Maximum likelihood, parametric component separation and CMB B-mode detection in suborbital experiments

F. Stivoli, J. Grain, S. M. Leach, M. Tristram, C. Baccigalupi, R. Stompor

Maximum likelihood, parametric component separation and CMB B-mode detection in suborbital experiments

F. Stivoli,[1][★] J. Grain,[2] S. M. Leach,[3] M. Tristram,[4] C. Baccigalupi[3] and R. Stompor[5]

[1]*INRIA, Laboratoire de Recherche en Informatique, Université Paris-Sud 11, Bâtiment 490, 91405 Orsay Cedex, France*
[2]*CNRS, Institut d'Astrophysique Spatiale, Université Paris-Sud 11, Bâtiments 120-121, 91405 Orsay Cedex, France*
[3]*SISSA/ISAS, Astrophysics Sector, Via Beirut, 2-4, and INFN, Sezione di Trieste, Via Valerio 2, I-34151 Trieste, Italy*
[4]*CNRS, Laboratoire de l'Accélérateur Linéaire, Université Paris-Sud 11, Bâtiment 200, 91898 Orsay Cedex, France*
[5]*CNRS, Laboratoire Astroparticule & Cosmologie, 10 rue A. Domon et L. Duquet, F-75205 Paris Cedex 13, France*

ABSTRACT

We investigate the performance of the parametric maximum likelihood component separation method in the context of the cosmic microwave background (CMB) B-mode signal detection and its characterization by small-scale CMB suborbital experiments. We consider high-resolution (FWHM $= 8'$) balloon-borne and ground-based observatories mapping low dust-contrast sky areas of 400 and 1000 square degrees, in three frequency channels, 150, 250, 410 GHz, and 90, 150, 220 GHz, with sensitivity of order 1 to 10 μK per beam-size pixel. These are chosen to be representative of some of the proposed, next-generation, bolometric experiments. We study the residual foreground contributions left in the recovered CMB maps in the pixel and harmonic domain and discuss their impact on a determination of the tensor-to-scalar ratio, r. In particular, we find that the residuals derived from the simulated data of the considered balloon-borne observatories are sufficiently low not to be relevant for the B-mode science. However, the ground-based observatories are in need of some external information to permit satisfactory cleaning. We find that if such information is indeed available in the latter case, both the ground-based and balloon-borne experiments can detect the values of r as low as ∼0.04 at 95 per cent confidence level. The contribution of the foreground residuals to these limits is found to be then subdominant and these are driven by the statistical uncertainty due to CMB, including E-to-B leakage, and noise. We emphasize that reaching such levels will require a sufficient control of the level of systematic effects present in the data.

Key words: cosmic background radiation – diffuse radiation.

1 INTRODUCTION

Astrophysical foregrounds are commonly recognized as one of the major obstacles on the way to first detecting and later exploiting the scientific potential of the cosmic microwave background (CMB) polarization signal. This is particularly the case with the so-called B-mode polarization (Zaldarriaga & Seljak 1997) due to its minute amplitude as compared to the foreground contributions as well as CMB total intensity and E-mode polarization signals. In fact current foreground models (Page et al. 2007) generally indicate that the foreground B-mode signal may be comparable or exceeding the CMB signal by a factor of a few in a broad range of angular scales even in the cleanest available sky areas. Some kind of foreground cleaning or separation procedure will therefore be necessary and its impact on the final 'cleaned' map of the presumed CMB sky needs to be understood and properly taken into account in its subsequent studies. Developing such an understanding is also already of importance for the designing and optimization of the future CMB experiments.

This has been recognized for some time and a number of studies have been performed and published, and which have treated the problem on different levels of generality and detail. The major challenge here is twofold. First, there is no general recipe for propagating errors incurred during the component separation step, i.e. for including both the statistical uncertainty and foreground residual uncertainty. Secondly, there is no easily calculable metric measuring the impact of the component separation on the B-mode measurement, as both the power spectrum and the tensor-to-scalar ratio, r, require a proper evaluation of the E-to-B leakage (Bunn et al. 2003).

Tucci et al. (2005) have performed a Fisher analysis of the problem treating the foreground residuals as an additional source of noise, and then estimated the expected limits on r. In their approach

[★]E-mail: federico.stivoli@gmail.com

the starting point was a single foreground contaminated science channel and a noisy foreground template channel from which the level of foreground residual was estimated. Although this allows to avoid specifying in great detail a foreground-cleaning technique, no direct connection exists between the noise values they assume and the properties of any specific experiment. They have also neglected the impact of the E-to-B leakage. A similar approach has been followed by Verde, Peiris & Jimenez (2006), who have attempted to link their Fisher matrix considerations to specific, fiducial, multifrequency data sets. The simplified error propagation they have adopted implicitly bypasses any realistic component separation approach, and so fails to include properly its effects on their final results. They also neglect the presence of the E-to-B leakage. Amarie, Hirata & Seljak (2005) performed a Fisher analysis as well, but used specific parameters anchored in those of the multifrequency data set assumed. This last work together with Amblard, Cooray & Kaplinghat (2007) and Betoule et al. (2009) come the closest in the spirit to what we discuss in this paper, although neither of the latter two works includes an actual power spectrum estimator accounting for the leakage, which is justified at least in part by their focus on full-sky observations. Stivoli et al. (2006) studied an application of an independent component analysis based approach to the component separation of polarized partial-sky maps, resorting to the cleaned map 'pseudo-spectra' as a basis for a *qualitative* assessment of its performance and relevance for the B-mode work. Dunkley et al. (2009) presented a review of most of those earlier approaches, including those incorporating a parametric approach similar to the one considered in this work, and presented their applications in the context of a potential future CMB B-mode satellite mission.

The approach we propose here is more specific. We focus on a particular component separation method and power spectrum estimation approach, which we then use to investigate the impact of the foreground separation on the CMB B-mode detection and characterization. The component separation method is a maximum likelihood (ML) parametric approach (Eriksen et al. 2006) in a two-step implementation of Stompor et al. (2009). The power spectrum estimator is a 'pure' pseudo-spectrum approach introduced by Smith (2006) (see also Smith & Zaldarriaga 2007) and elaborated on by Grain, Tristram & Stompor (2009). Strictly speaking, our results will therefore be specific to these two choices. However, given that these two methods are working, promising algorithms to be implemented in the data analysis pipelines of current and future CMB experiments, the results should be of practical relevance for many efforts currently going on in the field.

We also note that whenever the frequency-scaling laws can be assumed to be nearly perfectly known, as in one of the cases we study, and in particular in the small sky, and thus potentially statistics-starved limit, the parametric ML method would likely become a method of choice, potentially supplemented by some priors, e.g. spatial templates for all or some of the components (Efstathiou, Gratton & Paci 2009). The results derived here can therefore be regarded as representative and realistic expectations for the performance of classes of the future experiments we consider. Moreover, part of the analysis presented here can be straightforwardly applied to any component separation method in which foreground spectral and amplitude parameters are estimated in separate steps.

Our focus in this work is on suborbital experiments. Those have a potential advantage of selecting the cleanest sky areas, but suffer due to the cut-sky effects. They also usually have a limited number of frequency channels with which to observe the sky. We consider two kinds of experiments: those with an access to the high frequencies

($\gtrsim 250$ GHz) referred to as balloon-borne, and those with access limited to frequencies lower than 250 GHz, referred to as ground-based. We will also consider some combination and extensions of these two cases. We then apply our proposed analysis chain to the simulated data for different foreground case studies, allowing for different levels of mismatch between the assumptions made on the analysis and simulation stages, in order to evaluate the impact of the component separation residuals first on the recovered B-mode power spectrum and later on the value of an r which can be derived from such data.

The paper is organized as follows. In Sections 2 and 3, we first provide brief descriptions of the specific data analysis techniques and their implementation, used throughout this paper. In Section 4 we describe our simulated sky model, and in Section 5 we define the experimental characteristics and a set of foreground case studies. Our results are presented in Section 6, and their analysis, concerning the residuals and their impact on the cosmological B-mode detection, is given in Sections 7 and 8, respectively.

2 PARAMETRIC COMPONENT SEPARATION METHOD

In this section we briefly outline the parametric component separation algorithm proposed by Stompor et al. (2009). The multifrequency sky signal is modelled as

$$d_{\mathrm{p}} = \mathbf{A}_{\mathrm{p}} \, s_{\mathrm{p}} + n_{\mathrm{p}}, \tag{1}$$

where d_{p} is a vector containing the data from N_{freq} frequencies assumed to share a common angular resolution; s_{p} is a vector of N_{comp} signal amplitudes to be estimated; $\mathbf{A}_{\mathrm{p}} \equiv \mathbf{A}_{\mathrm{p}}(\boldsymbol{\beta})$ is a component 'mixing' or frequency-scaling matrix with a total of N_{spec} free 'spectral parameters' $\boldsymbol{\beta}$ also to be estimated; and n_{p} is the noise at each pixel p. We can write down a likelihood for the data of the form

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

where $\mathbf{N}$ is the noise covariance matrix of the data and we have now dropped the pixel index p. This likelihood reaches its maximum for the values of s and $\boldsymbol{\beta}$ fulfilling the relations

$$-\left(\mathbf{A}_{,\beta} \, s\right)^t \, \mathbf{N}^{-1} \, (d - \mathbf{A} \, s) = 0, \tag{3}$$

$$s = \left(\mathbf{A}^t \, \mathbf{N}^{-1} \, \mathbf{A}\right)^{-1} \, \mathbf{A}^t \, \mathbf{N}^{-1} \, d, \tag{4}$$

where $_{,\beta}$ denotes a partial derivative with respect to β_i. Under the assumption that the spectral parameters are the same for a collection of pixels, corresponding to the physical assumption that the spectral parameters vary more slowly in space than the signal amplitudes, it is possible to substitute the generalized least-squares solution equation (4) into the likelihood equation (2), thereby eliminating the sky signals s, in order to obtain a spectral index likelihood given by

$$-2 \ln \mathcal{L}_{\mathrm{spec}} \, (\boldsymbol{\beta}) = \mathrm{CONST}$$
$$- (\mathbf{A}^t \, \mathbf{N}^{-1} \, d)^t \, (\mathbf{A}^t \, \mathbf{N}^{-1} \, \mathbf{A})^{-1} (\mathbf{A}^t \, \mathbf{N}^{-1} \, d). \tag{5}$$

The spectral parameters that minimize equation (5) can be found using numerical techniques, and then substituted into equation (4) in order to find the corresponding signal amplitudes pixel by pixel. Finally, the noise covariance matrix, describing the properties of the noise contained in the data d, is propagated to the component estimates s:

$$\mathbf{N}_s \equiv (\mathbf{A}^t \, \mathbf{N}^{-1} \, \mathbf{A})^{-1}. \tag{6}$$

MIRAMARE,[1] which is our implementation of this two-step algorithm, uses codes from the COSMOMC package (Lewis & Bridle 2002) in order to perform an initial conjugate gradient descent to the minimum of the spectral index likelihood equation (5). This is followed by an estimation of the curvature of the spectral index likelihood on a regular grid, which then forms the basis of the 'proposal function' for drawing spectral index samples using the Markov chain Monte Carlo (MCMC) technique. Once the spectral parameters have been determined from an analysis of the MCMC samples, these values are substituted into equations (4) and (6) in order to obtain the signal amplitudes and their covariance.

We can check our assumption of the constancy of the spectral parameters and the overall 'goodness of fit' by evaluating the log-likelihood equation (2) at the maximum likelihood value, and comparing it with the number of degrees of freedom given by

$$N_{\text{dof}} = N_{\text{pix}} \times (N_{\text{freq}} - N_{\text{comp}}). \tag{7}$$

A poor fit, for instance, due to either the wrong parametrization assumed for the components present in the data or the spatially variability of the parameters, will be accompanied by an excessive log-likelihood. Other goodness-of-fit tests are also possible and will be discussed elsewhere.

The ML formalism allows us to straightforwardly incorporate the uncertainty due to errors on the calibration of the input maps. This can be done by replacing the mixing matrix, **A**, by a product of a diagonal matrix **C**, representing the calibration for each of the single channel maps, and the foreground mixing matrix, **A**. Following Stompor et al. (2009), we denote the diagonal elements of **C** as, ω_i, so,

$$\mathbf{C} = \text{diag}\,(\omega_i)_{i=0,\dots,n_f-1}\,. \tag{8}$$

Such a problem is clearly degenerate if no external constraints are imposed on the calibration parameters. In a case of actual observations, prior information on the calibration uncertainty is in fact usually available and can be often expressed as a Gaussian error centred around some most likely value. Mathematically, it just corresponds to multiplying the likelihood in equation (5) by the relevant priors.

We note that if the absolute calibration of the resulting component maps is not required, one could in principle reduce the number of calibration parameters by one by subsuming one of the ω_i, say that of the very first channel, ω_0, factors into the overall normalization of the sought-after component maps. This is the approach that we will use later in this paper, always assuming that the uncertainty of the relative calibration of the higher-frequency channels with respect to the lowest-frequency channel can be sufficiently well described as a Gaussian random variable with a known width.

3 POLARIZED POWER SPECTRUM ESTIMATION

In this section we briefly describe the approach that we use to measure the E- and B-mode power spectra which is based on the 'pure' pseudo-C_ℓ's method proposed in Smith (2006). Further details of the implementation and tests on simulations can be found in Smith (2006), Smith & Zaldarriaga (2007) and in Grain et al. (2009). The latter work also extends the pure approach to the case of cross-spectra. The pure pseudo-spectrum estimators retain speed and efficiency of the standard pseudo-spectrum methods and have been

devised to suppress the effect of the E-to-B leakage, thus ensuring the near-optimality of the estimated B-mode power spectrum. The performance of such an estimator is demonstrated in the last part of this section where the *statistical* uncertainties on the E- and B-mode reconstruction is evaluated thanks to Monte Carlo simulation.

3.1 The pure pseudo-C_ℓ's estimator

In general a polarization field on the partial sky can be divided into three classes of modes: pure E-modes, pure B-modes and ambiguous modes, which are a mixture of the true E- and B-modes (Bunn et al. 2003). The standard pseudo-spectrum approach consists of projecting the polarization fields

$$\mathbf{P} = \begin{pmatrix} Q \\ U \end{pmatrix} \tag{9}$$

on to the full-sky harmonic basis of B-modes

$$\mathbf{Y}^B_{\ell m} = \frac{1}{2}\sqrt{\frac{(\ell-2)!}{(\ell+2)!}} \begin{pmatrix} i(\partial^2 - \bar{\partial}^2) \\ \partial^2 + \bar{\partial}^2 \end{pmatrix} Y_{\ell m}, \tag{10}$$

where $\partial(\bar{\partial})$ stands for the spin-raising (lowering) operator. However, this basis contains both pure B-modes and ambiguous modes on the partial sky. Consequently, the standard pseudo-C_ℓ estimator includes these ambiguous modes in the B-mode power spectra estimates, thus resulting in contamination from the much larger E-mode contribution – an effect hereafter referred to as E-to-B leakage. This can be removed on the averaged spectra, but will still lead to a significant increase in the estimated spectrum variance.

The ambiguous modes can however be filtered out by projecting the polarization field on to a pure B-mode basis. Such a basis is constructed from the spherical harmonics and a particular window function W, such that W and ∂W vanish on the boundary of the observed sky:

$$\mathcal{Y}^B_{\ell m} = \sqrt{\frac{(\ell-2)!}{(\ell+2)!}} \begin{pmatrix} i(\partial^2 - \bar{\partial}^2) \\ \partial^2 + \bar{\partial}^2 \end{pmatrix} W Y_{\ell m}. \tag{11}$$

Pseudo-multipoles free of E-to-B leakage can then be computed from this basis by taking the dot product

$$\tilde{a}^B_{\ell m} = \int d\mathbf{n}\, \mathcal{Y}^{B\dagger}_{\ell m} \cdot \mathbf{P}, \tag{12}$$

from which a pseudo-power spectrum can be derived:

$$\tilde{C}^B_\ell = \frac{1}{2\ell+1} \sum_m \tilde{a}^B_{\ell m} \tilde{a}^{B\star}_{\ell m}. \tag{13}$$

The pseudo-power spectrum for the E-modes is identically derived by projecting the Stokes parameters on to the harmonic basis for the E-modes, defined as

$$\mathcal{Y}^E_{\ell m} = \frac{1}{2}\sqrt{\frac{(\ell-2)!}{(\ell+2)!}} \begin{pmatrix} \partial^2 + \bar{\partial}^2 \\ -i(\partial^2 - \bar{\partial}^2) \end{pmatrix} W Y_{\ell m}. \tag{14}$$

Unbiased estimates for both the E- and B-mode power spectra are finally obtained by solving the linear system

$$\begin{pmatrix} M^{(+)}_{\ell\ell'} & M^{(-)}_{\ell\ell'} \\ M^{(-)}_{\ell\ell'} & M^{(+)}_{\ell\ell'} \end{pmatrix} \begin{pmatrix} C^E_{\ell'} \\ C^B_{\ell'} \end{pmatrix} = \begin{pmatrix} \tilde{C}^E_\ell - N^E_\ell \\ \tilde{C}^B_\ell - N^B_\ell \end{pmatrix}, \tag{15}$$

where $N^{E/B}_{\ell'}$ stands for the noise bias and $\mathbf{M}^{(+/-)}_{\ell\ell'}$ for the different mode–mode coupling matrices. We emphasize that the pure formalism adopted hereafter corrects for the E-to-B leakage due to partial sky coverage only, leading to elements of $M^{(-)}$ much smaller

than those of $M^{(+)}$. However, such off-diagonal elements are not strictly zero because of pixel-induced E-to-B leakage. This remaining leakage between the two type of polarization is nevertheless very small and carefully taken into account in the computation of the mode–mode coupling matrices.

The required extensions are included in our implementation of the pure pseudo-C_ℓ's estimator, called XPURE, which we use in this work. XPURE is a generalization of the XSPECT and XPOL codes (Tristram et al. 2005). It can handle multiple maps for computing auto- and cross-power spectra. The different mode–mode couplings due to partial sky coverage (i.e. ℓ-to-ℓ' mixing) and pixelization (i.e. residual E-to-B leakage) are accounted for during the mode–mode coupling matrix computation. The code is based on the S^{2}HAT library^2 – a parallel library allowing for efficient computation of spin-weighted spherical harmonic transforms.

3.2 Sky apodization

The applicability of the pure pseudo-C_ℓ's estimator strongly depends on being able to compute the sky apodizations which are needed in the calculation of the relevant pure basis functions, equation (11), and which have to fulfil appropriate boundary conditions. Different functions have been proposed, ranging from those derived from an optimization procedure to those based on some analytic expressions. Their relative merit has been extensively discussed by Grain et al. (2009) who showed that the optimized sky apodization scheme proposed in Smith & Zaldarriaga (2007) leads to the lowest variance on the power spectrum estimation.

The underlying strategy for deriving such an optimized sky apodization is to search for the W functions which make the pure pseudo-C_ℓ as close as possible to optimal, quadratic power-spectrum estimator (see section V of Smith & Zaldarriaga 2007). An optimized weighting scheme therefore involves a specific sky apodization for each ℓ-band for which the power spectrum is to be estimated, according to the signal and noise power aliasing in each band. As shown in Smith & Zaldarriaga (2007), this optimization procedure can deal with both external and internal boundaries, due to limited sky coverage and holes induced by point source removal, respectively. Such an optimization procedure, however, assumes that noise and signal are known. Although this is a good assumption for noise and E-modes (which can be recovered precisely enough without optimization), this does not hold for the B-modes and an erroneous prior may introduce suboptimality in the B-modes estimate. However, it has been shown by Grain et al. (2009) that the B-modes prior only mildly affects the performance of the power-spectrum estimation using the optimized sky apodization. This weak dependence is in fact a direct consequence of the optimization process, which attempts to select an apodization to reduce first the sampling variance due to the E-to-B leakage and second the noise variance, with the B-mode variance itself usually being subdominant.

These optimized sky apodizations, leading to the lowest statistical uncertainties, will be used throughout this work to compute the polarized power spectra from the CMB maps estimated from the component separation process.

3.3 Statistical uncertainties

A complete characterization of the error budget incurred by both the foreground cleaning and power spectrum estimation processes

is mandatory for setting statistically meaningful constraints on the tensor-to-scalar ratio. We evaluate the sampling and noise variance induced by the power spectrum estimation stage using Monte Carlo (MC) simulation, allowing us to determine the statistical part of the total error budget and to demonstrate the performance of the pure pseudo-C_ℓ's estimator.

The observed sky area for the balloon-borne and ground-based experiments considered hereafter are shown in Fig. 2, where holes due to point-source removal are carefully taken into account. The sky fraction is roughly $f_{sky} = 1$ and 2.5 per cent for the balloon-borne and ground-based experiments, respectively. We assume homogeneous, white noise which is deduced from the noise per frequency channels using equation (6). This gives a noise level in the CMB maps of approximately 1.65 and 3.2 μK per 3.5′ pixel for the balloon-borne and the ground-based experiments, respectively. The simulations use the *WMAP* 5-yr best-fitting cosmological model (Dunkley et al. 2009) and the simulated B-mode includes the lensing and primordial component with a tensor-to-scalar ratio, r, equal to 0.05. The E- and B-mode power spectra are estimated by downweighting the simulated maps with the optimized sky apodization as described in the previous section and finally binned with a bandwidth of $\Delta\ell = 40$, with the lowest ℓ-bin starting at $\ell = 20$.

We demonstrate the performance of the pure estimator in Fig. 1 which shows the input E- and B-mode power spectra and the estimated variances of the reconstructed E- and B-mode power spectra derived as the standard deviation from 500 MC simulations. The three variances displayed on this figure are obtained from, (i) the Fisher estimate, or the so-called f_{sky}-formula; (ii) the standard deviation of 500 MC simulations using pure pseudo-C_ℓ's estimator; and (iii) the standard deviation of 500 MC simulations using standard pseudo-C_ℓ's estimator.

For the B-modes, the variance is significantly reduced by using the pure pseudo-C_ℓ's estimator at angular scales where sampling variance is dominating ($\ell \leqslant 700$), while the two approaches lead to similar variances at smaller angular scales where noise is the dominant contribution to the statistical uncertainties. More specifically, this technique is mandatory for the balloon-borne experiment to be able to extract the B-mode from the maps of the two Stokes parameters at intermediate and large angular scales ($\ell \leqslant 400$) while the variance is reduced by a factor of 2 at those scales for the ground-based experiment. Moreover, the pure pseudo-C_ℓ's estimator is required for both type of experiments to be able to *statistically* disentangle the inflationary gravitational waves at $\ell \leqslant 100$ from the secondary, lensing-induced B-mode.

For the E-modes, the standard pseudo-C_ℓ's estimator is preferred, for achieving higher accuracy (at least at large angular scales for the balloon-borne sky coverage). The reason for such a higher efficiency of the standard approach as compared to the pure one for E-mode is twofold: on the one hand, B-to-E leakage only mildly affects the variance in the standard approach and, on the other hand, ambiguous modes are mainly composed of E-modes, and so a significant amount of information may be lost by removing them using the pure approach.

4 SIMULATED SKY

In this section we describe the sky model we use for simulations in this work. We introduce the 'basic model', which is a simple model of the sky signal consistent with the currently available information on diffuse foregrounds, and then discuss a number of simple

---

^{2} S^{2}HAT: www.apc.univ-paris7.fr/~radek/s2hat.html
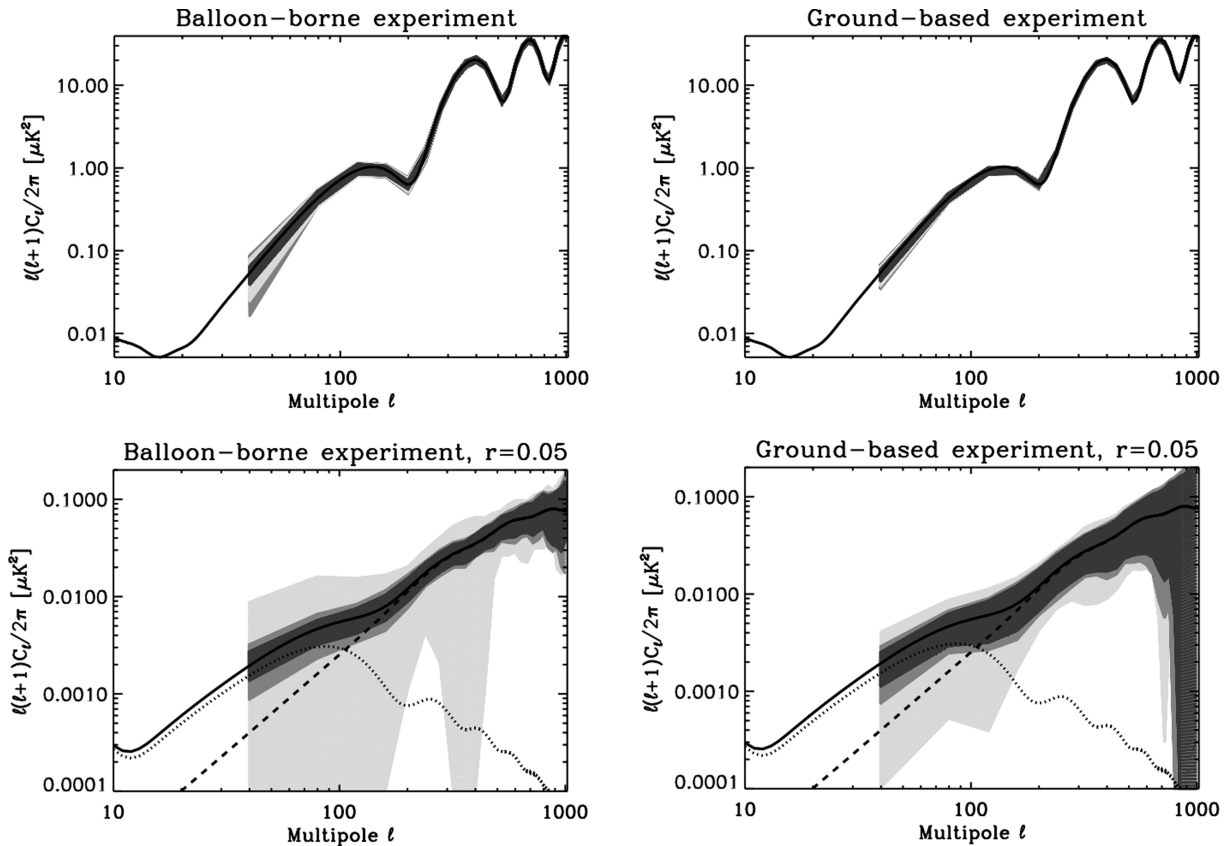PURES^{2}HAT: www.apc.univ-paris7.fr/~radek/pureS2HAT.html

Figure 1. Variance of the estimated E-mode (upper panels) and B-mode power spectrum (lower panels) for the balloon-borne (left panels) and ground-based (right panels) sky-coverage as shown in Fig. 2. From darkest to lightest grey: Fisher estimation (i.e. the so-called f_{sky} approximation), standard deviation of 500 Monte Carlo simulations using the pure pseudo-C_ℓ's estimator as implemented in XPURE and, standard deviation of 500 Monte Carlo simulations using the standard pseudo-C_ℓ's estimator as implemented in XPOL. We stress that XPURE provides lower variances than XPOL for the B-modes, while XPOL performs better for the E-modes.

extensions whose physical parameters are poorly constrained by current observations.

4.1 Sky signals – basic model

Our analysis area is centred around RA $= 62°$ and Dec. $= -45°$. This sky region, which is in the anti-Sun direction during the austral summer, has already been observed by several CMB polarization experiments including Boomerang (Montroy et al. 2006), ACBAR (Reichardt et al. 2009), QUAD (Brown et al. 2009) and will be targeted by future experiments like EBEX (Oxley et al. 2004; Grainger et al. 2008). We first note that *WMAP*-based estimates of the level of unresolved point source power together with conservative assumptions about the expected level of radio source polarization suggest that this signal will be negligible compared to the lensing B-mode signal (Ben Gold, private communication). Similar conclusions hold for the infrared sources. Certainly though, a few bright extragalactic sources will be present in the field both by chance and for the purposes of in-flight beam mapping and calibration. The process of masking out these sources is mimicked in our simulations, similar to what is done in this respect by Smith & Zaldarriaga (2007), as shown in Fig. 2.

Our Galactic sky model on this relatively high Galactic latitude region consists of two diffuse foreground components: synchrotron and thermal dust emission. The synchrotron total inten-

sity emission was simulated using the 408-MHz map of Haslam, Wielebinski & Priester (1982), with free–free emission subtracted and small-scale power added by Giardino et al. (2002). We extrapolated this template[3] up to 65 GHz, using the *WMAP* 5-yr maximum-entropy method derived synchrotron spectral index map (Gold et al. 2009).

The thermal dust total intensity emission is taken from the combined *COBE*-DIRBE and IRAS dust template of Schlegel, Finkbeiner & Davis (1998), extrapolated in frequency using the scaling accounting for temperature variations as described in Finkbeiner, Davis & Schlegel (1999). We first extrapolate the dust down to 65 GHz for matching our polarization model with *WMAP* data as indicated below, and then extrapolate it back into our chosen frequency range according to a uniform greybody scaling law inspired by FDS Model 3 (Finkbeiner et al. 1999):

$$A_{dust} \propto \frac{\nu^{\beta_d+1}}{\exp\frac{h\nu}{kT_d} - 1}, \qquad (16)$$

where $T_d = 18.0\,\mathrm{K}$ and $\beta_d = 1.65$.

In order to simulate the polarization in this region, we normalized the large-scale polarization amplitude to the E and B spectra estimated by Page et al. (2007). This normalization is achieved by first assuming that the polarized intensity of both synchrotron

---
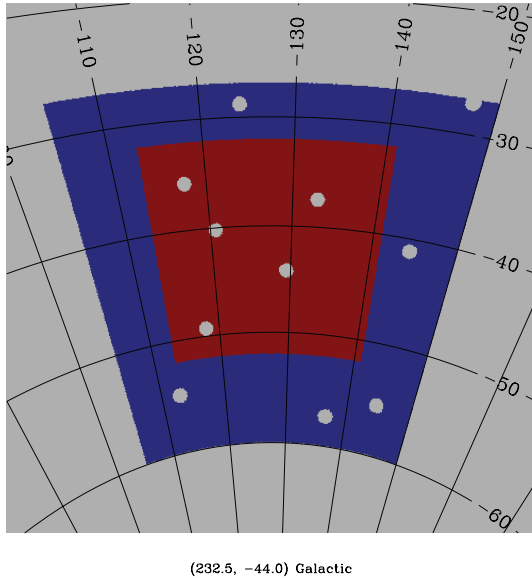
[3] ftp.rssd.esa.int/pub/synchrotron

Figure 2. Sky areas for the balloon-borne (inner) and ground-based (outer, larger patch) experimental set-ups. The holes mimic the effect of masking resolved point sources and are included in our optimized apodization calculation.

and dust is proportional to the total intensity, which introduces two free parameters, p_{dust} and p_{sync}. In order to obtain the large-scale polarization angles θ, we take the Q and U templates from the *WMAP* polarized dust template, which is based on information derived from starlight polarization and a geometric suppression factor taken from a three-dimensional Galactic magnetic field model described in Page et al. (2007). This template introduces a large-scale modulation to the polarization pattern of the synchrotron and dust. We then add Gaussian fluctuations to the polarization angles on smaller scales following the method of Giardino et al. (2002), assuming a model $C_l^{\cos 2\theta, \sin 2\theta} \propto l^\alpha$, where $\alpha = -3$.

With these templates of I, Q and U, we evaluate the power spectra for E and B modes in order to finally normalize the Q and U maps to an effective polarization fraction. A good match is found for $p_{dust} = p_{sync} = 0.1$. In Fig. 3 we show maps of the synchrotron and dust templates at 150 GHz. The large-scale modulation introduced by the *WMAP* polarized dust template is clear from the correlated appearance of the synchrotron and dust. A certain amount of correlation between these polarized components is expected because the Galac-

tic magnetic field (GMF) is a key ingredient common to both the physics of synchrotron emission and dust alignment. Indeed, theoretical modelling of the GMF is underway by several groups (Page et al. 2007; Miville-Deschênes et al. 2008; Waelkens et al. 2009; Jaffe et al. 2010). However, a global model of both synchrotron and dust polarization in a turbulent GMF at high galactic latitudes is currently unavailable [see, however, the recent work by Fauvet et al. (2010)], and this is why we adopt the semi-empirical approach to our GMF simulations just described. Such a modelling can potentially exaggerate the overall correlation level between these two components by extending it to smaller angular scales. This can in turn have important consequences for the performance of the component separation process, as discussed in Section 6.3.

In the range of frequencies relevant here, the synchrotron contribution is subdominant compared to the dust but increases monotonically with decreasing frequency. These two components become comparable at around 70 GHz, where the total foreground minimum is also found.

4.2 Sky signals – extensions

The problem of insufficient frequency coverage is exacerbated by effects that can generally be described as foreground complexity, which will increase the biases in the recovered components. We now introduce a few extensions to the basic model just described that we will investigate later in our analysis.

4.2.1 Extra small-scale power

We investigate how increasing the foreground power on small angular scales impacts on foreground cleaning by varying the parameter α. While in the basic model, the small-scale power in polarization is rapidly decaying with $\alpha = -3$, we also study a more extreme case with $\alpha = -2$.

4.2.2 Spatially varying frequency scaling

Foreground spectral index spatial variations, if poorly estimated, will compromise the estimation of the CMB signal. In turn this compromises the measurement of the CMB power spectrum, particularly on the angular scales on which the spectral index varies. In the basic model just described, the dust scales uniformly as equation (16). We will also investigate the impact of the FDS temperature variations (Schlegel et al. 1998), whose effective power-law spectral
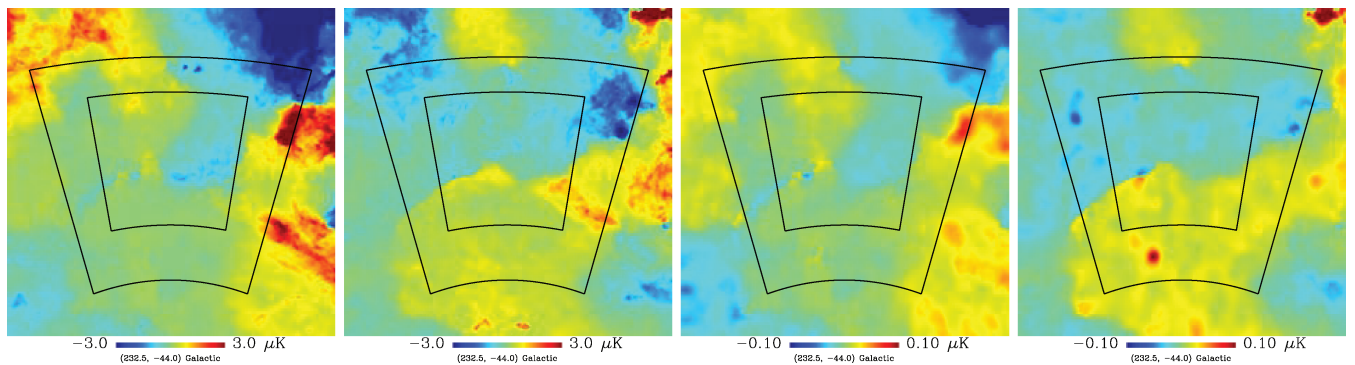

Figure 3. Simulated thermal dust (left) and synchrotron (right panels) Stokes Q and U templates at 150 GHz, for the 'basic model' described in Section 4. The high degree of correlation between the two components on large angular scales is imposed by polarization angle template we assume. The two foregrounds have roughly the same amplitude around 70 GHz.
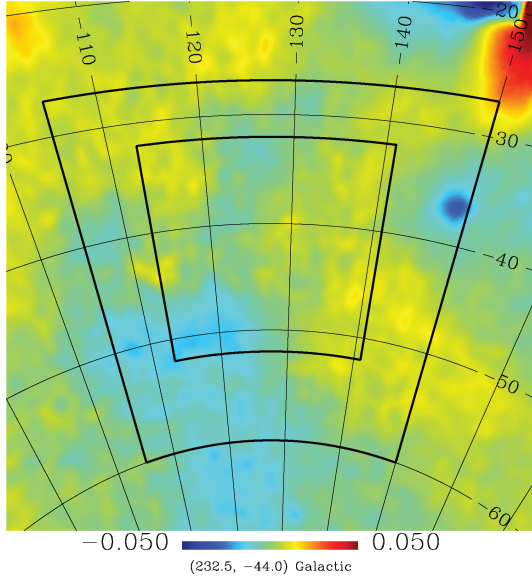
Figure 4. The dust spectral index variations at 150 GHz adopted in this paper (Schlegel et al. 1998).

index (minus the mean) at 150 GHz is shown in Fig. 4, and whose rms variation is 0.008.

Approaches for dealing with spectral index variations include Taylor expanding the foreground continuum emission and fitting the spectral index variations as a new component (Stolyarov et al. 2005), which requires extra frequency channels. Alternatively, the frequency channels can be degraded into many low-resolution, high signal-to-noise ratio pixels (Eriksen et al. 2006), in which case an approach for dealing with the high-ℓ modes is still required.

For the multipixel approach it is clear that the map could be divided up into many subregions in which the spectral index is estimated, thereby introducing extra foreground parameters at a cost of increased noise at low ℓ. Gaining a quantitative understanding of the impact of spectral index variations on the multipixel method is clearly required, since the constancy of the spectral parameters is a central assumption in deriving the spectral index likelihood equation (5), but is not the main focus of this work.

5 MOCK OBSERVATIONS AND FOREGROUND CASE STUDIES

We define here two fiducial experimental set-ups which we will use in the following analysis. They are chosen to reflect some general characteristics of bolometric experiments, but idealized and simplified for demonstration purposes. We emphasize that we do not attempt here to forecast a performance of any specific experiment, but rather, on one hand, to demonstrate the performance of the considered methods in the context of the future B-mode experiments, and on the other, to provide reference numbers quantifying the precision levels potentially achievable by the small-scale experiments.

5.1 Balloon-borne experiment

We consider a balloon-borne experiment data set with the following characteristics.

(i) Survey area: Approximately 1 per cent of sky, corresponding to around 126 000 3.5′ pixels, showed in Fig. 2.

(ii) Frequency coverage: Three channels at 150, 250 and 410 GHz each with the same Gaussian beam of FWHM = 8′.

(iii) Noise level: Homogeneous uncorrelated noise with an rms of 1.5, 4 and 40 μK on a 3.5′ pixel.

Such a set-up can be considered as a minimal choice for a balloon-borne observatory. It takes an advantage of not being limited in the frequency range by the atmosphere on one hand, and on the other it restricts the number of dominant foreground components to one. It can be considered as an idealization of the data set anticipated from, for instance, the EBEX experiment (Oxley et al. 2004; Grainger et al. 2008).

We combine these experimental specifications with the following three case studies.

(i) Basic: Based on the basic sky model and the observation characteristics as defined above.

(ii) Small-scale power: The basic sky model is augmented with extra small-scale power in the dust.

(iii) Varying spectral index: The basic sky model is augmented with a variation of the spectral index.

5.2 Ground-based experiment

Our ground-based experiment is characterized as follows:

(i) Survey area: Approximately 2.5 per cent of the sky corresponding roughly to 320 000 3.5′ pixels, showed in Fig. 2.

(ii) Frequency coverage: Three channels at 90, 150 and 220 GHz each with FWHM = 8′.

(iii) Noise level: Homogeneous uncorrelated noise with an rms of 3, 3 and 9 μK on a 3.5′ pixel.

The frequency range is chosen to fit in the window allowed by the atmosphere on one hand, and on the other to conform with the optimal working conditions for bolometric detectors. As a result, the assumed, covered range is quite limited. We note that going beyond the lower frequency bound can easily be imagined by combining radiometric and bolometric data. We will however use the 'minimal' set-up defined above as our standard case, and extend it on an as-needed basis. We finally note that such characteristics are not far from those planned for, for instance, the first deployment of the POLARBeaR experiment (Lee et al. 2008).

The qualitative difference between the ground-based and balloon-borne configurations is that the ground-based frequency channels are much closer to our foreground minimum at approximately 70 GHz where, by definition, two foregrounds are present. This leads us to consider the ground-based experimental configurations with a more involved set of foreground case studies as follows.

(i) Basic: Based on the basic sky model and the observation characteristics as defined above.

(ii) Dust spectral index prior from balloon-borne experiment: Here we assume the value of the dust spectral index determined by the balloon-borne experiment. We discuss two cases with all three or only two (150 and 220 GHz) channels included.

(iii) Synchrotron template: In this case each channel has been corrected for the presence of the synchrotron signal via subtraction of an external synchrotron template. We will assume that the subtraction is performed down to some predefined precision level.

(iv) Extra low-frequency channel: Here we add an extra channel centred at 40 GHz to the basic data set with the noise as in the 90 GHz map case.
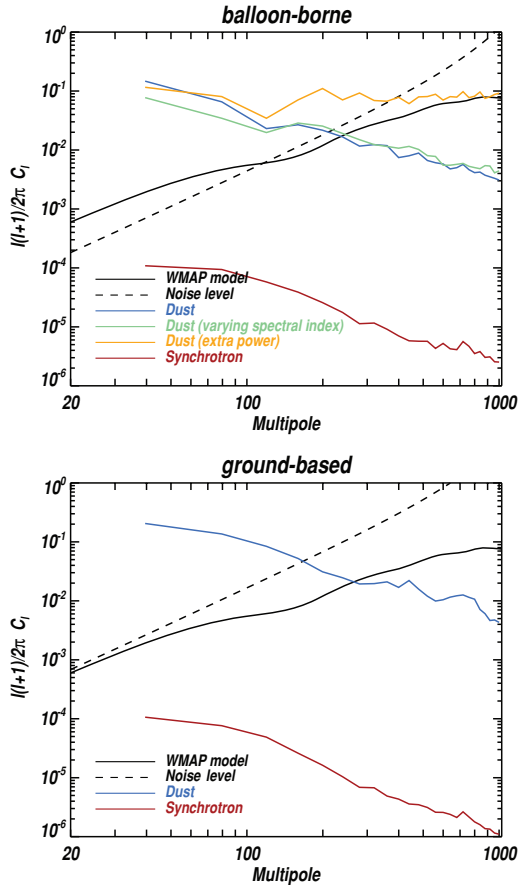
Figure 5. B-mode spectra of the input components at the 150 GHz channel, for the balloon-borne (upper) and ground-based (lower panel) experiments.

We point out that only one of the above cases (the basic case) is both realistic and self-contained. In all the other cases, the presence of some extra external knowledge is postulated.

In addition to the cases listed above, we have also performed the same analysis assuming a lower noise level of 1 μK per 3.5′ pixel in all three channels, as well as devised specific test runs designed to highlight some particular aspects of the performance of the component separation method. These include the following.

(i) No synchrotron: In this case the problematic synchrotron emission is removed from the sky model in order to understand the way in which it biases the dust estimation and subtraction.

(ii) Shuffled synchrotron template: In this case the spatial morphology of the synchrotron template is randomized. The idea here is to understand the effect of accidental correlations between the foreground components in this regime of a restricted number and coverage of the available frequency channels.

We show in Fig. 5 the B-mode power spectra of the input components at 150 GHz. Depending on the angular scale, the dust amplitude must be suppressed by a factor of between 5 and 25 to have successful measurement of the cosmological B-mode signal.

We note again that the noise levels of the derived CMB maps in the basic balloon-borne and ground-based cases assuming a perfect knowledge of the foreground frequency scaling properties are 1.6 and 3.2 μK per 3.5′ pixel, respectively (while in the low-noise ground-case we get 0.9 μK). These can be obtained from equation (6) and are used for the noise level shown in Fig. 5.

5.3 Foreground spectral modelling

Implicit in our approach is the assumption that the three channel configurations of the ground-based and balloon-borne just described will provide information on three, but no more than three, parameters: the CMB amplitude and the dust amplitude at each pixel, and the dust spectral index, as constrained by the ensemble of data and by the two Stokes parameters. In deriving results in the following section, we begin by assuming exactly the same smooth model, equation (16), for fitting the dust as was used in the simulations. To some degree this choice is made for expediency, in order to define a comparison benchmarks for our case studies, and in order to disentangle the effect of different factors on the quality of the CMB estimation. We do, however, attempt to gauge the strength of this assumption by studying two further case studies that are relevant in this context. These are map-level calibration uncertainties and 'spectral mismatch'. Intuitively, calibration uncertainties will downgrade our ability to infer useful information about the dust model and spectral index. Spectral mismatch refers to multiplicative factors applied to the dust model at each channel, breaking the smoothness of the dust emission spectra. This situation could physically occur when the dust scaling cannot be sufficiently well characterized by the three-channel experiment. For instance one can conceive of a case in which two dust components with warmer and colder temperatures also have different polarization fractions, leading to a sudden change in the dust scaling (and/or position angle) coincident with the frequency coverage of the experiment. However, this particular case is thought to be not particularly likely for nearby Galactic cirrus, as reviewed by Dunkley et al. (2009). Spectral mismatch could also result from poorly characterized experimental bandpasses (Church, Knox & White 2003).

6 RESULTS

6.1 Performance evaluation metric

To assess the performance of our component separation method, we will look both at the estimated foreground spectral indices and at the quality of the recovered CMB maps. The latter are clearly not expected to be perfect with potential contamination arising either due to the noise present already in the input, single channel maps, or a failure of the algorithm to perform the separation perfectly. This may result either in some level of non-CMB signal still present in the map, or in the CMB signal being compromised. These two effects are usually referred to as residuals. With the noise uncertainty being quite straightforward to characterize using equation (6), it is our aim to evaluate the level of the residuals expected in the foreground case studies and then to compare it with statistical uncertainties. The latter includes just pixel noise in the case of the maps, and both the noise and CMB signal variance on the level of the power spectra analysis.

For each of the case studies described in the previous section, we first estimate the best-fitting spectral parameters by maximizing the spectral likelihood, equation (5), including on occasions some extra prior information. Then, given the estimated values of the spectral parameters, we compute the map of residuals as

$$\Delta = \hat{s} - s_0 - (\hat{A}^t N^{-1} \hat{A})^{-1} \hat{A}^t N^{-1} n, \tag{17}$$

where s_0 are the input simulated components and the last term on the right-hand side subtracts away the noise in the recovered components $\hat{s}$, given the best-fitting values for the spectral parameter and the known input noise realization n. Directly subtracting the

input noise improves upon the metric used by Leach et al. (2008) in which residual noise was suppressed by smoothing. The CMB element of equation (17), $\mathbf{\Delta}^{\mathrm{CMB}}$, quantifies the residual foreground signal contained in the estimated CMB map, as well as part of the genuine CMB signal correctly interpreted by the algorithm as the CMB contribution. Finally, we compare the latter with the anticipated level of the genuine CMB B-mode signal as well as the level of its statistical uncertainty due to only CMB and noise sampling and cosmic variance. For the last step we use as a metric the B-mode power spectra calculated with the help of the pure estimator. As described in Section 3.3, the spectrum variance is estimated via 500 Monte Carlo simulations, for which we use the best-fitting *WMAP* 5-yr cosmology with $r = 0.05$ as the fiducial model. A satisfactory level of foreground cleaning is achieved when the foreground residual power spectrum is smaller than the statistical uncertainties, ensuring that the systematic errors due to foreground contamination are subdominant to the global error budget.

In Section 7 we perform an analysis of these residuals, and in Section 8 we express the results in terms of an *effective* detectable value of the tensor-to-scalar ratio, r.

6.2 Balloon-borne cases

In Table 1 we report the recovered values of the dust spectral index β for the three tests we have made: in all the cases this parameter was successfully estimated.

Table 1. Results for the dust spectral index estimation in the balloon-borne case. The last column gives the rms of the total foreground contamination left in the CMB map. The input dust spectral index is $\beta_d = 1.65$. For comparison, on the pixel scale the B-mode signal rms is 0.19 μK and the foreground level in the 150 GHz channel is 0.60 μK.

Balloon-borne ($\beta_0 = 1.65$)			
Case	β	$\Delta\beta$	rms (μK)
Basic	1.655	0.009	0.020
Small-scale power	1.655	0.009	0.021
Varying spectral index	1.657	0.011	0.024

In Fig. 6 we show the residuals, $\mathbf{\Delta}$, for the basic case. To appreciate the level reached by the cleaning, they can be compared with the input foreground contamination at 150 GHz, shown in Fig. 3. The resulting residuals for all the balloon-borne cases differ only in a minor way and they are always dominated by the unmodelled synchrotron because the accurate estimates of the dust spectral index, as derived earlier, allow the dust to be subtracted with a superior precision (see Fig. 7).

6.3 Ground-based cases

To perform the foreground cleaning, we again assume the presence of a single dust foreground because the limited number of channels prevents us from performing spectral modelling of the synchrotron present in the data.

The basic result, reported in Table 2, is that the estimated dust spectral index $\beta = 1.875 \pm 0.028$ is significantly biased away from the input value of 1.65, giving rise to residuals significantly higher than found in the balloon-borne cases, as shown in Fig. 8 and quantified by an rms larger by a factor of ~ 8. This bias of the dust spectral index can be explained qualitatively by the fact that in this case not only does the synchrotron component, which remains unmodelled and unsubtracted, have a higher amplitude than before due to a presence of the 90-GHz channel, but also because it has a significant spatial correlation with the dust component, as shown in Fig. 3 and discussed in Section 4.1. We have verified this explanation by making tests first with no synchrotron present and later including only the 'shuffled synchrotron'. The latter case erases the synchrotron–dust correlation, artificially converting the synchrotron to a white-noise-like component with less fluctuations on large angular scales. In both these artificial test cases, satisfactory spectral index estimates and foreground cleaning were obtained (see second and third panel of Fig. 9). Moreover, by changing the Gaussian realization of the foreground spectra (without reshuffling the synchrotron template then), it is possible to reduce the level of correlation between dust and synchrotron. We found that the pixel-by-pixel correlation between the dust and synchrotron, computed with the Pearson coefficient $C = \mathrm{cov}(X, Y)/\sigma_x\sigma_y$, has to drop below ~ 15 per cent to allow for satisfactory foreground cleaning.
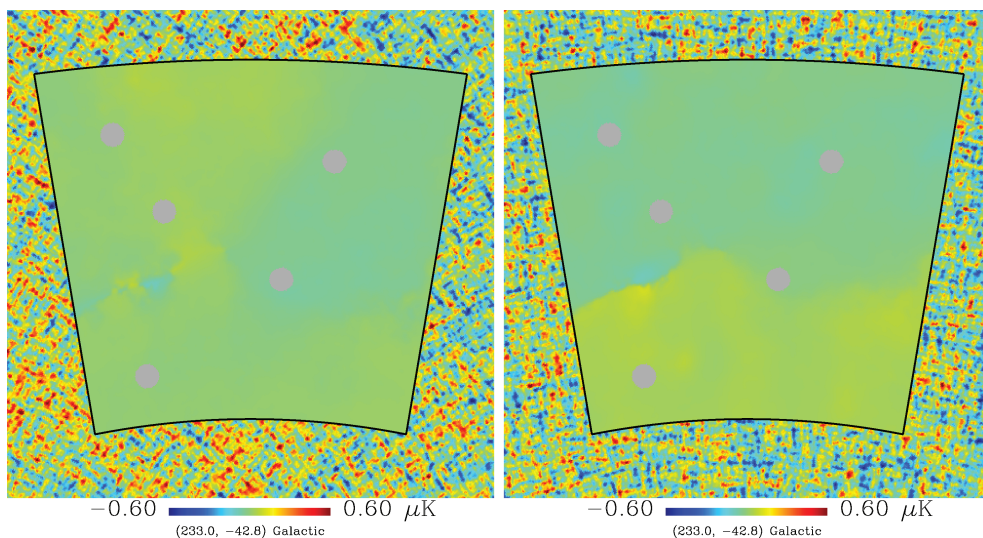

Figure 6. Balloon-borne experiment. Maps of residuals in the CMB map, as defined in equation (17), for the Stokes Q (left) and U (right) parameters in the basic case. For comparison, outside the survey border is shown a pure B-mode realization.
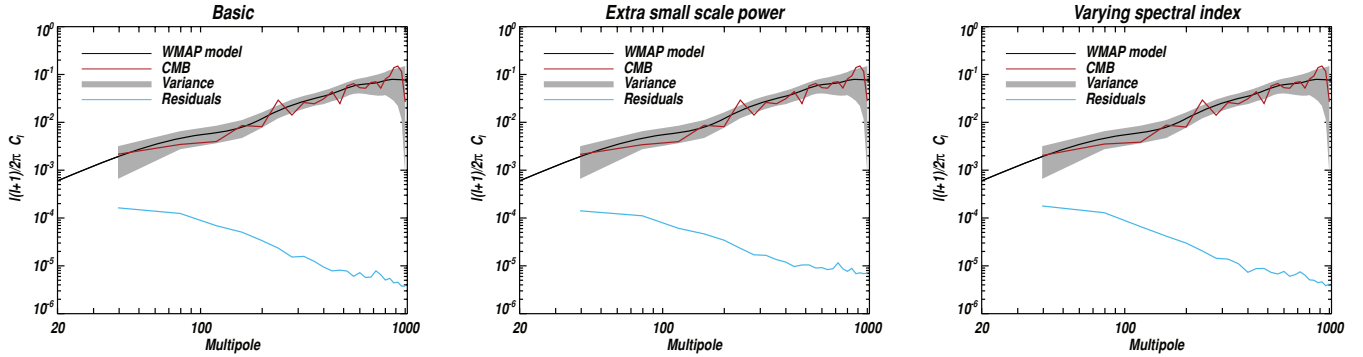
Figure 7. Balloon-borne experiment. B-mode power spectra of the residuals in the CMB map (cyan curves). As a reference, the input and estimated CMB B-mode power spectra are represented by solid black and the solid red curves. The statistical uncertainty of the CMB power spectrum estimation is shown by the grey band. For both E- and B-modes and for all these cases, the level of foreground residuals is smaller than the statistical uncertainties ensuring a precise enough foreground cleaning. From left to right: (i) The basic case, (ii) Extra small-scale power for the dust emission, (iii) Spatially varying spectral index for the dust emission.

Table 2. Results for the dust spectral index estimation in the ground-based case, as in Table 1. The asterisk indicates the cases in which the dust spectral index is fixed, based on the value recovered from the balloon-borne experiment. For comparison, on the pixel scale the B-mode signal rms is 0.19 μK and the foreground level in the 150 GHz channel is 1.3 μK.

Ground-based ($\beta_0 = 1.65$)			
Case	β	$\Delta\beta$	rms (μK)
Basic	1.875	0.028	0.170
No synchrotron	1.642	0.030	0.004
Shuffled synchrotron	1.667	0.028	0.074
90+150+220 + balloon expt.	1.655*	–	0.077
150+220 + balloon expt.	1.655*	–	0.034
External template	1.682	0.026	0.028

6.4 Ground-based cases with external information

The ground-based set-up discussed here is therefore not self-sufficient and thus unable to provide a appropriately cleaned CMB map. In this section we therefore investigate the effect of using 'external information', specifically priors on the foreground spectral indices or an external synchrotron template, on the analysis of this data set.

The first test we made, mimicking a possible real-life situation, was to impose strictly a dust spectral index prior with the value found in the balloon-borne basic case ($\beta = 1.655$). At this point, there are no free spectral parameters to estimate, and the corresponding least-squares components can be directly estimated. The residuals for this case are shown in the 'delta prior' panels of Fig. 10, which demonstrates that the dust spectral index prior does help reduce to some extent the residuals and the final spectral contamination of the B-mode spectrum. Thanks to the high precision of the estimation of the dust spectral index in the balloon-borne experiment, we find that those residuals are again due to the unmodelled synchrotron (on which we will elaborate in Section 7).

Knowing that the 90-GHz channel is contaminated by synchrotron, we also have investigated the impact of simply dropping this channel, fixing again the dust spectral index to the value determined by the balloon-borne experiment. Though clearly rather drastic, this choice could in principle provide a better foreground cleaning than the three channel set-up. Unfortunately, for the specific case analyzed in this work, we found that the remaining two channels are too noisy to produce a CMB cleaned map suitable for
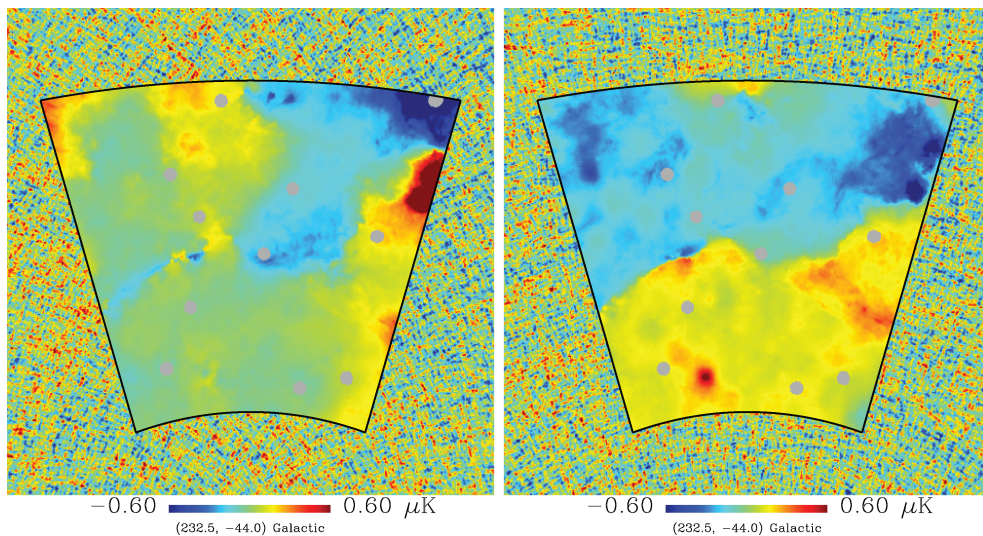

Figure 8. Ground-based experiment. Maps of residuals in the CMB map for the Stokes *Q* (left) and *U* (right) parameters in the basic case, as in Fig. 6.
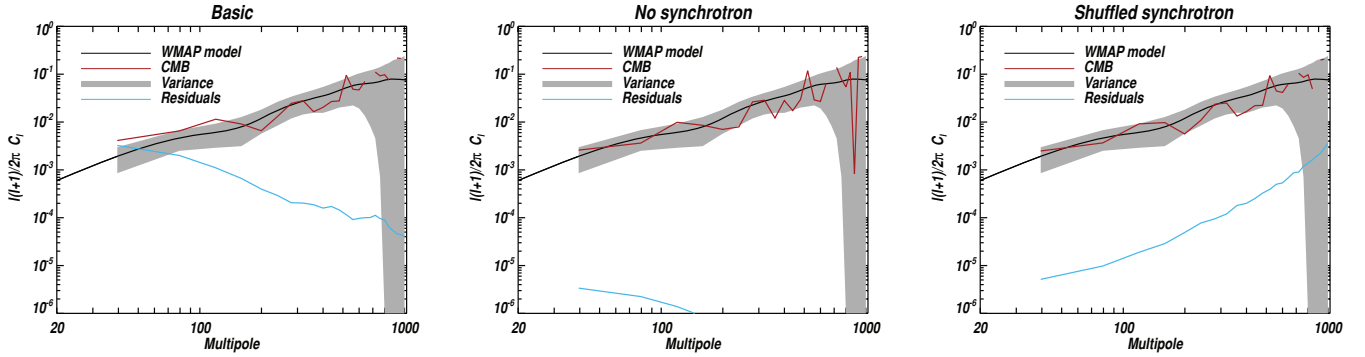
Figure 9. Ground-based experiment. B-mode power spectra of the residuals in the CMB map, as in Fig. 7. From left to right: (i) the basic case, (ii) a test case with no synchrotron in the simulation, (iii) a test case in which the pixels in the synchrotron map were reshuffled to remove the spatial correlation with the dust.
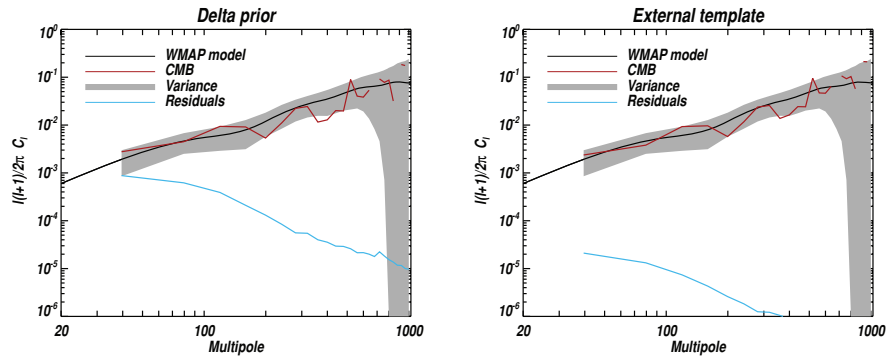

Figure 10. Ground-based experiment. B-mode power spectra of the residuals in the CMB map, as in Fig. 7. Left: The basic case, imposing the the dust spectral index value recovered by the balloon-borne experiment (to be compared with the first panel of Fig. 9). Right: Subtracting a synchrotron template with an amplitude known to be within a 10 per cent.

any B-mode work. We also, however, have found that a two-channel set-up could be a viable option if the noise in these two channels is suppressed down to a ~1 μK level – a rather challenging goal. Nevertheless, this result suggests that some specific attention may need be paid to finding the best trade-off between frequency bands choice and observation time for this kind of experiment.

We also attacked the problem from the other side of the foreground minimum, using information coming from lower frequencies. First, we made use of an external template for the synchrotron, whose amplitude is assumed to be known with a 10 per cent uncertainty, subtracting it directly from the data channels. The satisfactorily foreground cleaned result for this case, in which no prior on the dust spectral index was assumed, is shown in the second panel of Fig. 10. We note that though in this test we have assumed a high-resolution template as only the low-ℓ modes need to be corrected, a low-resolution synchrotron templates, as for example anticipated from the CBASS experiment[4] should be sufficient. Also, in cases where the overall calibration of the available template is considered less reliable than its morphology, the template marginalization could be a more robust technique to be used in this context (Jaffe et al. 2004).

Yet another option we have considered is to extend the covered frequency range by adding an extra frequency channel operating at 40 GHz. This could be achieved for example by co-analysing the data of two bolometric and radiometric experiments observing an overlapping sky area. In our analysis the extra 40 GHz was assumed
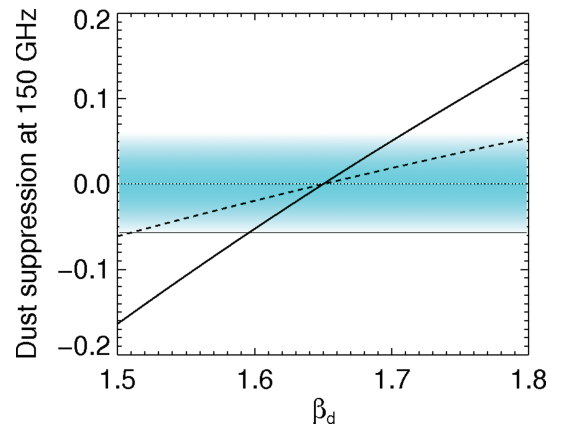
Figure 11. Dust suppression factor in the 150-GHz channel of the basic case of the balloon-borne (solid line) and ground-based (dashed line) experiments, as a function of the recovered dust spectral index β_d. The horizontal shaded band is indicative of the requirement on the dust suppression factor in order to do B-mode science.

to have a similar noise level as the one at 90 GHz. We have indeed found that such a combined analysis fares well in terms of residuals amplitude, which are comparable to those shown in the second panel of Fig. 10, in spite of the fact that the obtained best-fitting value of the spectral index, $\beta \simeq 1.75$, is still found to be significantly away from the true input value. This fact is illustrated in Fig. 11 and discussed in Section 7. We also note that the gain from the extra

channel in the total noise budget of the final CMB map, equation (6), turns out to be negligible.

To recap the results of this section so far, the balloon-borne experiment represents an example of a self-contained case where, owing to the fact that its channels are slightly displaced from the foreground minimum, a single simplistic foreground can be adequately characterized and subtracted. It also helps the ground-based experiment to alleviate the effect of biases caused by spatial correlation between the foregrounds. However, it is easy to envisage the need of the ground-based experiment for further external information on the synchrotron in the form of a template or a lower-frequency channel.

6.5 Miscalibration and spectral mismatch

Here we extend the study conducted so far by incorporating two specific systematic effects which are miscalibration of the data channels and spectral mismatch between the assumed and the real model of foregrounds. These two effects differ not only as to their origin, one being due the instrument properties and the other reflecting our ignorance of the physical phenomena relevant to the following case studies. They also appear differently within the discussed component separation formalism, within which a consistent description of only the miscalibration can be incorporated and thus its effects properly accounted for.

Miscalibration

We simulate a relative calibration error, uncorrelated between the channel input maps and applied directly at the map level (so that no leakage between the Stokes parameters occurs). Though this is clearly a simplification, we note that it is not completely unrealistic and may be expected for experiments implementing a fast polarization modulator, such as a continuously rotating half-wave plate used by MAXIPOL (Johnson et al. 2007) and under development for use with EBEX (Grainger et al. 2008). In such experiments the three Stokes parameters can be disentangled from data of any single detector, and the resulting maps co-added a posteriori in a noise-weighted fashion. The impact of calibration errors and uncertainties can be mitigated by modelling the calibration parameters at the same time as estimating the spectral parameters, as mentioned in Section 2.

We simulate several cases in which we impose different predefined miscalibration values, centred on $\omega_i = 1 + \delta\omega_i$, introducing Gaussian priors on the calibration parameters, centred on $\omega_i = 1$ with width $\delta\omega_i$. We then quantify the impact of calibration errors by calculating the ratio between the residuals in the basic, perfectly calibrated case, and the miscalibrated cases. We report these ratios in Table 3 which shows the effect of miscalibrating the 250- and the 410-GHz channels with respect to the 150-GHz channel, finding

Table 3. This table summarizes how miscalibration errors of the two higher-frequency channels relative to the lowest-frequency channel increases the foreground residuals Δ with respect to the basic case.

Channel miscalibration (per cent)	$\dfrac{\Delta_{\text{calibration}}}{\Delta_{\text{basic}}}$
0, 2, 2	~ 2
0, 5, 5	~ 3
0, 10, 10	~ 4

Table 4. This table reports how a mismatch applied to the input dust models, expressed in percentage per channel, increases the foreground residuals with respect to the basic case.

Spectral mismatch (per cent)	$\dfrac{\Delta_{\text{mismatch}}}{\Delta_{\text{basic}}}$
0, 0, 3	~ 1.5
0, 3, 0	~ 2.5
0, 0, 7	~ 4
0, 5, 0	~ 6

for instance that 5 per cent calibration errors in these two channels lead to foreground residuals that are amplified by a factor of 3. From Fig. 7, we can see that this enhancement of the foreground spectrum by a factor of approximately 10 would impact adversely on the large angular scale B-mode estimation.

Spectral mismatch

Here we consider situations where a mismatch between the true and postulated scaling laws for the dust component is present. In the studied cases, we use different dust scaling laws in the simulations, but during the separation process we always assume the same, simple dependence as defined in equation (16). The specific laws used in the simulations are: the Model 8 from Finkbeiner et al. (1999) and an arbitrary mismatch, where the dust amplitudes are changed by some factor from their values as expected in the model in equation (16). Model 8 is a two-temperature model with two specific spectral indices and two temperatures for the dust. Even if its functional form is different from Model 3 of equation (16), the two models are actually very close in the frequency range from 150 to 410 GHz. Fitting the greybody scaling to these Model 8 simulations, we found that the final foreground residual was small and comparable to the other successful cases already shown.

This motivated us to investigate cases with a larger spectral mismatch in which we inserted some discrete multiplicative factors in the dust scaling law used to simulate the frequency channels. We progressively broke the assumption of a perfect knowledge of the dust spectral behaviour, until the model is too far from the simulation, leading to B-mode biases. In Table 4 we report how the results deteriorate, in terms of larger residuals Δ^{CMB}, for some mismatch choices in different channels. The basic result is that mismatches of upwards of 5 per cent in the dust scaling can lead to an enhancement by a factor of 6 and upwards of the foreground residual level. It is perhaps not surprising that an effective modelling and subtraction of foregrounds using a few channels and a few free parameters depends on the underlying smoothness of the foreground frequency scaling.

We note that this test differs from the miscalibration case discussed earlier, as only one of the components amplitude is modified and no prior is used in the separation process.

7 ANALYSIS OF THE RESIDUALS

In the simulation environment we have the power to control details of all the aspects of the separation process. In this section we take the advantage of this fact and investigate the nature and origin of the residuals Δ shown so far in the paper. We emphasize that the considerations presented below do not depend on how the estimate of the spectral parameters has been obtained, and therefore they

apply more generally than just to the specific parametric ML estimator considered here. In fact, the analogous reasoning could be applied, and similar conclusions drawn, in a case of any two-step approach in which first the spectral indices estimates are derived and then the sky components are estimated via equation (4). This includes FastICA (see Bottino, Banday & Maino 2010, and reference therein), neural networks (Nørgaard-Nielsen & Jørgensen 2008), and correlated component analysis (Bonaldi et al. 2007).

As introduced in equation (4), the operator we apply to the data set d to recover the component estimates s is

$$W(\beta) = [\mathbf{A}^t(\beta)\,\mathbf{N}^{-1}\,\mathbf{A}(\beta)]^{-1}\,\mathbf{A}^t(\beta)\,\mathbf{N}^{-1}, \qquad (18)$$

where we explicitly highlight the dependence on the recovered dust spectral index β. Neglecting the presence of the noise, the data can be written as $d = \mathbf{A}(\beta_0)s_0$ (hereafter the subscript 0 refers to true, i.e. input rather than estimated quantities), and therefore,

$$s = W\,d = W\,\mathbf{A}(\beta_0)\,s_0 \equiv Z(\beta)\,s_0, \qquad (19)$$

and thus the residuals can be written down as

$$\mathbf{\Delta} = s - \hat{s}_0 = [\mathbf{Z}(\boldsymbol{\beta}) - \mathbf{I}]\,s_0, \qquad (20)$$

where the last definition of $\mathbf{\Delta}$ coincides with equation (17) if no noise is considered. Here $\hat{s}_0$ refers to the true, input component, which is modelled in the separation process and is thus a subvector of s_0. The matrix $\mathbf{I}$ is made of a square unit matrix, corresponding to all modelled components, supplemented by extra columns of zeros – one for each unmodelled component.

The size of the matrix, $\mathbf{Z}$, depends on the number of actual and derived sky components and not on the number of the observed frequency channels. In most of the examples shown in this paper, $\mathbf{Z}$ is a 2×3 matrix, since we attempt to recover only two (out of three) components. We note that once we have estimated $\boldsymbol{\beta}$, the operator $\mathbf{Z}(\boldsymbol{\beta})$, that transforms the input components into the output ones, can be readily computed since, in the simulations, we know the input scalings. In this case, equation (20) provides insight into the origin of the residuals and their relative amplitudes.

We first observe that

$$\mathbf{Z}(\boldsymbol{\beta} = \boldsymbol{\beta}_0) = \mathbf{1}, \qquad (21)$$

if the number of assumed and actual components is the same. If there are more components used for the simulations then subsequently recovered, this will no longer be the case. However, even then the maximal square block of the matrix $\mathbf{Z}(\boldsymbol{\beta} = \boldsymbol{\beta}_0)$ will be equal to a unit matrix. This is shown in the upper part of Table 5. Clearly, even perfect knowledge of the true dust spectral index does not assure the lack of the residual. Nevertheless, in such a case each of the recovered components contains only a contribution of this component plus one due to the unmodelled signal. In a specific case of the balloon-borne experiment considered here, nearly all of the unmodelled synchrotron is added mostly to the recovered CMB signal, given the similar scaling of both these components in the respective frequency bands.

The matrix $\mathbf{Z}(\beta)$ computed in a more general and realistic case, when the spectral index is unknown and needs to be estimated from the data, is shown in the middle part of Table 5. We note first that as before, and for the same reason, the unmodelled synchrotron contributes predominantly to the CMB residuals. Nevertheless, the dust is now divided between the two recovered components, though the dust signal found in the recovered CMB is very subdominant. Also in the previous case, the recovered CMB component contains the entire CMB signal, which is completely absent in the recovered dust template. This is reminiscent of equation (21), which in the

Table 5. This table reports the values of the $\mathbf{Z}$ matrix as defined in equation (19). It shows how the input components are weighted in the outputs, for the balloon-borne basic case and an ideal case where the dust spectral index is known. The rms values of the CMB, dust and synchrotron signal in the studied patch are 3.1, 0.6, 0.02 μK for CMB, dust and synchrotron respectively.

Input:	CMB	Dust	Synchrotron
Balloon-borne: Ideal case, $\beta = \beta_0 = 1.65$			
Output:			
CMB	1.000	0.000	1.003
Dust	0.000	1.000	−0.037
Basic case, $\beta = 1.655$			
Output:			
CMB	1.000	0.005	1.003
Dust	0.000	0.994	−0.036
5 per cent miscalibration case, $\beta = 1.639$			
Output:			
CMB	0.988	0.005	0.992
Dust	−0.0004	0.982	−0.037

Table 6. This table reports the values of the $\mathbf{Z}$ matrix. It shows how the input components are weighted in the outputs, for the studied ground-based cases. The rms values of the CMB, dust and synchrotron signal in the studied patch are 3.1, 1.3, 0.02 μK for CMB, dust and synchrotron, respectively.

Input:	CMB	Dust	Synchrotron
Ground-based: Ideal case, $\beta = \beta_0 = 1.65$			
Output:			
CMB	1.000	0.000	3.116
Dust	0.000	1.000	−1.45
Basic case, $\beta = 1.875$			
Output:			
CMB	1.000	0.067	2.993
Dust	0.000	0.919	−1.450
Synchrotron template, $\beta = 1.682$			
Output:			
CMB	1.000	0.004	3.108
Dust	0.000	0.995	−1.441

present case holds only for the CMB component reflecting the fact that the perfect blackbody derivative scaling is assumed on both the simulation and recovery stages. In this case again, the CMB residuals, $\mathbf{\Delta}^{\mathrm{CMB}}$, do not contain any CMB. This is no longer true if the calibration errors are allowed for as shown in the bottom of Table 5. In this case the recovered CMB component contains only part of the total CMB signal as determined by W-matrix weighted average of the relative calibration errors for each of the channels. The remainder of the CMB is then found in the recovered dust. We note, however, that though for the calibration errors considered here these effects are small, the recovered CMB residual is now indeed typically a mixture of the CMB, dust and synchrotron signals. We point out that, maybe somewhat counter-intuitively, the elements of $\mathbf{Z}$ in any column do not have to sum up to unity. This reflects the

fact that due to our wrong assumptions about the spectral parameters values, the estimated components contain overall 'more' of the input components than there really is. This effect is small, if the estimated β values are close to the true ones, as to first order in β, the columns of $\mathbf{Z}$ do sum up to (nearly) unity.

Finally, in all the cases we find that although the code leaves in most of the unmodelled synchrotron, it cleans the dust to better than ~0.5 per cent. The latter depends on the specific value of β assumed for the recovery as shown in Fig. 11, which shows the relative contamination of the recovered CMB due to the dust as a function of the recovered spectral index in the case of the balloon-borne and ground-based basic cases. The shaded band indicates the precision which is needed to avoid contamination in the cosmological B-mode recovery.

The results obtained in the ground-based cases are qualitatively similar to the ones described above as one can see in Table 6. However, in the basic case, due to the different assumed frequency coverage, even for the true value of β we find non-zero contributions of the synchrotron in both recovered components. Moreover, the contribution to the CMB is more than three times that of the actual synchrotron signal at 150 GHz. Assuming in turn the best-fitting value β, we find that the synchrotron levels in both recovered components remain essentially unchanged; however, the *absolute* dust contribution to the CMB template increases to become of the same order as that of synchrotron.

8 COSMOLOGICAL B-MODE DETECTION

The central question of this work is to understand how much the recovery of the cosmological B-modes is affected by the presence of the foregrounds or foreground residuals left over by some foreground-cleaning technique. This question is often phrased as a question about the detectable values of the tensor-to-scalar ratio, r, of the corresponding primordial power spectra, and the calculated limiting values of r are dependent on the implicit and explicit assumptions made in the course of its derivation.

In the context of this paper this question is, however, well defined as we study specific foreground separation and power spectrum estimation algorithms. Our goal here is to derive values of r, which can not only be detected by the considered experiments but also be convincingly argued for from a perspective of an observer as indeed being driven by the primordial signal of the cosmological origin. The limits we aim for here are not to be considered as some 'ultimate' lower limits on r, as often quoted in the literature, but rather as representative of the potential of the specific experiments and data analysis techniques considered here. Reaching values of r lower than our estimates by the actual experiments once they are deployed and operating, may not only be plausible but indeed is expected owing to the build-up, over time, of our knowledge of foreground properties and data analysis tools.

We start by developing a model to describe statistically the foreground residuals found in the recovered CMB map, thus extending the pixel-domain discussion of the last section into the power spectrum domain. We assume that our estimates of β are unbiased and that the obtained statistical uncertainty of the β recovery is small, as is indeed the case in the cases considered. Now, if all the actual components are included in our data model, as defined by $\mathbf{W}(\beta)$, it is accurate to write

$$\mathbf{Z}(\beta) - \mathbf{1} \simeq \delta\beta \, \frac{\partial Z(\beta_0)}{\partial \beta}, \tag{22}$$

where $\delta\beta (\equiv \beta - \beta_0)$ is assumed to be a Gaussian variable with a zero mean and the dispersion as derived earlier, and we used the fact that $\mathbf{Z}(\beta_0) = 1$, as in equation (21). Using equation (20) we can then express the residuals of all the modelled components as

$$\Delta^{\mathrm{mod}} = [\mathbf{Z}(\beta) - \mathbf{1}]\, s_0 \simeq \delta\beta \, \frac{\partial \mathbf{Z}(\beta_0)}{\partial \beta}\, s_0 \tag{23}$$

$$= \delta\beta \left[\frac{\partial \mathbf{W}(\beta_0)}{\partial \beta} \right] \mathbf{A}(\beta_0)\, s_0, \tag{24}$$

and use a first row of this matrix equation to compute the foreground residuals as found in the recovered CMB map. We first introduce a tensor, α_k^{ij}, defined as

$$\alpha_k^{ij} \equiv \frac{\partial \mathbf{Z}_{ij}(\beta_0)}{\partial \beta_k} \tag{25}$$

and we can then express a combined residual in the CMB map ($i = 0$) due to all the modelled, non-CMB components as

$$\Delta^{\mathrm{CMB}} = \sum_{k,j} \delta\beta_k \, \alpha_k^{0j} \, s_0{}^j. \tag{26}$$

This shows that the residuals behave like templates with amplitudes randomized due to the impact of the CMB and noise variance on the determination of the spectral parameters. Hereafter we will assume that the latter are Gaussian variables centred at the true value, β_0, and with dispersions as estimated from the data, $\Sigma^2 (\equiv \langle \delta\beta \, \delta\beta^t \rangle)$. Consequently, for pure spectra averaged over the statistical ensemble (of the noise and CMB), we can then write

$$\mathbf{C}_\ell^\Delta = \sum_{k,k'} \sum_{j,j'} \Sigma_{kk'}^2 \, \alpha_k^{0j} \, \alpha_{k'}^{0j'} \, \mathbf{C}_{0,\ell}^{jj'}, \tag{27}$$

where $\mathbf{C}_{0,\ell}^{jj'}$ are the (pure) auto- and cross-spectra for every pair of the actual j and j' non-CMB components.

In a case of foreground components which cannot be, and/or are not, modelled in the separation process, we can no longer use the procedure outlined above to estimate their residuals. From the previous section we note that such residuals depend only weakly on the assumed spectral parameters and thus will not be stochastic over the CMB plus noise realizations and should rather be treated as a bias. Moreover, in the specific cases considered in this paper, we note that the element of the matrix, $\mathbf{Z}$, which determines the weight with which the component is added to the cleaned CMB map, is on order of at most a few, which together with the fact that the overall expected level of the synchrotron is very low, leads to the conclusion that indeed it can be neglected for the estimations of r as derived here. This last statement can be phrased somewhat more formally by using the approach of Amblard et al. (2007), which allows us to quantify values of r, denoted by r_{res}, which will not be affected by a bias due to some present residual. Indeed, we define the two quantities as

$$s(r) = \sum C_l^{\mathrm{cmb}}(r), \tag{28}$$

$$u = \sum C_l^{\mathrm{res}}. \tag{29}$$

In the successful cases considered in the paper, we find that typically $u \sim s(r_{\mathrm{res}})$ for $r_{\mathrm{res}}^{(b)} = 0.005$ and $r_{\mathrm{res}}^{(g)} = 0.015$ for the balloon-borne and ground-based experiments, respectively. They represent the smallest values of r for which the unmodelled residuals due to synchrotron can be neglected.

We will now assume that the bias due to the unmodelled components is negligible and proceed to the estimation of the parameter

r accounting for the extra variance due to the modelled component residuals. For this we use a Fisher-like approach, which in our case reads

$$F(r) = \frac{\partial \boldsymbol{C}_b^{\mathrm{CMB}}}{\partial r} \, \boldsymbol{\Sigma}_{bb'}^{-1}(r) \, \frac{\partial \boldsymbol{C}_{b'}^{\mathrm{CMB}}}{\partial r}, \tag{30}$$

where b denotes the ℓ-bin number used for the power spectrum estimation, and

$$\begin{aligned}\boldsymbol{\Sigma}_{bb'}(r) &\equiv \mathrm{Var}\left(\boldsymbol{C}_{bb'}^{\mathrm{CMB+noise}}\right) + \boldsymbol{C}_b^{\mathrm{CMB+noise}} \boldsymbol{C}_{b'}^{\boldsymbol{\Delta}} \\ &+ \boldsymbol{C}_{b'}^{\mathrm{CMB+noise}} \boldsymbol{C}_b^{\boldsymbol{\Delta}} + \boldsymbol{C}_b^{\boldsymbol{\Delta}} \boldsymbol{C}_{b'}^{\boldsymbol{\Delta}}.\end{aligned} \tag{31}$$

Here the first term is the covariance of estimated pure CMB B-mode spectra, which we assume to be diagonal in the bin-space and which is estimated via MC simulations for a grid of values of r. $\boldsymbol{C}^{\mathrm{CMB+noise}}$ are the estimated pure CMB + noise spectra also computed on a grid of r values. The partial derivatives are computed using the binned theoretical spectra, obtained in this case from the CAMB code. This is justified given that the pure estimator is assumed to be unbiased. The last term in equation (31) describes the variance due to the residual foreground treated here as a template fully correlated between different bins. The third and fourth terms reflect the cross-terms between the foreground residuals and CMB + noise. These again are non-diagonal in the bin domain. We note that the off-diagonal template-like correction appearing in equation (31) corresponds roughly to the tensor-product term found in equation (A4) of Stompor et al. (2009), which indeed describes the correction to the noise correlation matrix, equation (6), due to the foreground residuals.

We then search for the values of r_{d} such that $r \geqslant r_{\mathrm{d}} \simeq 2 \, F^{-1/2}(r)$, which can hence be detected at a confidence level not smaller than 95 per cent. For the experimental set-ups described in this work, we find that $r_{\mathrm{d}} \sim 0.04$ both for the balloon-borne and the ground-based cases. This value of r is sufficiently high that the bias expected due to the unmodelled synchrotron is indeed much smaller than the detected values. We also find that the derived value of r is limited by the CMB + noise variance, with the foregrounds effects being subdominant. This observation is perhaps not surprising given the low level of the residuals as discussed in the previous sections. However, only an analysis as the one presented in this section, can properly account for the bin–bin correlations of the foreground residual templates.

We note that up to this point we have assumed that we are privy to some insights as to the true nature and morphology of the foreground signals and their scaling well beyond what is usually available to an actual CMB observer and data analyst. In a real-life situation we will lack some of that information. Specifically, we will be likely ignorant of the true value of the $\boldsymbol{\beta}$ parameters, i.e. $\boldsymbol{\beta_0}$, and cross-spectra of all the actual sky components, $\boldsymbol{C}_{0,\ell}^{j,j'}$. That may not look like a big issue given our conclusion above stating that the dominant uncertainty will be due to the CMB and noise variance. However, this conclusion may need to be corroborated case-by-case using data, analyzed together with some necessary external information. Whenever it turns out not to hold, a consistent procedure to account for the extra effects may be needed.

In the case in hand this can be done by replacing the needed information with their best estimates derived in the analysis process. We will thus use the best-fitting value in place of $\boldsymbol{\beta_0}$ and the pure spectra of the components estimated as a result of the separation process, and corrected for the noise bias using equation (6) in the case of the auto-spectra, $j = j'$. In Fig. 12 we show how this procedure fares in the case of the dust component in one of the ground-based case examples considered in this paper. To evaluate the bias due
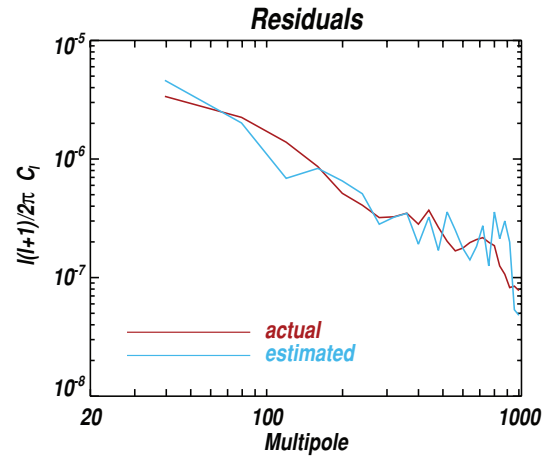

Figure 12. Comparison between actual and estimated residuals for the ground-based, no-synchrotron case, calculated using equation (27).

to the unmodelled synchrotron, we need to use external data. We could correct for such a bias, if the available data are of sufficiently high precision, and include the resulting variance in the final power spectrum error budget. However, more typically, one would rather aim to show that the bias is indeed negligible in the sense defined above in equation (29). For this task the required external data, however, need not be very precise, not least due the fact that the unmodelled components considered here are assumed to be truly subdominant.

Using the estimated quantities in our r-estimation procedure, we recover the limits on r essentially identical to those found before. We thus conclude that values of $r \gtrsim 0.04$ are not only *detectable* at greater than the 95 per cent confidence level, but can also be detected in the realm of an actual experiment and argued for as of a cosmological origin, all of that providing a sufficient control of systematic effects. Moreover, since the residual of the modelled foreground turns out to be much lower than the one induced by the bias of the unmodelled component, we can conclude that the foreground effects start to be important at those values of $r = r_{\mathrm{res}}$ quoted above, respectively $r_{\mathrm{res}}^{(b)} = 0.005$ and $r_{\mathrm{res}}^{(g)} = 0.015$ for balloon and ground case.

9 CONCLUSIONS

In this paper we study the performance of the maximum likelihood parametric component separation method from the point of view of its application to the CMB B-mode polarization analysis. We investigate the residuals left over from the separation in both the pixel and harmonic domains. We propose an efficient framework for evaluating the pixel domain residuals in the simulation, and show how it can be used to gain important insights into the separation process. We then compute the power spectra of the recovered CMB maps, as well as maps of the residuals, using the pure pseudo-spectra technique, and estimate their variances using Monte Carlo simulations. Finally, we propose a Fisher-like approach to evaluate the effects of the foreground residuals on the r parameter and use the latter to derive some estimates of typical values of r, which are potentially detectable by the considered experiments at the 95 per cent confidence level. The latter estimates thus include the uncertainties due to sampling variance, noise scatter, E-to-B leakage, and foreground residuals, all of which are consistently propagated through the proposed pipeline.

We focus here on the small-scale, bolometric experiments, broadly dividing them into two classes, referred to as balloon-borne and ground-based set-ups, both observing in three different frequency bands. We find that the balloon-borne case, with frequency bands at 150, 250 and 410 GHz, provides a robust experimental set-up for the detection of the B-mode polarization. The foreground residuals in the recovered CMB maps derived in this case are found to be usually subdominant. This is true whenever the assumed data model is indeed correct, but it also holds when some small systematic effects are permitted. Selected effects of this kind considered in this work include relative calibration errors, unmodelled spatial variation of the spectral parameters and spectral mismatch between assumed and true spectral scaling laws. We emphasize that all these systematics, though manageable if sufficiently small, may lead to spurious effects in general, and therefore need to be controlled in actual experiments with a sufficient precision, which need to be determined specifically for any experiment.

The success of the considered balloon-borne cases is related to the wide frequency range available to such experiments, which permits selecting frequency bands at the sufficiently high frequencies to avoid the unwanted residual synchrotron. The latter is found to be a dominant source of the bias for the ground-based experiments. In the balloon case we can also afford a long leverage arm between the lowest (CMB-dominated) and the highest (dust-dominated) frequency bands, which plays a pivotal role in setting tight constraints on the spectral parameters of the dust. From our Fisher-like analysis, we show that one could detect r values as low as 0.04 at the 95 per cent confidence level with such experiments, if both our models and measurements are sufficiently well characterized.

For the ground-based case the atmospheric loading restricting the available frequency window proves to be a significant limitation. We find that even in the absence of any systematic effects with three frequency bands set at 90, 150 and 220 GHz, it is generally not possible to produce sufficiently clean CMB maps. This is due to the unmodelled, and thus not separated, synchrotron contribution, which is significant enough (if the polarized emission is at the level suggested by *WMAP*) at these frequencies to bias the estimation of the dust spectral parameters. We point out that this contribution has been neglected in some earlier works, which consequently has arrived at a different conclusion. This therefore emphasizes the importance of accurate sky modelling for this kind of the analysis. Nevertheless, we find a satisfactory cleaning can be achieved in such a case if some external information is available. In particular, we discuss the extended ground experiment analysis allowing for the presence of the extra lower-frequency channels, rough synchrotron templates and priors on the dust scalings. We find that in such cases, and under realistic assumptions, the ground-based experiments should reach a sensitivity roughly matching those found in the balloon-borne case, in terms of a detectable r parameter. We also conclude that, for both these types of the experiments, the foreground contamination anticipated in low-contrast dust regions, should not be an obstacle preventing them from exploring the parameter space of r down to the values $\sim$0.04. Indeed, for the considered experimental set-ups this limit is determined by the uncertainty due to the CMB itself and the instrumental noise, with the effects of the residual foregrounds found to be subdominant. In the realm of the actual observations, whether these limits are reached will be crucially dependent on the control of systematic effects. We note here, however, that the limits derived on r are strongly dependent on the level of the noise assumed in the input single-channel maps. These can therefore be improved upon, if a deeper integration of the same field is performed. However, if no additional external information is used, those limits will remain appropriately higher than the r_{res} values obtained earlier, and below which foreground bias would become significant.

In this context we point out that the framework described in this paper provides a blueprint for similar ongoing or future studies focused on systematic effects. It can also be extended to perform a realistic experiment optimization procedure from the viewpoint of detection of the B-mode signal of cosmological origin.

REFERENCES

Amarie M., Hirata C., Seljak U., 2005, Phys. Rev. D, 72, 123006
Amblard A., Cooray A., Kaplinghat M., 2007, Phys. Rev. D, 75, 083508
Betoule M., Pierpaoli E., Delabrouille J., Le Jeune M., Cardoso J., 2009, A&A, 503, 691
Bonaldi A., Ricciardi S., Leach S., Stivoli F., Baccigalupi C., de Zotti G., 2007, MNRAS, 382, 1791
Bottino M., Banday A. J., Maino D., 2010, MNRAS, 402, 207
Brown M. L. et al., 2009, ApJ, 705, 978
Bunn E. F., Zaldarriaga M., Tegmark M., de Oliveira-Costa A., 2003, Phys. Rev. D, 67, 023501
Church S., Knox L., White M., 2003, ApJ, 582, L63
Dunkley J., Amblard A., Baccigalupi C. et al., 2009, in Dodelson S. et al., eds, AIP Conf. Ser. Vol. 1141, Prospects for Polarized Foreground Removal. Am. Inst. Phys., New York, p. 222
Dunkley J. et al., 2009, ApJS, 180, 306
Efstathiou G., Gratton S., Paci F., 2009, MNRAS, 397, 1355
Eriksen H. K. et al., 2006, ApJ, 641, 665
Fauvet L. et al., 2010, ArXiv e-prints
Finkbeiner D. P., Davis M., Schlegel D. J., 1999, ApJ, 524, 867
Giardino G., Banday A. J., Górski K. M., Bennett K., Jonas J. L., Tauber J., 2002, A&A, 387, 82
Gold B. et al., 2009, ApJS, 180, 265
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Grain J., Tristram M., Stompor R., 2009, Phys. Rev. D, 79, 123515
Grainger W., Aboobaker A. M., Ade P. et al., 2008, in Duncan W. D., Holland W. S., Withington S., Zmuidzinas J., eds, Proc. SPIE 7020, Millimeter and Submillimeter Detectors and Instrumentation for Astronomy IV. SPIE, Bellingham, WA
Haslam G., Wielebinski R., Priester W., 1982, Sky Telesc., 63, 230
Jaffe A. H. et al., 2004, ApJ, 615, 55
Jaffe T. R., Leahy J. P., Banday A. J., Leach S. M., Lowe S. R., Wilkinson A., 2010, MNRAS, 401, 1013

Johnson B. R., Collins J., Abroe M. E. et al., 2007, ApJ, 665, 42

Leach S. M., Cardoso J., Baccigalupi C. et al., 2008, A&A, 491, 597

Lee A. T., Tran H., Ade P. et al., 2008, in Kodama H., Ioka K., eds, AIP Conf. Ser. 1040, POLARBEAR: Ultra-high Energy Physics with Measurements of CMB Polarization. Am. Inst. Phys., New York, p. 66

Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511

Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473

Miville-Deschênes M., Ysard N., Lavabre A., Ponthieu N., Macías-Pérez J. F., Aumont J., Bernard J. P., 2008, A&A, 490, 1093

Montroy T. E., Ade P. A. R., Bock J. J., Bond et al., 2006, ApJ, 647, 813

Nørgaard-Nielsen H. U., Jørgensen H. E., 2008, Ap&SS, 318, 195

Oxley P. et al., 2004, in Strojnik M. ed., SPIE 5543, Infrared Spaceborne Remote Sensing XII. SPIE, Bellingham, WA, p. 320

Page L. et al., 2007, ApJS, 170, 335

Reichardt C. L. et al., 2009, ApJ, 694, 1200

Schlegel D. J., Finkbeiner D. P., Davis M., 1998, ApJ, 500, 525

Smith K. M., 2006, Phys. Rev. D, 74, 083002

Smith K. M., Zaldarriaga M., 2007, Phys. Rev. D, 76, 043001

Stivoli F., Baccigalupi C., Maino D., Stompor R., 2006, MNRAS, 372, 615

Stolyarov V., Hobson M. P., Lasenby A. N., Barreiro R. B., 2005, MNRAS, 357, 145

Stompor R., Leach S., Stivoli F., Baccigalupi C., 2009, MNRAS, 392, 216

Tristram M., Macías-Pérez J. F., Renault C., Santos D., 2005, MNRAS, 358, 833

Tucci M., Martínez-González E., Vielva P., Delabrouille J., 2005, MNRAS, 360, 935

Verde L., Peiris H. V., Jimenez R., 2006, J. Cosmol. Astropart. Phys., 1, 19

Waelkens A., Jaffe T., Reinecke M., Kitaura F. S., Enßlin T. A., 2009, A&A, 495, 697

Zaldarriaga M., Seljak U., 1997, Phys. Rev. D, 55, 1830

This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.