

Deconvolution map-making for cosmic microwave background observationsCharmaine Armitage^{1,*} and Benjamin D. Wandelt^{1,2,3,†}¹*Department of Physics, UIUC, 1110 West Green Street, Urbana, Illinois 61801, USA*²*Department of Astronomy, UIUC, 1002 West Green Street, Urbana, Illinois 61801, USA*³*Center for Advanced Studies, UIUC, 912 West Illinois Street, Urbana, Illinois 61801, USA*

(Received 6 October 2004; published 16 December 2004)

We describe a new map-making code for cosmic microwave background observations. It implements fast algorithms for convolution and transpose convolution of two functions on the sphere [B. Wandelt and K. Górski, Phys. Rev. D **63**, 123002 (2001)]. Our code can account for arbitrary beam asymmetries and can be applied to any scanning strategy. We demonstrate the method using simulated time-ordered data for three beam models and two scanning patterns, including a coarsened version of the WMAP strategy. We quantitatively compare our results with a standard map-making method and demonstrate that the true sky is recovered with high accuracy using deconvolution map-making.

DOI: 10.1103/PhysRevD.70.123007

PACS numbers: 98.70.Vc, 95.75.Pq

I. INTRODUCTION

Real microwave telescopes collect distorted information about the cosmic microwave background (CMB) anisotropies due to asymmetries in the beam shape [1] and stray light from sources such as the Galaxy [2,3]. To correct for these systematic errors, we must be able to remove the detector response at all orientations of the telescope over the whole sky. In an optimal treatment, this correction must be applied during the map-making step of the CMB data analysis pipeline, before the angular power spectrum can be reconstructed. The problem becomes increasingly important as new generations of CMB observations probe for ever fainter signals in the CMB sky, and especially as we are preparing to measure the polarization of the CMB with high sensitivity. We present a complete map-making algorithm, in which time-ordered data (TOD) is used to construct a temperature map and beam distortions are removed.

We call our approach *deconvolution map-making*, a generalization of existing CMB map-making techniques to solve the maximum likelihood map-making problem for arbitrary beam shapes. For sufficiently high signal to noise, this technique allows super-resolution imaging of the CMB from time-ordered scans. We implement our method using the exact algorithms for the convolution and transpose convolution of two arbitrary functions on the sphere—in this case the sky and the beam—as detailed by Wandelt and Górski in [4]. These fast methods for convolution and transpose convolution are efficient because they make use of the fast Fourier transform algorithm. They are guaranteed to work to numerical precision for bandlimited functions on the sphere.

Early work on the map-making problem has relied on the brute force method of direct matrix inversion. However, current and future CMB experiments, like the Wilkinson

Microwave Anisotropy Probe (WMAP) [5] and Planck satellite [6], return enormous data sets that render the brute force method useless. More recent advancements include map-making methods applicable to the latest experiments; however, many treat the beam like a perfect delta function (e.g., [7,8]) or assume a symmetric beam profile (e.g., [9]), and thereby relegate the problem of treating a non-Gaussian radial response of the beam to subsequent stages in the data analysis [10]. In this class, special techniques exist to deal with differential measurements like that of the Differential Microwave Radiometer (DMR) on the Cosmic Background Explorer (COBE) satellite [11] or WMAP [12]. A Fourier method has been developed [13] to perform deconvolution but only for nonrotating asymmetric beams. Last, [14] present a method to remove the main beam distortion over patches of the sky for asymmetric, rotating beams but operate in pixel space which is computationally more expensive than spherical-harmonic-space algorithms for the same level of accuracy [4,15].

We test our algorithm on a simulated foreground- and Galaxy-free sky using a standard Λ CDM power spectrum and simulated spherical-harmonic multipoles $a_{\ell m}$ up to $\ell = 128$. We also use the first-year WMAP Ka-band temperature map as our true sky containing galactic emission.

In Sec. II we present the deconvolution method and briefly review a standard map-making method. In Sec. III we detail the various test cases. Our results for the deconvolution method are given, discussed, and compared with the standard estimates in Sec. IV. We conclude in Sec. V and remark on future directions.

II. DECONVOLUTION MAP-MAKING

In order to define our notation, we will briefly review the path from observations to maps. A microwave telescope scans the CMB sky according to some scanning strategy, effectively convolving the true sky with a beam function, and returns a vector, \mathbf{d} , containing the n_{TOD} samples of the time-ordered data. We represent this by

*Electronic address: carmitag@uiuc.edu

†Electronic address: bwandelt@uiuc.edu

$$\mathbf{A} \mathbf{s} = \mathbf{d}, \quad (1)$$

where \mathbf{A} is the observation matrix, defined below, and \mathbf{s} is a n_{pix} vector containing the true sky.

The matrix \mathbf{A} encodes both the scanning strategy and the optics of the CMB instrument. Each sample of the TOD is modeled as the scalar product of a row of the matrix \mathbf{A} with the sky \mathbf{s} . Each of the n_{TOD} rows of \mathbf{A} contains a rotated map of the beam. In a given row, the beam rotation corresponds to the orientation of the antenna at the point in time when the sample is taken. We will assume the beam shape and pointing of the satellite to be known.

The observation matrix \mathbf{A} generalizes the notion of the pointing matrix which is often used in expositions of map-making algorithms by including both optics and scanning strategy. This generalization is necessary for any map-making method that accounts for beam functions with azimuthal structure.

The least-squares estimate of the true sky, $\hat{\mathbf{s}}$, is given by

$$\mathbf{A}^T \mathbf{A} \hat{\mathbf{s}} = \mathbf{A}^T \mathbf{d}. \quad (2)$$

The coefficient matrix in this system of equations, $\mathbf{A}^T \mathbf{A}$, is a smoothing matrix and, hence, ill-conditioned. Inverting it to solve Eq. (2) therefore poses a problem.

We describe here a regularization technique for dealing with this problem. We split off the ill-conditioned part of \mathbf{A} by factoring the convolution operator into $\mathbf{A} = \mathbf{B}\mathbf{G}$ where \mathbf{G} is a simple Gaussian smoothing matrix, represented in harmonic space by

$$G_\ell = \exp\left(\frac{-\sigma^2 \ell(\ell+1)}{2}\right), \quad (3)$$

where $\sigma = \text{FWHM}/\sqrt{8 \ln 2}$.

Substituting the factorization into Eq. (2), we get

$$\mathbf{G}^T \mathbf{B}^T \mathbf{B} \mathbf{G} \hat{\mathbf{s}} = \mathbf{G}^T \mathbf{B}^T \mathbf{d}, \quad (4)$$

$$\mathbf{B}^T \mathbf{B} \mathbf{x} = \mathbf{B}^T \mathbf{d}, \quad (5)$$

where we are solving for $\mathbf{x} = \mathbf{G} \hat{\mathbf{s}}$ so as not to reconstruct the sky at higher resolution than that of the instrument.

Equation (2) is exact if the noise is stationary and uncorrelated in the time-ordered domain. For a more general noise covariance matrix in the time-ordered domain, \mathbf{N} , the normal equation is modified as follows:

$$\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A} \hat{\mathbf{s}} = \mathbf{A}^T \mathbf{N}^{-1} \mathbf{d}. \quad (6)$$

We proceed, considering only white noise in this paper; however, it is straightforward to generalize to nonwhite noise as indicated in Eq. (6). Indeed, the matrix-vector operations required for this generalization have already been implemented in publicly available map-making codes (e.g., MADMAP [16]).

A. Fast convolution on the sphere

We now briefly review the relevant formalism for general convolutions on the sphere and refer the reader to [4] for the full details of the fast convolution and transpose convolution of two functions in the spherical-harmonic basis. The convolution of a bandlimited beam function $b(\vec{\gamma})$ with the sky $s(\vec{\gamma})$ is given by the following integral over all solid angles:

$$T(\Phi_2, \Theta, \Phi_1) = \int d\Omega_{\vec{\gamma}} [\hat{D}(\Phi_2, \Theta, \Phi_1) b](\vec{\gamma})^* s(\vec{\gamma}), \quad (7)$$

where \hat{D} is the operator of finite rotations.¹ In spherical-harmonic space this becomes

$$T_{mm'm''} = \sum_{\ell} s_{\ell m} d_{mm}^{\ell}(\theta_E) d_{m'm''}^{\ell}(\theta) b_{\ell m}^*. \quad (8)$$

Analogously, the transpose convolution of $T(\Phi_2, \Theta, \Phi_1)$ is given by

$$y^*(\vec{\gamma}) = \int d\Phi_2 d\Theta d\Phi_1 \times [\hat{D}(\Phi_2, \Theta, \Phi_1) b](\vec{\gamma})^* T(\Phi_2, \Theta, \Phi_1), \quad (9)$$

and in spherical harmonics

$$y_{\ell m}^* = \sum_{m'm''} d_{mm'}^{\ell}(\theta_E) d_{m'm''}^{\ell}(\theta) b_{\ell m}^* T_{mm'm''}. \quad (10)$$

An important feature of our approach is that it economizes the computational effort if the beam is nearly azimuthally symmetric. The parameter of the method that sets the degree to which asymmetries of the beam are taken into account is m_{\max} , the maximum m'' in Eqs. (8) and (10). For $m_{\max} = 0$, we recover the computational cost of simple spherical harmonics transforms, $\mathcal{O}(\ell_{\max}^3)$. Since m_{\max} is bounded from above by ℓ_{\max} , the computational cost of the method never scales worse than $\mathcal{O}(\ell_{\max}^4)$. For a mildly elliptical beam, we anticipate that just including the $m_{\max} = 0$ and $m_{\max} = 2$ terms will suffice, since the $m_{\max} = 1$ term vanishes by symmetry.

For clarity, we now rewrite Eq. (5) in the compact spherical-harmonic basis (summing over repeated indices):

$$\mathbf{A}_{L'M'mm'm''}^T \mathbf{B}_{mm'm''LM} \mathbf{x}_{LM} = \mathbf{A}_{L'M'mm'm''}^T T_{mm'm''}, \quad (11)$$

where \mathbf{A}^T acting on $T_{mm'm''}$ is given by Eq. (10) and \mathbf{B} acting on \mathbf{x}_{LM} is given by Eq. (8).

To make matters even more concrete, we now explicitly describe the steps required to simulate time-ordered data \mathbf{d} from a map (“simulation”). We convolve the beam $b_{\ell m}$ with the map $a_{\ell m}$ to obtain $T_{mm'm''}$. Then we inverse Fourier transform the $T_{mm'm''}$ to get $T(\Phi_2, \Theta, \Phi_1)$. Next, we must account for the scan path $[\Phi_2(t), \Theta(t), \Phi_1(t)]$, where Φ_2 and Θ specify the position on the sphere and

¹Our Euler angle convention is defined as active right-hand rotations about the z , y , and z axes by Φ_2 , Θ , Φ_1 , respectively

Φ_1 specifies the orientation of the beam. This is achieved by extracting those values in $T(\Phi_2, \Theta, \Phi_1)$ which fall on the scan path whenever the instrument samples the sky.

As a second example, we describe how to compute the right-hand side of Eq. (5). Start with the TOD \mathbf{d} . For each sample in \mathbf{d} , the scanning strategy specifies the orientation (Φ_2, Θ, Φ_1) . The sampled temperature is added into the element of an initially empty array which is identical in size and shape to the array which stored $T(\Phi_2, \Theta, \Phi_1)$. We have effectively binned the TOD \mathbf{d} , according to the position and orientation of the beam on the sky. Let us therefore refer to this operation as “binning.” In order to minimize discreteness effects due to the gridded representation of $T(\Phi_2, \Theta, \Phi_1)$, more sophisticated interpolation techniques could be implemented. Additionally, the resolution of the grid into which the data is binned may be increased.

B. Solving the deconvolution equations

To obtain the optimal map estimate, we numerically solve the linear system of equations in Eq. (5) for $x_{\ell m}$. We have a choice between direct and iterative solution methods. An iterative method is advantageous compared to a direct method (such as Cholesky inversion) if the cost per iteration times the number of iterations required to converge to sufficient accuracy is less than the cost of the direct method.

For the problem sizes of current and upcoming CMB missions, where the map contains a number of pixels $n_{pix} \sim 10^6-10^7$, direct solution methods would be prohibitive for two reasons. First, the required number of floating point operations scales as n_{pix}^3 . Second, the amount of space required to store the coefficient matrix and its inverse scales as n_{pix}^2 . Therefore direct solution exceeds the capabilities of modern supercomputers by several orders of magnitude. For the Planck mission direct solution would require of order 10^{21} floating point operations and hundreds of terabytes of random access memory.

We therefore advocate using an iterative technique, the conjugate gradient (CG) method [17]. The CG method is well suited to this problem. It solves linear systems with the symmetric positive definite coefficient matrix and has advantageous convergence properties compared to other iterative methods such as the Jacobi method [17]. In order to apply the CG method, we must be able to apply the coefficient matrix on the left-hand side of Eq. (5) to our current guess of the solution \mathbf{x} . In order to do so, we simply perform the two operations of simulation and “binning” in succession. The fast convolution and transpose convolution algorithms allow computing the action of the coefficient matrix on a map without ever having to store the matrix coefficient in memory.

It is desirable to minimize the number of iterations the CG method requires to converge to a given level of accuracy. This can be done by “preconditioning” the system of equations. Preconditioning amounts to multiplying on both

sides with an approximation of the inverse of the coefficient matrix and solving this modified system. As long as the preconditioner is nonsingular, the solution will be the same for the original and the preconditioned systems, but for a well-chosen preconditioner the number of iterations can be reduced significantly. A natural choice of the preconditioning matrix, which we used to obtain the results in this paper, is the diagonal matrix $[\text{diag}(\mathbf{A}^T \mathbf{A})]^{-1}$, which, for a delta-function beam, is just the inverse of the number of hits per pixel.

At every iteration we have an approximate solution $\tilde{\mathbf{x}}$ of Eq. (5). We assess convergence by computing the ratio of L_2 norms

$$\frac{L_2[\mathbf{B}^T \mathbf{B} \tilde{\mathbf{x}} - \mathbf{B}^T \mathbf{d}]}{L_2[\mathbf{B}^T \mathbf{d}]}, \quad (12)$$

where $L_2[\mathbf{x}] \equiv \sqrt{|\mathbf{x} \cdot \mathbf{x}|}$.

C. Standard map-making: Brief review and critique

In order to compare our results to traditional techniques, we also implemented a traditional map-making code that solves the normal equation [Eq. (2)] assuming an azimuthally symmetric beam. In this implementation, the observation matrix \mathbf{A} becomes the pointing matrix, containing only a single entry on each row corresponding to the direction in which the main beam lobe is pointing at the time of sampling. Standard map-making therefore reconstructs a map which is smoothed by an effective beam whose shape varies as a function of position on the map. This variation depends on the scanning strategy. More precisely, at any given position on the estimated map the effective beam shape depends on the various orientations of the beam as it passed through this position during the scan.

For uncorrelated noise and an azimuthally symmetric beam, the solution of the normal equation is simple to compute: Bin the TOD into discrete sky pixels, summing over repeated hits, and dividing through by the number of hits per pixel. Numerical implementations of this algorithm and its generalization to correlated noise have been described in the literature [10]. However, all of these treatments assume azimuthally symmetric beams. For experiments with highly asymmetric beams and where contamination from the Galaxy is picked up in the side-lobes, we expect that this method will not fare well against our deconvolution method which also removes artifacts due to these optical systematics. We use the same TOD and scan path as for our deconvolution method. Here, the data is binned into pixelized maps, rather than into the $T(\Phi_2, \Theta, \Phi_1)$ grid. Unless otherwise stated we use the HEALPix pixelization scheme [18] with resolution parameter $n_{side} = 64$. The angular scale of a pixel is therefore just under 1° . Recall that our regularization method returns a smoothed map with an effective, azimuthally symmetric Gaussian beam. Thus, in order to compare the

two methods we must make a similar modification to our standard map-making. We read out the resulting $a_{\ell m}$ of our standard map (using the HEALPix anafast routine), after the binning step, and modify them in the following way:

$$a'_{\ell m} = \frac{a_{\ell m}}{B_{\ell}} G_{\ell}, \quad (13)$$

where B_{ℓ} is the beam power spectrum and G_{ℓ} is given in Eq. (3).

III. TEST CASES

In this section we detail our tests and comparisons of the deconvolution and standard map-making methods. For the purposes of testing our method, we create several mock beam models $b_{\ell m}$. We test three possible beam shapes which break azimuthal symmetry progressively strongly, two scanning patterns, and skies with and without galactic emission.

The first beam is a simple model of a sidelobe; it is composed of a Gaussian beam of $\text{FWHM} = 180'$ rotated at 90° to another Gaussian beam of $\text{FWHM} = 180'$. Both the main beam and the sidelobe are normalized such that they integrate to one. The second beam models a (somewhat exaggerated) elliptical shape, composed of two identical Gaussian beams with $\text{FWHM} = 180'$ whose centers are on both sides of the optical axis, separated by $180'$. The third beam is composed of two identical Gaussian beams ($\text{FWHM} = 180'$) rotated at 140° from each other; we refer to this as the two-beam model. This case is motivated by the design of the WMAP observatory [19].

We set the asymmetry parameter m_{\max} for our three cases (sidelobe, elliptical, and two-beam) to 8, 38, and 128, respectively.

Following [4], we first considered a *basic scan path* (BSP) in which the beam scans the full sky on rings of constant longitude with no rotation about its outward axis. To be clear, for the case of the sidelobe beam, the smaller beam follows this ringed scan while the larger beam remains fixed at the equatorial longitude. Similarly, in the two-beam model, one beam follows the ring scan while the other rotates in smaller circles 140° away. The central lobe therefore covers the whole sky, while the offset beam remains within a band of $\pm 50^\circ$ centered on the ecliptic plane. The elliptical beam simply follows the ring scan, and is oriented such that its long axis remains perpendicular to the lines of longitude.

A more realistic observational strategy has a beam that revisits locations on the sky in different orientations. Therefore, we model the one-year WMAP scan path followed by one horn. The WMAP scan strategy also covers the full sky and includes a spin modulation of 0.464 revolutions per minute and a spin precession of one revolution per hour [19]. We used a scaled-down model of the WMAP scan in which the spin modulation is 0.00232

revolutions per minute and a step size of about 46 s (roughly 562 samples per period). This produces a pattern very similar to the spirograph-type pattern shown in Fig. 4 of [19]. We refer to this as the WMAP-like scan path (WSP). The WSP has about 6 times as many samples as the BSP. For this strategy, the spirograph pattern is followed by the small beam of the sidelobe, the elliptical beam, and *both* beams of the two-beam model. In the two-beam case both beams are offset from the spin axis of the satellite, to mimic the WMAP scanning geometry. It is not differential in nature, since both beams have positive weight.

We test each beam (sidelobe, elliptical, and two-beam) with both scanning patterns (BSP and WSP) on a sky without galactic emission. We refer to these as the six main test cases.

In reality, CMB experiments will also pick up signal from the Galaxy. We use the first-year WMAP Ka-band temperature map, degraded to an *nside* of 64 and smoothed with a Gaussian beam of $\text{FWHM} = 180'$ as our model of the true sky with galactic emission. For our last test case, we convolve it with the sidelobe beam.

For each test case, we assume that the beamshapes of the instrument are known and use the deconvolution method to deconvolve the map with the same beam with which the true sky was originally convolved. We attempt to recover features in the map corresponding to the smallest scale features of our test beams. We therefore set the width of our regularization kernel, represented by the matrix \mathbf{G} in Eqs. (3) and (4), to $\text{FWHM} = 180'$ in every case. We compare our map estimates to the true sky, smoothed by the regularization kernel. When we refer to the “true” sky in the following we mean this kernel-smoothed input sky.

IV. RESULTS AND DISCUSSION

We present the results of the deconvolution algorithm for the six main test cases in the form of residual maps. We compare these results to the results from standard map-making by examining their power spectra and by calculating the root mean square (RMS) difference between the estimated and true sky. For the tests that include the galactic signal, we show the actual map estimates.

In Fig. 1 we plot ratios of the power spectra of the residual maps (both standard and deconvolved) and the power spectrum of the input map. The BSP (WSP) results are plotted in the left (right) column. The solid (dashed) lines represent the relative difference in C_{ℓ} between the deconvolved (standard) map and the true-sky map. The standard map-making algorithm failed to give meaningful results for the two-beam test. We therefore excluded this case from the plot.

For all cases we chose to present the results after a fixed number of iterations to show the impact of scanning strat-




FIG. 1. Ratios of the spectra of the residual map to the spectrum of the input map for each of the beam models and both scanning strategies. The BSP results are plotted in the left column and the WSP results are plotted in the right column. Results for the sidelobe, elliptical, and two-beam beams are shown in the top, middle, and bottom panels, respectively. The solid lines correspond to deconvolved spectra and the dashed lines correspond to the standard spectra.

egy and beam pattern on the condition number of the map-making equations. We find that the deconvolution algorithm outperforms standard map-making by orders of magnitude in accuracy.

For a fixed number of iterations, the BSP tests performed less well than the WSP tests. The two-beam BSP and, to a lesser extent, the elliptical beam BSP test cases have not converged to sufficient accuracy.

There are several possible causes for this behavior. The BSP leads to an extremely nonuniform sky coverage. Also, the BSP visits each pixel in a narrow range of beam orientations. Further, the number of sky samples is smaller for the BSP case than for the WSP case (as noted in Sec. III). All of these aspects can contribute to increasing the condition number of the normal equation, which in turn leads to smaller error decay per iteration of the preconditioned CG solver.

In Table I, we summarize the RMS difference between the reconstructed and the true sky. The RMS values are computed using the standard deviations of the residual and true maps:

TABLE I. Fractional RMS error for each of the six main test cases.

Beam	BSP		WSP	
	Standard	Deconvolved	Standard	Deconvolved
Sidelobe	0.257 467	0.000 216 367	0.178 828	$3.137 41 \times 10^{-7}$
Elliptical	0.186 262	0.021 365 7	0.129 715	$2.927 14 \times 10^{-5}$
Two-beam	N/A	0.102 778	N/A	$1.085 79 \times 10^{-6}$

$$\text{RMS} = \frac{\text{stdev}(\text{residual map})}{\text{stdev}(\text{true map})}, \quad (14)$$

where

$$\text{residual map} = \text{estimated map} - \text{true map}. \quad (15)$$

The RMS values reflect the trends seen in the spectra in Fig. 1. The residual maps are shown in Fig. 2. In order that the scale of the axes on the residual maps are meaningful, we also show the true-sky map.

Achieving a stably converging iterative solution method for the deconvolution problem is a success of our regularization technique. The convergence of our iterative solver as a function of iteration number is plotted in Fig. 3.

In order to be able to compare the performance of our method for different beam patterns and scanning strategies, we make the deliberate choice of limiting the num-

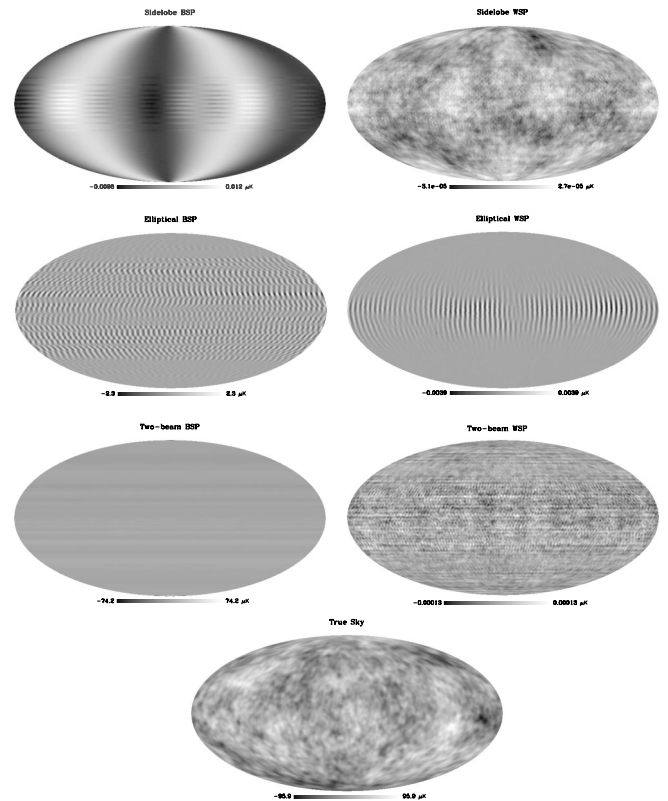


FIG. 2. Residuals after the first 100 iterations are shown in the first three rows. The figures on the left are for the basic scan path; the ones on the right for the WMAP-like scan path. First, second, and third rows correspond to the sidelobe, elliptical, and two-beam beams, respectively. The true-sky is shown in the fourth row. Note that solutions of the two-beam BSP and, to a lesser extent, the elliptical beam BSP test cases have not converged to sufficient accuracy. We chose to present the results for all cases after a fixed number of iterations to show the impact of scanning strategy and beam pattern on the condition number of the map-making equations.




FIG. 3. Convergence rates of the preconditioned conjugate gradient solver for each test case. The left panel refers to the basic scan path and the right panel to the WMAP-like scan path. The solid lines correspond to the sidelobe beam, dotted lines to the elliptical beam, and dot-dashed lines to the two-beam model.

ber of iterations to 100 and comparing the best results obtained up to this point. Since our error estimate continues to drop stably (except in two cases, where we reach the single precision numerical accuracy floor after ~ 25 and ~ 70 iterations), it is clear that the accuracy of the reconstruction can be improved by allowing the system to iterate further, or by choosing a more sophisticated preconditioner.

Our final test case consists of a model sky with Galaxy emission convolved with the sidelobe beam over the WSP. We present the output maps of both the standard and deconvolution methods in Fig. 4, where the maps are shown in ecliptic coordinates. One can see that the standard map contains a distorted image of the Galaxy and that the deconvolved map is virtually identical to the true map.

V. CONCLUSIONS AND FUTURE WORK

We have presented a deconvolution map-making method for data from scanning CMB telescopes. Our methods remove artifacts due to beam asymmetries and far sidelobes. We compare our technique with the standard map-making method and demonstrate that the true sky is recovered with greatly enhanced accuracy via the deconvolution method. Deconvolution map-making recovers features of the CMB sky on the smallest scale of the beam, thereby achieving a form of super-resolution imaging. This extracts more of the information content in CMB data sets.

One of the key difficulties encountered in deconvolution problems is that the systems of linear equations we need to solve are very nearly singular. We solve this problem by introducing a regularization method which allows us to

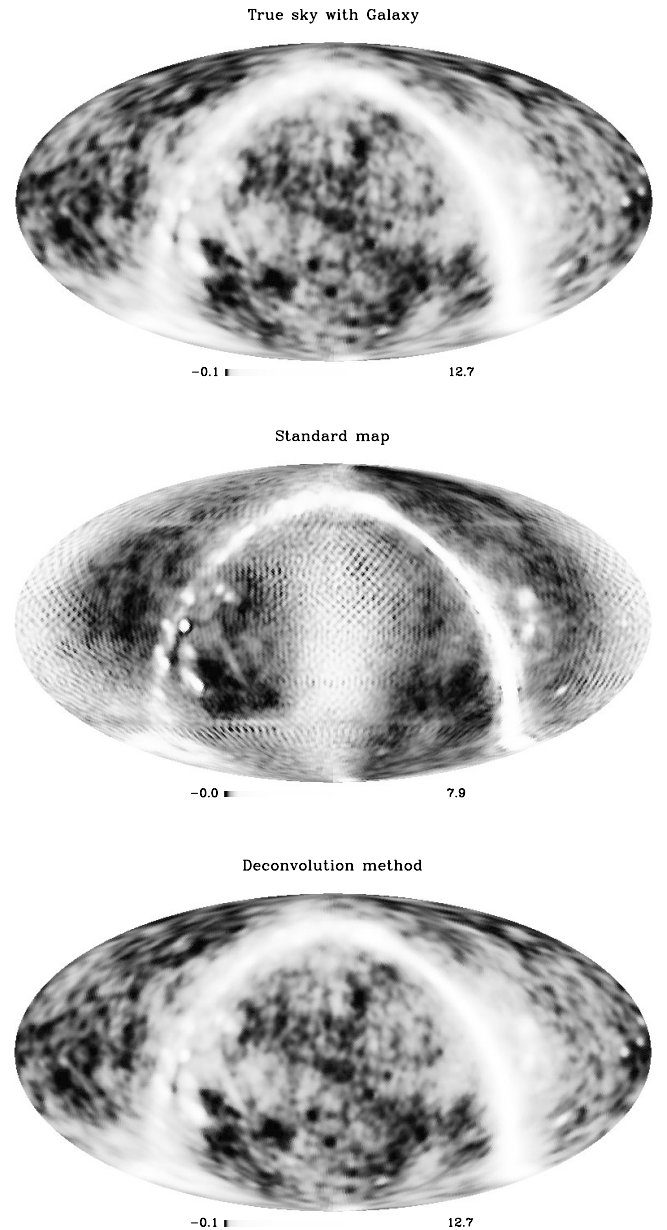


FIG. 4. Deconvolving the effects of a large sidelobe in simulated observations of the WMAP Ka-band map, using the coarsened WMAP scanning strategy described in the text. The top map is the input sky map, the middle map is the standard map-making result, and the bottom map is the deconvolved result.

solve the systems stably and recover maps at a uniform resolution and with an effective beam that is azimuthally symmetric and has a Gaussian profile.

We tested the convergence speed of two particular scanning strategies and found that the WMAP-like scan is superior to the basic scan in both rate of convergence and true-sky recovery. We hypothesize that this is due to the nature of the BSP, where the poles are the location of the only beam crossings and receive many more hits than the rest of the sky. In addition, our implementation of the BSP

had a smaller number of samples overall than our implementation of the WSP.

We have also shown the relevance of this algorithm to the WMAP mission by demonstrating its operation using a WMAP-like scanning strategy and a two-beam model which, while not differential, resembles the telescope orientations of the WMAP spacecraft. Our results underline the qualities of the WMAP scanning strategy compared to a BSP strategy for deconvolution map-making.

In order to decouple from issues that are not directly related to the optical performance of CMB instruments, we did not consider the effects of noise in our simulations. For a realistic assessment of the performance of our methods on real data this needs to be added. In particular, the choice of scale for the regularization kernel will depend on weighing the benefits of increased resolution against increased high frequency.

Recently, several groups have published CMB polarization results on the EE power spectrum [20]. WMAP released an all-sky analysis of the TE cross correlation in the first-year data [21]. We eagerly anticipate the large-angle polarization data from WMAP in the impending release of the second year data. In a few years' time, polarimeters on board the Planck satellite will collect data. Owing to the

difficulty of separating the two polarization modes, we expect that polarimetry experiments will be very sensitive to beam asymmetries and stray light. Future measurements of the tiny B-mode polarization will require both exquisite instruments and sophisticated analysis tools. We hope that the generalization of our methods to polarized map-making will be useful for making maps of the polarized microwave sky. It has already been shown [22] that little modification to the Wandelt-Górski method of fast all-sky convolution is needed to accommodate polarization data.

ACKNOWLEDGMENTS

We thank the performance engineering group at NCSA, in particular, Gregory Bauer, for stimulating conversations and help with optimizing our code. B. D. W. gratefully acknowledges funding from the National Center for Supercomputing Applications and the Center for Advanced Studies in the form of a Beckman Fellowship. This work was partially supported by NASA Contract No. JPL1236748, by the National Computational Science Alliance under AST300029N and the University of Illinois. We utilized the IBM pSeries 690 Cluster copper.ncsa.uiuc.edu.

-
- [1] J. V. Arnau, A. M. Aliaga, and D. Saez, *Astron. Astrophys.* **382**, 1138 (2002).
 - [2] C. Burigana *et al.*, *Astron. Astrophys.* **373**, 345 (2001).
 - [3] C. Barnes *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 51 (2003).
 - [4] B. Wandelt and K. Górski, *Phys. Rev. D* **63**, 123002 (2001).
 - [5] <http://map.gsfc.nasa.gov>
 - [6] <http://astro.estec.esa.nl/Planck>
 - [7] O. Dore *et al.*, *Astron. Astrophys.* **374**, 358 (2001).
 - [8] D. Yvon and F. Mayet, *Astron. Astrophys.* (to be published).
 - [9] P. Natoli *et al.*, *Astron. Astrophys.* **372**, 346 (2001).
 - [10] M. Tegmark, *Astrophys. J.* **480**, L87 (1997); X. Dupac and M. Giard, *Mon. Not. R. Astron. Soc.* **330**, 497 (2002); R. Stompor, *et al.*, in *Proceedings of the MPA/ESO/MPE Workshop held at Garching (astro-ph/0012418)*; see [7]; G. Efstathiou, *Mon. Not. R. Astron. Soc.* (to be published).
 - [11] C. Lineweaver *et al.*, *Astrophys. J.* **436**, 452 (1994).
 - [12] E. L. Wright, in *Proceedings of the IAS CMB Data Analysis Workshop, Princeton, 1996 (astro-ph/9612006)*; G. Hinshaw *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 63 (2003).
 - [13] J. V. Arnau and D. Saez, *New Astron. Rev.* **5**, 121 (2000).
 - [14] C. Burigana and D. Saez, *Astron. Astrophys.* **409**, 423 (2003).
 - [15] A. Challinor *et al.*, *Mon. Not. R. Astron. Soc.* **331**, 994 (2002).
 - [16] <http://crd.lbl.gov/~cmc/MADmap/doc/>
 - [17] J. Reid, in *Large Sparse Sets of Linear Equations*, edited by J. Reid (Academic, London, 1971) p. 231–254; W. Press *et al.*, *Numerical Recipes: the Art of Scientific Computing* (Cambridge University Press, Cambridge, England, 1992); J. R. Shewchuk, <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.ps> (1994).
 - [18] K. M. Gorski, E. Hivon, and B. D. Wandelt, in *Proceedings of the MPA/ESO Cosmology Conference*, edited by A. J. Banday, R. S. Sheth, and L. Da Costa (PrintPartners Ipskamp, Netherlands, 1998), pp. 37–42; K. M. Górski *et al.*, *astro-ph/0409513* [*Astrophys. J.* (to be published)]; <http://www.eso.org/science/healpix/>
 - [19] C. L. Bennett *et al.*, *Astrophys. J.* **583**, 1 (2003).
 - [20] E. M. Leitch *et al.*, *astro-ph/0409357* [*Astrophys. J.* (to be published)]; D. Barkats *et al.*, *astro-ph/0409380* [*Astrophys. J.* (to be published)]; S. T. Myers *et al.*, *AAS Meeting 204*, 91.07 (2004).
 - [21] A. Kogut *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 161 (2003).
 - [22] A. Challinor *et al.*, *Phys. Rev. D* **62**, 123002 (2000).