

## ESTIMATION OF POLARIZED POWER SPECTRA BY GIBBS SAMPLING

D. L. LARSON,<sup>1</sup> H. K. ERIKSEN,<sup>2,3,4,5</sup> B. D. WANDELT,<sup>1,6</sup> K. M. GÓRSKI,<sup>4,5,7</sup> GREG HUEY,<sup>1</sup>  
 J. B. JEWELL,<sup>4</sup> AND I. J. O'DWYER<sup>4,5</sup>

Received 2006 July 31; accepted 2006 September 29

### ABSTRACT

Earlier papers introduced a method of accurately estimating the angular cosmic microwave background temperature power spectrum based on Gibbs sampling. Here we extend this framework to polarized data. All advantages of the Gibbs sampler still apply, and exact analysis of megapixel polarized data sets is thus feasible. These advantages may be even more important for polarization measurements than for temperature measurements. While approximate methods can alias power from the larger *E*-mode spectrum into the weaker *B*-mode spectrum, the Gibbs sampler (or equivalently, exact likelihood evaluations) allows for a statistically optimal separation of these modes in terms of power spectra. To demonstrate the method, we analyze two simulated data sets: (1) a hypothetical future CMBPol mission, with the focus on *B*-mode estimation; and (2) a *Planck*-like mission, to highlight the computational feasibility of the method.

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

### 1. INTRODUCTION

Since the first detection of cosmic microwave background (CMB) polarization (Degree Angular Scale Interferometer [DASI]; Kovac et al. 2002; Leitch et al. 2002) and subsequent measurement of the temperature-gradient (*TE*) cross-power spectrum by the *Wilkinson Microwave Anisotropy Probe* (*WMAP*; Kogut et al. 2003), emphasis has shifted to the measurement and analysis of the full polarization angular power spectra. Many experiments (Leitch et al. 2005; Readhead et al. 2004; Barkats et al. 2005; Montroy et al. 2006; Page et al. 2006) have improved on those early findings, producing measurements with a considerable gain in raw sensitivity. However, an important concern with all such measurements is systematic errors, including not only instrumental effects, observing strategy effects, and astrophysical contaminants, but also statistical issues. It is essential to develop powerful and flexible data analysis tools to extract the desired information from the raw data reliably. In this paper we progress toward this goal by extending the previously introduced Gibbs sampling framework (Jewell et al. 2004; Wandelt et al. 2004; Eriksen et al. 2004) to polarization.

The scientific importance of CMB polarization power spectra is high. For example, our current understanding of the optical depth, amplitude, and scalar spectral index hinges on what we know about the magnitude of the low- $\ell$  temperature and polarization spectra from the *WMAP* 3 year data (Page et al. 2006). In addition, a detection of large-scale *B*-modes would give a very exciting insight into primordial gravitational waves.

Earlier Gibbs analyses of unpolarized CMB data were described by Wandelt et al. (2004), O'Dwyer et al. (2004), and Eriksen et al. (2004, 2006). These efforts demonstrated that ex-

act analyses are indeed feasible even for such large data sets as the *WMAP* data, which comprise several million pixels. This is possible due to the very favorable scaling of the Gibbs sampling algorithm. While brute-force likelihood evaluations scale as  $O(N_{\text{pix}}^3)$ ,  $N_{\text{pix}}$  being the number of pixels in the data set, the Gibbs sampler scales identically to the map-making operation. For the special case of uncorrelated noise and symmetric beams, this reduces further to  $O(N_{\text{pix}}^{3/2})$ . Thus, even *Planck*-sized data can be analyzed using these tools, as is demonstrated in the present paper.

Gibbs sampling thus provides an efficient route to the exact posterior (or likelihood). Moreover, it does not rely on any ad hoc approximations. Even for the analysis of temperature data, this proved to be both an important and subtle issue (Eriksen et al. 2006). However, it is critical for polarization measurements, because well-known approximate methods such as the pseudo- $C_\ell$  methods (e.g., Chon et al. 2004) can lead to aliasing of *E*-mode power into the much smaller *B*-mode power spectrum. Although it is possible to construct ways around this problem (Smith 2006), exact methods such as full likelihood evaluations or Gibbs sampling are clearly preferable solutions.

We start by discussing the algorithms used for polarized Gibbs sampling, extending the signal and power-spectrum sampling steps from temperature to polarization. Then we analyze simulated data to verify that the algorithm works and to determine the computational efficiency of the method.

### 2. ALGORITHMS

#### 2.1. Overview of Gibbs Sampling

Gibbs sampling in the polarization case is essentially the same as in the temperature case, with objects involved in the sampling redefined to account for the additional information. For full details on the methodology of Gibbs sampling as applied to CMB analysis, see Jewell et al. (2004), Wandelt et al. (2004), and Eriksen et al. (2004).

Specifically, the CMB signal is generalized to a vector of harmonics coefficients ( $a_{\ell m}^T, a_{\ell m}^E, a_{\ell m}^B$ ) for each  $\ell$  and  $m$ , where the letters *T*, *E*, and *B* stand for temperature, electric/gradient, and magnetic/curl, respectively. The covariance matrix **S** of the CMB

<sup>1</sup> Department of Physics, University of Illinois, Urbana, IL; dlaron1@uiuc.edu.

<sup>2</sup> Institute of Theoretical Astrophysics, University of Oslo, Oslo, Norway.

<sup>3</sup> Centre of Mathematics for Applications, University of Oslo, Oslo, Norway.

<sup>4</sup> Jet Propulsion Laboratory, Pasadena, CA.

<sup>5</sup> California Institute of Technology, Pasadena, CA.

<sup>6</sup> Astronomy Department, University of Illinois, Urbana-Champaign, IL.

<sup>7</sup> Warsaw University Observatory, Warszawa, Poland.

signal then becomes block diagonal, with an identical  $3 \times 3$  sub-matrix for each  $m$ -value at a given  $\ell$ :

$$\mathbf{C}_\ell = \begin{pmatrix} C_\ell^{TT} & C_\ell^{TE} & C_\ell^{TB} \\ C_\ell^{TE} & C_\ell^{EE} & C_\ell^{EB} \\ C_\ell^{TB} & C_\ell^{EB} & C_\ell^{BB} \end{pmatrix}. \quad (1)$$

The data are pixelized maps  $\mathbf{m}$  of the Stokes parameters  $I$ ,  $Q$ , and  $U$  of the form

$$\mathbf{m} = \mathbf{A}\mathbf{s} + \mathbf{n}, \quad (2)$$

where  $\mathbf{A}$  is a linear operator that includes convolution with an instrument beam and the transformation of the  $T$ ,  $E$ , and  $B$  components of the signal  $\mathbf{s}$  into the Stokes parameters. Note that for the rest of this paper, we assume both the instrumental beam to be symmetric and the noise  $\mathbf{n}$  to be uncorrelated, having a diagonal covariance matrix  $\mathbf{N}$ . These are the reasons we can work with maps instead of time-ordered data. However, to simplify the notation we disregard in the following all issues concerning data format, beam convolutions, multifrequency observations, etc. and model our data as a simple sum of a signal term and a noise term. For the full expressions, see the Appendix.

Application of a galactic mask is implemented by increasing the noise variance to infinity for masked pixels, or rather, by setting the inverse noise covariance to zero. For full details we refer the interested reader to Eriksen et al. (2004).

As in the temperature-only case discussed in Jewell et al. (2004) and Wandelt et al. (2004), we wish to sample from the  $P(\mathbf{S}|\mathbf{d})$  posterior. It is typically not easy to evaluate  $P(\mathbf{S}|\mathbf{d})$  directly because of a large and dense  $\mathbf{S} + \mathbf{N}$  covariance matrix, nor is it easy to sample from it directly. This is precisely the motivation for Gibbs sampling, which allows sampling from a joint density through the corresponding conditional densities. For the case of CMB power spectrum estimation, this is done by first sampling from  $P(\mathbf{S}, \mathbf{s}|\mathbf{d})$ , using  $P(\mathbf{S}|\mathbf{s}, \mathbf{d})$  and  $P(\mathbf{s}|\mathbf{S}, \mathbf{d})$  (neither of which requires inversion of dense  $\mathbf{S} + \mathbf{N}$  matrices), and then marginalizing over  $\mathbf{s}$ . Using the fact that, given a full-sky signal map the conditional density for the signal matrix is independent of the data  $P(\mathbf{S}|\mathbf{s}, \mathbf{d}) = P(\mathbf{S}|\mathbf{s})$ , the basic Gibbs sampling scheme can be written in the following form:

$$\mathbf{S}^{i+1} \leftarrow P(\mathbf{S}|\mathbf{s}^i, \mathbf{d}) \quad (3)$$

$$\mathbf{s}^{i+1} \leftarrow P(\mathbf{s}|\mathbf{S}^{i+1}). \quad (4)$$

Here the arrow symbol indicates sampling from the distribution on the right-hand side. The only remaining problem is to establish the correct sampling algorithms for each of the two conditional distributions for polarized data, and this is the topic of the following sections.

Note that if a continuous distribution for  $P(\mathbf{S}|\mathbf{d})$  is desired, as opposed to a set of individual samples, one can take advantage of the known analytical form of the distribution  $P(\mathbf{S}|\mathbf{s})$  by applying the Blackwell-Rao estimator. This procedure was discussed in detail by Wandelt et al. (2004) and Chu et al. (2005) for the temperature-only case, and the generalization to polarization is once again straightforward. The required modifications are written out in § 2.5.

## 2.2. Signal Sampling

The signal-sampling equations for polarization are identical to those for temperature-only data, taking into account the generalizations mentioned above. Specifically, the sky signal ( $\mathbf{s} = \mathbf{x} + \mathbf{y}$ )

is sampled (given the current covariance matrix  $\mathbf{S}$ ) by solving for the mean field,  $\mathbf{x}$ , and fluctuation,  $\mathbf{y}$ , maps:

$$(\mathbf{1} + \mathbf{S}^{1/2} \mathbf{N}^{-1} \mathbf{S}^{1/2}) \mathbf{S}^{-1/2} \mathbf{x} = \mathbf{S}^{1/2} \mathbf{N}^{-1} \mathbf{m}, \quad (5)$$

$$(\mathbf{1} + \mathbf{S}^{1/2} \mathbf{N}^{-1} \mathbf{S}^{1/2}) \mathbf{S}^{-1/2} \mathbf{y} = \boldsymbol{\xi} + \mathbf{S}^{1/2} \mathbf{N}^{-1/2} \boldsymbol{\chi}, \quad (6)$$

where  $\boldsymbol{\xi}$  and  $\boldsymbol{\chi}$  are random maps containing Gaussian unit variates (zero mean and unit variance) in each pixel for each of the  $I$ ,  $Q$ , and  $U$  components.<sup>8</sup> Note that the symbols in these equations can be interpreted either in terms of pixel-space or spherical-harmonic-space objects. In practice, this is implemented in terms of conversions between pixel and harmonic space with standard spherical harmonics transforms. For example, the inverse noise covariance matrix is given by  $\mathbf{N}^{-1}$  in pixel space and  $\mathbf{Y}^T \mathbf{N}^{-1} \mathbf{Y}$  in harmonic space, where  $\mathbf{Y}$  and  $\mathbf{Y}^T$  are the inverse and standard spherical harmonics transforms, respectively. For explicit details on such computations, see Eriksen et al. (2004).

The signal-sampling operation is by far the most demanding step of the Gibbs sampler, because it requires the solution of a very large linear system. Formally speaking, this corresponds to inverting an  $\sim 10^6 \times 10^6$  matrix, which clearly is not computationally feasible through brute-force methods. However, as described in detail by, e.g., Eriksen et al. (2004), the systems in equations (5) and (6) can be solved by means of conjugate gradients (CGs). The computational scaling is thus reduced to the most expensive step for applying the operator on the left-hand side of the equations, which for symmetric beams and uncorrelated noise is a standard spherical harmonic transform.

The efficiency of the CG technique depends critically on the condition number of the matrix under consideration. For our case, this is simply the highest signal-to-noise ratio of any mode in the system. As an example, for a fixed preconditioner it takes about 60 iterations to solve for the first-year *WMAP* data, about 120 iterations to solve for the three-year *WMAP* data, and about 300 iterations to solve for the *Planck* 100 GHz data.

This is a particularly serious issue for CMB polarization measurements. While these signatures by themselves have a very low signal-to-noise ratio, and therefore should be easy to determine on their own, the corresponding signal-to-noise ratio for temperature is tremendous. Consequently, if a main goal is to estimate the  $TE$  cross spectrum, by far most of the CPU time is spent on temperature map convergence. On the other hand, if all interest lies in  $E$ - and  $B$ -modes, the temperature data can be disregarded completely [or alternatively conditioned on by sampling from  $P(a_{\ell m}^E, a_{\ell m}^B | \mathbf{d}, a_{\ell m}^T)$ ], and convergence is then achieved rapidly even for CMBPol type missions. This will be explicitly demonstrated in § 3.1.

It is possible to reduce the computational expense of a CG search significantly by preconditioning. One approach that has proved successful so far is to precompute a subset of the coefficient matrix in equations (5) and (6) and multiply both sides of the equations by the inverted submatrix. Thus, by inverting the most problematic parts of the matrix by hand, the effective condition number is greatly reduced, and significant speed-up can be achieved.

Currently, our preconditioner is constructed independently for the temperature and polarization states. For the polarization components, it is a diagonal matrix in  $EE$  and  $BB$  independently,

<sup>8</sup> Note that  $\mathbf{S}^{1/2}$  and  $\mathbf{S}^{-1/2}$  must be symmetric for these equations to be valid. On the other hand,  $\mathbf{N}^{-1/2}$  only has to satisfy  $\mathbf{N}^{-1/2} (\mathbf{N}^{-1/2})^T = \mathbf{N}^{-1}$  and can be chosen to be the Cholesky decomposition.

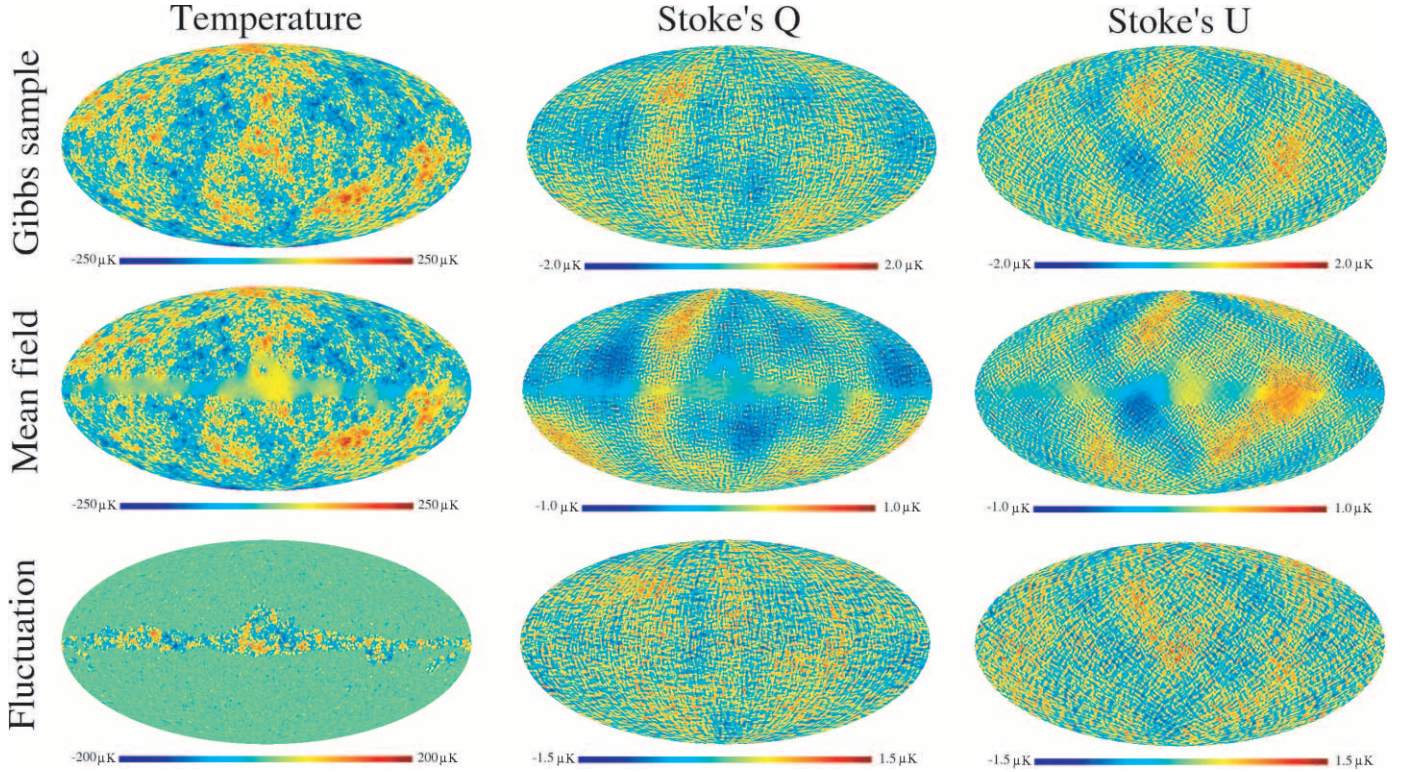


FIG. 1.—Gibbs sampled signal maps. The three columns show, from left to right, temperature and Stokes  $Q$ - and  $U$ -parameters. The three rows show, from top to bottom, the complete Gibbs samples, the mean field (Wiener-filtered) maps, and the fluctuation maps. The mean field map provides the information content of the data, and the fluctuation map provides a random complement such that the sum of the two is a full-sky, noiseless sky consistent with both the current power spectrum and the data.

while for the  $TT$  correlations, it consists of a low- $\ell$  matrix that includes all coefficients up to some  $\ell_{\max}$  and then the diagonal elements at higher  $\ell$ 's (Eriksen et al. 2004). For *WMAP*-type applications we typically used  $\ell_{\max} = 50$ , which requires 52 MB of memory and about 1 minute of CPU time for inversion. For upcoming *Planck* data, it will be desirable to use a significantly larger preconditioner, and more realistic numbers are  $\ell_{\max} \sim 150$  or 200. This will require extensive parallelization and has not yet been implemented in our codes. We therefore still use a serial preconditioner up to  $\ell_{\max} = 70$  in this paper and pay the extra cost in CG iterations.

### 2.3. Power Spectrum Sampling

Given the (full-sky) signal polarization map sampled from  $P(\mathbf{s}|\mathbf{S}, \mathbf{d})$  as described above, we must sample the signal covariance matrix from  $P(\mathbf{S}|\mathbf{s})$ , which is explicitly given by

$$P(\mathbf{S}|\mathbf{s}) \propto \prod_{\ell} \frac{1}{\sqrt{|\mathbf{C}_{\ell}|^{2\ell+1+2q}}} \exp\left(-\frac{1}{2} \text{tr} \sigma_{\ell} \mathbf{C}_{\ell}^{-1}\right). \quad (7)$$

Here we have assumed a prior of the form  $P(\mathbf{S}) \propto \prod_{\ell} |\mathbf{C}_{\ell}|^{-q}$  (i.e.,  $q = 0$  for a uniform prior, and  $q = 1$  for a Jeffreys prior), and we have defined

$$\sigma_{\ell} = \sum_{m=-\ell}^{\ell} \mathbf{s}_{\ell m} \mathbf{s}_{\ell m}^{\dagger}. \quad (8)$$

Each  $\mathbf{s}_{\ell m}$  represents the signal in harmonic space and is a three-dimensional complex-valued column vector containing the coefficients for the  $T$ -,  $E$ -, and  $B$ -modes at that  $\ell$  and  $m$ . This

distribution is known as the inverse Wishart distribution (Gupta & Nagar 2000).

Sampling from this conditional density can be done with a vector generalization of the sampling algorithm described in Wandelt et al. (2004). If  $p \times p$  is the size of the matrix being sampled (typically  $p = 3$  for polarization), then the required steps for sampling are: (1) sample  $n = 2\ell - p + 2q$  vectors from a Gaussian with covariance matrix  $\sigma_{\ell}$ , (2) compute the sum of outer products of these independently sampled vectors, and (3) invert this matrix. For full details on both the inverse Wishart distribution and the sampling algorithm, we refer the interested reader to chapter 3 of Gupta & Nagar (2000).

There is a caveat for  $\ell = 2$ . The Wishart distribution, from which we derive our sampling algorithm, is defined only if  $n \geq p$ ; if not, the sampled matrix is singular. This is a problem for  $\ell = 2$  and a flat prior, since we would only sample one vector to form a  $3 \times 3$  matrix. Thus, the algorithm breaks down for this particular case. Fortunately, this is not a major problem in practice. Three straightforward solutions are: (1) sample the  $2 \times 2$   $TE$  block and the  $B$  block of the matrix separately, assuming no  $TB$  or  $EB$  correlations; (2) use a Jeffrey's prior ( $q = 1$ ); or (3) bin the quadrupole and octopole together. Note that all other multipoles can be sampled individually by the above algorithm without modifications.

#### 2.3.1. Binning

As discussed by Eriksen et al. (2006) it is highly desirable for the Gibbs sampler to be able to bin several power-spectrum multipoles together. The main advantage of this is improved sampling efficiency. As currently implemented, the step size taken between two consecutive Gibbs samples is given by cosmic variance alone. The full posterior, however, is given by both cosmic variance and noise. Therefore, in the low signal-to-noise regime, one must take






FIG. 2.—Close-up of the Galactic center shown in Fig. 1, emphasizing how the algorithm separates  $E$ - and  $B$ -modes. Each of the sampled maps (the sum of the fluctuation and mean field map) are full-sky maps, so decomposing the polarization into  $E$ - and  $B$ -modes is straightforward. The images show temperature as color and polarization overlaid as a fingerprint pattern of stripes. The stripes are aligned with the direction of polarization. They are darkest where the polarization is strongest, and they disappear where the polarization goes to zero (Cabral & Leedom 1993). The maximum amplitude of the polarization is given in  $\mu\text{K}$  and centered under each image. The maps have been smoothed to  $1^\circ$ . This is an orthogonal projection of the sky, about  $60^\circ$  wide, centered on the Galactic center. The *WMAP* Kp0 galactic mask is visible in the fluctuation and mean terms.

a larger number of steps to obtain two independent samples. The easiest way of improving on this is simply to bin many multipoles together and thereby increase the signal-to-noise ratio of the power-spectrum coefficient. In practice we choose bins such that the signal-to-noise ratio is always larger than some limit, say 3.

Since the CMB power spectrum is roughly proportional to  $1/\ell(\ell+1)$ , it is convenient to define uniform bins in  $C_\ell\ell(\ell+1)$ . We therefore redefine  $\sigma_\ell$  for bin  $b = [\ell_{\min}, \ell_{\max}]$  as

$$\sigma_\ell = \sum_{\ell \in b} \sum_{m=-\ell}^{\ell} \ell(\ell+1) s_{\ell m} s_{\ell m}^\dagger. \quad (9)$$

Note that there are now

$$M = \sum_{\ell \in b} (2\ell+1) = (\ell_{\max}+1)^2 - \ell_{\min}^2 \quad (10)$$

independent spherical harmonic modes contributing to this power-spectrum coefficient. Thus, the inverse Wishart distribution has  $n = M - p - 1 + 2q$  degrees of freedom rather than  $n = 2\ell -$

$p + 2q$ . With this modification, the basic sampling algorithm remains unchanged, but since we have sampled  $\mathbf{C}_b = \ell(\ell+1)\mathbf{C}_\ell$  and not  $\mathbf{C}_\ell$ , the actual power spectrum coefficients are given by  $\mathbf{C}_\ell = \mathbf{C}_b/\ell(\ell+1)$  for each  $\ell$  in bin  $b$ .

#### 2.4. Separation of $E$ - and $B$ -Modes

We now make a brief comment on the so-called  $E$ - $B$  coupling problem that plagues most approximate methods, such as the pseudo- $C_\ell$  methods (see, e.g., Smith 2006). Briefly put, the problem lies in the fact that the spherical harmonics are not orthogonal on a cut sky, and this may result in leakage from the (much larger)  $E$ -mode power into the  $B$ -mode power spectrum.

Exact methods such as exact likelihood analyses or Gibbs sampling do not have this problem. This can be understood intuitively in terms of the signal-sampling process illustrated in Figures 1 and 2. Obtaining a complete sky sample for the Gibbs sampler is a two-step process. First, one filters out as much information as possible from the observed data using a Wiener filter. Second, one replaces the lost power due to noise and partial-sky coverage by a random fluctuation term. The sum of the two is

a full-sky, noiseless sample that is consistent with the data. Because it is a full-sky sample, no  $E$ - $B$  coupling arises.

### 2.5. Blackwell-Rao Estimator

The Gibbs sampler provides a set of samples of the signal covariance matrix  $\mathbf{S}$ . In practice, it is often preferable to have a smooth description of the probability density of  $\mathbf{S}$ . In such cases, one can use the Blackwell-Rao estimator, which takes advantage of the known analytical form of the probability distribution  $P(\mathbf{S}|\mathbf{s})$  and uses the set of signal samples  $\{\mathbf{s}\} = \{\mathbf{s}^1, \dots, \mathbf{s}^k\}$  to approximate  $P(\mathbf{S}|\mathbf{d})$ .

An intuitive understanding of the Blackwell-Rao estimator can be found in terms of the usual Gibbs sampling algorithm. Within the theory of Gibbs sampling (or more generally Markov chain Monte Carlo), it is perfectly valid to sample one parameter more often than others, so long as the sampling scheme is independent of the current “state” of the Markov chain. In particular, one can choose to sample  $\mathbf{S}$  one thousand times for each time one samples  $\mathbf{s}$ , and thereby obtain more power-spectrum samples (although not sky-signal samples) with negligible cost. The result is a smooth power-spectrum histogram. The Blackwell-Rao estimator takes this idea to the extreme and replaces the power-spectrum sampling step by the corresponding analytical distribution. The result is a highly accurate and smooth description of  $P(\mathbf{S}|\mathbf{d})$  that is very useful for, say, estimation of cosmological parameters (Wandelt et al. 2004; Chu et al. 2005; Eriksen et al. 2006).

For full details on this estimator for the temperature-only case, we refer the interested reader to Chu et al. (2005). But again, the generalization to polarization is indeed straightforward, and the generalized estimator reads

$$P(\mathbf{S}|\{\mathbf{s}\}) \propto \sum_j \prod_{\ell} \frac{1}{\sqrt{|\sigma_{\ell}|^{2\ell+2q-p}} \sqrt{|\mathbf{C}_{\ell}|^{2\ell+1+2q}}} \times \exp\left(-\frac{1}{2} \text{tr} \sigma_{\ell}^j \mathbf{C}_{\ell}^{-1}\right), \quad (11)$$

where  $\mathbf{C}_{\ell}$  is the  $3 \times 3$  submatrix of  $\mathbf{S}$  for a given  $\ell$ , and  $j$  is the index over Gibbs samples.

### 3. APPLICATION TO SIMULATED DATA

We now apply the methodology described in § 2 to simulated data. Two different cases are considered to highlight different features. In the first, we consider a low-resolution, high signal-to-noise ratio experiment aimed at detecting primordial  $B$ -modes. The main goal of this exercise is to demonstrate the fact that the so-called  $E$ - $B$  coupling problem that plagues approximate methods is not an issue for exact methods. Second, we consider a high-resolution simulation based on the *Planck* 100 GHz channel to demonstrate that Gibbs sampling is feasible even for very large CMB data sets.

#### 3.1. Low-Resolution $B$ -Mode Experiment (CMBPol)

Our first case corresponds to a possible future mission targeting the primordial  $B$ -modes that arise during the inflationary period. Such modes are expected to have a very low amplitude and to be limited to large angular scales. Some case studies for a  $B$ -mode mission therefore emphasize extreme sensitivity over angular resolution, and we adopt similar characteristics for this exercise.

As discussed in § 2.2, the convergence ratio for the CG search depends critically on the signal-to-noise ratio of the data. In order to achieve acceptable performance when analyzing temperature observations with the sensitivity required for detecting  $B$ -modes, a much better preconditioner than what we have currently im-

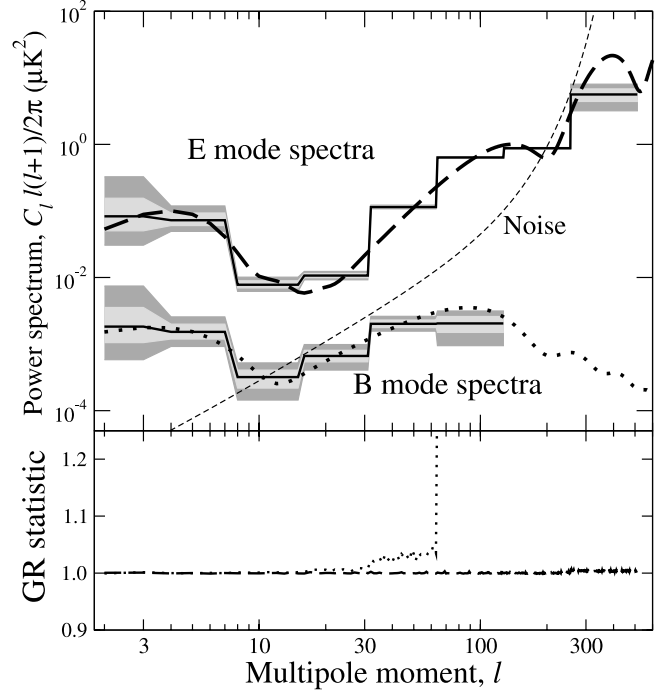


FIG. 3.—Reconstructed  $E$ - and  $B$ -mode power spectra from the low-resolution analysis. Input spectra are shown as dashed and dotted lines, respectively, while the reconstructed posterior distributions are indicated by solid lines (posterior maximum) and gray regions (1 and 2  $\sigma$  confidence regions). The corresponding noise spectrum is given by a thin dashed line. The Gelman-Rubin convergence statistic as a function of multipole is shown in the bottom frame.

plemented is required. We therefore only consider the  $E$ - and  $B$ -mode spectra here and not the temperature spectrum.

The simulated data set consists of the sum of a CMB component and a white-noise component. The CMB realization was drawn in harmonic space from a Gaussian distribution with a  $\Lambda$ CDM spectrum (downloaded from the WMAP3 parameter table at the Legacy Archive for Microwave Background Data Analysis [LAMBDA]) having a tensor contribution of  $r \simeq 0.03$ . Multipoles up to  $\ell_{\max} = 512$  were included. This realization was then convolved with a  $1^\circ$  FWHM Gaussian beam and  $N_{\text{side}} = 256$  pixel window and projected onto a HEALPix<sup>9</sup> grid. Next, uniform (and uncorrelated between  $Q$  and  $U$ ) noise of  $1 \mu\text{K rms}$  was added to each pixel. Finally, the WMAP3 polarization mask (Page et al. 2006) was applied, removing 26.5% of the sky from the analysis.

We adopted a binning scheme logarithmic in  $\ell$ , such that  $b_i = [2^i, 2^{i+1} - 1]$ . Note that this is not directly connected to the signal-to-noise ratio of the data themselves, and this will have consequences for the convergence properties of the high- $\ell$   $B$ -mode bins. However, our main focus in this paper is the method itself, and this scheme is chosen to illustrate the effect of both high and low signal-to-noise binning, not to obtain an optimal power spectrum.

The simulation was then analyzed with the Gibbs sampler described earlier, producing 1000 sky samples in each of five independent Markov chains. The CPU cost for producing one sample was 10 minutes, or a wall-clock time of 2.5 minutes when parallelized over four processors. The total running time was thus 42 hr using 20 processors. For each sky sample, 20 independent power-spectrum samples were drawn in order to obtain smoother  $C_{\ell}$  confidence regions. (See the discussion of the Blackwell-Rao estimator in § 2.5 for more details.)

<sup>9</sup> See <http://healpix.jpl.nasa.gov>.

We first consider the reconstructed autospectra, which are shown in the top panel of Figure 3. The input (unbinned) spectra are given by dashed and dotted lines for  $E$ - and  $B$ -modes, respectively, and the reconstructed (binned) posterior maximum spectra are shown by solid lines. Gray shading marks 1 and 2  $\sigma$  confidence regions. Finally, the beam deconvolved noise spectrum is indicated by a thin dashed line.

In the bottom panel we show the Gelman-Rubin convergence statistic (Gelman & Rubin 1992), as computed from the  $\sigma_\ell$  sky signal power spectra for each  $\ell$ . This is much more conservative than computing the same statistic from the  $C_\ell$  samples for two reasons. First, convergence in the binned power spectrum is achieved faster than convergence in each sky mode. Second, cosmic variance only contributes to the power spectrum and not the signal on the sky. Therefore, this may be accounted for either analytically through the Blackwell-Rao estimator or by resampling the  $C_\ell$  spectra given the  $\sigma_\ell$ 's. In other words, a small error in the sky signal variance does not affect the full posterior significantly if the desired distribution is anyway dominated by cosmic variance.

A general recommendation is that the Gelman-Rubin statistic  $R$  should be less than 1.1 or 1.2 to claim convergence, although the value depends on the particular application and initialization procedure, and should be compared against other methods such as jackknife tests. However, for the particular case shown in Figure 3, it is clear that the  $E$ -mode spectrum has converged very well everywhere, while the  $B$ -mode spectrum only has converged up to  $\ell \approx 60$ .

As discussed in § 2.3, this behavior can be understood intuitively in terms of signal-to-noise ratio. Since the step size between two signal samples is given by cosmic variance alone, while the full posterior distribution is given by both cosmic variance and noise, it takes a large number of Gibbs steps to diffuse efficiently in the very low signal-to-noise regime. Further, the noise spectrum is about 3 orders of magnitude larger than the  $B$ -mode spectrum at  $\ell \gtrsim 100$ , and the Gibbs sampler is therefore unable to probe the full distribution with a reasonable number of samples.

To resolve this issue we bin the power spectrum. However, the binning scheme was not tuned to obtain constant signal-to-noise in each bin but was rather arbitrary. The result is clearly seen in the Gelman-Rubin statistic. For  $\ell \lesssim 60$  the signal-to-noise per bin is high, and convergence is excellent. At  $\ell \gtrsim 60$ , it is low, and the convergence is very poor. The way to resolve this would have been to choose larger bins at higher  $\ell$ 's.

In Figure 4, we show the  $E \times B$  cross spectrum. As expected, this is nicely centered on zero.

We end this section by commenting on the applicability of this formalism to a possible future CMBPol-type mission. As is well known, the main problems for such a mission will not be primarily statistical issues of the type discussed above, but rather systematics in various forms. Two important examples are correlated noise and asymmetric beams. However, if it is possible to precompute the complete  $\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}$  matrix for the data set under consideration, then these two important effects can be fully accounted for using the methods described here. And for a low-resolution CMBPol mission this may be possible. For an upper multipole limit of, say,  $\ell_{\max} = 300$  there is a total of  $2 \times 90,000 = 180,000$  polarized spherical modes to account for. In other words, one has to store and invert a  $180,000 \times 180,000$  matrix in order to analyze such an experiment exactly. Although this is a considerable computational problem, it is quite tractable already with current computers. Thus, if it is possible to compute this matrix in the first place for a given experiment, an exact and




FIG. 4.—Plot of the  $E \times B$  cross spectrum from the low-resolution analysis.

complete analysis is feasible using the methods described in this paper.

### 3.2. High-Resolution $T+E$ Experiment (Planck)

In order to demonstrate the feasibility of this method for analyzing even the largest planned data set, we now consider a simulation with properties similar to those of the *Planck* 100 GHz instrument. Specifically, the grid resolution is chosen to be  $N_{\text{side}} = 1024$  (corresponding to a  $3.4'$  pixel size), the maximum multipole moment is  $\ell_{\max} = 1500$ , the beam size is  $9.5'$  FWHM, and the noise level is  $38.2 \mu\text{K}$  rms per pixel for temperature and  $61 \mu\text{K}$  per  $Q/U$  pixel. These noise levels are a factor of 2 higher than the goal levels for the 100 GHz channel given in the *Planck* bluebook.<sup>10</sup> No noise correlations between  $T$ ,  $Q$ , and  $U$  were included, the  $B$ -mode spectrum was set to zero, and the sky cut was chosen to be the *WMAP* Kp2 mask.

Such large data sets are certainly a challenge for the Gibbs sampling algorithm, and the computational requirements are considerable. Specifically, the CPU time for generating one sample (requiring  $\sim 250$ – $300$  CG iterations) is about 16 CPU hr when using the low- $\ell$  preconditioner described by Eriksen et al. (2004) up to  $\ell = 70$ . Better preconditioners will of course reduce this cost significantly.

However, it is important to note that even though this is an expensive operation, it is by no means prohibitive. To obtain a reasonably well-converged posterior distribution, one requires on the order of  $\sim 10^3$  independent samples, and this would then require  $\sim 10^4$  CPU hr. Of course, this number must be multiplied with a significant factor for an actual production analysis (e.g., number of frequency bands or data combinations), but considering the tremendous efforts spent on obtaining the *Planck* data in the first place, this amount of CPU time is a most reasonable cost for analyzing them.

For the high-resolution analysis presented in this paper, we produced a total of 800 sky samples, divided over eight independent chains. Again, 20 independent power-spectrum samples were then drawn from each of these for visualization purposes. The results from these computations are summarized in Figure 5, showing both the reconstructed power spectra and the corresponding convergence statistics.

With the chosen binning scheme and number of samples, we see that the  $TT$  spectrum has converged well everywhere, while

<sup>10</sup> See [http://www.rssd.esa.int/SA/PLANCK/docs/Bluebook-ESA-SI\(2005\)1\\_V2.pdf](http://www.rssd.esa.int/SA/PLANCK/docs/Bluebook-ESA-SI(2005)1_V2.pdf).




FIG. 5.—Reconstructed power spectra from the high-resolution *Planck* 100 GHz simulation. The true spectra are shown as dashed lines, and the reconstructed posterior distributions are given by a maximum posterior value (solid lines) and 68% and 95% confidence regions. The Gelman-Rubin convergence statistics are shown in the bottom frames.

the *TE* spectrum has some small problems at the end of the second bin. The *EE* spectrum would clearly have benefited from more samples, and even more importantly, slightly larger bins; increasing the bin size by, say, 20% would have resolved both the *TE* and *EE* issues.

However, as far as computational feasibility goes, the important part is the signal-sampling step and not binning or re-sampling issues; these can always be adjusted given some crude knowledge of the data set under consideration. Therefore, the fact that already this first implementation of the polarized Gibbs sampler is able to produce hundreds of sky samples with only a few days on a standard computer cluster is a direct demonstration of computational feasibility for even *Planck*-sized data sets.

#### 4. CONCLUSIONS

This paper extends the Gibbs sampling technique to polarized power spectrum estimation. We have detailed the necessary generalization steps relative to the original temperature-only descriptions given by Jewell et al. (2004), Wandelt et al. (2004), and Eriksen et al. (2004) and have considered computational aspects of polarized analysis.

The algorithm was demonstrated with two specific examples. First, considering a possible CMBPol-type mission, we showed that the Gibbs sampler cleanly separates *E*- and *B*-modes, and no special care is required. This is in sharp contrast to approximate methods such as the so-called pseudo- $C_\ell$  method, for which great care must be taken in order for the larger *E*-modes not to compromise the minute *B*-modes. Second, we analyzed a *Planck*-sized data set, demonstrating that the algorithm is useful for analyzing the quantity of data that will come from near-future CMB experiments.

The Gibbs sampling results presented here use symmetric beams and noise that is uncorrelated between pixels. However, the Gibbs sampling algorithm has potential to analyze considerably more complicated data sets than these. For *Planck*, the so-

lution lies in exploiting the very regular scanning strategy, which reduces the computational burden of a time-ordered data analysis. For a future CMBPol mission, the solution lies in the relatively large angular scales required. Since it is possible to invert the noise covariance matrix for multipoles up to several hundreds, one can precompute the all-important  $\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}$  matrix. After paying this high one-time cost, efficient and exact analysis is feasible using the methods described in this paper.

Finally, we reemphasize that the Gibbs sampler provides a direct route to the exact likelihood (and to the Bayesian posterior), and it is much more reliable than approximate methods. This issue has been demonstrated explicitly through the analysis of the 3 year *WMAP* data, where an approximate likelihood between  $\ell = 13$  and 30 caused a nonnegligible bias in the spectral index  $n_s$  (Eriksen et al. 2006). Using Gibbs sampling, such worries are greatly reduced. Further, this paper demonstrates that the method is in fact capable of analyzing the amount of data that will come from the *Planck* mission with reasonable computational resources. It therefore seems very likely that this method will play a significant role in the analysis of future *Planck* data.

We thank Kendrick Smith, Graca Rocha, and Charles Lawrence for useful and interesting discussions. We acknowledge use of the HEALPix software (Górski et al. 2005) and analysis package for deriving the results in this paper. We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). This work was partially performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. H. K. E. acknowledges financial support from the Research Council of Norway. B. D. W. and D. L. L. acknowledge support through NSF grant AST 05-07676 and NASA JPL subcontract 1236748.

#### APPENDIX

##### SIGNAL SAMPLING

Equations (5) and (6) have been written as simply as possible for clarity when describing the extension of our Gibbs sampling algorithm to polarization. A more realistic treatment will involve multiple channels, symmetric beams, the pixel window function, and a cutoff at some value of  $\ell$ . In this appendix we write out, for reference, the log likelihood for  $s$  and the sampling equations that one derives from this.



Let the index  $i$  run over channels. Let  $\mathbf{B}_i$  be the beam-smoothing function, and  $\mathbf{W}$  be the HEALPix pixel window-smoothing function. If all channels are at the same resolution, then there is only one pixel window function; otherwise  $\mathbf{W}$  will need an  $i$  index as well. Let  $\mathbf{P}$  be a projection operator that removes all modes with  $\ell$  above some cutoff. Note that  $\mathbf{P}$ ,  $\mathbf{W}$ , and  $\mathbf{B}_i$  all commute, and  $\mathbf{P}$  commutes with  $\mathbf{S}$ . As before,  $\mathbf{m}_i$  are the maps and  $s$  is the signal. For generality, we also include a foreground component  $f_i$ , which is not otherwise discussed in this paper. We obtain

$$-2 \log P(s|\mathbf{m}_i, f_i, \mathbf{S}, \mathbf{N}_i, \mathbf{B}_i, \mathbf{W}) = s^T \mathbf{S}^{-1} s + \sum_i (\mathbf{m}_i - \mathbf{B}_i \mathbf{W} s - \mathbf{B}_i \mathbf{W} f_i)^T \mathbf{P} \mathbf{N}_i^{-1} \mathbf{P} (\mathbf{m}_i - \mathbf{B}_i \mathbf{W} s - \mathbf{B}_i \mathbf{W} f_i) + \text{const.} \quad (\text{A1})$$

From the above equation, it is clear that  $\mathbf{W}$  can be absorbed into  $\mathbf{B}_i$ , so we do this and drop  $\mathbf{W}$  from the equations. The equations for sampling  $s = \mathbf{x} + \mathbf{y}$  become

$$\left[ \mathbf{P} + \mathbf{P} \mathbf{S}^{1/2} \sum_i (\mathbf{B}_i \mathbf{P} \mathbf{N}_i^{-1} \mathbf{P} \mathbf{B}_i) \mathbf{S}^{1/2} \mathbf{P} \right] \mathbf{S}^{-1/2} \mathbf{P} \mathbf{x} = \mathbf{P} \mathbf{S}^{1/2} \sum_i \mathbf{B}_i \mathbf{P} \mathbf{N}_i^{-1} \mathbf{P} (\mathbf{m}_i - \mathbf{B}_i f_i), \quad (\text{A2})$$

$$\left[ \mathbf{P} + \mathbf{P} \mathbf{S}^{1/2} \sum_i (\mathbf{B}_i \mathbf{P} \mathbf{N}_i^{-1} \mathbf{P} \mathbf{B}_i) \mathbf{S}^{1/2} \mathbf{P} \right] \mathbf{S}^{-1/2} \mathbf{P} \mathbf{y} = \mathbf{P} \boldsymbol{\xi} + \mathbf{P} \mathbf{S}^{1/2} \sum_i \mathbf{B}_i \mathbf{P} \mathbf{N}_i^{-1/2} \boldsymbol{\chi}_i, \quad (\text{A3})$$

where now we have several maps  $\boldsymbol{\chi}_i$  of Gaussian unit variates. Recall that these equations require the square root of  $\mathbf{S}$  to be symmetric.

#### REFERENCES

- Barkats, D., et al. 2005, *ApJ*, 619, L127  
 Cabral, B., & Leedom, L. 1993, in *Proc. 20th Ann. Conf. on Computer Graphics and Interactive Techniques*, ed. M. C. Whittom (New York: ACM), 263  
 Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *MNRAS*, 350, 914  
 Chu, M., Eriksen, H. K., Knox, L., Górski, K. M., Jewell, J. B., Larson, D. L., O'Dwyer, I. J., & Wandelt, B. D. 2005, *Phys. Rev. D*, 71, 103002  
 Eriksen, H. K., et al. 2004, *ApJS*, 155, 227  
 ———. 2006, *ApJ*, 641, 665  
 Gelman, A., & Rubin, D. 1992, *Stat. Sci.*, 7, 457  
 Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759  
 Gupta, A. K., & Nagar, D. K. 2000, *Matrix Variate Distributions* (Boca Raton: Chapman & Hall)  
 Jewell, J., Levin, S., & Anderson, C. H. 2004, *ApJ*, 609, 1  
 Kogut, A., et al. 2003, *ApJS*, 148, 161  
 Kovac, J. M., Leitch, E. M., Pryke, C., Carlstrom, J. E., Halverson, N. W., & Holzapfel, W. L. 2002, *Nature*, 420, 772  
 Leitch, E. M., Kovac, J. M., Halverson, N. W., Carlstrom, J. E., Pryke, C., & Smith, M. W. E. 2005, *ApJ*, 624, 10  
 Leitch, E. M., et al. 2002, *Nature*, 420, 763  
 Montroy, T. E., et al. 2006, *ApJ*, 647, 813  
 O'Dwyer, I. J., et al. 2004, *ApJ*, 617, L99  
 Page, L., et al. 2006, *ApJ*, submitted (astro-ph/0603450)  
 Readhead, A. C. S., et al. 2004, *Science*, 306, 836  
 Smith, K. M. 2006, *Phys. Rev. D*, 74, 083002  
 Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511