



SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments

Original

Maximum likelihood algorithm for parametric component separation in cosmic microwave background experiments / Stompor, R.; Leach, S.; Stivoli, F.; Baccigalupi, C.. - In: MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY. - ISSN 0035-8711. - 392:1(2009), pp. 216-232. [10.1111/j.1365-2966.2008.14023.x]

Availability:

This version is available at: 20.500.11767/16068 since:

Publisher:

Published

DOI:10.1111/j.1365-2966.2008.14023.x

Terms of use:

Testo definito dall'ateneo relativo alle clausole di concessione d'uso

Publisher copyright

note finali coverpage

(Article begins on next page)

Maximum Likelihood algorithm for parametric component separation in CMB experiments

Radek Stompor¹, Samuel Leach², Federico Stivoli², Carlo Baccigalupi²

¹ *CNRS, UMR 7164, Laboratoire Astroparticule & Cosmologie, 10 rue A. Domon et L. Duquet, F-75205 Paris Cedex 13, France*

² *SISSA/ISAS, Astrophysics Sector, Via Beirut, 4, and INFN, Sezione di Trieste, Via Valerio 2, I-34014 Trieste, Italy*

21 January 2009

ABSTRACT

We discuss an approach to the component separation of microwave, multi-frequency sky maps as those typically produced from Cosmic Microwave Background (CMB) Anisotropy data sets. The algorithm is based on the two step, parametric, likelihood-based technique recently elaborated on by Eriksen et al. (2006), where the foreground spectral parameters are estimated prior to the actual separation of the components. In contrast with the previous approaches, we accomplish the former task with help of an analytically-derived likelihood function for the spectral parameters, which, we show, yields estimates equal to the maximum likelihood values of the full multi-dimensional data problem. We then use these estimates to perform the second step via the standard, generalised-least-square-like procedure. We demonstrate that the proposed approach is equivalent to a direct maximization of the full data likelihood, which is recast in a computationally tractable form. We use the corresponding curvature matrices to characterise statistical properties of the recovered parameters. We incorporate in the formalism some of the essential features of the CMB data sets, such as inhomogeneous pixel domain noise, unknown map offsets as well as calibration errors and study their consequences for the separation. We find that the calibration is likely to have a dominant effect on the precision of the spectral parameter determination for a realistic CMB experiment. We apply the algorithm to simulated data and discuss the results. Our focus is on partial-sky, total-intensity and polarization, CMB experiments such as planned balloon-borne and ground-based efforts, however, the techniques presented here should be also applicable to the full-sky data as for instance, those produced by the WMAP satellite and anticipated from the Planck mission.

1 INTRODUCTION AND MOTIVATION

The measurement and characterization of the Cosmic Microwave Background (CMB) polarization anisotropies is an important goal of the present day observational cosmology. It poses a significant experimental challenge which is aimed at by an entire slew of currently operating and forthcoming experiments such as PolarBear, QUIET, and Clover, observing from the ground, EBEX and Spider, observing from the balloon platform, and WMAP and Planck observing from space¹.

The sensitivity of modern CMB experiments is rapidly increasing as the instruments incorporate more detectors and observe the sky for longer periods of time. Consequently, instrumental systematics and foreground emissions have come to the forefront as the major limiting factor in the effort of detecting and characterizing the primordial, cosmological signals. For total intensity measurements, recently extended down to arcminute scales by ACBAR (Reichardt et al. 2008), and at current levels of sensitivity, the diffuse galactic foregrounds have turned out luckily to be quite benign. However diffuse foreground emission remains a serious concern for CMB polarimetry experiments.

¹ PolarBear: <http://bolo.berkeley.edu/polarbear/science/science.html>, QUIET: <http://quiet.uchicago.edu/>, Clover: <http://www-astro.physics.ox.ac.uk/research/expcosmology/groupclover.html>, EBEX: <http://groups.physics.umn.edu/cosmology/ebex/>, Spider: http://www.astro.caltech.edu/~lgg/spider_front.htm, WMAP: <http://map.gsfc.nasa.gov/>, Planck: http://www.esa.int/esaSC/120398_index_0_m.html. See also <http://lambda.ngfc.nasa.gov> for a complete list of the ongoing and forthcoming CMB experiments.

CMB polarization signal can be decomposed into a gradient ‘ E mode’ component, which is expected to be dominated by adiabatic scalar perturbations, and into a curl ‘ B mode’ contribution, which can be generated by primordial gravitational waves (Kamionkowski et al. 1997; Zaldarriaga and Seljak 1997) and lensing of the E mode by intervening structures (Zaldarriaga and Seljak 1998; Lewis and Challinor 2006). Recent experiments targeting low foreground emission regions of sky are successfully measuring the E mode of CMB on degree and sub-degree angular scales, most recently by CAPMAP (for a compilation of recent results, see Bischoff et al. 2008). However on large angular scales WMAP have reported a foreground contamination exceeding both the E and B modes of the CMB in the frequency range of the minimum Galactic emission (Page et al. 2007). Recent analyses of the WMAP five-year data (Gold et al. 2008; Dunkley et al. 2008) have seen an renewed focus on studying the impact of foreground subtraction errors on the determination of cosmological parameters, using a variety of analysis methods and assumptions about the data.

Multi-frequency component separation methods differ in the extent to which they combine the data with ‘external’ sources of information, and the way in which the components are modeled. The approaches studied to date include methods based on the internal template subtraction (e.g. Bennett et al. 1992; Hansen et al. 2006), those exploiting statistical independence of (some of) the sky components (e.g. Delabrouille et al. 2003; Maino et al. 2007; Bonaldi et al. 2006), those invoking the maximum-entropy principle (e.g. Stolyarov et al. 2005), or performing a parametric fit to the data (Brandt et al. 1994; Eriksen et al. 2006). Those latter methods, called parametric, are distinguished by the explicit modeling of the frequency scaling of all components and typically by the large number of free parameters describing the component amplitudes. See the Planck Working Group 2 work, Leach et al. (2008), for more information about the component separation methods and a comparison of their performance.

A common requirement for any component separation technique for use in cosmological analyses is the capability of propagating errors due to foreground subtraction, while having the flexibility to incorporate foreground modeling and external constraints in a transparent way. Against the backdrop of these two requirements, we will consider the parametric approach to the foreground separation problem.

1.1 Parametric component separation

The defining premise of the parametric approach is an assumption that the functional form of the frequency scaling for all relevant components is known, and our ignorance can be quantified by means of relatively few, though non-linear, spectral parameters. The parameters of the model are determined via a fitting procedure, often performed on a pixel-by-pixel basis. The attractiveness of this approach lies in its simplicity, the flexibility of the parametrization schemes and possibility of phrasing the separation problem in a coherent maximum likelihood form (Jaffe et al. 2004). Its strength lies in exploiting nearly optimally the prior knowledge of the frequency scaling of different components, which in several instances is quite well known, while ignoring the information which is less known and more debatable. Its weakness is related to the difficulties in performing the non-linear and high-dimensional parameter fits especially in the realm of the CMB data, which are usually characterised by low signal-to-noise ratio on the pixel scale, and numerical efficiency, given a large number of pixels, $\mathcal{O}(10^7)$, anticipated from the next generation of the CMB experiments.

In a recent paper, Eriksen et al. (2006) reconsidered a parametric approach to the component separation task originally introduced by Brandt et al. (1994) and proposed its rendition which was shown to be both numerically efficient and stable in the application to the nearly full-sky simulated Planck and WMAP data. In their method, the numerical efficacy and stability is achieved by splitting an estimation of the parameters and component maps into two separate steps. The spectral parameter fitting is performed pixel-by-pixel on the input data which are first smoothed and underpixelized during a preprocessing stage. This enhances the signal-to-noise ratio of the data and thus improves the stability and robustness of the fitting procedure. It also makes the procedure more numerically tractable as fewer fits have to be performed. The fitting criterion, employed by Eriksen et al. (2006), uses a likelihood function, and involves the (non-linear) spectral parameters as well as (linear) component amplitudes of the smoothed input maps. In the modern twist to the method, they propose to use a Monte-Carlo Markov Chain (MCMC) sampling technique as a way to determine the best-fitting parameters, and which elegantly allows to marginalise over low resolution components leaving one with the estimates of the spectral parameters. The latter are taken to be the averages rather than peak values of the sampled likelihood. On the second step of the method, the recovered low-resolution parameters are then applied to the corresponding pixels of the full resolution maps to render the high resolution component map estimates. More details on the method can be found in Eriksen et al. (2006). In this paper, we modify the parametric method of Eriksen et al. (2006) and devise the genuinely maximum likelihood parametric algorithm, while preserving the efficiency of its two step design. We then discuss in some detail the general features of the technique, and also demonstrate how it lends itself to a number of generalizations relevant for the analysis of actual observational data, such as calibration errors and arbitrary offsets of the input maps.

The plan of the paper is as follows. In the remaining part of the Introduction, we define our notation. In Section 2 we describe the simulated sky data which we use in the examples discussed throughout the paper. In Section 3 we cast the parametric approach to component separation as a maximum likelihood problem, showing how a two step procedure allows




Figure 1. The RMS of the simulated signal amplitudes for the CMB (dashed) and thermal dust (dotted line) smoothed with an 8' beam, for total intensity I and Stokes Q parameter (**left and right panels**) as a function of frequency. The noise RMS per 3.4' pixel ($n_{\text{side}} = 1024$) is also shown for the frequency bands of the experiment we consider here.

to recover frequency scaling and spatial patterns of the components. In Section 4 we extend the formalism by considering the presence of offsets and calibration errors into the data, and Section 5 contains our conclusions.

Notation

We use bold typeface to mark vectors and matrices, written in lower and upper case letters respectively. A vector corresponding to a pixel, p , or a matrix – to a pair of pixels, p, q , is always given a subscript p or pq , respectively. (However, if $p = q$ in the latter case we usually drop one of the p .) Such an object will usually contain all the values for all relevant Stokes parameters and frequency channels (or alternately, sky components) characterizing a chosen pixel on the sky. Multi-pixel vectors are always composed of pixel-specific vectors concatenated together so that the information related to a first pixel precedes that of a second one etc. The multi-pixel matrices are made of pixel pairs specific blocks arranged in an analogous manner.

2 SIMULATED DATA SET

We simulate the CMB and thermal dust polarized emission on a patch of sky with an area of 350 square degrees centered at RA=60°, Dec=-50°, adopting the HEALPix² pixelisation scheme (Górski et al. 2005) with 3.4' pixels ($n_{\text{side}}=1024$). This region of sky is known to be characterised by low dust emission in total intensity, having been observed by Boomerang (Masi et al. 2006; Montroy et al. 2006) and QUaD (Ade et al. 2007). The polarized emission in this region of sky and in this frequency range is low but to date has not been as well characterised. This area is close to the anticipated EBEX field (Oxley et al. 2004), which we will take as our reference experiment. Specifically we simulate three channels at 150, 250 and 410 GHz each with 8' resolution, and a noise level with RMS per pixel of 0.56, 0.66, 1.13 μK respectively, in antenna (Rayleigh-Jeans) units; for Stokes Q and U the noise level is a factor $\sqrt{2}$ higher. The noise is uncorrelated between pixels, frequency channels and Stokes parameters. The adopted here noise levels correspond to the deep integration anticipated for B mode CMB experiments. The CMB sky is a Gaussian realisation of the Spergel et al. (2007) best-fitting C_ℓ , except with tensor perturbations added in at a level $T/S = 0.1$.

Thermal dust emission is simulated here using the following strategy: we use ‘Model 8’ of Finkbeiner et al. (1999) in order to extrapolate the combined COBE-DIRBE and IRAS dust template of Schlegel et al. (1998) to 65 GHz where a comparison with WMAP can be performed. We then assume a constant dust polarization fraction $p = 11\%$ across our patch. This value is in broad agreement with measurements from Archeops which found $p \simeq 5\%$ at low galactic latitudes (Benoit et al. 2004; Ponthieu et al. 2005), taking into account the fact that less depolarization along the line of sight is expected at higher galactic latitudes (Prunet et al. 1998). The model adopted here agrees also, as far as the slope and amplitude is concerned, with the WMAP results based on the large-scale diffuse polarized foreground emission at intermediate Galactic latitudes and as summarised in equation (25) of Page et al. (2007). The polarization angle is given by the WMAP dust template on large angular scales (Page et al. 2007; Kogut et al. 2007) and we add in small scale power decaying as $C_\ell \simeq 45 \times \ell^{-3} \mu\text{K}_{\text{CMB}}^2$ for the E and B mode spectra. In order to extrapolate the dust from 65GHz back up to our frequency range we follow Eriksen et al. (2006) in using ‘Model 3’ of Finkbeiner et al. (1999),

$$s_d(\nu) = A_d \frac{\nu}{\exp \frac{h\nu}{kT_d} - 1} \frac{\exp \frac{h\nu_0}{kT_d} - 1}{\nu_0} \left(\frac{\nu}{\nu_0} \right)^\beta, \quad (1)$$

² <http://healpix.jpl.nasa.gov>

with $T_d = 18.1\text{K}$ and $\beta = 1.65$, assumed constant on the patch. In Figure 1 we plot the frequency scaling of the CMB and dust RMS for total intensity and polarization. The total intensity observations are characterised by higher signal-to-noise and higher CMB-to-dust ratios than polarization observations.

3 MAXIMUM LIKELIHOOD PROBLEM FOR COMPONENT SEPARATION – BASIC FORMALISM

We will model the multi-frequency sky signal as,

$$\mathbf{d}_p = \mathbf{A}_p \mathbf{s}_p + \mathbf{n}_p. \quad (2)$$

Here, \mathbf{d}_p is a data vector containing the measured or estimated signals for all n_f frequencies and n_s Stokes parameters, which are to be analyzed simultaneously; \mathbf{s}_p is a vector of the underlying true, and thus to be estimated from the data, values of n_s Stokes parameters for each of the n_c components; $\mathbf{A}_p \equiv \mathbf{A}_p(\boldsymbol{\beta})$ is a component ‘mixing matrix’, which hereafter will be assumed to be parameterised by a set of unknown parameters $\{\boldsymbol{\beta}_i\}$. The mixing matrix is in general rectangular as the number of components can not be larger than the number of the frequency channels for which the data are available. In the examples discussed below, we allow only for one unknown spectral parameter, the dust slope index, β , fixing the dust temperature. Consequently, the elements of the mixing matrix, \mathbf{A} , corresponding to the dust component, are given by (equation 1),

$$\mathbf{A}_{\nu, \text{dust}} = \frac{s_d(\nu, \beta)}{s_d(\nu_0, \beta)}, \quad (3)$$

for each considered frequency band ν , where ν_0 stands for a fiducial, reference frequency set here to 150GHz. The expression for the CMB column of the matrix \mathbf{A} is analogous but does not involve any unknown parameters, as its frequency scaling is assumed to be perfectly that of the black-body (in antenna units).

With these definitions and models in hand we can write a single pixel log-likelihood which up to an irrelevant constant is given by,

$$-2 \ln \mathcal{L}(\mathbf{s}_p, \boldsymbol{\beta}_i) = \text{CONST} + (\mathbf{d}_p - \mathbf{A}_p \mathbf{s}_p)^t \mathbf{N}_p^{-1} (\mathbf{d}_p - \mathbf{A}_p \mathbf{s}_p). \quad (4)$$

Here, \mathbf{N}_p is a square, symmetric noise matrix of the frequency maps for a pixel p , with rank $n_s \times n_f$. This likelihood is the basis of the Eriksen et al. (2006) approach. It is clearly straightforward to introduce a multi-pixel version of this likelihood, and with the definition as above we have,

$$-2 \ln \mathcal{L}_{\text{data}}(\mathbf{s}, \boldsymbol{\beta}) = \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}), \quad (5)$$

where all the matrices and vectors now span over many pixels. Hereafter, we will refer to this likelihood as the full data likelihood. In the simple case of the matrix \mathbf{N} being block-diagonal, thus allowing only for correlations between different Stokes parameters at each pixel, equation (5) simplifies to read,

$$-2 \ln \mathcal{L}_{\text{data}}(\mathbf{s}, \boldsymbol{\beta}) = \text{CONST} + \sum_p (\mathbf{d}_p - \mathbf{A}_p \mathbf{s}_p)^t \mathbf{N}_p^{-1} (\mathbf{d}_p - \mathbf{A}_p \mathbf{s}_p). \quad (6)$$

This likelihood reaches its maximum for the values of \mathbf{s} and $\boldsymbol{\beta}$ fulfilling the relations,

$$-(\mathbf{A}_{,\beta} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) = 0 \quad (7)$$

$$\mathbf{s} = (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^t \mathbf{N}^{-1} \mathbf{d}, \quad (8)$$

where $_{,\beta}$ denotes a partial derivative with respect to $\boldsymbol{\beta}_i$. Solving this system of equations above can be problematic, in particular in the case of multiple spectral parameters, $\{\boldsymbol{\beta}_i\}$, due to their generally non-linear character (Brandt et al. 1994; Eriksen et al. 2006). However, once the values of the spectral parameters are found the second set of linear equations provides straightforward estimates of the pixel amplitudes. Equations (5), (6), & (8), are the starting point for our discussion in the following.

3.1 Constraining the spectral parameters

3.1.1 Maximum likelihood estimator

Equation (8) provides an estimate of the sky components for any assumed set of spectral parameter values, $\boldsymbol{\beta}$. This estimate maximizes locally the full data likelihood restricted to a subspace of the parameter space as defined by the assumed values of the spectral parameters. We can use that equation to eliminate the sky signals from the full data likelihood expression, equation (5), obtaining,

$$-2 \ln \mathcal{L}_{\text{spec}}(\boldsymbol{\beta}) = \text{CONST} - (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d})^t (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d}). \quad (9)$$




Figure 2. The one dimensional likelihoods for the slope index, β , with the sky component amplitudes marginalised over and computed for the total intensity data and the Stokes Q parameter (**left and right panels**). The histograms show the result of the MCMC sampling of the full data likelihood, equation (5). The thick solid and dashed lines show the result for the marginal and spectral likelihood expressions, equations (10) and (9) respectively. The dotted vertical lines show the β value corresponding to a peak of the full data likelihood equation (5), found through a direct grid search. The displayed results are calculated for a single pixel with $n_{side} = 1024$ for the total intensity, and $n_{side} = 16$ for the polarization. The true value of the slope index, $\beta = 1.65$, is indicated by the solid vertical lines.

The above expression permits a calculation of the ‘ridge’ values of the full data likelihood, i.e., the maximum value of the likelihood as attained for any chosen value of the spectral parameters and, therefore, is maximized by the true (global) maximum likelihood value of the spectral parameters. Hereafter, we will call this function a spectral likelihood and use it to determine the maximum likelihood estimates of the spectral parameters and obtain an insight into the resulting errors. We point out that the spectral likelihood is not a proper likelihood in a common sense, and is thus often referred to as a ‘ridge’ or ‘profile’ likelihood to emphasize this fact (e.g., Barndoff-Nielsen & Cox 1994). Nevertheless, it permits to retrieve the maximum likelihood values of the spectral parameters, and provides useful, if not strictly precise, estimates of their uncertainty.

In the alternate approach as implemented by Eriksen et al. (2006) the spectral parameters are estimated with help of the marginal distribution derived for the full data likelihood, equation (5). The marginalisation over the sky amplitudes is performed numerically, via MCMC, and assumes flat priors for the sky signals. We will show here that that approach is not only more computationally involved but also fails to reproduce the maximum likelihood estimates of the spectral parameters derived by maximizing equation (5).

We first observe that the marginalisation procedure can be done analytically and that the integration over the sky amplitudes leads to,

$$\begin{aligned} -2 \ln \mathcal{L}_{marg}(\beta) &= -2 \ln \int d\mathbf{s} \exp \left[-\frac{1}{2} (\mathbf{d} - \mathbf{A}\mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A}\mathbf{s}) \right] \\ &= \text{CONST} - (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d})^t (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d}) + \ln \left| (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \right|, \end{aligned} \quad (10)$$

where $|\dots|$ denotes the matrix determinant. This expression defines the likelihood function for the parameters, β , given the data, \mathbf{d} , assuming an explicit independence of any spectral parameter, β , on \mathbf{s} . (We remark here in passing that though this equation is derived under the assumption of the flat priors for the component amplitudes. However, any kind of a spatial template-like prior constraint with known, and Gaussian, uncertainty can be easily incorporated here, and treated as yet another data set in equation (5).) In Figure 2 we show that indeed our analytical result for the marginal likelihood reproduces perfectly an outcome of the numerical marginalisation (via MCMC) of the full data likelihood.

The marginalised and spectral likelihoods, equations (10) & (9), clearly differ. As the difference depends on the spectral parameters, we can conclude that the marginal likelihood peaks at their values, which do not coincide with their maximum likelihood estimates. We can see this more explicitly by calculating the first derivative of equation (10),

$$\begin{aligned} -2 \frac{\partial \ln \mathcal{L}_{marg}}{\partial \beta_i} &= -2 (\mathbf{A}_{,\beta_i}^t \mathbf{N}^{-1} \mathbf{d})^t \hat{\mathbf{N}} (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d}) + 2 (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d})^t \hat{\mathbf{N}} \mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta_i} \hat{\mathbf{N}} (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{d}) \\ &\quad - 2 \text{tr} \left[\hat{\mathbf{N}} (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta_i}) \right], \end{aligned} \quad (11)$$

which, in general, does not vanish for the maximum likelihood values of the spectral indices. Indeed, we have for $\beta = \beta_{ML}$, as defined by equations (7) & (8),

$$-2 \left. \frac{\partial \ln \mathcal{L}_{marg}}{\partial \beta_i} \right|_{\beta=\beta_{ML}} = -2 \text{tr} \hat{\mathbf{N}} \mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta_i} \Big|_{\beta=\beta_{ML}}. \quad (12)$$




Figure 3. The one dimensional likelihoods for the slope index, β , with the sky component amplitudes marginalised over and computed for the total intensity and Stokes Q parameter data (**two left and two right panels**). The results shown here are from the analysis involving 4096 high resolution pixels. The solid and dotted lines (the latter missing in the Q data case, see text) represent the spectral and marginalised likelihood cases, equations (9) and (10) respectively, computed using the *multi-pixel approach*. The dashed and dot-dashed lines (nearly perfectly overlapping) show the corresponding results derived in the *single-pixel approach*. The homogeneous and inhomogeneous noise cases are shown by the first and second panel, respectively, of each pair of panels. The true value, $\beta = 1.65$, is indicated by the solid vertical lines.

The two major observations that we have made in this Section are visualized in Figure 2. It shows the presence of a mismatch between the maximum likelihood value of the slope index derived from the marginalised likelihood and that obtained from the full data likelihood, as well as a lack of it if the spectral likelihood, equation (9), is utilized, instead of the marginalised one. Note that all these effects, though small in the presented examples due to our assumptions, are however clearly discernible. We present a more detailed discussion of the numerical results in the next Section.

Deriving the best possible constraints on a set of the spectral parameters requires using in the likelihood analysis all the available data points for which the frequency scaling of the component is, or can be assumed the same. When the scaling is identical for all the data analyzed, the mixing matrix \mathbf{A} is made of identical, diagonal blocks, a fact which can be additionally exploited to further speed up the numerical calculations and/or simplify the implementation of the formalism. However, the expressions derived above are fully general and can be applied in the cases with arbitrary mixing matrices for instance, such as those corresponding to multiple sets of the spectral parameters, each describing the frequency scaling of different subsets of all pixels. All the scaling parameters are then estimated simultaneously, and we will use this fact later in the paper.

3.1.2 Comparison and examples

We can apply the formulae presented here without any preprocessing to subsets of the available maps for which the assumption of a single set of spectral parameters is physically justifiable. Alternately, as proposed by Eriksen et al. (2006) we could first smooth the available data and downgrade it to a lower resolution, prior to calculating the spectral likelihood for each low-resolution pixel separately. The spectral parameters determined in this way are then applied to all the high resolution pixels falling into the low resolution one. Hereafter, we refer to these two implementations of the formalism as a *multi-pixel* and *single-pixel* approach, respectively.

Figure 3 shows the one dimensional likelihood for the spectral index β – both the marginalised, equation (5), and spectral, equation (9), renditions as computed in the single- and multi-pixel approaches. The two left panels show the total intensity cases for homogeneous and inhomogeneous noise distribution. The latter was obtained by allowing for a $\sqrt{3}$ noise RMS variation across the patch, characterised by the same value on zones with size of about $14'$ ($n_{\text{side}} = 256$), but changing randomly from zone to zone. The solid and dotted lines represent the spectral and the marginalised likelihood, respectively in the multi-pixel approach, while the nearly overlapping dashed and dot-dashed lines are for the single-pixel analysis. In the inhomogeneous noise case, the difference between the single- and multi-pixel approaches as well as marginal and spectral likelihoods is now much more evident and manifests itself as an overall shift and broadening, showing the macroscopic gain in adopting the expression in equation (9) throughout the analysis.

The two right panels of Figure 3 show analogous cases but for the Stokes Q parameter. The marginalised multi-pixel likelihoods (dotted lines) are, however, not shown as they do not fit the x -axis ranges. Note that unlike in the total intensity case the single and multi-pixel likelihoods differ even in the homogeneous noise case (middle panel, the solid and dashed lines). This reflects the fact that in the simulations used here there is more small-scale power in the dust polarization than in the total intensity, which leads to the loss of constraining power as a result of smoothing in the single pixel approach (see the discussion below). The effect is additionally enhanced if the inhomogeneous noise is present as shown in the right panel.




Figure 4. The likelihood constraint on the slope index for different noise levels and dust amplitudes. The **left panel** shows the spectral likelihoods calculated for a single pixel and three different noise levels: 50%, 100% and 200% of the nominal noise value corresponding to the low resolution total intensity pixel as described in Section 2. They are depicted with dot-dashed, solid and dashed lines respectively. In all three cases the same sky signal amplitudes for all components are assumed. In addition, the crosses correspond to the case with both the noise levels and the dust amplitude amplified by a factor 2. The **middle panel** shows the effects of changing the level of the dust signal level in the input data, while keeping all the other parameters constant. The displayed lines correspond to the cases with the dust amplitude rescaled by a factor of: 0.5 (dot-dashed), 1 (solid), 2 (dotted), and 5 (dashed). The **right panel** shows all the cases from the two previous panels but now rescaled by a factor proportional to a ratio of the dust amplitude and the RMS noise level and shifted to have a peak at the true value of $\beta = 1.65$, indicated in all panels with a vertical line.

In both cases the mismatch and the loss of precision are small, as expected given a relatively smooth variation of the dust foreground as adopted here.

The uncertainty involved in the determination of the spectral parameters depends on the noise level and the amplitude of the sky components. It is visualised for a single pixel in Figure 4, where the noise and dust amplitudes are rescaled in order to check their impact on the precision of the determination of β . In the left panel, the noise RMS and the actual noise level are varied as explained in the caption, while in the middle panel, only the dust amplitude is rescaled. In both cases the width of the likelihood changes accordingly. The right panel shows the same likelihoods shown in the two other panels but rescaled by a factor proportional to the ratio of the RMS of the dust and noise, and shifted to have the maximum at $\beta = 1.65$. In these examples the constraints on the slope parameter are found independent of the level of the CMB anisotropy. This is because the frequency scaling of the CMB component is assumed to be known a priori. The increase of the actual noise level in a given pixel, with the noise RMS kept constant, leads to an overall shift of the likelihood peak, thus changing (and thus potentially biasing) the best-fitting β value, but does not affect the likelihood width.

By using more pixels, the constraints can be usually improved. The rate at which that takes place, will however depend on the specific, and not known a priori, magnitude of the sky component in the newly added pixels, as well as, pixel instrumental noise levels, and therefore may not conform with the usual, $\sim 1/\sqrt{n_{pix}}$, expectation. Nevertheless, for the diffuse components, and in particular at high Galactic latitudes, as for example the dust foreground considered in the examples presented here, we find that once the number of included sky pixels is large enough the uncertainty of the slope determination falls roughly as expected, i.e., $\sim 1/\sqrt{n_{pix}}$. However, even in an extreme case when the newly added pixels happen to contain no information about a given component, the uncertainty, estimated using the multi-pixel approach, is guaranteed not to deteriorate. This may not be however the case with the single pixel approach, where including too many, for instance, dust-free pixels may suppress the noise power less than that of the dust and consequently increase the errors of the dust parameter determination. In a less extreme and more common case, the smoothing of the rapidly varying sky components whose amplitudes change across the low resolution pixel, will somewhat affect the precision of the spectral parameter determination, though will not bias the estimation result. The presence of the inhomogeneous noise on the scales smaller than the low resolution pixel will also have similar consequences (Figure 3).

Nevertheless, in the examples considered here, and concurred by those of Eriksen et al. (2006), the smoothing with a progressively larger window generally leads to an improvement of the constraints on the spectral parameters. This is because as the result of the smoothing the noise is typically suppressed more significantly than the dust component. In all the cases, the multi-pixel approach will produce nearly optimal constraints, independently of the actual sky distribution of the signals. It is also more flexible as it permits arbitrary pixel subsets for which identical spectral parameter values are assumed, and thus can better deal with the masked pixels and/or sky patch edge effects. Moreover, as discussed before, it also treats more optimally cases in which either the noise or the relative component content varies rapidly across the sky. For all these reasons the multi-pixel approach is therefore an approach of the choice here.

The evaluation of the spectral index is the first step toward the recovery of the actual maps of the components superposed in the data, with the latter step treated in the next Section. Before concluding this part, however, we draw a few more comments

on the numerical aspects of the likelihood treatment presented here, as well as on the relation between pixel size and constraint on the spectral index.

3.1.3 Numerical issues

The spectral likelihood equation (9) is more amenable to a numerical work than the marginalised expression in equation (10), since it does not require computation of a matrix determinant. In fact, it can be evaluated in a numerically efficient way, requiring at most on order of $\mathcal{O}(n_s^2 n_f^2 n_{pix}^2)$ floating point operations and often many fewer. For instance, the potentially time consuming, inversion of the matrix, $\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}$, involved in its computation can be sometimes simplified by exploiting its particular structure, for instance, block-diagonality, if the data noise \mathbf{N} is diagonal. Alternately and more generally, the computation of the inverse can be replaced by solving a linear set of equations and thus performed using fast, iterative techniques, e.g., preconditioned Conjugate Gradient (Golub and van Loan 1996; Shewchuk 1994), which involve only matrix-vector products. These can be efficiently done even if the noise in the sky map is correlated, a usual feature of the small-scale experiments. We emphasise that for many forthcoming CMB experiments even an explicit calculation of the full noise covariance matrix for the input maps \mathbf{N} may be already unfeasible. The computations involved in the evaluation of the spectral likelihood should be ultimately related to the time domain data as gathered by the CMB experiments. This can be done by exploiting techniques similar to those employed on the map-making stage of the CMB data analysis (e.g., Ashdown et al. 2007) and combining them with the formulae presented here. We leave the further development of the algorithm described here along these lines for future work, restricting hereafter ourselves to the cases when the noise correlation matrix \mathbf{N} is diagonal, and the calculations are done in the pixel domain. In such a case the numerical workload is brought down to merely $\mathcal{O}(n_{pix})$.

3.1.4 Pixel size and constraints on the spectral parameters

The rôle of a pixel and, in particular, its size may appear more difficult to grasp in the context discussed here than in, for example, the standard map making procedure. The pixel size determines the number of input data points and at the same time the number of independent degrees of freedom which are subsequently marginalised over. It may however, or may not, affect the number of spectral parameters which are to be constrained on this step.

Using too crude a pixelization may compromise the quality of the determination of the spectral parameters whenever the instrumental noise is inhomogeneous and the sky signal changes on the pixel scale (Figure 3). The pixel size should be therefore selected small enough to make such a variation negligible. This is in fact a usual criterion used in the production of the CMB maps, where the pixel size is usually fixed as some fraction of the resolution of the instrument. Also as in the map-making case, there is no loss of precision due overpixelization, if only we constrain the mixing matrix to be the same for all sub-pixels and thus we do not introduce any extra scaling parameters. That can be seen directly from equation (9) by noting that for all the sub-pixels of a given pixel, the sky signal is in fact the same, even though we introduce multiple parameters, i.e., sub-pixel amplitudes, to describe them. Clearly, in such circumstances overpixelizing does not lead to any gain and may affect a performance of the numerical calculations by unnecessarily exaggerating the size of the problem.

To conclude this discussion we note that for polarized map-making the pixel size may need to be fixed to alleviate the total intensity leakage to the Q and U Stokes parameter maps. This may result in a pixel size which is unnecessarily tiny for the purposes discussed here. In such a case, it may be useful to downgrade the pixelization during the component separation procedure, an operation which, in the formalism presented here, can be done by introducing a single set of sky component amplitudes for multiple pixels of the input maps.

3.2 Constraining the component amplitudes

We can calculate the maximum likelihood estimates of the component maps, given the maximum likelihood estimates of the spectral parameters derived in the previous section and expression for the peak of the data likelihood, equation (8). The two step algorithm, which incorporates these two procedures, is therefore a genuinely maximum likelihood approach, as we set to develop in this paper. We note that equation (8) is commonly employed in the component separation techniques and in particular, it also defines the second step in the approach of Eriksen et al. (2006). However, given that the spectral estimates which those methods obtained on the first step do not coincide with, or can not be straightforwardly related to, the maximum likelihood values, the recovered sky signal values correspond only to a likelihood ridge value (a maximum of the likelihood for given values of the spectral parameters) rather than a global maximum.

Equation (8) also corresponds to a standard general least square solution for the components, if the spectral parameters are perfectly known ahead of the time. In this case the solution error is quantified by the error correlation matrix $\hat{\mathbf{N}}$, given




Figure 5. **Left panel:** The likelihoods for the slope index, β , computed using equation (9) for the total intensity maps and 4 different resolutions: $n_{side} = 256, 128, 64, 16$, depicted with solid, dotted, dashed and short-long dashed lines respectively. The squares show the Gaussian approximation to each of the likelihoods with RMS computed using the curvature matrix formalism, equation (A5). The likelihoods were rescaled for the visualization reasons. **Middle panel:** The non-diagonality measure, κ_{pp} , as a function of a pixel number for the same four data renditions as in the left panel. The resolution decreases from bottom-up and the three bottom curves have been somewhat smoothed to facilitate a better comparison. **Right panel:** The covariance matrix for CMB and dust components recovered from the total intensity data with the pixel size $n_{side} = 64$. The two diagonal blocks correspond to CMB–CMB (bottom left), dust–dust (top right), while two off-diagonal blocks show their cross-correlation.

by,

$$\hat{N} \equiv (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1}. \quad (13)$$

Whenever the spectral parameters are not known perfectly, and need to be constrained from the data, we can get some insight into the structure of the error correlation matrix from the curvature matrix computed at the peak of the likelihood. The relevant formulae are listed in Appendix A. The results of their application to our test cases are visualised in Figure 5.

We first note that the Gaussian approximation to the spectral likelihood of the slope index with the RMS error derived from the curvature matrix formalism (equation A5) fares extremely well in reproducing the actual likelihood as calculated using equation (9). We show the comparison in the left panel of Figure 5. The four likelihoods shown there correspond to four different pixel sizes used to represent the same input data. The respective RMS uncertainties and peak positions are therefore very similar with small differences attributable to the presence of the small scale power as discussed earlier. In all cases the agreement between the approximate (squares) and exact computation is very good. This supports the expectation that the curvature matrix approach can be also successful in describing the ‘linear’ degrees of freedom such as the component maps.

The respective covariance matrix, $\mathbf{N}_{ss}^{\Lambda}$, (equation A4) for the recovered component maps is shown in the right panel of Figure 5. The four blocks shown in there correspond to the CMB–CMB (bottom left), dust–dust (top right) and CMB–dust and dust–CMB (bottom right and top left respectively) cases. The two most salient features of the matrix are the strong correlations between the amplitudes of different recovered components in any single sky pixel, as reflected by the prominent diagonals of the off-diagonal blocks displayed in the figure, and the presence of the off-diagonal correlations for each of the two components. The former feature is already encountered in the cases in which all the spectral parameters are known precisely and whenever the matrix $\mathbf{A}^t \mathbf{A}$ happens to be non-diagonal. It simply reflects the fact that the estimates are derived from the same input data. However, if the spectral parameters are also recovered from the data, and therefore, are known only with a limited precision, these may lead to a leakage between different sky component signals, resulting in the presence of the off-diagonal correlations between the errors with which the component amplitudes can be estimated. Both these are quantified by the expression for the covariance matrix of the sky component amplitudes in equation (A4).

The tartan-like pattern seen in the right panel of Figure 5 reflects the fact that the correction term in this equation has a form of a product of two low-rank matrices, with the brightest, vertical and horizontal stripes correspond to the pixels with the largest dust signal amplitude. This highlights the fact that the second ‘correction’ term on the rhs of equation (A4) depends on two factors: the error on the spectral parameters, and the amplitude of the actual sky signals. In the test cases considered here, we find the error correlation matrix, $\mathbf{N}_{ss}^{\Lambda}$, is strongly diagonal-dominated for the renditions of the data with small pixels but is becoming progressively less diagonal when the pixel-size increases. This is illustrated in the middle panel of Figure 5, where we show a non-diagonality measure, κ , (defined in Appendix B) for all considered pixels and different pixel resolutions, to quantify the importance of the off-diagonal elements of the matrix, $\mathbf{N}_{ss}^{\Lambda}$ and, therefore of the correction term in equation (A4). Whenever $\kappa \ll 1$, the off-diagonal elements are likely to be unimportant, and then the matrix \hat{N} may provide a good approximation for the overall pixel-pixel correlation pattern. If, however, $\kappa \lesssim 1$, the off-diagonal elements, and thus the correction term in equation (A4) may need to be taken into account. We see that for our multi-resolution rendition of the data the transition between the two regimes happens to be somewhere around $n_{side} \sim 64$. This suggests that though simple uncorrelated noise model (for each component separately), as implied in this case by the matrix \hat{N} , may be a sufficient

description of the component map error for high- ℓ part of the power spectrum, however, at its lower end – a more sophisticated treatment may be needed.

We point out that the increase in the importance of the non-diagonal correction term with the growing size of the pixels is not induced by the change in the error in the determination of the spectral index, as this remains roughly constant (see the left panel of Figure 5), but is due to the increasing signal-to-noise ratio of the component signals. We stress that the case discussed here is rather extreme, as it treats the total intensity signal with the noise level as expected for the experiments targeting the B mode polarization. In the cases with the lower signal-to-noise across the range of scales, the off-diagonal elements may well be less important even at the largest angular scales considered. Indeed, for the Q Stokes parameter maps considered in this paper the diagonal approximation seems to be breaking down only at the lowest resolution ($n_{side} = 16$). The significance of the off-diagonal elements will be, however, enhanced whenever more spectral parameters is considered, as this will unavoidably boost the spectral parameter determination uncertainty. This could be, for instance, the case when a multiple sets of the spectral parameters are introduced to account for their spatial variability, or whenever the calibration uncertainty of the input maps has to be included in the analysis. We discuss this latter case in more detail later.

We also emphasise that the question whether the covariance estimate based on the curvature matrix, equation (A4), provides a sufficiently good approximation to the actual covariance of the recovered sky components may in fact need to be re-evaluated application-by-application.

From the numerical point of view, as in the context of the marginalised likelihood for the spectral parameters, \mathcal{L}_{spec} , we can recast the algebraic operations involved in the calculations of the component amplitudes, equation (8) as matrix vector multiplications and which can be performed with help of the iterative solvers. The same is true of the correction terms in the expressions for the covariance matrices. The calculation of the matrix, $\tilde{\mathbf{N}}$, may however present a significant computational challenge. As our focus here is on the small-scale experiments this problem can be overcome with aid of the available supercomputers.

Summarizing, in this Section we have presented a likelihood based estimator of the spectral parameters, equation (9), and shown that it coincides with the maximum likelihood predictions derived from the maximization of the full likelihood. The new estimator does not require any pre-processing, handles optimally the inhomogenous noise cases and avoids problems with the observed sky edge and masked pixel effects. Combined with the standard determination of the component amplitudes, equation (8), it provides a two-step, maximum likelihood estimator of the sky components and the spectral parameters, which is computationally tractable. We also have introduced relevant curvature matrices to characterise statistical properties of the derived component maps and spectral parameters.

4 MAXIMUM LIKELIHOOD PROBLEM FOR COMPONENT SEPARATION – EXTENSIONS

In this Section we include extra parameters in the likelihood expression, which account for common instrumental characteristics, namely the presence of the undetermined offsets in the input data, and calibration errors.

4.1 Unconstrained single frequency map offsets

The single frequency maps produced from the data of the most CMB experiments have often arbitrary global offsets (e.g., Stompor et al. 2002), which for instance, could result from either instrumental $1/f$ -noise, differential scanning strategy or nearly sky synchronous, systematic effects (e.g. Johnson et al. 2007).

The data model for each map signal may need therefore to be more complex than it is assumed in equation (2),

$$\mathbf{d}_p = \mathbf{A}_p \mathbf{s}_p + \mathbf{o} + \mathbf{n}_p, \quad (14)$$

where, \mathbf{o} is a vector of $n_f \times n_s$ components, each of which describes an offset of a map at a given frequency and for each of the Stokes parameters. The likelihood function needs to be generalised accordingly,

$$-2 \ln \mathcal{L}_{data}(\mathbf{s}, \boldsymbol{\beta}, \mathbf{o}) = \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s} - \mathbf{o})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s} - \mathbf{o}). \quad (15)$$

Clearly, as we have just introduced $n_f \times n_s$ extra unknowns, the likelihood for a single pixel can not be uniquely maximised. However, as by definition the offset vector, \mathbf{o} , does not depend on a pixel, one could hope that the problem is becoming well-posed if a sufficient number of pixels is analyzed simultaneously. We will show that although such an expectation is not fully borne out, and a more subtle treatment of the arbitrary map offsets is needed, meaningful constraints can be derived for the spectra parameters and the component maps in such a case.

As mentioned above the offsets are completely spurious and consequently do not contain any cosmologically relevant information, we can therefore marginalise over them without any loss of information. Assuming the Gaussian priors for all

the offsets, with mean zero and a dispersion given by ξ_i , we arrive at

$$\begin{aligned} -2 \ln \mathcal{L}'_{data}(\mathbf{s}, \boldsymbol{\beta}) &\equiv -2 \ln \exp \int d\mathbf{o} \mathcal{L} \\ &= \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s})^t \left[\mathbf{N}^{-1} - (\mathbf{N}^{-1} \mathbf{U}) (\boldsymbol{\Xi}^{-1} + \mathbf{U}^t \mathbf{N}^{-1} \mathbf{U})^{-1} (\mathbf{N}^{-1} \mathbf{U})^t \right] (\mathbf{d} - \mathbf{A} \mathbf{s}). \end{aligned} \quad (16)$$

Here, \mathbf{U} is a ‘mixing matrix’ for the map offsets, i.e., a matrix of a rank $(n_f n_s n_{pix}) \times (n_f n_s)$, such as,

$$\mathbf{U}^t \equiv \left[\mathbf{I}_{(n_s n_f) \times (n_s n_f)}, \dots, \mathbf{I}_{(n_s n_f) \times (n_s n_f)} \right] \quad (17)$$

where $\mathbf{I}_{n \times n}$ is a square unit matrix of a rank n , and $\boldsymbol{\Xi} \equiv \text{diag}\{\xi_i\}$. Hereafter, we will refer to the matrix appearing in equation (16) as $\mathbf{M}_{\boldsymbol{\Xi}}^{-1}$. Note that its form could have been anticipated, and obtained as a result of the application of the Sherman-Morrison-Woodbury formula (Golub and van Loan 1996) to a matrix, $\mathbf{M}_{\boldsymbol{\Xi}}$,

$$\mathbf{M}_{\boldsymbol{\Xi}} \equiv \mathbf{N} + \mathbf{U} \boldsymbol{\Xi} \mathbf{U}^t. \quad (18)$$

The latter matrix is nothing else but the input map noise correlation matrix with an extra uncertainty due to the random offsets included (e.g., Stompor et al. 2002).

The likelihood in equation (16) has the standard (generalised) “ χ^2 ” form with the covariance given now by the matrix $\mathbf{M}_{\boldsymbol{\Xi}}$. Consequently, all the formulae derived in the previous sections could, in principle, be applied straightforwardly to our case at hand by replacing \mathbf{N}^{-1} by $\mathbf{M}_{\boldsymbol{\Xi}}^{-1}$. Nevertheless caution needs to be exercised whenever the matrix $\mathbf{M}_{\boldsymbol{\Xi}}^{-1}$ turns out to be singular. For instance, this is the case whenever no prior knowledge about the offsets is available, i.e., $\boldsymbol{\Xi}^{-1} = 0$. Readily, we then have (see equation 18),

$$\mathbf{M}^{-1} \mathbf{U} = 0, \quad (19)$$

and thus \mathbf{M}^{-1} is singular and has as many singular vectors as the arbitrary offsets. (Hereafter, \mathbf{M}^{-1} denotes the matrix $\mathbf{M}_{\boldsymbol{\Xi}}^{-1}$ in the no-prior limit: $\boldsymbol{\Xi}^{-1} \rightarrow 0$.)

Note that, although for a nearly full sky experiment one could hope to have a constraint on the relative offsets, at least in some cases, from considerations of the foreground-free, high Galactic latitude regions of the sky, for the small scale experiments such constraints will most likely not be available. The no prior case is therefore of significant, practical importance, and we will consider it and its consequences on the formalism developed here in the following sections.

4.1.1 The spectral parameters in presence of random offsets

From equation (16) we have in the limit of disappearing prior constraints on the offsets, i.e., $\boldsymbol{\Xi}^{-1} \rightarrow 0$,

$$-2 \ln \mathcal{L}'_{data}(\mathbf{s}, \boldsymbol{\beta}) = \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s})^t \mathbf{M}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}). \quad (20)$$

The singular vectors of the matrix \mathbf{M}^{-1} define flat directions of the likelihood, i.e., those for which the likelihood value remains the same. For any choice of the spectral parameters these flat directions correspond to the changes of the component maps, $\delta \mathbf{s}$, which, when projected to the map domain, $\mathbf{A} \delta \mathbf{s}$, result in vectors contained in the null space of \mathbf{M}^{-1} . The latter is spanned by the offset vectors, \mathbf{U}_i , (see equation 19). The presence of the flat directions in the likelihood however does not necessarily mean that no constraint can be set, first, on the spectral parameters and, later, on the component maps. We can avoid the unconstrained modes by restricting the allowed parameter space so that the only included solutions are those for which $(\mathbf{A}^t \mathbf{s})^t \mathbf{U} = 0$. We then take every such a solution to represent an entire class of solutions, which differ by some $\delta \mathbf{s}$, such that $\mathbf{A} \delta \mathbf{s}$ is contained within the subspace spanned by the offset vectors, \mathbf{U} , and as such, given the data model, are indistinguishable. We thus modify the likelihood in equation (20) by introducing a penalty function, which forces the likelihood to decay rapidly along the flat directions, thus effectively excluding undesired solutions. The relevant likelihood expression then reads,

$$\begin{aligned} -2 \ln \mathcal{L}''_{data}(\mathbf{s}, \boldsymbol{\beta}) &\equiv \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s})^t \mathbf{M}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) + \gamma^2 [(\mathbf{A} \mathbf{s})^t \mathbf{U}] [(\mathbf{A} \mathbf{s})^t \mathbf{U}]^t \\ &= \text{CONST} + (\mathbf{d} - \mathbf{A} \mathbf{s})^t \mathbf{M}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) + \gamma^2 \mathbf{s}^t (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t \mathbf{s}, \end{aligned} \quad (21)$$

where γ^2 is a positive and sufficiently large number. We can see that the above likelihood effectively restricts the allowed component maps, \mathbf{s} , to those which are orthogonal to the subspace spanned by $\mathbf{A}^t \mathbf{U}$. A similar approach to “weighing out” unwanted contributions to a likelihood has been proposed in other contexts, e.g., power spectrum estimations (Bond, Jaffe, & Knox 1998). The vectors $\mathbf{A}^t \mathbf{U}$ are not generally orthogonal, but they can always be replaced by their linear combinations which are, whenever it is required or convenient. That can be done using any of the standard orthogonalization procedures (e.g. Golub and van Loan 1996). These orthogonal vectors can be inserted in the formulae presented below, instead of the vectors, $\mathbf{A}^t \mathbf{U}$.




Figure 6. The one dimensional likelihoods for the slope index with the offsets of the input maps marginalised over (equations 22 and 25). The presented results are for the total intensity data. **Left panel:** The solid, dashed and dot-dashed lines correspond to the cases with $n_{pix} = 100, 500$ and 1000 high resolution pixels. **Center panel:** The same likelihoods as in the left panel, but this time the results have been rescaled by a factor of $\sqrt{n_{pix}}$ and recentered at the $\beta = 1.65$ value. This shows that the uncertainty on slope determination decays with a number of pixels as roughly $\sim 1/\sqrt{n_{pix}}$, once sufficiently many pixels are included. **Right panel:** A comparison of the two likelihoods computed with (dashed) and without (solid) marginalisation over the offsets. Both results are for $n_{pix} = 1000$ high resolution pixels and show the loss of precision incurred due to the marginalisation. The latter depends crucially on the amount of power at the offset scale as present in the input data, and it is consequently typically found much smaller in the applications to the (simulated) polarized data. The thin, vertical lines mark the true value of $\beta = 1.65$.

We can use equation (21) to derive the likelihood for the spectral parameters, which now reads,

$$-2 \ln \mathcal{L}''_{spec}(\beta) \equiv \text{CONST} - (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d})^t \left(\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t \right)^{-1} (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d}). \quad (22)$$

We note here that the matrix $\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A}$ may be singular given the singularity of \mathbf{M}^{-1} , and thus the second term, i.e., $\gamma^2 (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t$ has to correct for that if the likelihood in the above equation is to be well-defined. In fact, it is straightforward to show that the singular eigenvectors of $\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A}$ are indeed all contained in the subspace spanned by $\mathbf{A}^t \mathbf{U}$. This follows from the observation that if \mathbf{x} is a singular eigenvector of $\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A}$, i.e., $\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A} \mathbf{x} = 0$, and the mixing matrix, \mathbf{A} , is non-singular, i.e., $\mathbf{A} \mathbf{z} = 0 \iff \mathbf{z} = 0$, then $\mathbf{M}^{-1} (\mathbf{A} \mathbf{x}) = 0$, and $\mathbf{A} \mathbf{x}$ is indeed a singular vector of \mathbf{M}^{-1} and thus belongs to its null space spanned by $\mathbf{A}^t \mathbf{U}$. Consequently, the relevant matrix in equation (22) is invertible as long as γ^2 is non-zero. We also note that in general the singular eigenvectors of $\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A}$ do not have to span the entire subspace defined by $\mathbf{A}^t \mathbf{U}$. However, the latter is the case whenever a single set of spectral parameters is adopted for all considered pixels.

We thus conclude that the likelihood in equation (22) is indeed well-defined for any choice of the spectral parameters, what shows that, indeed, the degeneracy of the matrix, \mathbf{M}^{-1} , does not lead to any degeneracy in the parameter space of the spectral parameters. Nonetheless it does affect the final uncertainty of their determination. In the case of a single spectral parameter, slope index, the same for all pixels, this is shown in the left panel of Figure 6. In particular, if only one pixel is observed no limit can be set on the slope value (see also, Eriksen et al. 2008). For a fixed number of pixels, the loss of a precision with which the slope can be constrained is significant as seen in the right panel of Figure 6. The error, however, usually decreases quickly with a number of included pixels and follows a similar dependence on a number of the pixels as in a case with the offsets fixed (see the middle panel of Figure 6). Therefore, useful constraints on the spectral parameters can be usually derived even in the case of arbitrary offsets of the input maps.

From the numerical perspective, the likelihood calculations can be performed using an efficient iterative matrix-vector multipliers similar to those already discussed before. To facilitate that it is useful to rewrite the expression for the likelihood in equation (22). In order to do that, we first introduce matrices \mathbf{P} and \mathbf{R} defined as,

$$\mathbf{P} \equiv [\mathbf{A}^t \mathbf{N}^{-1} \mathbf{U}, \quad \mathbf{A}^t \mathbf{U}], \quad \mathbf{R} \equiv \begin{bmatrix} -(\mathbf{U}^t \mathbf{N}^{-1} \mathbf{U})^{-1} & \mathbf{0} \\ \mathbf{0} & \gamma^2 \mathbf{I}_{(n_s n_f) \times (n_s n_f)} \end{bmatrix}. \quad (23)$$

and note that,

$$\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t \equiv \mathbf{A}^t \mathbf{N}^{-1} \mathbf{A} + \mathbf{P} \mathbf{R} \mathbf{P}^t. \quad (24)$$

We then insert the above expression into equation (22) and use the Sherman-Morrison-Woodbury formula to obtain finally,

$$\begin{aligned} (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d})^t [\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t]^{-1} \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d} &= (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d})^t (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d} + \\ &- (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d})^t \left[(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right] \left[\mathbf{R}^{-1} + \mathbf{P}^t (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right]^{-1} \left[(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right]^t \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d}. \end{aligned} \quad (25)$$

This expression can be efficiently evaluated performing the matrix vector multiplications from outside-inwards. Moreover, the

likelihood calculation rewritten in this form can be freed of any ambiguity related to the choice of γ^2 , by simply setting γ^{-2} appearing in the expression for \mathbf{R}^{-1} directly to zero.

We also point out that if we allow for only one set of spectral parameters describing the frequency scalings of all considered pixels and assume homogeneous noise for all input maps, then $(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P}$ belongs to the subspace spanned by the vectors $\mathbf{A}^t \mathbf{U}$ and therefore the second term on the right hand side of equation (25) vanishes. In this case the calculations of the spectral likelihood with the offset amplitudes marginalised over is remarkably simple. This can be done applying formula from equation (9) but subtracting the offsets from the input maps prior to the computation.

4.1.2 Amplitude determination in the presence of the random offsets and fixed spectral parameters

With the spectral parameters fixed we can attempt to determine the amplitudes of the components, while accounting for the arbitrary offsets of the input maps. This can be done using the expression analogous to equation (8) but derived for the modified likelihood, equation (21). This leads to,

$$\mathbf{s} = \left(\mathbf{A}^t \mathbf{M}^{-1} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{U}) (\mathbf{A}^t \mathbf{U})^t \right)^{-1} \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d}, \quad (26)$$

where γ^2 is the same as before. Using the results derived in the previous Section (equation 25) we can rewrite the above equation as,

$$\begin{aligned} \mathbf{s} &= (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d} \\ &- (\mathbf{A}^t \mathbf{M}^{-1} \mathbf{d})^t \left[(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right] \left[\mathbf{R}_0^{-1} + \mathbf{P}^t (\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right]^{-1} \left[(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{P} \right]^t \mathbf{A}^t \mathbf{M}^{-1} \mathbf{d}, \end{aligned} \quad (27)$$

where the matrices \mathbf{P} and \mathbf{R} are defined in equation (23) and $\mathbf{R}_0^{-1} = \mathbf{R} (\gamma^{-2} = 0)$, and thus the above equation is strictly independent on a value of γ^2 .

As it is clear from our previous discussion the actual solution is not unique and any other, equivalent solution can be obtained by adding to equation (26) any vector from the subspace spanned by $\mathbf{A}^t \mathbf{U}$. These unconstrained modes in the recovered component maps depend on the mixing matrix as well as best-fitting spectral parameters found on the first step of the procedure. For example, in a case of the matrix \mathbf{A} being pixel-independent discussed already above they will be likewise pixel independent and will result in the total offset of all the derived component maps being unconstrained. For spatially-dependent frequency scalings the unconstrained modes will be more complex, potentially giving rise to some degree of arbitrariness in the spatial pattern of the recovered signals and therefore may have more pronounced consequences. The important fact is that the arbitrary modes can always be straightforwardly determined and therefore the relevant information passed to the next stages of the map processing such as power spectrum estimation, and the loss of precision due to their presence can be properly included in the estimated errors of the entire analysis chain (e.g., Bond, Jaffe, & Knox 1998).

It is interesting to note that the problem of the arbitrary offsets considered here is just a generalization of a simpler case involving combining multiple maps at the same frequency, which may completely or only partially overlap on the sky (Stompor et al. 2002). In such a case the constraint on relative offsets between any pair of the maps can be always derived if only the two maps overlap. However, their absolute offsets remain arbitrary.

As in Section 3.2 we can use the curvature matrix formalism to describe the covariance pattern of the recovered components. The relevant formulae are listed in Appendix A (equation A13). We stress that those equations do not incorporate the unconstrained modes, which need to be accounted for separately.

From the point of view of a numerical implementations, the same efficient techniques as discussed in the previous section can be applied in the computations required here.

The effects of the unknown offsets are also discussed in Eriksen et al. (2008). Those authors have introduced a prior to break the degeneracy between the sky parameters and the offsets in the data likelihood in equation (15) in order to estimate the offset values. This is in contrast with the approach proposed here, which marginalises over the offsets and then incorporates the additional uncertainty on the subsequent stages of the separation process.

4.2 Calibration errors

The overall calibration of the single frequency maps is usually known only with some non-zero precision which for the small scale experiments can be as large as 5 – 10% of the best-fitting value. It is therefore important to account for the calibration uncertainty in the likelihood considerations of the sort presented here.

We will denote hereafter the calibration factors at different frequencies and Stokes parameters, as ω_i . As offsets, these coefficients are assumed to be pixel-independent as they characterise the map as a whole. We can then write the initial likelihood as (cf. equation 5),

$$-2 \ln \mathcal{L}_{data}(\mathbf{s}, \boldsymbol{\omega}, \boldsymbol{\beta}) = \text{CONST} + (\mathbf{d} - \boldsymbol{\Omega} \mathbf{A} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \boldsymbol{\Omega} \mathbf{A} \mathbf{s}). \quad (28)$$

Here the noise correlation matrix \mathbf{N}^{-1} is assumed not to contain any correction due to the arbitrary offsets. Next we will consider the complete case. The calibration matrix, $\mathbf{\Omega}$, is then block-diagonal with each block, $\mathbf{\Omega}_p$, corresponding to a different sky pixel, being pixel independent, diagonal, and given by,

$$\mathbf{\Omega}_p \equiv \text{diag} \{ \omega_i \}_{i=0, \dots, n_f n_s - 1}. \quad (29)$$

Alternately, we also arrange all the calibration coefficients in a vector and refer to it as, ω .

In general, the two matrices, $\mathbf{\Omega}$ and \mathbf{A} , do not commute and therefore the calibration errors of the single frequency maps can not be straightforwardly incorporated as the calibration error of the components. However, we can always rescale all the components by some arbitrary factor at the same time rescaling the calibration matrix, $\mathbf{\Omega}$. This clearly gives rise to a presence of the degeneracy in the likelihood problem, which can be broken if the prior constraints are available. Consequently, in the following, we will assume that the calibration factors have been measured with some accuracy and their error distribution is well represented by Gaussian with a error matrix, $\mathbf{\Sigma}$. In the typical case, all ω_i are assumed to be measured independently with an RMS precision, $\sigma_{\omega_i}^2$, around some central value, $\bar{\omega}_i$. Then,

$$\mathbf{\Sigma} \equiv \text{diag} \left(\sigma_{\omega_i}^2 \right)_{i=0, n_f - 1}. \quad (30)$$

In a case when the calibration errors are correlated, e.g., due to the uncertainty on the flux of the common source used for this purpose, $\mathbf{\Sigma}$ may be non-diagonal. For example in a case when a single broad frequency source is used for the calibration of all channels, the relevant error matrix could read,

$$\mathbf{\Sigma} = \sigma_{source}^2 \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} \begin{bmatrix} 1 & \dots & 1 \end{bmatrix} + \text{diag} \left(\sigma_{\omega_i}^2 \right), \quad (31)$$

where the first term is due to the error in our knowledge of the brightness of the calibration source, and the second due to the instrumental noise for each frequency channel. In addition we will postulate that these (absolute) errors are Gaussian, and use $\bar{\mathbf{\Omega}}$ (or $\bar{\omega}$ for the vectors) to denote the central values of the calibration parameters. We also posit hereafter that $\mathbf{\Sigma}$ is invertible. The first term corresponds to a global calibration error and it will propagate without any changes all the way to the components amplitudes. Consequently, in the examples discussed in the following, we will discuss only the second term describing the uncorrelated, frequency channel specific errors. We emphasise though that the formalism presented here is fully general.

Here, we adopt an approach in which we treat the calibration parameters as part of the parametrization of the mixing matrix. In this case, the considerations from Section 3.1 are straightforwardly applicable. In particular, the spectral likelihood estimator, equation (9), reads in the present case as,

$$-2 \ln \mathcal{L}_{spec}(\omega, \beta) \equiv \text{CONST} - (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{N}^{-1} \mathbf{d})^t (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{N}^{-1} \mathbf{\Omega} \mathbf{A})^{-1} (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{N}^{-1} \mathbf{d}) + (\omega - \bar{\omega})^t \mathbf{\Sigma}^{-1} (\omega - \bar{\omega}), \quad (32)$$

where the last term is the prior constraints on the calibration. Note that the likelihood above is not degenerate only if some prior information is included. Using equation (32) we can find the maximum likelihood estimates of the spectral parameters and calibration coefficients and use those in the component amplitude estimation on the second step of the process. The relevant equation is identical to equation (8), but with the mixing matrix \mathbf{A} replaced by a product of the calibration and mixing matrices, $\mathbf{\Omega} \mathbf{A}$. We can also employ the curvature matrix approximation to estimate the extra uncertainty related to the calibration uncertainty.

4.3 Combining the two: unknown offsets plus calibration

We present here our final algorithm for the component separation problem. The algorithm looks for the maximum likelihood estimates of the spectral parameters, calibration factors, component maps, while allows for a marginalisation over the arbitrary input maps offsets. The likelihood maximization is done in two steps. First, we determine the spectral parameters and calibration factors. This is done by maximization (direct or via MCMC sampling) of the following spectral likelihood (this likelihood expression is a straightforward result of combining together equations (22) and (32)),

$$\begin{aligned} (\bullet) \quad -2 \ln \mathcal{L}_{spec}(\omega, \beta) = & - (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{M}^{-1} \mathbf{d})^t \left(\mathbf{A}^t \mathbf{\Omega}^t \mathbf{M}^{-1} \mathbf{\Omega} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{U}) (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{U})^t \right)^{-1} \mathbf{A}^t \mathbf{\Omega}^t \mathbf{M}^{-1} \mathbf{d} \\ & + (\omega - \bar{\omega})^t \mathbf{\Sigma}^{-1} (\omega - \bar{\omega}) + \text{CONST}. \end{aligned} \quad (33)$$

On the second step, we estimate the component amplitudes by solving,

$$(\bullet\bullet) \quad \mathbf{s} = \left(\mathbf{A}^t \mathbf{\Omega}^t \mathbf{M}^{-1} \mathbf{\Omega} \mathbf{A} + \gamma^2 (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{U}) (\mathbf{A}^t \mathbf{\Omega}^t \mathbf{U})^t \right)^{-1} \mathbf{A}^t \mathbf{\Omega}^t \mathbf{M}^{-1} \mathbf{d}. \quad (34)$$

(c.f., equation 26) The constant γ^2 is a sufficiently large, positive number, a value of which should be such that the matrix inversion required in the above equations is stable. The matrix, \mathbf{M}^{-1} , is defined in equation (16) (with $\mathbf{\Xi}^{-1} = 0$).




Figure 7. The constraints on the slope index β derived for four different sets of assumptions and total intensity and Stokes Q parameter (**left and middle panels**). These correspond to (from the most stringent to the most lax constraints): (1) the calibration parameters and map offsets are known perfectly, (2) the calibration parameters are known, the offsets are marginalised over, (3) the calibration parameters for two high frequency channels are known only with 2% precision, the remaining channel calibration is known, the offsets marginalised over; (4) the calibration parameters for all three channels are all known with 2% accuracy and the offsets are marginalised over. **Right panel:** the two-dimensional marginalised likelihood of the slope index and the calibration parameter for the lowest (150GHz) frequency channel calculated in the case (4).

The effects due to the calibration uncertainty in the cases with and without the offsets marginalisation, are illustrated in Figure 7. Clearly, its impact on the precision of the determination of the spectral parameters is significant, as demonstrated by the substantial broadening of the posterior likelihoods for the slope index in the left panels of Figure 7. This effect exceeds by a large factor the loss of precision due the unknown offsets. The loss of accuracy incurred in this case is due to the near degeneracy between the calibration factors and the slope index as displayed in the right panel of Figure 7 and which effectively is only broken by the prior constraints used for the calibration. We stress that the assumed here 2% calibration uncertainty is quite optimistic in the context of the current small-scale CMB experimentation.

We would like to minimise the effects due to the nuisance factors such as calibration and offset uncertainties. For that we would have to strive to use complete (or nearly so) maps as an input data vector \mathbf{d} in equation (33). That clearly can be done only if general mixing matrices \mathbf{A} , permitting multiple sets of spectral parameters are used and simultaneously estimated. As we have hinted at earlier, the formalism presented here is sufficiently general to accommodate such cases. A more important limitation here is more likely to arise due to the limitation in available computing resources. One may need therefore to develop strategies as to how to introduce the pixel subsets with different spectral parameter set and how large fraction of all pixels has to be used in the estimation process not to be significantly affected by our ignorance of the calibrations and offsets.

5 CONCLUSIONS

We have presented a parametric, maximum likelihood algorithm for separating the components of the sky maps produced by CMB experiments. The likelihood maximization is performed in two steps (c.f., Eriksen et al. 2006): on the first step, we perform a maximization of the spectral likelihood as given by formula equation (33); on the second step, we use the derived parameters to estimate the component maps via equation (34). We have shown that this procedure indeed leads to the maximum likelihood estimates derived directly from the full data likelihood (equation 5). We worked out example applications of the algorithm, considering a modeled portion of the sky containing a mixture of CMB and Galactic dust foreground, in three frequency channels at 150, 250, 410 GHz, in total intensity and polarization.

The likelihood approach, described here, is applicable as long as a suitable data model (e.g., equation (14)) and signal parametrization is available for each pixel. From the experience of the data analysis carried out for the past CMB experiments so far, as well as the current knowledge of the gross features of diffuse foregrounds, these two requirements are likely to be fulfilled. The approach is flexible and permits an inclusion of a variety of priors, in particular those constraining the allowed values of to the spectral parameters, as well a number of extensions accounting for the specificity of the input data. Any spatial templates for component amplitudes can be incorporated as additional data sets. Therefore, the approach promises to be useful for the analysis of the real data sets as those anticipated from satellite missions (see e.g., Eriksen et al. 2006) and in particular the balloon- and ground- based experiments. In fact we have shown how to amend the basic algorithm to account for the input map offset uncertainty and to deal with the calibration error. We have found that though both these effects may affect the precision of the spectral index recovery, it is the calibration uncertainty, which, for realistic prior constraints, is likely to be a dominant factor. The arbitrary offsets lead in addition to the presence of the unconstrained modes in the component maps. In a simple case of spatially independent mixing matrices these will only result in the free offsets of the




Figure 8. The recovered total intensity dust signal (**left panel**) calculated when allowing for the calibration uncertainty of 2% in all three frequency channels (i.e., the case (3) below). The **three following panels** show the maps of the recovered dust residual and correspond to three specific test cases in which it is assumed that: (1) the actual values of spectral index and calibration parameters are known a priori, (2) the calibration parameters are known, but the spectral index is determined from the data, (3) the calibration parameters are known only with a 2% error for each of the three frequency channels, and therefore their actual values, as well as that of the spectral index, are derived from the data. The RMS of the residual maps is found to be: $0.263\mu\text{K}_{\text{RJ}}$, $0.263\mu\text{K}_{\text{RJ}}$, and $0.260\mu\text{K}_{\text{RJ}}$, respectively. The slope index values used in these three cases are: $\beta = 1.65$, 1.651 and 1.685 .

recovered component maps. However, for more complex mixing matrices these may correspond to the more complicated modes in each of the component maps.

The examples of the end-to-end runs of the algorithm described here are shown in Figure 8. We show there a recovered total intensity map of the dust component and the maps of the residuals for three different sets of progressively more relaxed priors. In particular, we observe that in the rightmost map of the residuals the actual dust signal is lurking within the noise – an effect absent in the two other maps shown here. This is yet another demonstration of the loss of the precision of the spectral index determination due to assumed calibration uncertainties. We note that this effect is apparent only due to a very high signal-to-noise ratio for the total intensity maps considered here and would not be seen in the corresponding polarized maps with their lower sky signal and higher noise. We also note that RMS values of the map residuals are in good agreement with the values derived from the curvature matrix formalism, applied to the proposed here likelihood functions.

We note that it is possible to treat the method described here just as an algorithm for an efficient maximization of the full data likelihood, equation (5), avoiding therefore a need of interpreting the expression in equation (33) as a common likelihood. We have however shown in this paper that adopting the latter point of view is likely to be insightful and useful in the actual CMB data analysis practice, and does not affect the method results.

The implementation of the formalism presented in this paper has been pixel-based, however, we have pointed out that the algorithm can be straightforwardly and efficiently extended to operate on the time domain data streams, simplifying the task of the proper error estimate propagation between the map-making to the component separation stages.

In future works, we will specialise the present algorithm in the context of the forthcoming CMB experiments.

ACKNOWLEDGEMENTS

We would like to thank Jean-François Cardoso for a useful discussion on the ‘profile’ likelihoods. We also thank him, Jacques Delabrouille, Hans-Kristian Eriksen, Andrew Jaffe, and, the reviewer, Mark Ashdown for numerous comments on the manuscript.

Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005) package. We acknowledge the use of the CAMB (Lewis, Challinor & Lasenby 2000) and the CosmoMC (Lewis & Bridle 2002) packages. RS acknowledges support of the EC Marie Curie IR Grant (mirg-ct-2006-036614) and CB that of the NASA LTSA Grant NNG04CG90G. This research has been supported in part by the ASI Planck LFI Activity of Phase E2 contract.

REFERENCES

- Ade, P. et al., 2007, submitted to ApJ.
- Ashdown, M.A.J. et al., 2007, A&A, 467, 761
- Barndorff-Nielsen, O.E., & Cox, D.R., 1994, *Inference and Asymptotic*, Chapman & Hall, London
- Bennett, C.L., et al., 1992, ApJ, 396, L7
- Benoit A. et al., 2004, A&A, 424, 571
- Bischoff, C. et al., 2008, submitted to ApJ.

- Bonaldi, A., Bedini, L., Salerno, E., Baccigalupi, C., & de Zotti G., 2006 MNRAS, 373, 263
 Bond, J.R., Jaffe, A.H. & Knox, L., 1998, PRD, 57, 2117
 Brandt, W.N., Lawrence, C.R., Readhead, A.C.S., Pakianathan, J.N., & Fiola, T.M., 1994, ApJ, 424, 1
 Delabrouille, J., Cardoso, J.-F., & Patanchon, G., 2003, MNRAS, 346, 1089
 Dunkley, J., et al., 2008, submitted to ApJS
 Eriksen, H.K. et al., 2006, ApJ, 641, 665
 Eriksen, H.K., Jewell, J.B., Dickinson, C., Banday, A.J., Górski, K.M., & Lawrence, C.R., 2008, ApJ, 676, 10
 Finkbeiner, D.P., Davis, M., & Schlegel, D.J., 1999 ApJ, 524, 867
 Gold, P., et al., 2008, submitted to ApJS
 Golub, G. H., and van Loan, C. F., 1996, Matrix Computations, The Johns Hopkins University Press, 3rd edition
 Górski, K.M., Hivon, E., Banday, A.J., Wandelt, B.D., Hansen, F.K., Reinecke, M., & Bartelmann, M., 2005, ApJ, 622, 759
 Hansen, F.K., Banday, A.J., Eriksen, H.K., Górski, K.M., & Lilje, P.B., 2006, ApJ, 648, 784
 Jaffe, A.H., et al., 2004, ApJ, 615, 55
 Johnson, B., et al., 2007, 665, 42
 Kamionkowski, M., Kosowsky, A., & Stebbins, A., 1997, PRD, 55, 7368
 Kogut, A. et al., 2007, ApJ, 665, 355
 Leach, S.M. et al., 2008, submitted to A&A
 Lewis, A., & Bridle, S., PRD, 66, 103511, 2002
 Lewis, A., and Challinor, A., Phys. Rep., 429, 1, 2006
 Lewis, A., Challinor, A., & Lasenby, A., ApJ, 538, 473, 2000
 Maino, D., Donzelli, S., Banday, A.J., Stivoli, F., & Baccigalupi, C., 2007, MNRAS, 374, 1207
 Masi, S. et al., 2006, ApJ, 458, 687
 Montroy, T.E. et al., 2006, ApJ, 647, 813
 Oxley, P. et al., 2004, in Barnes, W.L. & Butler, J.J., eds, Proc. of the SPIE, Earth Observing Systems IX, 5543, 320
 Page, L. et al., 2007, ApJS, 170, 335
 Ponthieu, N. et al., 2005, A&A, 444, 327
 Prunet, S., Sethi, S. K., Bouchet, F. R., Miville-Deschênes, M.-A., 1998, A&A, 339, 187
 Reichardt, C.L. et al., 2008, submitted to ApJ
 Shewchuk, J.R., 1994, preprint, <http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.ps>
 Schlegel, D.J., Finkbeiner, D.P., & Davis, M., 1998 ApJ, 500, 525
 Spergel, D.N. et al., 2007, ApJS, 170, 377
 Stolyarov, V., Hobson, M.P., Lasenby, A.N., & Barreiro, R.B., 2005, MNRAS, 357, 145
 Stompor, R. et al., 2002, PRD, 65, 022003
 Zaldarriaga, M. and Seljak, U., 1997, PRD, 55, 1830
 Zaldarriaga, M. and Seljak, U., 1998, PRD, 58, 023003
 Zaldarriaga, M., 2001, PRD, 64, 103001

APPENDIX A: CURVATURE MATRICES FOR THE CONSIDERED LIKELIHOODS

We list here some useful curvature matrices for the data likelihoods as considered in this paper.

We start from the case of the full data likelihood, equation (5), applicable to the most basic data model with both calibrations and offsets known precisely. The relevant expression for the curvature matrix, $\mathbf{\Lambda}_{p,q} \equiv -\partial^2 \mathcal{L} / \partial p \partial q$, is given by,

$$-\frac{\partial^2 \ln \mathcal{L}_{data}}{\partial \mathbf{s}^2} = \mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}, \quad (\text{A1})$$

$$-\frac{\partial^2 \ln \mathcal{L}_{data}}{\partial \beta_i \partial \beta_j} = (\mathbf{A}_{,\beta_i} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{A}_{,\beta_j} \mathbf{s}) - (\mathbf{A}_{,\beta_i \beta_j} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}), \quad (\text{A2})$$

$$-\frac{\partial^2 \ln \mathcal{L}_{data}}{\partial \mathbf{s} \partial \beta_i} = \mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta_i} \mathbf{s} - \mathbf{A}_{,\beta_i}^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}). \quad (\text{A3})$$

We will interpret the inverse of the curvature matrix, computed at the peak of the likelihood, as a good approximation to the covariance matrix for the parameters estimated via maximization of the likelihood. Specifically, in this approximation, the $\mathbf{s} - \mathbf{s}$ diagonal block of the inverse curvature matrix, $\mathbf{N}_{\mathbf{s} \mathbf{s}}^{\mathbf{A}}$, describes the correlation pattern of the recovered sky components, while the $\beta - \beta$ block, $\mathbf{N}_{\beta \beta}^{\mathbf{A}}$, provides estimates of the errors, and their correlations, of the recovery of the spectral parameters. Using equations (A1-A3), and inverting the curvature matrix by partition, we can write explicitly relevant expressions for

both these matrix blocks. We obtain,

$$\tilde{\mathbf{N}}_{\mathbf{s}\mathbf{s}}^{\Lambda} = \hat{\mathbf{N}} + \left[\hat{\mathbf{N}} \left(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta} \mathbf{s} - \mathbf{A}_{,\beta}^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) \right) \right] \tilde{\mathbf{N}}_{\beta\beta}^{\Lambda} \left[\hat{\mathbf{N}} \left(\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta} \mathbf{s} - \mathbf{A}_{,\beta}^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) \right) \right]^t \quad (\text{A4})$$

and

$$\begin{aligned} \tilde{\mathbf{N}}_{\beta\beta}^{\Lambda} &= \left\{ (\mathbf{A}_{,\beta_i} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{A}_{,\beta_j} \mathbf{s}) - (\mathbf{A}_{,\beta_i \beta_j} \mathbf{s})^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) \right. \\ &\quad \left. - \left[\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta} \mathbf{s} - \mathbf{A}_{,\beta}^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) \right]^t \hat{\mathbf{N}} \left[\mathbf{A}^t \mathbf{N}^{-1} \mathbf{A}_{,\beta} \mathbf{s} - \mathbf{A}_{,\beta}^t \mathbf{N}^{-1} (\mathbf{d} - \mathbf{A} \mathbf{s}) \right] \right\}^{-1}. \end{aligned} \quad (\text{A5})$$

Here, $\hat{\mathbf{N}}$ is a ‘standard’ uncertainty estimate neglecting the error due to the spectral parameters defined in equation (13) and we treat all the matrices $\mathbf{A}_{,\beta}$ and $\mathbf{A}_{,\beta\beta}$ as respectively 3 and 4 dimensional with the extra dimensions having a size given by a number of spectral parameters to be determined. The matrix elements are evaluated at the maximum likelihood values of the parameters. We point out that the correction terms included in equations (A4) and (A5) involve the ‘map-making’ type calculations (which may need to be performed on order of n_{β} ($\ll n_{pix}$) times) and thus can be efficiently computed using the relevant iterative map-making solvers.

Similar expressions for the curvature matrix can be straightforwardly derived also in the case with calibration errors included (Section 4.2). In fact, if the calibration coefficients, $\boldsymbol{\omega}$, are considered a subset of all the spectral parameters $\boldsymbol{\beta}$, the curvature matrix, $\boldsymbol{\Lambda}'$, obtained in this case mostly coincides with the matrix $\boldsymbol{\Lambda}$, as discussed above, but an extra term, which needs to be added to its $\boldsymbol{\omega} - \boldsymbol{\omega}$ block to account for the calibration prior, equation (28), i.e.,

$$\boldsymbol{\Lambda}'_{\boldsymbol{\omega}\boldsymbol{\omega}} = \boldsymbol{\Lambda}_{\boldsymbol{\omega}\boldsymbol{\omega}} + \boldsymbol{\Sigma}^{-1}, \quad (\text{A6})$$

$$\boldsymbol{\Lambda}'_{i\,j} = \boldsymbol{\Lambda}_{i\,j}, \quad \text{otherwise.} \quad (\text{A7})$$

The curvature matrix, $\boldsymbol{\Lambda}''$, for the case with the input map offsets marginalised over can be derived in a similar fashion, but starting off from equation (21). First, we get,

$$-\frac{\partial^2 \ln \mathcal{L}''_{data}}{\partial \mathbf{s}^2} = \mathbf{A}^t (\mathbf{M}^{-1} + \gamma^2 \mathbf{U} \mathbf{U}^t) \mathbf{A}, \quad (\text{A8})$$

$$-\frac{\partial^2 \ln \mathcal{L}''_{data}}{\partial \beta_i \partial \beta_j} = (\mathbf{A}_{,\beta_i} \mathbf{s})^t (\mathbf{M}^{-1} + \gamma^2 \mathbf{U} \mathbf{U}^t) (\mathbf{A}_{,\beta_j} \mathbf{s}) + (\mathbf{A}_{,\beta_i \beta_j} \mathbf{s})^t \mathbf{M}^{-1} (\mathbf{A} \mathbf{s} - \mathbf{d}) \quad (\text{A9})$$

$$-\frac{\partial^2 \ln \mathcal{L}''_{data}}{\partial \mathbf{s} \partial \beta_i} = \mathbf{A}_{,\beta_i}^t \mathbf{M}^{-1} (\mathbf{A} \mathbf{s} - \mathbf{d}) + (\mathbf{A}_{,\beta_i} \mathbf{s})^t (\mathbf{M}^{-1} + \gamma^2 \mathbf{U} \mathbf{U}^t) \mathbf{A}, \quad (\text{A10})$$

where we have made an explicit use of the fact that $\mathbf{U}^t \mathbf{A} \mathbf{s} = 0$. Then we represent the curvature matrix as,

$$\begin{aligned} \boldsymbol{\Lambda}'' &= \boldsymbol{\Lambda}''_0 + [\mathbf{A}, \mathbf{A}_{,\beta} \mathbf{s}]^t (\mathbf{M}^{-1} + \gamma^2 \mathbf{U} \mathbf{U}^t) [\mathbf{A}, \mathbf{A}_{,\beta} \mathbf{s}] \\ &= \boldsymbol{\Lambda}''_0 + \mathbf{P} \mathbf{R} \mathbf{P}^t, \end{aligned} \quad (\text{A11})$$

where we have used the notation from equation (23) with the mixing matrix, \mathbf{A} , replaced by $[\mathbf{A}, \mathbf{A}_{,\beta} \mathbf{s}]$ and defined,

$$\boldsymbol{\Lambda}''_0 \equiv [\mathbf{A}, \mathbf{A}_{,\beta} \mathbf{s}]^t \mathbf{N}^{-1} [\mathbf{A}, \mathbf{A}_{,\beta} \mathbf{s}] + \begin{bmatrix} 0 & \mathbf{A}_{,\beta}^t \mathbf{M}^{-1} (\mathbf{A} \mathbf{s} - \mathbf{d}) \\ [\mathbf{A}_{,\beta}^t \mathbf{M}^{-1} (\mathbf{A} \mathbf{s} - \mathbf{d})]^t & (\mathbf{A}_{,\beta\beta} \mathbf{s})^t \mathbf{M}^{-1} (\mathbf{A} \mathbf{s} - \mathbf{d}) \end{bmatrix}. \quad (\text{A12})$$

Using similar arguments as those already used in the derivation of equation (25), we first note that the dependence on the γ^2 factor disappears in the final result for the covariance, which can be written down as,

$$\tilde{\mathbf{N}}^{\Lambda''} = \boldsymbol{\Lambda}''_0^{-1} - \left(\boldsymbol{\Lambda}''_0^{-1} \mathbf{P} \right) \left[\mathbf{R}_0^{-1} + \mathbf{P}^t \boldsymbol{\Lambda}''_0^{-1} \mathbf{P} \right] \left(\boldsymbol{\Lambda}''_0^{-1} \mathbf{P} \right)^t. \quad (\text{A13})$$

The inverse of the matrix $\boldsymbol{\Lambda}_0$ can be computed by partition as before, while the matrix products such as $\left(\boldsymbol{\Lambda}_0''^{-1} \mathbf{P} \right)$ can be calculated efficiently using the iterative linear equation solvers. We also note that the complete estimate of the uncertainty of the recovered component maps needs to take into account the unconstrained modes as defined by the subspace of the solutions span by the vectors $\mathbf{A}^t \mathbf{U}$.

APPENDIX B: MATRIX NON-DIAGONALITY

We define here a simple measure of non-diagonality of a square, invertible, diagonal-dominated matrix, \mathbf{B} . We first split the matrix, \mathbf{B} , into diagonal, \mathbf{B}_{diag} and non-diagonal part, \mathbf{B}_{off} , i.e.,

$$\mathbf{B} \simeq \mathbf{B}_{diag} + \mathbf{B}_{off} = \mathbf{B}_{diag} (\mathbf{1} + \mathbf{B}_{diag}^{-1} \mathbf{B}_{off}). \quad (\text{B1})$$

If the off-diagonal part is indeed small enough than the inverse of the matrix \mathbf{B} should be well described by,

$$\mathbf{B}^{-1} \simeq (\mathbf{1} - \mathbf{B}_{diag}^{-1} \mathbf{B}_{off}) \mathbf{B}_{diag}^{-1}. \quad (\text{B2})$$

This will be true whenever,

$$(\mathbf{1} - \mathbf{B}_{diag}^{-1} \mathbf{B}_{off}) \mathbf{B}_{diag}^{-1} \mathbf{B}_{diag} (\mathbf{1} + \mathbf{B}_{diag}^{-1} \mathbf{B}_{off}) = \mathbf{1} - [\mathbf{B}_{diag}^{-1} \mathbf{B}_{off}] [\mathbf{B}_{diag}^{-1} \mathbf{B}_{off}] \equiv \mathbf{1} - \boldsymbol{\kappa} \simeq \mathbf{1}, \quad (\text{B3})$$

where the matrix, $\boldsymbol{\kappa}$ introduced above, is our measure. We posit that the off-diagonal elements of the matrix \mathbf{B} are negligible, whenever, $\boldsymbol{\kappa}_{ij} \simeq 0$, with a precision as required by a problem at hand.