# FFTs for the 2-Sphere - Improvements and Variations* 

D.M. Healy Jr., D. Rockmore and P. J. Kostelec ${ }^{\dagger}$<br>Department of Mathematics<br>Dartmouth College<br>Hanover, NH 03755

Sean S. B. Moore<br>BBN Systems and Technologies<br>70 Fawcett Street<br>Cambridge, MA 02138

June 16, 1998


#### Abstract

Earlier work by Driscoll and Healy [16] has produced an efficient algorithm for computing the Fourier transform of band-limited functions on the 2 -sphere. In this paper we present a reformulation and variation of the original algorithm which results in a greatly improved inverse transform, and consequent improved convolution algorithm for such functions. All require at most $O\left(N \log ^{2} N\right)$ operations where $N$ is the number of sample points. We also address implementation considerations and give heuristics for allowing reliable and computationally efficient floating point implementations of slightly modified algorithms. These claims are supported by extensive numerical experiments from our implementation in C on DEC, HP and SGI platforms. These results indicate that variations of the algorithm are both reliable and efficient for a large range of useful problem sizes. Performance appears to be architecture-dependent. The paper concludes with a brief discussion of a few potential applications.


## 1 Introduction

### 1.1 History

The numerical calculation of Fourier expansions and convolutions of functions on the 2 -sphere are related problems which have been identified as important computational issues in many areas of applied science. For example, potential applications are found in astronomy [39], computer vision [31, 44], medical imaging [11], biology [37], statistical analysis of directional data [25, 24] and chemistry [34]. The paper [23] contains many references for possible applications in physics.

A significant set of applications comes from the fields of numerical weather prediction and global circulation modeling. In these areas much of the computational effort is directed towards the numerical solution of partial differential equations in spherical geometry [ $8,45,48]$. Use of spectral methods for these purposes requires efficient and reliable algorithms for computing spherical harmonic expansions - the lack of such an algorithm (until now) has been a serious bottleneck in pursuing a spectral method approach (see e.g. [45], p. 3416 and [48]).

[^0]These sorts of applications have motivated much of the research in fast algorithms for computing spherical harmonic expansions. Early work proposed approximate solutions to the problem of computing Fourier expansions on the 2 -sphere [41]. Other related work is the use of fast multipole methods for the computation of Legendre polynomial expansions, which gives part of the the full spherical harmonic expansion [3]. The complexity of these algorithms scales linearly in the desired accuracy.

A different approach was proposed in [16], which presents an exact, asymptotically fast approach to the problem of computing spherical harmonic coefficients. More specifically, an algorithm is given which in exact arithmetic permits efficient exact computation of the Fourier expansion and convolution of functions on the two-sphere, assuming that the functions have finite expansions in terms of spherical harmonics. The effects of finite precision arithmetic in the implementation were studied through a priori error estimates and numerical experiments for crucial steps in the algorithm. These results strongly suggest the possibility of an effective floating point implementation of the algorithm.

In a general setting, these algorithms can all be viewed as computational approaches to nonabelian harmonic analysis. Cast in this light, they have as their natural ancestor the now "classical" Fast Fourier Transform (FFT), first discovered by Gauss and later rediscovered and popularized by Cooley and Tukey (see [29] for a nice outline of much of the history). This family of algorithms efficiently computes the Fourier coefficients of a band-limited function on the circle, an abelian group. Its effective implementation has made possible a wealth of advances in many fields, most noticeably digital signal processing (see e.g. [18, 40]). A natural direction of generalization of the Cooley-Tukey FFT is the development of efficient and reliable algorithms to compute expansions of functions defined on finite or compact groups in terms of irreducible matrix coefficients. In this paper we present some particular results towards the development of this program. For a survey of other results and applications of "generalized FFTs" see [36, 43].

This paper continues and supplements the work in [16]. We give a reformulation and variation of that paper's algorithm in terms of a sparse structured factorization of the appropriate Fourier transform matrix. Reordering and transposing provides a more efficient inverse transform than that presented in [16]; we now attain the same order of complexity as that fast forward transform. Efficient inverse and forward transforms combine to yield a faster convolution algorithm. This is the most efficient such algorithm known to date. Our new presentation resembles the filterbank technology currently of interest in many digital signal processing applications (see e.g. [49]). The algorithms are of more than theoretical interest. By slightly varying the basic algorithm (at little theoretical cost) we have obtained numerically reliable and computationally efficient implementations that are competitive with other algorithms at useful problem sizes. The relative performance of the algorithms appears to be architecture-dependent.

### 1.2 Main idea

The Fourier transform of a function on the 2-sphere amounts to its $L^{2}$-projection onto the elements of a basis of spherical harmonic functions. This particular basis respects the rotational symmetries of the 2 -sphere in much the same way that the familiar sines and cosines are adapted to translations of periodic functions on the real line.

The primary algorithmic tool presented in this paper is an efficient algorithm for the computation of discrete Legendre transforms. For a given "bandwidth" $B>0$ (cf. Section 2) these are the sums of the form

$$
\begin{equation*}
\widehat{\mathbf{s}}(\ell, m)=\sum_{k=0}^{2 B-1} P_{\ell}^{m}\left(\cos \theta_{k}\right)[\mathbf{s}]_{k} \quad|m| \leq \ell=0,1, \ldots, B-1 \tag{1}
\end{equation*}
$$

where $P_{\ell}^{m}$ is the associated Legendre function of degree $\ell$ and order $m, \theta_{k}=\frac{\pi(k+1 / 2)}{2 B}$ and $\mathbf{s}$ is a data
vector with $k^{\text {th }}$ component $[\mathbf{s}]_{k}$, obtained from the samples of the original function which we wish to transform. Simply stated, these are inner products of a vector of sampled associated Legendre functions $\mathbf{P}_{\ell}^{m}$ against a data vector $\mathbf{s}$.

An obvious approach to evaluating the sums (1) is to compute them successively for the various degrees and orders. Computed in this way, each of these steps requires $2 B$ multiplications and $2 B-1$ additions. Since there are $n=B^{2}$ of these steps required to compute the full Fourier transform, this implies a total naive complexity of at most $4 n^{3 / 2}=4 B^{3}$ operations. ${ }^{1}$ We will refer to this as the direct algorithm.

In contrast, the results of this paper provide the basic tools for algorithms which improve the asymptotic complexity of the complete set of Legendre transforms from $O\left(n^{3 / 2}\right)$ to $O\left(n \log ^{2} n\right)(n=$ $\left.B^{2}\right)$. Variants of our algorithm still improve upon the better $O\left(n^{3 / 2}\right)$ exact algorithms, beginning at problem sizes as small as $B=256$ (cf. Section 6).

These fast algorithms use a divide and conquer approach, which may be familiar from the structure of many of the usual (abelian) FFT algorithms (see eg. [13, 50]). In such an approach, the problem of computing projections onto Legendre functions is decomposed into smaller subproblems of a similar form. The subproblems are solved recursively, by further subdivision, and their solutions are combined to solve the original problem. The advantage to this approach derives from the fact that the costs of the smaller subproblems, together with the cost of splitting will be less than the cost of the direct approach.

To insure that the splitting actually results in subproblems of reduced complexity we apply two main ideas:

- Using the Recurrence. The associated Legendre functions satisfy a three-term recurrence (as do many systems of special functions). This drives the divide and conquer strategy, expressing the higher degree inner products in terms of inner products with sampled trigonometric polynomials of lower degree, which according to the following observation, can be computed more efficiently.
- Smoothing and Subsampling. The inner products of a data vector against sampled low degree trigonometric polynomials may actually be accomplished in fewer than $B$ operations by using a "smoothed" data vector with fewer samples. In fact, only $\ell$ samples are needed to compute the inner product with a trigonometric polynomial of degree $\ell<B$.

Issues of numerical reliability are important in implementation and we will see how these considerations may be satisfied by developing variations of the basic algorithm. The main problem to guard against is that of pushing the recurrence too far.

### 1.3 Organization

The organization of the remainder of the paper is as follows. In Section 2 we briefly recall the notation and some necessary technical background material on Fourier analysis and fast algorithms. In particular, we review Fourier analysis on the 2-sphere i.e., the theory of spherical harmonics. We continue with a discussion of the recurrences satisfied by these functions and of the technique of smoothing and subsampling mentioned above. Section 3 contains the main algorithmic results of the paper. Here we explain how the techniques mentioned above (using the recurrence, smoothing and subsampling) may be combined to produce our basic divide and conquer transform algorithm. This algorithm has a natural formulation as a particular structured matrix factorization of matrix

[^1]containing sampled Legendre functions. This simultaneously provides a similar factorization of the transpose of the matrix for the forward transform, and as a result, we immediately obtain a fast inverse transform or synthesis algorithm and consequent fast convolution algorithm as well. This is all explained in Section 4. In Section 5 we discuss some simple variants of the basic algorithm, designed to give speed-ups in actual implementations. Numerical results supporting our claims of efficiency and reliability are presented in Section 6. In Section 7, we outline two possible applications of the algorithm. For the first application we show how our algorithm may be used towards the efficient computation of the bispectrum for band-limited functions. In the other application we present some experiments in using our algorithm for matched filtering of signals on the 2 -sphere. We conclude in Section 8 with a summary and brief discussion of future work.

Acknowledgement. We have benefitted greatly from close contact with the National Center for Atmospheric Research (NCAR). We thank NCAR's Scientific Computing Division for their hospitality on several occasions and access to their HP Exemplar, a machine named Sioux. In particular, we thank Mark Taylor for helping to educate us in the ways of scientific computing for climate modeling and helping us to run our comparisons against SPHEREPACK. Thanks also to Paul Swarztrauber for helpful discussions and encouragement.

## 2 Background

This section collects the basic tools for the formulation and solution of the problem of efficient Fourier analysis on the 2 -sphere. We first recall the definitions of Fourier analysis on the 2 -sphere and the spherical harmonic functions. We present the discretization of the Fourier transform and indicate the complexity of standard algorithms for its evaluation. We show how the problem of a fast transform is reduced to the question of fast discrete Legendre transforms, and detail the main tools (the use of recurrence relations as well as smoothing and subsampling) used for their efficient evaluation.

### 2.1 Fourier analysis on the 2 -sphere

As usual, $S^{2}$ denotes the 2-sphere or unit sphere in $\mathbf{R}^{3}$. A unit vector in $\mathbf{R}^{3}$ may be described by an angle $\theta, 0 \leq \theta \leq \pi$ measured down from the $z$-axis and an angle $\phi, 0 \leq \phi<2 \pi$ measured counterclockwise off the $x$-axis; this representation is unique for almost all unit vectors. Thus, if $\omega \in S^{2}$, then we may write $\omega(\theta, \phi)=(\cos \phi \sin \theta, \sin \phi \sin \theta, \cos \theta)$.


Figure 1: Parametrization of the 2-sphere.
Let $L^{2}\left(S^{2}\right)$ denote the Hilbert space of square integrable functions on the $S^{2}$. In coordinates, the usual inner product is given by

$$
\begin{equation*}
\langle f, h\rangle=\int_{0}^{\pi}\left[\int_{0}^{2 \pi} f(\theta, \phi) \overline{h(\theta, \phi)} d \phi\right] \sin \theta d \theta \tag{2}
\end{equation*}
$$

As is well-known (see e.g. [52]), the spherical harmonics provide an orthonormal basis for $L^{2}\left(S^{2}\right)$. For any nonnegative integer $\ell$ and integer $m$ with $|m| \leq \ell$, the $(\ell, m)$-spherical harmonic $Y_{\ell}^{m}$ is a harmonic homogeneous polynomial of degree $\ell$. The harmonics of degree $\ell$ span a subspace of $L^{2}\left(S^{2}\right)$ of dimension $2 \ell+1$ which is invariant under the rotations of the sphere. The expansion of any function $f \in L^{2}\left(S^{2}\right)$ in terms of spherical harmonics is written

$$
\begin{equation*}
f=\sum_{\ell \geq 0} \sum_{|m| \leq \ell} \widehat{f}(\ell, m) Y_{\ell}^{m} \tag{3}
\end{equation*}
$$

and $\hat{f}(\ell, m)$ denotes the $(\ell, m)$-Fourier coefficient, equal to $\left\langle f, Y_{\ell}^{m}\right\rangle$.
In the coordinates $(\theta, \phi), Y_{\ell}^{m}$ has a factorization,

$$
\begin{equation*}
Y_{\ell}^{m}(\theta, \phi)=k_{\ell, m} P_{\ell}^{m}(\cos \theta) e^{i m \phi} \tag{4}
\end{equation*}
$$

where $P_{\ell}^{m}$ is the associated Legendre function of degree $\ell$ and order $m$ and $k_{\ell, m}$ is a normalization constant.

Consequently, separating variables according to (4) shows that the computation of the spherical harmonic transform can be reduced to a regular Fourier transform in the longitudinal coordinate $\phi$ followed by a projection onto the associated Legendre functions

$$
\begin{equation*}
\widehat{f}(\ell, m)=\left\langle f, Y_{\ell}^{m}\right\rangle=k_{\ell, m} \int_{0}^{\pi}\left[\int_{0}^{2 \pi} e^{-i m \phi} f(\theta, \phi) d \phi\right] P_{\ell}^{m}(\cos \theta) \sin \theta d \theta \tag{5}
\end{equation*}
$$

The associated Legendre functions satisfy a characteristic three-term recurrence

$$
\begin{equation*}
(\ell-m+1) P_{\ell+1}^{m}(x)-(2 \ell+1) x P_{\ell}^{m}(x)+(\ell+m) P_{\ell-1}^{m}(x)=0 \tag{6}
\end{equation*}
$$

critical for the algorithms developed in this paper.
In analogy with the case of functions on the circle, we say that $f \in L^{2}\left(S^{2}\right)$ is band-limited with band-limit or bandwidth $B \geq 0$ if $\widehat{f}(\ell, m)=0$ for all $\ell \geq B$. For band-limited functions we have a simple quadrature (sampling) result which reduces the integrals (5) to finite weighted sums of a sampled data vector obtained from the integrand.

Theorem 1 [cf. [16], Theorem 3] Let $f \in L^{2}\left(S^{2}\right)$ have bandwidth $B$. Then for each $|m| \leq \ell<B$,

$$
\widehat{f}(\ell, m)=\frac{\sqrt{2 \pi}}{2 B} \sum_{j=0}^{2 B-1} \sum_{k=0}^{2 B-1} a_{j}^{(B)} f\left(\theta_{j}, \phi_{k}\right) e^{-i m \phi_{k}} P_{\ell}^{m}\left(\cos \theta_{j}\right)
$$

where the sample points are chosen from the equiangular grid: $\theta_{j}=\pi(2 j+1) / 4 B, \phi_{k}=2 \pi k / 2 B$; and the weights $a_{j}^{(B)}$ (cf. Figure 2) play a role analogous to the $\sin \theta$ factor in the integrals.

Remark. Although we will use this formulation of the Sampling Theorem to provide the starting point for our fast algorithms, it is actually possible to give a sampling theorem which uses only $B$ samples in the $\theta$ coordinate. These samples are then interpolated to give $2 B-1$ samples for use in the fast transform algorithms.


Figure 2: Plot of sample weights for range of problem sizes.

### 2.2 A discrete Fourier transform for $S^{2}$ and the Legendre transform

The Fourier transform of a function $f$ of bandwidth $B$ is the collection of its Fourier coefficients,

$$
\{\hat{f}(\ell, m)|0 \leq|m| \leq \ell<B\} .
$$

The Sampling Theorem (Theorem 1) expresses the $(\ell, m)$-Fourier coefficient of $f$ as the finite sum

$$
\begin{equation*}
\widehat{f}(\ell, m)=\frac{\sqrt{2 \pi}}{2 B} \sum_{j=0}^{2 B-1} \sum_{k=0}^{2 B-1} a_{j}^{(B)} f\left(\theta_{j}, \phi_{k}\right) e^{-i m \phi_{k}} P_{\ell}^{m}\left(\cos \theta_{j}\right) \tag{7}
\end{equation*}
$$

This method of computing the Fourier coefficients of $f$ is called the discrete Fourier transform, or DFT of $f$. Notice that direct computation of each $\hat{f}(\ell, m)$ would require $4 B^{2}$ operations and thus $O\left(B^{4}\right)$ in total.

More efficient algorithms use a separation of variables approach. We proceed by first summing over the $k$ index, computing the inner exponential sums which depend only the indices $j$ and $m$. We may do this efficiently for all $m$ between $-B$ and $B$ via the FFT (cf. [18]). The computation is completed by performing the requisite discrete Legendre transforms, which for a given order $m>0$ we define as a set of sums

$$
\begin{equation*}
\sum_{k=0}^{N-1}[\mathbf{s}]_{k} P_{\ell}^{m}\left(\cos \theta_{k}\right)=\left\langle\mathbf{s}, \mathbf{P}_{\ell}^{m}\right\rangle ; \quad \ell=m, m+1, \ldots, N-1, \tag{8}
\end{equation*}
$$

for an arbitrary input vector $\mathbf{s}$ with $k^{t h}$ component $[\mathbf{s}]_{k}$. Here we have introduced a discrete inner product notation and the convention that $\mathbf{P}_{\ell}^{m}$ denotes the vector comprised of the appropriate samples of the function $P_{\ell}^{m}(\cos \theta)$ :

$$
\mathbf{P}_{\ell}^{m}=\left(\begin{array}{c}
P_{\ell}^{m}\left(\cos \theta_{0}\right) \\
\vdots \\
P_{\ell}^{m}\left(\cos \theta_{N-1}\right)
\end{array}\right) .
$$

We may also say that (8) computes the projection of $\mathbf{s}$ onto $\mathbf{P}_{\ell}^{m}$.
The problem of a fast spherical harmonic transform is now reduced to the efficient calculation of these discrete Legendre transforms. Notice that even without a fast algorithm here, the separation of variables turns an $O\left(B^{4}\right)$ calculation into an $O\left(B^{3}\right)$ calculation.

In order to simplify the discussion of the basic idea for the efficient evaluation of the sums (8), we will specialize to the case $m=0$, for which $P_{\ell}^{m}=P_{\ell}^{0}=P_{\ell}$ is the Legendre polynomial of degree $\ell$. For higher orders the algorithm generalizes directly.

As remarked in the introduction, our fast algorithm relies on two basic ideas,
(1) Using the Recurrence; (2) Smoothing and Subsampling.

We proceed by developing each of these techniques separately (Sections 2.3 and 2.4) and then show how they are combined to yield our algorithm (Section 3).

### 2.3 Using the Legendre recurrence

The recurrence relation satisfied by $P_{\ell}(\cos \theta)$ is (for $\ell \geq 0$ )

$$
\begin{equation*}
(\ell+1) P_{\ell+1}(\cos \theta)-(2 \ell+1) \cos \theta P_{\ell}(\cos \theta)+(\ell) P_{\ell-1}(\cos \theta)=0 \tag{9}
\end{equation*}
$$

with initial conditions $P_{0}(\cos \theta)=1$ and $P_{-1}(\cos \theta)=0$. Consequently, the higher degree Legendre polynomials can be expressed in terms of those of lower degree as follows. For any fixed "level" $L$, iterating the recurrence formula (9) forward $r$ steps produces trigonometric polynomials $A_{r}^{L}$ and $B_{r}^{L}$ such that

$$
\begin{equation*}
P_{L+r}(\cos \theta)=A_{r}^{L}(\cos \theta) P_{L}(\cos \theta)+B_{r}^{L}(\cos \theta) P_{L-1}(\cos \theta), \tag{10}
\end{equation*}
$$

for $r \geq 1$. We refer to $A_{r}^{L}$ and $B_{r}^{L}$, as shifted Legendre polynomials, as they are generated by the following shifted form of the Legendre polynomial recurrence (9),

$$
\begin{equation*}
(L+r+1) p_{r+1}^{L}(\cos \theta)-(2 L+2 r+1) \cos \theta p_{r}^{L}(\cos \theta)+(L+r) p_{r-1}^{L}(\cos \theta)=0 \tag{11}
\end{equation*}
$$

with initial conditions $A_{0}^{L}=1, A_{-1}^{L}=0$ and $B_{0}^{L}=0, B_{-1}^{L}=1$, respectively. This is readily concluded by comparing a matrix formulation of (10),

$$
\binom{P_{L+r+1}}{P_{L+r}}=\left(\begin{array}{cc}
A_{r+1}^{L} & B_{r+1}^{L}  \tag{12}\\
A_{r}^{L} & B_{r}^{L}
\end{array}\right)\binom{P_{L}}{P_{L-1}},
$$

with the matrix form of the one-step recurrence (9)

$$
\begin{aligned}
\binom{P_{L+r+1}}{P_{L+r}} & =\left(\begin{array}{cc}
\frac{2(L+r)+1}{L+r+1} \cos \theta & -\frac{L+r}{L+r+1} \\
1 & 0
\end{array}\right)\binom{P_{L+r}}{P_{L+r-1}} \\
& =\left(\begin{array}{cc}
\frac{2 L+2 r+1}{L+r+1} \cos \theta & -\frac{L+r}{L+r+1} \\
1 & 0
\end{array}\right)\left(\begin{array}{cc}
A_{r}^{L} & B_{r}^{L} \\
A_{r-1}^{L} & B_{r-1}^{L}
\end{array}\right)\binom{P_{L}}{P_{L-1}} .
\end{aligned}
$$

A similar argument gives the general recurrence satisfied by the shifted Legendre functions

$$
\left(\begin{array}{ll}
A_{r+s}^{L} & B_{r+s}^{L}  \tag{13}\\
A_{r+s-1}^{L} & B_{r+s-1}^{L}
\end{array}\right)=\left(\begin{array}{ll}
A_{s}^{L+r} & B_{s}^{L+r} \\
A_{s-1}^{L+r} & B_{s-1}^{L+r}
\end{array}\right) \cdot\left(\begin{array}{ll}
A_{r}^{L} & B_{r}^{L} \\
A_{r-1}^{L} & B_{r-1}^{L}
\end{array}\right) .
$$

We make the convention that

$$
\begin{equation*}
A_{r}^{0}=P_{r}, \quad \text { and } \quad B_{r}^{0}=0 \tag{14}
\end{equation*}
$$

so that (13) subsumes the original Legendre recurrence (12).
From (11) we see that the degrees of the shifted Legendre polynomials $A_{r}^{L}$ and $B_{r}^{L}$ are $r$ and $r-1$ respectively. The point of introducing these polynomials is that they allow us to rewrite projections
onto high degree Legendre polynomials as sums of projections onto (shifted) Legendre polynomials of lower degree. More precisely, suppose that in the course of computing the inner products $\left\langle\mathbf{s}, \mathbf{P}_{j}\right\rangle$ (for $j \leq L$ ) we had stored the components of the vectors $\mathbf{s}^{j}=\mathbf{s} \mathbf{P}_{j}$, defined as the pointwise product of the vectors s and $\mathbf{P}_{j}$. The recurrence (10) allows us to re-use this data to compute the projection of $\mathbf{s}$ onto a higher order function, $P_{L+r}$ as follows:

$$
\begin{equation*}
\left\langle\mathbf{s}, \mathbf{P}_{L+r}\right\rangle=\left\langle\mathbf{s},\left(\mathbf{A}_{r}^{L} \mathbf{P}_{L}+\mathbf{B}_{r}^{L} \mathbf{P}_{L-1}\right)\right\rangle=\left\langle\mathbf{s}^{L}, \mathbf{A}_{r}^{L}\right\rangle+\left\langle\mathbf{s}^{L-1}, \mathbf{B}_{r}^{L}\right\rangle . \tag{15}
\end{equation*}
$$

If the vectors $\mathbf{s}^{L}=\mathbf{s} \mathbf{P}_{L}$ and $\mathbf{s}^{L-1}=\mathbf{s} \mathbf{P}_{L-1}$ have been stored, then (15) shows that the higher degree inner product can be computed as inner products of stored data and (precomputed) sampled values of the polynomials $A_{r}^{L}$ and $B_{r}^{L}$, each of which have degree at most $r$, which is necessarily less than the degree of $P_{L+r}$.

Since $A_{r}^{L}$ and $B_{r}^{L}$ also satisfy a recurrence, this procedure can be repeated. Following this through yields a divide and conquer scheme for performing the full computation.

The motivation for the successive reductions is that the projections of a data vector onto lower degree trigonometric polynomials can be computed more efficiently by first lowpass filtering (smoothing) and then subsampling the data vector, so that ultimately the projection may be computed by a summation with fewer terms (cf. Lemma 1). This is the next topic.

### 2.4 Smoothing and subsampling - working in the "cosine transform" domain

The "conquer" part of the divide and conquer algorithm described above is the computation of inner products of the form $\left\langle\mathbf{s}, \mathbf{Q}_{n}\right\rangle$ for a sequence of vectors $\mathbf{Q}_{n}$ varying over some range of $n$. Here $\mathbf{s}$ is a data vector, and $\mathbf{Q}_{n}$ is composed of samples of a trigonometric polynomial $Q_{n}$ of degree $n$ (in fact, a shifted Legendre polynomial). We examine the complexity of computing these inner products, and show that it grows linearly with $n$, modulo a fixed overhead. We will make use of a cosine transform representation of the vectors in the inner product. Put in the language of matrices, we will show that the matrix of the discrete Legendre transform may be brought to triangular form by means of the cosine transform matrix.

We begin by recalling that any vector of length $N$, s, may always be represented as uniform samples of a cosine series of degree less than $N$ :

$$
\begin{equation*}
[\mathbf{s}]_{k}=\sum_{n=0}^{N-1} \sigma_{n} \cos \left(n \theta_{k}\right), \quad k=0, \ldots, N-1 \tag{16}
\end{equation*}
$$

where, as before, $\theta_{k}=\frac{\pi(2 k+1)}{2 N}$.
The coefficients $\sigma_{n}$ are obtained by computing the discrete cosine transform (DCT) of s. Explicitly, we let $\mathcal{C}_{N}$ denote the $N$-dimensional orthogonal DCT matrix (see e.g. [18], p. 386) comprised of normalized sampled cosines:

$$
\begin{equation*}
\left(\mathcal{C}_{N}\right)_{j, k}=b(j) \cos \left(j \theta_{k}\right) \quad 0 \leq j, k \leq N-1, \tag{17}
\end{equation*}
$$

with normalization factors

$$
\begin{equation*}
b(0)=\sqrt{\frac{1}{N}} \quad \text { and } \quad b(j)=\sqrt{\frac{2}{N}} \text { for } j=1, \ldots, N \tag{18}
\end{equation*}
$$

In terms of $\mathcal{C}=\mathcal{C}_{N}$, the coefficients in (16) are given by

$$
\sigma_{n}= \begin{cases}\sqrt{\frac{2}{N}}[\mathcal{C} \mathbf{s}]_{n} & \text { if } n \neq 0  \tag{19}\\ \sqrt{\frac{1}{N}}[\mathcal{C} \mathbf{s}]_{n} & \text { if } n=0\end{cases}
$$

Remark. In practice, the set of cosine coefficients $\left\{[\mathcal{C} \mathbf{s}]_{n} \mid n=0, \ldots, N-1\right\}$ can be obtained efficiently (in at most $\frac{3}{2} N \log N$ operations for $N$ a power of 2 ) by a fast DCT algorithm, which amounts to a clever factorization of the matrix $\mathcal{C}$ (see [47] and the references contained therein).

The orthogonality of $\mathcal{C}$ implies $\langle\mathcal{C} \mathbf{s}, \mathcal{C} \mathbf{Q}\rangle=\langle\mathbf{s}, \mathbf{Q}\rangle$. The computational advantage of computing the inner product in the cosine transform domain comes from the fact that for any trigonometric polynomial $Q$ of degree at most $N-1$, the cosine coefficients $[\mathcal{C} \mathbf{Q}]_{n}$ vanish for $n>\operatorname{deg}(Q)$ (cf. Lemma 1) and this reduces the number of operations required to compute the inner product with $\mathcal{C}$ s. In particular, this applies to the various Legendre functions we use.

To illustrate, the Legendre polynomial $Q=P_{\ell}$ is a trigonometric polynomial of degree $\ell$. This is easily verified using the recurrence relation (9) with the initial conditions $P_{0}(\cos \theta)=1, P_{1}(\cos \theta)=$ $\cos \theta$. For example,

$$
P_{2}(\cos \theta)=\frac{5}{3} \cos \theta \cos \theta+\frac{2}{3} 1=\frac{5}{3} \frac{\cos (2 \theta)+1}{2}+\frac{2}{3} .
$$

Consequently, for $n>\ell,\left[\mathcal{C} \mathbf{P}_{\ell}\right]_{n}=0$ and the inner product $\operatorname{sum}\left\langle\mathcal{C} \mathbf{P}_{\ell}, \mathcal{C} \mathbf{s}\right\rangle=\left\langle\mathbf{P}_{\ell}, \mathbf{s}\right\rangle$ can be computed as a sum of only $\ell+1$ terms (instead of $N)$,

$$
\begin{align*}
\left\langle\mathbf{s}, \mathbf{P}_{\ell}\right\rangle & =\sum_{m=0}^{N-1}[\mathcal{C} \mathbf{s}]_{m}\left[\mathcal{C} \mathbf{P}_{\ell}\right]_{m} \\
& =\sum_{m=0}^{l}[\mathcal{C} \mathbf{s}]_{m}\left[\mathcal{C} \mathbf{P}_{\ell}\right]_{m}+\sum_{m=l+1}^{N-1}[\mathcal{C} \mathbf{s}]_{m} \cdot 0 \\
& =\sum_{m=0}^{l}[\mathcal{C} \mathbf{s}]_{m}\left[\mathcal{C} \mathbf{P}_{\ell}\right]_{m} \tag{20}
\end{align*}
$$

This shows that the computation of the low degree Legendre projections (i.e., projections of data onto $P_{\ell}$ for small $\ell$ ) can be accomplished with very few operations, after the overhead of computing the cosine expansions of the data vector $\mathbf{s}$ and the vector of sampled Legendre polynomials $\mathbf{P}_{\ell}$.

In terms of matrix arithmetic (20) is equivalent to the observation that the matrix representing the discrete Legendre transform has a factorization as a product of a cosine transform matrix and a triangular matrix comprised of the Legendre polynomial cosine coefficients,

$$
\left(\left(P_{\ell}\left(\cos \theta_{j}\right)\right)\right)=\left(\begin{array}{cccc}
{\left[\mathcal{C} \mathbf{P}_{0}\right]_{0}} & 0 & \cdots & 0  \tag{21}\\
{\left[\mathcal{C} \mathbf{P}_{1}\right]_{0}} & {\left[\mathcal{C} \mathbf{P}_{1}\right]_{1}} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
{\left[\mathcal{C} \mathbf{P}_{N-1}\right]_{0}} & {\left[\mathcal{C} \mathbf{P}_{N-1}\right]_{1}} & \cdots & {\left[\mathcal{C} \mathbf{P}_{N-1}\right]_{N-1}}
\end{array}\right) \cdot \mathcal{C}
$$

where $\mathcal{C}$ denotes the $N \times N$ orthogonal discrete cosine transform matrix defined in (17) If the cosine coefficients of the sampled Legendre polynomials $\left(\left[\mathcal{C} \mathbf{P}_{j}\right]_{k}\right)$ are prestored, this approach is an alternative to the direct computation of the Legendre transform which has the same asymptotic complexity, but is faster for even moderate sized transforms, assuming the use of a fast DCT routine. (See the previous remark and reference on this issue). Figures 3 and 4 illustrate both of these approaches.

Lemma 1 below gives a generalization of (20) appropriate for use in our fast algorithm, where we apply it to reduce the complexity of projections onto the lower degree Legendre functions. The lemma reformulates the earlier discussion of this section in terms of smoothing or filtering operators applied to the vectors involved in the inner products.

We define the critically sampled lowpass operator (of bandwidth $p$ ), denoted $\mathcal{L}_{p}^{N}($ for $p<N)$, by

$$
\begin{equation*}
\mathcal{L}_{p}^{N}=\mathcal{C}_{p}^{-1} \mathcal{T}_{p}^{N} \mathcal{C}_{N} \tag{22}
\end{equation*}
$$

a)

b)


Figure 3: Graphical representation of a direct non-DCT discrete Legendre transform of sampled signal s: (a) Input of length $N, \mathbf{s}$; (b) The vectors of sampled Legendre polynomials $\mathbf{P}_{k}$; (c) Pointwise multiplications $\mathbf{s} \mathbf{P}_{k}$ such that $\sum_{j}\left[\mathbf{s} \mathbf{P}_{k}\right]_{j}=\left\langle\mathbf{s}, \mathbf{P}_{k}\right\rangle$ in at most $N$ operations indicated by the $\Sigma$ notation.
where $\mathcal{T}_{p}^{N}$ is the truncation operator that only keeps the first $p$ coordinates of a given input vector. The effect of $\mathcal{L}_{p}^{N}$ is to first compute the cosine representation of a vector of length $N$, then remove all frequency components beyond $p$ from $\mathbf{s}$, (smoothing) and finally, to keep only those samples necessary to represent this smoothed version (subsampling). This is illustrated schematically in Figure 5 below. As indicated by the subscript and superscript, this operator takes sequences of length $N$ to sequences of length $p$.
a)

$$
\mathcal{C s}=
$$


b)

c)

Figure 4: Graphical representation of discrete Legendre transform of Figure 3, using the DCTs of the signal $\mathbf{s}$ and the sampled Legendre polynomials: (a) Input $\mathcal{C}$ s for $\mathbf{s}$ in Figure 3; (b) DCT of the vectors of sampled Legendre polynomials $\mathbf{P}_{k}, \mathcal{C} \mathbf{P}_{k}$. Notice that the DCT of a given Legendre polynomial has only as many nonzero coefficients as the degree; (c) Pointwise multiplications $\mathcal{C} \mathbf{s} \mathcal{C} \mathbf{P}_{k}$ such that $\sum_{j}\left[\mathcal{C} \mathbf{s} \mathcal{C} \mathbf{P}_{k}\right]_{j}=\left\langle\mathbf{s}, \mathbf{P}_{k}\right\rangle$ in at most $k$ operations indicated by the $\Sigma$ notation.

With the preceding notation, the necessary generalization of (20) is given by Lemma 1.
Lemma 1 Let $Q$ be a trigonometric polynomial of degree $p$,

$$
Q(\cos \theta)=\sum_{m=0}^{p} \gamma_{m} \cos m \theta
$$

and let $\mathbf{s}$ be any sequence of length $N$ with $N \geq p$. Then

$$
\langle\mathbf{s}, \mathbf{Q}\rangle=\sum_{k=0}^{N-1}[\mathbf{s}]_{k} Q\left(\cos \theta_{k}\right)=\sum_{j=0}^{p-1}\left[\mathcal{L}_{p}^{N} \mathbf{s}\right]_{j} Q\left(\cos \frac{N \theta_{j}}{p}\right)=\left\langle\mathcal{L}_{p}^{N} \mathbf{s}, \mathcal{L}_{p}^{N} \mathbf{Q}\right\rangle
$$

## Proof:

$$
\begin{aligned}
\langle\mathbf{s}, \mathbf{Q}\rangle & =\left\langle\mathcal{C}_{N} \mathbf{s}, \mathcal{C}_{N} \mathbf{Q}\right\rangle \\
& =\sum_{k=0}^{N-1}\left[\mathcal{C}_{N} \mathbf{s}\right]_{k}\left[\mathcal{C}_{N} \mathbf{Q}\right]_{k}
\end{aligned}
$$



Figure 5: Smoothing and subsampling.

$$
=\sum_{k=0}^{p}\left[\mathcal{C}_{N} \mathbf{s}\right]_{k}\left[\mathcal{C}_{N} \mathbf{Q}\right]_{k}
$$

since $\left[\mathcal{C}_{N} \mathbf{Q}\right]_{k}=0$ for $k>p$. This last sum is the same as

$$
\begin{equation*}
\left\langle\mathcal{T}_{p}^{N} \mathcal{C}_{N} \mathbf{s}, \mathcal{T}_{p}^{N} \mathcal{C}_{N} \mathbf{Q}\right\rangle=\left\langle\mathcal{L}_{p}^{N} \mathbf{s}, \mathcal{L}_{p}^{N} \mathbf{Q}\right\rangle \tag{23}
\end{equation*}
$$

and it is immediately verified that $\mathcal{L}_{p}^{N} \mathbf{Q}$ is the same as the function $Q$ sampled on the coarser grid.

Note that the rightmost inner product is of sequences of length $p$, and since $Q$ already has bandwidth $p$, the effect of applying $\mathcal{L}_{p}^{N}$ to $Q$ is to simply sample $Q$ on the coarser grid.

Lemma 2 Let $0<p \leq N$ be powers of 2 . For an arbitrary input $\mathbf{s}$ the computation $\mathbf{s} \mapsto \mathcal{L}_{p}^{N} \mathbf{s}$ can be accomplished in at most $\frac{3}{2} N \log N+\frac{3}{2} p \log p$ operations.

Proof: As described above, $\mathcal{L}_{p}^{N} \mathbf{s}$ can be accomplished by taking the DCT of $\mathbf{s}$, a vector of length $N$, truncating the result to the first $p$ coefficients, and then using the inverse DCT to evaluate the lowpassed function at $p$ samples. A fast DCT (and inverse DCT) algorithm due to Stiedl and Tasche [47] requires $\frac{3}{2} m \log m$ operations to compute a DCT of length $m$. Since truncation requires no additional operations, the lemma follows.

Lemma 3 Let $N$ be a power of 2 and $\mathbf{s}$ a vector of length $N$. Suppose $F_{\ell}(\cos \theta)(\ell=0, \ldots, N-1)$ satisfies a recurrence $a_{\ell} F_{\ell+1}(\cos \theta)-b_{\ell} \cos (\theta) F_{\ell}(\cos \theta)+c_{\ell} F_{\ell-1}(\cos \theta)$ with initial conditions $F_{-1}=0$ and $F_{0}=1$. Then assuming the prestorage of the $D C T$ of the $F_{\ell}$, for an arbitrary input $\mathbf{s}$, the collection of inner products $\left\langle\mathbf{F}_{\ell}, \mathbf{s}\right\rangle$ can be computed in at most $\frac{3}{2} N \log N+\frac{N(N+1)}{2}$ operations, versus $N^{2}$ required by direct computation.

Computation of any collection of discrete trigonometric polynomial projections by applying the DCT according to Lemma 1 will be referred to as computation by the semi-naive approach.

The key property of (20) is that it shows that only $\ell+1$ samples of a suitably modified data vector are needed to compute the inner products $\left\langle\mathbf{s}, \mathbf{P}_{\ell}\right\rangle$. The generalization of this to the various associated Legendre functions and their shifted forms permits us to reduce the complexity of the small Legendre transform problems obtained by applying a divide and conquer approach to the original problem. As we shall see in Section 3.1, these subproblems compute inner products of data with (shifted) Legendre polynomials of at most half the degree of the inner products in the original problem.

## 3 Fast Discrete Legendre Transforms

In Section 2 we saw that the problem of computing fast spherical harmonic expansions for band-limited functions may be reduced to the efficient calculation of discrete Legendre transforms defined in (8).

In this section we show how this is done. We simplify the discussion by presenting the case of the Legendre polynomial transform, $m=0$, For higher orders the algorithm generalizes directly.

The problem of interest is then to efficiently compute

$$
\begin{equation*}
\sum_{k=0}^{N-1}[\mathbf{s}]_{k} P_{\ell}\left(\cos \theta_{k}\right)=\left\langle\mathbf{s}, \mathbf{P}_{\ell}\right\rangle ; \quad \ell=0,1, \ldots, N-1 \tag{24}
\end{equation*}
$$

where $\mathbf{s}$ represents a data vector which we think of as sampled values of a function $f$ which we wish to transform. To do this we apply the two basic notions, the Legendre recurrence, and smoothing, which we have considered in Sections 2.3 and 2.4. In this section we show how they are combined in a divide and conquer algorithm for efficient computation of the Legendre transform. This is our main theoretical result and it can be found in Section 3.3 as Theorem 2.

### 3.1 Fast discrete Legendre transform via divide and conquer

Sections 2.3 and 2.4 provide the ingredients that make up our divide and conquer algorithm for computing discrete Legendre transforms. Schematically, the idea is quite simple to describe: The recurrence relations satisfied by the (shifted) Legendre polynomials (10) permit projection of data onto a collection of Legendre functions to be computed as linear combinations of similar projections onto Legendre functions of lower degree. The projections at lower degree require fewer samples (cf. Lemma 1), thereby taking on the form of problems of smaller size.

A uniform formulation of the problem allows us to divide the original problem (24) into a low degree and high degree subproblem, each of which can be shown to be equivalent to Legendre transforms of half the size. Let $T(N)$ denotes the complexity of a discrete Legendre transform of size $N$. Then this simple description gives rise to the usual divide and conquer recurrence $T(N)=2 T(N / 2)+S(N)$ (see e.g. [13]), where $S(N)$ represents the overhead cost of rewriting the problem of size $N$ as two subproblems of size $N / 2$. Efficiency derives from the ability to perform the reduction quickly. In particular we shall see that $S(N)$ is $O(N \log N)$, so that iterating this subdivision procedure yields an $O\left(N \log ^{2} N\right)$ algorithm for computing the original discrete Legendre transform.

Before giving the general framework, we illustrate the details of the basic idea by applying it once to split the original problem of computing the $N$ Legendre projections $\left\langle\mathbf{s}, \mathbf{P}_{\ell}\right\rangle,(0 \leq \ell<N)$ into two subproblems of size $N / 2$. To do this, we proceed by dividing the original problem into two separate computations:

- Low degree transform coefficients: $\left\langle\mathbf{s}, \mathbf{P}_{r}\right\rangle, 0 \leq r<\frac{N}{2}$.
- High degree transform coefficients: $\left\langle\mathbf{s}, \mathbf{P}_{\frac{N}{2}+r}\right\rangle, 0 \leq r<\frac{N}{2}$.

Although each subproblem is a collection of $N / 2$ projections, these are not as yet problems of size $N / 2$ since both use input of size $N$. (Recall that the original problem assumes input of size $N$ ). Reduction of the problem size will come from applying Lemma 1 which showed that any trigonometric polynomial projection $\langle\mathbf{s}, \mathbf{Q}\rangle$ can be computed as a sum of length $\operatorname{deg}(Q)$ by smoothing. This fact can be used immediately for the low degree projections, yielding the equivalent set of inner products $\left\langle\mathcal{L}_{\frac{N}{2}}^{N} \mathbf{s}, \mathcal{L}_{\frac{N}{2}}^{N} \mathbf{P}_{\ell}\right\rangle$ $\left(0 \leq \ell<\frac{N}{2}\right)$, which is a discrete Legendre transform of size $N / 2$. That is, we are projecting onto $N / 2$ Legendre functions using inner products of length $N / 2$.

To apply the same idea to the high degree projections, we use the recurrence (10) which allows $\left\langle\mathbf{s}, \mathbf{P}_{\frac{N}{2}+r}\right\rangle, 0 \leq r<\frac{N}{2}$ to be rewritten as a sum of projections onto lower degree shifted Legendre polynomials,

$$
\begin{align*}
\left\langle\mathbf{s}, P_{\frac{N}{2}+r}\right\rangle & =\left\langle\mathbf{s}, \mathbf{A}_{r}^{\frac{N}{2}} \mathbf{P}_{\frac{N}{2}}+\mathbf{B}_{r}^{\frac{N}{2}} \mathbf{P}_{\frac{N}{2}-1}\right\rangle \\
& =\left\langle\mathbf{s}^{\frac{N}{2}}, \mathbf{A}_{r}^{\frac{N}{2}}\right\rangle+\left\langle\mathbf{s}^{\frac{N}{2}-1}, \mathbf{B}_{r}^{\frac{N}{2}}\right\rangle \tag{25}
\end{align*}
$$

where we have maintained the notation from Section 2.3, writing $\mathbf{s}^{\frac{N}{2}}=\mathbf{s P}_{\frac{N}{2}}$ and $\mathbf{s}^{\frac{N}{2}-1}=\mathbf{s P}_{\frac{N}{2}-1}$. Since for $r=0, \ldots, \frac{N}{2}-1$ the shifted Legendre polynomials $A_{r}^{\frac{N}{2}}$ and $B_{r}^{\frac{N}{2}}$ all have degree less than $N / 2$, again Lemma 1 may be applied and (25) may be computed as a sum of inner products of length $N / 2$,

$$
\begin{equation*}
\left\langle\mathcal{L}_{\frac{N}{2}}^{N} \mathbf{s}^{\frac{N}{2}}, \mathcal{L}_{\frac{N}{2}}^{N} \mathbf{A}_{r}^{\frac{N}{2}}\right\rangle+\left\langle\mathcal{L}_{\frac{N}{2}}^{N} \mathbf{s}^{\frac{N}{2}-1}, \mathcal{L}_{\frac{N}{2}}^{N} \mathbf{B}_{r}^{\frac{N}{2}}\right\rangle . \tag{26}
\end{equation*}
$$

While the reductions at low and high degree look slightly different (the former requires only smoothing, but the latter requires both smoothing and the use of the recurrence (10)), Section 3.2 will provide a uniform formulation in which both sets of projections become instances of a Legendre transform of size $N / 2$.

Finally, it is critical that the reduction to a smaller problem size may be accomplished efficiently. Setting $\mathcal{L}=\mathcal{L}_{\frac{N}{2}}^{N}$, if the data $\mathcal{L} \mathbf{P}_{\ell}, \mathcal{L} \mathbf{A}_{r}^{\frac{N}{2}}$, and $\mathcal{L} \mathbf{B}_{r}^{\frac{N}{2}}$ are all stored, then only the quantities $\mathcal{L}$ s $\mathcal{L} \mathbf{s}^{\frac{N}{2}}$, and $\mathcal{L} \mathbf{s}^{\frac{N}{2}-1}$ need be computed. According to Lemma 2 this requires at most $3\left[\frac{3}{2}\left(N \log N+\frac{N}{2} \log \frac{N}{2}\right)\right]$ operations.
Remark. In order to simplify the remaining formulas, for the remainder of the paper we will continue to write $\mathcal{L}$ for $\mathcal{L}_{\frac{N}{2}}^{N}$, leaving $N$ to be determined by context.

### 3.2 General Legendre transforms - a uniform problem formulation

The preceding discussion gives the basic idea behind the divide and conquer approach. The process indicated above must be repeated, recursively subdividing the original problem into smaller and smaller subproblems. In order to see how the pieces fit together, we need a uniform description of the computational "unit" encountered at each division. This motivates the following definition.

Definition 1 For integers $M>0$ and $L \geq 0$, define the $M \times 2 M$ shifted Legendre transform
matrix, $L T_{M}^{L}$ by

$$
L T_{M}^{L}=\left(\begin{array}{cc}
\left(\mathbf{A}_{0}^{L}\right)^{t} & \left(\mathbf{B}_{0}^{L}\right)^{t}  \tag{27}\\
\left(\mathbf{A}_{1}^{L}\right)^{t} & \left(\mathbf{B}_{1}^{L}\right)^{t} \\
\vdots & \vdots \\
\left(\mathbf{A}_{M-1}^{L}\right)^{t} & \left(\mathbf{B}_{M-1}^{L}\right)^{t}
\end{array}\right)
$$

where $\mathbf{A}_{r}^{L}$ and $\mathbf{B}_{r}^{L}$ are $M$-vectors obtained as appropriately sampled versions of $A_{r}^{L}$ and $B_{r}^{L}$, the shifted Legendre polynomials defined in Section 2.3. (i.e. $\left[\mathbf{A}_{r}^{L}\right]_{k}=A_{r}^{L}\left(\cos \theta_{k}\right)$, where $\theta_{k}=\frac{\pi(2 k+1)}{2 M}$.)

The Legendre transform (LT) of size $M$ and shift $L$ computed from input data vectors $\mathbf{z}_{0}$ and $\mathbf{z}_{1}$ each of length $M$, is the matrix-vector multiplication

$$
\begin{equation*}
L T_{M}^{L} \cdot \mathbf{Z} \tag{28}
\end{equation*}
$$

for $\mathbf{Z}=\binom{\mathbf{z}_{1}}{\mathbf{z}_{0}}$.
The original Legendre transform of a data vector $\mathbf{s}$ of length $N$ defined in equation (24) may be written using Definition 1 as $L T_{N}^{0} \cdot\binom{\mathbf{s}}{\mathbf{0}}$ with $\mathbf{0}$ denoting the 0 -vector of length $N$. This follows from the definitions of the 0 -shifted Legendre polynomials (14),

$$
\begin{equation*}
\left\langle\mathbf{A}_{r}^{0}, \mathbf{s}\right\rangle+\left\langle\mathbf{B}_{r}^{0}, \mathbf{0}\right\rangle=\left\langle\mathbf{P}_{r}, \mathbf{s}\right\rangle+0=\left\langle\mathbf{P}_{r}, \mathbf{s}\right\rangle . \tag{29}
\end{equation*}
$$

Likewise, the description of the splitting of the problem into low and high degree transforms may now be restated formally: the LT of size $N$ may instead be computed as two LTs, each of size $N / 2$. In particular, the low and high order transforms are computed as $L T_{\frac{N}{2}}^{0} \cdot\binom{\mathcal{L} \mathbf{s}}{\mathcal{L} \mathbf{0}}$ and $L T_{\frac{N}{2}}^{\frac{N}{2}} \cdot\binom{\mathcal{L}^{\mathbf{s}^{\frac{N}{2}}}}{\mathcal{L} \mathbf{s}^{\frac{N}{2}-1}}$ respectively. In this way the original discrete Legendre transform may be reduced to the computation of two smaller LTs.

In order to actually carry this out, the input for the two smaller LTs must be calculated (efficiently) from the data for the original full sized problem. To give this step a uniform description we write the input as $\mathbf{Z}=\binom{\mathbf{s}^{0}}{\mathbf{s}^{-1}}=\binom{\mathbf{z}_{1}}{\mathbf{z}_{0}}$ with $\mathbf{s}^{0}=\mathbf{s}$ and $\mathbf{s}^{-1}=\mathbf{0}$. We define splitting operators $S^{\epsilon}(\epsilon=0,1)$ which appropriately weight and then lowpass the input to provide the data for the low and high order transforms respectively according to

$$
\begin{align*}
\mathbf{Z}^{\epsilon} & =\mathcal{S}^{\epsilon} \mathbf{Z}  \tag{30}\\
& =(\mathcal{L} \oplus \mathcal{L}) \cdot \mathcal{M}^{\epsilon} \mathbf{Z}  \tag{31}\\
& =\left(\begin{array}{cc}
\mathcal{L} & \mathbf{0} \\
\mathbf{0} & \mathcal{L}
\end{array}\right) \cdot\left(\begin{array}{cc}
\operatorname{diag} \mathbf{A}_{\epsilon \frac{N}{2}}^{0} & \operatorname{diag} \mathbf{B}_{\epsilon \frac{N}{2}}^{0} \\
\operatorname{diag} \mathbf{A}_{\epsilon \frac{N}{2}-1}^{0} & \operatorname{diag} \mathbf{B}_{\epsilon \frac{N}{2}-1}^{0}
\end{array}\right) \cdot\binom{\mathbf{z}_{1}}{\mathbf{z}_{0}} . \tag{32}
\end{align*}
$$

Thus, the matrix $M^{\epsilon}$ is a $2 \times 2$ block matrix with $N \times N$ diagonal blocks defined by sampled values of the shifted Legendre polynomials. By convention, $\mathcal{L}$ is the $N / 2 \times N$ lowpass operator. Notice that when $\epsilon=0, \mathcal{M}^{0}$ reduces to the identity matrix. Figure 6 represents this first divide and conquer step.

To count carefully the number of operations used to accomplish the reduction recall first that $\mathbf{A}_{0}^{0}$ is a vector of all 1's, while $\mathbf{z}_{0}, \mathbf{B}_{r}^{0}$ and all negative subscripts give 0 -vectors. Using Lemma 2 it is now easy


Figure 6: Illustration of the first split. The Legendre transform computed as two half-sized Legendre transforms of modified data.
to see that at most $2 N+3 N \log N+3 \frac{N}{2} \log \frac{N}{2}$ operations are required. If we stopped subdividing at this point and efficiently completed the computation by working with the cosine transforms, then assuming that the cosine transforms of the shifted Legendre polynomials were stored, at most $\frac{9}{4} N \log \frac{N}{2}+\frac{3 N^{2}}{8}-N$ operations would be needed to compute in this manner. This is opposed to $\frac{3}{2} N \log N+\frac{N}{2}(N-1)$ for a complete semi-naive approach. Hence we obtain an advantage for $N \geq 256$. Of course, the real algorithmic advantage is obtained by performing this split recursively. This is explained in the next subsection.

Before leaving this subsection, we record the complexity of the application of the shifted Legendre matrices. It follows from our observations in Section 2 that there is a semi-naive approach to this calculation, as observed in the following lemma.

Lemma 4 Assuming that the cosine transforms $\mathcal{C} \mathbf{A}_{r}^{L}$ and $\mathcal{C} \mathbf{B}_{r}^{L}$ are prestored $(0 \leq r<M)$, then an $L T$ of size $M$ and shift $L$ can be computed in at most $M^{2}+3 M \log M+2 M$ operations.

Proof: Using [47], at most $2 \cdot\left(\frac{3}{2} M \log M\right)$ operations are needed to compute the DCTs of $\mathbf{z}_{0}$ and $\mathbf{z}_{1}$. Having done that, an additional $2 \cdot \frac{M(M+1)}{2}$ are needed to compute the inner products $\left\{\left\langle\mathcal{C} \mathbf{A}_{r}^{L}, \mathcal{C} \mathbf{z}_{1}\right\rangle,\left\langle\mathcal{C} \mathbf{B}_{r}^{L}, \mathcal{C} \mathbf{z}_{0}\right\rangle \mid r=0, \ldots, M-1\right\}$ and an additional $M$ operations to add together each pair.

Remark. Lemma 4 is precisely the computation of a LT of size $M$ and shift $L$ by the semi-naive algorithm. Notice that it too may be cast as a particular matrix factorization,

$$
\begin{equation*}
L T_{M}^{L} \cdot \mathbf{Z}=\left(L T_{M}^{L}\left(\mathcal{C}_{M} \oplus \mathcal{C}_{M}\right)^{t}\right) \cdot\left(\mathcal{C}_{M} \oplus \mathcal{C}_{M}\right) \cdot \mathbf{Z} \tag{33}
\end{equation*}
$$

recalling that we use the orthogonal matrix $\mathcal{C}_{M}(17)$ to effect the DCT and assume that the first factor is precomputed.

### 3.3 Recursive subdivision

For a complete recursive subdivision we need the following lemma.
Lemma 5 (Splitting Lemma) Let $M$ be a positive integer divisible by two.

- (i) A Legendre transform of size $M$ can be computed as two Legendre transforms of size $M / 2$. Specifically, multiplication of $L T_{M}^{L}$ against a given input can instead be computed as the separate multiplications of $L T_{\frac{M}{2}}^{L}$ and $L T_{\frac{M}{2}}^{L+\frac{M}{2}}$ against new inputs obtained from the original input.
- (ii) Let $\mathbf{Z}=\binom{\mathbf{z}_{1}}{\mathbf{z}_{0}}$ be the initial data for a LT of size $M$ and shift $L$. Then at most $8 M+$ $6 M \log M+3 M \log \frac{M}{2}$ operations are needed to compute the necessary inputs for the pair of LTs at size $M / 2$ which together, compute the original LT. In particular, if $\mathbf{Z}^{\epsilon}=\binom{\mathbf{z}_{1}^{\epsilon}}{\mathbf{z}_{0}^{\epsilon}}(\epsilon=0,1)$ are the necessary input for the half-sized LTs which compute the low order $(\epsilon=0)$ and high order $(\epsilon=1)$ transforms, then

$$
\begin{align*}
\mathbf{Z}^{\epsilon} & =\mathcal{S}_{L, M}^{\epsilon} \cdot \mathbf{Z}  \tag{34}\\
& =(\mathcal{L} \oplus \mathcal{L}) \cdot \mathcal{M}^{\epsilon} \cdot \mathbf{Z}  \tag{35}\\
& =\left(\begin{array}{cc}
\mathcal{L} & \mathbf{0} \\
\mathbf{0} & \mathcal{L}
\end{array}\right) \cdot\left(\begin{array}{cc}
\operatorname{diag} \mathbf{A}_{\epsilon \frac{M}{2}}^{L} & \operatorname{diag} \mathbf{B}_{\epsilon \frac{M}{2}}^{L} \\
\operatorname{diag} \mathbf{A}_{\epsilon \frac{M}{2}-1}^{L} & \operatorname{diag} \mathbf{B}_{\epsilon \frac{M}{2}-1}^{L}
\end{array}\right)\binom{\mathbf{z}_{1}}{\mathbf{z}_{0}} . \tag{36}
\end{align*}
$$

- (iii) The matrix $L T_{M}^{L}$ has the factorization

where $\mathcal{S}_{L, M}^{0}$ and $\mathcal{S}_{L, M}^{1}$ are defined as in (ii).
Proof: The low degree projections in the size $M$ problem, $\left\langle\mathbf{A}_{k}^{L}, \mathbf{z}_{1}\right\rangle+\left\langle\mathbf{B}_{k}^{L}, \mathbf{z}_{0}\right\rangle, 0 \leq k<M / 2$ only involve shifted Legendre polynomials of degree less than $M / 2$. Consequently, according to Lemma 1 these inner products may be computed as inner products of the lowpassed data with subsampled versions of the shifted Legendre polynomials, thereby reducing it to a LT of size $M / 2$.

To reduce the set of high order projections to a LT of half the size, we apply the recurrence formula (13) to obtain

$$
\begin{aligned}
& \mathbf{A}_{\frac{M}{2}+k}^{L}=\mathbf{A}_{k}^{L+\frac{M}{2}} \mathbf{A}_{\frac{M}{2}}^{L}+\mathbf{B}_{k}^{L+\frac{M}{2}} \mathbf{A}_{\frac{M}{2}-1}^{L}, \\
& \mathbf{B}_{\frac{M}{2}+k}^{L}=\mathbf{A}_{k}^{L+\frac{M}{2}} \mathbf{B}_{\frac{M}{2}}^{L}+\mathbf{B}_{k}^{L+\frac{M}{2}} \mathbf{B}_{\frac{M}{2}-1}^{L} .
\end{aligned}
$$

Consequently, we rewrite the high order transfoms as

$$
\begin{equation*}
\left\langle\mathbf{A}_{\frac{M}{2}+k}^{L}, \mathbf{z}_{1}\right\rangle+\left\langle\mathbf{B}_{\frac{M}{2}+k}^{L}, \mathbf{z}_{0}\right\rangle=\left\langle\mathbf{A}_{k}^{L+\frac{M}{2}}, \mathbf{A}_{\frac{M}{2}}^{L} \mathbf{z}_{1}+\mathbf{B}_{\frac{M}{2}}^{L} \mathbf{z}_{0}\right\rangle+\left\langle\mathbf{B}_{k}^{L+\frac{M}{2}}, \mathbf{A}_{\frac{M}{2}-1}^{L} \mathbf{z}_{1}+\mathbf{B}_{\frac{M}{2}-1}^{L} \mathbf{z}_{0}\right\rangle . \tag{38}
\end{equation*}
$$

Since both $A_{k}^{L+\frac{M}{2}}$ and $B_{k}^{L+\frac{M}{2}}$ have degree less than $\frac{M}{2}$ we can lowpass both sides of each inner product in (38) and obtain

$$
\begin{equation*}
\left\langle\mathbf{A}_{\frac{M}{2}+k}^{L}, \mathbf{z}_{1}\right\rangle+\left\langle\mathbf{B}_{\frac{M}{2}+k}^{L}, \mathbf{z}_{0}\right\rangle=\left\langle\mathcal{L} \mathbf{A}_{k}^{L+\frac{M}{2}}, \mathcal{L}\left(\mathbf{A}_{\frac{M}{2}}^{L} \mathbf{z}_{1}+\mathbf{B}_{\frac{M}{2}}^{L} \mathbf{z}_{0}\right)\right\rangle+\left\langle\mathcal{L} \mathbf{B}_{k}^{L+\frac{M}{2}}, \mathcal{L}\left(\mathbf{A}_{\frac{M}{2}-1}^{L} \mathbf{z}_{1}+\mathbf{B}_{\frac{M}{2}-1}^{L} \mathbf{z}_{0}\right)\right\rangle . \tag{39}
\end{equation*}
$$

This completes the proof of (i). The proof of (ii) follows immediately. The stated form of $\mathbf{Z}^{\epsilon}$ follows from simply rewriting the collection of linear equations describing the new data. The complexity result follows from the fact that application of $\mathcal{M}^{\epsilon}$ requires at most $4 M$ operations and application of $\mathcal{L} \oplus \mathcal{L}$, an additional $3\left(M \log M+\frac{M}{2} \log \frac{M}{2}\right)$ operations (cf. Lemma 2).

Finally, (iii) is simply a restatement of (ii) in the setting of matrix arithmetic.

Using Lemma 5 we can now describe the full algorithm in a succinct way. Starting with the original Legendre transform written as $L T_{N}^{0}$, part (iii) implies that this matrix factors as

Now we apply Lemma 5 to the half-sized LT matrices, $L T_{\frac{N}{2}}^{0}$ and $L T_{\frac{N}{2}}^{M}$ producing the factorization of $L T_{N}^{0}$ as


To keep track of the factorization as we continue splitting, we'll use a binary tree-based indexing notation, indicated schematically in Figure 7. Given initial data s of length $N$, define the block vector of length $2 N, \mathbf{Z}=\binom{\mathbf{s}}{\mathbf{0}}$. We will factor the matrix $L T_{N}^{0}$ as a product of $\log N$ block diagonal matrices such that for each $k, 1 \leq k \leq \log N$, the $k^{t h}$ matrix has $2^{k}$ blocks made up of splitting matrices of the type described in the Splitting Lemma (Lemma 5). These splitting matrices will be indexed by binary $k$-tuples $\vec{\epsilon}=\left(\epsilon_{1}, \epsilon_{2}, \ldots, \epsilon_{k}\right), \vec{\epsilon} \in\{0,1\}^{k}$, denoted as $\mathcal{S}^{\vec{\epsilon}}$ and defined by

$$
\begin{equation*}
\mathcal{S}_{\vec{\epsilon}}=(\mathcal{L} \oplus \mathcal{L}) \cdot M^{\vec{\epsilon}} \tag{41}
\end{equation*}
$$

for

$$
\mathcal{M}^{\vec{\epsilon}}=\mathcal{M}^{\left(\epsilon_{1}, \epsilon_{2}, \ldots, \epsilon_{k}\right)}=\left(\begin{array}{cc}
\operatorname{diag} \mathbf{A}_{\epsilon_{k} \frac{N}{2^{k-1}}}^{\left.\stackrel{\sim}{1}, \ldots, \epsilon_{k-1}\right)} & 2 \cdot \frac{N}{2^{k-1}}  \tag{42}\\
\operatorname{diag} \mathbf{B}_{\epsilon_{k} \frac{N}{2^{k-1}}}^{\longrightarrow\left(\epsilon_{1}, \ldots, \epsilon_{k-1}\right)} \\
\operatorname{diag} \mathbf{A}_{\epsilon_{k} \frac{N}{2^{k-1}-1}}^{L\left(\epsilon_{1}, \ldots, \epsilon_{k-1}\right)} & \operatorname{diag} \mathbf{B}_{\epsilon_{k} \frac{N}{2^{k-1}-1}}^{L\left(\epsilon_{1}, \ldots, \epsilon_{k-1}\right)}
\end{array}\right) \quad \begin{gathered}
\\
2 \cdot \frac{N}{2^{k-1}}
\end{gathered} .
$$

and $L(\vec{\epsilon})=L\left(\epsilon_{1}, \epsilon_{2}, \ldots, \epsilon_{m}\right)=\epsilon_{1} \frac{N}{2}+\epsilon_{2} \frac{N}{4}+\ldots+\epsilon_{m} \frac{N}{2^{m}}$.


Figure 7: Schematic illustration of the computation of the Legendre transform by recursive splitting.
With this notation and Lemma 5 we now have the following.

Theorem 2 Let $N=2^{r}$ and let $\mathbf{s}$ be any vector of length $N$.

- (i) The Legendre transform of $\mathbf{s}$ may be computed from $L T_{N}^{0} \cdot\binom{\mathbf{s}}{\mathbf{0}}$ in $O\left(N \log ^{2} N\right)$ operations via the factorization: $L T_{N}^{0}=\mathcal{E} \mathcal{F}_{r-1} \cdots \mathcal{F}_{1}$, where
and the $2^{j}$ splitting matrices $\mathcal{S}^{\vec{\epsilon}}$ are defined as in (41), and

$$
\mathcal{E}=L T_{2}^{N-2} \oplus L T_{2}^{N-4} \oplus \cdots \oplus L T_{2}^{2}
$$

- (ii) Assuming the precomputation and storage of the masks $\mathcal{M}^{\vec{\epsilon}}$ (see (42)) as well as the cosine transforms of the shifted Legendre polynomials $A_{r}^{L(\bar{\epsilon})}, B_{r}^{L(\bar{\epsilon})}$, then the complexity of the Legendre transform of $\mathbf{s}$ is $O\left(N \log ^{2} N\right)$. More precisely, at most

$$
8 \cdot \frac{N}{2}+\sum_{a=1}^{r-1}\left(2^{a}-1\right) \cdot 2 \cdot \frac{3}{2}\left[\frac{N}{2^{a}} \log \frac{N}{2^{a}}+\frac{N}{2^{a-1}} \log \frac{N}{2^{a-1}}\right]+\left(2^{a}-2\right) \cdot 4 \cdot \frac{N}{2^{a-1}}
$$

operations are needed to compute the Legendre transform of $\mathbf{s}$.
Proof: Part (i) follows directly from a recursive application of Lemma 5. As for (ii), the first term in the computation follows from the fact that $\mathcal{E}$ is block diagonal with $\frac{N}{2} 2 \times 4$ matrices on the diagonal, each requiring 8 operations to multiply by a vector of length 4 . The summation which is the second term is the complexity of the successive multiplications of the $\mathcal{F}$. This follows directly from the form of the $\mathcal{S}^{\vec{\epsilon}}$. In particular noting that in general (for $\vec{\epsilon} \in\{0,1\}^{a}, \vec{\epsilon} \neq \mathbf{0}$ ), $2 \cdot \frac{3}{2}\left[\frac{N}{2^{a}} \log \frac{N}{2^{a}}+\frac{N}{2^{a-1}} \log \frac{N}{2^{a-1}}\right]+$ $4 \cdot \frac{N}{2^{a-1}}$ operations are required to apply $\mathcal{S}^{\vec{\epsilon}}$ to an arbitrary input. When $\vec{\epsilon}=\mathbf{0}$, we may take advantage of the fact that the lower half of the input vector will be all 0 and that the relevant diagonal blocks of $\mathcal{M}^{\vec{\epsilon}}$ will be the identity and all 0 's as well. Summing the complexities at each level gives (ii).

The inductive nature of Theorem 2 shows that the algorithm is simply given by the successive application of the splitting operators to the (changing) input. In practice, this splitting is applied as long as it offers computational advantage. At this point the calculation may be concluded by applying the shifted Legendre matrices according to the semi-naive algorithm.

## 4 Fast Fourier Transform, Inverse Transform, and Convolution for the 2-Sphere

The algorithm presented in the previous section enables us to write a fast algorithm for the efficient calculation of spherical harmonic expansion, or Fourier transform, of a function on $S^{2}$ in terms of its samples on a regular grid. This is readily converted into a fast inversion algorithm (transforming a set of Fourier coefficients into sample values) improving on the results of [16]. These two algorithms combine to give an improved convolution algorithm.

### 4.1 Fast Fourier transform for $S^{2}$

Collecting the results of previous sections gives the following theorem.
Theorem 3 Let $B$ be a power of 2 and $n=B^{2}$. If $f(\theta, \phi)$ is in the span of $\left\{Y_{\ell}^{m}| | m \mid \leq \ell<B,\right\}$, then the $n$ Fourier coefficients $\hat{f}(\ell, m)$ for $\ell<B,|m| \leq \ell$ can be computed in $O\left(n \log ^{2} n\right)$ operations from the $4 n$ sampled values $f((2 j+1) \pi / 4 B, 2 \pi k / 2 B), 0 \leq j, k<2 B-1$, using a precomputed data structure of size $O\left(n \log ^{2} n\right)$.

Proof: From the Sampling Theorem we know that the Fourier transform may be computed by means of finite sums,

$$
\hat{f}(\ell, m)=\sum_{j=0}^{2 B-1} \sum_{k=0}^{2 B-1} a_{j}^{(B)} f\left(\theta_{j}, \phi_{k}\right) Y_{\ell}^{m}\left(\theta_{j}, \phi_{k}\right)
$$

where $\theta_{j}=\pi(2 j+1) / 4 B$ and $\phi_{k}=2 \pi k / 2 B$, and the $a_{j}^{(B)}$ are as defined in the Sampling Theorem. Rewriting, we obtain

$$
\hat{f}(\ell, m)=q_{\ell}^{m} \sum_{j=0}^{2 B-1} a_{j}^{(B)} P_{\ell}^{m}\left(\cos \theta_{j}\right) \sum_{k=0}^{2 B-1} e^{-i m \phi_{k}} f\left(\theta_{j}, \phi_{k}\right),
$$

where the $q_{\ell}^{m}$ are the normalization coefficients for the spherical harmonics. The inner sums (defined as $\left.\breve{f}\left(\theta_{j}, m\right)\right)$ are computed for each fixed $j$ and for all $m$ in the appropriate range by means of a fast Fourier transform. This will require $O\left(B^{2} \log B\right)$ operations.

It remains to compute

$$
\hat{f}(\ell, m)=q_{\ell}^{m} \sum_{j=0}^{2 B-1} a_{j}^{(B)} \breve{f}\left(\theta_{j}, m\right) P_{\ell}^{m}\left(\cos \theta_{j}\right)
$$

which has the form of an associated Legendre transform for each fixed $m$. Each transform can be accomplished in $O\left(B \log ^{2} B\right)$ operations, and since $m$ ranges over $2 B-1$ values, the total number of operations needed to compute all values of $\hat{f}(\ell, m)$ is $O\left(B^{2} \log ^{2} B\right)$.

### 4.2 Fast inversion

For our purposes, the inverse transform is the map which takes as input a set of complex numbers $c_{\ell, m}$, interpreted as Fourier coefficients in a spherical harmonic expansion and returns a set of sample values. The discussion is restricted to band-limited functions so we consider only the case in which the $c_{\ell, m}$ vanish for $\ell \geq B$ for some bandwidth $B>0$.

Definition 2 The discrete inverse spherical Fourier transform with bandwidth $B=2^{b}$ maps a collection of complex coefficients $c_{\ell, m}, 0 \leq m \leq \ell<B$ to a collection of samples values via the formula

$$
f\left(\theta_{j}, \phi_{k}\right)=\sum_{\ell \geq 0} \sum_{|m| \leq \ell} c_{\ell, m} Y_{\ell}^{m}\left(\theta_{j}, \phi_{k}\right),
$$

where the recovery is on the equiangular grid of the sampling theorem, $\theta_{j}=\frac{(j+1 / 2) \pi}{2 B}, j=0, \ldots, 2 B-1$; and $\phi_{k}=\frac{(k+1 / 2) \pi}{B}, k=0, \ldots, 2 B-1$.

Setting $n=B^{2}$, this map transforms a collection of $O(n)$ Fourier coefficients to $O(n)$ samples of the function with the associated spherical harmonic expansion. Computed directly, this would require $O\left(n^{2}\right)$ calculations, or $O\left(n^{3 / 2}\right)$ by using a little reorganization (cf. [16], Theorem 9). Our new result is that in fact, this too may be accomplished in $O\left(n \log ^{2} n\right)$ operations.

Theorem 4 Let $B=2^{r}$ be a fixed bandwidth. The discrete inverse spherical Fourier transform for this bandwidth may be computed in $O\left(n \log ^{2} n\right)$ operations, where $n=B^{2}$.

Proof: Let the Fourier coefficients $c_{\ell, m}, 0 \leq m \leq \ell<B$ be given. Our goal is to compute the collection of samples of the inverse spherical Fourier transform

$$
f\left(\theta_{j}, \phi_{k}\right)=\sum_{\ell=0}^{B-1} \sum_{|m| \leq \ell} c_{\ell, m} Y_{\ell}^{m}\left(\theta_{j}, \phi_{k}\right)
$$

Using the definition of the $Y_{\ell}^{m}$ we rewrite this as

$$
\begin{equation*}
\sum_{|m|<B} e^{-i m \phi_{j}} \sum_{\ell=|m|}^{B-1} c_{\ell, m} q_{\ell}^{m} P_{\ell}^{m}\left(\cos \theta_{k}\right) \tag{44}
\end{equation*}
$$

where $q_{\ell}^{m}$ denotes the appropriate normalization constant. Let $h\left(\theta_{k}, m\right)$ denote the inner sum

$$
h\left(\theta_{k}, m\right)=\sum_{\ell=|m|}^{B-1} c_{\ell, m} q_{\ell}^{m} P_{\ell}^{m}\left(\cos \theta_{k}\right) .
$$

Computed directly, each of the $4 B^{2}$ values for $h$ would require a summation of $O(B)$ terms for a naive complexity of $O\left(B^{3}\right)$. Instead we show how a simple "adaptation" of the ideas of Section 3 yields a more efficient algorithm. Observe that the column vector $\mathbf{h}^{m}=\left(h\left(\theta_{0}, m\right), \ldots, h\left(\theta_{2 B-1}, m\right)\right)^{t}$ is obtained as a matrix-vector product

$$
\begin{aligned}
\mathbf{h}^{m}=\left(\begin{array}{c}
h\left(\theta_{0}, m\right) \\
h\left(\theta_{1}, m\right) \\
\vdots \\
h\left(\theta_{2 B-1}, m\right)
\end{array}\right) & =\left(\begin{array}{ccc}
P_{m}^{m}\left(\cos \theta_{0}\right) & \cdots & P_{B-1}^{m}\left(\cos \theta_{0}\right) \\
P_{m}^{m}\left(\cos \theta_{1}\right) & \cdots & P_{B-1}^{m}\left(\cos \theta_{1}\right) \\
\vdots & \vdots & \vdots \\
P_{m}^{m}\left(\cos \theta_{2 B-1}\right) & \cdots & P_{B-1}^{m}\left(\cos \theta_{2 B-1}\right)
\end{array}\right) \cdot\left(\begin{array}{c}
c_{m, m} q_{m}^{m} \\
c_{m+1, m} q_{m}^{m+1} \\
\vdots \\
c_{B-1, m} q_{m}^{B-1}
\end{array}\right) \\
& =\tilde{\mathcal{P}}_{B}^{m} \cdot \tilde{\mathbf{c}}^{\tilde{m}}
\end{aligned}
$$

where $\tilde{\mathcal{P}}_{B}^{m}$ and $\tilde{\mathbf{c}}^{\tilde{m}}$ are defined by the above equations. Notice that $\left(\tilde{\mathcal{P}}_{B}^{m}\right)^{t}$ is the Vandermonde-like matrix for the order $m$ associated Legendre functions, and as such is the transpose of the matrix whose application to a fixed set of sample values yields the associated Legendre transform.

The key observation we now make is that the algorithm described in Section 3 for computing the Fourier transform gives a factorization of the matrix $\tilde{\mathcal{P}}_{B}^{m}$ as a product of matrices whose structure admit efficient multiplication against an arbitary vector. More precisely, Section 3 shows that we have a factorization

$$
\tilde{\mathcal{P}}_{B}^{m}=\mathcal{M}_{r-1} \cdots \mathcal{M}_{1} \mathcal{M}_{0}
$$

such that a total of $O\left(B \log ^{2} B\right)$ operations are required to successively multiply the $\mathcal{M}_{i}$ against an arbitrary vector. A similar complexity result holds if instead, the transpose of each of the matrices is considered, and the product is reordered. Since $\left(\tilde{\mathcal{P}}_{B}^{m}\right)^{t}=\mathcal{M}_{0}^{t} \mathcal{M}_{1}^{t} \cdots \mathcal{M}_{r-1}^{t}$ and it follows that the inverse transform may be computed in $O\left(B \log ^{2} B\right)$ operations as well.

To complete the computation, an abelian FFT is performed to compute the sums $\sum_{m} h\left(\theta_{k}, m\right) e^{-i m \phi_{k}}$ in $O(B \log B)$ operations for each $k$, for a total additional cost of $O\left(B^{2} \log B\right)$. Denoting the total number of samples by $n=B^{2}$, we may write the total cost as $O\left(n \log ^{2} n\right)$, as desired.

### 4.3 A fast spherical convolution algorithm

Taken together, the fast forward and inverse transforms allow for the convolution of two band-limited functions in $L^{2}\left(S^{2}\right)$ to be computed efficiently - and exactly (in exact arithmetic).

Defining convolution for two functions in $L^{2}\left(S^{2}\right)$ uses the structure of $S^{2}$ as a quotient of the group $S O(3)$. Generalizing the case of the circle, the left convolution of $h$ by $f$ for $f, h \in L^{2}\left(S^{2}\right)$, is defined as

$$
\begin{equation*}
f \star h(\omega)=\int_{g \in S O(3)} f(g \eta) h\left(g^{-1} \omega\right) d g \tag{45}
\end{equation*}
$$

In (45) $d g$ denotes the (essentially) unique invariant volume form on $S O(3)$ and $\eta$ denotes the north pole.

By combining the fast expansion and synthesis algorithms of Sections 3 and 4.2, a fast convolution algorithm is obtained. Given the sample values of two functions $f, g \in L^{2}\left(S^{2}\right)$ of bandwidth $B$, we now give an $O\left(n \log ^{2} n\right)$ algorithm $\left(n=B^{2}\right)$ to compute the sample values of the convolution $f \star g$ thereby improving on the $O\left(n^{1.5}\right)$ algorithm in [16].

As in the more familiar case of convolution on the circle via the abelian FFT [7], the spherical convolution algorithm may be decomposed into three basic steps:

- (1) Computation of a forward transform;
- (2) Pointwise multiplication of the appropriate transforms;
- (3) Computation of the inverse transform of the result of step (2).

We have shown how to accomplish steps (1) and (3) efficiently. We need only satisfy step (2). This uses the following relationship between the transform and the convolution.

Theorem 5 [[16], Theorem 1] Let $f, h \in L^{2}\left(S^{2}\right)$. Then

$$
(\widehat{f \star h})(l, m)=2 \pi \sqrt{\frac{4 \pi}{2 l+1}} \hat{f}(l, m) \hat{h}(l, 0) .
$$

In particular, the convolution of functions of bandwidth $B$, yields a function of bandwidth $B$.

Theorem 6 Let $f, g \in L^{2}\left(S^{2}\right)$ such that by $f$ and $g$ are band-limited with bandwidth $B$. Then the $n=O\left(B^{2}\right)$ sample values of the convolution $f \star g$ at the points $\left(\theta_{j}, \phi_{k}\right)$ where $\theta_{j}=\frac{\pi(j+1 / 2)}{2 B}$ and $\phi_{k}=\frac{\pi(k+1 / 2)}{B}$ may be computed in $O\left(n \log ^{2} n\right)$ operations, versus the $O\left(n^{2}\right)$ operations required by direct computation.

Proof: Using the algorithm described in Section 3, compute the Fourier coefficients of $f$ and $g$. Compute the pointwise products according to Theorem 5 for the convolution $f \star g$ and finally, compute the inverse transform for the Fourier coefficients of the convolution according to Theorem 4.

## 5 Variations of the Basic Algorithm

The recursive splitting described in Section 3 computes the discrete Legendre transform of a vector of length $N$ in $O\left(N \log ^{2} N\right)$ operations, assuming that the splitting is carried out as far as possible, i.e., to $\log N$ levels. By applying the analogous algorithms to the samples of a function on $S^{2}$ with bandwidth $B$, for the remaining Legendre functions $P_{\ell}^{m}, m \neq 0$, the projections onto the spherical harmonics $Y_{\ell}^{m}$, for each $\ell$ in the range $0 \leq \ell<B$ and $|m| \leq \ell$ are then efficiently obtained. This is the "basic" Driscoll-Healy algorithm, which we subsequently refer to as the DH algorithm or DH. Our discussion shows that DH computes the Fourier coefficients of a function on $S^{2}$ of bandwidth $B$ in $O\left(n \log ^{2} n\right)$ operations for $n=B^{2}$.

This is an asymptotic result, exact in exact arithmetic. For application to actual problems of moderate size, we must consider issues of numerical reliability and computational efficiency in a floating point implementation. We must also demonstrate that an implementation can obtain real speed-ups over existing algorithms at useful problem sizes.

We now turn to a discussion of some simple variations of our basic approach which we use to obtain fast and reliable algorithms at various moderate problem sizes. Again we specialize to the Legendre polynomial case.

- Variation 1. Using the reverse recurrence: The DH-Mid algorithm

This simple variation reduces the complexity by approximately one-half. The three-term recurrence $(9)$ is a forward recurrence. This also gives rise to the reverse recurrence

$$
\begin{equation*}
P_{\ell-1}(x)=\frac{2 \ell+1}{\ell} x P_{\ell}(x)-\frac{\ell+1}{\ell} P_{\ell+1}(x) . \tag{46}
\end{equation*}
$$

Using (46) in a way analogous to the use of (9) we may define the reverse shifted Legendre polynomials in analogy with the functions $A_{r}^{L}$ and $B_{r}^{L}$ defined for the forward direction (10). For any fixed level $L$, iterating the recurrence formula (46) back $r$ steps produces trigonometric polynomials $A_{-r}^{L}$ and $B_{-r}^{L}$ such that

$$
\begin{equation*}
P_{L-r}(\cos \theta)=A_{-r}^{L}(\cos \theta) P_{L-1}(\cos \theta)+B_{-r}^{L}(\cos \theta) P_{L}(\cos \theta) \tag{47}
\end{equation*}
$$

for $r \geq 2$. This may be used in several ways to reduce computations.
First, this can provide a "balanced" version of the semi-naive algorithm. Section 3 shows that, given the vectors $\mathbf{s}^{L-1}$ and $\mathbf{s}^{L}$, the Legendre coefficient of degree $L+r$ can be computed in $O(r)$ operations by passing to the cosine transform domain and forming the inner product with the cosine transform vector of the shifted Legendre functions of degree at most $r$. Similarly, the Legendre coefficient of degree $L-r$ can be computed in $O(r)$ operations using reverse shifted Legendre polynomials of degree
bound $r$. Again, simply form the inner products of the cosine transforms of $\mathbf{s}^{L-1}$ and $\mathbf{s}^{L}$ with the cosine transforms of the reverse shifted Legendre polynomials. In this way, we can compute a seminaive algorithm for the Legendre coefficients in the range $[L-r, L+r]$ which is "balanced" in the sense that the number of operations required to compute any given transform in this interval depends only on the distance of the index from $L$. Furthermore, only the cosine transform coefficients of $\mathbf{s}^{L-1}$ and $\mathbf{s}^{L}$ of degree $r$ or less are used to compute these Legendre coefficients.

We may also use this approach to expedite the splitting steps of the algorithm by starting in the middle of the sequence and moving out to both the left and right from the initial data. This is done as follows. From the initial data $\mathbf{s}^{0}$ we compute the lowpassed subsampled sequences $\mathbf{s}^{N / 2}, \mathbf{s}^{N / 2+1}$, each of length $N / 2$. Then $\mathbf{s}^{3 N / 4}$ and $\mathbf{s}^{3 N / 4+1}$ are computed as before, but by using the three-term recurrence in the reverse direction (46) the sequences $\mathbf{s}^{N / 4}$ and $\mathbf{s}^{N / 4+1}$ can be computed. The additional savings come from using the initial data to compute four new sequences instead of two.

At the next stage, each of the pairs $\mathbf{s}^{N / 4}, \mathbf{s}^{N / 4+1}$ and $\mathbf{s}^{3 N / 4}, \mathbf{s}^{3 N / 4+1}$ then act similarly as initial data for obtaining the sequences $\mathbf{s}^{j}, j=N / 8, N / 8+1,3 N / 8,3 N / 8+1$ and $\mathbf{s}^{j}, j=5 N / 8,5 N / 8+$ $1,7 N / 8,7 N / 8+1$ respectively. The recursion continues down to some base case. We call this algorithm the DH-Mid algorithm.

Asymptotically and theoretically, a full divide and conquer strategy is optimal. However, in actual implementations overhead costs can accrue and it is often the case that for smaller problem sizes divide and conquer is no longer advantageous. Below this "breakeven point" a more direct approach may be faster. We may successfully address this issue with some of the simple variants discussed below. This basic idea uses a simple truncation of the splitting of the basic algorithm at an appropriate level.

- Variation 2. The Bounded DH-Mid Algorithm

The semi-naive algorithm described in Section 3.2 is quite competitive in speed for moderate problems ( $N \leq 256$ ). We need to take this into consideration when optimizing algorithms for problems of moderate size encountered in applications ( $256 \leq N \leq 1024$ ). In this range, we will find it useful to stop the splitting when the resulting subproblems reach a given size.

Consequently, in this further variation, we use DH-Mid but only split down to a fixed level $k$, i.e., to a point at which all of the most recently computed sequences $\mathbf{s}^{j}$ have length $2 k$. At this point we switch over to a semi-naive approach to compute the remaining Legendre transforms. For example, suppose the recursion halts after we have computed a group of length $2 k$ sequences which includes $\mathbf{s}^{r}, \mathbf{s}^{r+1}$ and $\mathbf{s}^{r+k}, \mathbf{s}^{r+k+1}$. By using a semi-naive method and the shifted Legendre and reverse shifted Legendre functions, $\mathbf{s}^{r}, \mathbf{s}^{r+1}$ can be used to compute the vectors $\mathbf{s}^{j}, j=r-$ $k / 2+1, r-k / 2+2, \ldots, r-1, r+2, \ldots, r+k / 2-1$. Similarly, $\mathbf{s}^{r+k}, \mathbf{s}^{r+k+1}$ can be used to compute $\mathbf{s}^{j}, j=r+k / 2, r+k / 2+1, \ldots, r+k-1, r+k+2, \ldots, r+k+k / 2-1$. The value of $k$ can be chosen for a given problem size so as to minimize the number of operations.

- Variation 3. Simple split algorithm

Here the idea is to immediately split the original problem of length $N$ into $C$ subproblems and run a semi-naive approach on each of the resulting subproblems. The input data for each of the subproblems is computed directly from the original input data by multiplying all $N$ samples of $\mathbf{s}$ onto the appropriate samples of the Legendre functions $\mathbf{P}_{\ell}$. In terms of the tree description of $\mathbf{D H}$ or DH-Mid (cf. Figure 7) this simply truncates the tree immediately by computing only one level. In contrast, the other algorithms we have described compute reduced size input data vectors for the subproblems recursively from the coarser splitting at the the previous level of the tree.

This has the advantage of simplicity, but at the expense of increased complexity. Indeed, we shall see that

$$
\begin{equation*}
O\left(N^{\frac{3}{2}} \log ^{\frac{1}{2}} N\right) \tag{48}
\end{equation*}
$$

are required for the full transform.
From our description of the backward recurrence given above, we know that, after computing the cosine transforms of the vectors $\mathbf{s}^{L-1}$ and $\mathbf{s}^{L}$, each of the Legendre coefficients with degree in the interval $[L-r, L+r]$ can be computed in $O(r)$ further operations. This is done by forming the inner product against the cosine transform vector of the appropriate shifted or reverse shifted Legendre functions, each of which has degree less than or equal to $r$. Furthermore, only the cosine transform elements of $\mathbf{s}^{j-1}$ and $\mathbf{s}^{j}$ of degree $r$ or less are used to compute the Legendre coefficients with degree in the interval $[L-r, L+r]$. In this way, we can obtain a very simple algorithm by splitting the original problem into subproblems in which matrices derived from (reverse) shifted Legendre functions are applied to some fixed number $C$ of cosine transform vector pairs $\mathcal{C} s_{L-1}, \mathcal{C} s_{L}$, where $L$ is evenly spaced in the range $[0, N-1]$. Using analysis and experimentation, a value of $C$ can be chosen which minimizes the execution time.

Specifically, given a fixed value for $C$, we can evenly space the values of $L$ by letting

$$
\begin{equation*}
L=\frac{(2 j+1) N}{2 C}+1, \quad j=0, \ldots C-1 \tag{49}
\end{equation*}
$$

It is easy to see that, given $\mathcal{C} s_{L-1}, \mathcal{C} s_{L}$ as well as all of the required cosine transformed shifted Legendre vectors, $\mathcal{C} \mathbf{A}_{r}^{L}, \mathcal{C} \mathbf{N}_{r}^{L},\left(0<|r|<\frac{N}{2 C}\right)$ then

$$
\frac{N^{2}}{4 C^{2}}+\frac{N}{2 C}-2
$$

operations are needed to compute the Legendre coefficients $\left\langle f, P_{L+r}\right\rangle$ for $|r|<\frac{N}{2 C}$. With $C$ vector pairs $\left\{\mathcal{C s}_{L-1}, \mathcal{C} s_{L}\right\}$ the total cost of computing Legendre coefficients using shifted Legendre functions is

$$
\begin{equation*}
\frac{N^{2}}{4 C}+\frac{N}{2}-2 C \tag{50}
\end{equation*}
$$

Note that the assumption that the data structure of cosine transformed shifted Legendre vectors is available, i.e., precomputed and stored, does not change the order of complexity of the algorithm. All of these vectors can be computed in $O\left(\frac{N^{2}}{C}\right)$, and stored in a data structure of size $O\left(\frac{N^{2}}{C}\right)$.

It remains to determine the complexity of computing the $C$ vector pairs $\left\{\mathcal{C} \mathbf{s}_{L-1}, \mathcal{C} \mathbf{s}_{L}\right\}$. These vectors are obtained by the fast cosine transform of $f \mathbf{P}_{L}$, where $f$ and $P_{L}$ are sequences of length $2 N$. Also, only the first $\frac{N}{2 C}$ coefficients of the cosine series are needed. Using a cosine transform algorithm derived from [47], each vector pair requires

$$
2 N \log \frac{N}{2 C}
$$

multiplications, which gives a total cost for all vector pairs of

$$
\begin{equation*}
2 N C \log \frac{N}{2 C} \tag{51}
\end{equation*}
$$

Combining Eqs. (50) and (51) gives the total cost of computing the Legendre coefficients as

$$
\begin{equation*}
\frac{N^{2}}{4 C}+2 N C \log \frac{N}{2 C}+\frac{N}{2}-2 C \tag{52}
\end{equation*}
$$

in terms of the number of multiplications. Minimizing for $C$ in Eq. (52) gives

$$
C=O\left(\frac{N^{\frac{1}{2}}}{\log ^{\frac{1}{2}} N}\right)
$$

and substituting this value of $C$ into Eq. (52) gives

$$
\begin{equation*}
O\left(N^{\frac{3}{2}} \log ^{\frac{1}{2}} N\right) \tag{53}
\end{equation*}
$$

for the cost of the algorithm.
This algorithm may be an attractive candidate for parallelization as the computation of each vector pair, and the Legendre coefficients derived from them, are independent of other vector pairs, but the same sequence of arithmetic operations is used.

- Variation 4. Hybrid algorithms

In this approach, the semi-naive algorithm is used to compute Legendre transforms of degrees $m$ (where $m$ is the order of the transform) through $r$ for some fixed bound $r$ and then the simple-split algorithm is used to compute the remaining Legendre transforms of degrees $r+1$ through $N-1$.

While the order of complexity for this algorithm is greater than the $\mathbf{D H}$ variants, the constants and the overhead are small, and we will show in Section 6 that this algorithm performs quite well in practice.

## 6 Numerical Results

Experiments were performed on three platforms: a DEC Alpha 500/200, an HP Exemplar X-Class SPP2000/64, and a SGI Origin 2000. Even though the HP and SGI machines both have parallel architectures, our code never took advantage of this. All our tests were run on one processor. In order to (easily) maximize computational efficiency, we used a slightly modified version of the freely available software package FFTPACK [19]. We found the FFTs and DCTs provided in this library (written in Fortran) to be very efficient. Apart from FFTPACK, all other code was written in C, and all code was compiled using the native compiler with available optimizations.

The semi-naive algorithm was chosen as the standard against which we would measure the performance of the various DH-based algorithms. Previous work [15] has shown the semi-naive algorithm to be both stable and faster than the direct algorithm.

When conducting timing tests we would execute the algorithm 1000 times. This would average out timing variations due to multiprocessing and discretization. Unless stated otherwise, it is the CPU time (in seconds) of the 1000 iterations which report in the tables. We assume all precomputation and storage of the cosine transforms of the shifted Legendre polynomials. To measure the error of the various algorithms, we employ the following procedure:

1. Select a bandwidth $B$ and order $m$.
2. Generate a set of random Legendre coefficients $\hat{f}_{m}^{m}, \hat{f}_{m+1}^{m}, \ldots, \hat{f}_{B-1}^{m}$, normally distributed with mean 0 and standard deviation 1, using the Mathematica ${ }^{\text {© }}$ package NormalDistribution.m.
3. Synthesize the function

$$
f\left(\cos \theta_{k}\right)=\sum_{l=m}^{B-1} \hat{f}_{l}^{m} \tilde{P}_{l}^{m}\left(\cos \theta_{k}\right)
$$

where $k=0, \ldots, 2 B-1$ and $\theta_{k}=\frac{\pi(2 k+1)}{4 B}$.
4. Apply the algorithm to this synthesized function, generating a new set of Legendre coefficients $\hat{g}_{m}^{m}, \hat{g}_{m+1}^{m}, \ldots, \hat{g}_{B-1}^{m}$.
5. Compute the error as

$$
\max _{l}\left\|\hat{f}_{l}^{m}-\hat{g}_{l}^{m}\right\|
$$

6. Repeat steps 2-5 ten times.
7. Compute the average and relative error over the ten trials.

The synthesis step (Step 3) will necessarily introduce some error which we have empirically determined to be on the order of $10^{-9}$ in both the mean and standard deviation when bandwidth $B=1024$.

We focused our attention on the semi-naive, simple split and hybrid algorithms. Initial testing of the various, non-direct, algorithms revealed that for moderate problem sizes, regardless of the platform, the basic DH algorithm was the slowest. Though theoretically the basic algorithm is optimal, we found the cost of applying the smoothing operator to be expensive. This was one reason why variations of the basic algorithm were developed. Another reason, which will be discussed later, concerned stability.

In the matter of the simple-split algorithm, we would always choose the number of splits which minimized runtime and average error. This would be determined by trial and error. The number of splits used would not necessarily minimize the total number of operations (as discussed in Section 5). A general lesson we learned is that theoretical, optimal algorithms do not always yield the most cache-friendly algorithms. Inefficient cache-usage can render "fast" algorithms slow.

With regards to the hybrid algorithm the switch point of the algorithm, i.e. the degree sw of the coefficient at which we switched from the semi-naive to simple split algorithm was determined (after much experimentation) by the formula

$$
s w= \begin{cases}m+\frac{B-m}{2} & m<\frac{B}{16} \\ m+\frac{3(B-m)}{4} & \text { otherwise }\end{cases}
$$

where $m$ is the order of the transform and $B \leq 1024$ is the bandwidth of the problem. Originally, the above formula was used on all platforms and at all bandwidths $B$. However, we discovered that at problem size $B=1024$ we could increase the simple-split portion of the hybrid algorithm on the HP and SGI (but not the DEC) and get a bigger "win" over the pure semi-naive algorithm. Therefore, we used the following separate formula on the HP and SGI for bandwidth $B=1024$ :

$$
s w=m+\frac{1024-m}{2}
$$

Choosing this $s w$ seemed to minimize the runtime. Of course, your mileage may vary.
We first give timing results for order $m=0$ transforms. In Figure 8 we plot the ratios of the running times of the various algorithms versus the semi-naive algorithm on the three platforms. The dependence of the algorithms's performance on architecture is quite apparent. On the DEC, the simple-split and hybrid algorithms gain an advantage over the semi-naive algorithm at approximately bandwidth $B=256$. On the HP and SGI, the simple-split and hybrid algorithms are competitive with the semi-naive algorithm starting at approximately $B=512$. A transform of size $B=1024$ is necessary for the basic algorithm to be competitive with the semi-naive algorithm on all three platforms. Note also that at $B=1024$, the simple-split and hybrid algorithms run more than twice as fast as the semi-naive algorithm on the SGI and HP as they do on the DEC. We plot the runtimes of the algorithms themselves in Figure 9. Table 1 gives the average and relative errors of the algorithms on the DEC Alpha. These errors are representative of what was measured on all platforms.

Figure 8 is remarkable evidence of how computer architecture can effect performance. Keep in mind that identical code was compiled on all machines. The code was not written to take advantage of any particular architectural features. In fact, the code was developed on the DEC.


Figure 8: DLT runtime ratios vs Semi-naive, Order $m=0$


Figure 9: DLT runtimes (in CPU seconds) for 1000 iterations, Order $m=0$

| Bandwidth | Semi-naive | Basic | Simple split | Hybrid |
| :---: | :--- | :--- | :--- | :--- |
| 128 | $5.9138 \mathrm{e}-11$ | $5.9146 \mathrm{e}-11(1)$ | $5.9130 \mathrm{e}-11(1)$ | $5.9120 \mathrm{e}-11(1)$ |
|  | $4.4843 \mathrm{e}-09$ | $4.4860 \mathrm{e}-09$ | $4.4802 \mathrm{e}-09$ | $4.4794 \mathrm{e}-09$ |
| 256 | $2.0897 \mathrm{e}-10$ | $2.0884 \mathrm{e}-10(1)$ | $2.0887 \mathrm{e}-10(1)$ | $2.0885 \mathrm{e}-10(1)$ |
|  | $3.6380 \mathrm{e}-08$ | $3.6388 \mathrm{e}-08$ | $3.6230 \mathrm{e}-08$ | $3.6128 \mathrm{e}-08$ |
| 512 | $5.8493 \mathrm{e}-10$ | $5.8475 \mathrm{e}-10(1)$ | $5.8496 \mathrm{e}-10(2)$ | $5.8499 \mathrm{e}-10(2)$ |
|  | $3.6618 \mathrm{e}-07$ | $3.6625 \mathrm{e}-07$ | $3.6670 \mathrm{e}-07$ | $3.6665 \mathrm{e}-07$ |
| 1024 | $3.0777 \mathrm{e}-09$ | $3.0760 \mathrm{e}-09(2)$ | $3.0764 \mathrm{e}-09(3)$ | $3.0771 \mathrm{e}-09(3)$ |
|  | $6.4634 \mathrm{e}-06$ | $6.4683 \mathrm{e}-06$ | $6.4645 \mathrm{e}-06$ | $6.4639 \mathrm{e}-06$ |

Table 1: Order $m=0$ Legendre transform average (first row) and relative (second row) errors on a DEC Alpha workstation. For the basic algorithm, the number in parentheses refers to how many levels of recursive splitting have been done. For the the simple split and hybrid algorithms, the number in parentheses refers to how many splits were done.

| Bandwidth | Semi-naive | Simple split | Hybrid |
| :---: | :--- | :--- | :--- |
| 128 | $7.4787 \mathrm{e}-11$ | $4.5475 \mathrm{e}-09(3)$ | $7.4795 \mathrm{e}-11(1)$ |
|  | $3.9475 \mathrm{e}-09$ | $1.3638 \mathrm{e}-07$ | $3.9471 \mathrm{e}-09$ |
| 256 | $2.4867 \mathrm{e}-10$ | $8.2479 \mathrm{e}-10(8)$ | $2.4866 \mathrm{e}-10(1)$ |
|  | $2.3683 \mathrm{e}-08$ | $2.2254 \mathrm{e}-08$ | $2.3683 \mathrm{e}-08$ |
| 512 | $1.1740 \mathrm{e}-09$ | $5.5436 \mathrm{e}-09(16)$ | $1.1755 \mathrm{e}-09(2)$ |
|  | $8.9257 \mathrm{e}-08$ | $9.9039 \mathrm{e}-08$ | $8.9252 \mathrm{e}-08$ |
| 1024 | $4.6725 \mathrm{e}-09$ | $1.9154 \mathrm{e}-07(32)$ | $4.6725 \mathrm{e}-09(7)$ |
|  | $9.9888 \mathrm{e}-07$ | $1.4598 \mathrm{e}-06$ | $1.0024 \mathrm{e}-06$ |

Table 2: Order $m=B / 2$ Legendre transform average (first row) and relative (second row) errors on an SGI Origin workstation. For the the simple split and hybrid algorithms, the number in parentheses refers to how many splits were done.


Figure 10: DLT runtime ratios vs Semi-naive, Order $m=B / 2$


(c) HP Exemplar

Figure 11: DLT runtimes (in CPU seconds) for 1000 iterations, Order $m=B / 2$


Figure 12: $A_{256}^{512}$ : sampled 1024 times in the interval $[-1,1]$

Next, in Figures 10 and 11 we show the timing results for order $m=\frac{B}{2}$ transforms. Errors are given in Table 2.

When doing transforms of order $m=B / 2$, the basic algorithm was not tested because it was numerically unstable. The shifted Legendre polynomials proved to be the source of the instabilities. To illustrate, suppose the bandwidth is $B=1024$ and the order is $m=512$. Using the basic algorithm, the first recursive splitting involves multiplying the weighted signal by the shifted Legendre polynomial $A_{256}^{512}$. As Figure 12 makes clear, multiplying by $A_{256}^{512}$ would effectively zero out the sample values in the middle and stretch those near the end points to absurd values. For moderate and larger problem sizes, beginning at approximately order $m=4$, too great a shift will introduce instabilities.

This is why we do not test the basic and DH-Mid algorithms. Their "divide-and-conquer" structure is not only expensive computationally (as opposed to theoretically), but it also requires use of shifted Legendre polynomials. At present, the means to achieve a stable DH and DH-Mid algorithm still eludes us. The moral of this story is: to use shifted Legendre polynomials, take care not to shift too far with them.

The simple split algorithm can still be used with reasonable accuracy. Provided that enough initial splits are made, the numerically unsound shifted Legendre polynomials can be avoided. More initial splits implies shifts of smaller distance. But the penalty for making more splits is increased execution time. More splits require more discrete cosine transforms which are expensive. This can be seen in the Figure 10. At order $m=\frac{B}{2}$, the simple-split algorithm takes significantly longer to run than the semi-naive algorithm. This difference comes completely from the cost of doing cosine transforms. Profiling the code on the DEC at bandwidth $B=1024$ reveals that the semi-naive algorithm runs in roughly $45 \%$ of the time the simple-split algorithm spends simply doing cosine transforms. To achieve a numerically sound answer with the simple-split algorithm at $B=1024, m=512,32$ splits must be performed.

The hybrid algorithm, however, does remains competitive with semi-naive at order $m=\frac{B}{2}$, both in terms of runtime and accuracy, on the SGI and HP platforms at $B=1024$. Even though the hybrid algorithm uses shifted Legendre polynomials, it uses only relatively "nice" shifted Legendre polynomials. The rate at which $\max _{x \in[-1,1]}\left|A_{r}^{L}(x)\right| \rightarrow \infty$ as $r \rightarrow \infty$ is slower for large values of $L$ than for small. This also holds for the other shifted Legendre polynomials. For example, assuming bandwidth $B=1024$ and order $m=512$, we have the following maximum values:

$$
\max _{x \in[-1,1]}\left|\mathbf{A}_{10}^{768}(x)\right| \approx 3612
$$

and

$$
\max _{x \in[-1,1]}\left|\mathbf{A}_{10}^{512}(x)\right| \approx 6 \times 10^{11}
$$

This is the case because the initial conditions for the recurrence relations which generate the shifted Legendre polynomials grow smaller the further $L$ is from $m$. By applying the semi-naive technique to compute the lower degree coefficients, the hybrid algorithm avoids precisely those shifted Legendre polynomials which contribute most to error.

Since, on the HP and SGI platforms, the hybrid algorithm is both faster than the semi-naive algorithm and stable for orders $m=0$ through $m=512$ at $B=1024$, a significant savings can be achieved when performing a forward spherical transform. We implemented a hybrid spherical transform which works as follows. For orders $m=0$ through $m=B / 2$, the hybrid algorithm is used in performing the discrete Legendre transforms. For the remaining orders, the semi-naive algorithm is used.

In Table 3 we give timing results for the semi-naive and hybrid spherical transform algorithms, and also for the freely available software package SPHEREPACK [46]. We should note that, unlike all the algorithms discussed up to this point, SPHEREPACK assumes the function is sampled (in $\theta$ ) on the Gaussian points (i.e. Legendre points) and not the Chebyshev points (i.e. $\cos \left(\frac{\pi(k+1 / 2)}{2 B}\right)$ ). Hence, when doing a forward spherical transform of size $B$, SPHEREPACK expects the function to be sampled on a $2 B \times B$ grid, while the semi-naive and hybrid spherical transform algorithms require the sampling to be on a $2 B \times 2 B$ grid.

|  | DEC Alpha |  |  | SGI Origin |  | HP Exemplar |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| B | SPHEREPACK | Seminaive | Hybrid | Seminaive | Hybrid | Seminaive | Hybrid |
| 128 | $8.33 \mathrm{e}-01$ | $6.83 \mathrm{e}-01$ | $7.67 \mathrm{e}-01$ | $6.00 \mathrm{e}-01$ | $6.60 \mathrm{e}-01$ | $6.30 \mathrm{e}-01$ | $6.40 \mathrm{e}-01$ |
| 256 | $6.21 \mathrm{e}+00$ | $4.25 \mathrm{e}+00$ | $4.20 \mathrm{e}+00$ | $3.67 \mathrm{e}+00$ | $3.76 \mathrm{e}+00$ | $3.52 \mathrm{e}+00$ | $3.46 \mathrm{e}+00$ |
| 512 | $4.77 \mathrm{e}+01$ | $2.87 \mathrm{e}+01$ | $2.79 \mathrm{e}+01$ | $2.23 \mathrm{e}+01$ | $2.25 \mathrm{e}+01$ | $1.99 \mathrm{e}+01$ | $1.94 \mathrm{e}+01$ |
| 1024 | NA | NA | NA | $3.21 \mathrm{e}+02$ | $2.41 \mathrm{e}+02$ | $2.72 \mathrm{e}+02$ | $1.77 \mathrm{e}+02$ |
| 1024 | NA | NA | NA | $3.30 \mathrm{e}+02^{*}$ | $2.52 \mathrm{e}+02^{*}$ | $5.16 \mathrm{e}+02^{*}$ | $3.19 \mathrm{e}+02^{*}$ |

Table 3: Time (in CPU seconds) for 10 iterations of a Forward Spherical Transform on the DEC and HP; *: Walltime

As the bandwidth increases, the difference between the SPHEREPACK and other timings grows larger. At bandwidth $B=512$, the semi-naive and hybrid algorithms are nearly $40 \%$ faster than SPHEREPACK. Though we were not able to test SPHEREPACK on the HP and SGI, we are very confident that SPHEREPACK would be slower than the semi-naive and hybrid algorithms.

Due to local hardware limitations (the source of which will be apparent in the following sentence), we did not run the $B=1024$ forward spherical transform on the DEC. At $B=1024$, the seminaive spherical transform requires a precomputed data structure on the order of 1.3 gigabytes. The hybrid spherical transform requires under 0.9 gigabytes. Both these data structures are far too large to be entirely contained in memory. Therefore we had to read the precomputed data off disk, and this is why we also report the walltime. We believe that this is a legitimate quantity to mention because, in practice, we would expect not to precompute the data structures every time prior to performing a spherical transform. In either measure, CPU or walltime, the hybrid spherical transform is approximately $25 \%$ faster than semi-naive on the SGI. When run on the HP, the difference is 35 to $40 \%$. For the record, the HP walltime given in Table 3 was obtained when the precomputed data was saved on a disk specially configured to allow $30 \%$ faster disk i/o than on a normally configured
disk. Experiments confirmed this when the data was saved on such a "normal" disk. We note that the CPU time measured in both experiments was virtually identical, as to be expected. Errors are given in Table 4.

|  | SGI Origin |  | HP Exemplar |  |
| :---: | :---: | :---: | :---: | :---: |
|  | Seminaive | Hybrid | Seminaive | Hybrid |
| Average Error | $8.1483 \mathrm{e}-09$ | $7.9922 \mathrm{e}-09$ | $7.9740 \mathrm{e}-09$ | $7.5689 \mathrm{e}-09$ |
| Std Dev | $4.8625 \mathrm{e}-10$ | $7.1511 \mathrm{e}-10$ | $3.6623 \mathrm{e}-10$ | $1.0120 \mathrm{e}-09$ |
| Relative Error | $7.6050 \mathrm{e}-06$ | $8.7717 \mathrm{e}-06$ | $1.6206 \mathrm{e}-06$ | $4.0648 \mathrm{e}-06$ |
| Std Dev | $1.3939 \mathrm{e}-05$ | $1.2036 \mathrm{e}-05$ | $1.2383 \mathrm{e}-06$ | $3.9662 \mathrm{e}-06$ |

Table 4: Errors, over 10 iterations, on the HP for the Semi-naive and Hybrid Forward Spherical transforms, bandwidth $B=1024$.

We emphasize that although the asymptotic optimality of the basic algorithm is a theoretical result, it establishes the utility of shifted Legendre polynomials in developing fast algorithms. Our code is freely available at www.cs.dartmouth.edu/~geelong/sphere.

## 7 Two Applications

As stated in Section 1, a fast Fourier transform for $S^{2}$, as well as a fast convolution algorithm have many ready-made applications in applied science (see e.g. [39, 31, 44, 11, 37, 25, 24, 34].) In this section we examine in a little more detail two of these, one for the efficient computation of the bispectrum, potentially of use for image processing insensitive to rotations, and the other to matched filtering on $S^{2}$.

### 7.1 Computation of the bispectrum and triple correlation

The techniques of multiple correlations and higher order spectra have been developed for nonabelian Lie groups and their homogeneous spaces by R. Kakarala [30]. Of particular interest is the triple correlation and its associated Fourier transform, the bispectrum.

For functions on the line, the triple correlation is the integral of the product of the function with two independently shifted copies of itself. The resulting function on $\mathbf{R}^{2}$ determines the original function up to translation. The usefulness of computing the triple correlation derives from the fact that it is (1) insensitive to additive Gaussian noise; (2) retains most of the phase information of the underlying signal and (3) is invariant under translation of the underlying signal. This makes it useful in recovering a signal from multiple observations in situations in which the signal may be translating on a noisy background.

Kakarala has been able to generalize many of the results for functions on the line to arbitrary locally compact groups and their homogeneous spaces. For particular examples of interest such as the sphere, a suitably defined triple correlation of a band-limited function is again unique up to translation, (assuming that the Fourier coefficients are nonsingular) and insensitive to additive Gaussian noise. This suggests possible applications for global rotational motion compensation and Kakarala goes on to suggest possible applications to imaging the heart [11].

The techniques which we have developed for fast, reliable spherical convolution admit almost immediate application to fast, reliable computation of the triple correlation or bispectrum on the sphere. A detailed explanation of this is beyond the intended scope of this paper. However, the following abbreviated discussion should give some indication of our ideas.

To get to the bispectrum on the sphere we must go through the bispectrum for functions on its cover $S O(3)$. If $f \in L^{2}(S O(3))$, then the triple correlation of $f$ is the function on $S O(3) \times S O(3)$ given by

$$
a_{3, f}(s, t)=\int_{S O(3)} f(g s) f(g t) f(g) d g
$$

where $d g$ denotes Haar measure on $S O(3)$. Assuming $f$ is integrable on $S O(3)$, then $a_{3, f}$ is integrable on $S O(3) \times S O(3)$.

The irreducible representations of $S O(3)$ are naturally indexed by nonnegative integers, one irreducible of dimension $2 l+1$ for each $l \geq 0$, which we denote as $\rho_{l}$. Consequently, the irreducible representations of $S O(3) \times S O(3)$ are given by all possible tensor products $\rho_{l} \otimes \rho_{l^{\prime}}$, so are indexed by all pairs $\left\{l, l^{\prime}\right\}$ with $l \geq l^{\prime} \geq 0$.

The Fourier transform of $f$ at $\rho_{l}$, denoted as $\hat{f}(l)$ is the integral

$$
\hat{f}(l)=\int_{S O(3)} f(g) \rho_{l}(g)^{\dagger} d g
$$

where $\dagger$ indicates conjugate transpose. The Fourier transform of $f$ is the collection $\{\hat{f}(l)\}_{l \geq 0}$.
Similarly, the Fourier transform of a function on $S O(3) \times S O(3)$ will be the analogously defined collection $\left\{\hat{f}\left(l, l^{\prime}\right)\right\}_{l^{\prime} \geq l \geq 0}$. The bispectrum of $f$ is Fourier transform of $a_{3, f}$. In [30] Kakarala shows how the bispectrum may be computed from the Fourier transform of $f \in L^{2}(S O(3))$. For this we need to introduce one more piece of notation. Notice that $S O(3)$ has a natural embedding in $S O(3) \times S O(3)$ as the diagonal subgroup. Considered as such, each representation $\rho_{l^{\prime}, l}$ when restricted to the diagonal will be equivalent to a direct sum of appropriate $\rho_{j}$. Thus, there exists an invertible matrix $C_{l, l^{\prime}}$ such that

$$
\rho_{l, l^{\prime}}(s, s)=C_{l, l^{\prime}}\left[\rho_{j_{1}\left(l, l^{\prime}\right)}(s) \oplus \rho_{j_{2}\left(l, l^{\prime}\right)}(s) \oplus \cdots \oplus \rho_{j_{m}\left(l, l^{\prime}\right)}(s)\right] C_{l, l^{\prime}}^{\dagger}
$$

for suitable indices $j_{i}\left(l, l^{\prime}\right)$.
Theorem ([30], Lemma 3.2.3) With the notation as above,

$$
\begin{equation*}
\widehat{a}_{3, f}\left(l, l^{\prime}\right)=\hat{f}(l) \otimes \hat{f}\left(l^{\prime}\right) C_{l, l^{\prime}}\left[\hat{f}\left(j_{1}\left(l, l^{\prime}\right)\right)^{\dagger} \oplus \hat{f}\left(j_{2}\left(l, l^{\prime}\right)\right)^{\dagger} \oplus \cdots \oplus \hat{f}\left(j_{m}\left(l, l^{\prime}\right)\right)^{\dagger}\right] C_{l, l^{\prime}}^{\dagger} \tag{54}
\end{equation*}
$$

When $f \in L^{2}(S O(3))$ comes from a function on the sphere, (i.e. $f \in L^{2}(S O(3))$ is right $S O(2)$ invariant) then the matrix $\hat{f}(l)$ will have entries all 0 except, possibly, for a single column which (up to a normalization constant) will contain the associated Legendre transforms

$$
\{\hat{f}(l,-l), \ldots, \hat{f}(l, 0), \ldots, \hat{f}(l, l)\}
$$

Thus, if $f$ is band-limited, then the bispectrum will only involve a finite number of Fourier transforms. For each $l, l^{\prime}, \widehat{a}_{3, f}\left(l, l^{\prime}\right)$ can then be computed directly as follows. Compute first the spherical harmonic expansion as described in Section 3. This precomputes all possible Fourier transforms $\hat{f}(l)$ for any $f \in L^{2}\left(S^{2}\right)$. The inner direct sum of the matrices

$$
\left[\hat{f}\left(j_{1}\left(l, l^{\prime}\right)\right)^{\dagger} \oplus \hat{f}\left(j_{2}\left(l, l^{\prime}\right)\right)^{\dagger} \oplus \cdots \oplus \hat{f}\left(j_{m}\left(l, l^{\prime}\right)\right)^{\dagger}\right]
$$

is then constructed by retrieving the appropriate associated Legendre transforms and organizing them together into a single sparse block diagonal matrix. This is then conjugated by the precomputed change of basis matrices $C_{l, l^{\prime}}$. Finally the lefthand factor

$$
\hat{f}(l) \otimes \hat{f}\left(l^{\prime}\right)
$$

simply requires the computation of all possible pointwise products $\hat{f}(l, m) \hat{f}\left(l^{\prime}, m^{\prime}\right)$ organized as the appropriate single nonzero column in some suitably defined matrix. These matrices are then all multiplied together, giving the relevant component of the bispectrum.

### 7.2 Matched filters

One simple application of the techniques of this paper may be found in certain problems of detection, estimation, and pattern matching for data defined on the sphere. This sort of data arises in geophysics, computer vision, or quality assurance for computer designed and manufactured parts.

A simple problem arising in this area may be stated as follows: Suppose we are considering a known signal or pattern in the directional data setting, described by a function, $f(\omega)$ on the sphere. In many situations, we are interested in determining the presence or absence of this signal in data coming from measurements of some real world phenomenon. This is often made more difficult by the presence of some random interference, or noise, in the measurements. In the simplest cases, we assume that one of two hypotheses obtains for the measured data, $\mathbf{y}(\omega)$ :

- $H_{0}: \mathbf{y}(\omega)=\mathbf{n}(\omega)$
- $H_{1}: \mathbf{y}(\omega)=f(\omega)+\mathbf{n}(\omega)$,
where $\mathbf{n}(\omega)$ is a random process on the sphere representing the noise. Our task is then to devise an algorithm which takes a particular instance of the measured data and returns an assessment of whether or not the signal is present in the data.

A more interesting version of this problem occurs when, in addition to the additive noise, the pattern signal $f$ may have undergone a rotation which is unknown to us. That is, when $f$ is present, the measured data has the form

$$
\mathbf{y}(\omega)=\Lambda(g) f(\omega)+\mathbf{n}(\omega)
$$

where $g$ is an unknown element of $S O(3)$, and $\Lambda(g)$ is the associated operator, $\Lambda(g) f(\omega)=f\left(g^{-1} \omega\right)$. In this case we have the more involved detection and estimation problem; determine if a rotated version of the pattern is present, and if so, estimate the value of the rotation parameter $g$. We will concern ourselves with this question.

The intuitive approach to this involves a template matching operation. That is, one correlates the data with the pattern one is looking for, which amounts to forming the inner product of the data with a large number of shifted versions of the pattern. Those shifts which produce a large inner product, or correlation, between the pattern and the data are regarded as good indicators that there really is a copy of the pattern shifted to the corresponding location and buried in the noise. This correlation process is known as matched filtering; it amounts to computing the function

$$
\chi(g)=\int_{S^{2}} y(\omega) \Lambda(g) f(\omega) d \omega .
$$

This matched filter is also indicated by a standard statistical analysis for these sorts of problems. This analysis yields optimal detection and estimation solutions involving a computation of the appropriate likelihood function, $[32,53,51]$. Basically, for a given measurement of data, the value of this function gives the likelihood of having made that particular observation given a certain hypothesis (signal present or signal absent) or a given value of the unknown parameter.

For example, suppose we know that a rotated version of the signal is present, in the data process $\mathbf{y}(\omega)$, and we wish to know where it is. This is the same as estimating the rotation parameter $g$. We assume that the additive noise $\mathbf{n}(\omega)$ is Gaussian and white. The latter term refers to the covariance structure of the noise, implying first that the covariance $R\left(\omega_{1}, \omega_{2}\right)=\mathbf{E}\left[\mathbf{n}\left(\omega_{1}\right) \mathbf{n}\left(\omega_{2}\right)\right]$ is actually rotation
independent, so that $R\left(\omega_{1}, \omega_{2}\right)=R\left(g \omega_{1}, g \omega 2\right)$ for any rotation $g \in S O(3)$. This property is sometimes referred to as "stationarity." A consequence of stationarity is that $R$ is determined by the values $R(\omega)$ $=R(\omega, \eta)$, for $\eta$ the north pole of the sphere. White noise is a particular stationary noise with point mass covariance; $R(\omega)=\sigma^{2} \delta_{\eta}(\omega)$. Strictly speaking, this requires the usual sorts of mathematical temporizing required when dealing with distributions; we'll assume that is familiar.

The likelihood $L(g)$ of a particular value of the parameter $g$ given the data $y$ is the probability density for the random variable $\mathbf{y}=\Lambda(g) f+\mathbf{n}$ evaluated at the particular observed measurement values $y_{o}$; this is the same as $p_{\mathbf{y}}\left(y_{o} ; g\right)=p_{\mathbf{n}}\left(y_{o}-\Lambda(g) f\right)$. Using our assumptions on $\mathbf{n}$ and some limiting arguments, we obtain for our likelihood:

$$
\begin{aligned}
L(g) & \propto e^{-\frac{\left\|y_{o}-\Lambda(g) f\right\|_{2}^{2}}{2 \sigma^{2}}} \\
& \propto e^{\frac{\left\langle y_{o}, \Lambda(g) f\right\rangle}{\sigma^{2}}},
\end{aligned}
$$

as the other terms which come from expanding the norm, $e^{-\frac{\left\|y_{o}\right\|_{2}^{2}}{2 \sigma^{2}}}$ and $e^{-\frac{\|\Lambda(g) f\|_{2}^{2}}{2 \sigma^{2}}}=e^{-\frac{\|f\|_{2}^{2}}{2 \sigma^{2}}}$ are constant, independent of $g$. Thus the maximum likelihood estimate of $g$ is

$$
\operatorname{Arg} \operatorname{Max}_{\mathrm{g} \in \mathrm{SO}(3)} \int_{\mathrm{S}^{2}} \mathrm{y}(\omega) \Lambda(\mathrm{g}) \mathrm{f} \omega \mathrm{~d} \omega
$$

Let us consider now a simple case in which the pattern signal is rotationally symmetric. In fact, we take $f$ to be the analog of the normal density on the sphere, the Fisher von-Mises density $C_{\kappa} \exp (\kappa \cos \theta)$. Here, $\kappa$ is a concentration parameter, $C_{\kappa}$ a normalizing factor. In this case, the matched filter expression actually reduces to a function defined on the sphere, rather than the entire group, due to rotation invariance. Below we show the results of some experiments in which $f$ is rotated and buried in white noise, and then passed through a matched filter. The results are shown in Figure 13.

## 8 Summary and Future Directions

We have presented a divide-and-conquer algorithm for the efficient and exact computation of the forward and inverse Fourier transform of a band-limited function on the 2 -sphere, which in addition provides a fast algorithm to compute the exact convolution of two such functions. We give evidence that by fine-tuning different variations of the basic algorithm, highly efficient and numerically reliable implementations can be obtained. These algorithms have a wide range of applicability in a great range of scientific disciplines.

We view this work as another step in the still nascent development of algorithms and applications for efficient nonabelian harmonic analysis. Given the small ratio of nonabelian to abelian papers, it seems that there are still many potentially fruitful directions to pursue (see also [36]).

1. Fine-tuning. Different computational environments will require different adaptations of the basic algorithm for maximum performance. This will undoubtedly require much more experimentation.
2. Vector and tensor harmonic expansions. For various applications in meteorolgy, it is also important to compute the expansions of vector and tensor fields in terms of vector and tensor harmonics. With an appropriate definition, this may be reduced to the computation of several individual Fourier transforms on the 2-sphere, so that our algorithms may be applied.
3. Paralellizabilty. The divide-and-conquer nature of the basic algorithm indicates that efficient parallel implementations may be possible. The technical report [26] is a first step in this direction.


Figure 13: (a) The pattern signal $f(\omega)$ is the Fisher-von Mises density $C_{\kappa} \exp (\kappa \cos \theta)$, with concentration parameter $\kappa=64$, and rotated by an arbitrary rotation $g$ on the sphere; (b) $\Lambda(g) f(\omega)$ has been buried in additive white noise $n(\omega)$ to simulate noisy measured data $y(\omega)$; (c) The likelihood function $L(g)$ computed by using the fast convolution algorithm to convolve $y(\omega)$ with a matched filter. The position $g_{\max }$ of the maximum value of $L(g)$ indicates the maximum likelihood estimate of the position of the pattern signal $f(\omega)$.
4. Improved complexity. Our algorithms only use one of the recurrences satisfied by the $\mathbf{P}_{\ell}^{m}$. Perhaps through the use of other recurrences the overall complexity can be reduced to $O(n \log n)$.
5. Other compact groups and their quotients. Due to its applicability, we have concentrated our efforts on developing algorithms for the 2-sphere. The basic ideas shown here work (in theory) for any compact group and its quotients (cf. [35, 36]). Identification of new applications in this setting would probably dictate the priorities of related software development. The papers [28, 27, 43] give some indication of the wide variety of applications being found for these generalized FFTs.

## References

[1] N. Ahmed, T. Natarajan, and K. R. Rao, Discrete cosine transforms, IEEE Transactions on Computers, 23 (1974), 90-93.
[2] A. Aho, J. Hopcroft, and J. Ullman, The Design and Analysis of Computer Algorithms, Addison Wesley, Reading, 1976.
[3] B. Alpert and V. Rokhlin, A fast algorithm for the evaluation of Legendre transforms, SIAM J. Sci. Statist. Comput. 12 (1991), 158-179.
[4] L. Auslander and R. Tolmieri, Is computing with the fast Fourier transform pure or applied mathematics?, Bulletin of the American Mathematical Society, (N. S.) 1 (1979), 847-897.
[5] P. Barrucand and D. Dickinson, On the associated Legendre polynomials, in Orthogonal Expansions and their Continous Analogues, Southern Illinois University Press, Carbondale, 1968.
[6] L. C. Biedenharn and J. D. Louck, Angular Momentum in Quantum Mechanics, Addison Wesley, Reading, 1981.
[7] A. Borodin and I. Munro, The Computational Complexity of Algebraic and Numeric Problems Elsevier, New York, 1975.
[8] J. Boyd, Chebyshev and Fourier Spectral Methods, Lecture Notes in Engineering, Vol. 49, SpringerVerlag, NY, 1989.
[9] N. Bshouty, M. Kaminski, and D. Kirkpatrick, Addition requirements for matrix and transposed matrix products. J. of Algorithms, 9 (1988), 354-364.
[10] D. Calvetti, A stochastic roundoff error analysis for the fast Fourier transform, Math. Comput., 56(194) (1991), 755-774.
[11] C. W. Chen and T. S. Huang, Epicardial motion and deformation estimation from coronary artery bifurcation points, in Proc. of Third. Int. Conf. on Comp. Vision, Dec. 4-7, 1990, IEEE Press (1990), 456-460.
[12] J. W. Cooley and J. W. Tukey, An algorithm for machine calculation of complex Fourier series, Math. Comput., 19 (1965), 297-301.
[13] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, Introduction to Algorithms, MIT Press, 1990.
[14] P. Diaconis, Average running time of the fast Fourier transform, J. Algorithms, 1 (1980), 187-208.
[15] G. A. Dilts, Computation of spherical harmonic expansion coefficients via FFTs, Journal of Computational Physics, 57(3) (1985), 439-453.
[16] J. R. Driscoll and D. Healy, Computing Fourier transforms and convolutions on the 2-sphere. (extended abstract) in Proc. $34^{\text {th }}$ IEEE FOCS, (1989) 344-349; Adv. in Appl. Math., 15 (1994), 202-250.
[17] J. R. Driscoll, D. Healy and D. Rockmore, Fast discrete polynomial transforms with applications to data analysis for distance transitive graphs, SIAM J. Comput., 26(4) (1997), pp. 1066-1099.
[18] D. F. Elliott and K. R. Rao, Fast Transforms: Algorithms, Analyses, and Applications. Academic Press, New York, 1982.
[19] FFTPACK is a freely available collection of FORTRAN programs for computing one-dimensional fast Fourier transforms developed at NCAR by Paul Swarztrauber.
[20] W. Freeden, On integral formulas of the (unit) sphere and their application to numerical computation of integrals, Computing, 25 (1980), 131-146.
[21] N. C. Gallagher, G. L. Wise, and J. W. Allen, A novel approach for the computation of Legendre polynomial expansions, IEEE Transactions on Acoustics, Speech, and Signal Processing, vol. ASSP-26(1) (1978), 105-106.
[22] F. A. Gilbert, Inverse problems for the earth's normal modes, in Mathematical Problems in the Geophysical Sciences, Vol I, American Mathematical Society, Providence, 1971.
[23] L. Greengard, Fast algorithms for classical physics, Science, 265 (1994), 909-914.
[24] D. Healy and P. Kim, An empirical Bayes approach to directional data and efficient computation on the sphere, Ann. Stat., 24(1) (1996), 232-254.
[25] D. Healy, H. Hendriks and P. Kim, Spherical deconvolution with application to geometric quality assurance, Technical Report, Department of Mathematics and Computer Science, Dartmouth College, (1993).
[26] D. Healy, S. Moore, and D. Rockmore, Efficiency and stability issues in the numerical convolution of Fourier transforms and convolutions on the 2-sphere, Technical Report PCS-TR94-222, Department of Computer Science, Dartmouth College, 1994.
[27] D. Healy, T. Olson, D. Rockmore, and G. Mirchandani, Wreath products for image processing, in Proceedings of 1996 ICASSP, Vol. 6, pp. 3582-3585.
[28] D. Healy, D. Rockmore, and S. Moore, An FFT for the 2-sphere and applications, in Proceedings of 1996 ICASSP, Vol. 3, pp. 1323-1326.
[29] M. T. Heideman, D. H. Johnson, and C. S. Burrus, Gauss and the history of the fast Fourier transform, Arch. for History of Exact Sciences 34(3) (1985), 265-277.
[30] R. Kakarala, Triple Correlation on Groups, Ph.D. Thesis, Dept. of Math. University of California, Irivine, 1992.
[31] K. Kanatani, Group-Theoretical Methods in Image Understanding, Springer-Verlag, NY, 1990.
[32] S. Kay, Fundamentals of Statistical Signal Processing, Prentice-Hall, Englewood Cliffs, NJ, 1993.
[33] K. Kobayashi, Solution of multi-dimensional neutron transport equation of the spherical harmonics method using the finite Fourier transformation and quadrature formula, Transport Theory and Stat. Phys., 14 (1985), 63-85.
[34] G. J. Martyna and B. J. Berne, Structure and energies of $\mathrm{Xe}_{n}^{-}$: Many body polarization effects, J. Chem. Phys. 90(7) (1989), 3744-3755.
[35] D. Maslen, Fast Transforms and Sampling for Compact Groups, Ph.D. Thesis, Department of Mathematics, Harvard University, 1993.
[36] D. Maslen and D. Rockmore, Generalized FFTs, in Groups and Computation II, DIMACS Series in Discrete Math and Computer Science, Vol. 28, L. Finkelstein and W. Kantor (eds.), (1997), pp. 183-237.
[37] M. Miller, S. Joshi, D. Maffitt, J. McNally, and U. Grenander, Membranes, mitochondria and amoebae: shape models, in Advances in Applied Statistics (1994), 137-159.
[38] S. Moore, D. Healy, and D. Rockmore, Symmetry stabilization for polynomial evaluation and interpolation. Lin. Alg. Appl., 192 (1993), 249-299.
[39] Siang Peng Oh, David N. Spergel and Gary Hinshaw, An Efficient Technique to Determine the Power Spectrum from Cosmic Microwave Background Maps. preprint, 1998.
[40] A. Oppenheim and R. Schafer, Discrete-Time Signal Processing. Prentice Hall, NJ, 1989.
[41] S. A. Orzag, Fast eigenfunction transforms, in Science and Computers, Academic Press, Orlando, 1986.
[42] G. U. Ramos, Roundoff error analysis of the fast Fourier transform, Math. Comp., 25 (1971), 757-768.
[43] D. Rockmore, Some applications of generalized FFTs (Appendix with D. Healy), in Groups and Computation II, DIMACS Series in Discrete Math and Computer Science, Vol. 28, L. Finkelstein and W. Kantor (eds.), (1997), pp. 329-369.
[44] J. T. Schwartz, Mathematics Addresses Problems in Computer Vision for Advanced Robotics, SIAM News, 18(3) (1985).
[45] P. Schwartztrauber, The vector harmonic transform method for solving partial differential equations in spherical geometry, Monthly Weather Review 121(12) (1993), 3415-3437.
[46] SPHEREPACK is a freely available collection of FORTRAN programs that facilitates computer modeling of geophysical processes developed at NCAR by John Adams and Paul Swarztrauber.
[47] G. Steidl and M. Tasche, A polynomial approach to fast algorithms for discrete Fourier-cosine and Fourier-sine transforms, Math. Comp., 56 (1991), 281-296.
[48] C. Temperton, On scalar and vector transform methods for global spectral models, Mon. Wea. Rev. 119 (1991), 1303-1307.
[49] P. P. Vaidyanathan, Multirate Systems and Filter Banks, Prentice Hall, Englewood Cliffs, NJ, 1993.
[50] C. Van Loan, Computational Framework for the Fast Fourier Transform, SIAM, Philadelphia, 1992.
[51] H. Van Trees, Detection, Estimation and Modulation Theory, Vol I, Wiley, New York, 1968.
[52] N. J. Vilenkin, Special Functions and the Theory of Group Representations, American Mathematical Society, Providence, 1968.
[53] P. M. Woodward, Probability and Information Theory, with Applications to Radar, Artech House, Dedham MA, 1980.


[^0]:    *Preliminary versions of some of these results have appeared in the Dartmouth College Department of Computer Science Technical Report PCS-TR94-222 and "An FFT for the 2-sphere and Applications", Proc. of ICASSP-96, Volume 3, pp. 1323-1326.
    ${ }^{\dagger}$ This work supported in part by ARPA as administered by the AFOSR under contract DOD-F4960-93-0567. D. R. also supported in part by NSF DMS Awards 9404275 and a Presidential Faculty Fellowship under which P. K. is supported

[^1]:    ${ }^{1}$ Here we assume the standard arithmetic complexity model which defines a single operation as a complex multiplication followed by a complex addition.

