

Remark 3.3. Needless to say, the same multilevel representation of A permits the rapid application of A both from the left to column vectors and from the right to row vectors. There is no need for constructing multilevel representations of both A and the transpose of A .

Remark 3.4. In practice, the IDs used for accurately approximating subblocks of A do not all have the same fixed rank k . Instead, for each subblock, we determine the minimal possible rank such that the associated ID still approximates the subblock to precision ε , and we use this ID in place of one whose rank is k . Determining ranks adaptively in this manner accelerates the algorithm substantially. For further details, see [11]. All our implementations use this adaptation.

3.3. Normalized associated Legendre functions

In this subsection, we discuss several classical facts concerning normalized associated Legendre functions. All of these facts follow trivially from results contained, for example, in [1] or [14].

For any nonnegative integers l and m such that $l \geq m$, we use \overline{P}_l^m to denote the normalized associated Legendre function of degree l and order m , defined on $(-1, 1)$ via the formula

$$\overline{P}_l^m(x) = \sqrt{\frac{2l+1}{2} \frac{(l-m)!}{(l+m)!}} (1-x^2)^{m/2} \frac{d^m}{dx^m} P_l(x), \tag{19}$$

where P_l is the Legendre polynomial of degree l ,

$$P_l(x) = \frac{1}{2^l l!} \frac{d^l}{dx^l} ((x^2 - 1)^l) \tag{20}$$

(see, for example, Chapter 8 of [1]). “Normalized” refers to the fact that the normalized associated Legendre functions of a fixed order m are orthonormal on $(-1, 1)$ with respect to the standard inner product. If $l - m$ is even, then $\overline{P}_l^m(-x) = \overline{P}_l^m(x)$ for any $x \in (-1, 1)$. If $l - m$ is odd, then $\overline{P}_l^m(-x) = -\overline{P}_l^m(x)$ for any $x \in (-1, 1)$.


The following lemma states that the normalized associated Legendre functions satisfy a certain self-adjoint second-order linear (Sturm–Liouville) differential equation.

Lemma 3.5. *Suppose that m is a nonnegative integer.*

Then,

$$-\frac{d}{dx} \left((1-x^2) \frac{d}{dx} \overline{P}_l^m(x) \right) + \left(\frac{m^2}{1-x^2} - l(l+1) \right) \overline{P}_l^m(x) = 0 \tag{21}$$

for any $x \in (-1, 1)$, and $l = m, m + 1, m + 2, \dots$



L indicates the left member of a pair; R indicates the right member.
 T indicates the top member of a pair; B indicates the bottom member.

Fig. 1. The partitionings in the multilevel decomposition for an 8×8 matrix with $C = 1$.

The following lemma states that the normalized associated Legendre function of order m and degree $m + 2n$ has exactly n zeros inside $(0, 1)$, and, moreover, that the normalized associated Legendre function of order m and degree $m + 2n + 1$ also has exactly n zeros inside $(0, 1)$.

Lemma 3.6. *Suppose that m and n are nonnegative integers with $n > 0$.*

Then, there exist precisely n real numbers $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ such that

$$0 < x_0 < x_1 < \dots < x_{n-2} < x_{n-1} < 1 \quad (22)$$

and

$$\bar{P}_{m+2n}^m(x_j) = 0 \quad (23)$$

for $j = 0, 1, \dots, n-2, n-1$.

Moreover, there exist precisely n real numbers $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ such that

$$0 < y_0 < y_1 < \dots < y_{n-2} < y_{n-1} < 1 \quad (24)$$

and

$$\bar{P}_{m+2n+1}^m(y_j) = 0 \quad (25)$$

for $j = 0, 1, \dots, n-2, n-1$.

Suppose that m and n are nonnegative integers with $n > 0$. Then, we define real numbers $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}, \sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$, and σ_n via the formulae

$$\rho_j = \frac{2(2m + 4n + 1)}{(1 - (x_j)^2) \left(\frac{d}{dx} \bar{P}_{m+2n}^m(x_j) \right)^2} \quad (26)$$

for $j = 0, 1, \dots, n-2, n-1$, where $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ are from (23),

$$\sigma_j = \frac{2(2m + 4n + 3)}{(1 - (y_j)^2) \left(\frac{d}{dx} \bar{P}_{m+2n+1}^m(y_j) \right)^2} \quad (27)$$

for $j = 0, 1, \dots, n-2, n-1$, where $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ are from (25), and

$$\sigma_n = \frac{2m + 4n + 3}{\left(\frac{d}{dx} \bar{P}_{m+2n+1}^m(0) \right)^2}. \quad (28)$$

The following lemma describes what are known as Gauss–Jacobi quadrature formulae corresponding to associated Legendre functions.

Lemma 3.7. *Suppose that m and n are nonnegative integers with $n > 0$.*

Then,

$$\int_{-1}^1 dx (1 - x^2)^m p(x) = \sum_{j=0}^{n-1} \rho_j (1 - (x_j)^2)^m p(x_j) \quad (29)$$

for any even polynomial p of degree at most $4n - 2$, where $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ are from (23), and $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ are defined in (26).

Furthermore,

$$\int_{-1}^1 dx (1 - x^2)^m p(x) = \sigma_n p(0) + \sum_{j=0}^{n-1} \sigma_j (1 - (y_j)^2)^m p(y_j) \quad (30)$$

for any even polynomial p of degree at most $4n$, where $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ are from (25), and $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_n$ are defined in (27) and (28).

Remark 3.8. Formulae (35) and (36) of [15] incorrectly omitted the factors $(1 - (x_j)^2)^m$ and $(1 - (y_j)^2)^m$ appearing in the analogous (29) and (30) above.

Suppose that m is a nonnegative integer. Then, we define real numbers $c_m, c_{m+1}, c_{m+2}, \dots$ and $d_m, d_{m+1}, d_{m+2}, \dots$ via the formulae

$$c_l = \sqrt{\frac{(l - m + 1)(l - m + 2)(l + m + 1)(l + m + 2)}{(2l + 1)(2l + 3)^2(2l + 5)}} \quad (31)$$

for $l = m, m + 1, m + 2, \dots$, and

$$d_l = \frac{2l(l+1) - 2m^2 - 1}{(2l-1)(2l+3)} \quad (32)$$

for $l = m, m + 1, m + 2, \dots$

The following lemma states that the normalized associated Legendre functions of a fixed order m satisfy a certain three-term recurrence relation.

Lemma 3.9. *Suppose that m is a nonnegative integer.*

Then,

$$x^2 \bar{P}_l^m(x) = d_l \bar{P}_l^m(x) + c_l \bar{P}_{l+2}^m(x) \quad (33)$$

for any $x \in (-1, 1)$, and $l = m$ or $l = m + 1$, and

$$x^2 \bar{P}_l^m(x) = c_{l-2} \bar{P}_{l-2}^m(x) + d_l \bar{P}_l^m(x) + c_l \bar{P}_{l+2}^m(x) \quad (34)$$

for any $x \in (-1, 1)$, and $l = m + 2, m + 3, m + 4, \dots$, where $c_m, c_{m+1}, c_{m+2}, \dots$ are defined in (31), and $d_m, d_{m+1}, d_{m+2}, \dots$ are defined in (32).

4. Precomputations for the butterfly scheme

In this section, we discuss the preprocessing required for the butterfly algorithm summarized in Section 3.2. We will be using the notation detailed in Section 3.2.

Perhaps the most natural organization of the computations required to construct the multilevel representation of an $n \times n$ matrix A is first to process all blocks having n rows (Level 1 in Fig. 1 above), then to process all blocks having about $n/2$ rows (Level 2 in Fig. 1), then to process all blocks having about $n/4$ rows (Level 3 in Fig. 1), and so on. Indeed, [11] uses this organization, which amounts to a “breadth-first traversal” of the control-flow graph for the program applying A to a vector (see, for example, [3] for an introduction to “breadth-first” and “depth-first” orderings). This scheme for preprocessing is efficient when the entries of A can be efficiently computed on-the-fly, individually. (Of course, we are assuming that A has a suitable rank property, that is, that there are positive real numbers ε and C , and a positive integer k , such that any contiguous rectangular subblock of A containing at most Cn entries can be approximated to precision ε by a matrix whose rank is k , using the Frobenius/Hilbert–Schmidt norm to measure the accuracy of the approximation. Often, taking C to be roughly proportional to k suffices, and ideally k and ε are small.) If the entries of A cannot be efficiently computed individually, however, then the “breadth-first traversal” may need to store $\mathcal{O}(n^2)$ entries at some point during the precomputations, in order to avoid recomputing entries of the matrix.

If individual columns of A (but not necessarily arbitrary individual entries) can be computed efficiently, then “depth-first traversal” of the control-flow graph requires only $\mathcal{O}((k^2/C)n \log(n))$ floating-point words of memory at any point during the precomputations, for the following reason. We will say that we “process” a block of A to mean that we merge it with another, and split and recompress the result, producing a pair of new, unprocessed blocks. Rather than starting the preprocessing by constructing all blocks having n rows, we construct each such block only after processing as many blocks as possible which previous processing creates, but which have not yet been processed. Furthermore, we construct each block having n rows only after having already constructed (and possibly processed) all blocks to its left. To reiterate, we construct a block having n rows only after having exhausted all possibilities for both creating and processing blocks to its left.

For each *processed* block B , we need only store the interpolation matrix $\widetilde{B}^{(2k)}$ and the indices of the columns chosen for the ID; we need not store the k columns of $B^{(k)}$ selected for the ID, since the algorithm for applying A (or its transpose) to a vector never explicitly uses any columns of a block that has been merged with another and split, but instead interpolates from (or anterpolates to) the shorter blocks arising from the processing. Conveniently, the matrix $\widetilde{B}^{(2k)}$ that we must store is small – no larger than $k \times (2k)$. For each *unprocessed* block B , we do need to store the k columns in $B^{(k)}$ selected for the ID, in addition to storing $\widetilde{B}^{(2k)}$ and the indices of the columns chosen for the ID, facilitating any subsequent processing. Although $B^{(k)}$ may have many rows, it has only as many rows as B and hence is smaller when B has fewer rows. Thus, every time we process a pair of tall blocks, producing a new pair of blocks having half as many rows, the storage requirements for all these blocks together nearly halve. By always processing as many already constructed blocks as possible, we minimize the amount of memory required.

5. Spherical harmonic transforms via the butterfly scheme

In this section, we describe how to use the butterfly algorithm to compute fast spherical harmonic transforms, via appropriate modifications of the algorithm of [15].

We substitute the butterfly algorithm for the divide-and-conquer algorithm of [7] used in Section 3.1 of [15], otherwise leaving the approach of [15] unchanged. Specifically, given numbers $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$, we use the butterfly scheme to compute the numbers $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ defined via the formula

$$\alpha_i = \sum_{j=0}^{n-1} \beta_j \sqrt{\rho_i} \bar{P}_{m+2j}^m(x_i) \quad (35)$$

for $i = 0, 1, \dots, n-2, n-1$, where m is a nonnegative integer, $\bar{P}_m^m, \bar{P}_{m+2}^m, \dots, \bar{P}_{m+2n-2}^m, \bar{P}_{m+2n}^m$ are the normalized associated Legendre functions of order m defined in (19), $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ are the positive zeros of \bar{P}_{m+2n}^m from (23), and $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ are the corresponding quadrature weights from (29). Similarly, given numbers $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$, we use the butterfly scheme to compute the numbers $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ satisfying (35). The factors $\sqrt{\rho_0}, \sqrt{\rho_1}, \dots, \sqrt{\rho_{n-2}}, \sqrt{\rho_{n-1}}$ ensure that the linear transformation mapping $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ to $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ via (35) is unitary (due to (19), (29), and the orthonormality of the normalized associated Legendre functions on $(-1, 1)$), so that the inverse of the linear transformation is its transpose.

Moreover, given numbers $v_0, v_1, \dots, v_{n-2}, v_{n-1}$, we use the butterfly scheme to compute the numbers $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ defined via the formula

$$\mu_i = \sum_{j=0}^{n-1} v_j \sqrt{\sigma_i} \bar{P}_{m+2j+1}^m(y_i) \quad (36)$$

for $i = 0, 1, \dots, n-2, n-1$, where m is a nonnegative integer, $\bar{P}_{m+1}^m, \bar{P}_{m+3}^m, \dots, \bar{P}_{m+2n-1}^m, \bar{P}_{m+2n+1}^m$ are the normalized associated Legendre functions of order m defined in (19), $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ are the positive zeros of \bar{P}_{m+2n+1}^m from (25), and $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$ are the corresponding quadrature weights from (30). Similarly, given numbers $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$, we use the butterfly scheme to compute the numbers $v_0, v_1, \dots, v_{n-2}, v_{n-1}$ satisfying (36). As above, the factors $\sqrt{\sigma_0}, \sqrt{\sigma_1}, \dots, \sqrt{\sigma_{n-2}}, \sqrt{\sigma_{n-1}}$ ensure that the linear transformation mapping $v_0, v_1, \dots, v_{n-2}, v_{n-1}$ to $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ via (36) is unitary, so that its inverse is its transpose.

Computing spherical harmonic transforms requires several additional computations, detailed in [15]. (See also Remark 5.1 below.) The butterfly algorithm replaces only the procedure described in Section 3.1 of [15].

In order to use (35) and (36) numerically, we need to precompute the positive zeros $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ of \bar{P}_{m+2n}^m from (23), the corresponding quadrature weights $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ from (29), the positive zeros $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ of \bar{P}_{m+2n+1}^m from (25), and the corresponding quadrature weights $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$ from (30). Section 3.3 of [15] describes suitable procedures (based on integrating the ordinary differential equation in (21) in “Prüfer coordinates”). We found it expedient to perform this preprocessing in extended-precision arithmetic, in order to compensate for the loss of a couple of digits of accuracy relative to the machine precision.

To perform the precomputations described in Section 4 above associated with (35) and (36), we need to be able to evaluate efficiently all n functions $\bar{P}_m^m, \bar{P}_{m+2}^m, \dots, \bar{P}_{m+2n-4}^m, \bar{P}_{m+2n-2}^m$ at any of the precomputed positive zeros $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ of \bar{P}_{m+2n}^m from (23), and, similarly, we need to be able to evaluate efficiently all n functions $\bar{P}_{m+1}^m, \bar{P}_{m+3}^m, \dots, \bar{P}_{m+2n-3}^m, \bar{P}_{m+2n-1}^m$ at any of the precomputed positive zeros $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ of \bar{P}_{m+2n+1}^m from (25). For this, we may use the recurrence relations (33) and (34), starting with the values of $\bar{P}_m^m(x_i), \bar{P}_{m+1}^m(y_i), \bar{P}_{m+2}^m(x_i)$, and $\bar{P}_{m+3}^m(y_i)$ obtained via (19). (We can counter underflow by tracking exponents explicitly, in the standard fashion.) Such use of the recurrence is a classic procedure; see, for example, Chapter 8 of [1]. The recurrence appears to be numerically stable when used for evaluating normalized associated Legendre functions of order m and of degrees at most $m+2n-1$, at these special points $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ and $y_0, y_1, \dots, y_{n-2}, y_{n-1}$, even when n is very large. We did not need to use extended-precision arithmetic for this preprocessing.

Remark 5.1. The formula (88) in [15] that is analogous to (35) of the present paper omits the factors $\sqrt{\rho_0}, \sqrt{\rho_1}, \dots, \sqrt{\rho_{n-2}}, \sqrt{\rho_{n-1}}$ included in (35). Obviously, the vectors

$$(\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1})^\top \quad (37)$$

and

$$\left(\frac{\alpha_0}{\sqrt{\rho_0}}, \frac{\alpha_1}{\sqrt{\rho_1}}, \dots, \frac{\alpha_{n-2}}{\sqrt{\rho_{n-2}}}, \frac{\alpha_{n-1}}{\sqrt{\rho_{n-1}}} \right)^\top \quad (38)$$

differ by a diagonal transformation, and so we can obtain either one from the other efficiently. In fact, the well-conditioned matrix A from Section 3.1 of [15] represents the same diagonal transformation, mapping (37) to (38). Similar remarks apply to (36), of course.

6. Numerical results

In this section, we describe the results of several numerical tests of the algorithm of the present paper. (Computing spherical harmonic transforms requires several additional computations, detailed in [15] — see Section 5 for further information. The butterfly algorithm replaces only the procedure described in Section 3.1 of [15].)

Tables 1–4 report the results of computing from real numbers $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ the real numbers $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ defined by the formula

$$\alpha_i = \sum_{j=0}^{n-1} \beta_j \sqrt{\rho_j} \bar{P}_{m+2j}^m(x_i) \tag{39}$$

for $i = 0, 1, \dots, n-2, n-1$, where $\bar{P}_m^m, \bar{P}_{m+2}^m, \dots, \bar{P}_{m+2n-2}^m, \bar{P}_{m+2n}^m$ are the normalized associated Legendre functions defined in (19), $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ are the positive zeros of \bar{P}_{m+2n}^m from (23), and $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ are the corresponding quadrature weights from (29). We will refer to the map via (39) from $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ to $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ as the forward transform, and the map via (39) from $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ to $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ as the inverse transform (the inverse is also the transpose, due to (19), (29), and the orthonormality of the normalized associated Legendre functions on $(-1, 1)$). The values of m differ in Tables 1, 2 and Tables 3, 4.

Tables 5 and 6 report the results of computing from real numbers $v_0, v_1, \dots, v_{n-2}, v_{n-1}$ the real numbers $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ defined by the formula

$$\mu_i = \sum_{j=0}^{n-1} v_j \sqrt{\sigma_j} \bar{P}_{m+2j+1}^m(y_i) \tag{40}$$

for $i = 0, 1, \dots, n-2, n-1$, where $\bar{P}_{m+1}^m, \bar{P}_{m+3}^m, \dots, \bar{P}_{m+2n-1}^m, \bar{P}_{m+2n+1}^m$ are the normalized associated Legendre functions defined in (19), $y_0, y_1, \dots, y_{n-2}, y_{n-1}$ are the positive zeros of \bar{P}_{m+2n+1}^m from (25), and $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$ are the corresponding quadrature weights from (30). We will refer to the map via (40) from $v_0, v_1, \dots, v_{n-2}, v_{n-1}$ to $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ as the forward transform, and the map via (40) from $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ to $v_0, v_1, \dots, v_{n-2}, v_{n-1}$ as the inverse transform (as above, the inverse is also the transpose).

For the test vectors β and v whose entries appear in (39) and (40), we used normalized vectors whose entries were pseudorandom numbers drawn uniformly from $(-1, 1)$, normalized so that the sum of the squares of the entries is 1.

As described in Remark 3.4, we compute for each block in the multilevel representation of A an ID whose rank is as small as possible while still approximating the block to nearly machine precision.

The following list describes the headings of the tables:

Table 1
Ranks and running-times for even degrees.

n	m	k_{\max}	k_{avg}	k_{σ}	t_{dir}	t_{fwd}	t_{inv}
1250	1250	170	65.3	29.5	.25E-2	.18E-2	.15E-2
2500	2500	168	67.0	32.6	.98E-2	.46E-2	.38E-2
5000	5000	195	70.5	35.9	.39E-1	.12E-1	.96E-2
10000	10000	247	73.5	38.7	.15E-0	.28E-1	.23E-1
20000	20000	308	75.9	41.5	(.60E-0)	.67E-1	.55E-1
40000	40000	379	78.0	43.8	(.24E+1)	.16E-0	.13E-0

Table 2
Precomputation times, memory requirements, and accuracies for even degrees.

n	m	t_{quad}	t_{comp}	m_{\max}	ϵ_{fwd}	ϵ_{inv}
1250	1250	.46E2	.96E0	.86E6	.62E-14	.19E-13
2500	2500	.92E2	.41E1	.20E7	.37E-14	.25E-13
5000	5000	.18E3	.17E2	.50E7	.59E-14	.43E-13
10000	10000	.37E3	.82E2	.14E8	.32E-14	.57E-13
20000	20000	.74E3	.39E3	.29E8	.30E-14	.88E-13
40000	40000	.15E4	.17E4	.64E8	.24E-14	.13E-12

Table 3
Ranks and running-times for even degrees.

n	m	k_{\max}	k_{avg}	k_{σ}	t_{dir}	t_{fwd}	t_{inv}
1250	0	110	67.0	23.3	.25E-2	.18E-2	.15E-2
2500	0	110	70.0	25.0	.98E-2	.46E-2	.40E-2
5000	0	111	73.9	26.1	.39E-1	.12E-1	.10E-1
10000	0	111	77.3	26.8	.15E-0	.29E-1	.24E-1
20000	0	112	80.2	27.1	(.60E-0)	.68E-1	.57E-1
40000	0	169	82.7	27.4	(.24E+1)	.16E-0	.13E-0

Table 4

Precomputation times, memory requirements, and accuracies for even degrees.

n	m	t_{quad}	t_{comp}	m_{max}	ε_{fwd}	ε_{inv}
1250	0	.44E2	.96E0	.86E6	.49E–14	.12E–12
2500	0	.88E2	.41E1	.20E7	.35E–14	.14E–12
5000	0	.18E3	.18E2	.51E7	.23E–14	.35E–12
10000	0	.36E3	.82E2	.14E8	.18E–14	.63E–12
20000	0	.74E3	.40E3	.29E8	.20E–14	.22E–11
40000	0	.14E4	.18E4	.66E8	.16E–14	.37E–11

Table 5

Ranks and running-times for odd degrees.

n	m	k_{max}	k_{avg}	k_{σ}	t_{dir}	t_{fwd}	t_{inv}
1250	1250	170	65.3	29.5	.25E–2	.17E–2	.15E–2
2500	2500	169	67.0	32.6	.98E–2	.48E–2	.39E–2
5000	5000	196	70.5	35.9	.39E–1	.12E–1	.97E–2
10000	10000	247	73.5	38.7	.15E–0	.28E–1	.24E–1
20000	20000	308	75.9	41.4	(.60E–0)	.68E–1	.56E–1
40000	40000	379	78.0	43.8	(.24E+1)	.16E–0	.13E–0

Table 6

Precomputation times, memory requirements, and accuracies for odd degrees.

n	m	t_{quad}	t_{comp}	m_{max}	ε_{fwd}	ε_{inv}
1250	1250	.46E2	.94E0	.86E6	.41E–14	.19E–13
2500	2500	.92E2	.40E1	.20E7	.41E–14	.29E–13
5000	5000	.18E3	.17E2	.50E7	.40E–14	.51E–13
10000	10000	.37E3	.80E2	.14E8	.31E–14	.62E–13
20000	20000	.74E3	.39E3	.29E8	.34E–14	.10E–12
40000	40000	.15E4	.18E4	.64E8	.25E–14	.14E–12

- n is the size of the transform, the size of the vectors α and β whose entries are given in (39), and of the vectors μ and ν whose entries are given in (40).
- m is the order of the normalized associated Legendre functions used in (39) and (40).
- k_{max} is the maximum of the ranks of the IDs for the blocks in the multilevel representation.
- k_{avg} is the average of the ranks of the IDs for the blocks in the multilevel representation.
- k_{σ} is the standard deviation of the ranks of the IDs for the blocks in the multilevel representation.
- t_{dir} is the time in seconds required to apply an $n \times n$ matrix to an $n \times 1$ vector using the standard procedure. We estimated the last two entries for t_{dir} by multiplying the third-to-last entry in each table by 4 and 16, since the large matrices required to generate those entries cannot fit in the available 2 GB of RAM. We indicate that these entries are estimates by enclosing them in parentheses.
- t_{fwd} is the time in seconds required by the butterfly algorithm to compute the forward transform via (39) or (40), mapping from $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$ to $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$, or from $\nu_0, \nu_1, \dots, \nu_{n-2}, \nu_{n-1}$ to $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$.
- t_{inv} is the time in seconds required by the butterfly algorithm to compute the inverse transform via (39) or (40), mapping from $\alpha_0, \alpha_1, \dots, \alpha_{n-2}, \alpha_{n-1}$ to $\beta_0, \beta_1, \dots, \beta_{n-2}, \beta_{n-1}$, or from $\mu_0, \mu_1, \dots, \mu_{n-2}, \mu_{n-1}$ to $\nu_0, \nu_1, \dots, \nu_{n-2}, \nu_{n-1}$.
- t_{quad} is the time in seconds required in the precomputations to compute the quadrature nodes $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ or $y_0, y_1, \dots, y_{n-2}, y_{n-1}$, and weights $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ or $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$, from (29) or (30), used in (39) and (40).
- t_{comp} is the time in seconds required to construct the compressed multilevel representation used in the butterfly algorithm, after having already computed the quadrature nodes $x_0, x_1, \dots, x_{n-2}, x_{n-1}$ or $y_0, y_1, \dots, y_{n-2}, y_{n-1}$, and weights $\rho_0, \rho_1, \dots, \rho_{n-2}, \rho_{n-1}$ or $\sigma_0, \sigma_1, \dots, \sigma_{n-2}, \sigma_{n-1}$, from (29) or (30), used in (39) and (40).
- m_{max} is the maximum number of floating-point words of memory required to store entries of the transform matrix during any point in the precomputations (all other memory requirements are negligible in comparison).
- ε_{fwd} is the maximum difference between the entries in the result of the forward transform computed via the butterfly algorithm and those computed directly via the standard procedure for applying a matrix to a vector. (The result of the forward transform is the vector α whose entries are given in (39) or the vector μ whose entries are given in (40).)
- ε_{inv} is the maximum difference between the entries in a test vector and the entries in the result of applying to the test vector first the forward transform and then the inverse transform, both computed via the butterfly algorithm. (The result of the forward transform is the vector α whose entries are given in (39) or the vector μ whose entries are given in (40). The

result of the inverse transform is the vector β whose entries are given in (39) or the vector v whose entries are given in (40.) Thus, ε_{inv} measures the accuracy of the butterfly algorithm without reference to the standard procedure for applying a matrix to a vector (unlike ε_{fwd}).

For the first level of the multilevel representation of the $n \times n$ matrix, we partitioned the matrix into blocks each dimensioned $n \times 60$ (except for the rightmost block, since n is not divisible by 60). Every block on every level has about the same number of entries (specifically, $60n$ entries). We wrote all code in Fortran 77, compiling it using the Lahey–Fujitsu Linux Express v6.2 compiler, with optimization flag `-o2` enabled. We ran all examples on one core of a 2.7 GHz Intel Core 2 Duo microprocessor with 3 MB of L2 cache and 2 GB of RAM. As described in Section 5, we used extended-precision arithmetic during the portion of the preprocessing requiring integration of an ordinary differential equation, to compute quadrature nodes and weights (this is not necessary to attain high accuracy, but does yield a couple of extra digits of precision). Otherwise, our code is compliant with the IEEE double-precision standard (so that the mantissas of variables have approximately one bit of precision less than 16 digits, yielding a relative precision of about $.2E-15$).

Remark 6.1. Tables 1, 3, and 5 indicate that the linear transformations in (39) and (40) satisfy the rank property discussed in Section 3.2, with arbitrarily high precision, at least in some averaged sense. Furthermore, it appears that the parameter k discussed in Section 3.2 can be set to be independent of the order m and size n of the transforms in (39) and (40), with the parameter C discussed in Section 3.2 roughly proportional to k . The acceleration provided by the butterfly algorithm thus is sufficient for computing fast spherical harmonic transforms, and is competitive with the approach taken in [15] (though much work remains in optimizing both approaches in order to gauge their relative performance). Unlike the approach taken in [15], the approach of the present paper does not require the use of extended-precision arithmetic during the precomputations in order to attain accuracy close to the machine precision, even while accelerating spherical harmonic transforms about as well. Moreover, the butterfly can be easier to implement.

Remark 6.2. The values in Tables 2, 4, and 6 vary with the size n of the transforms in (39) and (40) as expected. The values for t_{quad} are consistent with the expected values of a constant times n . The values for t_{comp} are consistent with the expected values of a constant times n^2 (with the constant being proportional to k_{avg}). The values for m_{max} are consistent with the expected values of a constant times $n \log(n)$ (again with the constant being proportional to k_{avg}); these modest memory requirements make the preprocessing feasible for large values of n such as those in the tables.

Remark 6.3. In the current technological environment, neither the scheme of [15] nor the approach of the present paper is uniformly superior to the other. For example, the theory from [15] is rigorous and essentially complete, while the theory of the present article ideally should undergo further development, to prove that the rank properties discussed in Remark 6.1 are as strong as numerical experiments indicate, yielding the desired acceleration. In contrast, to attain accuracy close to the machine precision, the approach of [15] requires the use of extended-precision arithmetic during its precomputations, whereas the scheme of the present paper does not. Implementing the procedure of the present article can be easier. Finally, an anonymous referee kindly compared the running-times of the implementations reported in [15] and the present paper, noticing that the newer computer system used in the present article is about 2.7 times faster than the old system; the algorithms of [15] and of the present article are roughly equally efficient — certainly neither appears to be more than twice faster than the other. However, both implementations are rather crude, and could undoubtedly benefit from further optimization by experts on computer architectures; also, we made no serious attempt to optimize the precomputations. Furthermore, with the advent of multicore and distributed processors, coming changes in computer architectures might affect the two approaches differently, as they may parallelize and utilize cache in different ways. In the end, the use of one approach rather than the other may be a matter of convenience, as the two methods are yielding similar performance.

7. Conclusions

This article provides an alternative means for performing the key computational step required in [15] for computing fast spherical harmonic transforms. Unlike the implementation described in [15] of divide-and-conquer spectral methods, the butterfly scheme of the present paper does not require the use of extended precision during the compression precomputations in order to attain accuracy very close to the machine precision. With the butterfly, the required amount of preprocessing is quite reasonable, certainly not prohibitive.

Unfortunately, there seems to be little theoretical understanding of why the butterfly procedure works so well for associated Legendre functions (are the associated transforms nearly weighted averages of Fourier integral operators?). Complete proofs such as those in [11,15] are not yet available for the scheme of the present article. By construction, the butterfly enables fast, accurate applications of matrices to vectors when the precomputations succeed. However, we have yet to prove that the precomputations will compress the appropriate $n \times n$ matrix enough to enable applications of the matrix to vectors using only $\mathcal{O}(n \log(n))$ floating-point operations (flops). Nevertheless, the scheme has succeeded in all our numerical tests. We hope to produce rigorous mathematical proofs that the precomputations always compress the matrices as much as they did in our numerical experiments.

The precomputations for the algorithm of the present article require $\mathcal{O}(n^2)$ flops. The precomputations for the algorithm of [15] also require $\mathcal{O}(n^2)$ flops as implemented for the numerical examples of that paper; however, the procedure of [7] leads naturally to precomputations for the approach of [15] requiring only $\mathcal{O}(n \log(n))$ flops (though these “more efficient” precomputations do not become more efficient in practice until n is absurdly large, too large even to estimate reliably). We do not expect to be able to accelerate the precomputations for the algorithm of the present article without first producing the rigorous mathematical proofs mentioned in the previous paragraph. Even so, the current amount of preprocessing is not unreasonable, as the numerical examples of Section 6 illustrate.

Acknowledgements

We would like to thank V. Rokhlin for his advice, for his encouragement, and for the use of his software libraries, all of which have greatly enhanced this paper and the associated computer codes. We are also grateful to R.R. Coifman and Y. Shkolnisky. We would like to thank the anonymous referees for their useful suggestions.

References

- [1] M. Abramowitz, I.A. Stegun (Eds.), *Handbook of Mathematical Functions*, Dover Publications, New York, 1972.
- [2] J.C. Adams, P.N. Swarztrauber, SPHEREPACK 3.0: A model development facility, *Mon. Wea. Rev.* 127 (1999) 1872–1878.
- [3] A. Aho, J. Hopcroft, J. Ullman, *Data Structures and Algorithms*, Addison-Wesley, 1987.
- [4] E. Candès, L. Demanet, L. Ying, A fast butterfly algorithm for the computation of Fourier integral operators, *Multiscale Model. Simul.* 7 (4) (2009) 1727–1750.
- [5] H. Cheng, Z. Gimbutas, P.-G. Martinsson, V. Rokhlin, On the compression of low rank matrices, *SIAM J. Sci. Comput.* 26 (4) (2005) 1389–1404.
- [6] S.A. Goreinov, E.E. Tyrtyshnikov, The maximal-volume concept in approximation by low-rank matrices, in: V. Olshevsky (Ed.), *Structured Matrices in Mathematics, Computer Science, and Engineering I: Proceedings of an AMS-IMS-SIAM Joint Summer Research Conference*, University of Colorado, Boulder, June 27–July 1, 1999, vol. 280 of *Contemporary Mathematics*, AMS Publications, Providence, RI, 2001, pp. 47–51.
- [7] M. Gu, S.C. Eisenstat, A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem, *SIAM J. Matrix Anal. Appl.* 16 (1995) 172–191.
- [8] M. Gu, S.C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM J. Sci. Comput.* 17 (4) (1996) 848–869.
- [9] P.-G. Martinsson, V. Rokhlin, M. Tygert, On interpolation and integration in finite-dimensional spaces of bounded functions, *Commun. Appl. Math. Comput. Sci.* 1 (2006) 133–142.
- [10] E. Michielssen, A. Boag, A multilevel matrix decomposition algorithm for analyzing scattering from large structures, *IEEE Trans. Antennas Propag.* 44 (8) (1996) 1086–1093.
- [11] M. O’Neil, F. Woolfe, V. Rokhlin, An algorithm for the rapid evaluation of special function transforms, *Appl. Comput. Harmon. Anal.* 28 (2) (2010) 203–226.
- [12] M.G. Reuter, M.A. Ratner, T. Seideman, A fast method for solving both the time-dependent Schrödinger equation in angular coordinates and its associated “m-mixing” problem, *J. Chem. Phys.* 131 (2009) 094108-1–094108-6.
- [13] P.N. Swarztrauber, W.F. Spitz, Generalized discrete spherical harmonic transforms, *J. Comput. Phys.* 159 (2) (2000) 213–230.
- [14] G. Szegő, *Orthogonal Polynomials*, eleventh ed., vol. 23, Colloquium Publications, American Mathematical Society, Providence, RI, 2003.
- [15] M. Tygert, Fast algorithms for spherical harmonic expansions, II, *J. Comput. Phys.* 227 (8) (2008) 4260–4279.
- [16] E.E. Tyrtyshnikov, Incomplete cross approximation in the mosaic-skeleton method, *Computing* 64 (4) (2000) 367–380.
- [17] L. Ying, Sparse Fourier transform via butterfly algorithm, *SIAM J. Sci. Comput.* 31 (3) (2009) 1678–1694.