

Analysis of cosmic microwave background data on an incomplete sky

Daniel J. Mortlock,^{1,2★} Anthony D. Challinor^{1★} and Michael P. Hobson^{1★}

¹*Astrophysics Group, Cavendish Laboratory, Madingley Road, Cambridge CB3 0HE*

²*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*

Accepted 2001 October 19. Received 2001 June 20; in original form 2000 August 10

ABSTRACT

Measurement of the angular power spectrum of the cosmic microwave background is most often based on a spherical harmonic analysis of the observed temperature anisotropies. Even if all-sky maps are obtained, however, it is likely that the region around the Galactic plane will have to be removed as a result of its strong microwave emissions. The spherical harmonics are not orthogonal on the cut sky, but an orthonormal basis set can be constructed from a linear combination of the original functions. Previous implementations of this technique, based on Gram–Schmidt orthogonalization, were limited to maximum Legendre multipoles of $l_{\max} \lesssim 50$, as they required all the modes have appreciable support on the cut-sky, whereas for large l_{\max} the fraction of modes supported is equal to the fractional area of the region retained. This problem is solved by using a singular value decomposition to remove the poorly supported basis functions, although the treatment of the non-cosmological monopole and dipole modes necessarily becomes more complicated. A further difficulty is posed by computational limitations – orthogonalization for a general cut requires $O(l_{\max}^6)$ operations and $O(l_{\max}^4)$ storage and so is impractical for $l_{\max} \geq 200$ at present. These problems are circumvented for the special case of constant (Galactic) latitude cuts, for which the storage requirements scale as $O(l_{\max}^2)$ and the operations count scales as $O(l_{\max}^4)$. Less clear, however, is the stage of the data analysis at which the cut is best applied. As convolution is ill-defined on the incomplete sphere, beam-deconvolution should not be performed after the cut and, if all-sky component separation is as successful as simulations indicate, the Galactic plane should probably be removed immediately prior to power spectrum estimation.

Key words: methods: analytical – methods: numerical – cosmic microwave background.

1 INTRODUCTION

Since the first measurements of the temperature anisotropy of the cosmic microwave background (CMB) by the *Cosmic Background Explorer (COBE)* satellite (Smoot et al. 1992), a number of sophisticated experiments have been undertaken to measure the fluctuations at higher resolutions and sensitivities (e.g. Scott et al. 1996; Tanaka et al. 1996; Netterfield et al. 1997; de Oliveira-Costa et al. 1998; Coble et al. 1999; de Bernardis et al. 2000; Wilson et al. 2000; Padin et al. 2001; Lee et al. 2001; Halverson et al. 2002; and Netterfield et al. 2002). The primary result of these experiments has been the measurement of the angular power spectrum of the CMB to Legendre multipoles of up to $l \approx 1000$, which places strong constraints on a number of cosmological parameters (Lineweaver 1998; Efstathiou et al. 1999; de Bernardis et al. 2000; Netterfield et al. 2001; Wang, Tegmark & Zaldarriaga 2001 and references therein). In the future the *Microwave Anisotropy Probe (MAP)*; e.g.

Jarosik et al. 1998) and the *Planck* satellite (e.g. Bersanelli et al. 1996) will produce maps of the microwave sky with resolutions of between 5 and 30 arcmin at a number of frequencies. Such extraordinary data sets, consisting of millions of independent measurements, will clearly require novel analysis techniques.

One of the many difficulties is the treatment of the non-cosmological contributions to the observed microwave sky. Dust, synchrotron and free–free emission from the Galaxy (e.g. Haslam et al. 1982; Schlegel, Finkbinder & Davies 1998); radio galaxies and other extra-Galactic ‘point’ sources (e.g. Toffolatti et al. 1998); and the Sunyaev–Zel’dovich (Sunyaev & Zel’dovich 1970) effect caused by galaxy clusters (e.g. Birkinshaw 1999) all obscure the CMB at some level [see Hu, Sugiyama & Silk (1997) or Barreiro (2000) for more complete reviews], although these components have quite distinct spectral properties and so can be separated using multi-frequency observations (e.g. Bennett et al. 1992; Tegmark & Efstathiou 1996; Hobson et al. 1998; Bouchet & Gispert 1999; Jones, Hobson & Lasenby 1999; Baccigalupi et al. 2000). However, these techniques are not likely to be able to extract the Galactic emissions completely (Stolyarov et al. 2001), leaving the removal of the

★E-mail: mortlock@ast.cam.ac.uk (DJM); a.d.challinor@mrao.cam.ac.uk (ADC); mph@mrao.cam.ac.uk (MPH)

Galactic plane as the only option. The Galaxy contributes relatively little at high latitudes (e.g. Haslam et al. 1982; Schlegel et al. 1998), so this is an acceptable, if not optimal, solution. For instance, Górski et al. (1994) removed the band within 20° of the Galactic plane to estimate the power spectrum of the 2-yr *COBE* Differential Microwave Radiometer sky maps, and similar cuts have been proposed by both the *MAP* and *Planck* collaborations. An essentially equivalent problem is posed if the sky coverage of the survey is incomplete, although there is less choice about the geometry of the cut in this case.

A number of aspects of the analysis become more difficult on an incomplete sphere, one of the most obvious reasons being that the spherical harmonics are no longer an orthonormal basis set. The most successful component separation techniques to date (e.g. Hobson et al. 1998; Bouchet & Gispert 1999) rely on a mode-by-mode analysis which explicitly utilizes the orthogonality of the spherical harmonics, although it may be preferable to remove the Galactic plane only when estimating the CMB power spectrum. Unbiased power spectrum estimation using the spherical harmonics is possible on the cut sky (Wandelt, Górski & Hivon 2001), but the covariance structure of the resulting pseudo-harmonics is far from ideal, so analysis using an orthonormal basis set is preferable. In particular the noise covariance matrix remains diagonal in the case of spatially uniform (and uncorrelated) noise.

It is possible to construct an orthonormal basis set from linear combinations of the spherical harmonics, and an elegant implementation of this, based on Cholesky decomposition of the coupling matrix of the spherical harmonics on the cut sky, was described by Górski (1994). However the coupling matrix becomes ill-conditioned for $l_{\max} \gtrsim 50$, and so this method cannot be used to perform cut-sky orthogonalization for either the *MAP* experiment (with $l_{\max} \approx 1500$) or the *Planck* mission (with $l_{\max} \approx 2500$).

A general formalism for orthogonalization of the spherical harmonics is presented in Section 2, although implementation to high l_{\max} is only possible at present in the special case of a constant latitude cut. The relationship between the various harmonic coefficients is discussed in Section 3, and the extension of these results to CMB analysis techniques (specifically component separation and power spectrum estimation) is covered in Section 4. The results are summarized and future possibilities are discussed in Section 5. Finally, the chosen conventions for the spherical harmonics are defined in Appendix A; formulae for integrals of the products of Legendre functions are given in Appendix B; and the treatment of the non-cosmological monopole and dipole modes is discussed in Appendix C.

2 ORTHOGONALIZATION OF SCALAR BASIS FUNCTIONS

The physics of the CMB is most naturally expressed in Fourier space, and it is standard practice to represent sky maps by their harmonic coefficients. The basis functions chosen here are the real spherical harmonics, $Y_{l,m}(\hat{r})$ (as defined in Appendix A), which form an orthonormal basis on the complete sphere, S . In general, $l \geq 0$ and $-l \leq m \leq l$, although in practice a finite l_{\max} must be used, which implies a band limit. It is convenient to combine the two indices, allowing the basis set to be expressed as a vector, $\mathbf{Y}(\hat{r}) = [Y_1(\hat{r}), Y_2(\hat{r}), \dots, Y_{i_{\max}}(\hat{r})]^T$, where $i_{\max} = (l_{\max} + 1)^2$. There are several reasonable choices for the indexing, $i(l,m)$, most notably grouping coefficients in l or m , as defined in Appendix A. Grouping in l is most natural for power spectrum estimation, but grouping in m is more efficient computationally in cases of azimuthal symmetry (Section 2.3).

The spherical harmonics are not orthogonal on the incomplete sphere S' , as can be seen from the structure of their coupling matrix (Section 2.1). A decomposition of the coupling matrix can be used to construct an orthonormal basis set (Section 2.2), but implementation to high resolution is currently possible only in the special case of constant latitude cuts (Section 2.3).

2.1 The coupling matrix

The coupling matrix of a set of functions encodes their orthogonality and normalisation properties over a given range. In the case of the spherical harmonics on the incomplete sphere it is given by

$$\mathbf{C} = \int_{S'} \mathbf{Y}(\hat{r}) \mathbf{Y}^T(\hat{r}) d\Omega. \quad (1)$$

If $S' = S$ then the harmonics are orthonormal and $\mathbf{C} = \mathbf{I}$; otherwise the off-diagonal elements are non-zero, indicating that the basis functions are non-orthogonal.

An alternative formulation is to introduce a window function, $w(\hat{r})$, so that

$$\mathbf{C} = \int_S w^2(\hat{r}) \mathbf{Y}(\hat{r}) \mathbf{Y}^T(\hat{r}) d\Omega. \quad (2)$$

In some ways this approach is more flexible, as $w(\hat{r})$ can either be a smoothly-varying apodizing function (cf. Tegmark 1997) or take the form

$$w_{S'}(\hat{r}) = \begin{cases} 1, & \text{if } \hat{r} \in S', \\ 0, & \text{if } \hat{r} \in S - S', \end{cases} \quad (3)$$

mimicking the effect of the sharp cut defined above. However, this definition of the window function can lead to inconsistencies if a band-limited analysis is carried out, as $w_{S'}(\hat{r})$ cannot be properly represented by a finite analysis (see Section 3.2). It is for this reason that the first formalism is used here, although most of the subsequent results can also be derived using window functions.

For a pixel-based analysis, the coupling matrix can be defined by replacing the integral in equation (1) by a sum over points on the sphere (i.e. pixel centres), \hat{r}_p , where $1 \leq p \leq N_p$ and N_p is the number of pixels. In this case

$$\mathbf{C} = \sum_{p=1}^{N_p} \mathbf{Y}(\hat{r}_p) \mathbf{Y}^T(\hat{r}_p) \Omega_p, \quad (4)$$

where Ω_p is the area of the p th pixel and there is no pixel smoothing (cf. Górski 1994). In the limit $N_p \rightarrow \infty$, equations (1) and (4) become equivalent and pixelization issues become irrelevant. If the points are uniformly distributed over the sphere \mathbf{C} should be close to the identity, the small discrepancies being due to the approximation of the integral as a sum; otherwise \mathbf{C} reflects the spatial distribution of the points much as in the continuum case, but there is freedom to represent apodizing filters as well as discrete cuts.

Assuming $\Omega_{S'} > 0$, the coupling matrix is formally symmetric, positive definite and invertible, irrespective of which of the above definitions is used. However, \mathbf{C} rapidly becomes numerically singular: e.g. if $l_{\max} = 50$, the condition number¹ of \mathbf{C} is $\sim 5 \times 10^9$

¹ The condition number of a matrix is (the absolute value of) the ratio of its greatest and smallest eigenvalues; it is large for ill-conditioned matrices, and infinite for singular matrices.

for a constant latitude cut of $b_{\text{cut}} = \pm 20^\circ$. This can be further understood in terms of the eigenstructure of the coupling matrix.

2.1.1 Eigenstructure

The coupling matrix has $i_{\text{max}} = (l_{\text{max}} + 1)^2$ eigenvectors, \mathbf{v}_i , and eigenvalues, λ_i , which satisfy

$$\mathbf{C}\mathbf{v}_i = \lambda_i\mathbf{v}_i. \quad (5)$$

Premultiplying by $\mathbf{Y}^T(\hat{\mathbf{r}})$ and expanding out the implicit summations gives

$$\int_{S'} \sum_{k=1}^{i_{\text{max}}} Y_k(\hat{\mathbf{r}}')(v_i)_k \sum_{j=1}^{i_{\text{max}}} Y_j(\hat{\mathbf{r}})Y_j(\hat{\mathbf{r}}') d\Omega' = \lambda_i v_i(\hat{\mathbf{r}}), \quad (6)$$

where $v_i(\hat{\mathbf{r}}) = \mathbf{Y}^T(\hat{\mathbf{r}})\mathbf{v}_i$ is the i th eigenfunction of the coupling matrix. The completeness of the spherical harmonics in the $l_{\text{max}} \rightarrow \infty$ limit implies that $\mathbf{Y}^T(\hat{\mathbf{r}})\mathbf{Y}(\hat{\mathbf{r}}') = \delta(\hat{\mathbf{r}} - \hat{\mathbf{r}}')$ [where $\delta(x)$ is the Dirac delta function], so that equation (6) reduces to

$$\int_{S'} v_i(\hat{\mathbf{r}}')\delta(\hat{\mathbf{r}} - \hat{\mathbf{r}}') d\Omega' = \lambda_i v_i(\hat{\mathbf{r}}). \quad (7)$$

For a given i this must be true at all $\hat{\mathbf{r}}$, which implies that either: $v_i(\hat{\mathbf{r}}) = 0$ in the cut region, $S - S'$, in which case $\lambda_i = 1$; or $v_i(\hat{\mathbf{r}}) = 0$ in S' , in which case $\lambda_i = 0$. In other words these eigenfunctions are completely localised in either the cut sphere or the removed region. This bimodality is only strictly true in the $l_{\text{max}} \rightarrow \infty$ limit, but, as shown in Fig. 1, is a good approximation for $l_{\text{max}} \gtrsim 500$.

As the coupling matrix is symmetric, those eigenvectors with different eigenvalues are orthogonal, and those with the same eigenvalues can be made orthogonal by a rotation in the subspace defined by the eigenvalue in question (e.g. Arfken 1985). Thus the eigenfunctions with $\lambda_i = 1$ represent an orthogonal basis set on S' , whereas those with $\lambda_i = 0$ have no support in this region and so cannot be orthogonal (or normalized) on the cut sky. The freedom in choice of basis does not extend to mixing the $\lambda_i = 0$ modes (i.e. those corresponding to the null-space of \mathbf{C}) with the $\lambda_i = 1$ modes (i.e. those in the range of \mathbf{C}), and so the number of supported modes is determined by a combination of the band limit and the cut.

The number of orthonormal basis functions (i.e. the rank of \mathbf{C}) is proportional to the area of the sphere retained, $\Omega_{S'}$. Hence it is possible to define only

$$i'_{\text{max}} \approx \frac{\Omega_{S'}}{\Omega_S} i_{\text{max}} = \frac{\Omega_{S'}}{4\pi} i_{\text{max}} \quad (8)$$

orthonormal functions on the cut sphere for a given (large) band limit. The relative reduction in the basis set is the same as would occur in the equivalent pixel analysis: the number of pixels retained is also given by $N'_p \approx \Omega_{S'}/\Omega_S N_p$. For low l_{max} these arguments do not hold, and it is possible to create a basis set with more than $\Omega_{S'}/\Omega_S i_{\text{max}}$ elements. Moreover, all these functions are required to ensure that the cut-sphere basis set is complete (as well as orthonormal) in the case of a low band limit.

2.2 Construction of an orthonormal basis

The construction of an orthonormal basis set from a set of linearly independent functions is a well-established mathematical technique, and a number of orthogonalization methods are possible. The most basic is Gram–Schmit orthogonalization (e.g. Arfken 1985), in which the new basis functions are built-up sequentially,

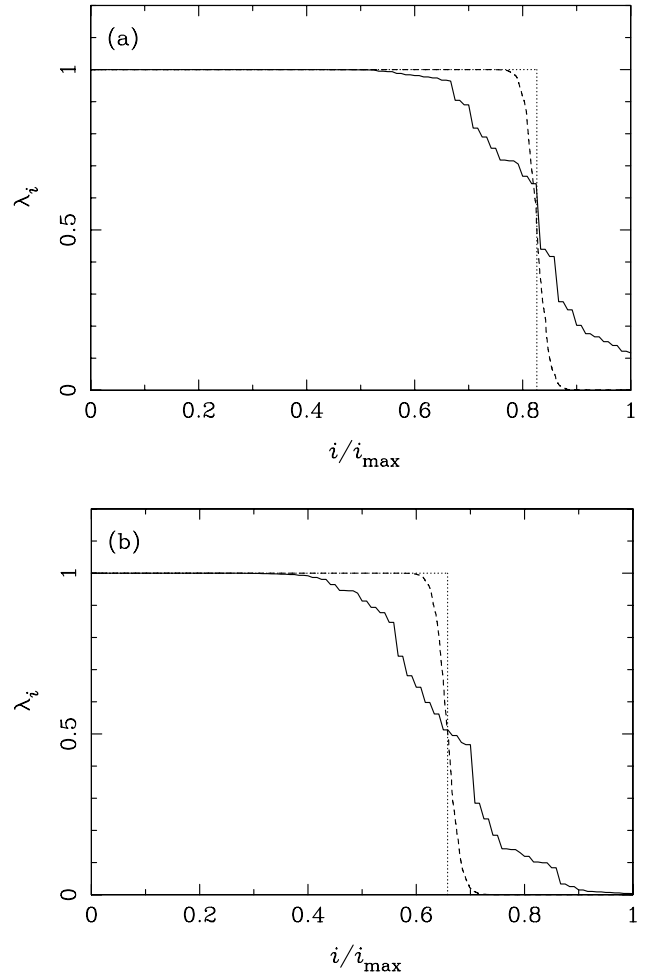


Figure 1. The distribution of eigenvalues of \mathbf{C} , shown for symmetric constant latitude cuts of $b_{\text{cut}} = 10^\circ$ in (a) and $b_{\text{cut}} = 20^\circ$ in (b). In both panels the eigenvalues are sorted into decreasing order, and distributions are shown for $l_{\text{max}} = 10$ (solid line), $l_{\text{max}} = 100$ (dashed line) and $l_{\text{max}} = 1000$ (dotted line). The limiting cases of $l_{\text{max}} \rightarrow \infty$ [in which $\Omega_{S'}/(4\pi)$ of the eigenvalues are unity and the remainder zero] are almost indistinguishable from the $l_{\text{max}} = 1000$ distributions.

but this algorithm is numerically unstable. The modified Gram–Schmidt algorithm (e.g. Golub & van Loan 1996) is stable, but it is generally preferable to use matrix techniques to create all the new basis functions simultaneously.

Starting with the spherical harmonics, $\mathbf{Y}(\hat{\mathbf{r}})$, the task is to find a set of functions $\mathbf{Y}'(\hat{\mathbf{r}})$ which are orthonormal on the incomplete sphere. In terms of a conversion matrix, \mathbf{B} , the two sets of functions are related by

$$\mathbf{Y}'(\hat{\mathbf{r}}) = \mathbf{B}\mathbf{Y}(\hat{\mathbf{r}}). \quad (9)$$

Note that \mathbf{B} has dimensions $i'_{\text{max}} \times i_{\text{max}}$, where $i_{\text{max}} = (l_{\text{max}} + 1)^2$ and $i'_{\text{max}} \leq i_{\text{max}}$ is determined by the band limit and the cut, as described in Section 2.1. It is also important to note that the indexing of the $Y'_i(\hat{\mathbf{r}})$ is qualitatively different from the $Y_i(\hat{\mathbf{r}})$. The latter are really two-index objects, with their characteristic scale given by $\sim \pi/l$ and m relating to ‘orientation’. However the new basis functions include contributions from spherical harmonics with different l -values, and thus do not have a well-defined angular scale. Hence their single index contains no physical information, and the ordering or grouping of the new basis functions is arbitrary.

From equation (1), the coupling matrix of these new functions is

$$\mathbf{C}' = \int_{\mathcal{S}'} \mathbf{Y}'(\hat{r}) \mathbf{Y}'^T(\hat{r}) d\Omega = \mathbf{C} \mathbf{B} \mathbf{C}^T. \quad (10)$$

Hence any conversion matrix which satisfies²

$$\mathbf{C} \mathbf{C} \mathbf{B}^T = \tilde{\mathbf{I}}, \quad (11)$$

yields basis functions which are orthonormal on the cut sphere, and the task of orthogonalization is reduced to finding a solution for \mathbf{B} given \mathbf{C} . Whilst such a solution does not exist for arbitrary \mathbf{C} , in all cases of practical interest a suitable conversion matrix can be constructed from the coupling matrix. One possible method is direct calculation of the eigenstructure of \mathbf{C} , which yields a conversion matrix with elements given by $B_{\hat{r},i} = (\mathbf{v}_{\hat{r}})_i \lambda_{\hat{r}}^{-1/2}$, where the $\mathbf{v}_{\hat{r}}$ are the eigenvectors of \mathbf{C} and the $\lambda_{\hat{r}}$ its positive eigenvalues. However it is advantageous to include the symmetry of the coupling matrix explicitly, which leads to a factorization of the form

$$\mathbf{C} = \mathbf{A} \mathbf{A}^T, \quad (12)$$

where \mathbf{A} is an $i_{\max} \times i'_{\max}$ matrix, the form of which is determined by the decomposition method. Combining equations (11) and (12), the task of orthogonalization is reduced to finding \mathbf{B} such that

$$\mathbf{B} \mathbf{A} = \tilde{\mathbf{O}}, \quad (13)$$

where $\tilde{\mathbf{O}}$ is an $i'_{\max} \times i'_{\max}$ orthogonal matrix (i.e. $\tilde{\mathbf{O}} \tilde{\mathbf{O}}^T = \tilde{\mathbf{O}}^T \tilde{\mathbf{O}} = \tilde{\mathbf{I}}$).

Whilst equations (12) and (13) are general expressions which must be satisfied by the conversion matrix, they do not define a definite algorithm for the orthogonalization. In practice it is simplest to choose $\tilde{\mathbf{O}} = \tilde{\mathbf{I}}$, leading to the requirement that

$$\mathbf{B} \mathbf{A} = \tilde{\mathbf{I}}. \quad (14)$$

However the optimal choice of decomposition method used to generate \mathbf{A} depends on whether the coupling matrix is (numerically) invertible, and hence on the band limit of the analysis.

2.2.1 Low-resolution analysis

If $l_{\max} \leq 50$ and most of the sphere is retained (i.e. $\Omega_{\mathcal{S}'} \approx \Omega_{\mathcal{S}}/2 = 2\pi$), then the coupling matrix is numerically invertible and can be treated as positive definite in practice. Consequently \mathbf{A} (defined in equation 12) is also invertible, and $\mathbf{B} = \mathbf{A}^{-1}$, so that, from equation (9), the orthonormal basis set is given by

$$\mathbf{Y}'(\hat{r}) = \mathbf{A}^{-1} \mathbf{Y}(\hat{r}). \quad (15)$$

The form of \mathbf{A} depends on the factorization method; of the wide variety available (e.g. Golub & van Loan 1996), the two most useful here are singular value decomposition (SVD) and Cholesky decomposition.

The SVD of the covariance matrix is defined in terms of equation (1) by $\mathbf{A} = \mathbf{V} \mathbf{W}^{1/2}$ (i.e. $\mathbf{C} = \mathbf{V} \mathbf{W} \mathbf{V}^T$), where \mathbf{V} is orthogonal and \mathbf{W} is diagonal.³ The diagonal elements of \mathbf{W} are the eigenvalues of \mathbf{C} and, as their ordering is arbitrary, \mathbf{W} can be

² Here $\tilde{\mathbf{I}}$ is the $i'_{\max} \times i'_{\max}$ identity matrix, as distinct from the (potentially larger) $i_{\max} \times i_{\max}$ identity matrix, \mathbf{I} .

³ If \mathbf{M} is diagonal then the notation $\mathbf{M}^{\pm 1/2}$ is used here to denote the matrix defined by $(M^{\pm 1/2})_{ij} = \delta_{ij} M_{ii}^{\pm 1/2}$, where δ_{ij} is the Kronecker delta function. Thus $\mathbf{M}^{1/2}$ only exists if the diagonal elements of \mathbf{M} are non-negative and $\mathbf{M}^{-1/2}$ only exists if the diagonal elements of \mathbf{M} are strictly positive.

defined such that $W_{i,i} \geq W_{i+1,i+1}$, provided the columns of \mathbf{V} are permuted in the same way. The columns of \mathbf{V} , in turn, are the eigenvectors of \mathbf{C} , and \mathbf{V} is an orthogonal matrix (i.e. $\mathbf{V}^{-1} = \mathbf{V}^T$). Hence the conversion matrix is given by $\mathbf{B} = \mathbf{A}^{-1} = \mathbf{W}^{-1/2} \mathbf{V}^T$, which is trivially computed once the SVD has been performed. Note that this approach is effectively the same as the direct calculation of the eigenstructure of \mathbf{C} mentioned above in Section 2.2.

Whilst SVD is a powerful technique, it is computationally expensive – a Cholesky decomposition is approximately 10 times faster, although it can only be performed on symmetric matrices which are numerically positive definite. The Cholesky decomposition of the covariance matrix takes the form $\mathbf{C} = \mathbf{L} \mathbf{L}^T$ (i.e. $\mathbf{A} = \mathbf{L}$), where \mathbf{L} is lower triangular. Hence the conversion matrix $\mathbf{B} = \mathbf{L}^{-1}$ can be computed quickly from the initial factorization in this case as well. The triangular structure of the conversion matrix also ensures that the new basis functions are the same as those formed by a numerically stable Gram–Schmidt orthogonalization (Górski 1994).

Despite the fact that the SVD and the Cholesky decomposition result in quite different sets of basis functions, there is no reason to prefer one over the other, in general. In the case of CMB analysis, however, the triangular structure of \mathbf{A} and \mathbf{B} as generated by the Cholesky decomposition is preferable as it ensures that the non-cosmological monopole and dipole modes are kept separate from the $l \geq 2$ modes, assuming l -ordering is used (Górski 1994). If the SVD route is taken (or another indexing scheme used) the separation of the $l = 0$ and $l = 1$ modes can be ensured using the partial Householder transform described in Appendix C. None the less, if the coupling matrix is sufficiently non-singular, a Cholesky decomposition should be used to create the orthonormal basis set, as a result of both its computational efficiency and the simplicity with which the non-cosmological modes are handled.

2.2.2 High-resolution analysis

If $l_{\max} \gtrsim 50$, the coupling matrix is numerically singular, and thus \mathbf{A} (defined in equation 1) is non-invertible. Cholesky decomposition of \mathbf{C} is thus impractical and, whilst an SVD is possible, the conversion matrix as defined in Section 2.2.1 cannot be computed, as the smallest elements of \mathbf{W} (i.e. the smallest eigenvalues of \mathbf{C}) are so close to zero. This implies that the corresponding columns of \mathbf{V} do not contribute to the reconstruction of \mathbf{C} and can be ignored. Hence it is possible to perform an approximate SVD of the coupling matrix, defined by $\mathbf{C} \approx \tilde{\mathbf{V}} \tilde{\mathbf{W}} \tilde{\mathbf{V}}^T$ (i.e. $\mathbf{A} \approx \tilde{\mathbf{V}} \tilde{\mathbf{W}}^{1/2}$), where $\tilde{\mathbf{W}}$ is an $i'_{\max} \times i'_{\max}$ diagonal matrix containing the largest elements of \mathbf{W} and $\tilde{\mathbf{V}}$ is an $i_{\max} \times i'_{\max}$ matrix consisting of the corresponding columns of \mathbf{V} . The value of i_{\max} is determined by the choice of W_{\min} used to truncate \mathbf{W} , but the bimodality of the eigenvalue distribution means that any value between $\sim 10^{-5}$ and ~ 0.1 is acceptable. The resultant conversion matrix is $\mathbf{B} = \tilde{\mathbf{W}}^{-1/2} \tilde{\mathbf{V}}^T$ (satisfying $\mathbf{B} \mathbf{A} \approx \tilde{\mathbf{I}}$) and the i'_{\max} new basis functions are given by

$$\mathbf{Y}'(\hat{r}) = \tilde{\mathbf{W}}^{-1/2} \tilde{\mathbf{V}}^T \mathbf{Y}(\hat{r}). \quad (16)$$

These basis functions represent an orthonormal basis set on the incomplete sphere, but they are not formally complete to the nominal band limit because of the slightly approximate nature of the reduced SVD. The decomposition becomes exact in the limit $l_{\max} \rightarrow \infty$ as $\tilde{\mathbf{W}} \rightarrow \tilde{\mathbf{I}}$ (from the eigenvalue arguments described in Section 2.1.1) and the reduced SVD becomes $\mathbf{C} \rightarrow \tilde{\mathbf{V}} \tilde{\mathbf{V}}^T$ (i.e. $\mathbf{A} \rightarrow \tilde{\mathbf{V}}$ and $\mathbf{B} \rightarrow \tilde{\mathbf{V}}^T$) in this limit. Note also that these basis functions are




Figure 2. Orthonormal cut-sphere basis functions, $Y_i(\mathbf{r})$, as described in Section 2.2.2. In all cases $l_{\max} = 100$, $W_{\min} = 0.01$, and a Mollweide projection is used. The colour map varies from black (large negative values), through grey (zero), to white (large positive values), in each case being scaled to cover the dynamic range of the relevant basis function. For those in the left column [(a), (c), (e) and (g)] the cut (shown by the dashed lines) is symmetric, with $b_{\text{cut}} = \pm 20^\circ$; for those in the right column [(b), (d), (f) and (h)] the cut (again shown by the dashed lines) is asymmetric, the region between $b_1 = 0^\circ$ and $b_2 = -45^\circ$ having been removed. Within each column the indexing of the basis functions is arbitrary, but they are displayed so that their fractional support in the removed region increases from (a) to (g) and (b) to (h), respectively.

still orthogonal on the full sphere, although they are no longer normalized to unity; this is another potential advantage of the SVD-based method.

Several examples of orthonormal cut-sphere basis functions are shown in Fig. 2, for both symmetric and asymmetric constant latitude cuts. The link between these functions and the spherical harmonics is apparent – they have the same cellular structure, but, for the most part, are combined in such a way that their support is localized to S' . However, the functions shown in Fig. 2(g) and (h)

are from the small subset with intermediate values of $W_{i,i}$ and, as such, have considerable support in the removed region.

In the case of CMB analysis, one shortcoming of this approach is that the non-cosmological monopole and dipole modes are not distinguished from the higher moments (see Section 3.2), although this separation can be achieved post facto by using a partial Householder transform, as described in Appendix C. In doing this, some of the useful properties of the SVD are lost, but this operation need only be performed as the last step in generating the

orthonormal basis set, by which stage all the computationally intensive matrix operations have already been performed.

Despite the inconvenience caused by the mixing of the monopole and dipole modes, SVD is clearly the most flexible and general method of orthogonalization. In part, this stems from the fact that it can be applied to the coupling matrix without any prior knowledge of its singular properties. However, in the high- l_{\max} limit, the number of supported basis functions is determined by the area of the cut and given, to a good approximation, by equation (8). Hence faster, if less powerful, techniques can be used to orthogonalize the spherical harmonics for $l_{\max} \gtrsim 1000$ as the coupling matrix is guaranteed to have at least $(l_{\max} + 1)^2 \Omega_S / (4\pi)$ positive eigenvalues. An example of this idea would be to modify the pivoting algorithm in the psuedo-Cholesky decomposition described in Section 4.2.9 of Golub & van Loan (1996), so that the decomposition halts when the predetermined number of basis functions have been generated, as opposed to using the less robust threshold based on the values of the diagonal elements of \mathbf{C} .

2.3 Constant latitude cuts

In principle the method presented above is a complete solution to the problem of constructing orthonormal bases on the cut sphere, but the coupling matrix requires $O(l_{\max}^4)$ storage, limiting a general implementation to $l_{\max} \approx 200$ on most current computers. Furthermore, the SVD of an $n \times n$ matrix requires $O(n^3)$ operations, and so the orthogonalization operation count scales as $O(l_{\max}^6)$. Similar difficulties are encountered in merely evaluating \mathbf{C} , regardless of whether numerical integration or recursive techniques are used (e.g. Hivon et al. 2001).

Fortunately, all these difficulties are significantly reduced in the case of a constant latitude cut (cf. Oh, Spergel & Hinshaw 1999; Wandelt, Hivon & Górski 2001), defined by ignoring all θ for which $\theta_1 \leq \theta \leq \theta_2$. This could be the symmetric removal of the Galactic plane (i.e. $\theta_1 = \pi/2 - b_{\text{cut}}$ and $\theta_2 = \pi/2 + b_{\text{cut}}$, where b_{cut} is the latitude of the cut) or the absence of data round one pole (i.e. $\theta_1 = 0$ and $\theta_2 = \theta_{\text{cut}}$). The formalism derived below can also be trivially extended to include multiple cuts, as would be required for a CMB experiment which did not observe either ecliptic pole.

Explicitly including the constant latitude cut in equation (1), the elements of the coupling matrix are given by

$$C_{i(l,m),i(l',m')} = \int_0^{2\pi} s_m(\phi) s_{m'}(\phi) d\phi \times \left[\int_{-1}^{\cos(\theta_2)} \lambda_{l,|m|}(x) \lambda_{l',|m'|}(x) dx + \int_{\cos(\theta_1)}^1 \lambda_{l,|m|}(x) \lambda_{l',|m'|}(x) dx \right], \quad (17)$$

where $s_m(\phi)$ is defined in equation (A2), and the $\lambda_{l,m}(x)$ are normalized associated Legendre functions, given in equation (A3). From Appendix A, the first integral in equation (17) reduces to $2\pi \delta_{m,m'}$, and so

$$C_{i(l,m),i(l',m')} = \delta_{m,m'} \left[\delta_{l,l'} - 2\pi \int_{\cos(\theta_2)}^{\cos(\theta_1)} \lambda_{l,|m|}(x) \lambda_{l',|m'|}(x) dx \right]. \quad (18)$$

The remaining integral can be evaluated using a combination of analytical formulæ and recursion relations, as described in Appendix B.

The most important aspect of equation (18) is that the coupling

matrix \mathbf{C} is extremely sparse (only one element in $\sim l_{\max}$ is non-zero) and, if stored using the indexing scheme defined in equation (A12) (i.e. grouped into sub-matrices of fixed m), is block diagonal. \mathbf{C} can thus be stored in the form of $2l_{\max} + 1$ sub-matrices, the m th of which has $(l_{\max} + 1 - |m|)^2$ elements, and the storage requirements thus scale as $O(l_{\max}^2)$ rather than $O(l_{\max}^4)$. Whilst it is convenient to store all the blocks simultaneously, there is no need to do so, which can further reduce the storage requirements to $O(l_{\max}^2)$. It is also clear from equation (18) that only the $m \geq 0$ terms need be treated explicitly and that l and l' are interchangeable, decreasing the storage requirements by a further factor of 4. Finally, in the case of a symmetric cut (i.e. $\theta_2 = \pi - \theta_1$), the parity of $\lambda_{l,m}(x)$ is such that all terms for which $l + l'$ is odd vanish, resulting in an additional halving of the memory requirements.

The orthogonalization can be performed by decomposing each sub-matrix separately, reducing the operation count from $O(l_{\max}^6)$ to $O(l_{\max}^4)$. The removal of the poorly-supported basis functions is achieved in the same manner as described in Section 2.2, although the book-keeping is more complicated. Similarly the partial Householder transform required to separate the $l = 0$ and $l = 1$ modes need only be applied to the $m = 0$ and $m = \pm 1$ blocks of the resultant conversion matrix (Appendix C). An important side-effect of the separation in m is that the $Y'_i(\hat{r})$ have the same trigonometric ϕ -dependence as the full-sky spherical harmonics (Appendix A). This also implies that the $Y'_i(\hat{r})$ can be treated as two-index quantities, defined by m and a second, arbitrary index in place of l .

The algorithms described here were implemented on the Cambridge Centre for Mathematical Science's COSMOS 64-processor Silicon Graphics Origin 2000 and the evaluation and decomposition of the coupling matrix at the highest *Planck* resolution of $l_{\max} \approx 2500$ required about an hour. The majority of the time was spent factorizing the sub-matrices, and thus significant accelerations are unlikely, the highly optimized LINEAR ALGEBRA PACKAGE (LAPACK; Anderson et al. 1992) routines having been used for all the decompositions. For a given choice of θ_1 and θ_2 , the decomposition of \mathbf{C} need only be performed once, so orthogonalization of the spherical harmonics on an incomplete sky should comprise only a small fraction of the analysis required for the forthcoming *MAP* and *Planck* missions.

3 HARMONIC ANALYSIS

Methods for constructing an orthonormal basis set on the incomplete sphere from the spherical harmonics were discussed in Section 2, but in most cases of data analysis it is the harmonic coefficients, representing functions on the sphere, that are of interest. There are at least three useful harmonic expansions of a general function on the sphere, and the relationships between these coefficients, which are summarized in Table 1, are derived here.

A band-limited function, $a(\hat{r})$, can be completely specified by a finite number of harmonic coefficients as (cf. Appendix A)

$$a(\hat{r}) = \mathbf{Y}^T(\hat{r}) \mathbf{a}, \quad (19)$$

where it is assumed that l_{\max} is greater than or equal to the band limit of $a(\hat{r})$ and the harmonic coefficients are defined by

$$\mathbf{a} = \int_S \mathbf{Y}(\hat{r}) a(\hat{r}) d\Omega. \quad (20)$$

The invertibility of these transformations is due to the

Table 1. The conversions between the various harmonic coefficients defined in Section 3: \mathbf{a} are the standard coefficients of the spherical harmonics (equation 20); $\tilde{\mathbf{a}}$ are the pseudo-harmonic coefficients (equation 21); \mathbf{a}' are the cut-sphere harmonic coefficients (equation 23); and $\hat{\mathbf{a}}$ are the reconstructed spherical harmonic coefficients. The two basis functions are the spherical harmonics, $\mathbf{Y}(\hat{\mathbf{r}})$ (defined in Appendix A), and the orthogonalized harmonics, $\mathbf{Y}'(\hat{\mathbf{r}})$ (defined in Section 2). If the coupling matrix of the spherical harmonics, \mathbf{C} , is invertible, then the expressions for $\hat{\mathbf{a}}$ following the \rightarrow can be used as exact inversions; otherwise the ‘estimators’ are approximate projections onto the cut region. \mathbf{A} can be any matrix such that $\mathbf{A}\mathbf{A}^T = \mathbf{C}$, and \mathbf{B} can be any matrix such that $\mathbf{B}\mathbf{A} = \mathbf{I}$.

$$\begin{array}{lll}
 \mathbf{a} = \int_S \mathbf{Y}(\hat{\mathbf{r}})a(\hat{\mathbf{r}}) d\Omega & \hat{\mathbf{a}} = \mathbf{B}^T\mathbf{B}\tilde{\mathbf{a}} \rightarrow \mathbf{C}^{-1}\tilde{\mathbf{a}} & \hat{\mathbf{a}} = \mathbf{B}^T\mathbf{a}' \rightarrow (\mathbf{A}^T)^{-1}\mathbf{a}' \\
 \tilde{\mathbf{a}} = \int_{S'} \mathbf{Y}(\hat{\mathbf{r}})a(\hat{\mathbf{r}}) d\Omega & \tilde{\mathbf{a}} = \mathbf{C}\mathbf{a} & \tilde{\mathbf{a}} = \mathbf{A}\mathbf{a}' \\
 \mathbf{a}' = \int_{S'} \mathbf{Y}'(\hat{\mathbf{r}})a(\hat{\mathbf{r}}) d\Omega & \mathbf{a}' = \mathbf{A}^T\mathbf{a} & \mathbf{a}' = \mathbf{B}\tilde{\mathbf{a}}
 \end{array}$$

orthonormality of the spherical harmonics on S and the fact that they represent a complete basis set given the band limit.

If $a(\hat{\mathbf{r}})$ is only known over some fraction of the sphere $S' \leq S$, then \mathbf{a} cannot be determined as above, as the integral in equation (20) is incomplete. In this case, the pseudo-harmonics

$$\tilde{\mathbf{a}} = \int_{S'} \mathbf{Y}(\hat{\mathbf{r}})a(\hat{\mathbf{r}}) d\Omega \quad (21)$$

fully specify $a(\hat{\mathbf{r}})$ in S' due to the band limit. From equation (19) they are related to the full harmonic coefficients by

$$\tilde{\mathbf{a}} = \mathbf{C}\mathbf{a}, \quad (22)$$

where \mathbf{C} is the coupling matrix, defined in equation (1).

The pseudo-harmonics are useful quantities, but it is preferable to work with basis functions that are orthonormal on S' . Denoted $\mathbf{Y}'(\hat{\mathbf{r}})$ in Section 2, their harmonic coefficients are given by

$$\mathbf{a}' = \int_{S'} \mathbf{Y}'(\hat{\mathbf{r}})a(\hat{\mathbf{r}}) d\Omega. \quad (23)$$

The relationship between $\mathbf{Y}(\hat{\mathbf{r}})$ and $\mathbf{Y}'(\hat{\mathbf{r}})$ given in equation (9) flows through to the harmonic coefficients, and applying equations (9), (12) and (19) to equation (23) gives

$$\begin{aligned}
 \mathbf{a}' &= \mathbf{B} \int_{S'} \mathbf{Y}(\hat{\mathbf{r}})[\mathbf{Y}'^T(\hat{\mathbf{r}})\mathbf{a}]d\Omega \\
 &= \mathbf{B}\mathbf{C}\mathbf{a} \\
 &= \mathbf{A}^T\mathbf{a}. \quad (24)
 \end{aligned}$$

The form of the above transformation depends on the decomposition used to generate \mathbf{A} (cf. Section 2.2), but, for CMB analysis, it is desirable to separate the non-cosmological modes. This amounts to demanding that the only four of the \mathbf{a}'_l have any contribution from the $l=0$ and $l=1$ spherical harmonic coefficients. The upper triangular structure of $\mathbf{A}^T = \mathbf{L}^T$ as generated by a Cholesky decomposition inherently satisfies this requirement, but in general the conversion matrix needs to be transformed explicitly. One option is to use successive partial Householder transforms, as described in detail in Appendix C. As can be seen from Fig. 3, the index-ordering and decomposition method combine to give a wide variety of conversion matrices; which of these is most suitable depends on the application.

It is also possible to convert between the pseudo-harmonics and the cut-sky harmonics, as they both contain information about $a(\hat{\mathbf{r}})$ in S' alone. Combining equations (12), (21) and (23) implies that $\tilde{\mathbf{a}} = \mathbf{A}\mathbf{a}'$ and $\mathbf{a}' = \mathbf{B}\tilde{\mathbf{a}}$. However it is not always possible to

determine \mathbf{a} from either \mathbf{a}' or $\tilde{\mathbf{a}}$. In the low- l_{\max} limit these inversions are defined (Section 3.1), but for appreciable band limits only a projection onto the cut sphere is possible (Section 3.2).

3.1 Low-resolution analysis

If the coupling matrix is numerically non-singular (i.e. $l_{\max} \lesssim 50$) then equation (24) can be inverted to give

$$\mathbf{a} = (\mathbf{A}^T)^{-1}\mathbf{a}' = \mathbf{B}^T\mathbf{a}', \quad (25)$$

and equation (21) implies that

$$\mathbf{a} = \mathbf{C}^{-1}\tilde{\mathbf{a}}. \quad (26)$$

These are specific examples of the fact that a band-limited function is completely defined if it is known over any finite portion of the sphere, and a cut-sky analysis serves no purpose – any apparently localized contaminants infect the entire sky. However, any measurement of a field on the sphere is subject to noise which is not band-limited, in which case the application of a cut has the effect of greatly amplifying the noise in the removed region (see Section 4.4), justifying the use of a cut-sky analysis in the low-resolution limit.

3.2 High-resolution analysis

If $l_{\max} \gtrsim 50$ and the coupling matrix is numerically singular, it is impossible to reconstruct (even) a band-limited function that is known only on S' . The loss of information about modes constrained to the cut makes it clear that the analysis has the desired effect of removing contaminated (or otherwise problematic) regions, but the most appropriate transformation from the cut-sphere basis to conventional harmonics is less obvious.

A least-squares-like approach leads to a definition of the reconstructed full-sky coefficients as

$$\hat{\mathbf{a}} \approx \mathbf{B}^T\mathbf{a}' = (\mathbf{A}\mathbf{B})^T\mathbf{a}. \quad (27)$$

Similarly, equation (22) implies that

$$\hat{\mathbf{a}} \approx \mathbf{B}^T\mathbf{B}\tilde{\mathbf{a}} = (\mathbf{A}\mathbf{B})^T\tilde{\mathbf{a}}, \quad (28)$$

where $(\mathbf{A}\mathbf{B})^T$ is a projection operator⁴ onto the range of \mathbf{C} , which in real space is S' . If $l_{\max} \rightarrow \infty$, it is possible to write $\hat{a}(\hat{\mathbf{r}}) = w_{S'}(\hat{\mathbf{r}})a(\hat{\mathbf{r}})$, where $w_{S'}(\hat{\mathbf{r}})$ is the sharp window function defined in equation (3).

⁴The definition of \mathbf{B} given in Section 2.2 implies that $[(\mathbf{A}\mathbf{B})^T]^2 = (\mathbf{A}\mathbf{B})^T$ and it is hence a projection operator if $l_{\max} \rightarrow \infty$.

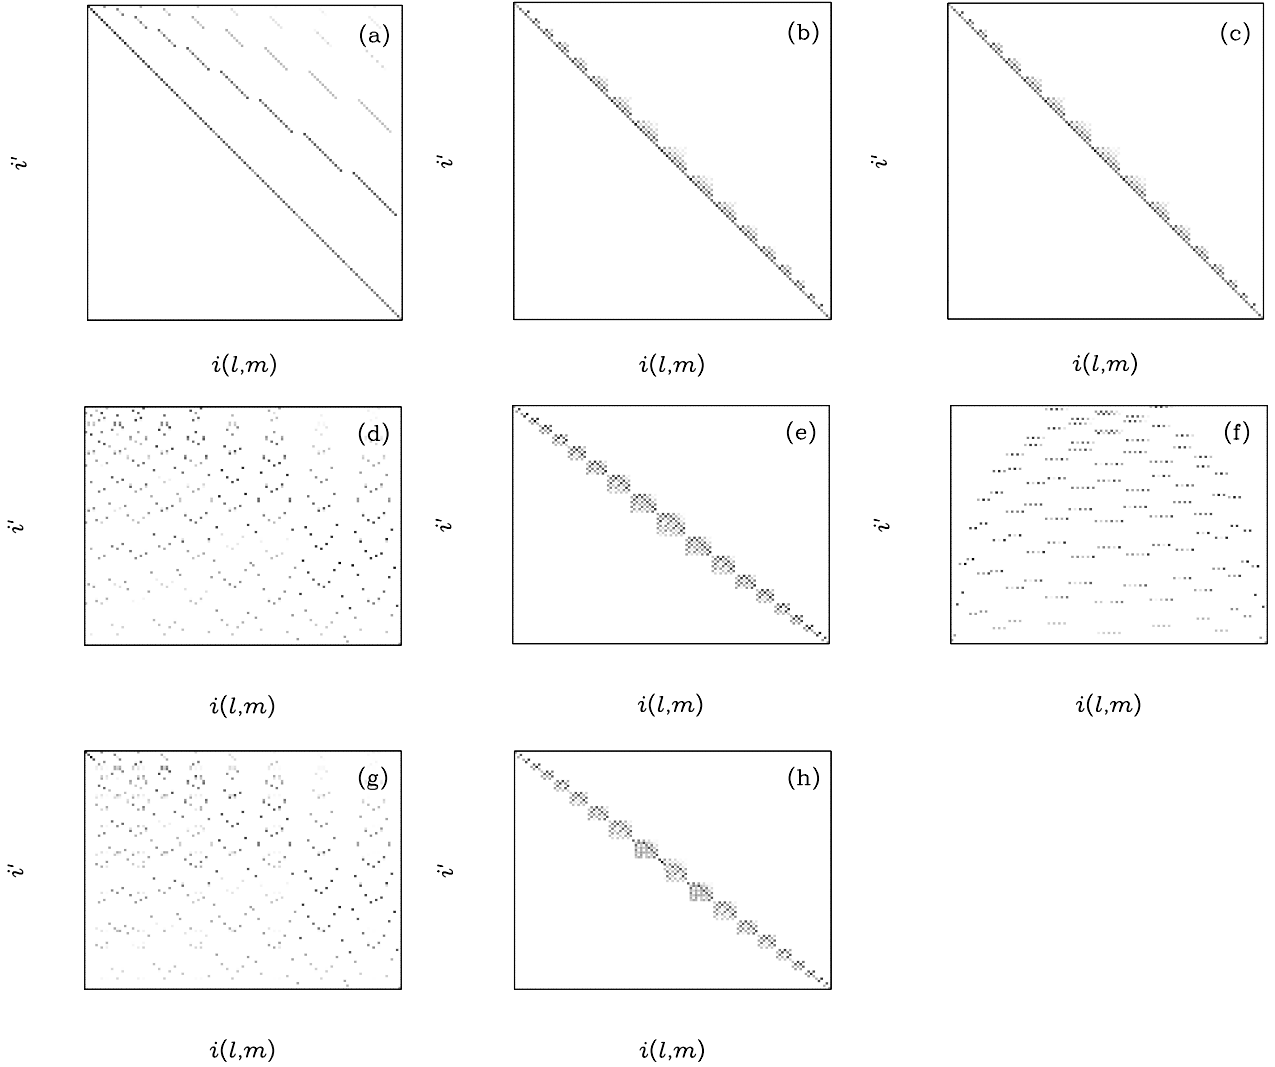


Figure 3. Several examples of the conversion matrices, \mathbf{A}^T , that relate the orthonormal cut sphere harmonic coefficients, a' , to the conventional spherical harmonic coefficients, a , according to equation (24). In all cases a symmetric, constant latitude of $b_{\text{cut}} = \pm 20^\circ$ has been applied and $l_{\text{max}} = 10$. The row index, i' , corresponds to the cut-sky harmonic coefficients and the column index, $i(l,m)$, corresponds to the original spherical harmonic coefficients. The colour map shows the absolute value of the elements of \mathbf{A}^T varying from zero (white) to their maximum value (black), which is normalised separately in each case. The spherical harmonic coefficients are indexed using l -ordering in the left column [(a), (d) and (g)]; m -ordering has been used in order to utilize the decoupling of different m -modes resulting from the azimuthal symmetry by using block-by-block storage in the central column [(b), (e) and (h)]; and m -ordering is used, but the structure of the coupling matrix is not utilized, in the right column [(c) and (f)]. In the case of l -ordering, the non-cosmological modes are the four left-most columns of the conversion matrix; in the case of m -ordering the non-cosmological modes are the left-most columns of the three central blocks. The conversion matrices shown in the top row [(a), (b) and (c)] result from a Cholesky decomposition; those in the middle row [(d), (e) and (f)] are produced by an SVD (in which all modes with $W_{i,i} \leq 0.1$ have been removed); finally, in the bottom row [(g) and (h)], the partial Householder transform described in Appendix C is applied to the conversion matrices shown in the above panels [(d) and (e), respectively].

It is at this point that the subtle distinctions between the use of a discrete cut and a window function become apparent. These results only hold for band-limited functions, but, as defined above, $\hat{a}(\hat{r})$ is not band-limited, and so cannot be analysed self-consistently. Whether a discrete cut or an apodizing function is to be preferred depends on the situation in which the incomplete sky analysis is required.

4 CMB DATA ANALYSIS

In order to determine the properties of the CMB from noisy observations of the microwave sky (Section 4.1) a number of non-trivial analysis steps are required, including: map-making

(Section 4.2); component separation (Section 4.3) and power spectrum estimation (Section 4.4). Several algorithms have been suggested for all these steps, and these are discussed briefly below, but the main focus is on when and how to apply a sky cut. Further, whilst it is possible to analyse the data in either real or Fourier space, the latter approach is emphasized here as it is more directly related to the formalism described in Sections 2 and 3, as well as being the focus of a related series of papers (van Leeuwen et al. 2001; Challinor et al. 2001; Stolyarov et al. 2001). Note that the term ‘map’ is used here to denote any representation of a field on the sky and can imply either a set of real-space pixel values or a vector of spherical harmonic coefficients.

4.1 Observations

Observations of the CMB can be made using a number of quite distinct techniques. Data have been obtained from the ground, high altitude balloons, and satellites, but the more important distinction is the type of telescope. The experiments listed in Section 1 include: straightforward single-dish telescopes, such as BOOMERanG (e.g. Netterfield et al. 2001) and *Planck* (e.g. Bersanelli et al. 1996); differencing experiments, such as *COBE* (e.g. Smoot et al. 1992) and *MAP* (e.g. Jarosik et al. 1998); and interferometers, such as the Cambridge Anisotropy Telescope (CAT; Scott et al. 1996), the Cosmic Background Imager (CBI; Padin et al. 2001) and the Degree Angular Scale Interferometer (DASI; Halverson et al. 2001). The interferometry surveys inevitably cover only a small fraction of the sky, and so a flat-sky Fourier analysis becomes possible. However the both the differencing and single-dish surveys can, in principle, cover most of the celestial sphere, and should yield maps of the microwave sky that are limited only by (the combined effects of) instrument noise and the finite telescope beam.

4.1.1 Noise

The typical receivers used in the above experiments have two main noise contributions: random white noise; and a correlated low-frequency (i.e. ‘ $1/f$ ’) component. The latter is potentially troublesome, leading to ‘stripey’ maps with correlated errors, and is the main reason for the popularity of differencing experiments which remove low frequency noise at the moment of observation. However, data from single-dish surveys can be ‘de-stripped’ if the scan strategy includes sufficiently many multiply observed points (e.g. Tegmark 1997; Delabrouille 1998; Maino et al. 1999) or the time–time noise covariance matrix can be fully included in the map-making process (Wright, Hinshaw & Bennett 1996; Natoli et al. 2001; Challinor et al. 2001). Hence correlated errors are ignored in the simple analysis presented here.

This leaves only the white component, which can be analysed most simply in the case of a single beam experiment. Following Knox (1995), a receiver is characterized by its sensitivity, s (generally chosen to have units of temperature time^{1/2}). Assuming the noise is Gaussian it has expectation values $\langle n \rangle = 0$ and $\langle n^2 \rangle = s^2 t$ over an integration time t . The manner in which this noise projects on to a sky map depends on the map-making algorithm, the scan strategy, and the beam.

4.1.2 Beam convolution

All telescopes necessarily have a finite point spread function or beam, which, for a given detector, can be characterized by $b(\hat{r})$, the fraction of photons from direction \hat{r} that are registered, given a nominal orientation towards the north pole (i.e. $\theta = 0$). The harmonic expansion of the beam in this orientation is denoted $\mathbf{b} = b_{i(l,m)}$, with the band limit being related to the nominal resolution of the detector. For a given type of telescope the resolution improves with frequency as a result of diffraction effects; this places limitations on the component-separation algorithms that are used on the incomplete sky (Section 4.3).

Most experiments have beams that are manifestly asymmetric, a fact which must be accounted for explicitly by the data analysis algorithms, but the cut-sky issues of interest here can be explored more clearly if the beam is approximated by its azimuthally

averaged counterpart (e.g. Challinor et al. 2001). Defined by

$$\bar{b}(\theta) = \frac{1}{2\pi} \int_0^{2\pi} b(\theta, \phi) d\phi, \quad (29)$$

its harmonic coefficients are simply

$$\bar{b}_{l,m} = \delta_{m,0} b_{l,0} = \delta_{m,0} \sqrt{(2l+1)\pi} \int_{-1}^1 \bar{b}[\arccos(x)] P_l(x) dx, \quad (30)$$

where $P_l(x)$ is a Legendre polynomial (Appendix A). The use of $\bar{b}(\hat{r})$ allows the definition of a beam-smoothed sky, $\bar{s}(\hat{r})$, given in terms of the true sky, $s(\hat{r})$, by

$$\bar{s}(\hat{r}) = \int_{S'} \bar{b}[\arccos(\hat{r} \cdot \hat{r}')] s(\hat{r}') d\Omega'. \quad (31)$$

This convolution is much simpler in harmonic space, and applying equation (A7) to equation (31) yields

$$\bar{\mathbf{s}} = \bar{\mathbf{B}} \mathbf{s}, \quad (32)$$

where $\bar{\mathbf{B}}$ (as distinct from the conversion matrix, \mathbf{B}) is a diagonal ‘convolution matrix’ with

$$\bar{B}_{i,i} = \sqrt{\frac{4\pi}{2l(i)+1}} \bar{b}_{l(i),0} = 2\pi \int_{-1}^1 \bar{b}[\arccos(x)] P_{l(i)}(x) dx, \quad (33)$$

$l(i)$ being defined in Appendix A. The simple form of equation (32) is often utilized explicitly in CMB analysis algorithms (e.g. Knox 1995; Hobson et al. 1998; Oh et al. 1999; Stolyarov 2001), but in all cases full sky coverage is – sometimes implicitly – assumed.

Turning to convolution on the incomplete sphere, S' , application of equation (24) to equation (32) yields

$$\bar{\mathbf{s}}' = \mathbf{A}^T \bar{\mathbf{B}} \mathbf{s}, \quad (34)$$

where \mathbf{A} is defined implicitly in equation (12). In the low-resolution limit, equation (25) then gives the cut-sky analogue of equation (32) as

$$\bar{\mathbf{s}}' = \bar{\mathbf{B}}' \mathbf{s}', \quad (35)$$

with the new convolution matrix defined by $\bar{\mathbf{B}}' = \mathbf{A}^T \bar{\mathbf{B}} \mathbf{A}$.

For higher band limits, no such relation exists, the loss of modes in the cut region rendering the convolution ill-defined. This is simply understood in real space, as the value of $\bar{s}(\hat{r})$ near the edge of S' is given by an integral that extends several beamwidths into the removed region. Thus it is impossible to relate $\bar{s}(\hat{r})$ to $s(\hat{r})$ with $\hat{r} \in S'$. These arguments are true independent of the representation chosen, but in harmonic space they mean that it is impossible to relate $\bar{\mathbf{s}}'$ to \mathbf{s}' .

Whilst equation (35) is formally incorrect in high- l_{\max} cases, it is potentially useful as a practical approximation. It is equivalent to assuming that the signal is given by $\mathbf{Y}'(\hat{r}) \mathbf{s}'$, which implies that $s(\hat{r}) \approx 0$ in $S - S'$. This is particularly inaccurate if the removed region contains anomalously strong sources, such as the Galactic plane. None the less, equation (35) gives $s'(\hat{r})$ correctly for all \hat{r} more than a few beamwidths away from the edge of the cut region. However, even if this is an acceptable approximation, there is the further inconvenience that the effective cut-sky beam, $\bar{\mathbf{B}}'$, is not diagonal, introducing couplings between all the modes.

In short, it is preferable to avoid performing any sort of convolution (or deconvolution) on the cut sky, although it is clear that this situation is encountered in any survey with incomplete sky coverage. The one, albeit trivial, exception to this rule is if the beam is a delta function, or at least the closest approximation to a

delta function possible given the band limit under consideration. In this case, $\bar{\mathbf{B}} = \mathbf{I}$, and hence, from equation (14), $\bar{\mathbf{B}}' = \bar{\mathbf{I}}$ as well. equation (35) then implies that $\hat{s}' (= s')$ is the true, unconvolved sky map, estimation of which is addressed next.

4.2 Map making

Some of the most important products of the next generation of CMB survey will be high-resolution maps of the sky at each of several frequencies. Such maps can be created in a number of ways, but care must be taken to account for a huge variety of systematics whilst retaining as much information as possible. Both real space (e.g. Wright, Hinshaw & Bennett 1996; Bond, Jaffe & Knox 1998; Natoli et al. 2001) and Fourier space (van Leeuwen et al. 2001; Challinor et al. 2001) algorithms have been proposed as being suited to particular aspects of the map-making problem. The resultant uncertainties in the pixel values or harmonic coefficients depend on both the data itself (i.e. the scan strategy, noise properties, etc.) and the map-making algorithms used, and can vary quite markedly from experiment to experiment.

Here only the idealized case of uniform sky coverage is considered as the discussion which follows is not significantly changed by this useful simplification. Under this assumption the optimal estimator for the unsmoothed sky, $\hat{s} (= \mathbf{s} + \mathbf{n})$, would be unbiased (i.e. $\langle \hat{s} \rangle = \mathbf{s}$) and have covariance given by

$$\mathbf{N} = \langle \mathbf{nn}^T \rangle = \langle (\hat{s} - \mathbf{s})(\hat{s} - \mathbf{s})^T \rangle = \sigma^2 \bar{\mathbf{B}}^{-2}, \quad (36)$$

where (cf. Knox 1995)

$$\sigma^2 = 4\pi s^2 / (N_d t_{\text{obs}}), \quad (37)$$

t_{obs} is the total observation time of the survey and N_d is the number of detectors at the frequency in question (all of which are assumed to have the same beam). An important issue at this point is the band limit chosen. Clearly only a finite analysis is possible in practice, and $\bar{\mathbf{B}}$ becomes increasingly singular as $l_{\text{max}} \rightarrow \infty$; these two points are related in so far as the sky can never be reconstructed with infinite resolution. The choice of l_{max} is somewhat arbitrary, although any value a factor of a few greater than the effective beamwidth will ensure that $\bar{\mathbf{B}}^{-1}$ exists whilst discarding only multipoles that are noise-dominated. The fact that, unlike the useful signal, the noise is not subject to any band limit is critical to the understanding of the low- l_{max} cut-sky power spectrum estimation discussed in Section 4.4.

Note also that, as a result of the assumption of uniform sky coverage, the covariance matrix is diagonal. Transforming this estimator into real space yields maps with covariance given by

$$\langle n(\hat{r}_1) n(\hat{r}_2) \rangle = \sigma^2 \mathbf{Y}^T(\hat{r}_1) \bar{\mathbf{B}}^{-2} \mathbf{Y}(\hat{r}_2). \quad (38)$$

As the noise term in the data is not beam-convolved the removal of the beam results in spatial correlations of the noise (as encoded in $\bar{\mathbf{B}}$), as well as correlations due to the finite resolution analysis (the sums over spherical harmonics), which are essentially equivalent to pixel smoothing. In the more realistic case of non-uniform sky coverage, the covariance matrix is non-diagonal in both bases, a point discussed further by Oh et al. (1999).

The above estimator for the true sky is closely linked to the more commonly used estimator for the smoothed sky, $\hat{\hat{s}}$. Being related by $\hat{\hat{s}} = \bar{\mathbf{B}}^{-1} \hat{s}$, it is clear they contain the same information (under the assumption of a symmetric beam). The covariance structure of $\hat{\hat{s}}$ is simpler as the correlations discussed above are not introduced, but $\hat{\hat{s}}$ is a more natural data object in the context of this discussion as it

is the true sky that is of interest. In particular, unsmoothed maps allow more flexibility in applying a sky cut, as the problems with convolution on the incomplete sphere described in Section 4.1.2 do not arise. In practice the best compromise may be to reconstruct the sky convolved with the azimuthally averaged beam, thus creating maps with the simplest covariance structure possible without information loss. This can be done in either real space (e.g. Bond et al. 1998) or harmonic space (Challinor et al. 2001), although the real-space pointing matrix is more complicated if beam asymmetry information is included.

In summary, if a survey covers the entire celestial sphere it is preferable to use full-sky frequency maps. However it is possible that small parts of the sky will be missed due to either the scan strategy (cf. Maino et al. 1999) or hardware problems during the survey itself. If this is the case the best unsmoothed map that could be constructed would be larger than the actual observed region, but the errors around the boundary of this area would be very high. An inferential approach is possible, but significant difficulties are encountered, especially in Fourier space (Challinor et al. 2001). Fortunately, it is probable that both *MAP* and *Planck* will produce full sky maps at several frequencies, which can then be used to construct maps of the various astrophysical components.

4.3 Component separation

The microwave sky consists of several distinct astrophysical components, as listed in Section 1. Fortunately they have sufficiently distinct spectra that they can be separated using multi-frequency data. Given that *MAP* and *Planck* will produce maps in five and 10 bands, respectively, it should be possible to produce maps of the various components (particularly the CMB) that are relatively free of contamination. As with map making, a number of algorithms have been put forward for this stage of the data processing, although the main focus has been on Fourier space methods (e.g. Tegmark & Efstathiou 1996; Hobson et al. 1998; Bouchet & Gispert 1999; Prunet et al. 2001; Stolyarov et al. 2001). Aside from the expected statistical isotropy of the CMB signal, one of the reasons for this emphasis has been the simplicity of beam convolution in harmonic space (Section 4.1.2). This is critical if smoothed frequency maps are used as the effective smoothing scale will vary with frequency if the telescope is (close to) diffraction limited. However if unsmoothed maps are used real space component separation methods (e.g. Baccigalupi et al. 2000) must also come into consideration, the optimal choice of basis being less clear.

One common aspect of all the separation techniques cited above is that much of the (prior) information about both signal and noise correlations is disregarded in order to render the problem computationally feasible. In real space the correlations between nearby pixels are ignored, and in Fourier space it is the mode-mode couplings that are neglected. Surprisingly, these approximations appear to be unimportant in practice – even the Galactic components have been recovered with striking accuracy. The most relevant result to this discussion is the all-sky component separation to $l_{\text{max}} \approx 2500$ presented by Stolyarov et al. (2001), as it provides clear evidence that whatever correlations are present in the full-sky harmonic basis are unimportant – there are some errors close to the Galactic centre, but they are localized, and there is no sign of this affecting the reconstruction globally.

If the Galactic plane is removed prior to component separation this one troublesome region is no longer present in the analysis, but new problems arise. First, smoothed maps (with frequency-dependent beamwidths) cannot be used as input data without

inducing errors around the edges of S' because of the ill-defined nature of convolution on the cut sky (Section 4.1.2). Even if such errors are deemed acceptable (e.g. Prunet et al. 2001) or beam-deconvolved maps are used, the transformation described in equation (24) completely changes the correlation structure of the harmonics. In particular, the signal–signal correlation matrices are non-diagonal for all components, including random fields such as the CMB (see Section 4.4). Whereas the couplings between the spherical harmonic coefficients can apparently be disregarded, this has not been demonstrated for these induced correlations in the orthonormal basis. Prunet et al. (2001) performed cut-sky component separation which included them in full, but were thus limited to $l_{\max} \approx 500$, the computational task being made considerably larger.

In real space the application of a cut is trivial, provided that beam-deconvolved maps are used, as it simply requires that pixels in the removed region be ignored. Thus the Baccigalupi et al. (2000) method should be well suited to a cut-sky analysis.

Given that realistic component separation simulations have only recently become available, it is likely that important developments in this field will be made in the near future. For the moment, however, it appears that separation can usefully be performed on either the full or cut sky without introducing catastrophic errors. Provided the models of microwave emissions from the Galactic plane used in the above simulations are sufficiently realistic, it may thus be preferable to generate the full-sky maps of the various astrophysical components, retaining the option of masking unwanted regions at a later stage.

4.4 Power spectrum estimation

If the fluctuations in the early Universe were the result of inflation (e.g. Linde 1990) then the CMB is expected to be a Gaussian random field, the statistical properties of which can be specified completely by its angular power spectrum, C_l . Even if this is not the case, the power spectrum should encode much of the cosmological information present. It is thus unsurprising that, as with map making and component separation (Sections 4.2 and 4.3, respectively), many different methods of power spectrum estimation have been developed (e.g. Tegmark 1997; Górski 1994; Bond et al. 1998; Oh et al. 1999; Szapudi et al. 2001; Wandelt et al. 2001; Hivon et al. 2001). Further, sky cuts have been incorporated into many of these algorithms as it seems certain that the strength of the Galactic microwave emissions will prevent the CMB from ever being accurately measured in this region. As a result of the proliferation of papers on this subject, this discussion of power spectrum estimation is limited to a description of a maximum likelihood formalism using the orthonormal basis functions described in Section 2, with reference to how their behaviour differs in the low- and high-resolution regimes.

4.4.1 Maximum likelihood formalism

The most powerful method of power spectrum estimation is maximum likelihood (e.g. Press et al. 1992), although this has only been implemented to *MAP* resolution to date (Oh et al. 1999). By invoking Gaussian statistics for both the CMB and the noise, it is possible to write down the exact likelihood for the observed map (e.g. Górski 1994; Borrill 1999). On the full sky the effective data vector is $\hat{s} = s + n$ (i.e. the estimator for the true sky, of the form discussed in Section 4.2 and not the quantity being estimated here) with the non-cosmological $l = 0$ and $l = 1$ modes removed. The

full likelihood is given by

$$\frac{dp}{d\hat{s}} = \frac{1}{\sqrt{(2\pi)^{l_{\max}^2-4} \det(\mathbf{S} + \mathbf{N})}} \exp\left[-\frac{1}{2}\hat{s}^T(\mathbf{S} + \mathbf{N})^{-1}\hat{s}\right], \quad (39)$$

where \mathbf{S} and \mathbf{N} are the signal and noise covariance matrices, respectively. The assumption of Gaussianity implies that $S_{i,i'} = \delta_{i,i'} C_{l(i)}$, where C_l is the CMB power spectrum and $l(i)$ is defined in Appendix A. The form of \mathbf{N} is determined by a combination of the survey method and the data processing up to this point, but is unlikely to have the simple form of equation (36) as a result of the imperfect component separation. The maximum likelihood calculation consists of finding an estimator for the underlying power spectrum, \hat{C}_l , such that equation (39) is maximized, and there are a number of algorithms for finding this quantity (e.g. Bond et al. 1998; Oh et al. 1999).

The maximum likelihood formalism on the cut sky takes the same form as on the full sky, but with the data vector $\hat{s}' = \mathbf{A}^T \hat{s}$ and the covariance matrices suitably transformed to give (cf. Górski 1994)

$$\frac{dp}{d\hat{s}'} = \frac{1}{\sqrt{(2\pi)^{l_{\max}^2-4} \det(\mathbf{S}' + \mathbf{N}')}} \times \exp\left[-\frac{1}{2}\hat{s}'^T(\mathbf{S}' + \mathbf{N}')^{-1}\hat{s}'\right], \quad (40)$$

with the four modes containing information on the monopole and dipole (Section 2.2) again excluded. The signal covariance matrix is subject to a simple similarity transform, $\mathbf{S}' = \mathbf{A}^T \mathbf{S} \mathbf{A}$, but the same is not true for the noise covariance matrix as the noise field is not band-limited (a fact critical to the use of a cut-sky analysis at low resolution, as discussed below). The coupling of the cut-sky modes makes the maximization of equation (40) non-trivial (cf. Oh et al. 1999), even if $\mathbf{S} + \mathbf{N}$ is diagonal on the full sky. None the less, it is useful to work under this idealized assumption in order to see how the application of the cut ensures that the maximum likelihood solution is independent of the Galactic signal; the manner in which this is achieved is quite different in the low- and high-resolution cases.

4.4.2 Low-resolution analysis

The effect of a sky cut on power spectrum estimation is not entirely obvious in the low- l_{\max} case in which the coupling matrix (Section 2.1) is invertible. The effective band limit, produced by the combined effects of the beam and noise (Section 4.2), means the signal over the whole sky (including e.g. the Galactic plane) is encoded in the cut-sky coefficients. Thus the application of a cut would be redundant were it not for the presence of non-band-limited noise which cannot be characterized properly a finite harmonic analysis.

Applying the incomplete spherical transform defined in equation (23) to a purely white noise field $n(\hat{r})$ [i.e. $\langle n(\hat{r}) \rangle = 0$ and $\langle n(\hat{r}_1)n(\hat{r}_2) \rangle = \delta(\hat{r}_1 - \hat{r}_2)\sigma^2$; cf. equation (36)] gives cut-sky harmonic coefficients with $\langle n' \rangle = \mathbf{0}$ and

$$\mathbf{N}' = \langle n' n'^T \rangle = \sigma^2 \mathbf{I}. \quad (41)$$

Projecting back into real space gives a field $n'(\hat{r})$ which satisfies $\langle n'(\hat{r}) \rangle = 0$ and

$$\langle n'(\hat{r}_1)n'(\hat{r}_2) \rangle = \sigma^2 \mathbf{Y}^T(\hat{r}_1) \mathbf{C}^{-1} \mathbf{Y}(\hat{r}_2). \quad (42)$$

On the full sky the same procedure (i.e. a finite spherical harmonic

analysis followed by a transformation back into real space) would yield a noise field with covariance structure given by

$$\begin{aligned} \langle \hat{n}(\hat{r}_1)\hat{n}(\hat{r}_2) \rangle &= \sigma^2 \mathbf{Y}^T(\hat{r}_1)\mathbf{Y}(\hat{r}_2) \\ &= \frac{\sigma^2}{4\pi} \sum_{l=0}^{l_{\max}} (2l+1)P_l(\hat{r}_1 \cdot \hat{r}_2), \end{aligned} \quad (43)$$

where $P_l(x)$ is a Legendre polynomial (Appendix A). Taking the limit $\hat{r}_2 \rightarrow \hat{r}_1$, this implies that $\langle \hat{n}^2(\hat{r}) \rangle = i_{\max} \sigma^2 / (4\pi)$, which represents smoothing relative to the original noise field caused by the use of a finite analysis.

Whilst this smoothing occurs in both the cut- and full-sky formalisms, the presence of \mathbf{C}^{-1} in the former case (equation 42) implies a spatial dependence. As can be seen from Fig. 4, the noise in the cut region is greatly increased, which is a natural way of formally encoding the qualitative fact that, for whatever reason, the data in the cut is contaminated by more than just the white noise field. Thus, despite the invertibility of the coupling matrix (and the band-limited cut-sky analysis), the application of a cut has the desired effect of




Figure 4. The variance of an (initially) uniform Gaussian noise field as a function of latitude, b , after application of the cut-sky orthogonalization described in Section 3.1. Constant latitude cuts of $b_{\text{cut}} = 10^\circ$ (a) and $b_{\text{cut}} = 20^\circ$ (b) were applied (as indicated by the dashed vertical lines) and results are shown for $l_{\max} = 20$ and $l_{\max} = 40$, as labelled. (The oscillations near the peak of the latter curve are indicative of the limited accuracy of the decomposition of the ill-conditioned coupling matrix.)

greatly reducing the impact of any spurious signal, such as the Galaxy. However Fig. 4 also implies that a similar effect could be achieved without performing a cut (and hence leaving the signal unchanged), instead adding a high level of artificial noise in the offending region(s). Finally, it is important to note that the dual assumptions used in the derivation of equation (42) – uniform noise and no beam – are unrealistic, but the manner in which a low-resolution cut-sky analysis works is the same in less idealised scenarios.

4.4.3 High-resolution analysis

The high-resolution case is more straightforward, as the application of the cut results in a data vector, \hat{s}' , which contains little information about the removed region. This is quite distinct from the low-resolution case discussed above, in that here it is the predominantly the signal that is changed, rather than the noise. That said, the noise close to the boundary of the cut is increased in the same manner as explained above. This has the same effect as the apodizing function formalism described by Tegmark (1997), down-weighting points around which there is not full correlation information.

Another difference between the low- and high-resolution analyses is that \hat{s}' is smaller than \hat{s} , from Section 2.1.1. Although this does not result in any significant computational saving, it serves to emphasise the information loss associated with removing part of the sky, and is an independent derivation of the fact that the uncertainties in the estimated power spectrum increase as $4\pi/\Omega_S$ (cf. Hobson & Magueijo 1996; Tegmark 1997).

5 CONCLUSIONS

The upcoming microwave surveys will require a cut-sky analysis to prevent the strong Galactic emissions from contaminating the CMB signal. The spherical harmonics are non-orthogonal on the cut sphere, but an orthonormal basis set can be constructed from them using SVD-based techniques (Section 2). The application of the resultant conversion matrix to the conventional multipoles results in cut-sphere harmonics that contain only the desired information. In the low-resolution case the influence of the Galaxy is reduced by increasing the effective noise in the cut; in the high-resolution limit the orthonormal basis functions can model the infinitely sharp cut sufficiently well that they have no support in the removed region. It is also important to note that the cut should probably only be applied after beam-deconvolution has been attempted, as convolution is ill-defined on the incomplete sphere.

The algorithms described here were implemented to Legendre multipoles of $l_{\max} \approx 2500$ for a constant latitude cut, in which case the coupling matrix of the spherical harmonics is block-diagonal. At present, computational limitations make a general orthogonalization impractical for $l_{\max} \gtrsim 200$, although there are some possibilities to extend this. For instance, only ~ 1 per cent of the coupling matrix contains significant information if the cut is well-chosen (e.g. rectangular in θ and ϕ) and so sparse matrix techniques should thus allow orthogonalization to $l_{\max} \approx 1000$ in this case.

Another requirement is orthogonalization of tensor basis functions on the incomplete sphere, as both the *MAP* and *Planck* satellites will measure polarization. The resultant formalism is more complicated, but the same general principles hold; this issue is explored further by Lewis, Challinor & Turok (2001).

ACKNOWLEDGMENTS

This paper benefited from useful discussions with several members

of the *Planck* collaboration, in particular Mark Ashdown, François Bouchet, Martin Bucher, Rob Crittenden, Jacques Delabrouille, George Efstathiou, Krzysztof Górski, Floor van Leeuwen and Ben Wandelt. DJM was funded by PPARC. ADC acknowledges a PPARC Postdoctoral Research Fellowship. MPH acknowledges a PPARC Advanced Fellowship.

REFERENCES

- Abramowitz M., Stegun I. A., 1971, *Handbook of Mathematical Functions*. 9th ed Dover Publications, New York
- Anderson E. et al., 1992, *LAPACK User's Guide*. Soc. for Indust. and App. Math., Philadelphia
- Arfken G., 1985, *Mathematical Methods for Physicists*. 3rd edn Academic Press, Boston
- Baccigalupi C. et al., 2000, *MNRAS*, 318, 769
- Barreiro R. B., 2000, *New Astron Rev.*, 44, 179
- Bennett C. L. et al., 1992, *ApJ*, 396, L7
- Bersanelli M. et al., 1996, *COBRAS/SAMBA*, The Phase A Study for an ESA M3 Mission. ESA Report D/SCI(96)3
- Birkinshaw M., 1999, *Phys. Rep.*, 310, 97
- Bond J. R., Jaffe A. H., Knox L., 1998, *Phys. Rev. D*, 57, 2117
- Borrill J., 1999, in Maiani L., Mechiorri F., Vittorio N., eds, *3K Cosmology: EC-TMR Conference*. Am. Inst. Phys., New York, p. 277
- Brink D. M., Satchler G. R., 1993, *Angular Momentum*. 3rd edn Clarendon Press, Oxford
- Bouchet F. R., Gispert R., 1999, *New Astron.*, 4, 443
- Challinor A. D., Mortlock D. J., van Leeuwen F., Lasenby A. N., Hobson M. P., Ashdown M. A. J., Efstathiou G. P., 2002, *MNRAS*, submitted
- Coble K. et al., 1999, *ApJ*, 519, L5
- de Bernardis P. et al., 2000, *Nat*, 404, 955
- Delabrouille J., 1998, *A&AS*, 127, 555
- de Oliveira-Costa A., Kogut A., Devlin M. J., Netterfield C. B., Page L. A., Wollack E. J., 1997, *ApJ*, 482, L17
- Efstathiou G., Bridle S. L., Lasenby A. N., Hobson M. P., Ellis R. S., 1999, *MNRAS*, 303, L47
- Golub G. H., van Loan C. F., 1996, *Matrix Computations*. 3rd edn The Johns Hopkins Univ. Press, Baltimore
- Górski K. M., 1994, *ApJ*, 430, L85
- Górski K. M. et al., 1994, *ApJ*, 430, L89
- Gradshteyn I. S., Ryzhik I. M., 2000, *Table of Integrals, Series and Products*. 6th ed. Academic Press, New York
- Halverson N. W. et al., 2002, *ApJ*, submitted
- Haslam C. G. T., Klein U., Salter C. J., Stoffel H., Wilson W. E., Cleary M. N., Cooke D. J., Thomasson P., 1982, *A&A*, 100, 209
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, submitted
- Hobson M. P., Maguiejo J., 1996, *MNRAS*, 283, 1133
- Hobson M. P., Jones A. W., Lasenby A. N., Bouchet F. R., 1998, *MNRAS*, 300, 1
- Hu W., Sugiyama N., Silk J., *Nat*, 386, 37
- Jarosik N. et al., 1998, in Trân Tranh Vân T., Giraud-Héraud Y., Bouchet F., Damour T., Mellier Y., eds, *Fundamental Parameters in Cosmology*. Editions Frontieres, Paris, p. 249
- Jones A. W., Hobson M. P., Lasenby A. N., 1999, *MNRAS*, 305, 898
- Knox L., 1995, *Phys. Rev. D*, 52, 4307
- Landau L. D., Lifshitz E. M., 1975, *Quantum Mechanics*. Pergamon Press, Oxford
- Lee A. T. et al., 2001, *ApJ*, 561, L1
- Lewis A. M., Challinor A. D., Turok N. G., 2001, *Phys. Rev. D*, submitted
- Linde A., 1990, *Particle Physics and Inflationary Cosmology*. Harwood Academic Publishers, Dubbo
- Lineweaver C. H., 1998, *ApJ*, 505, L69
- Maino D. et al., 1999, *A&AS*, 140, 383
- Natoli P., de Gasperis G., Gheller C., Vittorio N., 2001, *A&A*, 372, 346
- Netterfield C. B., Devlin M. J., Jarosik N., Page L., Wollack E. J., 1997, *ApJ*, 474, 47
- Netterfield C. B. et al., 2002, *ApJ*, submitted
- Oh S. P., Spergel D. N., Hinshaw G., 1999, *ApJ*, 510, 551
- Padin S. et al., 2001, *ApJ*, 549, L1
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes: The Art of Scientific Computing*. 2nd edn Cambridge Univ. Press, Cambridge
- Prunet S., Teyssier R., Scully S. T., Bouchet F. R., Gispert R., 2001, *A&A*, 374, 358
- Schlegel D., Finkbinder D., Davies M., 1998, *ApJ*, 500, 525
- Scott P. F. et al., 1996, *ApJ*, 461, L1
- Smoot G. F. et al., 1992, *ApJ*, 396, L1
- Stolyarov V., Hobson M. P., Ashdown M. A. J., Lasenby A. N., 2002, *MNRAS*, submitted
- Sunyaev R. A., Zel'dovich Y. B., 1970, *AP&SS*, 7, 3
- Szapudi I., Prunet S., Pogossyan D., Szalay A. S., Bond J. R., 2001, *ApJ*, 548, 115
- Tanaka S. T. et al., 1996, *ApJ*, 468, L81
- Tegmark M., 1997, *Phys. Rev. D*, 56, 4514
- Tegmark M., Efstathiou G. P., 1996, *MNRAS*, 281, 1297
- Toffolatti L., Argüeso Gómez F., de Zotti G., Mazzei P., Francheschini A., Danese L., Burigana C., 1998, *MNRAS*, 297, 117
- van Leeuwen F. et al., 2001, *MNRAS*, submitted
- Varshalovich D. A., Moskalev A. N., Khersonskii V. K., 1988, *Quantum Theory of Angular Momentum*. World Scientific, Singapore
- Wandelt B. D., Górski K. M., 2001, *Phys. Rev. D*, 63, 123002
- Wandelt B. D., Hivon E., Górski K. M., 2001, *Phys. Rev. D*, 64, 083003
- Wang X., Tegmark M., Zaldarriaga M., 2002, *Phys. Rev. D*, submitted
- Wilson G. W. et al., 2000, *ApJ*, 532, 57
- Wright E. L., Hinshaw G., Bennett C. L., 1996, *ApJ*, 458, L53

APPENDIX A: SPHERICAL HARMONICS

The spherical harmonics form a complete set of orthonormal basis functions over the entire sphere. They are most commonly defined as complex functions (e.g. Landau & Lifshitz 1976; Brink & Satchler 1993), but it is more convenient to use real harmonics in this application. Adapting the notation of Górski (1994), the real spherical harmonics are given by

$$Y_{l,m}(\hat{\mathbf{r}}) = Y_{l,m}(\theta, \phi) = \lambda_{l,|m|}[\cos(\theta)]s_m(\phi), \quad (\text{A1})$$

where $l \geq 0$ and $|m| \leq l$ and

$$s_m(\phi) = \begin{cases} \sqrt{2}\sin(|m|\phi), & \text{if } m < 0, \\ 1, & \text{if } m = 0, \\ \sqrt{2}\cos(|m|\phi), & \text{if } m > 0, \end{cases} \quad (\text{A2})$$

implying that $\int_0^{2\pi} s_m(\phi)s_{m'}(\phi) d\phi = 2\pi\delta_{m,m'}$. For $0 \leq m \leq l$ and $-1 \leq x \leq 1$ the normalized associated Legendre functions are defined by

$$\lambda_{l,m}(x) = \sqrt{\frac{2l+1(l-m)!}{4\pi(l+m)!}} P_{l,m}(x). \quad (\text{A3})$$

Hence $\int_{-1}^1 \lambda_{l,m}(x)\lambda_{l',m}(x) dx = \delta_{l,l'}/(2\pi)$. Under the same conditions the (unnormalized) associated Legendre functions are given by⁵

$$P_{l,m}(x) = (-1)^m (1-x^2)^{m/2} \frac{d^m}{dx^m} P_l(x), \quad (\text{A4})$$

⁵This definition, with the $(-1)^m$ term, corresponds to that given by Abramowitz & Stegun (1971) and Gradshteyn & Ryzhik (2000) but differs from that used by Arfken (1985) and Brink & Satchler (1993).

with the Legendre polynomials given by

$$P_l(x) = \frac{(-1)^l}{2^l l!} \frac{d^l}{dx^l} (1-x^2)^l. \quad (\text{A5})$$

A real field on the sphere, $a(\hat{r})$, can be expanded in terms of spherical harmonic coefficients, given by

$$a_{l,m} = \int_S Y_{l,m}(\hat{r}) a(\hat{r}) d\Omega. \quad (\text{A6})$$

This can be inverted to give

$$a(\hat{r}) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l a_{l,m} Y_{l,m}(\hat{r}), \quad (\text{A7})$$

provided that $l_{\max} \rightarrow \infty$, due to the orthonormality of the spherical harmonics on the full sphere:

$$\int_S Y_{l,m}(\hat{r}) Y_{l',m'}(\hat{r}) d\Omega = \delta_{l,l'} \delta_{m,m'}. \quad (\text{A8})$$

If a finite l_{\max} is used this inversion is no longer possible for general $a(\hat{r})$, although it does still hold for band-limited functions.⁶

Whilst the two indices l and m have quite distinct interpretations, it is convenient to combine them into a single index, i , which allows the definition of vectors $\mathbf{Y}(i) = Y_{i(l,m)}(\hat{r})$ and $\mathbf{a} = a_{i(l,m)}$. Two obvious indexing schemes present themselves: grouping in l and m . The first, as introduced by Górski (1994), is natural for power spectrum estimation and very simple:

$$i(l, m) = l^2 + l + m + 1. \quad (\text{A9})$$

The two ‘inverses’ of this relationship are

$$l = \text{int}[(i-1)^{1/2}] \quad (\text{A10})$$

and

$$m = i - (l^2 + l + 1). \quad (\text{A11})$$

The second choice of ordering is useful in cases of azimuthal symmetry in which the orthogonality expressed in equation (A2) is maintained, and grouping in m is achieved by defining

$$i(l, m) = \begin{cases} l + m + 1 & \\ + (l_{\max} + m)(l_{\max} + m + 1)/2, & \text{if } m \leq 0, \\ l - l_{\max} + (l_{\max} + 1)^2 & \\ - (l_{\max} - m)(l_{\max} - m + 1)/2, & \text{if } m > 0. \end{cases} \quad (\text{A12})$$

The ‘inverses’ in this case are given by

$$m = \begin{cases} \text{int}[-l_{\max} + \frac{1}{2}] & \text{if } i \leq (l_{\max} + 1) \\ \times (\sqrt{8i + 1} - 3), & \times (l_{\max} + 2)/2, \\ \text{int}[-l_{\max} - \frac{1}{2}] & \text{if } i > (l_{\max} + 1) \\ \times [\sqrt{8[(l_{\max} + 1)^2 - i + 1] + 1} - 3], & \times (l_{\max} + 2)/2 \end{cases} \quad (\text{A13})$$

⁶ A band-limited function can, by (somewhat circular) definition, be constructed from a finite sum over spherical harmonics.

and

$$l = \begin{cases} i - [m + 1 & \\ + (l_{\max} + m)(l_{\max} + m + 1)/2], & \text{if } m \leq 0, \\ i - [-l_{\max} + (l_{\max} + 1)^2 & \\ - (l_{\max} - m)(l_{\max} - m + 1)/2], & \text{if } m > 0. \end{cases} \quad (\text{A14})$$

Other indexing schemes have been used in the more specific case of simulated *Planck* data-sets in which the sky coverage is periodic in azimuth (van Leeuwen, private communication), but are beyond the scope of this paper.

APPENDIX B: INTEGRATION OF PRODUCTS OF ASSOCIATED LEGENDRE FUNCTIONS

In Section 2.3 integrals of the form

$$I_{l,l',m}(x_1, x_2) = \int_{x_1}^{x_2} \lambda_{l,m}(x) \lambda_{l',m}(x) dx \quad (\text{B1})$$

arose; here the $\lambda_{l,m}(x)$ are the normalised associated Legendre functions, defined in equation (A3), and m is assumed to be non-negative. These integrals can be evaluated quickly and accurately using a combination of closed formulæ and recursion relations.

The associated Legendre functions, $P_{l,m}(x)$ (defined in equation A4), are solutions of the ordinary differential equation (e.g. Arfken 1985)

$$\frac{d}{dx} \left[(1-x^2) \frac{dP_{l,m}}{dx} \right] + \left[l(l+1) - \frac{m^2}{1-x^2} \right] P_{l,m}(x) = 0. \quad (\text{B2})$$

Multiplying this equation by $P_{l',m}(x)$ and integrating (from x_1 to x_2) by parts twice yields

$$(l-l')(l+l'+1) \int_{x_1}^{x_2} P_{l,m}(x) P_{l',m}(x) dx = \left[(1-x^2) P_{l,m}(x) \frac{dP_{l',m}}{dx} - (1-x^2) P_{l',m}(x) \frac{dP_{l,m}}{dx} \right] \Big|_{x=x_1}^{x=x_2}. \quad (\text{B3})$$

This is a reflection of the standard result that integrals of solutions of a self-adjoint differential equation (as equation B2 is) can be expressed as boundary terms (e.g. Arfken 1985). The derivatives in equation (B2) can be removed by using the standard recursion relationship (e.g. Gradshteyn & Ryzhik 2000)

$$(1-x^2) \frac{dP_{l,m}}{dx} = (l+m)P_{l-1,m}(x) - lxP_{l,m}(x) \quad (\text{B4})$$

to yield, for $l \neq l'$,

$$\int_{x_1}^{x_2} P_{l,m}(x) P_{l',m}(x) dx = \frac{1}{(l-l')(l+l'+1)} \times [(l'+m)P_{l,m}(x)P_{l'-1,m}(x) + (l-l')xP_{l,m}(x)P_{l',m}(x) - (l+m)P_{l-1,m}(x)P_{l',m}(x)] \Big|_{x=x_1}^{x=x_2}. \quad (\text{B5})$$

Note that the first term must be omitted if $l' = m$ and that the third term must be omitted if $l = m$; these Legendre functions are implicitly zero from equation (A4). Finally, this can be normalized

according to equation (A3), giving

$$\begin{aligned}
 I_{l,l',m}(x_1, x_2) &= \frac{1}{(l-l')(l+l'+1)} \\
 &\times \left[\sqrt{\frac{2l'+1}{2l'-1}} (l^2 - m^2) \lambda_{l,m}(x) \lambda_{l'-1,m}(x) \right. \\
 &+ (l-l') x \lambda_{l,m}(x) \lambda_{l',m}(x) \\
 &\left. - \sqrt{\frac{2l+1}{2l-1}} (l^2 - m^2) \lambda_{l-1,m}(x) \lambda_{l',m}(x) \right] \Bigg|_{x=x_2}^{x=x_1}. \quad (\text{B6})
 \end{aligned}$$

An alternative derivation of this result was presented by Wandelt et al. (2001); it is also in principle equivalent to equation (13) in section 5.9 of Varshalovich, Moskalev & Khersonskii (1988), but their application of equation (B4) is in error.

For the case $l = l'$, a recursion relation is required, starting with $l = m$. Combining equations (A3) and (A4),

$$\lambda_{m,m}(x) = (-1)^m \sqrt{\frac{2m+1}{4\pi}} (2m-1)!! (1-x^2)^{m/2}, \quad (\text{B7})$$

where $n!! = 1 \times 3 \times \dots \times (n-2) \times n$ for odd n . Integrating by parts and using equation (B7) again gives

$$I_{m,m,m}(x_1, x_2) = \begin{cases} \frac{x_2 - x_1}{4\pi}, & \text{if } m = 0, \\ I_{m-1,m-1,m-1}(x_1, x_2) + \frac{x \lambda_{m,m}^2(x) \Big|_{x_1}^{x_2}}{2m+1}, & \text{if } m > 0. \end{cases} \quad (\text{B8})$$

Moving to $l = m + 1$, the standard relationship (e.g. Gradshteyn & Ryzhik 2000) that

$$\lambda_{m+1,m}(x) = (2m+3)x \lambda_{m,m}(x) \quad (\text{B9})$$

combines with equation (B7) to give

$$\begin{aligned}
 I_{m+1,m+1,m}(x_1, x_2) &= I_{m,m,m}(x_1, x_2) \\
 &- \frac{2m+2}{2m+3} x \lambda_{m+1,m+1}^2(x) \Bigg|_{x=x_1}^{x=x_2}, \quad (\text{B10})
 \end{aligned}$$

where the first term is given in equation (B8).

The last step is to derive a recursion relation relating $I_{l,l,m}(x_1, x_2)$ to $I_{l-1,l-1,m}(x_1, x_2)$ and $I_{l-2,l-2,m}(x_1, x_2)$. Equation (C22) of Wandelt et al. (2001) gives a four-term recursion to obtain $I_{l,l',m}(x_1, x_2)$; it can be applied successively (once swapping l and l') to obtain

$$\begin{aligned}
 I_{l,l,m}(x_1, x_2) &= -\frac{2m^2 - 2l^2 + 2l - 1}{(l-m)(l+m)} I_{l-1,l-1,m}(x_1, x_2) \\
 &- \frac{(l-1-m)(l-1+m)}{(l-m)(l+m)} I_{l-2,l-2,m}(x_1, x_2) \\
 &- \sqrt{\frac{2l-1}{2l+1}} \frac{1}{(l-m)(l+m)} \\
 &\times (1-x^2) \lambda_{l,m}(x) \lambda_{l-1,m}(x) \Bigg|_{x=x_1}^{x=x_2} \\
 &+ \sqrt{\frac{2l-1}{2l-3}} \frac{(l-1-m)(l-1+m)}{(l-m)^2(l+m)^2} \\
 &\times (1-x^2) \lambda_{l-1,m}(x) \lambda_{l-2,m}(x) \Bigg|_{x=x_1}^{x=x_2}. \quad (\text{B11})
 \end{aligned}$$

In summary, equation (B6) can be used to evaluate all $I_{l,l',m}(x_1, x_2)$

for which $l \neq l'$, and equations (B8), (B10) and (B11) combine to give all $I_{l,l,m}(x_1, x_2)$ recursively.

APPENDIX C: TREATMENT OF NON-COSMOLOGICAL MODES

All the cosmological information encoded in the CMB is expected to be contained in the $l \geq 2$ modes; the $l = 0$ mode in an isotropic universe can be normalised arbitrarily and the $l = 1$ modes can be set to zero by adopting an appropriate reference frame. None the less, observations of the microwave sky will yield non-zero monopole and dipole values for a number of reasons (e.g. the observer's motion; Galactic emission; extragalactic point sources). Hence these low-order modes must be included in the analysis of CMB data, but should be kept separate from the cosmological modes, as is naturally the case if spherical harmonic coefficients are used to describe the data. It is also important to note that the properties of the basis functions themselves are unimportant – the essential requirement is that only four of the cut-sky harmonic coefficients contain information on the unwanted modes.

The method of orthogonalization summarized in equations (11), (12) and (14) does not explicitly impose any particular structure on the conversion matrix, \mathbf{A} (which relates harmonic coefficients on the incomplete sphere to those on the full sphere by $\mathbf{a}' = \mathbf{A}^T \mathbf{a}$; equation 24). The non-cosmological modes are kept separate from the cosmological modes if the first four columns of \mathbf{A}^T have only zeros from the fifth row on, assuming the full-sky harmonic coefficients are indexed using l -ordering (Appendix A). This is achieved naturally if \mathbf{A}^T is constructed to be upper triangular, as in the case of the Cholesky decomposition described in Section 2.2.1. The other decomposition methods discussed in Section 2.2 do not share this property, and so the resultant conversion matrices must be adjusted explicitly.

One way of achieving this is to use a partial Householder transform (e.g. Press et al. 1992). The last $i'_{\max} - i$ elements of the i th column of a general $i'_{\max} \times i_{\max}$ matrix can be set to zero by the transformation $\mathbf{M}' = \mathbf{P}_i \mathbf{M}$, with the orthogonal Householder matrix defined by

$$\mathbf{P}_i = \mathbf{I} - 2 \frac{\mathbf{m}_i \mathbf{m}_i^T}{\mathbf{m}_i^T \mathbf{m}_i}, \quad (\text{C1})$$

where \mathbf{m}_i is given by

$$(m_i)_j = \begin{cases} 0, & \text{if } j < i, \\ M_{j,i} + \sqrt{\sum_{k=i}^{i'_{\max}} M_{k,i}^2}, & \text{if } j = i, \\ M_{j,i}, & \text{if } j > i, \end{cases} \quad (\text{C2})$$

and $i \leq \min(i'_{\max}, i_{\max})$ is assumed. Provided that the Householder matrix applied to \mathbf{B} is that generated from \mathbf{A}^T , the transformations $\mathbf{A}'^T = \mathbf{P}_i \mathbf{A}^T$ and $\mathbf{B}' = \mathbf{P}_i \mathbf{B}$ leave equations (12) and (14) unaffected as \mathbf{P}_i is orthogonal by construction. Applying \mathbf{P}_1 , \mathbf{P}_2 , \mathbf{P}_3 and then \mathbf{P}_4 to the successively updated \mathbf{A}^T ensures that the $l = 0$ and $l = 1$ modes influence only the first four cut-sky harmonic coefficients, as required. This procedure could be continued, moving \mathbf{A}^T successively closer to upper triangular form, although this cannot be achieved in full as \mathbf{A}^T has more columns than rows.

Special mention must be made of the constant latitude cut case, the symmetry of which can only be utilized if the spherical harmonics are indexed using m -ordering (Appendix A). In this case only the $m = 0$ and $m = \pm 1$ blocks have any contribution from the

monopole or dipole, and each can be treated separately. Further, the ordering within these blocks is such that the non-cosmological modes are in the first rows, and so the above algorithm can be applied to each of three blocks as is. The only slight inconvenience is that it is no longer the first four cut-sky modes that contain the

non-cosmological information, and the relevant modes must be flagged explicitly.

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.