

COSMIC MICROWAVE BACKGROUND COMPONENT SEPARATION BY PARAMETER ESTIMATION

H. K. ERIKSEN,^{1,2,3} C. DICKINSON,⁴ C. R. LAWRENCE,³ C. BACCIGALUPI,⁵ A. J. BANDAY,⁶ K. M. GÓRSKI,^{3,7} F. K. HANSEN,^{1,2}
P. B. LILJE,^{1,2} E. PIERPAOLI,⁸ M. D. SEIFFERT,³ K. M. SMITH,⁹ AND K. VANDERLINDE¹⁰

Received 2005 August 11; accepted 2005 December 13

ABSTRACT

We propose a method for CMB component separation based on standard Bayesian parameter estimation techniques. We assume a parametric spectral model for each signal component and fit the corresponding parameters pixel by pixel in a two-stage process. First we fit for the full parameter set (e.g., component amplitudes and spectral indices) in low-resolution and high signal-to-noise ratio maps using MCMC, obtaining both best-fit values for each parameter and the associated uncertainty. The goodness of fit is approximated by a χ^2 statistic. Then we fix all nonlinear parameters at their low-resolution best-fit values and solve analytically for high-resolution component amplitude maps. This likelihood approach has many advantages: the fitted model may be chosen freely, and the method is therefore completely general; all assumptions are transparent; no restrictions on spatial variations of foreground properties are imposed; the results may be monitored by goodness-of-fit tests; and, most importantly, we obtain reliable error estimates on all estimated quantities. We apply the method to simulated *Planck* satellite and 6 year *WMAP* data based on realistic models and show that separation at the microkelvin level is indeed possible in these cases. We also outline how the foreground uncertainties may be rigorously propagated through to the CMB power spectrum and cosmological parameters using a Gibbs sampling technique.

Subject headings: cosmic microwave background — cosmology: observations — methods: numerical

1. INTRODUCTION

As experimental techniques improve rapidly and new high-sensitivity ground-based, balloon-borne, and satellite missions are being planned and implemented, the main problem in cosmic microwave background (CMB) measurement has changed from instrumental noise to separation of the cosmological CMB signal from noncosmological foreground signals. This problem will become even more important as our focus shifts from observations of temperature anisotropies to polarization anisotropies; while a simple template-fitting approach proved adequate for the first-year *Wilkinson Microwave Anisotropy Probe* (*WMAP*) analysis (Bennett et al. 2003a, 2003b; Hinshaw et al. 2003), no such hopes can be held for future polarization experiments.

While component separation is a difficult problem, it is not intractable. Since the cosmological CMB radiation follows a virtually perfect blackbody spectrum (Mather et al. 1999), whereas all known noncosmological signals have nonthermal spectra, it should be possible to disentangle the various contributions using spectral information. This fact motivated multiple frequencies on the *Cosmic Background Explorer* (*COBE*) Differential Microwave Radiometer (DMR) experiment (three bands between 31 and 90 GHz), the current *WMAP* experiment (five bands between 23 and 94 GHz), and the future *Planck* satellite experiment (nine bands between 30 and 857 GHz).

While the necessity of multifrequency observations has been recognized in the cosmological community for a long time, there has been uncertainty about how those observations should be used. Many different methods have been proposed, including the maximum entropy method (Barreiro et al. 2004; Bennett et al. 2003b; Hobson et al. 1998; Stolyarov et al. 2002, 2005), the internal linear combination method (Bennett et al. 2003b; Tegmark et al. 2003; Eriksen et al. 2004a), Wiener filtering (Bouchet & Gispert 1999; Tegmark & Efstathiou 1996), and the independent component analysis method (Maino et al. 2002, 2003; Donzelli et al. 2005). Some of these methods (e.g., Baccigalupi et al. 2004 and Stivoli et al. 2005) have been applied to polarization data. These methods all have their origins in the field of computational image processing.

In this paper, we advocate a more direct approach to the component separation problem, following in the footsteps of Brandt et al. (1994). We choose a physically motivated parametric model for each significant signal component and fit the free parameters by means of well-established algorithms, such as Markov Chain Monte Carlo (MCMC) and nonlinear searches. (For an alternative spectral matching algorithm, see Delabrouille et al. 2003.) The advantages of this approach are many: no assumptions about the spatial structure of the foreground properties are imposed; the method scales proportionally to the number of pixels; the results may be verified by means of goodness-of-fit tests; and, most importantly, the method yields accurate uncertainties for all estimated quantities.

¹ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, 0315 Oslo, Norway; h.k.k.eriksen@astro.uio.no, f.k.hansen@astro.uio.no, per.lilje@astro.uio.no.

² Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, 0316 Oslo, Norway.

³ Jet Propulsion Laboratory, Mail Stop 169-327, 4800 Oak Grove Drive, Pasadena, CA 91109; charles.r.lawrence@jpl.nasa.gov, Krzysztof.M.Gorski@jpl.nasa.gov, Michael.D.Seiffert@jpl.nasa.gov.

⁴ Department of Astronomy, Mail Stop 105-24, California Institute of Technology, 1200 East California Boulevard, Pasadena, CA, 91125; cdickins@astro.caltech.edu.

⁵ SISSA/ISAS, Astrophysics Sector, via Beirut, 4, 34014 Trieste; INFN, Sezione di Trieste, via Valerio 2, 34014 Trieste, Italy; Institut für Theoretische Astrophysik, Universität Heidelberg, Albert-Berle Strasse 2, 69120 Heidelberg, Germany; bacci@sisssa.it.

⁶ Max-Planck-Institut für Astrophysik, Karl-Schwarzschildstrasse 1, Postfach 1317, 85741 Garching, Germany; banday@MPA-Garching.MPG.DE.

⁷ Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warsaw, Poland; California Institute of Technology, Pasadena, CA 91125.

⁸ Theoretical Astrophysics, Mail Stop 130-33, California Institute of Technology, Pasadena, CA 91125; pierpa@caltech.edu.

⁹ Department of Physics, University of Chicago, Chicago, IL 60637; kmsmith@uchicago.edu.

¹⁰ Kavli Institute for Cosmological Physics and Enrico Fermi Institute, University of Chicago, Chicago, IL 60637; kvand@uchicago.edu.

We begin by developing a simple algorithm that is able to analyze real-world data in the presence of realistic noise. Whereas Brandt et al. (1994) relied exclusively on nonlinear fitting and was therefore quite unstable with respect to noise, we take advantage of more recent developments (most importantly MCMC) to make the algorithm both more robust and also to produce accurate errors. We also suggest a procedure for propagating these results into final data products, namely, the CMB power spectrum and cosmological parameters, by means of a Gibbs sampling method.

After establishing the algorithm, we apply it to realistic simulations corresponding to the future *Planck* and 6 year *WMAP* data, taking into account the predicted noise distributions of each experiment. We attempt to model the foregrounds as accurately as possible, given our current understanding of the involved foregrounds.

2. FORMULATION OF THE COMPONENT SEPARATION PROBLEM

Assume that the observed data take the form of a multifrequency set of sky maps, \mathbf{d}_ν , each of which may be written in the form

$$\mathbf{d}_\nu = \mathbf{A}\mathbf{s}_\nu + \mathbf{n}_\nu, \quad (1)$$

where ν identifies frequency bands and sky maps 1 through N , \mathbf{s}_ν is the true sky signal at the corresponding frequency band, \mathbf{n}_ν is instrumental noise, and \mathbf{A} denotes convolution with the instrument beam.

We assume the noise component \mathbf{n}_ν to be Gaussian distributed with vanishing mean and variance $\sigma_\nu^2(p)$, where p is the pixel number. Thus, the noise is uncorrelated both between pixels and between frequency channels, but spatial variations in the variance are allowed.

Note that there is no frequency index on the beam operator in equation (1), indicating that all channels are assumed to have the same beam response. For multiresolution experiments, this implies that the sky maps must be smoothed to a common resolution prior to analysis. Equation (1) is no longer strictly valid, since the noise is then correlated; however, in practice it works reasonably well to approximate the noise term as uncorrelated between pixels, with rms levels determined by Monte Carlo simulations of processed noise.

The signal \mathbf{s}_ν may be decomposed into a sum of components, $\mathbf{s}_\nu = \sum_i \mathbf{s}_\nu^i$, in which the most important ones are the cosmological CMB signal and three galactic foregrounds: synchrotron, free-free, and dust emission. Compact (unresolved) Galactic and extragalactic sources could also be included in the list, but these are more conveniently detected by other methods, such as wavelets (e.g., Vielva et al. 2003). In this paper we consider compact source removal as a part of the preprocessing stage and assume that resolved sources are either masked, median filtered, or fitted prior to analysis. Thus, we include only diffuse foregrounds in the following, but note that more work is needed on this issue.

Assume also that the frequency spectrum of each signal component may be parameterized by a small number of free parameters, for instance an amplitude and a spectral index $S = A(\nu/\nu_0)^\beta$. Given a set of multifrequency CMB sky maps as described above and a parametric signal model $S_\nu(\theta)$ with free parameters θ , we now establish both a point estimate $\hat{\theta}$ for the free parameters and the corresponding uncertainties. To do so, we use standard Bayesian parameter estimation methods.

Our goal is to establish the posterior distribution for the parameters given the data, $P(\theta|\mathbf{d})$. Using Bayes's theorem, this may be written as

$$P(\theta|\mathbf{d}) \propto P(\mathbf{d}|\theta) P(\theta) = \mathcal{L}(\theta)P(\theta), \quad (2)$$

where $\mathcal{L}(\theta) = P(\mathbf{d}|\theta)$ is the likelihood and $P(\theta)$ is a prior. In this paper, we adopt uniform priors between two (not necessarily finite) limits for all parameters, although it is in general straightforward to implement nonuniform priors if desired. (However, the precise interpretation of the goodness-of-fit test will be affected by doing so.) Specifically, we impose a positivity constraint on the foreground amplitudes ($A > 0$) and sufficiently generous limits on the spectral indices to never exclude physically realizable values ($-4 < \beta_s < -2.2$ for synchrotron and $1 < \beta_d < 3$ for thermal dust).

Since the noise is assumed to be Gaussian distributed and uncorrelated between pixels, the likelihood reduces to the usual χ^2 , independent between pixels,¹¹

$$\ln \mathcal{L} = -\frac{1}{2} \sum_{\nu=1}^N \left[\frac{d_\nu - S_\nu(\theta)}{\sigma_\nu} \right]^2 = -\frac{1}{2} \chi^2. \quad (3)$$

The problem of component separation is thus reduced to mapping out the posterior as given above by some numerical technique, for instance grid computation, MCMC, or nonlinear searches.

We illustrate the procedure in Figure 1, in which the results from an analysis of one arbitrarily chosen pixel are shown (see § 6 for details). The observed data points are marked by black circles, and the fitted components are shown as smooth, colored curves. The dashed black curve shows the sum of all components. Four signal components are included in this model, CMB, synchrotron, free-free, and thermal dust emission. The parametric models for the foregrounds are perfect power laws for synchrotron and free-free emission and a one-component model for dust (see eq. [14]).

Although we focus on temperature anisotropies in the present paper, the method can handle polarization anisotropies equally well. In that case, the χ^2 takes the form

$$\chi^2 = \sum_{\nu=1}^N [\bar{d}_\nu - \bar{S}_\nu(\theta)]^T \mathbf{N}_\nu^{-1} [\bar{d}_\nu - \bar{S}_\nu(\theta)], \quad (4)$$

where

$$\bar{d}_\nu = \begin{pmatrix} d_\nu^I \\ d_\nu^Q \\ d_\nu^U \end{pmatrix}, \quad \bar{S}_\nu(\theta) = \begin{pmatrix} S_\nu^I(\theta) \\ S_\nu^Q(\theta) \\ S_\nu^U(\theta) \end{pmatrix} \quad (5)$$

are the Stoke's I , Q , and U parameters for the data and model, respectively, and \mathbf{N}_ν is the 3×3 (I , Q , U) noise correlation matrix for the pixel. (We still assume uncorrelated noise between pixels and frequency bands but not between the three Stoke's parameters for each individual pixel.)

3. PARAMETRIC MODEL FITS

In this section we introduce the parametric model that we fit to the data, starting with a review of the currently favored parametric

¹¹ In practice this is not strictly correct, since the CMB and foreground components are indeed correlated between pixels, and we will at some point work with smoothed sky maps; nevertheless, it is a good approximation for component separation purposes. Further, as described in § 5.6, at least spatial CMB correlations may be taken into account by Gibbs sampling.

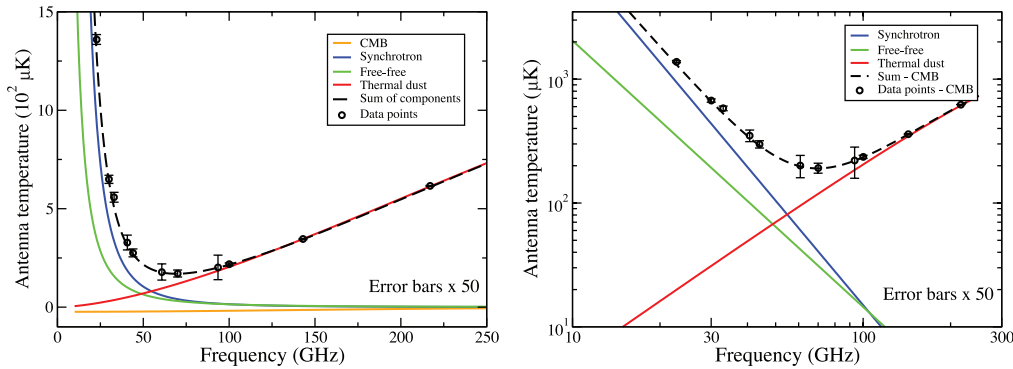


FIG. 1.—Component separation using multifrequency measurements (*left*: linear units; *right*: logarithmic units). Most signal components has a well-defined frequency spectrum that may be parameterized by one or a few parameters, and component separation may therefore be viewed as a standard parameter estimation problem. The example shown here is based on one single pixel in a simulated data set corresponding to the 6 year *WMAP* and the *Planck* experiments, as discussed in § 6. The error bars on the data points are multiplied by a factor of 50 in order to make them visible on this scale. (Due to modeling errors, this particular fit has a χ^2 of 44, and with 5 dof, it is rejected at the 99.9999% confidence level.)

signal models s_i^j for each signal component. We emphasize that the procedure as such is general, and any parametric model may be included in the analysis. Then we discuss how to take into account the effect of nonzero instrumental bandwidths and discuss implementation details that lead to more transparent computer code. Note that the parametric signal models we adopt for *fitting* the various components may differ from the parametric models that are used for *modeling* the components later in § 6.1, as will certainly be the case in dealing with real data.

3.1. Signal Components

CMB anisotropies.—The cosmological background component, due to its blackbody nature (Mather et al. 1999), is most easily characterized by its thermodynamic temperature T_{CMB} or, equivalently, the anisotropy temperature $\Delta T_{\text{CMB}} = T_{\text{CMB}} - T_0$, where $T_0 = 2.725$ K (Fixsen & Mather 2002) is the average CMB temperature. However, as discussed later, the foreground components are more easily described in terms of antenna temperatures, and therefore we choose to convert the CMB signal accordingly. The CMB signal model then reads

$$s_{\text{CMB}}(\nu) = \Delta T_{\text{CMB}} \frac{x^2 e^x}{(e^x - 1)^2}, \quad (6)$$

where $x = h\nu/kT_0$, and h is the Planck constant and k is Boltzmann's constant.

Synchrotron emission.—Synchrotron emission from the Galaxy originates in relativistic cosmic-ray (CR) electrons spiraling in the Galactic magnetic field. The morphology of the observed emission depends on the distribution of the relativistic electrons in the Galaxy and the Galactic magnetic field structure. In the Galactic plane, the latter exhibits a large-scale ordering with the field parallel to the spiral arms (the regular component). Superposed on this is real small-scale structure (the irregular component) which shows variations between the arm and inter-arm regions and with gas phase. The regular and irregular components seem to be of comparable magnitude. At high latitudes, there is a contribution from the Galactic halo and specific nearby structures (e.g., the North Polar Spur). Variations in the frequency spectral index of the synchrotron continuum emission arise from variations in the CR electron energy spectrum, which has a range of distributions depending on age and the environment of origin (e.g., supernova explosions or diffuse shocks in the interstellar medium).

The synchrotron emission may be accurately modeled by means of a simple power law over a considerable range of frequencies,

$$s_s(\nu) = A_s \left(\frac{\nu}{\nu_{0,s}} \right)^{\beta_s}. \quad (7)$$

Here A_s is the synchrotron amplitude¹² (measured in antenna temperature μK) at some reference frequency $\nu_{0,s}$, and β_s is the synchrotron spectral index. Since the spectral index varies with both frequency and position on the sky, at least two free parameters are required to describe the synchrotron emission properly in a given direction.

Lawson et al. (1987) studied the spectral index variation based on low-frequency radio surveys and found that the brighter regions away from the Galactic plane have typical values of β at 100 and 800 MHz of 2.55 and 2.8, respectively. Reich & Reich (1988) used radio continuum surveys of the northern sky at 408 and 1420 MHz to demonstrate a range of spectral index values between 2.3 and 3.0, with a typical dispersion $\Delta\beta = \pm 0.15$. The steepest spectra were observed toward the North Polar Spur, and there was a flattening in spectral index toward higher latitudes in the Galactic anticenter direction. Such behavior has been confirmed over the full sky by Reich et al. (2003), who find that spectral flattening is particularly pronounced in the southern hemisphere

At higher frequencies, the brightness temperature spectral index is expected to steepen by ~ 0.5 due to electron energy losses (Platania et al. 1998). Banday et al. (2003) derived a mean spectral index between 408 MHz and 19.2 GHz from the Cottingham (1987) survey and between 31.5, 53, and 90 GHz from the *COBE* DMR data. The steep spectral index of ~ 3.1 for Galactic latitudes $|b| > 15^\circ$ is consistent with expectations. Bennett et al. (2003b) claim that the spectral break occurs near the K band. Spectral indices above 10 GHz are likely between 2.7 to 3.2.

Free-free emission.—Free-free emission is the bremsstrahlung radiation resulting from the Coulomb interaction between free electrons and ions in the Galaxy. Free-free emission is weaker than synchrotron emission at frequencies below ~ 1 GHz over most regions of the sky, with exceptions in the Galactic

¹² Note that italic font A is used for component amplitudes, while bold font \mathbf{A} is used for beam convolution.

ridge and the brighter features of the local Gould Belt system (Dickinson et al. 2003). Free-free almost never dominates at high latitudes in any frequency band, and is therefore difficult both to observe and to simulate.

Using the relations summarized in Dickinson et al. (2003), it can be shown that the brightness temperature of the free-free emission is described by the relation

$$T_{\text{ff,b}} \propto \nu^{-2} T_e^{-0.5} [\ln(0.04995\nu^{-1}) + 1.5 \ln T_e] \quad (8)$$

where T_e is the electron temperature.

Shaver et al. (1983) used radio recombination lines to show that the electron temperature of H II regions at the Galactocentric radius of the Sun ($R_0 = 8.5$ kpc) is 7200 ± 1200 K. Paladini et al. (2005) found similar results from a larger sample containing many weaker sources. At high Galactic latitudes, the ionized hydrogen typically will be within ~ 1 kpc of the Sun (Dickinson et al. 2003), and we expect the electron temperature to be in the range $T_e = 7000\text{--}8000$ K, although it is possible that the diffuse emission at a given Galactocentric distance may differ from that of the higher density H II regions on the plane.

The frequency dependence is therefore well constrained, with an effective spectral index of $\beta_{\text{ff}} = -2.14$ at the frequencies of interest, the range over 10–100 GHz being of order -2.1 to -2.2 and steepening still further to -2.3 at hundreds of gigahertz.

Finkbeiner (2004) has analyzed the *WMAP* data and found a significant component with a free-free–like spectrum within 30° of the Galactic center. The component is uncorrelated with H α emission and may be indicative of hot ($\sim 10^6$ K) gas. We do not attempt to include such a component here.

In this paper, we model free-free emission as a simple power law with a fixed spectral index,

$$s_{\text{ff}}(\nu) = A_{\text{ff}} \left(\frac{\nu}{\nu_{0,\text{ff}}} \right)^{-2.14}. \quad (9)$$

Thus, only one parameter is required for free-free emission. Future experiments may need to estimate the electron temperature directly from the data. In that case equation (8) should be used directly, at the cost of introducing one extra free parameter into the fit.

Thermal dust emission.—The thermal dust emission that contributes to the frequencies of interest for CMB analysis arises from grains large enough to be in thermal equilibrium with the interstellar radiation field, and it is known from analysis of the *IRAS* and *COBE* DIRBE data to peak at a wavelength of approximately 140 μm . At higher frequencies, there is a contribution from the optically active modes of PAH molecules, but these are not of interest here.

Currently preferred dust emission models (Finkbeiner et al. 1999) extrapolate from high-frequency *COBE* FIRAS and DIRBE observations to CMB frequencies using combinations of modified blackbody fits and accounting for dust temperature variations. Such fits approximate the integrated contributions to the emission from multiple components of dust, i.e., with different grain properties (chemical composition and size) and equilibrium temperatures. The best-fit model (model 8 of Finkbeiner et al. 1999) assumes two main components:

$$s_{\text{d}}(\nu) = F \left(\frac{\nu}{\nu_{0,\text{d}}} \right)^{\beta_{\text{d}}(\nu)}. \quad (10)$$

Here F represents the combined *COBE* DIRBE and *IRAS* template (Schlegel et al. 1998), $\nu_{0,\text{d}} = 3000$ GHz, and $\beta_{\text{d}}(\nu)$ is dependent on the frequency, as discussed above,

$$\beta_{\text{d}}(\nu) = \frac{\log[d(\nu)/d(\nu_{0,\text{d}})]}{\log(\nu/\nu_{0,\text{d}})}, \quad (11)$$

$$d(\nu) = \frac{q_1}{q_2} f_1 \left(\frac{\nu}{\nu_{0,\text{d}}} \right)^{3+\alpha_1} \frac{1}{e^{(h\nu/kT_1)} - 1} \quad (12)$$

$$+ f_2 \left(\frac{\nu}{\nu_{0,\text{d}}} \right)^{3+\alpha_2} \frac{1}{e^{(h\nu/kT_2)} - 1}, \quad (13)$$

with best-fit parameters $f_1 = 0.0363$, $q_1/q_2 = 13$, $\alpha_1 = 1.67$, $\alpha_2 = 2.70$, $T_1 = 9.4$ K, $T_2 = 16.2$ K, and $f_2 = 1 - f_1$ (Finkbeiner et al. 1999).

In principle, these equations may serve as our parametric model for fitting the dust emission spectrum. However, few (current or future) CMB experiments have sufficient power to constrain six parameters for dust alone, and simplifications are therefore unavoidable. Rather than fitting the full form, as given above, we therefore choose the simpler model 3 of Finkbeiner et al. (1999), setting $f_1 = 1$ and $T_1 = 18.1$ K, but letting α_1 vary freely. Equation (10) may then be simplified to

$$s_{\text{d}}(\nu) = A_{\text{d}} \frac{\nu}{e^{(h\nu/kT_1)} - 1} \frac{e^{(h\nu_{\text{d},0}/kT_1)} - 1}{\nu_{\text{d},0}} \left(\frac{\nu}{\nu_{\text{d},0}} \right)^{\beta_{\text{d}}}, \quad (14)$$

where A_{d} is the thermal dust amplitude at a reference frequency $\nu_{\text{d},0}$ and β_{d} is a free parameter. Thus, the fitted model is a power law modulated by a slowly decreasing function of order unity over the frequencies of interest.

Such a parameterization does not allow for the spectral break at approximately 500 GHz in the *COBE* FIRAS, which may reflect either the emissivity of different grain components, frequency dependence of the emissivity of the dominant grain component, or possibly a population of cold dust grains mixed with the warmer dust (Reach et al. 1995). However, at 500 GHz the signal of dust anisotropies is so strong compared with CMB anisotropies that it is useless for purposes of component separation.

Anomalous dust emission.—Cross-correlation of the *COBE* DMR data with the DIRBE map of thermal dust emission at 140 μm in Kogut et al. (1996) revealed an anomalous component with rising spectrum from 53 to 31.5 GHz. Banday et al. (2003), again using the DMR data together with a survey at 19.2 GHz and independently the *WMAP* team (Bennett et al. 2003b), suggested that this component was well described by a power-law spectrum with index -2.5 for frequencies in the range $\sim 20\text{--}60$ GHz. The latter proposed that the emission originates in star-forming regions close to the Galactic plane. However, the favored model to explain this anomalous dust-correlated component is in terms of the rotational emission from very small grains.

Draine & Lazarian (1998) have developed a three-component model of this “spinning dust,” which contains contributions from the three phases of the interstellar medium—the cold neutral medium, the warm neutral medium, and the warm ionized medium. The characteristic spectral behavior of the model includes a rising spectrum up to a turnover in the range 10–20 GHz, then a rapidly falling spectrum that can be characterized by an effective spectral index in excess of 3 beyond 30 GHz.

Recent observations by Finkbeiner et al. (2004) of dust-correlated emission outside of H II regions between 8 and 14 GHz shows a rising spectral slope and amplitude far exceeding that associated with thermal dust emission. More importantly, Watson

et al. (2005) show that observations of the Perseus molecular cloud made between 11 and 17 GHz and augmented with the *WMAP* data can be adequately fitted by a spinning-dust model. Nevertheless, although the case is compelling for describing the anomalous emission by such models, the detailed morphology of the emission remains uncertain, and no unambiguous template to trace it exists. It is also possible that the emission may be confined to specific clouds at relatively low Galactic latitude, leading to a more patchy distribution than for the diffuse thermal dust contribution with some additional sensitivity to environment. Indeed, Lagache (2003) presents evidence that the excess emission is associated with small transiently heated dust particles, which may be destroyed under certain physical conditions. Given these uncertainties, we do not include this foreground component in our simulations.

3.2. Nonzero Bandwidths

The previous sections describe the basic behavior of each signal component as a function of well-defined frequencies. However, real experiments integrate over a range of frequencies (a “frequency band”), typically with unequal weights, and the observed signal strength does not equal that given by the central frequency alone.

We take into account this effect through the concept of an “effective frequency” ν_{eff} , defined by

$$S(\nu_{\text{eff}}) = \int f(\nu) S_\nu d\nu, \quad (15)$$

where S_ν is the frequency spectrum of the signal, and $f(\nu)$ is the frequency response profile of the detector. Thus, the spectrum at the effective frequency equals the average over the frequency band. The advantage of doing this is simply that computationally expensive integrals are replaced by single point computations.

In this paper we assume for simplicity that all frequency response functions correspond to flat bandpass filters with sharp frequency cutoffs at ν_a and ν_b . For simple power-law models, such as those of synchrotron and free-free emission as described above, the effective frequency of a signal component with spectral index β is then given by

$$\nu_{\text{eff}} = \left(\frac{1}{\beta + 1} \frac{\nu_b^{\beta+1} - \nu_a^{\beta+1}}{\nu_b - \nu_a} \right)^{1/\beta}. \quad (16)$$

For more complicated spectra, equation (15) must be solved numerically. Fortunately, it is straightforward to precompute a grid of the effective frequencies prior to the full analysis, since they only depend on the frequency scalings and not the component amplitudes, and computational speed is not compromised.

3.3. Implementation Details

To simplify the computer code, it is convenient to introduce some general notation. For instance, if we can write all signal models in a common form, we do not have to consider a list of special cases, but instead can handle all cases with the same code.

Indeed, all frequency spectra discussed above may be written in a common form, namely, that of a power law modulated by an arbitrary frequency-dependent function,

$$s_\nu(p) = \sum_{i=1}^{N_{\text{comp}}} A_i(p) c_{i,\nu_{\text{eff},i}} \left(\frac{\nu_{\text{eff},i}}{\nu_{0,i}} \right)^{\beta_i(p)}. \quad (17)$$

Here $A_i(p)$ and $\beta_i(p)$ are the “free” amplitude and index parameters for component i in each pixel p , and $c_{i,\nu_{\text{eff},i}}$ is an arbitrary function only dependent on frequency.

Specifically, synchrotron and free-free emission are included simply by setting $c_{i,\nu} = 1$, while the CMB component is defined by $\beta_{\text{CMB}} \equiv 0$ and $c_{\text{cmb},\nu} = x^2 e^x / (e^x - 1)^2$, as discussed above. For dust, $c_{d,\nu}$ is given by equation (14).

Even anomalous dust could be included within this notation. One option is simply to tailor the correction factors $c_{i,\nu}$ to match the predicted spinning-dust spectrum (Draine & Lazarian 1998), fix the corresponding spectral index at zero, and then fit for the amplitude only. Another is to merge the spectrum with that of the thermal dust emission and thereby enforce identical spatial templates.

4. SUMMARY OF PREVIOUS RESULTS

This paper may be seen as a natural continuation of the work started by Brandt et al. (1994), who considered how well future experiments could reconstruct the CMB signal in presence of noise and foregrounds. Their approach, parameter estimation, was the same as ours; however, they relied solely on maximum likelihood estimation (i.e., nonlinear fitting), and their results were therefore less stable with respect to noise than the ones we present here, as is shown below. Nevertheless, several of the conclusions drawn by Brandt et al. (1994) are still valid for our work and are well worth repeating:

1. The number of frequencies must equal or exceed the number of fitted parameters, or else the problem is mathematically degenerate. This is obvious but not trivial: no experiment to date has had the minimum number of frequencies required to separate CMB, synchrotron, free-free, and dust fluctuations, even in their simplest form.
2. One should attempt to reduce the number of free parameters in the problem, as this gives greater stability with respect to noise. Seemingly gross simplifications, such as approximating both synchrotron and free-free emission by a single power law, can often yield improvements in the reconstruction.
3. It is usually advantageous to fit spectral parameters to reduced-resolution and low-noise data and then solve for the amplitudes in the full-resolution data, fixing the indices at the smoothed values.
4. Due to the similarity between the synchrotron and free-free emission, better results are obtained whenever the latter is not a significant contaminant. Thus, if the free-free contamination could be constrained by radiation physics knowledge, it would be well worth trying.

While working on the present analysis, we have reproduced all of these conclusions, and most of them have been taken into account when establishing the prescription described below. However, we do not elaborate further on these issues here, but rather refer the interested reader to Brandt et al. (1994) and come back to the above points as they are needed in the analysis.

5. METHOD

In this section we propose an algorithm for solving the parameter estimation problem with sufficient speed and accuracy to be useful for practical analysis of current and future data. Each step of the algorithm consists of well-established methods, and the approach should seem quite familiar.

5.1. Overview

The goal is to establish both a point estimate of all interesting parameters and their uncertainties. Our prescription for doing so is as follows:

1. If required by noise levels or computational resources, downgrade all sky maps both in pixel and beam resolution.
2. For each low-resolution pixel:
 - a) choose a parametric model,
 - b) solve for all parameters jointly by MCMC,
 - c) estimate nonlinear parameters and corresponding uncertainties by marginalizing over all other parameters,
 - d) find the goodness of fit.
3. For each high-resolution pixel within a low-resolution pixel:
 - a) *either* (fast but approximate analytical approach)
 - i) fix the nonlinear parameters at the low-resolution pixel values,
 - ii) find maximum-likelihood estimates for all linear parameters (i.e., component amplitudes) by solving a linear equation,
 - iii) find corresponding uncertainties by analytic error propagation formula,
 - iv) estimate the goodness of fit.
 - b) *or* (exact but expensive Gibbs sampling approach)
 - i) use the low-resolution MCMC distributions to sample nonlinear parameters numerically,
 - ii) given a set of nonlinear parameters, sample amplitudes from their corresponding Gaussian distribution,
 - iii) given foregrounds, sample CMB sky map and power spectrum.

The route outlined in step 3*b* holds the promise of a complete solution to the foreground problem in a CMB context, since, if successful, the foreground uncertainties are propagated all the way from noisy observations to the CMB power spectrum and cosmological parameters. However, in this paper we only present the basic ideas and leave the details for a future more comprehensive study. All high-resolution results presented in the following are thus based on the analytical approach described in 3*a*.

5.2. Nonlinear Parameters and Large-Scale Smoothing

One of the main themes of Brandt et al. (1994) was the instability of a nonlinear fit with respect to noise. This is not hard to understand. In estimating multiple parameters from a limited number of frequency bands, the maximum likelihood point may easily slide along some degeneracy ridge on the likelihood surface in the presence of realistic noise. For all currently planned experiments, additional degree-scale smoothing is a requirement in order to reach reconstruction errors at the microkelvin level.

Another and more practical issue stems from the fact that our main algorithm relies on MCMC analysis of each individual pixel. This takes on the order of 100 s per pixel. Even though the algorithm parallelizes trivially because the pixels are analyzed individually, a complete analysis at full *Planck* resolution (~ 50 million pixels) would be unfeasible.

It is important to realize, however, that a full MCMC analysis is required only for estimating nonlinear parameters, such as spectral indices or dust temperatures. If all parameters in the problem are linear (i.e., component amplitudes), an analytical computation is equally good. We therefore compute the complete probability distributions from reduced-resolution maps, fix the high-resolution nonlinear parameters at the corresponding low-resolution values, and then solve for the high-resolution component amplitudes with alternative methods, for instance analytically or by Gibbs sampling.

When adequate data on foregrounds are in hand, the validity of this smoothing can be tested. If preliminary indications turn out to be true, and spectral indices vary more slowly on the sky

than amplitudes, the smoothing process will not lead to significant loss of information. In any case, the smoothing scale can be optimized within the bounds of computational resources.

5.3. Model Selection

A second main theme of Brandt et al. (1994) was the importance of model selection. They clearly demonstrated that a large number of parameters does not necessarily yield a better CMB reconstruction. Quite the contrary, too many parameters often yield unphysical results. In general, one should never fit more parameters than required by the data.

In principle, it would be useful to have an automated prescription to identify the optimal model for a given pixel. To some extent, such a procedure is provided by means of the so-called information criterion (IC), an idea that was introduced to CMB analysis by Liddle (2004). The fundamental idea in this approach is not to maximize the likelihood alone, but rather to minimize the IC, defined as follows:

$$\text{IC} = -2 \ln \mathcal{L} + \alpha k. \quad (18)$$

Here k is the number of parameters in the fit and α is a penalty factor. (Two useful choices for α are $\alpha_A = 2$, the Akaike information criterion, and $\alpha_B = \ln N$, the Bayesian information criterion, N being the number of data points.) Within this framework, a new parameter must prove its usefulness by returning a significant improvement in the χ^2 fit to be included in the model.

We implemented this approach in our codes and obtained reasonable results. However, the model sky map had a clear tendency to be patchy and not necessarily well correlated with physical structures. In the current implementation, we therefore only use the information criterion approach to inform our model choices and tailor the model map manually according to known structures. For instance, in the example given in § 6, we use a full four-component model (CMB, synchrotron, free-free, and dust) inside an expanded Kp0 Galactic cut (Bennett et al. 2003b), as well as in a few selected patches (e.g., the LMC), but we ignore free-free otherwise. However, we expect that the information criterion approach may be developed further and should be a valuable tool for future experiments.

Realistically, model selection is likely to be an iterative process, as is demonstrated in the worked example of § 6. Typically, one constructs an initial physically motivated model map and performs the analysis with that model set. Based on the results, one then evaluates the goodness of fit for each pixel and compares the results with the appropriate distribution function. If the agreement is poor, the model set can be modified, or the offending pixels may be rejected from further analysis.

Formally speaking, evaluating the goodness of fit in a Bayesian setting is a more delicate issue than in the frequentist case, due to the presence of priors. The usual frequentist approach is simply to compute the χ^2 (as given by eq. [3]) and compare this with a χ^2 distribution with the appropriate number of degrees of freedom (dof; the number of frequencies minus the number of free parameters). For nonuniform priors or nonlinear parameters, this is not strictly correct in the Bayesian case, since the reference distribution may be different from a χ^2 density.

However, for the purposes of this paper, such concerns are of little importance. In practice, the preferred model map will be established by means of a combination of external data (e.g., foreground templates) and some estimate of the goodness of fit following the first analysis. However, the model selection process is not likely to be a search for $2.5\text{--}3\sigma$ outliers. Therefore, for our

crude purposes, the naive χ^2 statistic as defined by equation (3) is adequate.

In practice, model selection is likely to be a relatively straightforward (although somewhat tedious) process for future experiments. In the present paper we are satisfied with a very simple choice of models, based on established sky masks.

5.4. Parameter Estimation by MCMC

By now, parameter estimation by Markov Chain Monte Carlo is a well established technique within the CMB community, with its most visible application being estimation of cosmological parameters from the CMB power spectrum. To our knowledge, it has not yet been applied to component separation, and we therefore briefly describe the algorithm here. For more details, we refer the interested reader to, e.g., Lewis & Bridle (2002) or Verde et al. (2003).

5.4.1. Algorithm

Suppose we want to estimate a set of parameters and corresponding uncertainties from a set of observed data and that we know how to compute the likelihood given an arbitrary combination of parameters. The MCMC algorithm is then given by the following simple steps. (Note that because of the uniform priors adopted in this paper, the posterior is numerically equal to the likelihood.)

1. Choose any initial point in parameter space, θ_0 , and compute the corresponding likelihood, $\mathcal{L}_0 = \mathcal{L}(\theta_0)$.
2. Define a *stochastic* function f that, given parameters θ_i , returns a new set of parameters $\theta_{i+1} = f(\theta_i)$.
3. Compute θ_{i+1} given by f and the corresponding likelihood, $\mathcal{L}_{i+1} = \mathcal{L}(\theta_{i+1})$.
4. Set $\theta_{i+1} = \theta_i$ (i.e., reject the proposal) with probability $p = 1 - \min(\mathcal{L}_{i+1}/\mathcal{L}_i, 1)$.
5. Go to step 3 and iterate as long as necessary.

This procedure returns a chain of parameter samples θ_i , $i = 1, \dots, N_{\text{samples}}$, and their multidimensional histogram equals the posterior in the limit of an infinite number of samples.

Some intuition for the process may be gained by noticing the form of the “jump probability” given in step 4 of the algorithm: If the likelihood of the proposed point is larger than that of the old point, we never reject the proposed point; we always move toward more likely solutions when proposed. However, if the likelihood is smaller, we still accept the proposed point with probability $p = \mathcal{L}_{i+1}/\mathcal{L}_i$. This guarantees that we spend most of the time around the peak position but still explore less likely points. Indeed, it may be proven that the time spent at a given point in parameter space is proportional to the posterior itself.

5.4.2. Automated MCMC in Practice

In practice, there are several problems connected to MCMC parameter estimation; usually, most of these may be identified (and often solved) by simple visual inspection of the Markov chains. However, since we want to analyze many thousands of independent pixels, finding automated and yet robust solutions to the same problems is of critical importance; a solution that works in 99% of all cases is not good enough.

Burn-in.—Although it is true that the initial guess may be chosen arbitrarily and that the chain eventually will burn in to the right solution, it is difficult to construct a truly reliable automated prescription for determining when burn-in has occurred. Moreover, since computational speed is of critical importance, it is not acceptable to spend a long time in the burn-in phase. For both of these reasons, we choose to initialize the chains at the maximum

likelihood point, which we find using a standard nonlinear fitting algorithm. (We have found that a sequential quadratic programming [SQP] method works very well for this task.)

Proposal function.—In step 2 of the algorithm, we must establish a proposal function f . A simple example of such a function would be $\theta_{i+1}^j = \theta_i^j + \delta\theta^j\eta^j$, where j is a parameter index, $\delta\theta^j$ is a predefined rms step size, and η^j is a Gaussian stochastic variate of zero mean and unit variance. However, since most parameters of interest are usually strongly correlated, this choice is quite inefficient.

Our current best solution is to run a preliminary chain (using the univariate Gaussian proposal function described above, with manually set step sizes) and compute the covariance matrix $C_{ij} = \langle \delta\theta^i \delta\theta^j \rangle$ of the resulting samples. We then Cholesky-decompose this matrix, $\mathbf{C} = \mathbf{LL}^T$, and define our new proposal function to be $\theta_{i+1} = \theta_i + \alpha\mathbf{L}\boldsymbol{\eta}$, where $\boldsymbol{\eta}$ now is a vector of Gaussian variates and α is an overall scale factor, typically initialized at ~ 0.3 . This ensures that the proposed samples have approximately the correct covariance structure, and the overall sampling efficiency is thereby greatly improved.

To avoid too large or too small step sizes, we also impose the requirements that the acceptance ratio (the ratio of accepted to rejected proposals) must be higher than 5% and lower than 80%. If one of these two criteria is violated, we divide or multiply α by 2 and restart the MCMC analysis.

Convergence.—Finally, we must decide when a sufficient number of samples have been accumulated. No general solutions are available. We have adapted a good working solution proposed by Gelman & Rubin (1992), as follows. Run m independent MCMC chains in parallel for the same pixel, each producing n samples. Compute the following quantities,

$$W = \frac{1}{m(n-1)} \sum_{j=1}^m \sum_{i=1}^n (\theta_i^j - \hat{\theta}_i^j)^2, \quad (19)$$

$$B = \frac{n}{m-1} \sum_{j=1}^m (\theta^j - \hat{\theta})^2, \quad (20)$$

$$V = \left(1 - \frac{1}{n}\right)W + \frac{1}{n}B, \quad (21)$$

$$R = \frac{V}{W}, \quad (22)$$

where $\hat{\theta}$ is the average overall $m \times n$ samples, and $\hat{\theta}^j$ is the average of the samples within chain number j . W estimates the variance within each chain individually, while B estimates the variance between the chains. When the chains have converged properly, V and W should be identical, and R should be close to unity.

Gelman & Rubin (1992) make the general recommendations that the initial points for the m chains should be overdispersed relative to the true distribution and that the chains should be run until $R < 1.2$. However, as discussed above, we initialize the chains at the maximum likelihood value in order to avoid burn-in problems, and thus the first point is certainly not fulfilled in our approach. Consequently, the numerical value of the convergence criterion they give does not apply to our prescription.

To remedy this situation, we impose two alternative criteria. First, we require that the chains run for a minimum number of samples (typically on the order of 10^7 , but only storing every, say, 500th sample, to reduce correlations). Second, we require that the largest value of R , individually computed for all parameters included in the model, must be smaller than 1.01. In most cases, we find that the latter criterion is fulfilled long before

the former, indicating that the overall CPU time may be decreased somewhat. However, since we cannot inspect the chains manually for more than a few pixels, we consider safety to be more important than speed and adopt a very conservative approach. With the criteria discussed here, we have found excellent convergence in all cases we have inspected, and the Monte Carlo error (the error on the error due to a finite number of samples) is typically less than 1%.

5.4.3. Point Estimators and Uncertainties

The MCMC algorithm provides us with a large number of multivariate samples drawn from the posterior, and these may be used to form a great variety of useful statistics. Here we focus on the univariate distributions for each parameter, marginalizing over all others. (Marginalization with MCMC samples is straightforward: simply disregard the “uninteresting” parameters and make a histogram of the “interesting” parameter sample values.)

Our point estimate for each parameter value is then defined to be the mean of the MCMC samples, with an uncertainty given by variance of the samples. As shown below in practical examples, this Gaussian approximation is quite good for both the CMB temperature and the spectral indices, while the distributions for the synchrotron and free-free amplitudes tend to be non-Gaussian due to a combination of strong internal correlations and a positivity prior.

One of the most important differences between our approach and that taken by Brandt et al. (1994) is that they chose the maximum likelihood value as their point estimate, whereas we choose the posterior mean. This makes our estimate considerably more stable with respect to noise, since it takes more to shift the entire likelihood volume than to change its shape. Therefore, not only does our method yield accurate error bars on all relevant quantities, but in addition the point estimates are more reliable.

5.5. Analytic Estimation of Linear Parameters

Due to computational and sensitivity limitations, the MCMC analysis described in the previous section is only practical at relatively low angular resolutions. In order to produce full-resolution results, approximations are therefore unavoidable, and the approach we take in the present paper is to fix the spectral indices at their low-resolution values and solve for only the (linear) component amplitudes at high resolution.

As we discuss in the next section, a complete Bayesian solution to this problem is available by means of a Gibbs sampling algorithm. The implementation of this approach is left for a future publication, and in the present work we instead adopt the following simple analytic approximation.

Recall that our posterior distribution (assuming uniform priors) takes the form of a standard χ^2 ,

$$\chi^2 = \sum_{k=1}^{N_{\text{band}}} \frac{1}{\sigma_k^2} \left[d_k - \sum_{i=1}^{N_{\text{comp}}} A_i c_{ik} \left(\frac{\nu_{ik}}{\nu_{i0}} \right)^{\beta_i} \right]^2. \quad (23)$$

Since the only free parameters are now the component amplitudes A_i and the noise is assumed to be Gaussian, this is simply a multivariate Gaussian distribution. Therefore, the mean of the distribution equals the maximum likelihood value and may be determined simply by equating the derivatives of the χ^2 with respect to the parameters to zero. In a matrix form, this reads

$$\frac{\partial \chi^2}{\partial \mathbf{a}} = 0 \quad \Rightarrow \quad \mathbf{a} = \mathbf{M}^{-1} \mathbf{d}, \quad (24)$$

where $\mathbf{a}^t = (A_1, \dots, A_{N_{\text{comp}}})^t$,

$$M_{ij} = \sum_{k=1}^{N_{\text{band}}} \frac{c_{ik} c_{jk}}{\sigma_k^2} \left(\frac{\nu_{ik}}{\nu_{i0}} \right)^{\beta_i} \left(\frac{\nu_{jk}}{\nu_{j0}} \right)^{\beta_j}, \quad (25)$$

and

$$d_i = \sum_{k=1}^{N_{\text{band}}} \frac{d_k c_{ik}}{\sigma_k^2} \left(\frac{\nu_{ik}}{\nu_{i0}} \right)^{\beta_i}. \quad (26)$$

While the above formulae yield excellent results for the high-resolution parameter point estimates (as long as all amplitudes are nonnegative), reliable error estimation within this framework is complicated. One problem is that we need to propagate the errors in the spectral indices into the final data products, taking into account the strong correlations between the errors. A second problem is introduced by the positivity prior on the foreground amplitudes, which leads to strongly non-Gaussian distributions when active. Nevertheless, for well-behaved pixels (i.e., those with clear detection of all components individually) a rough approximation may be established by means of the usual error propagation formula.

Suppose we are interested in a quantity $z = f(x, y, \dots)$ that depends on a set of measured quantities x, y, \dots , each with *independent* and *Gaussian* errors $\Delta x, \Delta y, \dots$. In this case, the uncertainty Δz may be approximated by

$$\Delta z^2 = \left(\frac{\partial f}{\partial x} \right)^2 \Delta x^2 + \left(\frac{\partial f}{\partial y} \right)^2 \Delta y^2 + \dots \quad (27)$$

This may be applied to our case by making the identification $A_i = f(\mathbf{x}) = \sum_{j=1}^{N_{\text{comp}}} M_{ij}^{-1} d_j$. The uncertain quantities are both the observed data and the nonlinear parameters, $\mathbf{x}^T = (d_\nu, \beta)^T$.

To compute the uncertainties, we need the partial derivatives, which by equation (24) read

$$\frac{\partial \mathbf{a}}{\partial \mathbf{x}} = \mathbf{M}^{-1} \frac{\partial \mathbf{M}}{\partial \mathbf{x}} \mathbf{M}^{-1} \mathbf{d} + \mathbf{M}^{-1} \frac{\partial \mathbf{d}}{\partial \mathbf{x}}. \quad (28)$$

The derivatives $\partial \mathbf{M} / \partial \mathbf{x}$ and $\partial \mathbf{d} / \partial \mathbf{x}$ are obtained from equations (25) and (26).

Great care must be taken when applying this method to the high-resolution data—it is only valid under the assumptions that the uncertain quantities are both Gaussian distributed and internally independent, neither of which is true for our problem. Nevertheless, while the formal requirements are not strictly fulfilled, the approximation may still be useful for establishing the order of magnitude of the uncertainties.

5.6. Propagation of Errors to the CMB Power Spectrum and Cosmological Parameters

The analytical approach described in the previous section yields good point estimates for the desired parameters, but only approximate uncertainties. Furthermore, it is not straightforward to propagate the errors further into higher level data products such as the CMB power spectrum or cosmological parameters.

A much more powerful solution may be devised by combining the methods described in the present paper with the Gibbs sampling approach of Jewell et al. (2004), Wandelt et al. (2004), and Eriksen et al. (2004b). Whereas most other techniques only provide the user with a very simple description of the power spectrum probability distribution (e.g., a maximum likelihood estimate and a Fisher matrix), the Gibbs sampling approach

yields the complete multivariate probability density $P(C_\ell|\mathbf{d})$, C_ℓ being the CMB power spectrum, and \mathbf{d} the data. Furthermore, it is straightforward to introduce new sources of uncertainty into the framework, and such uncertainties are then seamlessly propagated through to the final data products. We outline here how foreground uncertainties may be propagated to the CMB power spectrum and cosmological parameters, but leave the details for a future publication.

The Gibbs sampling approach is similar in philosophy to the MCMC method that is used extensively in this paper: the target density is established by drawing samples from it. In our case, we are interested in the joint probability distribution $P(C_\ell, \mathbf{s}_{\text{CMB}}, \mathbf{s}_s, \mathbf{s}_{\text{ff}}, \mathbf{s}_d|\mathbf{d})$, where \mathbf{s}_i are the four interesting signal components discussed earlier, CMB, synchrotron, free-free, and thermal dust emission. While it is difficult to sample from this distribution directly, the Gibbs sampling algorithm provides a neat solution. Suppose we want to draw samples from a joint distribution $P(x, y)$ but only know how to sample from the conditional densities $P(x|y)$ and $P(y|x)$. In that case, the theory of Gibbs sampling says that samples (x, y) can be drawn by iterating the following sampling equations:

$$x^{i+1} \leftarrow P(x|y^i), \quad (29)$$

$$y^{i+1} \leftarrow P(y|x^{i+1}). \quad (30)$$

The arrow symbol indicates that a random number is drawn from the distribution on the right-hand side. After some burn-in period, the samples will converge to being drawn from the required joint distribution.

Suppose now that we want to analyze a data set that, for simplicity, only includes CMB and synchrotron emission, the latter being parameterized by an amplitude A_s and a spectral index β_s for each pixel. Suppose further that we already have run an MCMC analysis for each pixel as described earlier and have access to the corresponding probability distributions. In that case, the Gibbs sampling algorithm may be applied by means of the following sampling chain:

$$\beta_s^{i+1} \leftarrow P(\beta_s|C_\ell^i, A_s^i, \mathbf{s}_{\text{CMB}}^i, \mathbf{d}), \quad (31)$$

$$A_s^{i+1} \leftarrow P(A_s|C_\ell^i, \beta_s^{i+1}, \mathbf{s}_{\text{CMB}}^i, \mathbf{d}), \quad (32)$$

$$\mathbf{s}_{\text{CMB}}^{i+1} \leftarrow P(\mathbf{s}_{\text{CMB}}|C_\ell^i, A_s^{i+1}, \beta_s^{i+1}, \mathbf{d}), \quad (33)$$

$$C_\ell^{i+1} \leftarrow P(C_\ell|\mathbf{s}_{\text{CMB}}^{i+1}). \quad (34)$$

(The CMB power spectrum C_ℓ only depends on the CMB signal, not the foregrounds, and therefore the other components are omitted from the right-hand side in the last equation.) The first two rows are to be performed for each pixel individually, while the last two rows are performed in harmonic space, reflecting the intuitively pleasing idea that foregrounds should be handled in pixel space, while CMB fluctuations are better handled in harmonic space.

To perform the analysis as described above, we have to be able to sample from all involved conditional distributions. Sampling the CMB signal and power spectrum parts is detailed by, e.g., Eriksen et al. (2004b). Sampling the foreground amplitudes given the spectral indices is straightforward, since the corresponding distributions are simple Gaussians.

However, sampling the spectral indices is a priori not trivial—their distributions are highly non-Gaussian and no analytical expressions exist. However, given that we already have run an MCMC analysis whose product is precisely the joint density $P(\mathbf{s}_{\text{CMB}}, A_d, \beta_s|\mathbf{d})$, the problem is mostly solved. We may sim-

ply generate a full multidimensional histogram from the MCMC samples (for each pixel separately) and pick out the synchrotron index column that corresponds to the other currently fixed parameter values. Given this one-dimensional distribution, we can then sample numerically using standard techniques.

For completeness, we note that adding more than one foreground component is a straightforward extension of this scheme. Component amplitudes are added individually, while there is a choice for spectral indices—one may either sample these individually, as done above for the synchrotron index, or for greater efficiency, one may also exploit the multivariate information given by the MCMC analysis.

With the above prescription, it is finally possible to propagate the foreground uncertainties rigorously all the way from the observed data through to the CMB power spectrum and therefore to the cosmological parameters. Furthermore, with this approach one also obtains full-resolution sampled uncertainties of the component amplitudes, as opposed to the analytical approximations discussed in the previous section, and a complete probabilistic description of the system is thereby established.

While Gibbs sampling as currently implemented by Jewell et al. (2004), Wandelt et al. (2004), and Eriksen et al. (2004b) has problems with probing the low signal-to-noise ratio (S/N) regime properly, it works very well for S/N's larger than unity (Eriksen et al. 2004b), and this is exactly where the foreground uncertainties dominate. Therefore, it seems reasonable to use the approach presented here to analyze the high and intermediate S/N regimes, propagating foreground uncertainties to the final products, and a standard MASTER-type analysis (Hivon et al. 2002) for the low S/N regime, at the cost of reducing the reliability of foreground uncertainty estimates at these angular scales.

6. EXAMPLE: APPLICATION TO *PLANCK* AND 6 YEAR *WMAP* DATA

We now apply the MCMC component separation method described in § 5 to simulations of experiments similar to the current *WMAP* and the future *Planck* missions. We first give a detailed presentation of the simulations and data. We then study the behavior of the algorithm for one arbitrarily chosen pixel, before considering the full-sky map solutions.

We point out that our main goal in this paper is to study the algorithm itself and not to simulate an actual data release. We therefore choose examples both with and without modeling errors, in order to illustrate problems that may be encountered in an analysis of real data.

6.1. Simulations and Models

The simulations used in the following are constructed as a sum of a cosmological CMB signal, three foreground components (synchrotron, free-free, and thermal dust emission) and instrumental noise. We include 5 bands (centered at 23, 33, 41, 61, and 94 GHz) from *WMAP*, 3 bands (30, 44, and 70 GHz) from the *Planck* Low Frequency Instrument (LFI), and 3 bands (100, 143, and 217 GHz) from the High Frequency Instrument (HFI), for a total of 11 bands between 23 and 217 GHz¹³. Specifications for each detector are given in Table 1.

CMB.—The CMB component is assumed to be Gaussian distributed, with variances given by the best-fit *WMAP* power-law power spectrum (Bennett et al. 2003a; Hinshaw et al. 2003;

¹³ The three highest HFI frequency bands are not included in the analysis because they would introduce significant dust modeling errors; we simulate dust with a two-component model, but fit for a one-component model.

TABLE 1
FREQUENCY BAND SPECIFICATIONS

Experiment	Center Frequency (GHz)	Bandwidth (GHz)	Beam FWHM (arcmin)	rms per 7' Pixel (μ K)	rms per 60' Beam (μ K)
<i>WMAP</i>	23	5	52.8	50 \pm 7	7.6 \pm 1.1
LFI.....	30	6	33.0	33 \pm 7	3.0 \pm 0.6
<i>WMAP</i>	33	8	39.6	51 \pm 7	5.0 \pm 0.7
<i>WMAP</i>	41	11	30.6	49 \pm 8	4.3 \pm 0.7
LFI.....	44	8	24.0	33 \pm 6	2.7 \pm 0.5
<i>WMAP</i>	61	16	21.0	60 \pm 9	4.8 \pm 0.7
LFI.....	70	14	14.0	31 \pm 6	2.4 \pm 0.5
<i>WMAP</i>	94	24	13.2	73 \pm 11	5.6 \pm 0.8
HFI.....	100	33	9.5	14 \pm 3	1.1 \pm 0.2
HFI.....	143	48	7.1	8 \pm 1	0.6 \pm 0.1
HFI.....	217	72	5.0	11 \pm 2	0.9 \pm 0.2

NOTES.—Specifications for each frequency band used in the simulation. All beams are assumed to be Gaussian. The rms values for 7' pixels are computed taking into account the scanning strategy of each detector, but neglecting noise correlations. The rms values per 60' beam are estimated from 1000 Monte Carlo simulations by drawing Gaussian random numbers corresponding to the rms level of 7' pixels, deconvolving the instrument beam, and finally convolving with a 60' FWHM Gaussian beam. The values also take into account reduced pixel resolution, from 7' to 14' pixels. The *Planck* rms values are requirement levels, not goals.

Spergel et al. 2003), including multipoles between $\ell = 2$ and 1024. The signal realization is filtered through the HEALPix pixel window function and the instrument specific beam windows. (Since all the *Planck* beam windows are not available, we choose for simplicity to model even the *WMAP* beams as Gaussians with appropriate FWHM's.)

Noise.—The noise is assumed to be Gaussian and uncorrelated, but nonuniform according to the scanning strategy of each detector. For *WMAP*, we assume a 6 year mission and rescale the published first-year sensitivity levels by $1/\sqrt{6}$. For *Planck*, we adopt the requirement levels, which are a factor of 2 worse than the goals, for the baseline 1 year mission.

Synchrotron emission.—The only all-sky map currently available to provide a template of Galactic synchrotron emission in HEALPix format is the 408 MHz survey by Haslam et al. (1982). This has a resolution of only 51', and thus additional power must be added on smaller angular scales for our purposes in this paper. We adopt the model of Giardino et al. (2002), who estimated the amplitude and slope of the synchrotron angular power spectrum at low Galactic latitudes for $l \geq 150$. A Gaussian realization was then generated from that power spectrum to which was applied a Galactic modulation, multiplying the signal in each pixel by the ratio between the Haslam et al. (1982) template in that pixel and the maximum. The final template is added to the original Haslam et al. (1982) map. An all-sky template for the synchrotron spectral index was estimated combining the all-sky data from Haslam et al. (1982), with the northern sky observations of Reich & Reich (1986) at 1420 MHz and the southern sky counterpart of Jonas et al. (1998) at 2300 MHz.¹⁴ This constitutes our synchrotron model.

Free-free emission.—Current models of the free-free emission exploit the expected correlation with H α emission (see Dickinson et al. 2003 and references therein) and predict free-free emission (given in antenna temperature) as follows:

$$s_{\text{ff}}(\nu) = (A_{\text{H}\alpha}) 0.1366 \times 10^{0.029(10^4 \text{ K}/T_e)} \text{G} \quad (35)$$

$$\times \left(\frac{T_e}{10^4 \text{ K}} \right)^{0.517} \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{-2}. \quad (36)$$

¹⁴ This hybrid spectral index model should ultimately be superseded by the full-sky 1420 MHz survey described in Reich et al. (2003).

Here $A_{\text{H}\alpha}$ is the H α amplitude, $G = 3.96 T_e^{0.21} (\nu/\nu_{0,\text{ff}})^{-0.14}$ is the Gaunt factor (Finkbeiner 2003), T_e is the electron temperature, and $\nu_{0,\text{ff}} = 40$ GHz is a reference frequency.

A major uncertainty when using H α as a template is due to the absorption of H α by foreground dust, but this can be estimated using the 100 μ m maps from Schlegel et al. (1998). The largest uncertainty, however, is related to the fraction of dust (f_d) lying in front of the H α -emitting region. Dickinson et al. (2003) show that for regions $30^\circ < l < 60^\circ$ and $5^\circ < |b| < 15^\circ$, $f_d \approx 0.3$, while for local high-latitude regions such as Orion and the Gum Nebula, there is little or no absorption by dust ($f_d \sim 0$). The latter is supported by the cross-correlation analysis of Banday et al. (2003) of the H α data with *COBE* DMR, which contrasts with the value of $f_d \sim 0.5$ expected under the assumption that the ionized gas and dust are coextensive along the line of sight (i.e., uniformly mixed), as is assumed in the *WMAP* analysis of Bennett et al. (2003b) and Finkbeiner (2004).

We correct for dust absorption, by assuming a single-component dust model, with a temperature of 18.3 K, and an absorption fraction of 0.33 up to a flux corresponding to 1 mag. We assume an electron temperature of $T_e = 7000$ K and therefore an effective frequency scaling close to $\beta^{-2.14}$ over the range of frequencies considered here. For the future, a more accurate model, accounting for the steepening spectral index, could be implemented as a correction to the simple power-law model.

Thermal dust.—We adopt model 8 of Finkbeiner et al. (1999) for thermal dust emission, with parameters $f_1 = 0.0363$, $q_1/q_2 = 13$, $\nu_{0,d} = 3000$ GHz, $\alpha_1 = 1.67$, $\alpha_2 = 2.70$, $T_1 = 9.4$ K, and $T_2 = 16.2$ K (see eqs. [10]–[13]). However, this is too many parameters to fit individually, and we therefore adopt a simpler model for reconstruction (see eq. [14]). Modeling errors of the sort to be expected with real data will result.

Data processing.—All simulations are initially made at a pixel resolution of $N_{\text{side}} = 512$, corresponding to a pixel size of 7'. However, since our method requires identical beam sizes for all frequency bands, we downgrade each band separately to 1° FWHM (determined by the 52'8 FWHM beam of the 23 GHz *WMAP* band) and reduce the pixel resolution to $N_{\text{side}} = 256$ (by deconvolving the original beam and pixel windows, and convolving the common 1° FWHM beam and lower resolution pixel window).

By downgrading the data, the noise specifications are also modified. To estimate the effective noise levels after degradation,

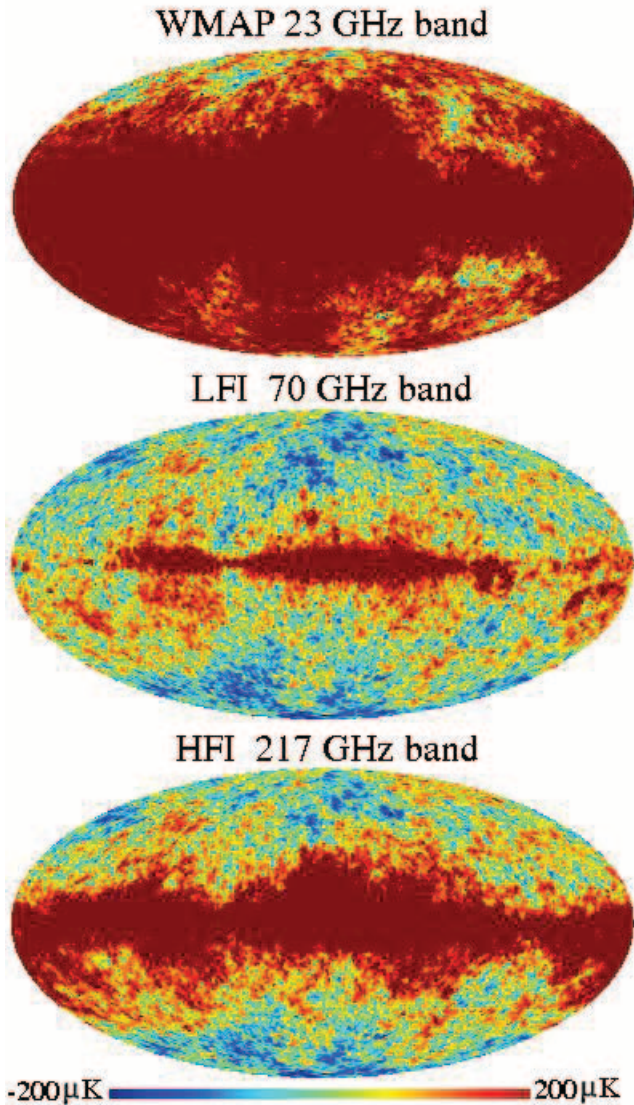


FIG. 2.—“High-resolution” simulations used in this paper. Shown are the 23 GHz channel from the *WMAP* experiment, the 70 GHz channel from the LFI experiment, and the 217 GHz channel from the HFI experiment. All maps are smoothed to a common resolution of 1° FWHM.

we therefore generate 1000 noise realizations for each band and downgrade these in the same manner as the actual data maps. The effective noise levels of the downgraded maps are then estimated by taking the standard deviation of the 1000 realizations.

The data set described above constitutes our main simulation and is referred to in the following as “high-resolution data.” Examples are shown in Figure 2. For the nonlinear parameter estimation step using MCMC, the noise levels must be lower, as discussed above. Therefore, we smooth all maps with an additional 6° FWHM Gaussian beam and downgrade the pixel resolution to $N_{\text{side}} = 32$. (This smoothing scale is not optimized, but it is sufficient for the purposes of the present paper.) Again, the effective noise levels are determined by Monte Carlo simulations. This smoothed data set is referred to as “low-resolution data.”

Initial model map.—The initial model map is based on the *WMAP* Kp0 mask (Bennett et al. 2003b). First the excluded region of the original mask is expanded by 10° in all directions. Then all accepted pixels (i.e., the high latitude region) are assigned the model that includes CMB, synchrotron, and dust

(both with free amplitude and spectral index), while the model for the rejected pixels includes in addition a free-free amplitude.

6.2. Results

We now apply the method of § 5 to the simulated data set described above.

6.2.1. Single-Pixel Results

We first examine the performance of the MCMC algorithm by studying one single pixel in the low-resolution data set, namely, pixel number 6100, which is located inside the Galactic plane at $l = 58^\circ$, $b = 0^\circ$. The reasons for choosing this pixel (or one like it) are twofold. First, the model for this pixel includes all three foreground components and has thus a complicated probability structure. Second, the model is rejected by the goodness-of-fit test, and this example therefore illustrates the modeling error problem.

As discussed above, the MCMC algorithm basically performs a random walk on the likelihood surface, producing a set of samples from which the likelihood may be estimated by constructing single or multidimensional histograms. Examples of such histograms are shown in Figures 3 and 4.

The first figure shows the probability distributions for each of the six included parameters, marginalized over all other parameters. Comparing with the true input values (*vertical lines*), we see that the algorithm reproduces the correct values and also that the uncertainties are reasonable compared to the true errors.

In Figure 4 we show two-dimensional probability distributions for the same parameters. The true values are marked by a box. Several points are worth noticing in this figure. First, all parameters are clearly correlated and some specific pairs very tightly so. Examples of the latter are dust amplitude versus dust spectral index and synchrotron amplitude versus free-free amplitude.

Second, many of the distributions are clearly non-Gaussian, and it is clear that a Gaussian approximation at this stage will not yield reliable errors. Still, the structures appear to be reasonably well behaved, and in principle it may be possible to find analytical parameter transformations that could ease the computational burden.

Third, while most of the true values lie inside the 3σ confidence regions, in one case, namely, the synchrotron amplitude versus synchrotron index, it lies far outside the acceptable region. Another perspective on this is provided by the χ^2 , which for this pixel is 44. With 5 dof (11 frequencies and 6 free parameters), this particular model is thus formally ruled out at the 99.9999% confidence level. This is because we fit for a simpler model than the one used in the simulation: the data are smoothed by a wide 6° beam, and the thermal dust is fitted with a one-component model, whereas the simulation was based on a two-component model. This may also be seen Figure 1, where the fitted spectra for each component for this pixel are plotted. At low frequencies, the data points lie systematically above the fitted model, resulting in a clear rejection.

However, even though the model is strongly rejected by the goodness-of-fit test, it is important to note all of the univariate distributions are still reasonable, and the CMB reconstruction is still useful. Therefore, a high χ^2 does not necessarily imply that the pixel has to be discarded from further analysis, but rather that extra care has to be taken. Preferentially, the extra information indicated by the high χ^2 should be used to improve the model.

6.2.2. Low-Resolution Full-Sky Maps

We now consider the reconstructed full-sky maps, starting with the low-resolution maps as computed by the MCMC analysis. The individual component maps are shown in Figure 5,

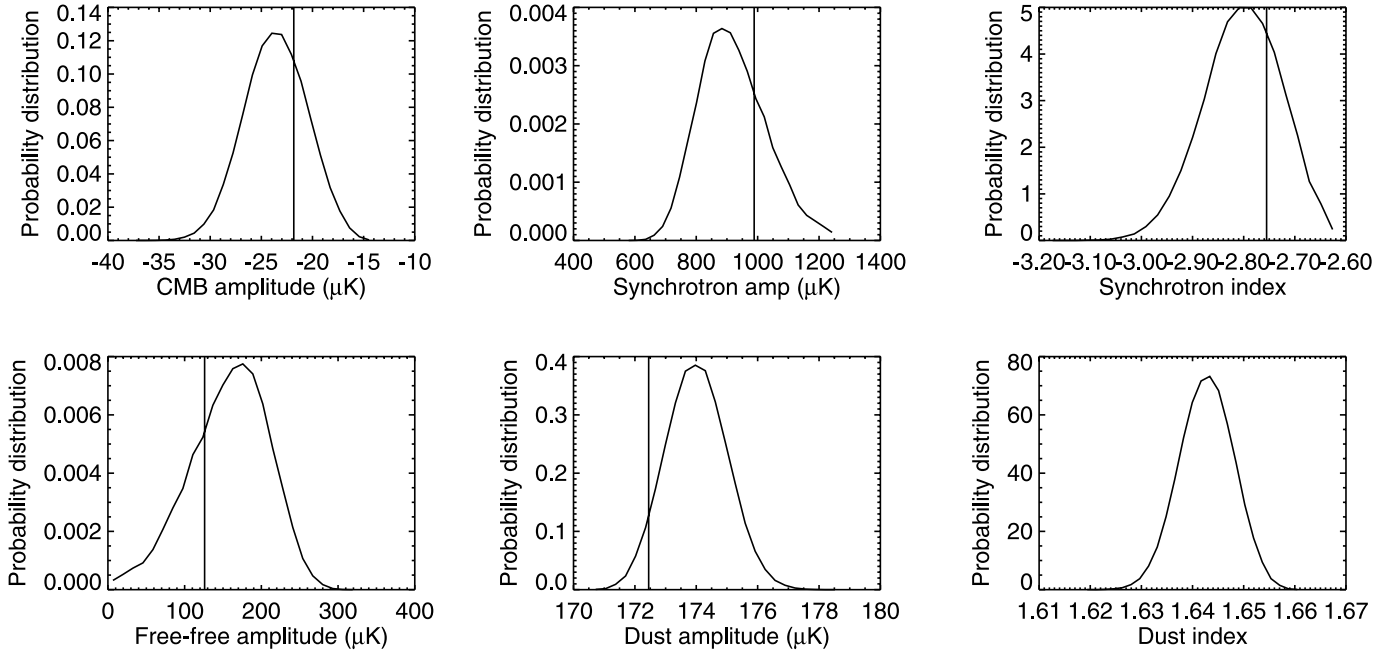


FIG. 3.—Marginalized parameter probability distributions for an arbitrarily chosen pixel inside the Galactic plane, generated by MCMC as described in the text. The vertical lines show the true input value for the pixel. (The true value is not well defined for the dust spectral index, since dust is modeled by a two-component spectrum, while a one-component model is fitted. However, the stronger of the two dust components has an index parameter of $\alpha_1 = 1.67$.)




FIG. 4.—Marginalized two-dimensional probability distributions for the same pixel as in Fig. 3, computed by MCMC. Boxes indicate the true input values, and the contours mark the peak and the 68%, 95%, and 99.7% confidence levels.

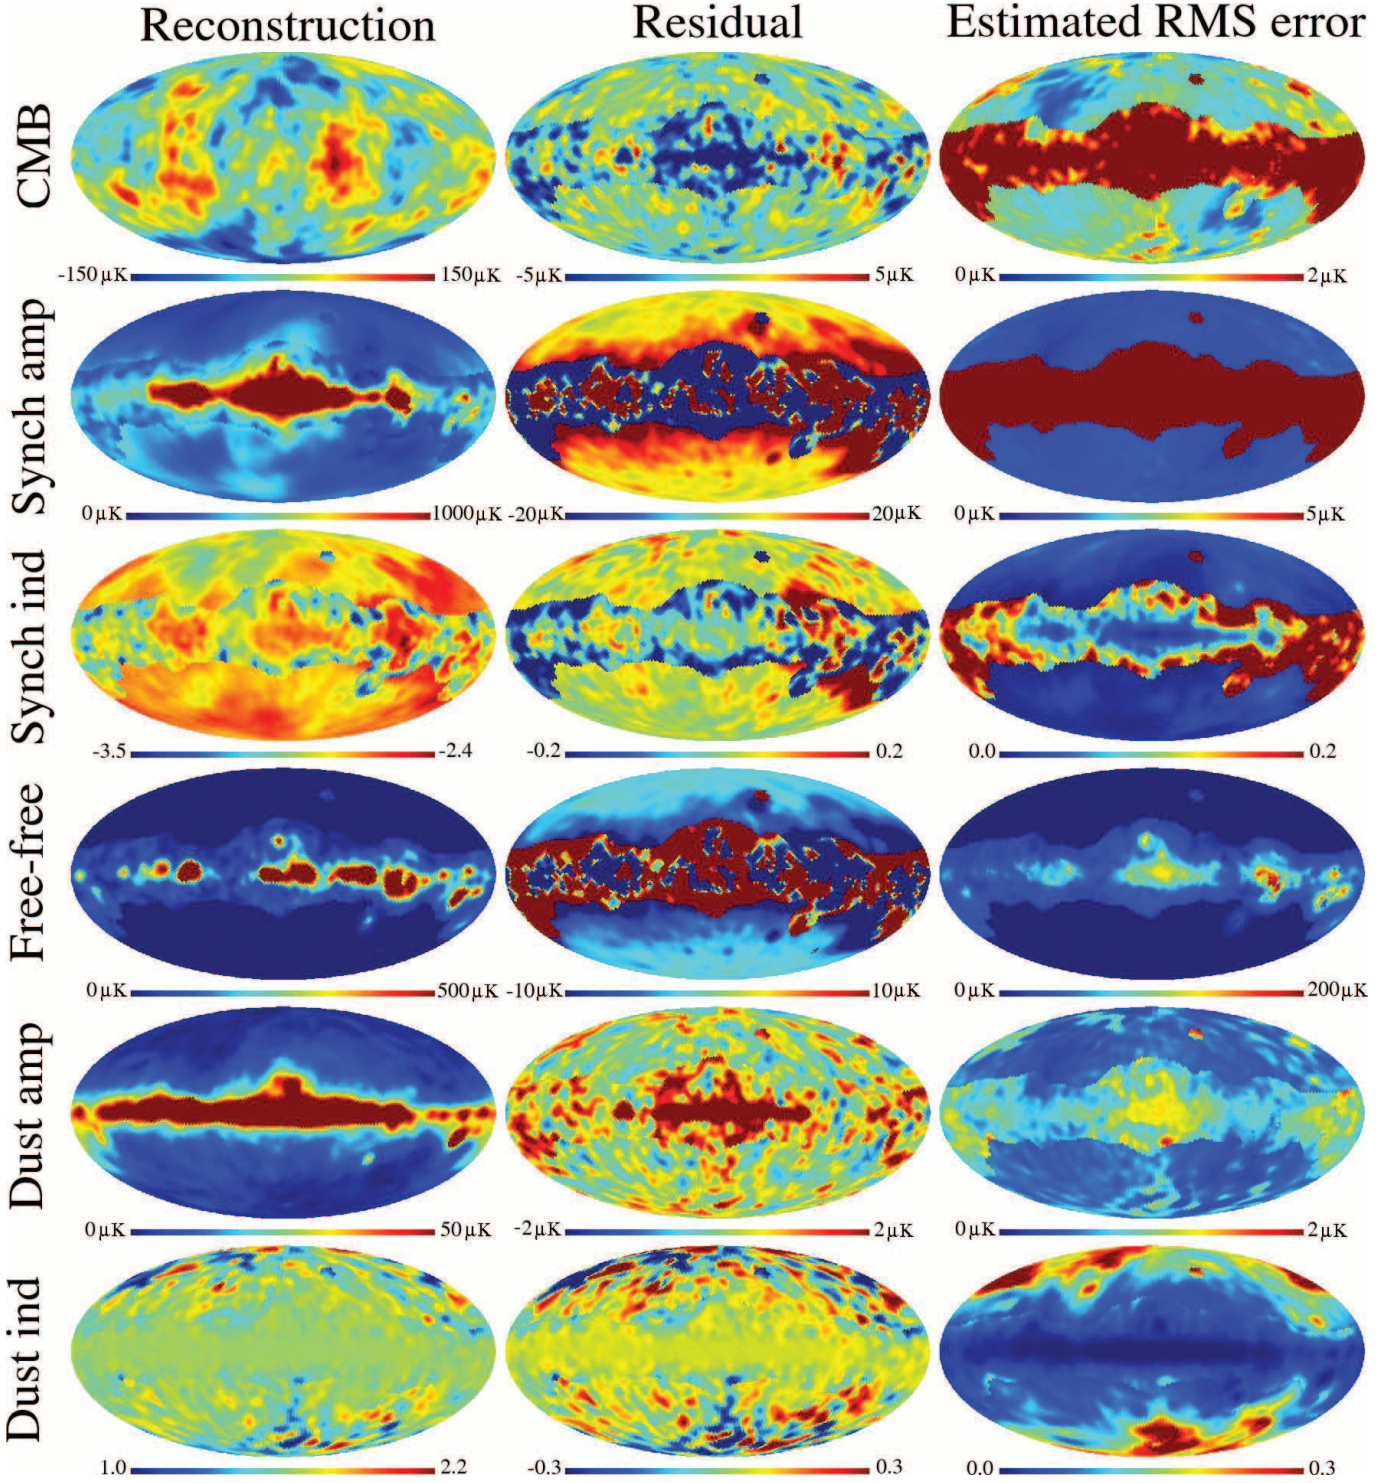


FIG. 5.—Low-resolution parameter maps reconstructed by MCMC. The left column shows the parameter estimates, the middle column shows the difference between reconstructed and input maps, and the right column shows the rms errors estimated by MCMC. *Top to bottom*: (1) the thermodynamic CMB temperature; (2) the synchrotron emission amplitude relative to 23 GHz; (3) the synchrotron spectral index; (4) the free-free emission amplitude relative to 33 GHz; (5) the thermal dust emission amplitude relative to 90 GHz; and (6) the thermal dust spectral index.

with reconstructions given in the left column, differences between reconstructed and input maps in the middle column, and estimated errors in the right column. In the left column of Figure 6, we show the model map used in the analysis and the resulting goodness-of-fit χ^2 distribution.

Starting with the goodness-of-fit map, we first note that we should expect $\chi^2 \lesssim 13$ at 2σ confidence at high latitudes, since

the model has 6 dof in this region. This is indeed the case for two wide bands on each side of the Galactic plane, and both the model and the estimated parameters may therefore be accepted as they stand. However, at very high latitudes and, less surprisingly, at low latitudes, the goodness of fit is poor.

In the left column of Figure 5 we show the six reconstructed parameter maps and in the middle column the actual output

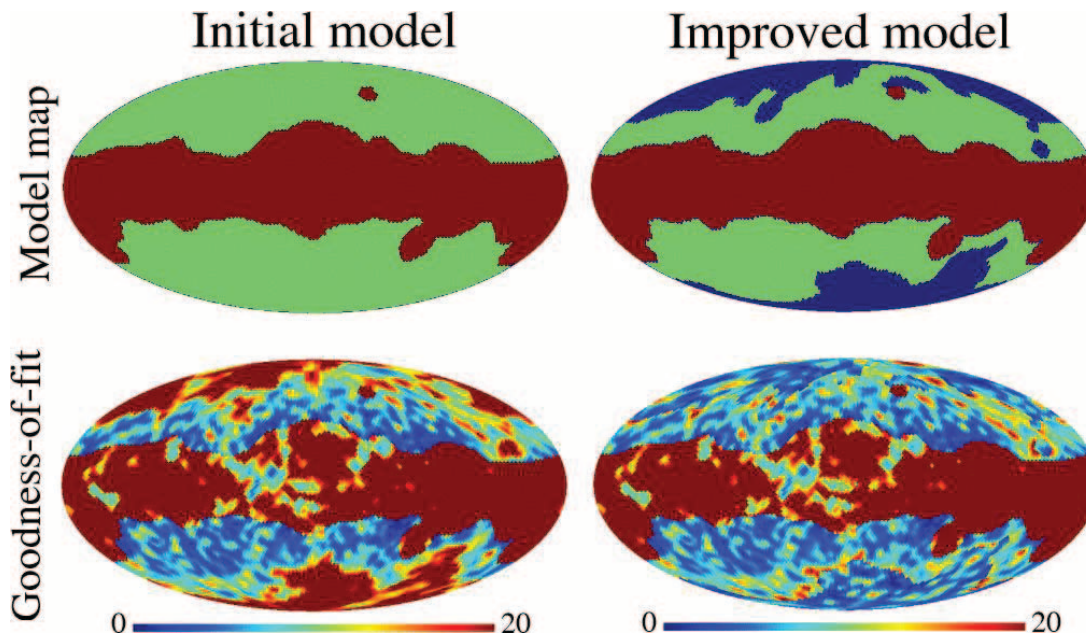


FIG. 6.—Example of iterative improvement of the model map. Two MCMC runs using the same data were made. The first run included CMB, synchrotron, and dust (both with free amplitude and spectral index) at high latitudes (*green*) and also free-free at low latitudes (*red*). The second run fixed the dust index at 1.55 at very high latitudes (*blue*). The bottom row shows the χ^2 distribution in the two cases; note the significant improvement at high latitudes resulting from *removing* a (noncritical) parameter from the system. Similar improvements at low latitudes could be made by trial and error.

versus input errors. Clearly, the method works very well, as the CMB sky map is virtually free of artifacts, with residuals less than 10% even in the inner Galactic plane. And with the exception of sharp boundaries in the foreground reconstruction, due to different models used in different regions, the foreground results also look encouraging.

However, as good as these results are, we warn the reader against interpreting them as an expected performance level for future missions. Even though our simulations are as realistic as possible given our current understanding of foreground properties, they are certainly not as complicated as the real sky. Considerable modeling errors must be expected for real data sets, and sky cuts are very likely still required for future work.

One note about the sharp boundaries seen in the foreground maps is in order. If the reconstructed maps are intended for foreground studies, such features are clearly not acceptable. In such cases, postprocessing may be required, for instance by smoothing the boundary by a Gaussian beam. On the other hand, if the maps are to be used for CMB power spectrum or cosmological parameter estimation, it is better to use the maps as they are and propagate the pixel errors reliably; the boundaries are mainly due to different noise properties in the various regions. However, we point out that the distinct boundaries seen in Figure 5 are at least partially due to a poorly chosen model map, constructed from a *WMAP* Galactic mask, rather than the specific simulation under consideration; manual tweaking would surely improve the results considerably.

Returning for a moment to the goodness-of-fit map shown in Figure 6 and comparing with the rms maps shown in the right column of Figure 5, we see that the very high latitude region with high χ^2 corresponds directly to the thermal dust spectral index map. Furthermore, we also see that the dust amplitude is very low in the same region. The interpretation is clear: thermal dust is not well constrained in these regions because of its low amplitude, leading to poorly constrained spectral indices. This again propagates into the CMB component, and the overall fit is unacceptable.

The solution to this problem seems obvious. Since the main problem is unconstrained dust spectral indices, we should manually fix them at some reasonable value. The potential bias introduced in the CMB and other components by doing so is very small because of the small dust amplitude found by the first analysis. We implement this by assigning a new model that fixes the dust spectral for all pixels with a dust spectral index rms larger than 0.15 in the bottom right panel of Figure 5. The fixed spectral index value is somewhat arbitrarily chosen to be 1.55. The modified model map is shown in the top right panel of Figure 6.

We now repeat the analysis and obtain the goodness-of-fit map shown in the bottom right panel of Figure 6. Clearly, introducing a new model at high latitudes had a very beneficial impact on the results. In principle, we could now proceed with similar considerations at low latitudes and obtain reasonable fits over the full sky. However, since our main purpose in this paper is to illustrate the method, we are content with the slightly revised model map shown in the top right panel of Figure 6, and we use this map in the rest of the paper.

We now consider the error estimation accuracy of the MCMC algorithm. In the left panel of Figure 7 we plot a histogram of the relative CMB reconstruction error $\delta = (\Delta T_{\text{est}} - \Delta T_{\text{in}})/\sigma_{\text{est}}$, where ΔT_{est} is the estimated CMB temperature, ΔT_{in} is the true value, and σ_{est} is the estimated error. If both the amplitude and the error are perfectly estimated, the pixel histogram will match a Gaussian distribution with vanishing mean and unit variance. (In this plot, we include only pixels in the intermediate latitude region with a goodness of fit $\chi^2 < 13$.) Obviously, the algorithm works very well, as the bias is very small indeed and the estimated error is very close to the true error.

Finally, in Figure 8 we plot a histogram of the χ^2 values of the same pixels and compare it to a χ_6^2 distribution. Clearly, there is a small shift toward high values. This is likely due to a combination of two effects. On the one hand, poor fits tend to bias the distribution to high values, and such an effect is therefore expected when considering the impact of modeling errors. On the




FIG. 7.—CMB reconstruction accuracy shown as a histogram of $\delta = (\Delta T_{\text{est}} - \Delta T_{\text{in}})/\sigma_{\text{est}}$ for pixels in the intermediate-latitude region. Results from the low-resolution MCMC analysis are shown in the left panel and from the high-resolution analytical analysis in the right column. For perfect reconstruction of both CMB amplitude and error, both curves would match a Gaussian distribution with vanishing mean and unit variance (*dashed curve*).

other hand, we also know that the actual distribution may not be a perfect χ^2 density, due to the presence of the nonlinear parameters and their uniform priors.

Nevertheless, taking into account all these issues, the main conclusion to be drawn from this figure is that the distribution is *sufficiently close* to the theoretical expectations for the goodness-of-fit maps to be useful in the model selection process; a strong outlier relative to the analytical distribution is also a strong outlier relative to the empirical distribution. (Note that for restrictive priors, this may no longer be true, but in such cases one is well advised to thoroughly check the priors.) The minor observed discrepancy is of no practical importance for the model definition process.

We conclude this section by making a few comments on the computational cost of the method. Running the MCMC analysis for each pixel is by far the most expensive step of the algorithm. For well-behaved pixels, we find that it takes on the order of 100

CPU seconds (divided over four processors per pixel) to reach the convergence criteria described above. For an $N_{\text{side}} = 32$ map with 12,288 pixels, it therefore takes about 350 CPU hours per run. For clusters with of order 10^2 processors, this is not a major problem. Further, since the algorithm scales with the number of pixels and parallelizes trivially, it is not unreasonable to apply it at higher resolutions, say, at $N_{\text{side}} = 128$ for 6000 CPU hours.

6.2.3. High-Resolution Full-Sky Maps

Having estimated the nonlinear parameters by MCMC, the next step is to estimate the component amplitudes from the full-resolution sky maps. As discussed above, this can be done either with a Gibbs sampling approach or with an analytic approach. In this paper, we choose the latter route and leave the former to a future publication.

The results from applying the method described in § 5.5 are shown in Figure 9. Once again, we see that the reconstructed parameter maps look visually compelling. There are few visible signs of contamination in the CMB reconstruction, and, indeed, even inside the central Galactic plane the errors are only a few tens of microkelvins.

In the right panel of Figure 7 we plot the relative CMB reconstruction error for the high-resolution map, as we did for the low-resolution map in the previous section. Two facts are clear from this plot. First, the bias is small, indicating that the analytic point estimator is quite accurate. Second, the histogram does not match the Gaussian distribution well, but is rather focused around smaller values. In other words, the errors are overestimated by some small factor by the analytic error propagation formula. This should not be surprising, given the assumptions that went into those calculations. Nevertheless, the estimated errors are in fact of the correct order of magnitude, and they can therefore be used as a mental guide, although not for quantitative work.

In Figure 10 we plot the power spectrum of the reconstructed high-resolution CMB map, the true realization specific input spectrum, and the ensemble-averaged spectrum. The reconstructed spectrum was computed by full-sky integrals without noise weights. The results are therefore excellent—the reconstruction is virtually perfect up to $\ell = 200$, after which a small noise term starts to make an impact, before the 1° beam renders the reconstruction arbitrary at $\ell = 300$. From this plot it seems clear that we were too conservative when choosing a 6° beam for the




FIG. 8.—Distribution of χ^2 from the low-resolution MCMC analysis, shown for the intermediate-latitude region, with 6 dof. The dashed curve shows the expected χ^2_6 distribution. Note the good agreement between the two distributions—a strong outlier relative to the theoretical curve is also an outlier by the empirical curve. For the crude purposes of the χ^2 statistic in this paper (assisting in the model definition process), this goodness-of-fit statistic is quite adequate.

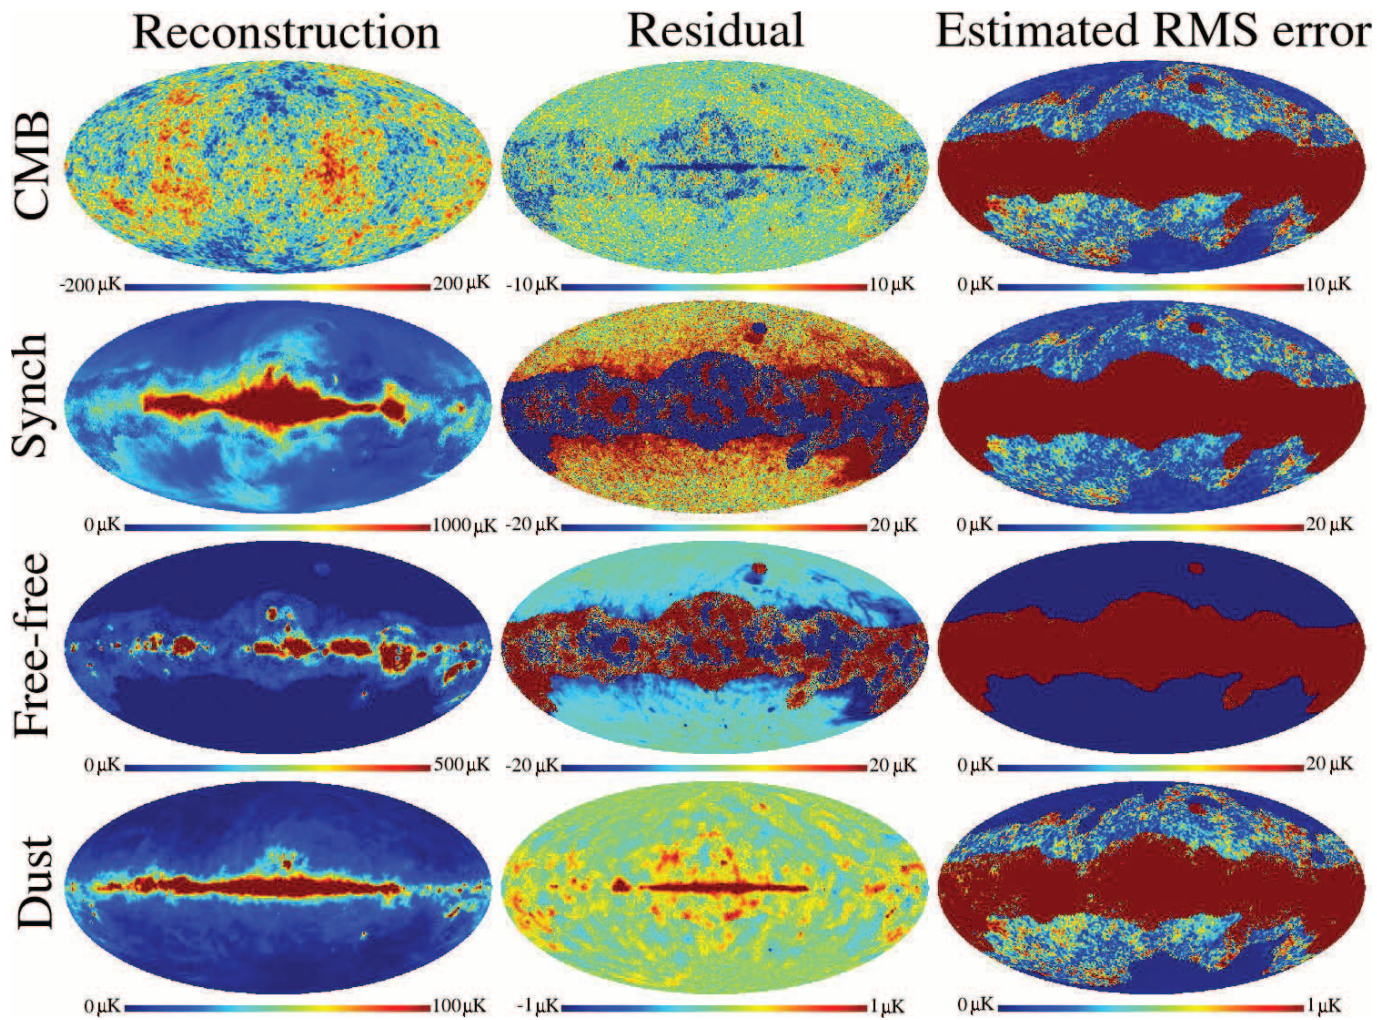


FIG. 9.—High-resolution parameter maps reconstructed by direct solution of linear systems, fixing spectral indices at low-resolution map values. The left column shows the parameter estimates, the middle column shows the difference between output and input maps, and the right column shows the analytically estimated errors. *Top to bottom*: (1) the thermodynamic CMB temperature; (2) the synchrotron emission relative to 23 GHz; (3) the free-free emission relative to 33 GHz; and (4) the thermal dust emission relative to 90 GHz.

low-resolution analysis and that the combined *Planck* and 6 year *WMAP* data can easily handle higher resolutions.

7. OUTSTANDING PROBLEMS

As demonstrated in the previous section, the component separation method presented in this paper works well on simulated data. However, there are a few outstanding issues that we have not taken into account but that will have to be addressed prior to analysis of real data. The most important of these are gain and zero-point calibration, dipole corrections, noise correlations, and beam asymmetries.

First, our algorithm requires all sky maps to be properly calibrated with respect to gain and zero point. Usually, the gain is calibrated using the CMB dipole, but this can be difficult for channels that are highly foreground contaminated. Moreover, zero-point calibration is never easy.

Closely related to these issues is dipole subtraction. The CMB dipole itself is hard to observe because of the large Doppler dipole induced by the motion of the solar system through space; it is usually subtracted in the map-making process. Nevertheless, residual dipoles may cause serious problems for our algorithm unless accounted for.

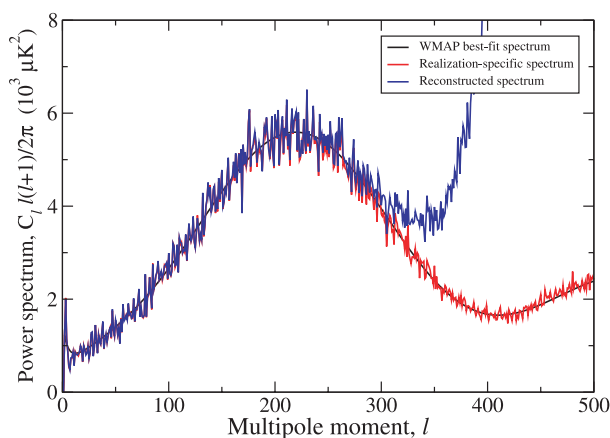


FIG. 10.—Reconstructed (blue curve) and input (red curve) CMB power spectra. The ensemble-averaged spectrum is shown as a black smooth curve. The reconstructed spectrum was computed by full-sky integration without noise weights or sky cut.

Finding reliable calibration methods for each of the above problems is clearly essential. Fortunately, the number of degrees of freedom represented by these issues is quite small, and it may be possible to include them in the analysis by replacing the signal model S_ν in equation (3) with

$$S_\nu(p) \leftarrow g_\nu S_\nu(p) + \sum_{\ell=0}^1 \sum_{m=-\ell}^{\ell} a_{\ell m, \nu} Y_{\ell m}(p). \quad (37)$$

Following a rough calibration with external techniques, one could then use methods similar to those described in this paper to optimize the gains g_ν and monopole and dipole coefficients $a_{\ell m, \nu}$.

From a conceptual point of view, correlated noise poses a more serious problem. For *Planck*, for example, the main effect will be to introduce stripes in the sky maps along the scanning path of the detectors, and locally this has the same effect as an overall offset. Properly speaking, correlated noise is a problem for map making more than it is for component separation; however, residual effects can be expected. Only when actual data are in hand will it be clear how serious a problem it is.

Finally, in this paper we have assumed that all detectors have identical beam response functions. This obviously is not true for any real system, and corresponding errors are unavoidable. Fortunately, this is likely to have a negligible effect on the low-resolution analysis, since we smooth with an additional degree-scale beam, strongly suppressing small-scale asymmetries. Only in the high-resolution analysis is this effect likely to be important.

8. DISCUSSION

In this paper, we approach the problem of component separation with CMB data from the perspective of parameter estimation. Our goal is to propagate foreground uncertainties all the way from observed data through to the final products, most importantly to the CMB power spectrum and cosmological parameters. This is more easily facilitated with standard parameter techniques than with image processing techniques.

We proposed and implemented one particular algorithm for performing this task, based on multifrequency parametric model fits established by means of a hybrid of MCMC and analytic methods. The method was then shown to work very well on simulated data, with properties corresponding to those of the future *Planck* and 6 year *WMAP* experiments.

We also outlined how to propagate the foreground-induced errors to the CMB power spectrum and cosmological parameter

errors, using the output from the MCMC analysis presented here as the input in a Gibbs sampling algorithm. As always, only an actual implementation will prove whether this method works or not, but the theoretical groundwork appears to be sound, and no insurmountable computational problems have been identified. Therefore, if this approach proves successful, we will have a complete, mathematically consistent, end-to-end solution to the foreground problem in CMB analysis.

While we only considered temperature anisotropy observations in the present paper, the method is completely general and can equally well handle polarization measurements, as will be demonstrated in a future study. We will also apply the method to two specific problems. First, we will study the optimization of frequency coverage and S/N in future polarization experiments. Since our method provides error bars on all estimated quantities, it is straightforward to compare different experiment designs. Modeling errors will be an integral part of this work, since such uncertainties have a direct impact on the optimal frequency range to observe. Second, we will apply the method to the currently available *WMAP* data.

We thank Jeff Jewell for interesting and useful discussions. H. K. E. thanks the Center for Long Wavelength Astrophysics at the Jet Propulsion Laboratory (JPL) for its hospitality during the period in which part of this work was performed, and acknowledges financial support from the Research Council of Norway, including a Ph.D. studentship. C. D. thanks Stanley and Barbara Rawn for funding a fellowship at the California Institute of Technology. We acknowledge use of the HEALPix¹⁵ software (Górski et al. 2005) and analysis package for deriving the results in this paper. We also acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). This work was partially performed at JPL, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. C. B. was supported in part by the NASA Long Term Space Astrophysics (LTSA) grant NNG04GC90G. E. P. is an ADVANCE fellow (NSF grant AST 03-40648) and is also supported by NASA grant NAG5-11489. K. M. S. was supported by NASA NAG5-10840, the Department of Energy and the Packard Foundation. This research was supported by a Marie Curie European Reintegration Grant within the 6th European Community Framework Programme.

¹⁵ At <http://healpix.jpl.nasa.gov>.

REFERENCES

- Baccigalupi, C., Perrotta, F., de Zotti, G., Smoot, G. F., Burigana, C., Maino, D., Bedini, L., & Salerno, E. 2004, *MNRAS*, 354, 55
- Banday, A. J., Dickinson, C., Davies, R. D., Davis, R. J., & Górski, K. M. 2003, *MNRAS*, 345, 897
- Barreiro, R. B., Hobson, M. P., Banday, A. J., Lasenby, A. N., Stolyarov, V., Vielva, P., & Górski, K. M. 2004, *MNRAS*, 351, 515
- Bennett, C. L., et al. 2003a, *ApJS*, 148, 1
- . 2003b, *ApJS*, 148, 97
- Bouchet, F. R., & Gispert, R. 1999, *NewA*, 4, 443
- Brandt, W. N., Lawrence, C. R., Readhead, A. C. S., Pakianathan, J. N., & Fiola, T. M. 1994, *ApJ*, 424, 1
- Cottingham, D. A. 1987, Ph.D. thesis, Princeton Univ.
- Delabrouille, J., Cardoso, J.-F., & Patanchon, G. 2003, *MNRAS*, 346, 1089
- Dickinson, C., Davies, R. D., & Davis, R. J. 2003, *MNRAS*, 341, 369
- Donzelli, S., et al. 2005, *MNRAS*, submitted (astro-ph/0507267)
- Draine, B. T., & Lazarian, A. 1998, *ApJ*, 494, L19
- Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004a, *ApJ*, 612, 633
- Eriksen, H. K., et al. 2004b, *ApJS*, 155, 227
- Finkbeiner, D. P. 2003, *ApJS*, 146, 407
- . 2004, *ApJ*, 614, 186
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
- Finkbeiner, D. P., Langston, G. I., & Minter, A. H. 2004, *ApJ*, 617, 350
- Fixsen, D. J., & Mather, J. C. 2002, *ApJ*, 581, 817
- Gelman, A., & Rubin, D. 1992, *Stat. Sci.*, 7, 457
- Giardino, G., Banday, A. J., Górski, K. M., Bennett, K., Jonas, J. L., & Tauber, J. 2002, *A&A*, 387, 82
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. 1982, *A&AS*, 47, 1
- Hinshaw, G., et al. 2003, *ApJS*, 148, 135
- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Hobson, M. P., Jones, A. W., Lasenby, A. N., & Bouchet, F. R. 1998, *MNRAS*, 300, 1
- Jewell, J., Levin, S., & Anderson, C. H. 2004, *ApJ*, 609, 1

- Jonas, J. L., Baart, E. E., & Nicolson, G. D. 1998, *MNRAS*, 297, 977
- Kogut, A., Banday, A. J., Bennett, C. L., Hinshaw, G. F., Górski, K. M., & Reach, W. T. 1996, *ApJ*, 460, 1
- Lagache, G. 2003, *A&A*, 405, 813
- Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *MNRAS*, 225, 307
- Lewis, A., & Bridle, S. 2002, *Phys. Rev. D*, 66, 103511
- Liddle, A. R. 2004, *MNRAS*, 351, L49
- Maino, D., Banday, A. J., Baccigalupi, C., Perrotta, F., & Górski, K. M. 2003, *MNRAS*, 344, 544
- Maino, D., et al. 2002, *MNRAS*, 334, 53
- Mather, J. C., Fixsen, D. J., Shafer, R. A., Mosier, C., & Wilkinson, D. T. 1999, *ApJ*, 512, 511
- Paladini, R., De Zotti, G., Davies, R. D., & Giard, M. 2005, *MNRAS*, 360, 1545
- Platania, P., Bensadoun, M., Bersanelli, M., de Amici, G., Kogut, A., Levin, S., Maino, D., & Smoot, G. F. 1998, *ApJ*, 505, 473
- Reach, W. T., et al. 1995, *ApJ*, 451, 188
- Reich, P., & Reich, W. 1986, *A&AS*, 63, 205
- . 1988, *A&AS*, 74, 7
- Reich, P., Reich, W., & Testori, J. C. 2003, in *The Magnetised Interstellar Medium*, ed. B. Uyaniker, W. Reich, & R. Wielebinski (Katlenburg-Lindau: Copernicus), 63
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Shaver, P. A., McGee, R. X., Newton, L. M., Danks, A. C., & Pottasch, S. R. 1983, *MNRAS*, 204, 53
- Spergel, D. N., et al. 2003, *ApJS*, 148, 175
- Stivoli, F., Baccigalupi, C., Maino, D., & Stompor, R. 2005, *MNRAS*, submitted (astro-ph/0505381)
- Stolyarov, V., Hobson, M. P., Ashdown, M. A. J., & Lasenby, A. N. 2002, *MNRAS*, 336, 97
- Stolyarov, V., Hobson, M. P., Lasenby, A. N., & Barreiro, R. B. 2005, *MNRAS*, 357, 145
- Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, *Phys. Rev. D*, 68, 123523
- Tegmark, M., & Efstathiou, G. 1996, *MNRAS*, 281, 1297
- Verde, L., et al. 2003, *ApJS*, 148, 195
- Vielva, P., Martínez-González, E., Gallegos, J. E., Toffolatti, L., & Sanz, J. L. 2003, *MNRAS*, 344, 89
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511
- Watson, R. A., et al. 2005, *ApJ*, 624, L89