

ALGORITHM FOR THE EVALUATION OF REDUCED WIGNER MATRICES

G. PRÉZEAU^{1,2} AND M. REINECKE³

¹ Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

² California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125, USA

³ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

Received 2010 January 14; accepted 2010 August 4; published 2010 September 16

ABSTRACT

Algorithms for the fast and exact computation of Wigner matrices are described and their application to a fast and massively parallel 4π convolution code between a beam and a sky is also presented.

Key words: cosmic background radiation – cosmology: observations – instrumentation: adaptive optics – methods: data analysis – techniques: image processing

Online-only material: color figures

1. INTRODUCTION

Wigner matrices have a ubiquitous presence in science; from the computation of molecular quantum states, through the description of solitons in particle physics and the convolution of beam and sky algorithms in astronomy, they are needed to sometimes very high quantum numbers. Algorithms that calculate Wigner matrices quickly and accurately are therefore far-reaching in their applicability. Other methods have been developed that calculate these matrices exactly but with suboptimal performance to very high angular momenta (Risbo 1996), or approximately but very efficiently (Rowe et al. 2001), but none of the methods calculate them exactly and quickly to almost arbitrarily high angular momentum. Two such methods are presented in this paper and applied to a convolution algorithm between beam and sky. The following section gives some basic properties of Wigner matrices, and this is followed by a section describing the algorithm. The fourth section describes its application to convolution and a summary is presented at the end.

2. WIGNER MATRICES

Wigner matrix elements,⁴ typically denoted by $D_{mm'}^l(\alpha, \beta, \gamma)$, are the eigenfunctions of the Schrödinger equation for a symmetric top and form an irreducible basis of the Lie group SU(2) and the rotation group SO(3); the angles α , β , and γ are the Euler angles that define the orientation of the top. As basis functions of SU(2), the $D_{mm'}^l(\alpha, \beta, \gamma)$ satisfy the standard angular momentum relations

$$\hat{J}^2 D_{mm'}^l(\alpha, \beta, \gamma) = l(l+1)D_{mm'}^l(\alpha, \beta, \gamma), \quad (1)$$

$$\hat{J}_z D_{mm'}^l(\alpha, \beta, \gamma) = m D_{mm'}^l(\alpha, \beta, \gamma), \quad (2)$$

$$\hat{J}_{z'} D_{mm'}^l(\alpha, \beta, \gamma) = m' D_{mm'}^l(\alpha, \beta, \gamma), \quad (3)$$

where l labels the irreducible representation of SU(2) and also corresponds to the quantum number representing the total angular momentum of the eigenfunction; $-l \leq m, m' \leq l$ are the quantum numbers representing the projections of the total angular momentum on two z -axes rotated with respect to each other as described below.

The Euler angles are defined as three rotations: a rotation γ about the z -axis that rotates the x - and y -axes $\rightarrow x'$ and y' ; this first rotation is followed by a rotation β about the new y' -axis rotating the x' - and z -axes $\rightarrow x''$ and z' ; the final rotation α is about z' . In the basis we are using as defined by Equations (2) and (3), the operators \hat{J}_z and $\hat{J}_{z'}$ are diagonal and $D_{mm'}^l(\alpha, \beta, \gamma)$ has the form

$$D_{mm'}^l(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mm'}^l(\beta) e^{-im'\gamma}, \quad (4)$$

where

$$d_{mm'}^l(\beta) = \langle lm | \exp\left[-i\frac{\beta}{\hbar}\hat{J}_y\right] | lm' \rangle. \quad (5)$$

Here, $d_{mm'}^l(\beta)$ is called the reduced Wigner matrix element and consists of the overlap of a spherical harmonic with another spherical harmonic that has been rotated by an angle β about the y -axis. The differential equation satisfied by $d_{mm'}^l(\beta)$ is

$$\begin{aligned} \frac{d^2 d_{mm'}^l(\beta)}{d\beta^2} + \cot\beta \frac{d d_{mm'}^l(\beta)}{d\beta} \\ + \left(\frac{2mm' \cos\beta - m^2 - m'^2}{\sin^2\beta} + l(l+1) \right) d_{mm'}^l(\beta) = 0. \end{aligned} \quad (6)$$

From the Schrödinger equation in Equation (6), it is possible to extract three-term recursion relations that relate reduced Wigner matrix elements that differ in their quantum numbers. In principle, it is possible to use such relations to calculate the $d_{mm'}^l(\beta)$. Three-term recursion relations can be unstable, which limits their usefulness unless the potential pitfalls are identified and avoided. Two examples of these relations are

$$\begin{aligned} \frac{-m + m' \cos\beta}{\sin\beta} d_{mm'}^l(\beta) = \frac{1}{2} \sqrt{(l+m')(l-m'+1)} d_{mm'-1}^l(\beta) \\ + \frac{1}{2} \sqrt{(l-m')(l+m'+1)} d_{mm'+1}^l(\beta) \end{aligned} \quad (7)$$

and

$$\begin{aligned} \left[\cos\beta - \frac{mm'}{l(l+1)} \right] d_{mm'}^l(\beta) = \frac{\sqrt{(l^2-m^2)(l^2-m'^2)}}{l(2l+1)} d_{mm'}^{l-1}(\beta) \\ + \frac{\sqrt{[(l+1)^2-m^2][(l+1)^2-m'^2]}}{(l+1)(2l+1)} d_{mm'}^{l+1}(\beta). \end{aligned} \quad (8)$$

⁴ For a nice review of Wigner matrices, see Varshalovich et al. (1988).

Generally, three-term recursion relations will have two linearly independent solutions, f_n and g_n (Press et al. 1988); these solutions can be oscillatory or exponentially decreasing or increasing. In the non-oscillatory case, f_n is the *minimal* solution if

$$\frac{f_n}{g_n} \rightarrow 0 \text{ as } n \rightarrow \infty, \quad (9)$$

while g_n is the *dominant* solution. For solutions to the Schrödinger equation, exponentially increasing/decreasing solutions appear only in the region where a particle cannot classically exist because of energy conservation, but where a wave function can be non-zero in quantum mechanics. In the case of a rigid rotor (Edmonds 1957), the kinetic energy of a spherically symmetric rotor is

$$2IT = p_\beta^2 + \frac{1}{\sin^2 \beta} (p_\gamma^2 + p_\alpha^2 - 2p_\alpha p_\gamma \cos \beta). \quad (10)$$

In classical mechanics, $p_\beta^2 > 0$. In quantum mechanics, the quantization of Equation (10) means substituting $p_\alpha \rightarrow -i\partial/\partial\alpha$, $p_\beta \rightarrow -i\partial/\partial\beta$, and $p_\gamma \rightarrow -i\partial/\partial\gamma$. These substitutions combined with an eigenfunction of the form (4) and the additional substitution $2IT \rightarrow l(l+1)D_{mm'}^l(\alpha, \beta, \gamma)$ inferred from Equation (1) yields Equation (6). Since p_β^2 corresponds to the first two terms of Equation (6), one concludes that classically we would have

$$l(l+1) + \frac{2mm' \cos \beta - m^2 - m'^2}{\sin^2 \beta} \geq 0. \quad (11)$$

Wherever Equation (11) is not satisfied, the solutions will be exponentially suppressed or divergent. When solving the Schrödinger equation for the physical solutions, the divergent solutions are simply put to zero. When using the three-term recursion relations, the divergent solution can be “sniffed” out because of round-off errors and the recursions quickly fail. One special case where this cannot happen is when $m, m' = 0$ where Equation (11) is always satisfied since $l \geq 0$. In that case, Equation (8) is stable and can be used to calculate $d_{00}^l(\beta)$ to very high l extremely accurately.

For the cases where $m, m' \neq 0$, we can still use three-term recursion relations provided we do so in the right direction in the quantum number being varied. For example, looking at Equation (7), one can either calculate each $d_{mm'}^l(\beta)$ for increasing m' or decreasing m' . In one direction, the divergent solution will be growing while in the other it will be shrinking. To determine the direction in which Equation (7) is stable, one only needs to consider Equation (11). Assume you are interested in evaluating all the reduced Wigner matrix elements $d_{0m'}^l(\beta)$ for $0 \leq m' \leq l$ using Equation (7); you can choose to begin your recurrence with $d_{00}^l(\beta)$ and increasing m' or begin from $d_{0l}^l(\beta)$ and decreasing m' . To use Equation (7) in a stable manner, you need to start from $d_{0l}^l(\beta)$ and decrease m' . Putting $m = 0$ in Equation (11) yields the new condition

$$l(l+1) - \frac{m'^2}{\sin^2 \beta} \geq 0, \quad (12)$$

where it is seen that as m' increases from 0 (taking, for example, $\beta = \pi/4$), we approach the non-classical region and violate Equation (12) at $m' \geq \sin \beta \sqrt{l(l+1)}$; increasing m' further means sampling the divergent dominant solution of the

Schrödinger equation from which Equation (7) is derived. It is then clear that the stable direction to use Equation (7) is for decreasing $|m'|$. From Equation (11), it is seen quite generally that the recursion relations (7) and (8) will be stable provided they are used in the direction of decreasing $|m'|$ and increasing l , respectively.

In addition to Equations (7) and (8), a third recursion relation in β can be derived by discretizing the derivatives in Equation (6) with the relations

$$f'(x) \cong \frac{f(x+\epsilon) - f(x-\epsilon)}{2\epsilon} + O(\epsilon^2 f'''), \quad (13)$$

$$f''(x) \cong \frac{f(x+\epsilon) + f(x-\epsilon) - 2f(x)}{\epsilon^2} + O(\epsilon f'''). \quad (14)$$

Substituting into Equation (6) yields

$$\begin{aligned} & \left[\epsilon^2 \left(\frac{2mm' \cos \beta - m^2 - m'^2}{\sin^2 \beta} + l(l+1) \right) - 2 \right] d_{mm'}^l(\beta) \\ & \cong \left(\frac{\epsilon \cot \beta}{2} - 1 \right) d_{mm'}^l(\beta - \epsilon) - \left(\frac{\epsilon \cot \beta}{2} + 1 \right) d_{mm'}^l(\beta + \epsilon) \\ & + O(\epsilon^3 d_{mm'}^{l'''}(\beta)). \end{aligned} \quad (15)$$

From Equation (11), it is seen that this recursion relation should be used for increasing β if $0 < \beta < \pi/2$ and decreasing β if $\pi/2 < \beta < \pi$. Examples of these conclusions are given in Figure 1. The top plot shows the change in the behavior of $d_{mm'}^l(\beta)$ with increasing l as one moves from the non-classical to classical regions; in that case, the angle β was chosen so that Equation (11) is satisfied only when $l \geq 100$. The middle plot shows the variation of $d_{mm'}^l(\beta)$ with m' . With $\beta = 0.52302$, $l = 500$, and $m = 0$, the transition from non-classical to classical regimes occurs at $m' = 250$. The last plot shows the variation of $d_{mm'}^l(\beta)$ with β ; the value of $m' = 71$ was chosen so that the transition from non-classical to classical occurs at $\beta = \pi/4$. A noticeable feature of all three plots is that the tallest peak is always the first peak after the transition to the classical region. This is qualitatively understandable from Equation (5) where the reduced Wigner matrix is seen to characterize the overlap between two spin states after a rotation. In the classical limit of large l , the angle ω of the spin direction of a quantum object with the z -axis is given by

$$|lm\rangle: \cos \omega \approx \frac{m}{\sqrt{l(l+1)}}. \quad (16)$$

One might expect that the overlap would be greatest when the “classical” spins are aligned after the rotation about the y -axis. From Equations (11) and (16), we can show that this is the case when

$$\beta = \arccos \left(\frac{m'}{\sqrt{l(l+1)}} \right) - \arccos \left(\frac{m}{\sqrt{l(l+1)}} \right). \quad (17)$$

In our example, $m = 0$ and Equation (17) reads $\sin^2(\beta) = m'^2/[l(l+1)]$ and the overlap is greatest at the transition point.

3. ALGORITHM

The evaluation of the $d_{mm'}^l(\beta)$ from Equations (7) and (8) requires starting values for the recursions. For large l , however, those values are often vanishingly small and cannot be represented by any of the IEEE 754 floating-point data formats which are used on practically all current computer hardware. Starting the recursions with 0 and 1 does not help because there will come a point where the $d_{mm'}^l(\beta)$ become too big to be represented numerically. Two solutions to this problem are presented here.

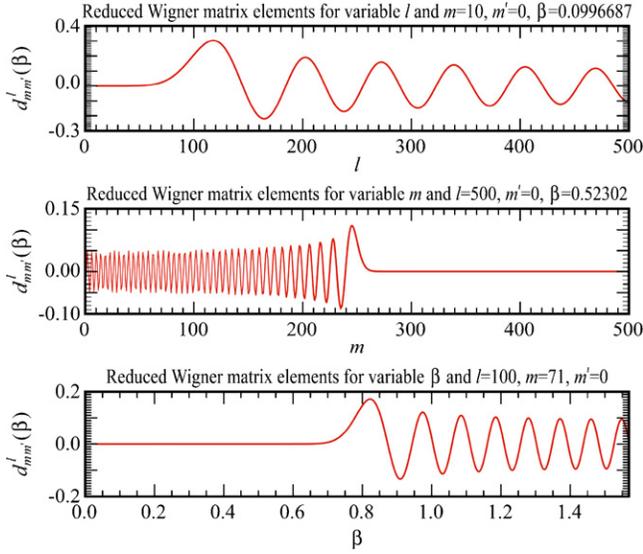


Figure 1. Top plot shows the variation of $d_{010}^l(0.0996687)$ for $10 \leq l \leq 500$, the middle plot shows the variation of $d_{0m}^{500}(0.52331)$ for $0 \leq m \leq 500$, and the bottom plot shows the variation of $d_{071}^{100}(\beta)$ for $0 \leq \beta \leq \pi/2$. (A color version of this figure is available in the online journal.)

3.1. $d_{mm'}^l(\beta)$ Ratios

From Equation (6) and the plots of Figure 1, it is clear that the $d_{mm'}^l(\beta)$ vary smoothly with varying l , m , and β . As a result, a recursion relation of ratios should always be finite in the non-classical region where the $d_{mm'}^l(\beta)$ are not oscillatory, and one only has to worry about singularities in the ratios in the classical/oscillatory region where the denominators could vanish if evaluated at a zero of the $d_{mm'}^l(\beta)$. In this ratio-based method, Equations (7) and (8) can be rewritten as

$$\frac{d_{mm'}^l}{d_{mm'-1}^l} = \frac{\sqrt{(l+m')(l-m'+1)}}{m \operatorname{cosec} \beta - m' \cotan \beta - \sqrt{(l-m')(l+m'+1)} \frac{d_{mm'+1}^{l-1}}{d_{mm'}^l}}, \quad (18)$$

$$\frac{d_{mm'}^l}{d_{mm'}^{l+1}} = \frac{(l+1)\sqrt{(l^2-m'^2)(l^2-m^2)}}{(2l+1)(mm' - \cos \beta) - l\sqrt{[(l+1)^2-m'^2][(l+1)^2-m^2]} \frac{d_{mm'}^{l-1}}{d_{mm'}^l}}. \quad (19)$$

Using

$$\frac{d_{ml}^l}{d_{ml-1}^l} = \frac{\sqrt{l/2} \sin \beta}{l \cos \beta - m} \quad \text{and} \quad (20)$$

$$\frac{d_{ml}^l}{d_{ml}^{l+1}} = \sqrt{\frac{(l+m+1)(l-m+1)}{2l+1}} [(l+1) \cos \beta - m]^{-1}, \quad (21)$$

the ratios can be calculated recursively down to $m' = 0$ if using Equation (18) or up to $l = l_{\max}$ if using Equation (19). For example, in the case where all the $d_{0m'}^l$ for $l \geq m' \geq 0$ are required, one would start with Equation (18) to calculate

$$\frac{d_{0l}^l(\beta)}{d_{0l-1}^l(\beta)}, \frac{d_{0l-1}^l(\beta)}{d_{0l-2}^l(\beta)}, \dots, \frac{d_{02}^l(\beta)}{d_{01}^l(\beta)}, \frac{d_{01}^l(\beta)}{d_{00}^l(\beta)}. \quad (22)$$

Then, to calculate the $d_{0m'}^l$, one would need to know d_{00}^l . Fortunately, the d_{00}^l are easy to calculate because their recursion relation does not contain exponential solutions as remarked under Equation (11):

$$\cos \beta d_{00}^l(\beta) = \frac{l}{2l+1} d_{00}^{l-1} + \frac{l+1}{2l+1} d_{00}^{l+1}. \quad (23)$$

Once $d_{00}^l(\beta)$ has been calculated, d_{01}^l can be calculated from the ratios; in order, each $d_{0m}^l(\beta)$ can be calculated by multiplying adjacent ratios until $d_{0l}^l(\beta)$ has been evaluated. In the case where $d_{0m'}^l$ for $l \geq m' \geq 0$ are also needed, the set of ratios

$$\frac{d_{1l}^l(\beta)}{d_{1l-1}^l(\beta)}, \frac{d_{1l-1}^l(\beta)}{d_{1l-2}^l(\beta)}, \dots, \frac{d_{12}^l(\beta)}{d_{11}^l(\beta)}, \frac{d_{11}^l(\beta)}{d_{10}^l(\beta)} \quad (24)$$

is computed next. To then calculate the $d_{1m'}^l$, one needs to know d_{10}^l . Fortunately, $d_{01}^l = -d_{10}^l$ has previously been calculated and all the $d_{1m'}^l$ can be obtained up to d_{1l}^l . In this fashion, all the $d_{1m'm'}^l(\beta)$ can be calculated up to a desired $m = m_{\max}$.

The special and extremely rare case where a ratio $d_{mm'+1}^l/d_{mm'}^l$ is infinite can be handled by using Equation (7) where $d_{mm'}^l(\beta)$ is set to zero and substituting an infinite ratio and a null ratio for a single finite ratio:

$$\frac{d_{mm'+1}^l}{d_{mm'}^l}, \frac{d_{mm'}^l}{d_{mm'-1}^l} \rightarrow \frac{d_{mm'+1}^l}{d_{mm'-1}^l} = -\frac{\sqrt{(l+m')(l-m'+1)}}{\sqrt{(l-m')(l+m'+1)}}. \quad (25)$$

Note that in contrast to the method described in Risbo (1996), the column of matrix elements $d_{ml}^l(\beta) \rightarrow d_{m0}^l(\beta)$ can be evaluated without having to calculate every single $d_{mm'}^l(\beta)$ for $l' < l$. The same tricks can be applied to the recursion relation in l for the evaluation of $d_{lm'}^l(\beta) \rightarrow d_{lm'}^{\max}(\beta)$. First calculate the column of elements $d_{ll}^{\max}(\beta) \rightarrow d_{l0}^{\max}(\beta)$ and then calculate the ratios

$$\frac{d_{lm'}^l(\beta)}{d_{lm'}^{l+1}(\beta)}, \frac{d_{lm'}^{l+1}(\beta)}{d_{lm'}^{l+2}(\beta)}, \dots, \frac{d_{lm'}^{\max-2}(\beta)}{d_{lm'}^{\max-1}(\beta)}, \frac{d_{lm'}^{\max-1}(\beta)}{d_{lm'}^{\max}(\beta)}. \quad (26)$$

Knowing $d_{lm'}^{\max}(\beta)$ allows us to evaluate $d_{lm'}^{\max-1}(\beta) \rightarrow d_{lm'}^l(\beta)$. For a particular β , it is sufficient to compute the elements of $d_{mm'}^l(\beta)$ for $0 \leq m \leq l$ and $-l \leq m' \leq l$ to know the entire matrix $d^l(\beta)$. To fill out the rest of the matrix, the symmetry relations in the Appendix can be used.

3.2. Wigner Matrix Elements by l -Recursion

Another way to deal with the underflow problem is to start from Equation (8) with the following initialization values:

$$d_{l,m}^l(\beta) = A \left(\cos \frac{\beta}{2} \right)^{l+m} \left(-\sin \frac{\beta}{2} \right)^{l-m}, \quad (27)$$

$$d_{-l,m}^l(\beta) = A \left(\cos \frac{\beta}{2} \right)^{l-m} \left(\sin \frac{\beta}{2} \right)^{l+m}, \quad (28)$$

$$d_{m,l}^l(\beta) = A \left(\cos \frac{\beta}{2} \right)^{l+m} \left(\sin \frac{\beta}{2} \right)^{l-m}, \quad (29)$$

$$d_{m,-l}^l(\beta) = A \left(\cos \frac{\beta}{2} \right)^{l-m} \left(-\sin \frac{\beta}{2} \right)^{l+m}, \quad (30)$$

where $A = \sqrt{(2l)!/(l+m)!(l-m)!}$. As far as Equations (27)–(30) are concerned, the underflow problem can be avoided by simply calculating the logarithm of the absolute value of the matrix element and storing its sign separately. Equation (27), e.g., then transforms to

$$\begin{aligned} \ln |d_{l,m}^l(\beta)| &= 0.5 (\ln((2l)!) - \ln((l+m)!) - \ln((l-m)!)) \\ &\quad + (l+m) \ln |\cos(\beta/2)| + (l-m) \ln |\sin(\beta/2)|. \end{aligned} \quad (31)$$

In cases where one of the last two terms is $-\infty$, the recursion in l can be stopped immediately, since all subsequent values will be zero.

The logarithms of the factorials are easily precomputed, so that the seed value for the recursion can be obtained in $\mathcal{O}(1)$ operations.

Since the result of Equation (31) is in some circumstances much smaller than the individual terms on the right-hand side, cancellation errors may reduce the number of significant digits of the result. In order to have the highest accuracy that can be achieved without sacrificing too much performance, the computation of the seed value is carried out with extended IEEE precision (corresponding to the C++ data type `long double`).

The recursion relation (8) itself unfortunately cannot be computed conveniently in logarithms; therefore, a way must be found to represent floating-point values with an extreme dynamic range, which does not incur a high performance penalty.

This was implemented by representing a floating-point number v using an IEEE double-precision value d and an integer scale n , such that

$$v = d \cdot S^n, \quad (32)$$

where either $d = 0$, or $S^{-1} \leq |d| \leq S$ and S (the “scale factor”) is a positive constant that can be represented as a double-precision IEEE value. Using this prescription, v does not have a unique representation as a (d, n) -pair, but this is not a problem.

Similar techniques have been in use since at least three decades in numerical algorithms; for a recent example, see the spherical harmonic transform routines of the HEALPix package.

It is advantageous to choose a scale factor which is an integer power of 2, because multiplying or dividing by such a factor only affects the exponent of a floating-point value stored in binary format, and is therefore exact (ignoring possible under- or overflows). In order to avoid frequent re-scaling of d , the scale factor should also be rather large; the value adopted for our implementation is 2^{90} .

Using this representation for the $d_{mm'}^l(\beta)$, the recursion is performed until either l_{\max} is reached, or the matrix element has become large enough to be safely represented by a normal double-precision variable (the threshold value used in the code is 2^{-900}). In the latter case, the remaining computations up to l_{\max} are done with standard floating-point arithmetic, which is significantly faster.

4. CONVOLUTION

One area where fast and efficient techniques of computing $d_{mm'}^l(\beta)$ are particularly valuable is in 4π convolution (Wandelt & Gorski 2001). For the convolution of two fields $b(\Omega)$ and $s(\Omega)$

defined on a sphere, the following integral must be calculated:

$$c = \int d\Omega b^*(\Omega) s(\Omega). \quad (33)$$

In the physical application where $b(\Omega)$ is a beam from a horn located on a slowly rotating space telescope that scans the sky (denoted $s(\Omega)$) as it orbits the Sun (e.g., *WMAP* or *Planck* missions), a large number of convolutions must be performed to account for every possible orientation (α, β, γ) of the satellite

$$c(\Omega') = \int d\Omega [R(\Omega') b(\Omega)]^* s(\Omega), \quad (34)$$

where $R(\Omega')$ is a rotation matrix that rotates the beam to a particular orientation of the satellite and is defined

$$R(\alpha, \beta, \gamma) Y_{lm}(\theta, \phi) = \sum_{m'=-l}^l D_{m'm}^l(\alpha, \beta, \gamma) Y_{lm'}(\theta, \phi), \quad (35)$$

and where the $Y_{lm}(\theta, \phi)$ are spherical harmonics. The $Y_{lm}(\theta, \phi)$ are related to the $D_{mm'}^l(\alpha, \beta, \gamma)$ through the relation

$$Y_{lm}(\theta, \phi) = (-1)^m \sqrt{\frac{2l+1}{4\pi}} D_{0m}^l(0, \theta, \phi) \quad (36)$$

and can therefore be calculated using the methods described above.

For beams with significant side-lobes stemming from reflections of light far away from the line of sight as is the case for both the *WMAP* and *Planck* missions, the beams can cover a significant portion of the sky and full-sky convolutions are necessary; as shown in Wandelt & Gorski (2001), such full-sky convolutions are much faster when performed in harmonic space instead of pixel space. We now describe a very fast and massively parallel method to perform full-sky convolutions in harmonic space.

4.1. Conviqt

Conviqt (CONvolution VIa the Quantum Top equation) is a fast 4π convolution algorithm that relies on fast computational methods for reduced Wigner matrix elements. Starting from Equations (34) and (35), the beam and sky fields can be expanded on the spherical harmonic basis to yield

$$c(\alpha, \beta, \gamma) = \sum_{m_b=-m_{b\max}}^{m_{b\max}} \sum_{m_{sky}=-l_{\max}}^{l_{\max}} e^{im_b\alpha} e^{im_{sky}\gamma} C_{m_b m_{sky}}(\beta), \quad (37)$$

$$C_{m_b m_{sky}}(\beta) \equiv \sum_{l=0}^{l_{\max}} b_{lm_b}^* d_{m_b m_{sky}}^l(\beta) s_{lm_{sky}}, \quad (38)$$

where b_{lm_b} and $s_{lm_{sky}}$ are the spherical components of the beam and sky fields.⁵ Typically, the b_{lm_b} are negligible for some $m_{b\max} < m_b \ll l_{\max}$. Noting that the number of β angles needed for the convolution scales as l_{\max} , the evaluation of Equation (37) scales as $\mathcal{O}(l_{\max}^2 m_{b\max} \log(l_{\max}))$ after the use

⁵ For ease of reading, only the scalar case is described since the generalization to polarized maps and beams is easily accomplished by evaluating Equation (38) for the additional pairs $(b_{lm_b}^G, s_{lm_{sky}}^G)$ and $(b_{lm_b}^C, s_{lm_{sky}}^C)$.

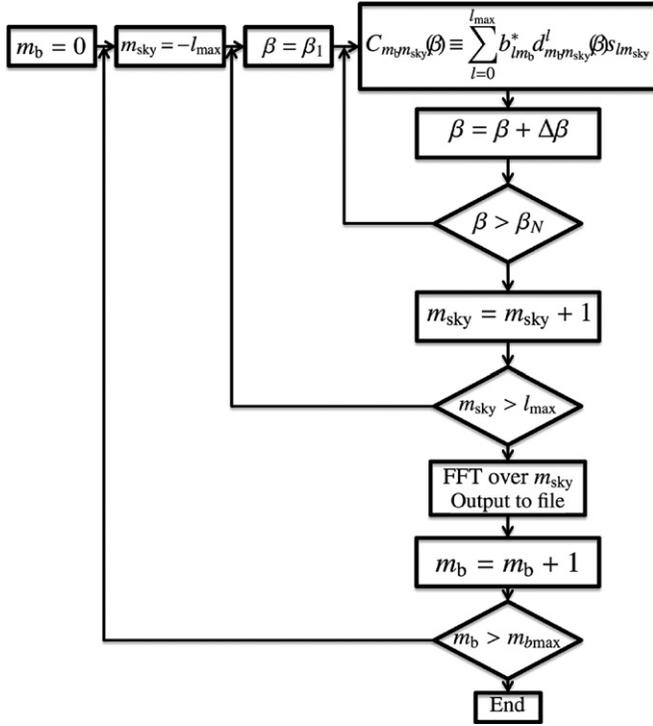


Figure 2. Calculation of the $C_{m_b, m_{sky}}(\beta)$ for an interval of angles $\beta_1 \leq \beta \leq \beta_N$ requires three nested loops over β , m_{sky} , and m_b as shown above.

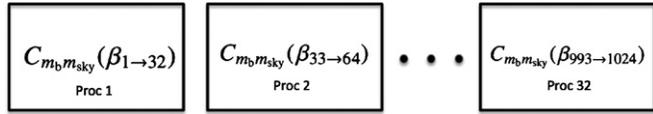


Figure 3. Generic parallelization of *conviqt* where the convolution is performed on 32 processors and the range $[0, \pi/2]$ is split into $N - 1$ intervals with $N = 1024$. The first 32 $C_{m_b, m_{sky}}(\beta)$ are calculated on processor 1 for 32 β 's with $0 \leq \beta \leq 31 \times \pi/(N - 1)$, and so on for each processor. Note that the range $\pi/2 < \beta \leq \pi$ can be obtained from the symmetries of the $d_{mm'}^l(\beta)$.

of the Fast Fourier Transform algorithm to perform the summations. The numerically expensive part of Equation (37) is the computation of Equation (38) which scales as $O(l_{max}^3 m_{bmax})$. Two separate computations of Equation (38) scale as $O(l_{max}^3 m_{bmax})$: the computation of the $d_{m_b, m_{sky}}^l(\beta)$, and the evaluation of the sum over $0 \leq l \leq l_{max}$ for every single m , m_{sky} , and β which form nested loops as seen in Figure 2. The fast methods described in the previous section are used to compute the $d_{m_b, m_{sky}}^l(\beta)$. To evaluate the $C_{m_b, m_{sky}}(\beta)$ for each β , a massively parallel MPI-based approach is used since the $C_{m_b, m_{sky}}(\beta)$ are uncorrelated between the different β and can be computed by different tasks as shown in Figure 3. As seen in Section 4.3.2, *conviqt* scales very well with the number of processors and its scaling can be described as

$$\text{conviqt scaling} \sim l_{max}^3 m_{bmax}/n, \quad (39)$$

where n is the number of processors. Additional acceleration techniques for both the computation of the $d_{m_b, m_{sky}}^l(\beta)$ and for the evaluation of the sums over l are described in the following subsection.

4.2. Acceleration Techniques for the $d_{m_b, m_{sky}}^l(\beta)$

In simulations of the measurement of cosmic microwave background, convolutions appear repeatedly especially if Monte

Carlo are required. Since the generation of Wigner matrix elements is typically the most computationally intensive part of the convolution algorithm, a large effort was made to increase its efficiency. This has two aspects: first to compute the matrix elements as quickly as possible, but also to decide (if possible) which matrix elements are too small to contribute measurably to the result and skip their generation altogether.

4.2.1. Skipping Unneeded Calculations

When performing convolutions, especially within the context of Monte Carlo simulations where many convolutions with the same l_{max} and m_{bmax} are needed, it is computationally profitable to skip unneeded calculations.⁶ Some terms in the sum of Equation (38) need not be included because their $d_{m_b, m_{sky}}^l(\beta)$ are vanishingly small. To determine which terms to exclude, we turn to Equation (11) where three general possibilities are considered:

1. m_b and m_{sky} are of similar magnitude and much smaller than l .
2. m_b and m_{sky} are of similar magnitude and of the same order of magnitude as l .
3. m_b and m_{sky} are of widely differing magnitude with one much smaller than l and one of similar magnitude.

In each of these possibilities, the neglected $d_{m_b, m_{sky}}^l(\beta)$ are those evaluated at β angles that correspond to the non-classical, exponentially suppressed region. Before explaining how these $d_{m_b, m_{sky}}^l(\beta)$ are identified, *conviqt*'s nested structure should be described. In *conviqt*, the outermost loop deals with m_b (which ranges from 0 to m_{bmax} ; nested into the m_b -loop is the m_{sky} -loop ranging from $-m_{sky}$ to m_{sky} ; nested in the m_{sky} -loop is the loop over the β processed by that particular task.⁷ Finally, the innermost loop is that over l which is where the condition is applied. To derive the minimum l such that outside the parameter space defined by $l_{min} \leq l \leq l_{max}$, $d_{m_b, m_{sky}}^l(\beta)$ is negligible, we use Equation (11) to write

$$l_{min} = -\frac{1}{2} + \frac{1}{2} \left[1 + \frac{4}{\sin^2 \beta} (m_{sky}^2 + m_b^2 - 2m_{sky}m_b \cos \beta) \right]^{1/2} \\ \cong \frac{\sqrt{m_{sky}^2 + m_b^2 - 2m_{sky}m_b \cos \beta - \text{offset}}}{\sin \beta}, \quad (40)$$

where $\text{offset} > 0$ and ensures that the $d_{m_b, m_{sky}}^l(\beta)$ neglected are well within the non-classical region and suppressed. Calling m_{big} the larger of $|m_b|$ or $|m_{sky}|$, it is noted that $l \geq m_{big}$. To determine the offset , we go back to the three possibilities listed above; of those, l_{min} will generally equal m_{big} in the cases where m_b and m_{sky} are of similar magnitude; only in the third case will we generally have $l_{min} > m_{big}$, namely when m_b and m_{sky} are of widely differing magnitude. From Rowe et al. (2001), this case can approximately be written as a harmonic oscillator wave function:

$$d_{m_b, m_{sky}}^l(\beta) \rightarrow (-1)^{l-m_{sky}} (\sqrt{l} \sin \beta_m)^{-1/2} u_{l-m_{sky}}(\sqrt{l}(\beta - \beta_m)) \quad (41)$$

$$u_\nu(x) = (\sqrt{\pi} 2^\nu \nu!)^{-1/2} H_\nu(x) e^{-x^2/2}, \quad (42)$$

⁶ Note that it is generally not more efficient to evaluate all of the $d_{m_b, m_{sky}}^l(\beta)$ beforehand because of the disk space required and the large amount of time needed to read them in; it is more efficient to calculate them on the fly.

⁷ A note to remind the reader that *conviqt* is parallelized in β ; each task will perform the convolutions in a subset of all the complete set of β where $C_{m_b, m_{sky}}(\beta)$ must be calculated.

where $H_\nu(x)$ is a Hermite polynomial and $\cos \beta_m = m_b/l$. Since we are dealing with orders of magnitude, it is not necessary to exactly evaluate Equation (41) to determine `offset`, only to calculate an estimate from the factor $\exp(-x^2/2)$. We have found that using `offset` $\geq l_{\max}/20$ gives extremely accurate results. Finally, it is noted that $d_{m_b m_{\text{sky}}}^l(0) = \delta_{m_b m_{\text{sky}}}$ and for that special case we put $l_{\min} > l_{\max}$ when $m_b \neq m_{\text{sky}}$ and no sum over l is performed. Thus, Equation (40) is used to estimate whether the absolute values of all $d_{m_b m_{\text{sky}}}^l(\beta)$ for a given combination of l , m_b , m_{sky} , and β lie below a certain threshold; if this is the case, the generation of these values can be skipped entirely. In particular, the most efficient version of the code written was one where the $d_{m_b m_{\text{sky}}}^{l_{\min}}$ and $d_{m_b m_{\text{sky}}}^{l_{\min}+1}$ were precalculated for $0 \leq m_b \leq m_{b\max}$, $-l_{\max} \leq m_{\text{sky}} \leq l_{\max}$, and β subset for a particular task, and read-in as seeds to the recursion relation of Equation (8). That way, none of the unneeded $d_{m_b m_{\text{sky}}}^l(\beta)$ were calculated during the convolution. This required an extra code to precompute the $d_{m_b m_{\text{sky}}}^l(\beta)$. In the end, we opted for a single code based on the evaluation of the reduced Wigner matrix elements as described in Section 3.2 because of the low overhead and maintenance as well as the high efficiency.

In this approach, we calculate all the $d_{m_b m_{\text{sky}}}^l(\beta)$ on the fly, but only include the relevant $d_{m_b m_{\text{sky}}}^l(\beta)$ in the final l -loop. This code is self-contained and easier to maintain at a very minimal cost in performance. As the $d_{m_b m_{\text{sky}}}^l(\beta)$ recursion is performed, the code checks the absolute values of the generated $d_{mm}^l(\beta)$ and records the l index at which a predefined threshold ε (typically set to 10^{-30}) is crossed for the first time. Due to the limited dynamic range of IEEE data types, values below this threshold have no measurable influence on the convolution result and can therefore be neglected during the final summation loop, which saves a significant amount of CPU time.

4.2.2. Precomputed Values

In this single-code approach, a precomputation strategy well suited to the loop structure was adopted.

1. At the beginning of each run, we compute just once $\ln(\cos(\beta/2))$, $\ln(\sin(\beta/2))$, and $\cos \beta$ for all β at which we need the Wigner matrix, and as mentioned above, $\ln n!$ up to $n = 2l_{\max}$.
2. Also at the beginning we compute the tables

$$P_i = \sqrt{1/(i+1)} \quad \text{and} \quad Q_i = \sqrt{i/(i+1)}$$

for i in the range of 0 to $2l_{\max} + 1$, which are needed for the next precomputation step.

3. Inside the second loop (i.e., for every combination of m_b and m_{sky}), we compute the tables

$$\begin{aligned} F_{0,l} &= (l+1)(2l+1)P_{l+m_b}P_{l-m_b}P_{l+m_{\text{sky}}}P_{l-m_{\text{sky}}}, \\ F_{1,l} &= m_b m_{\text{sky}} / (l(l+1)), \text{ and} \\ F_{2,l} &= Q_{l+m_b}Q_{l-m_b}Q_{l+m_{\text{sky}}}Q_{l-m_{\text{sky}}}(l+1)/l. \end{aligned}$$

After all these preparations, Equation (8) boils down to

$$d_{m_b m_{\text{sky}}}^{l+1}(\beta) = F_{0,l}(\cos \beta - F_{1,l})d_{m_b m_{\text{sky}}}^l(\beta) - F_{2,l}d_{mm}^{l-1}(\beta), \quad (43)$$

which corresponds to only five quick-to-compute floating-point operations.

The space overhead for the additional tables is $\mathcal{O}(l_{\max})$, which is insignificant compared to the $\mathcal{O}(l_{\max}^2)$ memory requirement

of the whole convolution code. This also means that for the reasonable assumption of $l_{\max} \lesssim 10^4$ all data required for the recursion fit conveniently into current processors' Level-2 caches.

4.2.3. Use of $d_{mm}^l(\beta)$ Symmetries

The use of the $d_{mm}^l(\beta)$ symmetries considerably cut the computational cost of the full sky convolution. In particular, Equation (A6) relates the computed values of $C_{m_b m_{\text{sky}}}(\beta)$ at $\beta < \pi/2$ to those at $\beta > \pi/2$. In Equation (A6), the phase factor $(-1)^l$ is accounted for by splitting the sum in Equation (38) into even and odd l . The phase factor $(-1)^{-m}$ is accounted for by using the relations

$$b_{lm} = (-1)^m b_{l-m}^*, \quad s_{lm} = (-1)^m s_{l-m}^* \quad (44)$$

and further splitting the odd and even sums of Equation (38) into real and imaginary parts. In addition, the symmetry of Equations (A1) and (44) can be used to show

$$C_{-m_b -m_{\text{sky}}}(\beta) = C_{m_b m_{\text{sky}}}(\beta)^* \quad (45)$$

speeding up the computation of $C_{m_b m_{\text{sky}}}(\beta)$ by another factor of 2.

4.3. Example Simulations

To determine the accuracy and performance of `convigt`, a detailed comparison with the stable release of the `LevelS totalconvolver` (Wandelt & Gorski 2001; Reinecke et al. 2006) currently compiled on the `planck` cluster at the National Energy Research Science Council (NERSC) was performed. `LevelS` is a simulation package for the generation of time-ordered data (TOD) by the *Planck* satellite (<http://www.esa.int/SPECIALS/Planck/index.html>). `Totalconvolver` and `convigt` both calculate a data cube that is fully compatible with `LevelS`. For both codes, data cube is composed of convolved points calculated at a polar angle θ , a longitudinal angle ϕ , and a particular beam orientation (a rotation about the beam axis) ψ . In Section 4.3.1, we compare the output data of both codes in order to detect potential significant discrepancies hinting at code bugs or numerical problems. In Section 4.3.2, we first compare the performance of both codes when run on a single processor. Afterward the scaling behavior of `convigt` to very high task numbers is demonstrated. We cannot provide a comparison with `totalconvolver` for this experiment, since due to its internal structure it can only run on shared-memory architectures and is therefore limited to a relatively low degree of parallelism.

4.3.1. Data Cube Comparison

For $l_{\max} = 2000$, GRASP beams for LFI-19a, $m_{b\max} = 9$, `offset` = 30 (see Equation (40)) and a polarized CMB map, there were 153,634,399 points in the data cube. Taking the difference between the `totalconvolver` and `convigt` data cubes (residual values below), we obtained Table 1 where “avr” refers to the average difference of the two data cubes (`convigt`(θ, ϕ, ψ) - `totalconvolver`(θ, ϕ, ψ)), σ is the variance of that residual data cube, “Max” refers to the maximum value found in the residual data cube, and “rel” refers to the ratio of σ to the variance of the `totalconvolver` data cube. We see that with `offset` = 2000, the two data cubes agree to approximately eight significant digits; for `offset` = 30 and 15, they agree to six and four significant digits, respectively.

Table 1Convolution Output Comparison Between `convqt` and `Totalconvolver` for Different Offsets

Offset	Avr	σ	Max	Rel
2000 (exact)	$-1.7e-16$	$4.1e-13$	$1.4e-11$	$4.0e-8$
30	$-1.2e-16$	$1.3e-11$	$5.6e-10$	$1.3e-6$
15	$-1.6e-16$	$3.0e-9$	$1.1e-7$	$2.9e-4$

Table 2Comparison in Timing and Memory Consumption Between `convqt` and `Totalconvolver` for $l_{\max} = 2000$, $m_{b\max} = 9$, and `Offset = 30`

Code	Clock (s)	GBytes
<code>Convqt</code>	349	0.37
<code>Totalconvolver</code>	1120	0.41

Table 3Timing and Memory Consumption Comparison Between `convqt` and `Totalconvolver`

l_{\max}	<code>Convqt</code>		<code>Totalconvolver</code>	
	s	GBytes	s	GBytes
256	1.9657	$1.51806e-02$	11.726	$2.32534e-02$
512	9.8691	$1.76401e-02$	54.634	$5.00412e-02$
768	27.982	$2.16856e-02$	144.51	$9.46579e-02$
1024	61.797	$2.73190e-02$	294.71	$1.57092e-01$
1280	117.13	$3.45383e-02$	513.29	$2.37350e-01$
1536	199.29	$4.33445e-02$	809.51	$3.35554e-01$
1792	313.27	$5.37376e-02$	1264.1	$4.51582e-01$
2048	621.26	$6.57196e-02$	1765.2	$5.85376e-01$

4.3.2. Performance

On a single processor on Jacquard, for $l_{\max} = 2000$ with a GRASP beam (LFI-19a) of $m_{b\max} = 9$, `offset = 30`, MC = T including polarization, `convqt` had the performance shown in Table 2 where “Clock” refers to the time it took to complete the convolution according to the wall clock, while “GBytes” refers to the total memory consumed. These numbers were obtained using NERSC’s Integrated Performance Monitoring (IPM) on a 712-CPU Opteron cluster called Jacquard running a Linux operating system; each processor on Jacquard runs at a clock speed of 2.2 GHz with a theoretical peak performance of 4.4 GFlop s^{-1} . Unlike `totalconvolver`, `convqt` is a massively parallel code which can be run on machines with distributed memory; running it on a single processor shows that `convqt` is intrinsically faster and more efficient than `totalconvolver`. The above table is for the case where `offset < l_{\max}`, i.e., the case where `convqt` sacrifices precision for the sake of a speedier convolution.

For the general case where `convqt` and `totalconvolver` calculate the same thing (`offset` $\rightarrow \infty$), we use the single-code approach to obtain Table 3: these numbers are plotted in Figure 4. The top plot shows that `convqt` is considerably faster than `totalconvolver` for $l_{\max} < 2048$; however, because both codes scale as l_{\max}^3 as $l_{\max} \rightarrow \infty$, the gap between their total wall clock times will narrow. It is also seen that `convqt` consumes significantly less memory.

The scaling of `convqt` timings as a function of the total number of processors is very good as can be seen from Table 4

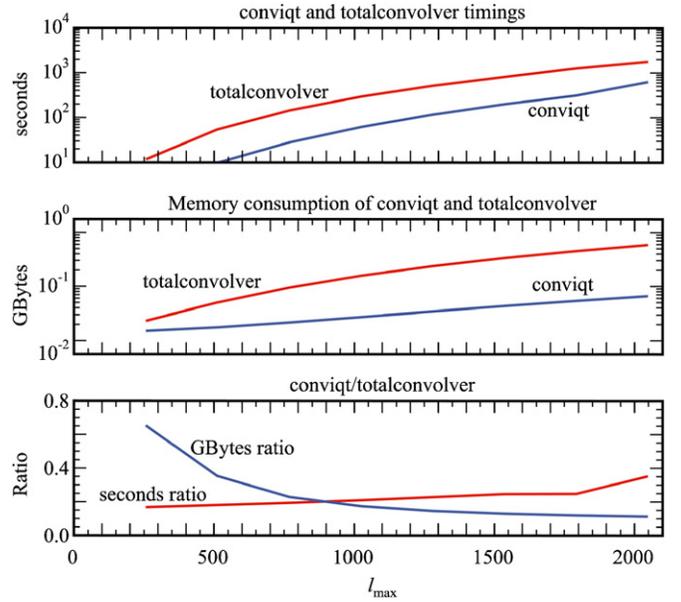


Figure 4. Top plot compares the wall clock performance of `convqt` and `totalconvolver` at l_{\max} intervals of 256 starting at $l_{\max} = 256$ and the middle plot compares the memory needs in GBytes of the two codes. The lower plot shows the ratio of wall clock and GBytes of the two codes.

(A color version of this figure is available in the online journal.)

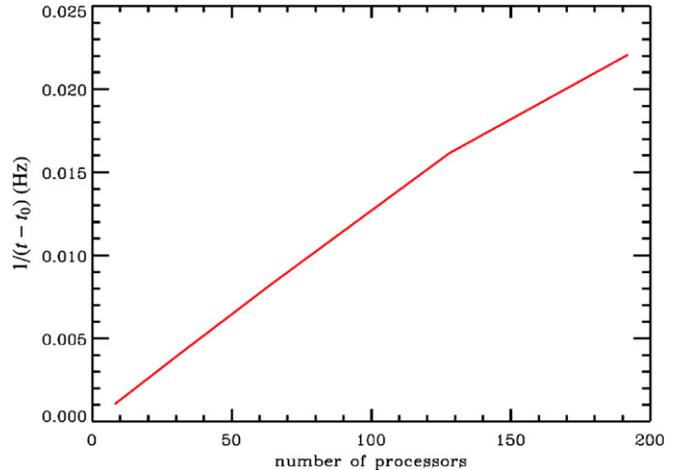


Figure 5. Timings of `convqt` as the number of processors is increased with $l_{\max} = 4096$ and $m_{b\max} = 14$.

(A color version of this figure is available in the online journal.)

Table 4Convqt Scaling with the Number of Processors, with $t_0 = 7.15$ s

Number of Processors	Seconds	$1/(t - t_0)$ (Hz)
8	965.46	0.0010435037
16	486.29	0.0020870727
32	247.83	0.0041548945
64	128.43	0.0082453826
128	69.06	0.016152479
192	52.47	0.022065313

and the plot in Figure 5. For $l_{\max} = 4096$ and $m_b = 14$ with polarized beam and sky, the plot in Figure 5 shows an inverse relationship between the timings t and the number of processors n

$$n = t_1/(t - t_0), \quad (46)$$

where $t_0 = 7.15$ s and $t_1 = 7666.72$ s are parameters that can be thought of, respectively, as the time needed for the intrinsically scalar part of the code and the time needed to perform the convolution on a single processor if $t_0 = 0$. This relationship is satisfied up to a convolution distributed on 128 processors. This plot was obtained using the single-code approach run on the NERSC cluster called `Planck`, a 256 cores cluster of Opteron 2350 2.0 GHz processors. To measure the scaling behavior of `convigt`, no output file was created to avoid skewing the scaling law with the time it takes to write the file (tens of seconds for a 4 GB file). As the number of processors increases and the time required to perform the convolution diminishes to less than a minute, the timings become dominated with operations that have nothing to do with the convolution; among these are the reading of the input data sets (which are read in full by all MPI tasks), the interprocess communication and various calculations which are performed redundantly on all tasks, because communicating the results would be more expensive. Increasing the number of tasks (while keeping the problem size constant) also means a smaller number of β angles per task, which decreases the achievable quality of load balancing. In addition, different runs with identical inputs show variations of a few seconds in wall clock timings that have an increasing relative impact on the decreasing timings stemming from using larger numbers of processors; the most likely explanation for this is differences in the exact nature of process startup and disk access, which is not exactly reproducible in this kind of computing environment.

5. SUMMARY

New algorithms for the efficient and accurate calculation of Wigner matrix elements were presented. These algorithms were used in a full sky convolution, massively parallel algorithm called `convigt` that was shown to be significantly more efficient and much faster than the only other algorithm currently available.

G.P. thanks Maura Sandri for the GRASP 8 beams used in these simulations and Charles Lawrence for useful comments on this manuscript. We gratefully acknowledge support by the NASA Science Mission Directorate via the US Planck Project. The research described in this paper was partially carried out at the Jet propulsion Laboratory, California Institute of Technology, under a contract with NASA. This research used

resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract DE-AC02-05CH11231. M.R. is supported by the German Aeronautics Center and Space Agency (DLR), under program 50-OP-0901, funded by the Federal Ministry of Economics and Technology.

APPENDIX:

SYMMETRIES OF $d_{mm'}^l(\beta)$

$$d_{mm'}^l(\beta) = (-1)^{m-m'} d_{-m-m'}^l(\beta), \quad (\text{A1})$$

$$d_{mm'}^l(\beta) = (-1)^{m-m'} d_{m'm}^l(\beta), \quad (\text{A2})$$

$$d_{mm'}^l(\beta) = d_{-m'-m}^l(\beta), \quad (\text{A3})$$

$$d_{mm'}^l(-\beta) = d_{m'm}^l(\beta), \quad (\text{A4})$$

$$d_{mm'}^l(-\beta) = (-1)^{m-m'} d_{mm'}^l(\beta), \quad (\text{A5})$$

$$d_{mm'}^l(\pi - \beta) = (-1)^{l-m'} d_{-mm'}^l(\beta), \quad (\text{A6})$$

$$d_{mm'}^l(\pi - \beta) = (-1)^{l+m'} d_{m-m'}^l(\beta). \quad (\text{A7})$$

REFERENCES

- Edmonds, A. R. 1957, *Angular Momentum in Quantum Mechanics* (Princeton, NJ: Princeton Univ. Press)
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1988, *Numerical Recipes in C++* (Cambridge: Cambridge Univ. Press)
- Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Ensslin, T. 2006, *A&A*, **445**, 373
- Risbo, T. 1996, *J. Geod.*, **70**, 383
- Rowe, D. J., de Guise, H., & Sanders, B. C. 2001, *J. Math. Phys.*, **42**, 2315
- Varshalovich, D. A., Moskalev, A. N., & Khersonskii, V. K. 1988, *Quantum Theory of Angular Momentum* (Singapore: World Scientific)
- Wandelt, B. D., & Gorski, K. M. 2001, *Phys. Rev. D.*, **63**, 123002