ANALYSIS OF *WMAP* 7 YEAR TEMPERATURE DATA: ASTROPHYSICS OF THE GALACTIC HAZE

Davide Pietrobon[1], Krzysztof M. Górski[1,2], James Bartlett[1,3], A. J. Banday[4,5], Gregory Dobler[6],
Loris P. L. Colombo[1,7], Sergi R. Hildebrandt[8], Jeffrey B. Jewell[1], Luca Pagano[1], Graca Rocha[1],
Hans Kristian Eriksen[9,10], Rajib Saha[11], and Charles R. Lawrence[1]

[1] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA; davide.pietrobon@jpl.nasa.gov
[2] Warsaw University Observatory, Aleje Ujazdowskie 4, 00478 Warszawa, Poland
[3] Laboratoire AstroParticule & Cosmologie (APC), Universiè Paris Diderot, CNRS/IN2P3, CEA/lrfu, Observatoire de Paris,
Sorbonne Paris Cité, 10, rue Alice Domon et Léonie Duquet, F-75205 Paris Cedex 13, France
[4] Universié de Toulouse, UPS-OMP, IRAP, Toulouse, France
[5] CNRS, IRAP, 9 Av. Colonel Roche, BP 44346, F-31028 Toulouse Cedex 4, France
[6] Kavli Institute for Theoretical Physics, University of California, Santa Barbara Kohn Hall, Santa Barbara, CA 93106, USA
[7] USC Dana and David Dornsife College of Letters, Arts and Sciences, University of Southern California, University Park Campus, Los Angeles, CA 90089, USA
[8] Division of Physics, Mathematics and Astronomy, California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125, USA
[9] Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029, Blindern, N-0315 Oslo, Norway
[10] Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053, Blindern, N-0316 Oslo, Norway
[11] Physics Department, Indian Institute of Science Education and Research Bhopal, Bhopal, MP 462023, India

ABSTRACT

We perform a joint analysis of the cosmic microwave background (CMB) and Galactic emission from the *WMAP* 7 year temperature data. Using the Commander code, based on Gibbs sampling, we simultaneously derive the CMB and Galactic components on scales larger than $1°$ with improved sensitivity over previous work. We conduct a detailed study of the low-frequency Galactic foreground, focusing on the "microwave haze" emission around the Galactic center. We demonstrate improved performance in quantifying the diffuse Galactic emission when including Haslam 408 MHz data and when jointly modeling the spinning and thermal dust emission. We examine whether the hypothetical Galactic haze can be explained by a spatial variation of the synchrotron spectral index, and find that the excess of emission around the Galactic center is stable with respect to variations of the foreground model. Our results demonstrate that the new Galactic foreground component—the microwave haze—is indeed present.

Key words: cosmic microwave background – cosmology: observations – diffuse radiation – Galaxy: center – methods: data analysis – methods: numerical – methods: statistical

Online-only material: color figures

1. INTRODUCTION

While data from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*; see Jarosik et al. 2011; Komatsu et al. 2011 and references therein) have enabled unprecedented advances in the understanding of cosmology over the past decade, it has also opened a unique window into the fundamental physical processes of the interstellar medium. The choice of observing bands for *WMAP* ensured that multiple emission mechanisms would be observed across the frequency coverage. In particular, there are at least three distinct physical processes at low frequencies (23–33–41 GHz): free–free, synchrotron, and anomalous microwave emission (AME), which falls with frequency and is highly correlated with $100\,\mu$m thermal dust emission. At high frequencies (61–94 GHz), Galactic emission is completely dominated by thermal dust emission (Gold et al. 2011).

Free–free emission (or thermal bremsstrahlung) is generated by scattering of ionized electrons off the proton nuclei in hot ($\sim$5000 K) gas, and has a brightness temperature that scales as $T \propto \nu^{-2.15}$ through the *WMAP* bands, where ν represents frequency. The bulk of the synchrotron emission observed by *WMAP* is seen to closely follow a power law $T \propto \nu^{-3}$ (Kogut et al. 2007; Dobler 2012). Lastly, spinning dust (whose presence has been observed in small, dusty clouds (Casassus et al. 2008; Planck Collaboration 2011), as well as hotter diffuse regions (Dobler & Finkbeiner 2008b; Dobler et al. 2009)) is the likely cause of the AME. Small dust grains with non-zero dipole moments are spun up by a variety of mechanisms such as ion

collisions, plasma density fluctuations, photon fields, etc., and produce spinning dipole radiation (see Erickson 1957; Draine & Lazarian 1998 for the original theoretical realization of this idea).

Using simple template regression techniques (see Section 3), Bennett et al. (2003) and Finkbeiner (2004) showed that the Galactic emissions are highly spatially correlated with maps at other frequencies. Free–free emission is morphologically correlated with Hα recombination line emission, synchrotron with low-frequency radio emission (e.g., at 408 MHz), and spinning (and thermal) dust with total dust column density (e.g., Schlegel et al. 1998 evaluated using models for the thermal emission to 94 GHz by Finkbeiner et al. 1999). After removing emission correlated with these templates, Finkbeiner (2004) found that there was an excess signal centered on the Galactic center (GC) and extending out roughly $\sim$30 deg. A more detailed study of this Galactic "haze" by Dobler & Finkbeiner (2008a) (and more recently with the seven year data by Dobler 2012) showed that its spectrum ($T \propto \nu^{\beta}$ with $\beta \sim -2.5$) was too soft to be free–free emission and too hard to be synchrotron emission associated with acceleration by supernovae shocks (after taking into account cosmic-ray diffusion). The origin of this residual component remains a mystery and has been the object of intense theoretical scrutiny (e.g., Hooper et al. 2007; Cholis et al. 2009; Biermann et al. 2010; Dobler et al. 2011; Crocker et al. 2011; Guo et al. 2011).

However, the subject of the haze has not been without controversy, most notably due to the claim by the *WMAP*

team (as well as others, e.g., Dickinson et al. 2009) that the haze is not detected in their analyses. In particular, Gold et al. (2011) claimed lack of evidence for the haze based on *WMAP* polarization data, though it is unclear that a polarization signal would be detectable (see Dobler 2012). A comprehensive study of component separation (and cosmic microwave background (CMB) cleaning) was performed by Eriksen et al. (2006, 2007b) and Dickinson et al. (2009) utilizing Bayesian inference of foreground amplitudes and spectra via Gibbs sampling. These studies have never claimed a significant excess of emission toward the GC. Nevertheless, with the release of gamma-ray data from the *Fermi Gamma-Ray Space Telescope* in 2009, a corresponding feature at high energies was discovered (Dobler et al. 2010), likely generated by the haze electron inverse Compton scattering starlight, infrared, and CMB photons up to *Fermi* energies.

With few exceptions (see, for example, Bottino et al. 2010), foreground studies using template fitting pre-subtract a CMB estimate which imprints a bias in the foreground spectra (see Dobler & Finkbeiner 2008a) since no CMB estimate is completely clean of foregrounds. In this work, we attempt to minimize this effect by solving jointly for the CMB map and angular power spectrum and Galactic emission parameters, within a Bayesian framework. We fit foreground parameters in every pixel, thus taking into account spatial variation of foreground spectra, and we test the stability of the solution against several Galactic emission models. Seeking greater independence from correlation with external templates, we instead treat other data sets as input channel maps and process them through the Gibbs sampler: to this end we add Haslam at 408 MHz (Haslam et al. 1982) to the *WMAP* data set.

In Section 2 we describe the data and our foreground model, while in Section 3 we describe its application in our Commander analysis. Our approach enables us to improve upon previous studies. In Section 4, we discuss our results and refine the methodology by performing a joint analysis of the *WMAP* 7 year data set and Haslam at 408 MHz. Lastly, in Section 5 we summarize and draw our conclusions.

2. DATA SETS AND FOREGROUND MODEL

Our approach uses the Gibbs sampling algorithm introduced by Jewell et al. (2004) and Wandelt et al. (2004), and further developed by Eriksen et al. (2004, 2007b), O'Dwyer et al. (2004), Chu et al. (2005), Jewell et al. (2009), Rudjord et al. (2009), and Larson et al. (2007). The method has been numerically implemented in a computer program called Commander, and successfully applied to previous releases of the *WMAP* data as reported in Eriksen et al. (2007a, 2008a) and Dickinson et al. (2009). Commander is a maximum-likelihood method, which generates samples from the joint posterior density for the CMB map and angular power spectrum, as well as foreground components for the chosen sky model. A detailed description of the algorithm and its validation on simulated data is provided by Eriksen et al. (2007b and references therein). For the interested reader, we summarize the technical aspects of the sampling procedure in Appendix A.

The main advantage of the approach is full characterization of the (posterior) probability distribution of the parameter space spanned by the adopted sky model. Moreover, we can evaluate the goodness of fit of each sample, i.e., of a given set of model parameters. Any parametric model can be encoded, leaving the freedom to vary parameters pixel by pixel as well as to fit templates.

In the following, we describe the data used in our analysis and their processing.

2.1. WMAP 7 Year Data Set

Our aim is to perform a comprehensive study of the *WMAP* 7 year temperature data (Jarosik et al. 2011), focusing not only on estimating the CMB signal and corresponding power spectrum, but also on characterizing the properties of the foreground emission, with particular attention to evidence for the haze signal.

The *WMAP* data set comprises data from ten differential assemblies (DAs) covering five frequencies from 23 to 94 GHz. The *WMAP* team provided maps centered at frequencies of 23, 30, 40, 60, and 90 GHz [K, Ka, Q, V, W bands] resulting from the co-addition of the DAs. We smooth the frequency maps to an effective 60 arcmin resolution by deconvolving the maps with the appropriately co-added transfer functions provided by the *WMAP* team and convolving with a Gaussian beam of 1 deg FWHM. We then downgraded the sky maps to a working resolution of $N_{side} = 128$ in the HEALPix[13] scheme (Górski et al. 2005). The choice of angular resolution, and subsequently of the pixelization, was dictated by both the angular resolution of the available foreground templates and by the need to resolve the first acoustic peak of the CMB power spectrum. This represents a novel element of our analysis compared to previous work (Eriksen et al. 2008a; Dickinson et al. 2009).

In a further break with earlier analyses, we do not add uniform white noise to regularize the maps (Eriksen et al. 2007a), but instead add a noise component proportional to the actual noise variance in the smoothed, processed maps. To do so, we compute the rms noise per pixel of the maps at the lower resolution via a Monte Carlo approach. We generate 1000 non-uniform white-noise realizations for each channel following the prescription given by the *WMAP* team. In practice, we draw random Gaussian noise maps for each frequency band with zero mean and a variance given by σ_ν^2/N_{obs} at the native *WMAP* resolution $N_{side} = 512$, where N_{obs} represents the number of observations in a given pixel. We then smooth and downgrade the noise maps as above and finally recompute the variance per pixel averaged over 1000 simulations. Noise is then added to each frequency based on the computed noise variance, and chosen so that the signal-to-noise ratio of the masked smoothed map at $\ell \simeq 2N_{side}$ is of the order of unity. The scaling factors applied are [16,12,12,6,6] for the [K, Ka, Q, V, W] channels, respectively. One of the advantages of this approach is that it preserves the noise structure imposed by the scanning strategy and describes at least the diagonal part of the instrumental noise, which has been correlated by the smoothing procedure.

2.2. Foregrounds

The diffuse Galactic emission consists of three contributions from well-understood foregrounds—synchrotron radiation from cosmic ray electrons losing energy in the Galactic magnetic field, free–free emission in the diffuse ionized medium, and radiation from dust grains heated by the interstellar radiation field. In addition, there is a strong contribution at frequencies in the range 20–70 GHz referred to as AME that is strongly correlated with the thermal dust emission and which has been explained by rotational excitation of small grains—the so-called spinning dust emission (see, for example, Hoang et al. 2010,

---

[13] http://healpix.jpl.nasa.gov/

2011 for recent studies on the subject and references therein). In the frequency range spanned by *WMAP*, the synchrotron emission is reasonably well described by a power-law brightness temperature emissivity with spectral index $\beta \sim -3$. The free–free emission also follows an approximate power law, well described by a spectral index $\alpha \sim -2.15$. The thermal dust component is usually modeled as a gray body with emissivity of the form $T(\nu) \propto \nu^\epsilon B(\nu, T_d)$ with typical values for the parameters given by $\epsilon \simeq 1.6$–2.0 and $T_d \approx 18$ K. Since the highest frequency *WMAP* band *W* has a nominal central frequency of 94 GHz, the overall thermal dust contribution is small and can be approximated by a simpler power-law model with a spectral index $\epsilon \simeq 1.7$. The AME is not completely well characterized as yet, but falls rapidly below 20 GHz. Fits to the high-latitude sky suggest that it may be reasonably well approximated by a power-law emissivity over the *WMAP* range of wavelengths, although this may reflect the combination of multiple spinning dust populations in different physical conditions along a given line of sight.

The total foreground intensity observed by the ν channel in a given direction p can then be summarized as follows:

$$T_\nu(p) = M + \sum_{d=x,y,z} D_{\nu,d}(p) + \left(\frac{\nu}{\nu_0}\right)^{\beta(p)} A_{\text{synch}}(p)$$
$$+ \left(\frac{\nu}{\mu_0}\right)^{\alpha(p)} A_{\text{f–f}}(p) + \left(\frac{\nu}{\lambda_0}\right)^\epsilon A_{\text{dust}}(p)$$
$$+ \text{AME}(\nu), \qquad (1)$$

where M and D represent monopole and dipole residuals, respectively, A_i is the amplitude in antenna temperature units of the ith foreground component at the reference frequency $(\nu_0, \mu_0, \lambda_0)$, and β, α, and ϵ describe the spectral response of synchrotron, free–free, and thermal dust emission, and an AME contribution has also been included.

In principle, there is no reason why these parameters should be constant over the sky, so we should allow them to vary pixel by pixel and let `Commander` solve for the most likely value. In practice, we are limited by the number of frequency bands, five in the case of *WMAP*, from which we also wish to determine the CMB contribution. Therefore, there remain only four maps that can be used to infer the foreground emission. Addressing the full problem as posed in Equation (1) in an independent and self-consistent way is therefore not possible. Since our present analysis is motivated by a desire to assess the presence of a hard synchrotron component in the GC, the microwave haze, we prefer to disentangle the foreground contributions in the low-frequency range.

A plausible and widely used alternative, which allows for more realistic foreground modeling, is the use of external templates. At very low frequency (<1 GHz), the observed sky signal is dominated by the synchrotron emission from our Galaxy, and is little contaminated by free–free emission, at least away from the Galactic plane. A full-sky map of the synchrotron emission is provided by the 408 MHz map of Haslam et al. (1982). The optical Hα line is known to be a tracer of the free–free continuum emission at microwave wavelengths. Finkbeiner (2003) produced a full-sky Hα map as a composite of various surveys in both the northern and southern hemispheres. Finally, Finkbeiner et al. (1999) predicted the thermal dust contribution at microwave frequencies from a series of models based on the *COBE*-DIRBE 100 and 240 μm maps tied to *COBE*-FIRAS spectral data. We use predicted emission at 94 GHz from the preferred Model 8 (FDS8) as our reference
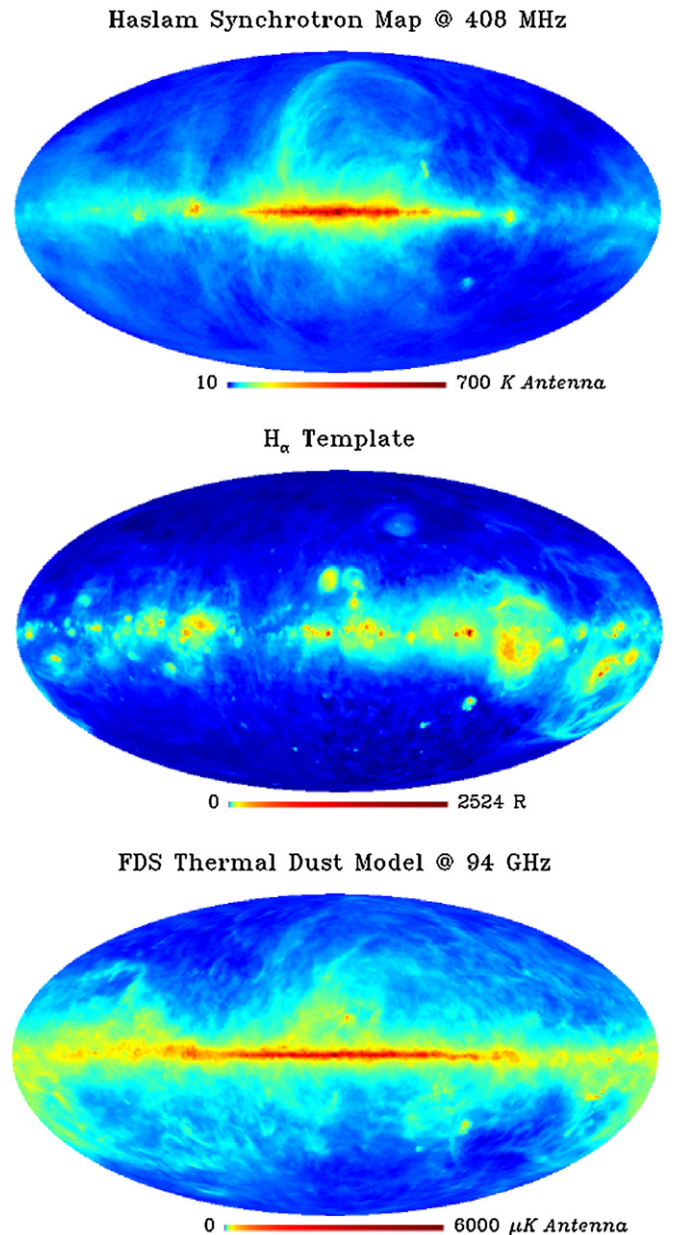

Figure 1. Top to bottom: the three templates used to trace synchrotron (Haslam 408 MHz, top panel), free–free (Hα, middle panel), and dust, both thermal and spinning components (FDS, lower panel).

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

template for dust emission. These templates, smoothed to 60 arcmin and downgraded to $N_{\text{side}} = 128$, are shown in Figure 1.

Since the thermal dust contribution to the *WMAP* bands is small compared to synchrotron and free–free emission, we can describe it by means of the FDS template with a fixed spectral index $\beta = 1.7$, allowing for an overall amplitude, b. The same can be done for the bremsstrahlung emission, assuming the Hα template as a sufficiently accurate description at 23 GHz and rescaling the amplitude according to a power law with index $\alpha = -2.15$. The synchrotron emission is expected to vary across the sky, and the template we have is at 408 MHz, quite a large stretch from the first band of *WMAP*. A wise choice is to let `Commander` solve for an amplitude and spectral index at every pixel, choosing 23 GHz as pivot. It should be noted that variations in the gas temperature will be imprinted onto the

synchrotron component, as will deviations in the spectrum of thermal dust (which is a much smaller contribution).

A detailed study of the foreground emission in the frequency range 23–94 GHz employing external templates can be found in Dickinson et al. (2009). Note that the value of the spectral index of the dust, 1.7, is somewhat dependent on the specific method used to derive it. A direct fit to the predicted dust templates at *WMAP* frequencies using the FDS8 model yields a lower value of $\beta = 1.55$, which is in agreement with what we find when applying template fitting procedure to *WMAP* maps (see Figure 8).

3. *WMAP* ANALYSIS

Previous work on the Galactic haze relied heavily on template regression techniques, raising a debate as to whether the assumption of constant spectral behavior across the sky is the cause of the excess emission near the GC. To address this issue, we follow an alternate approach and solve instead for a simple model describing Galactic emission with two power laws. One is a low-frequency component with a falling spectrum to account for synchrotron and free–free emission, as well as AME. The second is a higher frequency component with a rising spectrum to represent thermal dust emission. This appears to be well motivated in studies of the *WMAP* Maximum Entropy Method (MEM) foreground solutions, as in Park et al. (2007). While we can solve for the amplitude and spectral index at every pixel at of the low-frequency component, where the variability of the signal is large, we fix the dust emissivity at high frequency to $\epsilon = 1.7$, consistent with previous applications of the Gibbs sampling technique to the *WMAP* data. We choose 22.8 GHz and 94 GHz as pivot frequency for the low-frequency and thermal dust components, respectively. The drawback of this approach is that the three low-frequency emission mechanisms are combined into a single power-law component that may not provide a fully adequate description of their physical complexity. However, we will attempt to separate these components post-sampling in Section 3.1.

We stress that our approach differs from that followed by Cumberbatch et al. (2009), where Commander was only used to extract a CMB map while modeling the foreground components with templates. The main focus of that work was comparison between the *WMAP* Internal Linear Combination (ILC) CMB map, previously adopted in studies of the haze, and the posterior mean CMB map derived from Gibbs sampling. The presence of the haze was found to be stable with respect to the CMB map choice. In the present paper, we take full advantage of Commander to jointly sample from foregrounds and CMB. In addition, we choose a less aggressive smoothing of the *WMAP* maps and mask, both factors that are likely to enhance the signal-to-noise ratio of the haze component.

The Commander result for the CMB map is shown in Figure 2. As a diagnostic, we show the mean χ^2 map in the lower panel of Figure 2. The number of pixels with $\chi^2 > 20$, corresponding to the 99.9% confidence level for 5 degrees of freedom, is 0.7%. Most of these pixels lie near point sources and we conclude that they correspond to point source bleeding outside the mask after convolution with a large beam. The overall lack of features in the χ^2 map is an indication of the goodness of fit of the model and that the residuals at each frequency are compatible with our noise description.

We checked that the Commander power spectrum is consistent with the best estimate provided by the *WMAP* team (Larson et al. 2011) at the 1σ level up to $\ell = 200$. Beyond that, the Commander
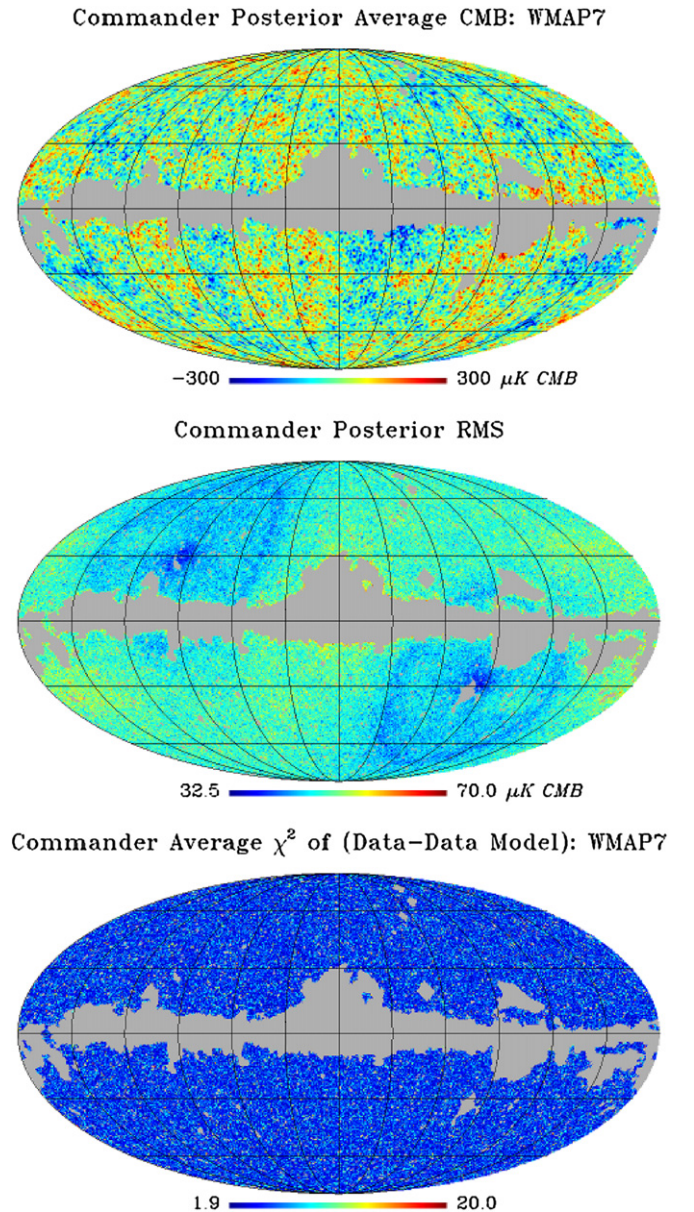

Figure 2. CMB Commander posterior average, rms, and mean χ^2 map. No particular features are present, meaning that the model is a very good fit to the *WMAP* 7 year data. The lower limit results 1.9, whereas the upper limit, 20, corresponds to 0.001 probability for a χ^2 distribution with 5 degrees of freedom. Within the Gibbs sampling framework, we do not search for the maximum-likelihood (ML) point, but rather sample from the posterior: in practice, we add a scatter term around the ML point, drawn from the noise model of the data. This results into a number that is distributed with 5 degrees of freedom.

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

solution has significantly more scatter due to the regularizing noise added to the smoothed, binned maps (see Section 2.1). In addition, while the *WMAP* team also used Gibbs sampling techniques for low multiples, their high ℓ power spectrum was obtained by applying a quadratic estimator to the cross-spectra of *V* and *W* bands. We emphasize that the agreement between the two power spectra below $\ell = 200$ is what is important for our purposes since we are concentrating on features in the map that are much larger than $1°$. The possible presence of a residual monopole and dipole in the ΔT data has been taken into account and the mean values obtained are displayed in Table 1. They are small and compatible with those found by Dickinson et al. (2009) in the *WMAP* 5 year data, though it is
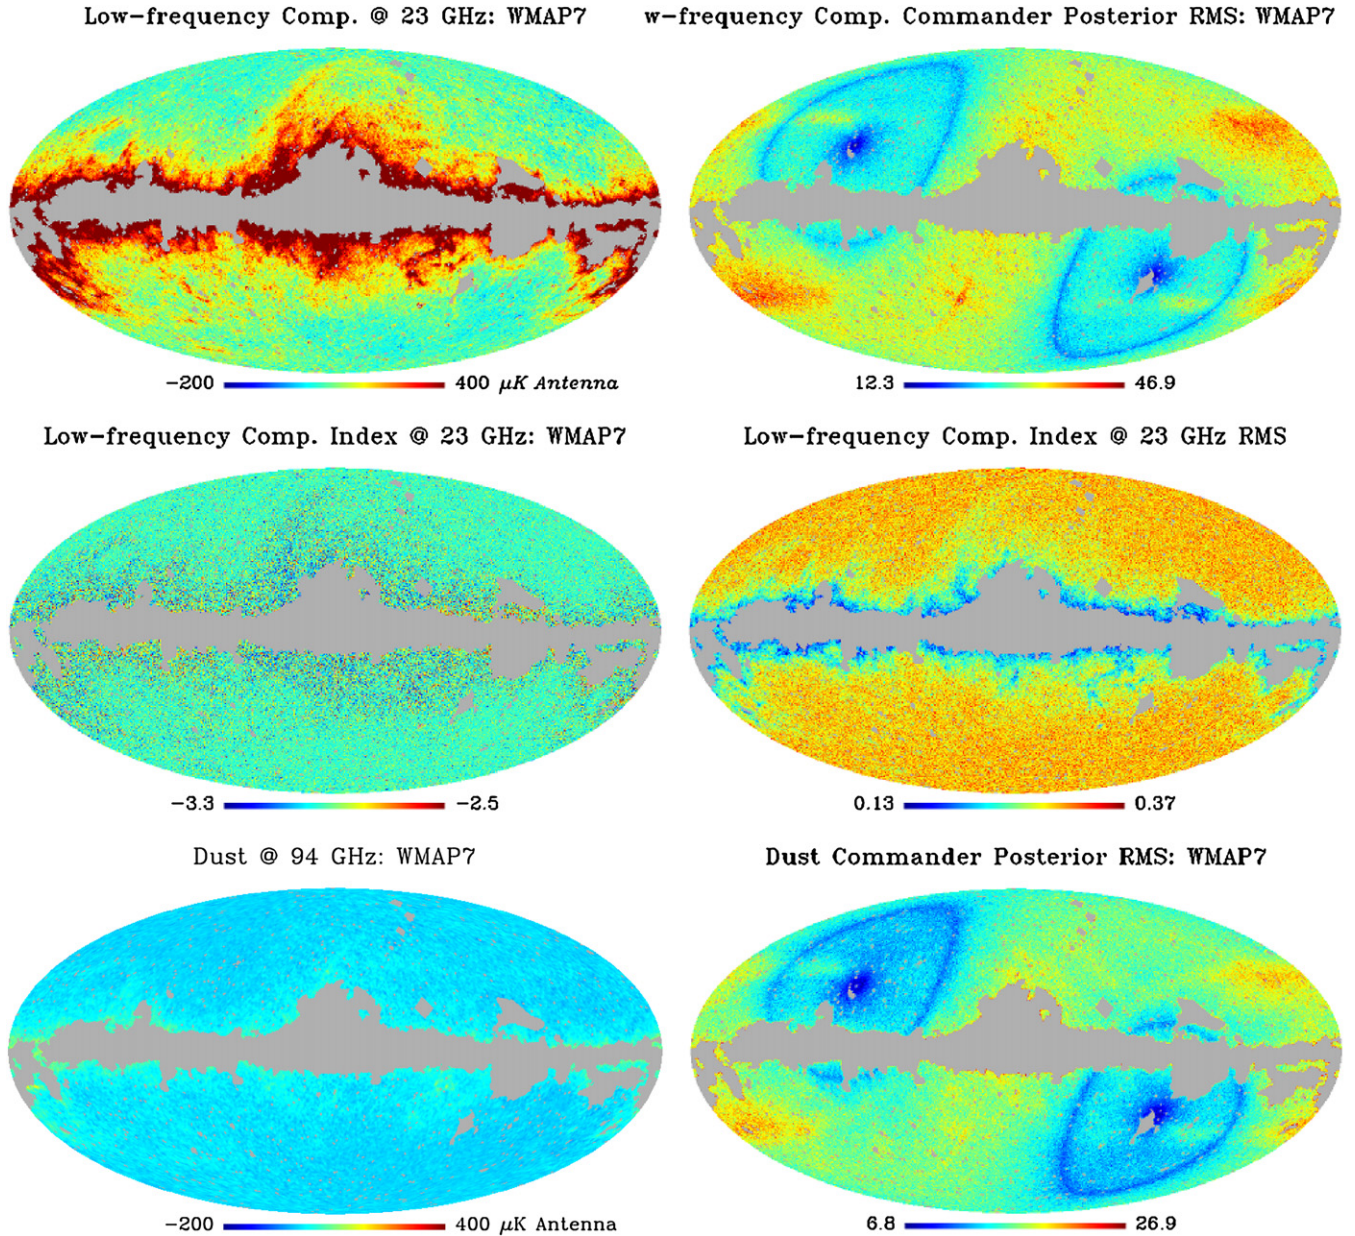
Figure 3. Mean field map (left column) and rms (right column) for low-frequency component amplitude (upper panel) and spectral index (middle panel) and dust amplitude (lower panel) for the foreground model with two power laws.

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

Table 1
Mean Values (μK) for the Monopole and Dipole Residuals at Every Frequency

	K band (23 GHz)	Ka band (33 GHz)	Q band (41 GHz)	V band (61 GHz)	W band (94 GHz)
M	2.4 ± 0.9	5.6 ± 1.0	3.5 ± 1.1	2.7 ± 1.0	3.9 ± 1.0
D_x	-2.6 ± 1.2	-3.4 ± 1.2	-2.7 ± 1.2	-2.7 ± 1.2	-3.0 ± 1.2
D_y	-3.6 ± 0.9	-4.3 ± 0.9	-4.4 ± 0.9	-3.5 ± 0.9	-4.0 ± 0.9
D_z	1.9 ± 0.2	1.9 ± 0.2	2.4 ± 0.2	1.8 ± 0.2	2.0 ± 0.2

important to keep in mind that monopole and dipole features become strongly coupled to foregrounds in CMB analyses, as discussed by Eriksen et al. (2008b).

Foreground maps (left column) and associated errors (right column) for low frequency amplitude and spectral index measured at 23 GHz and the thermal dust contribution at 94 GHz are shown in Figure 3.

It is clear that our derived low-frequency component represents a combination of multiple emission mechanisms, which we now attempt to disentangle. The low-frequency component shows a strong correlation with thermal dust emission at 94 GHz as modeled by FDS, a correlation that is interpreted as signature of a spinning dust contribution. The spectral index map, second row in Figure 3, is particularly informative in this respect since departures from the prior assumed in Commander, -3.0 ± 0.3 are driven by the data. The mean value at high Galactic latitudes turns out to be a bit higher than the prior ($\simeq -2.9$) as consequence of the superposition along the line of sight of multiple components. It is remarkable how very bright free–free regions show up clearly, requiring a spectral index close to -2.15, which saturates the scale in Figure 4. The visual correlation of these diffuse gas clouds with those present in the Hα map is striking. Spectral index values lower than the prior are visible in regions where both synchrotron and dust
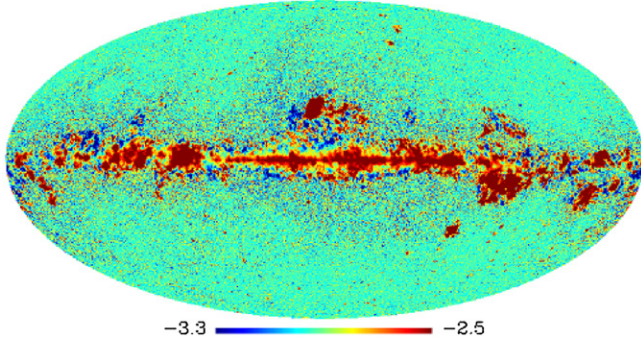
Figure 4. Full-sky spectral index mean field map. The Galactic plane shows strong variation corresponding to regions where one component among synchrotron, free–free, and spinning dust emission dominates.

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

Table 2
Regression Coefficients of the `Commander` Foreground Amplitude Maps

Data Set	WMAP 7 yr			
	Haslam	Hα	FDS	r
Low-frequency component	$(3.6 \pm 1.2) \times 10^{-6}$	6 ± 6	7 ± 2	0.94
Thermal dust	$(0.02 \pm 1.0) \times 10^{-6}$	-0.6 ± 5	1.1 ± 1.8	0.58

Notes. While the thermal dust correlates with the FDS model only, the low-frequency component is indeed a mixture of all three templates, synchrotron and dust emission being the strongest. Moreover, while the thermal dust residuals are featureless and consistent with noise, the low-frequency component residuals show a clear excess of power around the Galactic center.

correlated emission are present, and distinguishing between the two is difficult.

3.1. Template Regression

Following a common approach to the problem (see, for example, Dobler & Finkbeiner 2008b; Bottino et al. 2010 and references therein), we regress our amplitude maps against the three templates shown in Figure 1.
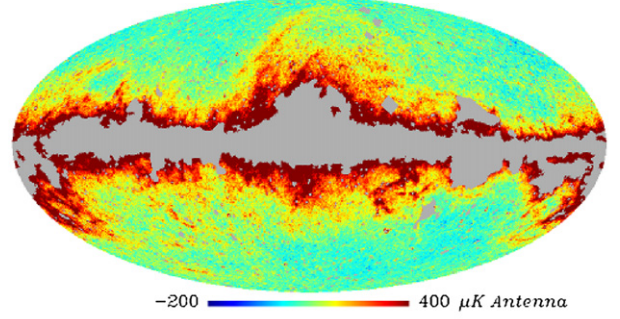
Assuming the amplitude map to be a linear combination of these templates, the coefficients A_i^ν are computed from a χ^2 minimization of the form

$$\chi^2 = \sum_\nu \sum_p \left(S_\nu(p) - \sum_i A_i^\nu T_i(p) \right)$$
$$N_\nu^{-1}(p) \left(S_\nu(p) - \sum_i A_i^\nu T_i(p) \right);$$
$$A_i^\nu : \frac{\partial \chi^2}{\partial A_i^\nu} = 0;$$
$$\Delta^2 A_i^\nu = \left(\frac{\partial^2 \chi^2}{\partial \left[A_i^\nu \right]^2} \right)^{-1}, \quad (2)$$
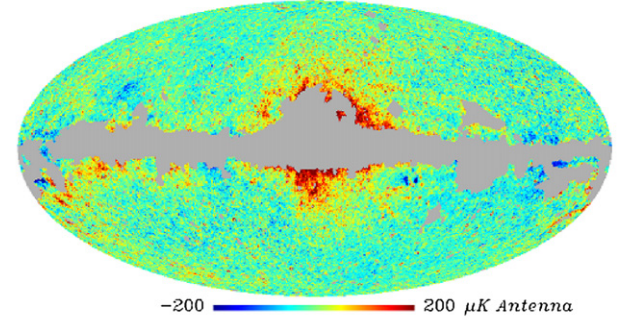
where T_i are the foreground templates and S_ν is the `Commander` amplitude solution. The explicit expression for the regression solution is given by

$$A_i^\nu = \sum_j \mathbf{T}^{\nu\,-1}_{ij} B_j^\nu,$$
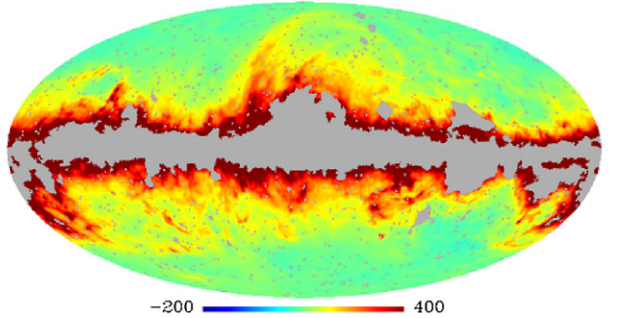$$T_{ij}^\nu = \sum_p \frac{T_i(p) T_j(p)}{N_\nu^2(p)},$$

Figure 5. Low-frequency foreground amplitude maps (top) compared to the linear combination of templates (bottom). The difference (i.e., residuals) is shown in the middle panel.

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

$$B_j^\nu = \sum_p \frac{S_i^\nu(p) T_j(p)}{N_\mu^2(p)}, \quad (3)$$

where T_{ij}^ν and B_j^ν describe the noise-weighted correlation between foreground templates and frequency maps and templates, respectively. The error ΔA_i^μ is given by $\sqrt{(2\mathbf{T}^\mu)_{ii}^{-1}}$. See also Fernández-Cerezo et al. (2006) and Hildebrandt et al. (2007). Since templates are not reliable in the Galactic plane, the fit is performed outside the Kq85 mask (Larson et al. 2011), which removes the plane and detected point sources. In addition, the mask covers regions of high dust column density where extinction makes our Hα template a poor tracer of free–free template emission.

The coefficients of the fit are quoted in Table 2 and the resulting template and residuals in Figures 5 and 6. Contrary to the usual technique, where the regression is performed on the *WMAP* maps, we regress out instead the `Commander` posterior average maps, which have noise contributed by the sampling procedure. Note that the error we quote is obtained by integrating

Dust @ 94 GHz: WMAP7

−100.0 ▬▬▬▬▬ 100.0 *μK Antenna*

Residual Dust @ 94 GHz: WMAP7

−100.0 ▬▬▬▬▬ 100.0 *μK Antenna*

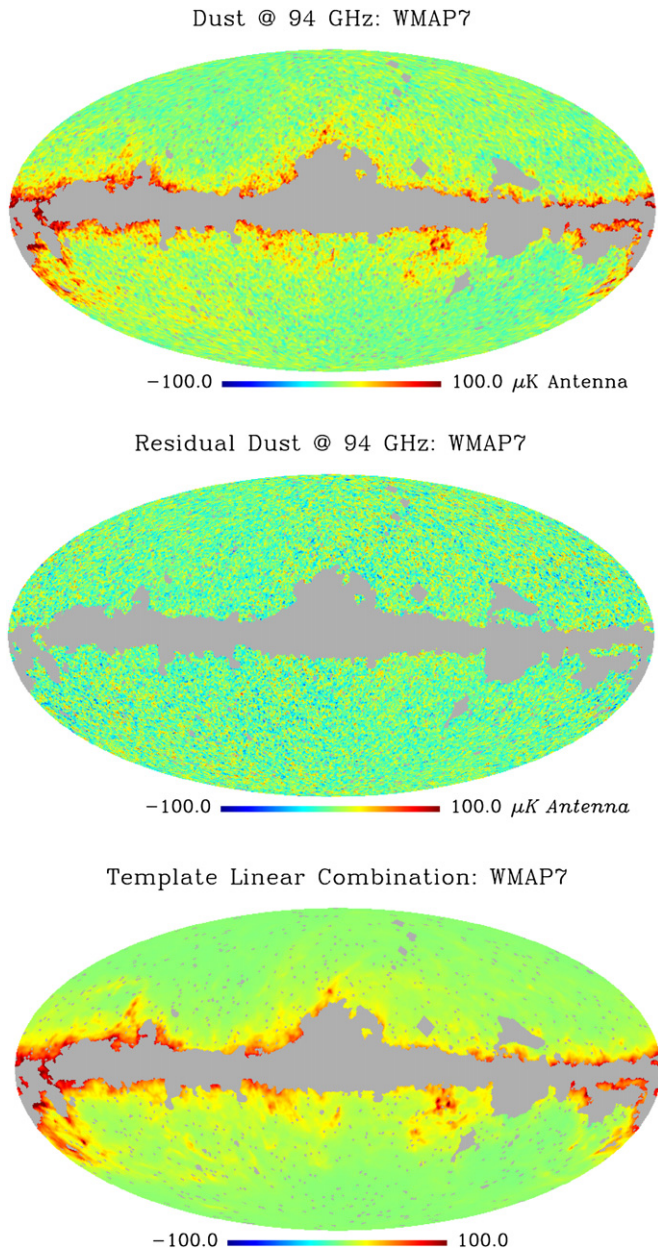Template Linear Combination: WMAP7

−100.0 ▬▬▬▬▬ 100.0

Figure 6. Thermal dust amplitude maps (top) compared to the linear combination of templates (bottom); the difference (i.e., residuals) is shown in the middle panel.

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

Low-freq. Comp. @ 23 GHz: WMAP7

Dust @ 94 GHz: WMAP7

Figure 7. For the low-frequency foreground (top) and dust amplitude (bottom), we show the pixel-to-pixel comparison between the best-fitting template combination (y-axis) and the `Commander` values (x-axis). The red line marks the $y = x$ scaling. Blue and green dots denote residuals with a significance larger than 3σ, positive and negative, respectively. Orange points are pixels corresponding to the haze: 3σ positive excess within a disk of 36 degrees around the Galactic center.

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

masked out. To further quantify the goodness of fit, we show scatter plots of the mean field `Commander` solution as function of the derived linear template shown in Figure 7.

Although the correlation is clearly present and follows the $y = x$ line, looking at the residual map is very instructive. In the case of dust, the residuals are compatible with noise and this explains why the points cluster close to $y = 0$ line. The situation is more intriguing for the low-frequency component where the fit is not perfect and leaves an excess of power in the proximity of the GC—the Galactic haze—which has been advocated to be a distinct contribution. Fainter positive and negative regions exist as well, perhaps suggesting an overly simplistic modeling of the dust component.

4. FURTHER ANALYSIS

The most troublesome foreground is the dust, both thermal and spinning. While the former remains weak at *WMAP* frequencies outside the Galactic mask, 94 GHz being only slightly contaminated, the latter is poorly characterized and traced through the correlation between the lowest *WMAP* channels and the FDS dust model, under the assumption that thermal dust traces spinning dust reasonably well. We now include this correlation in our foreground model by first regressing *WMAP* channels against Haslam, Hα, and FDS to obtain a spectral energy distribution (SED) for each foreground, and then use the resulting phenomenological dust SED with `Commander` to solve for the dust amplitude and a low-frequency component, described by a power law, together with the CMB. We recognize that this spinning dust model is simple and unlikely to capture the full

the posterior distribution over the other parameters, indices and component maps, and hence accounts for the uncertainty associated with the foreground fit as well as the instrumental noise. We also remind the reader that the input maps analyzed by `Commander` are noisier than the expected from simply smoothing them because we add regularizing noise required by the sampling algorithm (see Section 2.1). These two factors increase the uncertainty on the template amplitudes. A detailed investigation is presented in Appendix B.

As expected, the high-frequency foreground component correlates with the FDS model of thermal dust only. The low-frequency component is mainly a combination of synchrotron emission and correlated dust, the latter possibly due to spinning dust. We note that the contribution of free–free emission is weak because the Galactic plane, where it is strongest, has been
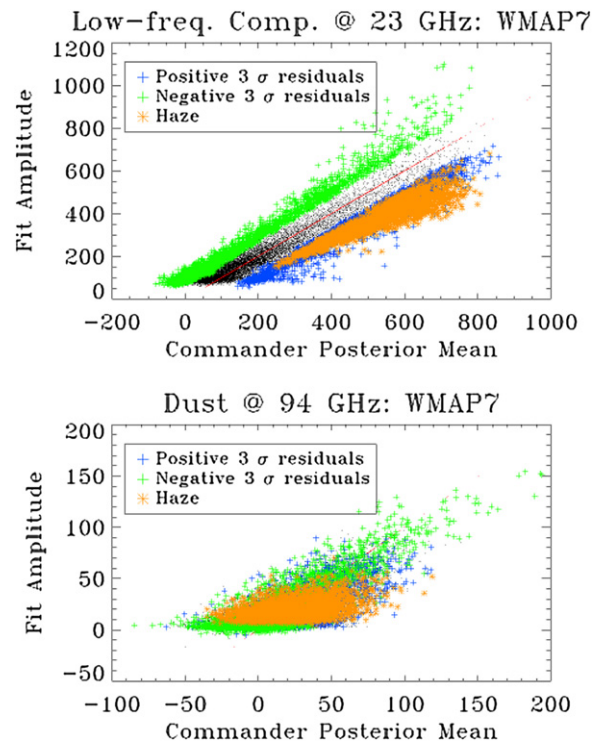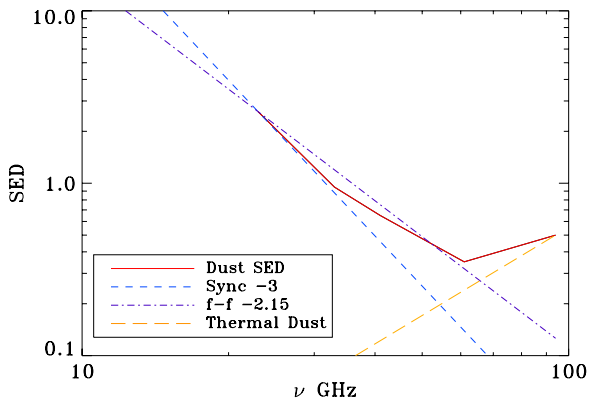
Figure 8. SED for the phenomenological dust, compared to those of other components of interest: synchrotron, free–free, and thermal dust emission.

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

complexity of the dust emission. However, more complicated models, like spatially varying spinning dust index or multiple component spinning dust, although physically motivated (such as those proposed by Hoang et al. 2010, 2011), would require larger frequency coverage for practical use. We therefore adopt our simple model as an adequate description for the available *WMAP* data set.

4.1. Low-frequency Component and Phenomenological Dust

The first `Commander` run with this model returned a very steep spectral index for the low-frequency component, close to the limit of our prior for some regions in the Galactic plane known to have strong spinning dust emission. The χ^2 of our solution in those regions was high, a sign that the foreground model fails. We then tuned the amplitude of our dust model, increasing it at low frequencies until we obtained a shallower spectral index for the low-frequency component and smaller χ^2 values for the same regions. In practice, this means that we modify the SED template to increase the relative amount of spinning dust to thermal dust. The resulting spectral response is shown in Figure 8, together with those of the soft synchrotron ($\beta = -3$), free–free ($\beta = -2.15$), and thermal dust components ($\beta = 1.7$).
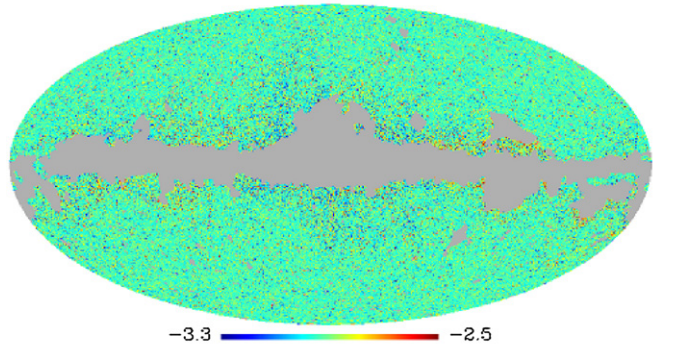
Figure 9 shows the mean field amplitude maps of the thermal/spinning dust, and of the low-frequency component along with its spectral index. At high Galactic latitudes, where the signal to noise is very low, the spectral index is consistent with the Gaussian prior, $\beta = -3.0 \pm 0.3$, -2.9 being the mean value, whereas closer to the plane the solution is driven by the data. In high-latitude regions of strong synchrotron features, such as Loop I, the spectral index is noticeably softer than close to the plane, where spinning dust and free–free emission become more important. Finally, it is interesting to note that the dust map recovered at 94 GHz is remarkably similar to the FDS prediction.

As noted above, the low-frequency `Commander` solution represents the combination of several different emission mechanisms known to coexist at 23–41 GHz: synchrotron, free–free, and possible spinning dust residuals. This is confirmed by the low-frequency spectral index map that presents shallower values at high Galactic latitude, but clearly emphasizes regions in the Galactic plane with a very steep spectral index, likely to be spinning dust clouds. A comparison between the low-frequency component spectral index and amplitude for the two models is shown in Figure 10. The evidence for dust-correlated emission
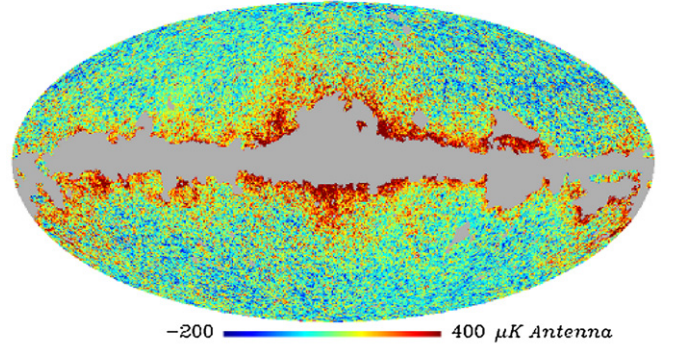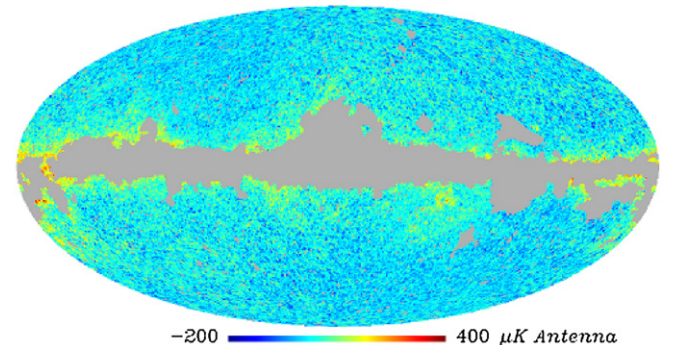

Figure 9. Low-frequency component mean field amplitude (upper panel) and spectral index (middle panel) map at 23 GHz and dust-correlated emission (lower panel).

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

is compelling, implying that our phenomenological dust SED is a good approximation of the sky.

For the most part, all of these emission mechanisms decrease in intensity above 23 GHz (but see Dobler & Finkbeiner 2008b; Dobler et al. 2009), and it is because of the noise in the data that a single power-law results in a good χ^2. In an effort to disentangle these primary sources of emission, we again apply template regression to the `Commander` outputs using the Haslam 408 MHz map, the Hα map, and the FDS map as tracers of synchrotron, free–free, and dust (thermal and spinning) emission, respectively.

The regression coefficients are quoted in Table 3 and indeed they confirm our success: outside the applied mask, dust is positively correlated with the FDS map only, whereas the low-frequency component map can be described as a linear

Table 3
Regression Coefficients for `Commander` Foreground Amplitudes

Data Set	Haslam	Hα	FDS
WMAP 7 yr			
Low-frequency component	$(3.4 \pm 3.3) \times 10^{-6}$	9 ± 15	0.6 ± 6
Dust	$(0.1 \pm 1.3) \times 10^{-6}$	-1 ± 6	2.3 ± 2.3
Haslam+*WMAP* 7 yr			
Soft synchrotron	$(5.9 \pm 0.4) \times 10^{-6}$	-0.4 ± 1.3	$(-0.6 \pm 7) \times 10^{-1}$
Low-frequency component	$(-1.7 \pm 3.5) \times 10^{-6}$	8 ± 15	0.4 ± 6
Dust	$(-0.3 \pm 1.2) \times 10^{-6}$	-0.25 ± 6	2.4 ± 2.2

Notes. Top: *WMAP* 7 yr Run; Bottom: *WMAP* 7 yr and Haslam. The Galactic emission is decomposed into physical components, each of which shows a clear correlation with one template only.
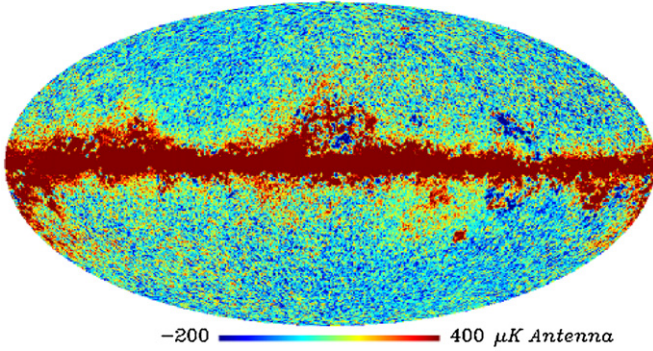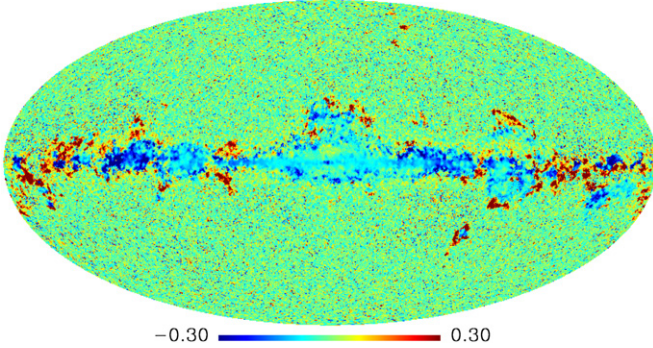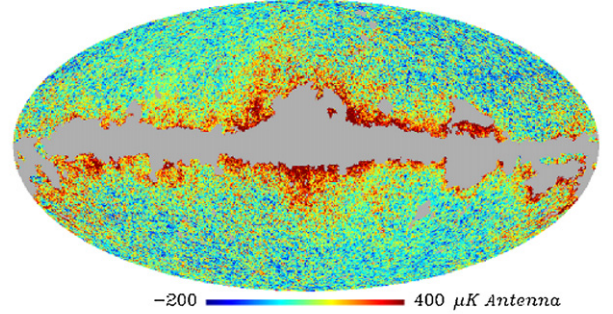

Figure 10. Difference between two `Commander` posterior mean low-frequency component amplitudes (upper panel) and spectral indices (lower panel) assuming two different foreground models: two power laws, and a power law with the phenomenological dust SED.

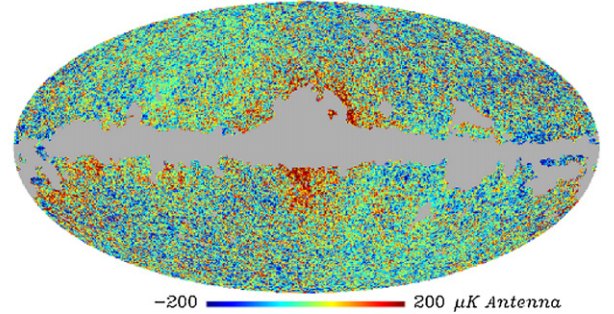(A color version of this figure is available in the online journal.)
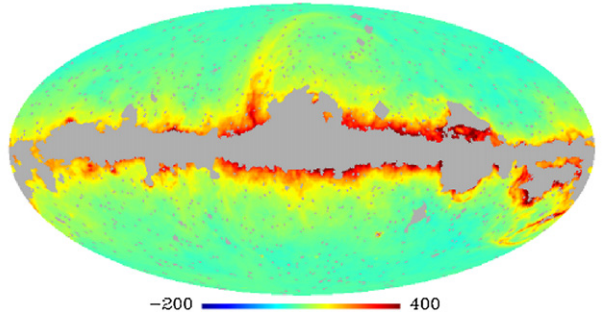
Figure 11. *WMAP* 7 yr comparison between low-frequency foreground component amplitude solution and linear fit of the templates: an excess of power around the Galactic center is present.

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

combination of Haslam and Hα. The goodness of the regression is expressed by the coefficient.

As in the former run with thermal dust described by a power law, the FDS template shows itself a remarkable tracer of the dust map recovered by `Commander`, with the residual showing no evidence for large-scale deviations from the template. The same is true for the low-frequency component with the *significant* exception of the "haze region" around the GC. Figure 11 shows the residuals after removing the template-correlated emission from the `Commander` mean field low-frequency foreground map. In other words, synchrotron emission elsewhere (in Loop I and in the Galactic plane) as well as free–free emission are traced well by our templates while the haze is definitely not, indicating that this large structure is not morphologically correlated with

any of the templates. We find that the haze emission is present in the data and contributes significantly to the total emission toward the GC (and particularly in the south).

Our findings are in agreement with what has been described by several authors and in particular by Dobler & Finkbeiner (2008a), who claim the presence of an additional foreground component characterized by a spectral behavior harder than a synchrotron component, but compatible with neither free–free (because the spectrum is too soft) nor spinning dust emission (because of lack of a thermal feature). Previous analyses were based on pre-subtraction of CMB cleaned maps from the data, which is problematic due to the fact that no CMB estimator is completely clean of foregrounds. This leads to a bias in the inferred spectrum of the Galactic emissions. To the extent that it is possible with five bands, we have attempted to reduce this systematic by simultaneously solving for the cleaned CMB map and the foreground maps while also determining a spectral index on the sky, within a Bayesian framework, where the goodness of the fit is controlled pixel by pixel through a χ^2 evaluation.

4.2. WMAP 7 Year Data Combined with Ancillary Data

Our results in the previous section indicate that the haze represents either a new component not present in the Haslam map or a variation of the spectral index associated with Haslam 408 MHz data. In order to assess this, we include the 408 MHz Haslam map together with *WMAP* data and run `Commander` on *six* bands covering a factor of ∼250 in frequency. This allows us to fit for an additional component and separate soft synchrotron from other emission mechanisms, the former being described by a power law with fixed spectral index $\beta = -3$, the latter by a spatially varying spectral index. We expect the soft contribution to be mainly driven by the Haslam map. Our foreground model in this second `Commander` run becomes CMB, dust described by a spatially constant spectrum which accounts for both spinning and thermal emission, soft synchrotron with fixed spectral response, and an additional low-frequency component with spatially varying spectral index

$$T_\nu(p) = M_\nu + \sum_{d=x,y,z} D_{\nu,d}(p) + \left(\frac{\nu}{\nu_0}\right)^\beta (p) A_{\text{lowfreq}}(p)$$
$$+ \left(\frac{\nu}{\lambda_0}\right)^{-3} A_{\text{soft synch}}(p) + s(\nu) A_{\text{dust}}(p), \qquad (4)$$

where $\nu_0 = 22.8$ GHz and $\lambda_0 = 408$ MHz, and $s(\nu)$ is the fixed effective dust spectrum, normalized to $\mu_0 = 33$ GHz.

The mean posterior CMB we obtain with this `Commander` model is slightly different compared to the five-band case showing less power close to the Galactic plane. This suggests that the improved flexibility of our model has resulted in better component separation. The overall χ^2 of these six-band run does not change, since we add one map and we fit for one more foreground component. It is interesting to notice that we retrieve a value of the monopole in the Haslam map of ≃3.2 K.

In Figure 12 we report the mean field map of the foreground amplitudes: by visual inspection, the three amplitude maps look strongly correlated with the foreground templates. The coefficients of the regression are summarized in Table 3.

The residuals obtained after regressing out the templates from the `Commander` solution obtained for this more flexible foreground model are shown in Figure 13. The residuals are quite striking and demonstrate that the soft synchrotron component (i.e., emission with a spectral index of $\nu^{-3.0}$) is almost perfectly correlated with Haslam and is not present, while the low-frequency component clearly has a $\nu^{-2.15}$ free–free component that is strongly correlated with Hα and the haze
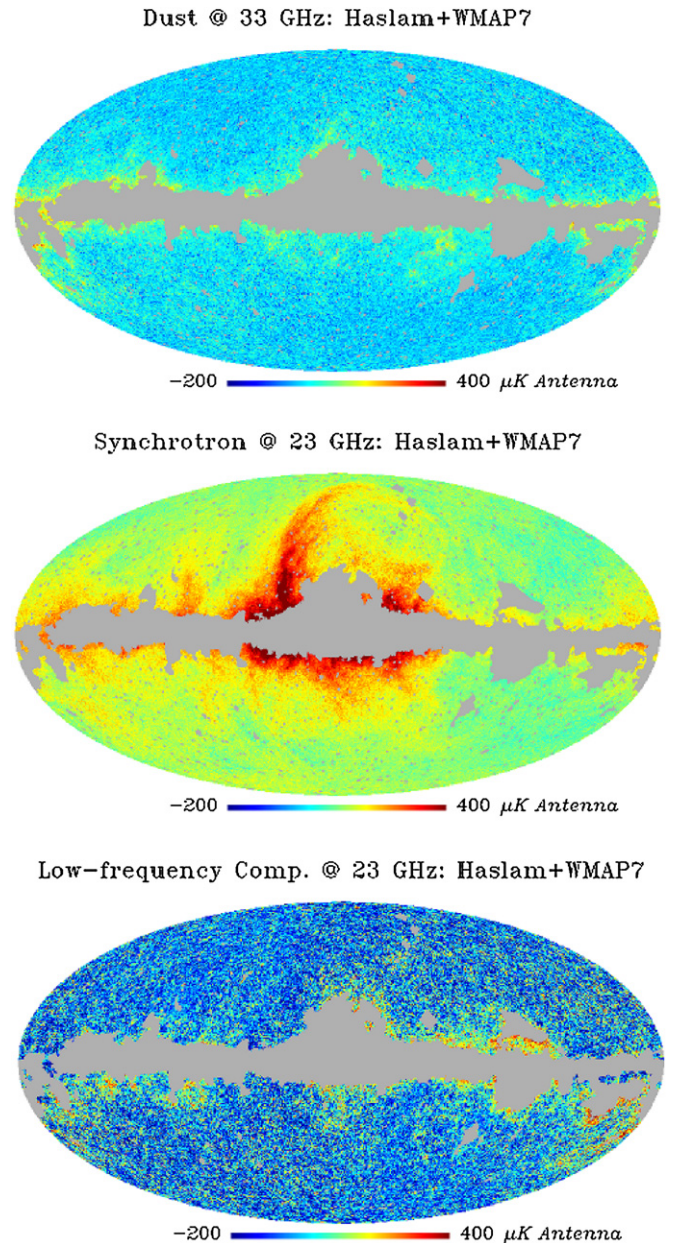

Figure 12. *WMAP* 7 yr+Haslam. `Commander` posterior mean foreground amplitude maps.

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

component, which is softer. This may suggest that the haze is not merely a spectral index variation of emission present in the Haslam map. In fact, the low-frequency component is anti-correlated with Haslam (driven primarily by the small negative residuals of Loop I). It is plausible that the prior we set for the spectral index of the soft component, $\beta = -3$, is not steep enough to characterize the Spur ($\beta = -3.05$), although at high Galactic latitudes a flatter behavior is measured ($\beta = -2.9$). Driven by this consideration, we ran `Commander` setting the spectral index of the synchrotron component to $\beta = -2.9$ and $\beta = -3.05$: the former resulted in more negative residuals in the synchrotron component, the latter reduced without completely removing them.

The haze is a relatively weak residual compared to the soft synchrotron component, thus a spectral index change is required to generate it (Mertsch & Sarkar 2010). If this were interpreted

Low-frequency Comp. @ 23 GHz: Haslam+WMAP7
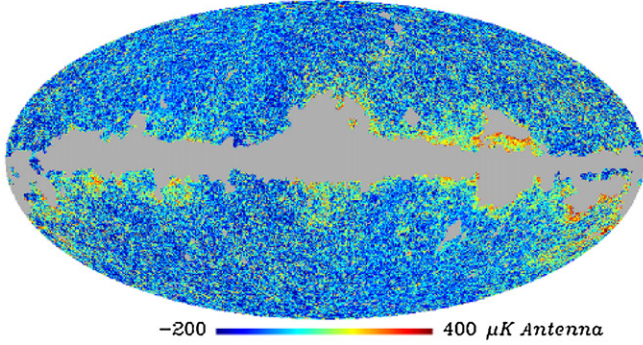
Residual Low-freq. Comp. @ 23 GHz: Haslam+WMAP7
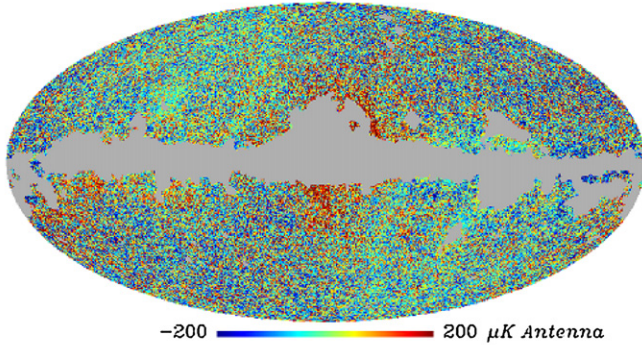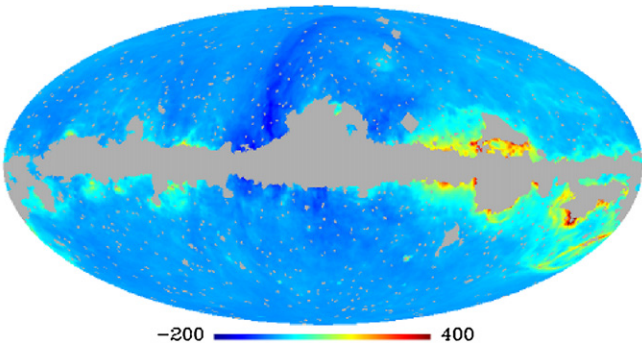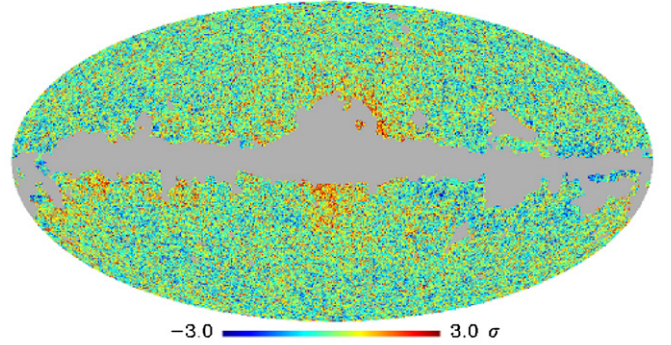

Template Linear Combination: Haslam+WMAP7

Figure 13. *WMAP* 7 yr+Haslam. Comparison between low-frequency component amplitude solution and linear fit of the templates. The Haslam-uncorrelated low-frequency component shows an excess of power which is not present in any of the templates.

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

al Foregrounds Low-freq. Comp. @ 23 GHz: Haslam+
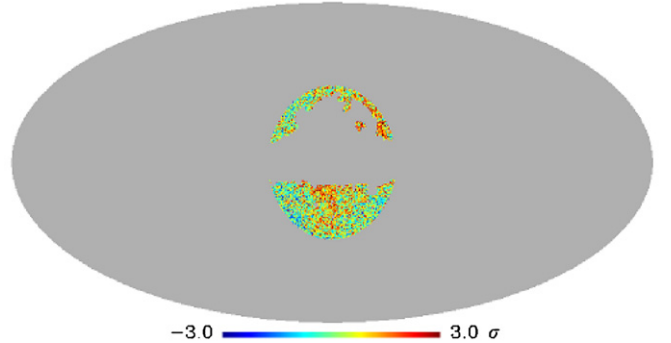

Galactic Haze Region at 23 GHz

Figure 14. Significance of the residuals for the soft synchrotron component (top panel) and pixel selected for the study of the correlation on a small region around the Galactic center.

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

as a simple variation of the spectral index of Haslam rather than a separate component not visible in Haslam, then this would imply that the spectral index is harder from 23–60 GHz than from 0.408–23 GHz; i.e., the spectrum would be concave up. However, cooling of electrons always generates spectra that are concave down because energy losses of cosmic-ray electrons are proportional to the electron energy squared. Thus, power-law injection of electrons results in concave down spectra as they cool. In other words, a different population of electrons with a harder spectral index must be present; a simple variation of the spectral index of the same electron population is not what we measured, since the amplitude of the soft component that we retrieved is strongly correlated with the Haslam map at 408 MHz. It is true that the assumption of a power law for the

synchrotron component can be too simplistic, and a (negative) curvature might be present, as a result of a break in the spectral behavior. However, given the number of frequencies available, sparse between 0.408 and 23 GHz and relatively narrow range of the *WMAP* data, allowing an extra parameter to be determined is not possible.

Again, the dust component is only strongly correlated with the FDS map, even more so than in the previous model, again likely indicating a better foreground separation.

Lastly, we point out that the spectral index map of the haze clearly shows that its spectral behavior is not compatible with free–free, but it is driven by the prior ($\beta = -2.5 \pm 0.3$). Unfortunately, this may suggest that the regularizing noise we added is too high to enable us to be definitive about the spectral index of the haze.

Up until now we have explored template fits in which the spectra are taken to be uniform across the sky. This is obviously not the case, and so to demonstrate the effects of this assumption we break the sky into regions and perform individual fits on these. Examples of this technique can be found in Hildebrandt et al. (2007), Dobler & Finkbeiner (2008a), Dobler (2012), and Ghosh et al. (2012). In practice, it means computing the correlation coefficients given in Equation (3) on a subset of pixels for the chosen regions. Here we will not perform such a detailed analysis, but restrict ourselves only to the haze region around the GC. The idea behind this approach is that if the haze corresponds to a region of global variation of the spectral index of the synchrotron component, the correlation between our residual map and the Haslam map computed on the haze pixels would be high. We can select the brightest pixels based on the
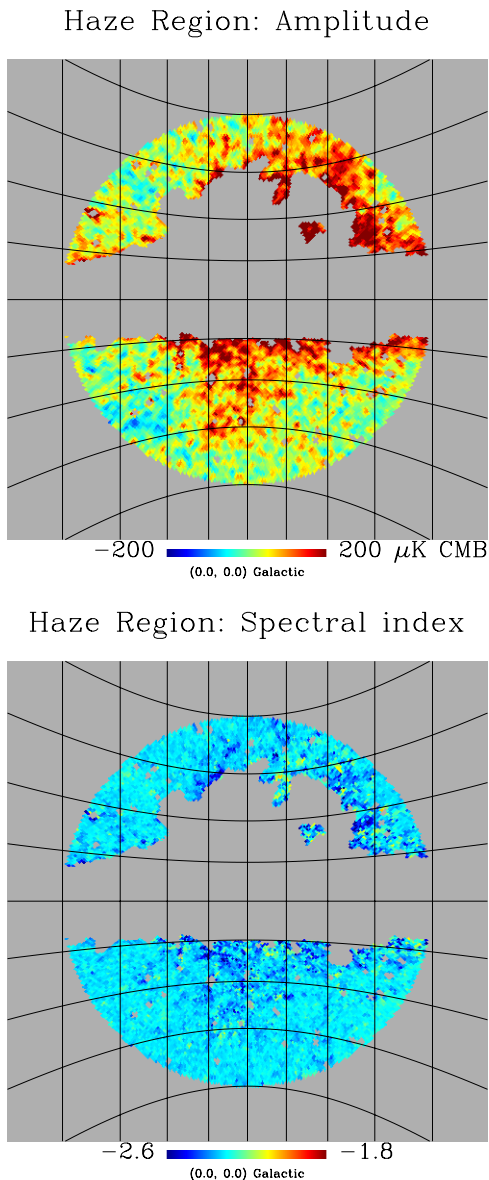
Haze Region: Amplitude

$$-200 \quad\blacksquare\blacksquare\quad 200 \; \mu\mathrm{K} \; \mathrm{CMB}$$
(0.0, 0.0) Galactic

Haze Region: Spectral index

$$-2.6 \quad\blacksquare\blacksquare\quad -1.8$$
(0.0, 0.0) Galactic

Figure 15. Haze amplitude and spectral index obtained after regressing `Commander` posterior mean amplitudes found without adding regularization noise.

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

rms map we derived from `Commander`: in Figure 14 we show the significance of the residuals and the pixels above 3σ which we chose within 36 deg of the GC. The correlation coefficients we obtained for the haze are $(3.3 \pm 3.5) \times 10^{-6}$, 0.1 ± 21, and 3 ± 7, respectively for Haslam, Hα, and FDS templates, showing no evidence for strong correlation. It is interesting to notice that while the amplitudes for free–free and dust emission remain stable with the threshold, this is not the case of the Haslam map, whose correlation decreases with the lower bound of the chosen significance cut. We tried 2.0σ, 2.5σ, 3.0σ, and 3.5σ finding more and more correlation: this is not surprising since the number of selected pixels decreases and the statistics degrades. This is a due to the regularizing noise.

This approach will be more effective when larger frequency coverage becomes be available.

To improve the constraints on the spectral index of the hard low-frequency component, we re-ran commander without adding noise to the input maps and adopting a different prior, $\beta = -2.15 \pm 0.3$, with the idea that the higher signal-to-noise ratio would allow the data to drive the solution. (In order to speed converge, we also dropped the sample of C_ls.) Except for regions with strong free–free emission, we now observe a mean spectral index of -2.3 ± 0.27 for $|b_{\mathrm{Gal}}| > 30°$, whereas we find -2.4 ± 0.22 for the haze region. Though not conclusive, this does support the hypothesis of a harder spectral index for the haze component proposed by Dobler (2012). The regression coefficients for this case are $(0.4 \pm 2) \times 10^{-7}$, 9 ± 7, and 0.08 ± 2 for the Haslam, Hα, and FDS template, respectively, which show a clear improvement of the component identification. Figure 15 shows the posterior mean amplitude and spectral index of the low-frequency component with a 35° latitude around the GC.

5. CONCLUSIONS

We applied the Gibbs sampling technique implemented in the `Commander` software to *WMAP* 7 year data aiming at characterizing, simultaneously, CMB (map and angular power spectrum) and foreground emission. Our analysis improves previous work in a number of aspects. We pushed the analysis to higher angular resolution, 1°, which enabled us to directly compare the foreground solutions to available templates, resulting in excellent agreement. This represents a success on its own but also confirms the power of the Bayesian approach, which is likely to perform better when the larger Planck data set becomes available.

Higher angular resolution requires better modeling of the noise in the maps after the smoothing procedure: we added scan-modulated regularizing noise estimated via Monte Carlo simulations. Despite the limited number of frequencies available, we did not directly use foreground templates in the Gibbs sampling run, but rather let the parameters of the model vary across the sky. This allowed us to distinguish regions in the sky where a specific foreground mechanism dominates on the basis of the its spectral behavior. The presence of a strongly dust correlated emission at low frequencies, explained by invoking spinning dust grains, emerges not only through the amplitude map but also in the spatial variation of the spectral index of the low-frequency component, which mainly results from a mixture of synchrotron emission and dust-correlated emission. The number of available frequencies still limits us, since it forces us to use a constant spectral index for the dust.

Regressing `Commander` solutions against foreground templates is a complementary way to disentangle the various emission mechanisms. By applying this procedure, we confirm the presence of excess of signal localized around the GC, the microwave haze. A hint of such a emission has been found in the Fermi data (Dobler et al. 2010, 2011). The study of this component has stimulated a rich production, both pointing out possible systematics in the applied regression procedures (see, for instance, Linden & Profumo 2010; Mertsch & Sarkar 2010) and proposing possible physical mechanisms (Bottino et al. 2010; Dobler et al. 2011; Guo et al. 2011; Guo & Mathews 2011; Biermann et al. 2010; Crocker & Aharonian 2011; Mertsch & Sarkar 2011). The *WMAP* team remains skeptical about the presence of the haze (Gold et al. 2011), their main counter argument being the lack of evidence for such a signal in the polarization maps (though Dobler 2012 showed that such a signal would likely not be detectable in the *WMAP* data given the noise). Regarding the spatial correspondence between the *WMAP* haze and the Fermi haze/bubbles, we note that the increased noise in our

output maps makes a direct comparison difficult. Comparison with earlier *WMAP* data was presented by Dobler et al. (2010) and Su et al. (2010), and more recently with *WMAP* 7 year data by Dobler (2012). We find, as in those studies, that the morphological correlation between the two is reasonable at $|b_{\text{Gal}}| < 30°$, where the signal in microwaves is most unambiguous.

We emphasize that our analysis overcomes the problem of CMB subtraction and the circularity argument arising from the removal of an internal linear combination of the channels used in previous analyses. Moreover, we addressed the issue of possible contamination of spinning dust as an explanation of the excess signal by including a correlation between the low-frequency emission and thermal dust. We computed the SED of the dust as described by the FDS model, performing a template fitting of the *WMAP* channels as a preprocessing step. We input the resulting SED to Commander, solving for an amplitude map, together with a low-frequency component. The resulting dust map is highly correlated with the FDS model, and the low-frequency component results from a sum of synchrotron and free–free emission only. When regressing this map against foreground templates, the Galactic haze is still present. One criticism to our approach could be the simplicity of the spinning dust model, but since one of our goals was to be independent of external data sets within a completely Bayesian framework, adding more complexity to the dust model is not possible given the available degrees of freedom.

We also argued that the haze is not an artifact of the spectral variation of the synchrotron component across the sky. To this end, we included the Haslam map at 408 MHz in the Commander run and expanding the foreground model: dust, combing both thermal and spinning contribution, synchrotron emission with fixed spectral response $\beta = -3$, and an additional component with a free spectral index, which would describe free–free, and any other contribution. We found that we could completely separate synchrotron and dust emission, and that we were left with an amplitude map that can be easily characterized by its spectral index. Free–free emission clearly shows up in the Galactic plane, together with the haze that seems to have a harder spectrum than the synchrotron component. Unfortunately, outside the Galactic plane, the index map is quite noisy and the haze does not have a distinctive signature compared to the rest of the sky, although it is definitely not compatible with free–free emission.

Our analysis has improved previous knowledge of the foreground and shed new light on the nature of the Galactic haze. We addressed at least three criticisms often attached to haze studies: coherent CMB removal, contamination from spinning dust, and the spatial variation of the synchrotron component spectral index. The excess of signal seems to be stable with respect to them. However, we have made assumptions on the dust spectral behavior, forced by the number of frequencies available, and we did not rule out the possibility of the presence of curvature in the synchrotron component spectral index or a model with a broken power law. Planck data will address these remaining issues and enable a more complete foreground model.

APPENDIX A

THE CMB MAP AND POWER SPECTRUM POSTERIOR

Here we review the basic concept behind Gibbs sampling. Let us first focus on the case of one frequency map and no foregrounds. The data model for this case is

$$\mathbf{d} = \mathbf{s} + \mathbf{n}, \tag{A1}$$

where $\mathbf{d}$ is the data, $\mathbf{s}$ the CMB sky signal, and $\mathbf{n}$ instrumental noise. We assume both the CMB signal and noise to be Gaussian random fields with vanishing mean and covariance matrices $\mathbf{S}$ and $\mathbf{N}$, respectively. The CMB sky can be written in spherical harmonics as $\mathbf{s} = \sum_{\ell,m} a_{\ell m} Y_{\ell m}$, with the CMB covariance matrix then fully characterized by the angular power spectrum C_ℓ according to $\mathbf{C}_{\ell m, \ell'm'} = \langle a_{\ell m}^* a_{\ell'm'} \rangle = C_\ell \delta_{\ell\ell'} \delta_{mm'}$. The noise matrix $\mathbf{N}$ is left unspecified for now, but we note that for white noise it is diagonal in pixel space, $N_{ij} = \sigma_i^2 \delta_{ij}$, for pixels i and j and noise variance σ_i^2.

Our goal is to sample from the posterior density for both the sky signal $\mathbf{s}$ and the power spectrum C_ℓ, given by

$$P(\mathbf{s}, C_\ell | \mathbf{d}) \propto P(\mathbf{d}|\mathbf{s}, C_\ell) P(\mathbf{s}, C_\ell) \tag{A2}$$

$$\propto P(\mathbf{d}|\mathbf{s}, C_\ell) P(\mathbf{s}|C_\ell) P(C_\ell). \tag{A3}$$

In what follows, we assume the prior $P(C_\ell)$ is uniform. Since we have assumed Gaussianity, the joint posterior distribution may be written as

$$P(\mathbf{s}, C_\ell | \mathbf{d}) \propto e^{-\frac{1}{2}(\mathbf{d}-\mathbf{s})^t \mathbf{N}^{-1}(\mathbf{d}-\mathbf{s})} \prod_\ell \frac{e^{-\frac{2\ell+1}{2}\frac{\sigma_\ell}{C_\ell}}}{C_\ell^{\frac{2\ell+1}{2}}} P(C_\ell), \tag{A4}$$

where we have defined the quantity $\sigma_\ell \equiv \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2$ as the angular power spectrum of the full-sky CMB signal.

For the case here with the CMB signal assumed to be a Gaussian field, one can integrate over the CMB sky signal and analytically solve for the marginalized posterior $P(C_\ell|\mathbf{d})$. However, evaluating the posterior numerically for any specific angular power spectrum is computationally prohibitive as it involves the computation of the inverse and determinant of very large matrices. We therefore sample from the posterior using a Gibbs sampling algorithm.

A.1. Gibbs Sampling

One procedure to sample from the joint density $P(\mathbf{s}, C_\ell | \mathbf{d})$, as proposed by Jewell et al. (2004) and Wandelt et al. (2004), is to alternately sample from the respective conditional densities

$$\mathbf{s}^{i+1} \leftarrow P\left(\mathbf{s}|C_\ell^i, \mathbf{d}\right) \tag{A5}$$

$$C_\ell^{i+1} \leftarrow P(C_\ell|\mathbf{s}^{i+1}, \mathbf{d}). \tag{A6}$$

Here $\leftarrow$ indicates sampling from the distribution on the right-hand side. After some "burn-in" period, the joint samples $(\mathbf{s}^i, C_\ell^i)$ will be distributed from the joint posterior. Thus, the problem is reduced to that of sampling from the two *conditional* densities $P(\mathbf{s}|C_\ell, \mathbf{d})$ and $P(C_\ell|\mathbf{s}, \mathbf{d})$.

The conditional density $P(C_\ell|\mathbf{s}, \mathbf{d})$ in this case is independent of the data, $P(C_\ell|\mathbf{s}, \mathbf{d}) = P(C_\ell|\mathbf{s})$, simply because the underlying CMB sky signal provides all the information needed to estimate the ensemble angular power spectrum C_ℓ. Under the assumption of Gaussianity and isotropy, this conditional is given by the inverse Gamma distribution, S_ℓ:

$$P(C_\ell|\mathbf{s}) \propto \frac{e^{-\frac{1}{2}\mathbf{s}_\ell^t \mathbf{S}_\ell^{-1}\mathbf{s}_\ell}}{\sqrt{|\mathbf{S}_\ell|}} = \frac{e^{-\frac{2\ell+1}{2}\frac{\sigma_\ell}{C_\ell}}}{C_\ell^{\frac{2\ell+1}{2}}}. \qquad (A7)$$

In order to sample from this conditional density, we first draw $2\ell - 1$ normal random variates ρ_ℓ^k, compute the sum $\rho_\ell^2 = \sum_{k=1}^{2\ell-1} |\rho_\ell^k|^2$, and finally set

$$\mathcal{C}_\ell = \frac{\sigma_\ell}{\rho_\ell^2}, \qquad (A8)$$

giving a sample distributed according to the inverse Gamma distribution.

The conditional density for the sky map given the angular power spectrum and data follows directly from the form of the joint Bayes posterior in Equation (A4), and given by

$$P(\mathbf{s}|C_\ell, \mathbf{d}) \propto P(\mathbf{d}|\mathbf{s}, C_\ell)P(\mathbf{s}|C_\ell) \qquad (A9)$$

$$\propto e^{-\frac{1}{2}(\mathbf{d}-\mathbf{s})^t \mathbf{N}^{-1}(\mathbf{d}-\mathbf{s})} \, e^{-\frac{1}{2}\mathbf{s}^t \mathbf{S}^{-1}\mathbf{s}} \qquad (A10)$$

$$\propto e^{-\frac{1}{2}(\mathbf{s}-\hat{\mathbf{s}})^t(\mathbf{S}^{-1}+\mathbf{N}^{-1})(\mathbf{s}-\hat{\mathbf{s}})}, \qquad (A11)$$

where we have defined the so-called mean-field map (or Wiener filtered data) $\hat{\mathbf{s}} = (\mathbf{S}^{-1} + \mathbf{N}^{-1})^{-1}\mathbf{N}^{-1}\mathbf{d}$. Thus, $P(\mathbf{s}|C_\ell, \mathbf{d})$ is a Gaussian distribution with mean equal to $\hat{\mathbf{s}}$ and a covariance matrix equal to $(\mathbf{S}^{-1} + \mathbf{N}^{-1})^{-1}$. In order to sample from this conditional, we first generate two independent white-noise maps ω_0 and ω_1, and solve

$$[\mathbf{S}^{-1} + \mathbf{N}^{-1}]\mathbf{s} = \mathbf{N}^{-1}\mathbf{d} + \mathbf{S}^{-\frac{1}{2}}\omega_0 + \mathbf{N}^{-\frac{1}{2}}\omega_1. \qquad (A12)$$

The resulting map $\mathbf{s}$ is exactly a sample from the conditional $P(\mathbf{s}|C_\ell, \mathbf{d})$. The addition of the white-noise maps simply reflects our uncertainty in the true but unobserved CMB sky—there are many CMB maps that are consistent with the data and power spectrum estimate, and we are simply making a random choice from this set of maps.

The discussion so far was limited to a single band and no modeling of instrumental response. The generalization is straightforward and can be found in Eriksen et al. (2007b), together with a discussion on the actual numerical implementation which reduces the round-off errors. Here we quote the main result, which generalizes Equation (A12) to

$$\left[\mathbf{S}^{-1} + \sum_\nu \mathbf{A}_\nu^t \mathbf{N}_\nu^{-1}\mathbf{A}_\nu\right]\mathbf{s}$$
$$= \sum_\nu \mathbf{A}_\nu^t \mathbf{N}_\nu^{-1}\mathbf{d}_\nu + \mathbf{S}^{-\frac{1}{2}}\omega_0 + \sum_\nu \mathbf{A}_\nu^t \mathbf{N}_\nu^{-\frac{1}{2}}\omega_\nu, \qquad (A13)$$

where $\mathbf{A}_\nu$ describes the beam response of the detector ν and $\mathbf{N}$ its noise properties. Note that we now draw one white-noise map for each frequency band, ω_ν. The sampling procedure for $P(C_\ell|\mathbf{s})$ is unchanged.

A.2. The Foreground Sampler

The Gibbs sampling algorithm can be easily extended to include foregrounds described by a parametric model, $F(\theta)$. Including the foreground model parameters into the joint posterior, the sampling procedure generalizes to

$$\mathbf{s}^{i+1} \leftarrow P(\mathbf{s}|C_\ell^i, \theta^i, \mathbf{d})$$
$$C_\ell^{i+1} \leftarrow P(C_\ell|\mathbf{s}^{i+1}, \mathbf{d})$$
$$\theta^{i+1} \leftarrow P(\theta|C_\ell^{i+1}, \mathbf{s}^{i+1}, \mathbf{d}).$$

We note that for the foreground models of interest, we do not have an algorithm to produce an exact sample from the conditional density $P(\theta|C_\ell^{i+1}, \mathbf{s}^{i+1}, \mathbf{d})$ but we instead generalize Gibbs sampling to an MCMC algorithm, the details of which can be found in Eriksen et al. (2008b).

The parametric family of data models including foregrounds implemented in `Commander` is of the form

$$\mathbf{d}_\nu = \mathbf{A}_\nu \mathbf{s} + \sum_{m=1}^{M} a_{\nu,m}\mathbf{t}_m$$
$$+ \sum_{n=1}^{N} b_n f_n(\nu)\mathbf{f}_n + \sum_{k=1}^{K} \mathbf{c}_k \, \mathbf{g}_k(\nu; \theta_k) + \mathbf{n}_\nu. \qquad (A14)$$

We may identify three main classes of foregrounds.

1. $\mathbf{t}_m$ are M templates multiplied by an amplitude at every frequency, $a_{\nu,m}$.
2. $\mathbf{f}_n$ are N templates whose spectral behavior is known and described by the function $f_n(\nu)$; we allow for an overall rescaling, b_n, for each of them.
3. K foregrounds are described by a map of coefficients, $\mathbf{c}_k$, multiplied by the spectral response which is function of the frequency and the parameters of the foreground model, θ_k.

We remind the reader that the bold face notation means an array of size N_{pix}. An example of the first class of foregrounds is given by monopole and dipole residuals; a special case of the second class is a free–free template (Figure 1, top panel) whose spectral behavior is known and follows the relation $(\nu/\nu_0)^{-2.15}$; for the third type we may quote synchrotron emission described by an amplitude we solve for at the reference frequency, e.g., $\mu_0 = 23$ GHz, and a spectral response given by $(\nu/\mu_0)^\beta$.

It is instructive to compute the degrees of freedom for such a foreground model applied to *WMAP* data. We solve three maps, $\mathbf{s}$, β, and $\mathbf{A}_{\mathrm{synch}}$, five monopoles and dipoles, and two overall amplitudes if we describe free–free emission and thermal dust by means of the second class of foreground models. In total we look for 3 $N_{\mathrm{pix}} + 22$ parameters. This is already pretty close to the maximum number of parameters allowed by the *WMAP*, $\leqslant 5 N_{\mathrm{pix}}$. We could ask for one more map, either dust or free–free, assuming a single power law and fixing the spectral index, or allowing a curvature term in the synchrotron model. Since a simple power law for thermal dust and free–free emission has been shown to be consistent with the data in the frequency range spanned by *WMAP* (Dickinson et al. 2009), this may be used to disentangle the spinning dust contribution from thermal dust.

We notice that to solve for the spectral index of a given component, the instrumental beam response of each channel must be taken into account. Up until now, `Commander` is able to work with maps at the same angular resolution only. Smoothing all frequency maps and ancillary data to a common angular scale is then a necessary preprocessing step.
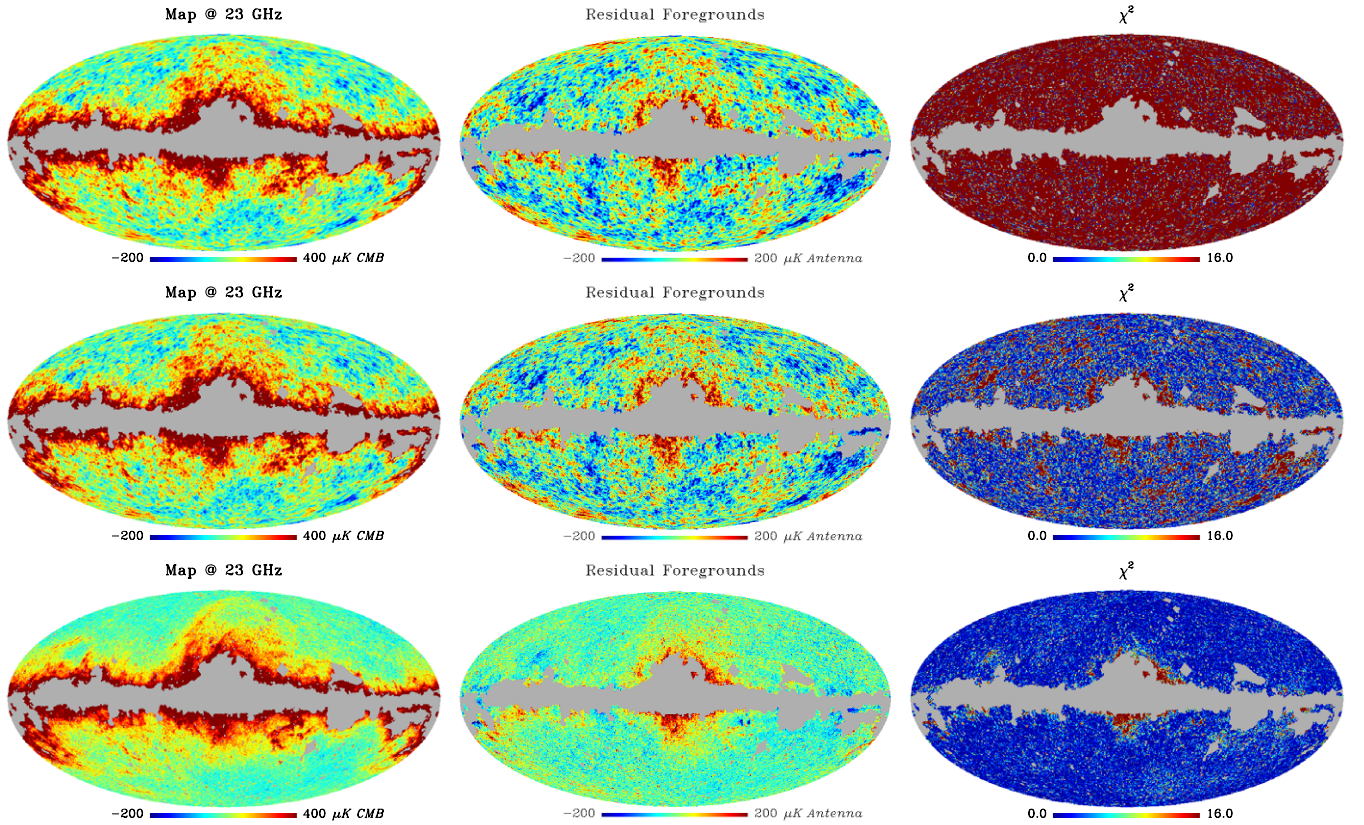
Figure 16. Example of three regressions performed on the *WMAP* K band smoothed to 60 arcmin (top), `Commander` input K band (middle row), and output foreground amplitude (bottom).

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

Table 4
Regression Coefficients of the `Commander` Foreground Amplitude Maps
Compared to Those Obtained when Regressing Smoothed *WMAP*
Channel and `Commander` Input *WMAP* Maps

	WMAP 7 yr			
Data Set	Haslam	Hα	FDS	r
Smoothed K band	$(3.7 \pm 0.17) \times 10^{-6}$	$6.4 \pm .6$	$7.0 \pm .25$	0.91
`Commander` input K band	$(3.7 \pm 1.1) \times 10^{-6}$	6 ± 5	7 ± 2	0.85
`Commander` output @ 23 GHz	$(3.6 \pm 1.2) \times 10^{-6}$	6 ± 6	7 ± 2	0.94

As discussed in Eriksen et al. (2008b), it turns out to be more efficient to sample all the map amplitudes at once, followed by the spectral response parameter, and finally the angular power spectrum, following the iterative scheme:

$$\{\mathbf{s}^{i+1}, a_m, b_n, \mathbf{c}_k\}^{i+1} \leftarrow P(\mathbf{s}, a_m, b_n, \mathbf{c}_k | C_\ell^i, \theta^i, \mathbf{d})$$
$$\theta^{i+1} \leftarrow P(\theta | C_\ell^i, \mathbf{s}^{i+1}, a_m^{i+1}, b_n^{i+1}, \mathbf{c}_k^{i+1}, \mathbf{d})$$
$$C_\ell^{i+1} \leftarrow P(C_\ell | \mathbf{s}^{i+1}, \mathbf{d}). \qquad (A15)$$

Sampling from the conditional density $P(\mathbf{s}, a_m, b_n, \mathbf{c}_k | C_\ell^i, \theta^i, \mathbf{d})$ is a generalization of Equation (A13), whereas the sampling of the angular power spectrum remains unchanged, since $\mathcal{C}_\ell$ are functions of the sky signal only. Sampling of the nonlinear degrees of freedom is through a standard inversion sampler: first, compute the conditional probability density $P(x|\theta)$, where x is the currently sampled parameter and θ denotes the set of all other parameters in the model, assuming the likelihood to be independent pixel by pixel: $-2\ln\mathcal{L}(x) = \chi^2 =$

$\sum_\nu (d_\nu - s_\nu(X, \theta))^2 / \sigma_\nu^2$. Then, the corresponding cumulative distribution is computed, $F(x|\theta) = \int_{-\infty}^{x} P(y|\theta)dy$, and the value of the x variable is chosen by drawing a uniformly distributed random number, u, and reading $F(x|\theta) = u$.

This concludes the review on the implementation of Gibbs sampling as implemented in the computer code `Commander`.

APPENDIX B

NOISE IMPACT ON THE REGRESSION PROCEDURE

We have observed that the regression coefficients we found when fitting foreground templates to `Commander` posterior mean amplitudes are consistent with those discussed in other works, but they have much larger errors. We argued in the text that this is the result of the sampling procedure and has two causes: (1) our input maps are noisier because of the additional noise term and (2) our uncertainties on the foreground amplitudes take into account the error on the other parameters of the model: CMB, other foreground amplitudes, and spectral indices. In this respect, our errors are more conservative.

To clearly show this, we perform the same regression described in Equations (2) and (3) on two different maps: (1) smoothed *WMAP* K band and (2) `Commander` input K band, which have to be compared to `Commander` output amplitude performance. The three maps are shown in Figure 16, together with the fit residuals and the corresponding χ^2.

Table 4 compares the regression coefficients for the signal at 23 GHz. Consistently, the amplitudes are the same but the error bars increase dramatically due to the noise added to the maps. In particular, we move from 10σ–30σ detection to 1σ–3σ, which

is driven by the scaling factor applied to the *K* band (16; see Section 2.1).

This comparison suggests that a lower level of noise added to the input maps is useful to better characterize diffuse foreground emission. We will further investigate this issue in a forthcoming work.

REFERENCES

Bennett, C., Hill, R. S., Hinshaw, G., et al. 2003, ApJS, 148, 97
Biermann, P. L., Becker, J. K., Caceres, G., et al. 2010, ApJ, 710, L53
Bottino, M., Banday, A. J., & Maino, D. 2010, MNRAS, 402, 207
Casassus, S., Dickinson, C., Cleary, K., et al. 2008, MNRAS, 391, 1075
Cholis, I., Dobler, G., Finkbeiner, D. P., Goodenough, L., & Weiner, N. 2009, Phys. Rev. D, 80, 123518
Chu, M., Eriksen, H. K., Knox, L., et al. 2005, Phys. Rev. D, 71, 103002
Crocker, R. M., & Aharonian, F. 2011, Phys. Rev. Lett., 106, 101102
Crocker, R. M., Jones, D. I., Aharonian, F., et al. 2011, MNRAS, 413, 763
Cumberbatch, D. T., Zuntz, J., Kamfjord Eriksen, H. K., & Silk, J. 2009, arXiv:0902.0039
Dickinson, C., Eriksen, H. K., Banday, A. J., et al. 2009, ApJ, 705, 1607
Dobler, G. 2012, ApJ, 750, 17
Dobler, G., Cholis, I., & Weiner, N. 2011, ApJ, 741, 25
Dobler, G., Draine, B., & Finkbeiner, D. P. 2009, ApJ, 699, 1374
Dobler, G., & Finkbeiner, D. P. 2008a, ApJ, 680, 1222
Dobler, G., & Finkbeiner, D. P. 2008b, ApJ, 680, 1235
Dobler, G., Finkbeiner, D. P., Cholis, I., Slatyer, T., & Weiner, N. 2010, ApJ, 717, 825
Draine, B. T., & Lazarian, A. 1998, ApJ, 508, 157
Erickson, W. C. 1957, ApJ, 126, 480
Eriksen, H. K., Dickinson, C., Jewell, J. B., et al. 2008a, ApJ, 672, L87
Eriksen, H. K., Dickinson, C., Lawrence, C. R., et al. 2006, ApJ, 641, 665
Eriksen, H. K., Huey, G., Banday, A. J., et al. 2007a, ApJ, 665, L1
Eriksen, H. K., Huey, G., Saha, R., et al. 2007b, ApJ, 656, 641
Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008b, ApJ, 676, 10
Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, ApJS, 155, 227
Fernández-Cerezo, S., Gutiérrez, C. M., Rebolo, R., et al. 2006, MNRAS, 370, 15
Finkbeiner, D. P. 2003, ApJS, 146, 407
Finkbeiner, D. P. 2004, ApJ, 614, 186
Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, ApJ, 524, 867
Ghosh, T., Banday, A. J., Jaffe, T., et al. 2012, MNRAS, 422, 3617
Gold, B., Odegard, N., Weiland, J. L., et al. 2011, ApJS, 192, 15
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Guo, F., & Mathews, W. G. 2011, arXiv:1103.0055
Guo, F., Mathews, W. G., Dobler, G., & Oh, S. P. 2011, arXiv:1110.0834
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
Hildebrandt, S. R., Rebolo, R., Rubiño-Martín, J. A., et al. 2007, MNRAS, 382, 594
Hoang, T., Draine, B. T., & Lazarian, A. 2010, ApJ, 715, 1462
Hoang, T., Lazarian, A., & Draine, B. T. 2011, ApJ, 741, 87
Hooper, D., Finkbeiner, D. P., & Dobler, G. 2007, Phys. Rev. D, 76, 083012
Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, ApJS, 192, 14
Jewell, J., Levin, S., & Anderson, C. H. 2004, ApJ, 609, 1
Jewell, J. B., Eriksen, H. K., Wandelt, B. D., et al. 2009, ApJ, 697, 258
Kogut, A., Dunkley, J., Bennett, C. L., et al. 2007, ApJ, 665, 355
Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, ApJS, 192, 18
Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, ApJS, 192, 16
Larson, D. L., Eriksen, H. K., Wandelt, B. D., et al. 2007, ApJ, 656, 653
Linden, T., & Profumo, S. 2010, ApJ, 714, L228
Mertsch, P., & Sarkar, S. 2010, J. Cosmol. Astropart. Phys., JCAP10(2010)019
Mertsch, P., & Sarkar, S. 2011, Phys. Rev. Lett., 107, 091101
O'Dwyer, I. J., Eriksen, H. K., Wandelt, B. D., et al. 2004, ApJ, 617, L99
Park, C.-G., Park, C., & Gott, J. R. 2007, ApJ, 660, 959
Planck Collaboration 2011, A&A, 536, A20
Rudjord, Ø., Groeneboom, N. E., Eriksen, H. K., et al. 2009, ApJ, 692, 1669
Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, ApJ, 724, 1044
Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, Phys. Rev. D, 70, 083511