

Component separation methods for the PLANCK mission

S. M. Leach^{1,2}, J.-F. Cardoso^{3,4}, C. Baccigalupi^{1,2}, R. B. Barreiro⁵, M. Betoule³, J. Bobin⁶, A. Bonaldi^{7,8}, G. de Zotti^{7,1}, J. Delabrouille³, C. Dickinson⁹, H. K. Eriksen^{10,11}, J. González-Nuevo¹, F. K. Hansen^{10,11}, D. Herranz⁵, M. Le Jeune³, M. López-Caniego¹², E. Martínez-González⁵, M. Massardi¹, J.-B. Melin¹³, M.-A. Miville-Deschênes¹⁴, G. Patanchon³, S. Prunet¹⁵, S. Ricciardi^{7,16}, E. Salerno¹⁷, J. L. Sanz⁵, J.-L. Starck⁶, F. Stivoli^{1,2}, V. Stolyarov¹², R. Stompor³ and P. Vielva⁵

¹ SISSA-ISAS, Astrophysics Sector, via Beirut 4, Trieste, 34014, Italy

² INFN, Sezione di Trieste, Via Valerio 2, I-34014 Trieste, Italy

³ CNRS & Université Paris 7, Laboratoire APC, 10 rue A. Domon et L. Duquet, 75205 Paris Cedex 13, France

⁴ Laboratoire de Traitement et Communication de l'Information (CNRS and Telecom ParisTech), 46, rue Barrault, 75634 Paris, France

⁵ Instituto de Física de Cantabria (CSIC-UC), Avda. de los Castros s/n, 39005 Santander, Spain

⁶ CEA - Saclay, SEDI/Service d'Astrophysique, 91191 Gif-Sur-Yvette, France

⁷ INAF-Osservatorio Astronomico di Padova, vicolo dell'Osservatorio 5, I-35122 Padova, Italy

⁸ Dipartimento di Astronomia, vicolo dell'Osservatorio 5, I-35122 Padova, Italy

⁹ Infrared Processing and Analysis Center, California Institute of Technology, M/S 220-6, 1200 E. California Blvd, Pasadena, 91125, U.S.A.

¹⁰ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

¹¹ Centre of Mathematics for Applications, University of Oslo, P.O. Box 1053 Blindern, N-0316 Oslo, Norway

¹² Astrophysics Group, Cavendish Laboratory, J J Thomson Avenue, Cambridge CB3 0HE, United Kingdom

¹³ DSM/Irfu/SPP, CEA/Saclay, F-91191 Gif-sur-Yvette Cédex, France

¹⁴ Institut d'Astrophysique Spatiale, Bâtiment 121, F-91405 Orsay, France

¹⁵ Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, F-75014 Paris, France

¹⁶ Space Sciences Laboratory, University of California Berkeley, Computational Cosmology Center, Lawrence Berkeley National Laboratory, CA 94720 USA

¹⁷ Istituto di Scienza e Technologie dell'Informazione, CNR, Area della ricerca di Pisa, Via G. Moruzzi 1, I-56124 Pisa, Italy

May 5, 2008

ABSTRACT

Context. The PLANCK satellite will map the full sky at nine frequencies from 30 to 857 GHz. The CMB intensity and polarization that are its prime targets are contaminated by foreground emission.

Aims. The goal of this paper is to compare proposed methods for separating CMB from foregrounds based on their different spectral and spatial characteristics, and to separate the foregrounds into 'components' of different physical origin (Galactic synchrotron, free-free and dust emissions; extra-galactic and far-IR point sources; Sunyaev-Zeldovich effect, etc).

Methods. A component separation challenge has been organized, based on a set of realistically complex simulations of sky emission. Several methods including those based on internal template subtraction, maximum entropy method, parametric method, spatial and harmonic cross correlation methods, and independent component analysis have been tested.

Results. Different methods proved to be effective in cleaning the CMB maps from foreground contamination, in reconstructing maps of diffuse Galactic emissions, and in detecting point sources and thermal Sunyaev-Zeldovich signals. The power spectrum of the residuals is, on the largest scales, four orders of magnitude lower than that of the input Galaxy power spectrum at the foreground minimum. The CMB power spectrum was accurately recovered up to the sixth acoustic peak. The point source detection limit reaches 100 mJy, and about 2300 clusters are detected via the thermal SZ effect on two thirds of the sky. We have found that no single method performs best for all scientific objectives.

Conclusions. We foresee that the final component separation pipeline for PLANCK will involve a combination of methods and iterations between processing steps targeted at different objectives such as diffuse component separation, spectral estimation and compact source extraction.

Key words. Methods: data analysis; Cosmology: cosmic microwave background

1. Introduction

PLANCK is a European Space Agency space mission whose main objective is to measure the cosmic microwave background (CMB) temperature and polarization anisotropies with high accuracy, high angular resolution and with unprecedented frequency coverage (The Planck Collaboration 2005). In anticipation of the launch, PLANCK is stimulating much research and

development into data processing methods that are capable of addressing an ambitious science programme enabled by these multi-frequency observations. It is expected that PLANCK will break new ground in studies of the CMB, of the interstellar medium and Galactic emission mechanisms on scales down to a few arcminutes, as well as of the emission from many extra-galactic objects.

The processing of such multi-frequency data depends on both the science goals, as well as on the signal to noise regime and on the overall level and complexity of foreground contamination. This observation is borne out by a brief historical perspective on CMB data processing.

An example of the low foreground level and complexity regime is provided by the observations made by the Boomerang at 145, 245 and 345 GHz (Masi et al. 2006), which targeted a region of sky with low emission from a single dust foreground. Here the two higher frequency channels acted as foreground monitors for the 145 GHz CMB deep survey, and were used to estimate that the foreground contamination at 145 GHz was at an RMS level of less than $10\mu\text{K}$ on $11.5'$ scales (Table 10, Masi et al. (2006)). The 145 GHz CMB maps were then used for the purpose of power spectrum estimation in both temperature and polarization, after masking away a handful of compact sources (Jones et al. 2006). Masi et al. (2006) estimate that the cleanest 40% of the sky have a level of dust brightness fluctuations similar to those of the Boomerang observations, and that the cleanest 75% of the sky have brightness fluctuations less than three times larger.

An example of the high foreground level and complexity regime is available with the all-sky observations of the WMAP mission in five frequency channels from 23 to 94 GHz (Bennett et al. 2003a; Hinshaw et al. 2008). In this frequency range, the emission from at least three Galactic components (synchrotron, free-free and dust), as well as contamination by unresolved point sources must be contended with. WMAP also gives a clear example of science goal dependent data processing: CMB maps for use in non-Gaussianity tests are obtained from a noise-weighted sum of frequency maps at differing angular resolution, for which the regions most contaminated by foregrounds are masked (Komatsu et al. 2003); The analysis leading to the WMAP cosmological parameter estimation involves foreground cleaning by template subtraction; masking of the most contaminated 15% of sky; and subtracting a model of the contribution of unresolved point sources from the CMB cross power spectra (Hinshaw et al. 2003, 2007); For an improved understanding of galactic emission, the WMAP team have used a number of methods including template fits, the maximum entropy method, and the direct pixel-by-pixel fitting of an emission model (Gold et al. 2008; Dunkley et al. 2008).

1.1. Component Separation

Component separation is a catch-all term encompassing any data processing that exploits correlations in observations made at separate frequencies, as well as external constraints and physical modeling, as a means of discriminating between different physical sources of emission.

PLANCK has a number of different scientific objectives: the primary goal is a cosmological analysis of the CMB, but important secondary goals include obtaining a better understanding of the interstellar medium and Galactic emission, measurement of extragalactic sources of emission and the generation of a Sunyaev-Zeldovich (SZ) cluster catalogue. These planned objectives will lead to a set of data products which the PLANCK consortium is committed to delivering to the wider community some time after the completion of the survey. These data products include maps of the main diffuse emissions and catalogues of extragalactic sources, such as galaxies and clusters of galaxies.

In this context, it is worth remembering that PLANCK is designed to recover the CMB signal at the level of a few microkelvin per resolution element of $5'$ (and less than one mi-

crokelvin per square degree). Numbers to keep in mind are the RMS of CMB smoothed with a beam of $45'$ FWHM, which is around $70\mu\text{K}$, while white noise RMS at the same scale is around $0.7\mu\text{K}$. This level of sensitivity sets the ultimate goals for data processing—and component separation in particular—if the full scientific potential of PLANCK is to be realised. However, less stringent requirements may be acceptable for statistical analyses such as power spectrum estimation, in particular on large scales where cosmic variance dominates the error of total intensity observations.

1.2. WG2

PLANCK is designed to surpass previous CMB experiments in almost every aspect. Therefore, a complete and timely exploitation of the data will require methods that improve upon foreground removal via template subtraction and masking. The development and assessment of such methods is coordinated within the PLANCK ‘Component Separation Working Group’ (WG2). Another working group in the PLANCK collaboration, the C_ℓ temperature and polarization working group (WG3 or ‘‘CTP’’), investigates other critical data analysis steps, in particular, map-making (Poutanen et al. 2006; Ashdown et al. 2007) and power spectrum estimation.

The present paper reports the results of the WG2 activity in the framework of a *component separation challenge* using a common set of simulated PLANCK data.¹ In turn, this exercise provides valuable feedback and validation during the development of the PLANCK Sky Model.

This is the first time, within the PLANCK collaboration, that an extensive comparison of component separation methods is attempted on simulated data based on models of sky emissions of representative complexity. As will be seen and emphasised throughout this paper, this aspect is critical for a meaningful evaluation of the performance of any separation method. In this respect, the present work significantly improves on the semi-analytical estimates of foreground contamination obtained by Bouchet & Gispert (1999) for the PLANCK phase A study, as well as on previous work by Tegmark et al. (2000).

The paper is organized as follows: In Section 2 we describe the sky emission model and simulations that were used, and describe the methodology of the Challenge. In Section 3 we give an overview of the methods that have been implemented and took part in the analysis. In Section 4 we describe the results obtained for CMB component separation and power spectrum estimation. The results for point sources, SZ and Galactic components are described in Section 5 and in Section 6 we present our summary and conclusions. In Appendix A we provide a more detailed description of the methods, their implementation details and strengths and weaknesses.

2. The challenge

The objective of the component separation challenge discussed herein is to assess the readiness of the PLANCK collaboration to tackle component separation, based on the analysis of realistically complex simulations. It offers an opportunity for comparing the results from different methods and groups, as well as to develop the expertise, codes, organisation and infrastructure necessary for this task.

¹ A similar data challenge has been undertaken in the past in the context of simulated WMAP and sub-orbital CMB data (the WOMBAT challenge; Gawiser et al. 1998).

This component separation challenge is designed so as to test on realistic simulated data sets, component separation methods and algorithms in a situation as close as possible to what is expected when actual PLANCK data will be analysed. Hence, we assume the availability of a number of ancillary data sets. In particular, we assumed that six-year WMAP observations will be available as additional data sets. Although WMAP is significantly less sensitive than PLANCK, it provides very useful complementary information for the separation of low-frequency Galactic components. This section describes our simulations, the challenge setup, and the evaluation methodology.

2.1. Sky emission

Sky simulations are based on an early development version of the PLANCK Sky Model (PSM, in preparation), a flexible software package developed by PLANCK WG2 for making predictions, simulations and constrained realisations of the microwave sky.

The CMB sky is based on the observed WMAP multipoles up to $\ell = 70$, and on a Gaussian realisation assuming the WMAP best-fit C_ℓ at higher multipoles. It is the same CMB map used by Ashdown et al. (2007).

The Galactic interstellar emission is described by a three component model of the interstellar medium comprising free-free, synchrotron and dust emissions. The predictions are based on a number of sky templates which have different angular resolution. In order to simulate the sky at PLANCK resolution we have added small scale fluctuations to some of the templates. The procedure used to add small scales is the one presented in Miville-Deschênes et al. (2007) which allows to increase the fluctuation level as a function of the local brightness and therefore reproduce the non-Gaussian properties of the interstellar emission.

Free-free emission is based on the model of Dickinson et al. (2003) assuming an electronic temperature of 7000 K. The spatial structure of the emission is estimated using a H α template corrected for dust extinction. The H α map is a combination of the Southern H-Alpha Sky Survey Atlas (SHASSA) and the Wisconsin H-Alpha Mapper (WHAM). The combined map was smoothed to obtain a uniform angular resolution of 1° . For the extinction map we use the $E(B-V)$ all-sky map of Schlegel et al. (1998) which is a combination of a smoothed IRAS $100\ \mu\text{m}$ (resolution of $6.1'$) and a map at a few degrees resolution made from DIRBE data to estimate dust temperature and transform the infrared emission in extinction. As mentioned earlier, small scales were added in both templates to match the PLANCK resolution.

Synchrotron emission is based on an extrapolation of the 408 MHz map of Haslam et al. (1982) from which an estimate of the free-free emission was removed. In any direction in the sky, the spectral emission law of the synchrotron is assumed to follow a power law, $T_b^{sync} \propto \nu^\beta$. We use a pixel-dependent spectral index β derived from the ratio of the 408 MHz map and the estimate of the synchrotron emission at 23 GHz in the WMAP data obtained by Bennett et al. (2003b) using a Maximum Entropy Method technique. A limitation of this approach is that this synchrotron model also contains any ‘anomalous’ dust correlated emission seen by WMAP at 23 GHz.

The thermal emission from interstellar dust is estimated using model 7 of Finkbeiner et al. (1999a). This model, fitted to the FIRAS data (7° resolution), makes the hypothesis that each line of sight can be modeled by the sum of the emission from two dust populations, one cold and one hot. Each grain population is in thermal equilibrium with the radiation field and thus has a




Fig. 1. Synchrotron and dust (effective powerlaw) spectral indices evaluated between 30 and 44 GHz, and 143 and 217 GHz respectively (in μK_{RJ}). A spatially varying spectral index corresponds to a foreground morphology that varies with frequency.

grey-body spectrum, so that the total dust emission is modelled as

$$I_\nu \propto \sum_{i=1}^2 f_i \nu^{\beta_i} B_\nu(T_i) \quad (1)$$

where $B_\nu(T_i)$ is the Planck function at temperature T_i . In model 7 the emissivity indices are $\beta_1 = 1.5$, $\beta_2 = 2.6$, and $f_1 = 0.0309$ and $f_2 = 0.9691$. Once these values are fixed, the dust temperature of the two components is determined using only the ratio of the observations at two wavelengths, $100\ \mu\text{m}$ and $240\ \mu\text{m}$. For this purpose, we use the $100/240\ \mu\text{m}$ map ratio published by Finkbeiner et al. (1999b). Knowing the temperature and β of each dust component at a given position on the sky, we use the $100\ \mu\text{m}$ brightness at that position to scale the emission at any frequency using Eq. 1. We emphasise that the emission laws of the latter two components, synchrotron and dust, vary across the sky, as shown in Figure 1. The spectral index of free-free is taken as uniform on the sky as it only depends on the electronic temperature, taken as a constant here.

Point sources are modeled with two main categories: radio and infra-red. Simulated radio sources are based on the NVSS or SUMSS and GB6 or PMN catalogues. Measured fluxes at 1 and/or 4.85 GHz are extrapolated to PLANCK frequencies assuming a distribution in flat and steep populations. For each of these two populations, the spectral index is randomly drawn within a set of values compatible with the typical average and dispersion. Infrared sources are based on the IRAS catalogue, and modelled as dusty galaxies (Serjeant & Harrison 2005). IRAS coverage gaps were filled randomly adding sources with a flux distribution consistent with the mean counts. Fainter sources were as-




Fig. 2. Upper panel: Hit counts for the 143 GHz channel. The inhomogeneities at the ecliptic poles are characteristic of PLANCK’s cycloidal scanning strategy. Lower panel: The masking scheme separating the sky in three regions of different foreground contamination. The grey region at high Galactic latitudes is Zone 1, covering $f_{\text{sky}} = 74\%$. The darker region at lower Galactic latitudes is Zone 2 and covers $f_{\text{sky}} = 22\%$. The remaining region (green) along the Galactic ridge is Zone 3. The point source mask (red) covers 4% of Zone 1. The SZ mask (yellow) cuts detected SZ clusters at Galactic latitudes above 20 degrees, covering 1.4% of sky.

sumed to be mostly sub-millimeter bright galaxies, such as those detected by SCUBA surveys. These were modelled following Granato et al. (2004) and assumed to be strongly clustered, with a comoving clustering radius $r_0 \approx 8 h^{-1}$ Mpc. Since such sources have a very high areal density, they are not simulated individually but make up the sub-mm background.

We also include in the model a map of thermal SZ spectral distortion from galaxy clusters, based on a cluster catalogue randomly drawn using a mass-function compatible with present-day observations and with Λ CDM parameters $\Omega_m = 0.3$, $h = 0.7$ and $\sigma_8 = 0.9$ (Colafrancesco et al. 1997; de Zotti et al. 2005).

Component maps are produced at all PLANCK and WMAP central frequencies. They are then co-added and smoothed with Gaussian beams as indicated in Table 1. A total of fourteen monochromatic maps have been simulated.

Finally, inhomogeneous noise is obtained by simulating the hit counts corresponding to one year of continuous observations by PLANCK, using the Level-S simulations tool (Reinecke et al. 2006). An example of a hit count map is shown in the upper panel of Figure 2. WMAP six year hit counts, obtained from scaling up the observed WMAP three year hit count patterns, are used to generate inhomogeneous noise in the simulated WMAP observations. The RMS noise level per hit for both experiments is given in Table 1.

2.2. Challenge setup

The simulated data sets were complemented by a set of ancillary data including hitmaps and noise levels, IRAS, 408 MHz, and H α templates, as well as catalogues of known clusters from ROSAT and of known point sources from NVSS, SUMSS, GB6, PMN and IRAS.

The Challenge then proceeded first with a blind phase lasting around four months between August and November 2006, when neither the exact prescription used to simulate sky emission from these ancillary data sets, nor maps of the input components alone, were communicated to challenge participants.

After this phase and an initial review of the results at the WG2 meeting in Catania in January 2007, the Challenge moved to an open phase lasting from January to June 2007. In this phase the input data—CMB maps and power spectrum, Galactic emission maps, SZ Compton y parameter map, point source catalogues and maps, noise realisations—were made available to the participating groups.

All of the results presented here have been obtained after several iterations and improvements of the methods, both during the comparison of the results obtained independently by the various teams, and after the input data was disclosed. Hence, the challenge has permitted significant improvement of most of the methods and algorithms developed within the PLANCK collaboration. The analysis of the Challenge results was led by the simulations team, with involvement and discussion from all participating groups.

Deliverables

A set of standard deliverables were defined. These included: a CMB map with $1.7'$ pixels (Healpix $N_{\text{side}} = 2048$) together with a corresponding map of estimated errors; the effective beam F_ℓ , which describes the total smoothing of the recovered CMB map due to a combination of instrumental beams and the filtering induced by the component separation process; a set of binned CMB power spectrum estimates (band averages of $\ell(\ell + 1)C_\ell$) and error bars; maps of all the diffuse components identified in the data; catalogues of the infrared and radio sources, and SZ clusters; a map of the SZ Compton y parameter.

Masks

Different separation methods are likely to perform differently in either foreground-dominated or noise-dominated observations. Also, they may be more or less sensitive to different types of foregrounds. Since the level of foreground contamination varies strongly across the sky, we use a set of standard masks throughout this work, and they are shown in the lower panel of Figure 2.

The sky is split into three distinct Galactic ‘Zones’: Zone 1 is at high Galactic latitudes and covers 74% of sky, similar to the WMAP Kp0 mask with smoother edges and small extensions. Zone 2 is at lower Galactic latitudes and covers 22% of sky. The remaining 4% of sky covered by Zone 3, which is similar to the WMAP Kp12 mask.

The point source mask is the product of nine masks, one for each channel, each constructed by excluding a two FWHM region around every source with a flux greater than 200 mJy at this frequency. This point source mask covers 4% of sky in Zone 1. For comparison, the WMAP point source masks of Bennett et al. (2003b) excludes a radius of 0.7° around almost 700 sources with fluxes greater than 500 mJy, covering a total of 2% of sky.

Table 1. Characteristics of PLANCK one year simulations (upper) and WMAP six year simulations (lower). PLANCK and WMAP hit counts correspond to 1.7' (Healpix $n_{\text{side}}=2048$) and 6.8' ($n_{\text{side}}=512$) pixels respectively. N_ℓ is the white noise level calculated from the inhomogeneous distribution of hits.

Channel	30 GHz	44 GHz	70 GHz	100 GHz	143 GHz	217 GHz	353 GHz	545 GHz	857 GHz
FWHM [arcmin]	33	24	14	10	7.1	5	5	5	5
$\sigma_{\text{hit}} [\mu\text{K}_{\text{RJ}}]$	1030	1430	2380	1250	754	610	425	155	72
$\sigma_{\text{hit}} [\mu\text{K}_{\text{CMB}}]$	1050	1510	2700	1600	1250	1820	5470	24700	1130000
Mean; Median hits per pixel	82; 64	170; 134	579; 455	1010; 790	2260; 1790	2010; 1580	2010; 1580	503; 396	503; 396
$N_\ell^{1/2} [\mu\text{K}_{\text{CMB}}]$	0.066	0.065	0.063	0.028	0.015	0.023	0.068	0.62	28.4

Channel	23 GHz (K)	33 GHz (Ka)	41 GHz (Q)	61 GHz (V)	94 GHz (W)
FWHM [arcmin]	52.8	39.6	30.6	21	13.2
$\sigma_{\text{hit}} [\mu\text{K}_{\text{RJ}}]$	1420	1420	2100	2840	5210
$\sigma_{\text{hit}} [\mu\text{K}_{\text{CMB}}]$	1440	1460	2190	3120	6500
Mean; Median hits per pixel	878; 792	878; 790	2198; 1889	2956; 2577	8873; 7714
$N_\ell^{1/2} [\mu\text{K}_{\text{CMB}}]$	0.10	0.10	0.10	0.12	0.14

The SZ mask is constructed by blanking out small circular regions centered on 1625 SZ clusters detected with the needlet-ILC + matched filter method (see Section 5.2). For each of them, the diameter of the cut is equal to the virial radius of the corresponding cluster.

2.3. Comments about the sky emission simulations

A note of caution about these simulations of sky emission is in order. Although the PSM, as described above, has a considerable amount of sophistication, it still makes some simplifying assumptions – and cannot be expected to describe the full complexity of the real sky. This is a critical issue, as component separation methods are very sensitive to these details. We mention four of them.

First, Galactic emission is modeled with three components only, with no anomalous emission at low frequencies. This affects the spectral behavior of components in the lower frequency bands below 60 GHz where the anomalous emission is thought to be dominant (Davies et al. 2006; Bonaldi et al. 2007; Miville-Deschenes et al. 2008).

Second, even though variable spectral emission laws are used for synchrotron and dust emission, this is still an idealisation: for the synchrotron, the emission law in each pixel is described by a single spectral index without any steepening. For dust, emission is modeled as a superposition of two populations, with distinct but fixed temperature and emissivity. These approximations impact component separation, since almost perfect estimation of the relevant parameters of a given foreground emission is possible at frequencies where this foreground dominates, thereby allowing perfect subtraction in the cosmological channels.

Third, it is worth mentioning that only low resolution ($\sim 1^\circ$) templates are available for synchrotron and free-free emissions. Hence, addition of small scale power is critical: if such scales were absent from the simulations, but actually significant in the real sky, one might get a false impression that no component separation is needed on small scales. Also, the detection of point sources as well as of galaxy clusters would be significantly easier, hence not representative of the actual problem. Here, missing small scale features are simulated using a non-stationary coloured Gaussian random field. Although quite sophisticated, this process can not generate for instance, filamentary or patchy structures known to exist in the real sky.

Fourth, our simulations are somewhat idealised in the sense that we use perfect Gaussian beams, assume no systematic ef-

fects, and assume that the noise is uncorrelated from pixel to pixel and from channel to channel. Also, the effect of the finite bandpass of the frequency channels is not taken into account, and we assume that the calibration and zero levels of each channel is perfectly known.

In spite of these simplifications, component separation remains a difficult task with our simulated data because of pixel-dependent spectral emission laws of dust and synchrotron, and of the presence of more than a million point sources with different emission laws, of hundreds of thousands of unresolved or extended SZ clusters, and of significant emission from a complex IR background. It is fair to say that this simulated sky is far more complex than anything ever used in similar investigations.

In closing this Section, we show in Figure 3 the angular power spectra of the basic components for the 70 and 100 GHz channels, close to the foreground emission minimum. The spectra of CMB, noise and thermal SZ are compared to the total Galactic emission spectra evaluated at high and low Galactic latitudes, on Zone 1 and 2 respectively. The point source spectra are evaluated in Zone 1, both with and without the brightest sources above 200 mJy masked. Figure 3 shows the obvious impact on CMB studies of masking the most foreground-contaminated regions. It also indicates that there is a significant region of sky, Zone 2, for which Galactic emission and CMB power are comparable. In the following sections, results are evaluated independently in both Zones 1 and 2.

3. Outline of the methods

In this section we present a brief overview of the methods that have been used in this challenge. The section is divided in three parts, one for diffuse component separation methods, one for point source extraction, and one for SZ cluster extraction.

3.1. Diffuse component separation

The spirit of each method tested on the challenge data is outlined here. A more detailed description, including some details of their implementations and a discussion of their strengths and weaknesses is presented in Appendix A.

First we define some relevant terminology. The data model for a given channel ν is

$$d_\nu = b_\nu * x_\nu + n_\nu \quad (2)$$

Table 2. Some characteristics of the diffuse component separation methods used in the challenge

	Channels used	Components modeled	Resources and runtime
COMMANDER	WMAP, PLANCK 30–353 GHz,	CMB, dust, sync, FF, mono-,dipoles	1000 CPU hr, 2 day
CCA	PLANCK, Haslam 408 MHz	CMB, dust, sync, FF	70 CPU hr, 1.5 day
GMCA	PLANCK	CMB, SZ, sync., FF	1200 CPU hr, 6 day
FastICA	143–353 GHz	2 components (CMB and dust)	21 CPU min, 20 sec
FastMEM	PLANCK	CMB, SZ, dust, sync, FF	256 CPU hr, 8 hr
SEVEM	PLANCK	CMB	30 CPU hr, 30 hr
SMICA	PLANCK, WMAP	CMB, SZ, dust, total galaxy	8 CPU hr, 4 hr
WI-FIT	70–217 GHz	CMB	400 CPU hr, 8 hr




Fig. 3. Spectra of the simulated microwave sky components near the foreground minimum. CMB, noise with the effect of beam deconvolution, and the thermal SZ effect are evaluated on the full sky; point source power is evaluated on Zone 1+2, and on Zone 1+2+PS; the galaxy power spectra are evaluated on Zone 1 and on Zone 2. The well-known importance of masking is evident, as is the fact that there is a significant proportion of sky (Zone 2, $f_{\text{sky}} = 22\%$) for which Galactic emission is comparable to CMB power.

where d_ν , x_ν , n_ν are respectively the observation map, the sky emission map and the noise map at frequency ν while b_ν is the instrumental beam of channel ν , assumed to be Gaussian symmetric, and $*$ denotes convolution on the sphere. The sky emission itself, x_ν , is a superposition of components. Most methods assume (implicitly or explicitly), that it can be written as a linear

mixture

$$x_\nu = \sum_c A_{\nu c} s_c \quad (3)$$

where the sum runs over the components. In matrix-vector format, this reads $\mathbf{x} = \mathbf{A}\mathbf{s}$ where \mathbf{A} is referred to as the ‘mixing matrix’. Vector \mathbf{s} is the vector of components. Vectors \mathbf{d} and \mathbf{n} are defined similarly. When this model holds, Eq. (2) becomes

$$d_\nu = b_\nu * \left(\sum_c A_{\nu c} s_c \right) + n_\nu \quad (4)$$

In simple models, matrix \mathbf{A} is constant over the sky; in more complex models, it varies over patches or even from pixel to pixel.

We now briefly describe each of the methods that performed component separation of the CMB (and possibly other diffuse components), and also mention how the CMB angular power spectrum is estimated.

- **Gibbs sampling** (Commander; Eriksen et al. 2008). The approach of Commander is to fit directly an explicit parametric model of CMB, foregrounds and noise to the antenna temperature of low-resolution map pixels. For the Challenge, Commander was used to analyse the data smoothed to 3° resolution at each channel with a pixel size of $54'$ (Healpix $N_{\text{side}}=64$). For a given foreground model, Commander provides an exact foreground-marginalised CMB C_ℓ distributions using the Gibbs (conditional) sampling approach.
- **Correlated component analysis** (CCA; Bedini et al. 2005). The CCA approach starts with an estimation of the mixing matrix on patches of sky by exploiting spatial correlations in the data, supplemented by constraints from external templates and foreground scaling modeling. Estimated parameters are then used to reconstruct the components by Wiener filtering in the harmonic domain. The C_ℓ are estimated from the recovered CMB map.
- **Independent component analysis** (FastICA; Maino et al. 2002). The FastICA method is a popular approach to blind component separation. No assumptions are made about the frequency scaling or mixing matrix. Instead, assuming statistical independence between CMB and foregrounds, the mixing matrix is estimated by maximizing the non-gaussianity of the 1-point distribution function of linear combinations of input data. The inferred mixing matrix is used to invert the linear system of Eq. 4. The C_ℓ ’s are estimated from the recovered CMB map.
- **Harmonic-space maximum entropy method** (FastMEM; Hobson et al. 1998; Stolyarov et al. 2002). The FastMEM method estimates component maps given frequency scaling models and external foreground power spectra (and cross-power spectra) with adjustable prior weight. It is a non-

blind, non-linear approach to inverting Eq. 4, which assumes a maximum-entropy prior probability distribution for the underlying components. The C_ℓ 's are estimated from the recovered CMB component.

- **Generalised morphological component analysis** (GMCA; Bobin et al. 2007). Generalized Morphological Component Analysis is a semi-blind source separation method which GMCA disentangles the components by assuming that each of them is sparse in a fixed appropriate waveform dictionary such as wavelets. For the Challenge two variants of GMCA were applied, GMCA-blind optimised for separation of the CMB component, and GMCA-model optimised for separation of galactic components. The C_ℓ 's are estimated from the recovered CMB map from the GMCA-blind method.
- **Spectral estimation via expectation maximisation** (SEVEM; Martínez-González et al. 2003). SEVEM performs component separation in three steps. In a first step, an internal template subtraction is performed in order to obtain foreground-reduced CMB maps in three centre channels (100-217 GHz). Then the CMB power spectrum is estimated from these maps, via the EM algorithm, assuming a signal plus (correlated) noise model. A final CMB map is obtained using a harmonic Wiener filter on the foreground-reduced maps.
- **Spectral matching independent component analysis** (SMICA; Delabrouille et al. 2003; Cardoso et al. 2008). The SMICA method estimates model parameters using observation correlations in the harmonic domain (auto- and cross-spectra). The estimated parameters typically are some mixing coefficients and the power spectra of independent components. For the challenge, the correlations between Galactic components are taken into account. The estimated parameters are then used to Wiener-filter the observations to obtain component maps. At small scales the C_ℓ 's are one of the estimated parameters. At large scales $\ell \leq 100$ the C_ℓ 's are estimated from the CMB map.
- **Wavelet based high resolution fitting of internal templates** (WI-FIT; Hansen et al. 2006). The WI-FIT method computes CMB-free foreground + noise templates from differences of the observations in different channels, and use those to fit and subtract foregrounds from the CMB dominated channels in wavelet space. The C_ℓ 's are estimated from the recovered CMB map.

Some characteristics of these methods are summarised in Table 2, which shows the data used, the components modeled and an indication of the computational resources required.

Note that many different approaches to diffuse component separation are represented here: blind, non-blind, semi-blind; methods based on linear combinations for foreground extraction; likelihood based methods which estimate parameters of a model of the foregrounds and the CMB; a maximum entropy method; methods based on cross correlations; a method based on sparsity. Also they rely on very diverse assumptions and models.

3.2. Point source extraction

In the present challenge, point sources are detected in all PLANCK channels independently. Two methods are used, the first based on a new implementation of matched filtering, and the second using the second member of the Mexican Hat Wavelet Family of filters (González-Nuevo et al. 2006). Point sources are detected by thresholding on the filtered maps.

This corresponds to a first step for effective point source detection. It does not exploit any prior information on the position of candidate sources. Such information can be obtained from external catalogues as in López-Caniego et al. (2007), or from detections in other PLANCK channels. Neither does this approach exploit the coherence of the contaminants throughout PLANCK frequencies, nor try to detect point sources *jointly* in more than one channel. Hence, there is margin for improvement.

- **Matched Filter (MF)**: The high spatial variability of noise and foreground emission suggests using local filters (for instance on small patches). The sky is divided into 496 overlapping circular regions 12 degrees in diameter. Matched filtering is independently applied on each patch. A local estimate of the power spectrum of the background is obtained from the data themselves by averaging the power in circular frequency bins. A first pass is performed to detect and remove the brightest sources (above 20σ), to reduce the bias in background power estimation and to reduce possible artifacts in the filtered maps. The 5σ level catalogue is obtained by a second application of the whole procedure after removal of the brightest sources.
- **Mexican Hat Wavelet (MHW2)**: In a similar way, the sky is divided into 371 square patches. The size of each patch is 14.65×14.65 degrees, with a 3 degrees overlapping among patches. Each patch is then individually filtered with the MHW2. For each patch, the optimal scale of the wavelet is obtained by means of a fast maximization of the wavelet gain factor. This step requires only a straightforward estimation of the variance of the patch, excluding the border of the patch. A 5σ level catalogue is obtained by simple thresholding in a single step.

3.3. SZ cluster extraction

In the present data challenge, we address both the questions of building an SZ catalogue, and of making a map of thermal SZ emission.

SZ map: Three methods successfully produce SZ maps: ILC in harmonic space, ILC on a needlet frame, and SMICA. For ILC methods, the data are modelled as $\mathbf{d} = \mathbf{a}s + \mathbf{n}$ where \mathbf{d} is the vector of observations (nine maps here, using PLANCK data only), \mathbf{a} is the SZ spectral signature at all frequencies (a vector with nine entries) and \mathbf{n} is the noise. The ILC provides an estimator \widehat{s}_{ILC} of s using

$$\widehat{s}_{\text{ILC}} = \frac{\mathbf{a}' \widehat{\mathbf{R}}^{-1} \mathbf{d}}{\mathbf{a}' \widehat{\mathbf{R}}^{-1} \mathbf{a}} \quad (5)$$

where $\widehat{\mathbf{R}}$ is the empirical correlation of the observations, i.e. a 9×9 matrix, with entries $R_{\nu\nu'}$. In practice, the filter is implemented in bands of ℓ (ILC in harmonic space) or on subsets of needlet coefficients (ILC in needlet space). The needlet ILC adapts to the local background to recover the SZ sky.

SZ catalogue: Three main methods are used to obtain the cluster catalogue:

- The first one uses a single frequency matched filter (Melin et al. 2006) to extract clusters for the ILC needlet map.
- The second one uses SExtractor (Bertin & Arnouts 1996) to extract clusters for the ILC needlet map. Then, a single frequency matched filter is used to estimate cluster fluxes.

- The third is a Matched MultiFilter (Herranz et al. 2002), which implements cluster detections using the full set of input observations rather than an intermediate SZ map.
- This third method is implemented independently in Saclay and in Santander.

The performances of these four methods are given in Table 4. The comparison is done at the same contamination level ($\sim 10\%$), which corresponds to $S/N > 4.7$ for the ILC needlet + MF catalogue, $S/N > 3.8$ for the ILC needlet + SExtractor catalogue, $S/N > 4.3$ for the Matched Multifilter (MMF) Saclay catalogue and $S/N > 4.6$ for the MMF Santander catalogue.

This comparison is being extended to other cluster extraction methods in collaboration with the PLANCK working group ‘Clusters and Secondary Anisotropies’ (WG5). Some improvements are obtained using SExtractor as the extraction tool after the component separation step. There is still some place for other improvements in increasing the studied area to lower Galactic latitude and in combining the SZ extraction methods with CMB and Galactic extraction methods more intimately.

4. Results for CMB

We now turn to the presentation and discussion of the results of the challenge, starting with the CMB component. We evaluate performance based on residual errors at the map and spectral level, and on residual errors at the power spectrum estimation level.

The first point to be made is that all methods have produced CMB maps in Zones 1 and 2. Foreground contamination is barely visible. A small patch representative of CMB reconstruction at intermediate Galactic latitude, is shown in Figure 9. In the following, we focus on the analysis of the reconstruction error (or residual).

Since the various methods produce CMB maps at different resolutions, the recovered CMB maps are compared both against the input CMB sky smoothed only by the $1.7'$ pixels, and against a $45'$ smoothed version of it to emphasize errors at large scales.

4.1. Map-level residual errors

Maps of the CMB reconstruction error, with all maps smoothed to a common $45'$ resolution, are shown in Figure 5 for all of the methods (excluding Commander, which produced maps at 3° resolution). The remaining Galactic contamination is now visible at various levels for most methods, in particular close to regions with the strongest levels of free-free emission. There is also evidence of contamination by SZ cluster decrements, which are visible as distinct negative sources away from the Galactic plane. As can be seen, significant differences between methods exist.

- At high Galactic latitudes, at this $45'$ scale, the lowest contamination is achieved by SMICA, GMCA-BLIND and FastICA.
- In Zone 2, CCA, GMCA-MODEL, and FastMEM seem to filter out Galactic emission best while FastICA and WI-FIT are strongly contaminated.

A quantitative measure of the raw residual of the CMB map (reconstructed CMB minus unsmoothed input CMB) is provided by its RMS, calculated for 18 zonal bands, each 10 degrees wide in Galactic latitude, excluding pixels in Zone 3 and the point source mask. The results are shown in the upper panel of Figure 4. This quantity, denoted $\sigma_{\Delta\text{CMB}}$, gives a measure of

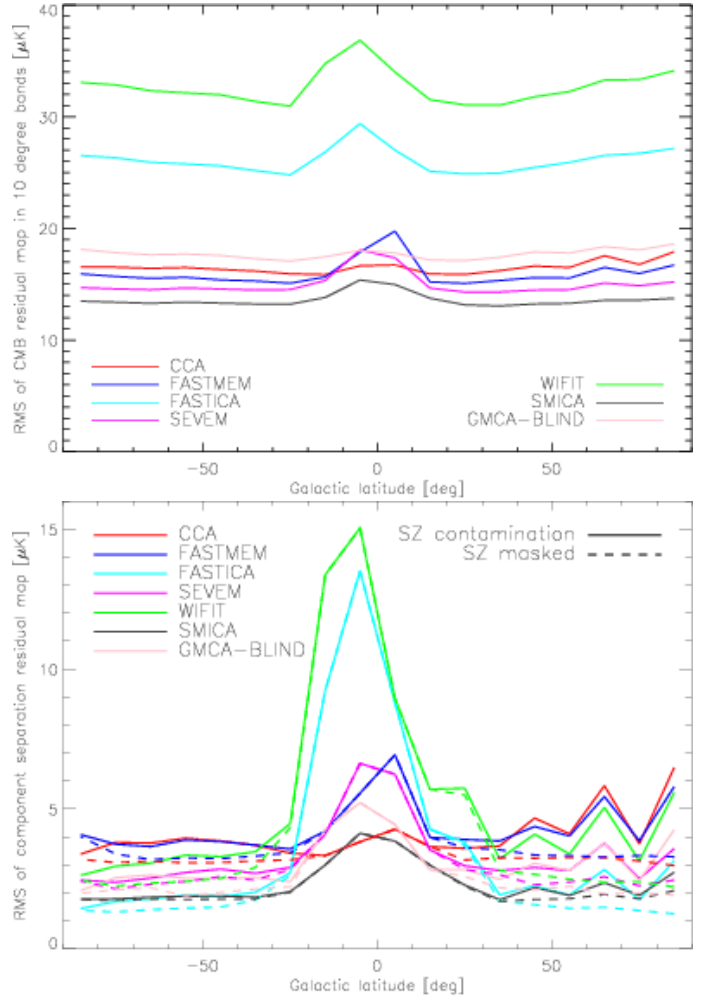


Fig. 4. (Upper) RMS of the residual error of the CMB map, calculated for each of 18 bands of 10 degrees wide in Galactic latitude. For comparison, $\sigma_{\text{CMB}} = 104.5\mu\text{K}$ and $\sigma_{\text{noise}} = 29.3\mu\text{K}$, for the 143 GHz channel ($1.7'$ pixels). (Lower) RMS of this residual map calculated at $45'$ resolution. For comparison, $\sigma_{\text{CMB}}(45') = 69.8\mu\text{K}$ and $\sigma_{\text{noise}}(45') = 0.7\mu\text{K}$ for the 143 GHz channel. The corresponding residual maps are shown in Figure 5.

the sum of the errors due to residual foreground contamination, noise, as well as from residual CMB (due to non unit response on small scales, for instance). For orientation, we can see that the ensemble of methods span the range $13\mu\text{K} < \sigma_{\Delta\text{CMB}} < 35\mu\text{K}$, which can be compared with $\sigma_{\text{CMB}} = 104.5\mu\text{K}$ and $\sigma_{\text{noise}} = 29.3\mu\text{K}$, for the 143 GHz channel.

Similarly, the lower panel of Figure 4 shows the RMS of the smoothed residual errors shown in Figure 5. Depending on the method, the typical level of foreground contamination (plus noise) has an RMS from 2 to $5\mu\text{K}$ on this smoothing scale. For comparison, $\sigma_{\text{CMB}}(45') = 69.8\mu\text{K}$ and $\sigma_{\text{noise}}(45') = 0.7\mu\text{K}$ for the 143 GHz channel.

4.2. Spectral residual errors

Next we calculate the spectra of the CMB raw residual maps, both on Zone 1 and Zone 2 (high and low Galactic latitudes), masked for the brightest point sources. Results are shown in Figure 6.

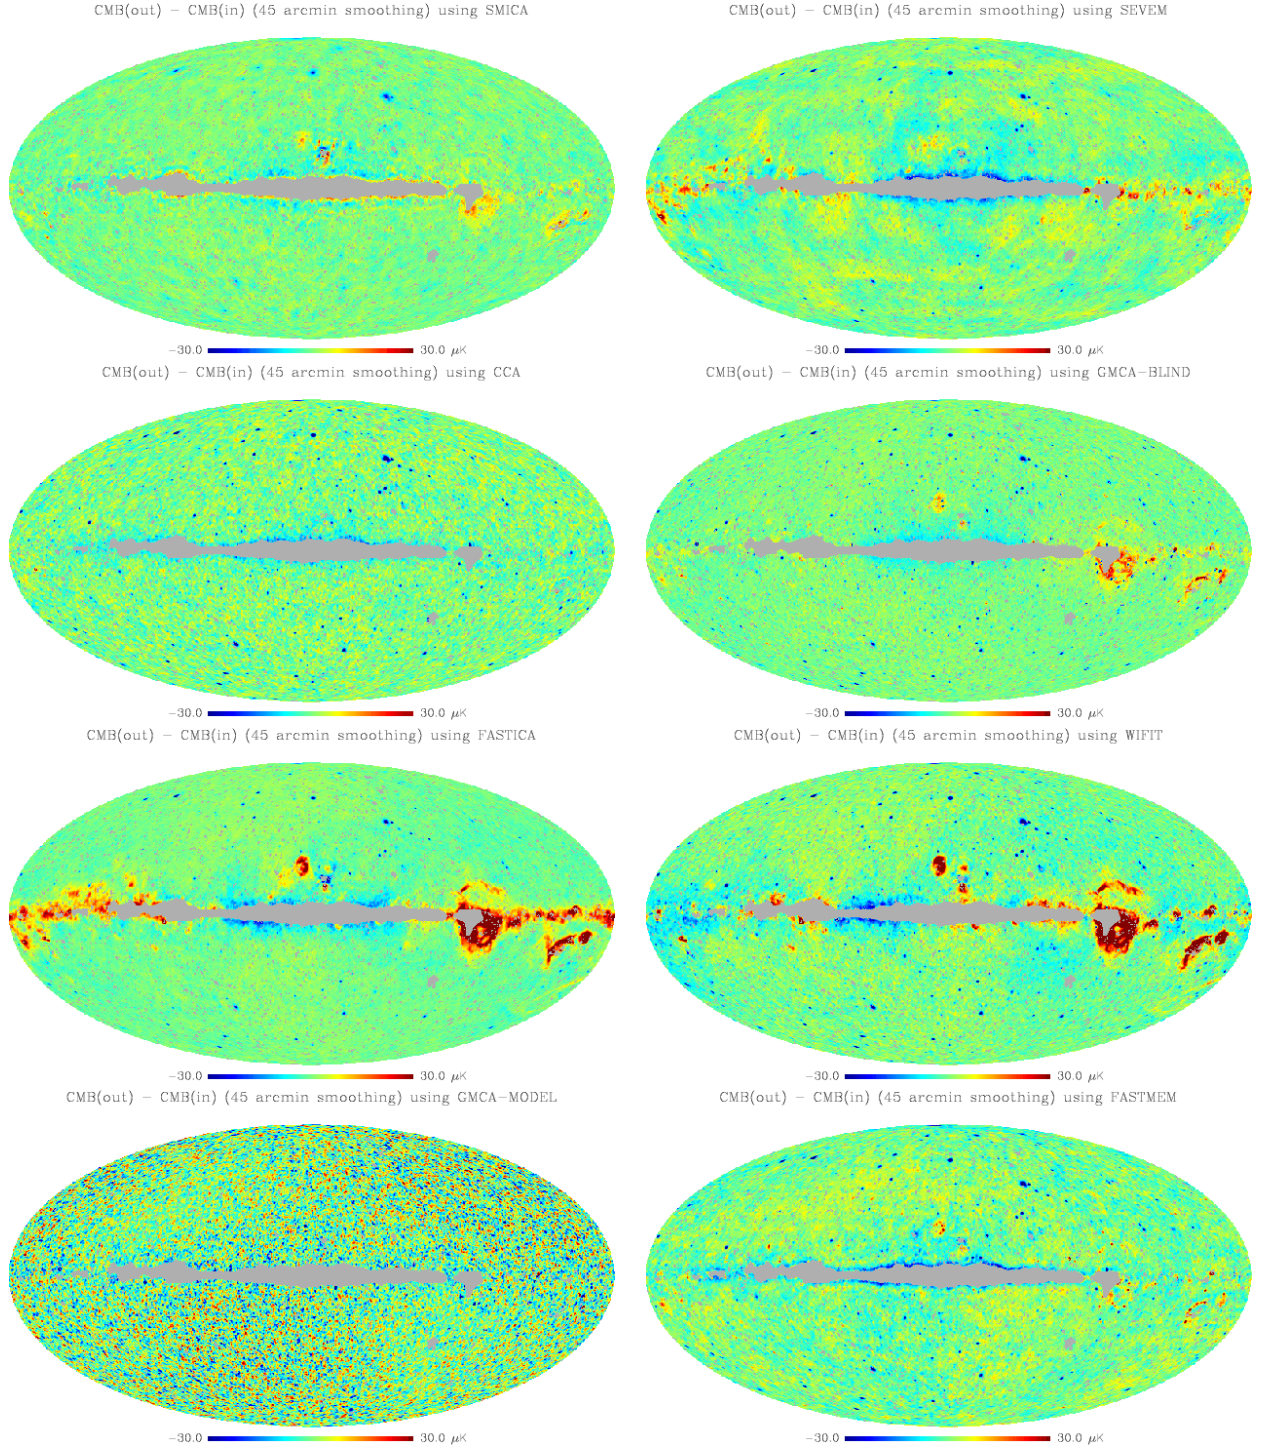


Fig. 5. CMB reconstruction error smoothed at 45' resolution. These maps are described in Section 4.1, and their RMS in Galactic latitude strips of 10° are shown in Figure 4.

Comparing the spectra of the residuals with the original level of diffuse foreground contamination shown in Figure 3, we see that a considerable degree of diffuse foreground cleaning has been attained. There seems however to be a ‘floor’ approached by the ensemble of methods, with a spread of about a factor of ten indicating differences in performance. This floor appears to be mostly free of residual CMB signal which would be visible as acoustic oscillations.

Its overall shape is not white: at high Galactic latitudes the residual spectra bottom out at very roughly $A = 0.015 \times \ell^{-0.7} \mu\text{K}^2$, while at low Galactic latitudes the spectra bottom out at $A = 0.02 \times \ell^{-0.9} \mu\text{K}^2$. This limit to the level of residuals is considerably higher than the ‘foreground-free’ noise limit displayed as a dashed line.

It is, however, also significantly lower than the CMB cosmic variance, even with 10% binning in ℓ . This comforts us in the impression that component separation is effective enough

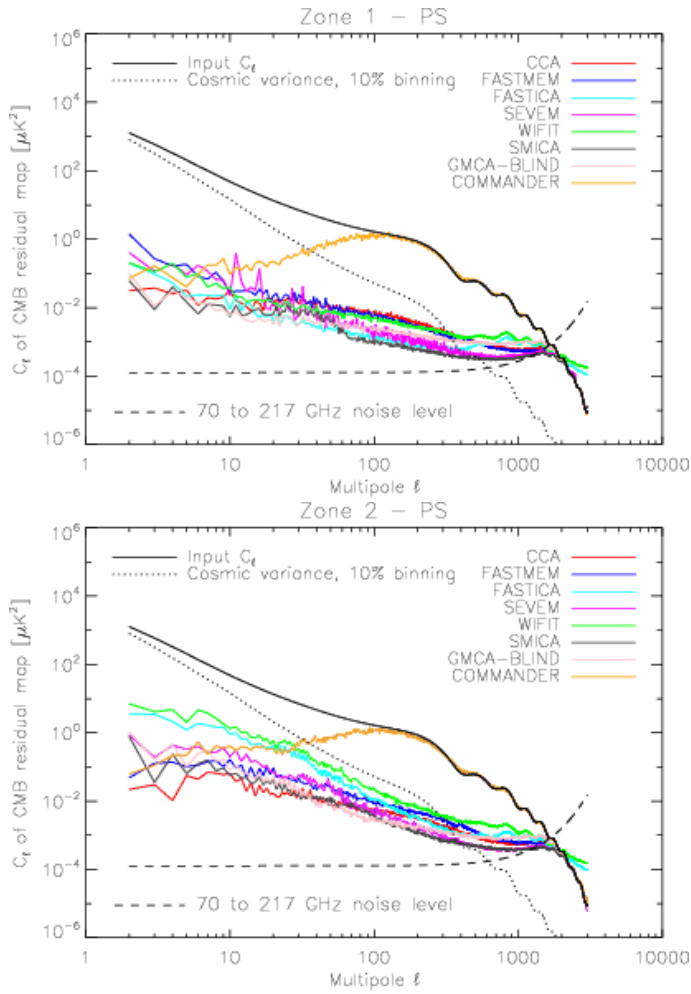


Fig. 6. Spectra of the CMB residual maps, evaluated on Zone 1 (high Galactic latitudes) and Zone 2 (low Galactic latitudes), both regions with point sources masked. Comparison with Figure 3 shows the extent to which the Galactic contamination has been removed from the CMB on large angular scales.

for CMB power spectrum estimation (discussed next in this paper), although it may remain a limiting issue for other type of CMB science. In particular, it suggests that the component separation residuals, with these channels and the present methods, will dominate the error in PLANCK CMB maps.

Recently Huffenberger et al. (2007) performed a reanalysis of the impact of unresolved point source power in the WMAP three-year data. They found that cosmological parameter constraints are sensitive to the treatment of the unresolved point source power spectrum beyond $\ell = 200$ characterised by a white noise level of $A = 0.015 \pm 0.005 \mu\text{K}^2$. By comparison, the residual foreground contamination obtained in our simulations is as low as $4 \times 10^{-4} \mu\text{K}^2$ at $\ell = 200$.

4.3. Power spectrum estimation errors

Although not the main focus of effort for the Challenge, each group provided their own bandpower estimates of the CMB power spectrum, which in many cases showed obvious acoustic structure out to the sixth or seventh acoustic peak at $\ell \sim 2000$. As an illustration of this result, we show in the upper and middle panels of Figure 7 the power spectrum estimates from the Commander and SMICA methods respectively.

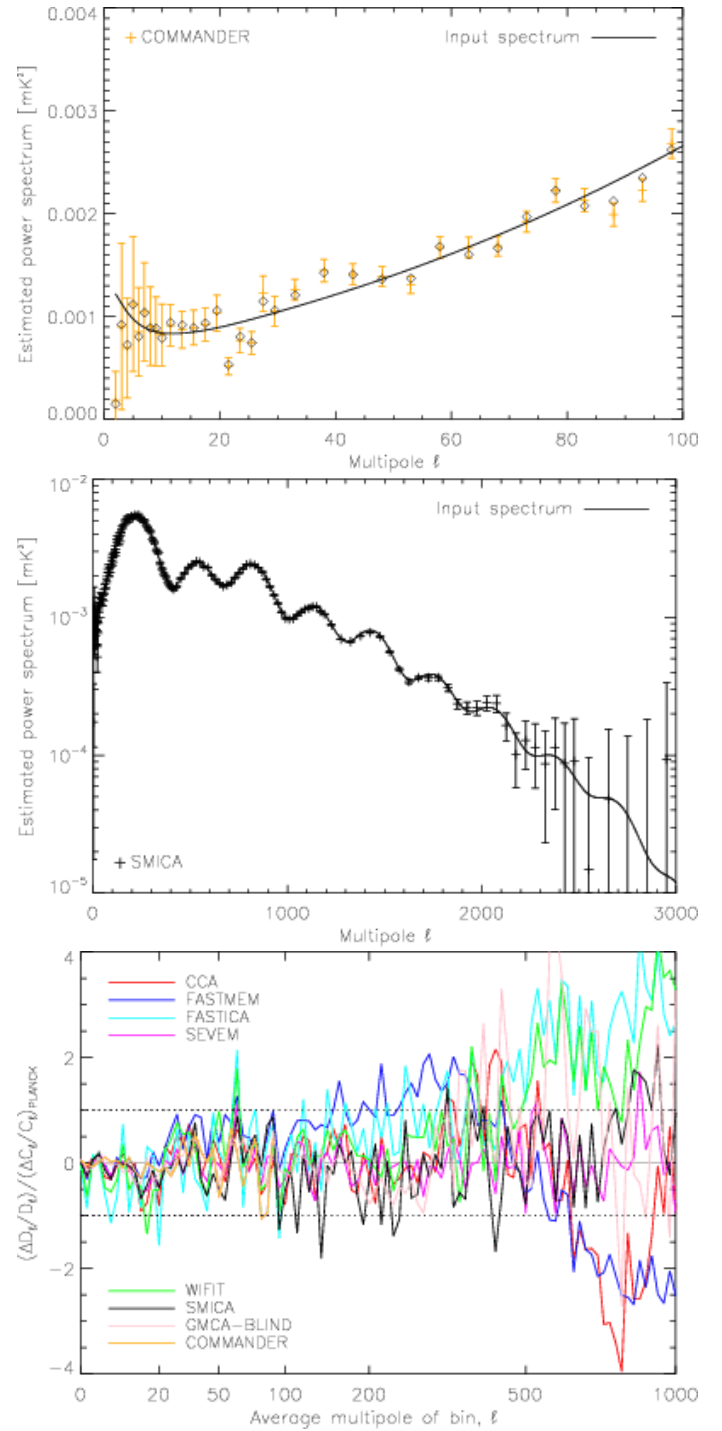


Fig. 7. (Upper) Power spectrum estimates (PSE) using Commander on large angular scales. The diamonds show the C_ℓ of the input CMB realisation. (Middle) PSE of the recovered CMB map using the SMICA method. (Lower) PSE compared with the estimates derived from the input CMB C_ℓ , and with the expected PLANCK sensitivity, assuming $f_{\text{sky}} = 0.8$. Beyond $\ell = 500$ biases in the PSE set in in some of the methods.

To make a quantitative estimate of the accuracy of the power spectrum estimates D_ℓ of $\ell(\ell + 1)C_\ell$, we calculate the quantity

$$\text{FoM} = \frac{\Delta D_\ell / D_\ell}{\Delta C_\ell / C_\ell} \quad (6)$$

Table 3. Results of point source detection on the present data challenge.

Channel	30 GHz	44 GHz	70 GHz	100 GHz	143 GHz	217 GHz	353 GHz	545 GHz	857 GHz
method	MF	MF	MF	MF	MF	MF	MF	MHW	MHW
flux limit (mJy, 5% cont.)	420	430	360	220	130	100	190	500	880
# detections	655	591	623	1103	2264	2597	1994	1233	2156

where ΔD_ℓ is the bias in the PSE compared to the PSE derived from the input CMB sky, and where $\Delta C_\ell/C_\ell$ is the expected accuracy of PLANCK, obtained from Eq. (7) below. This figure of merit penalises biases in the power spectrum estimates without taking into account the error bars claimed by each group.

In the absence of foregrounds, an approximate lower bound on the relative standard deviation in estimating the power spectrum is:

$$\frac{\Delta C_\ell}{C_\ell} \approx \sqrt{\frac{2}{N_{\text{modes}}}} \left(1 + \frac{\bar{N}_\ell}{C_\ell}\right), \quad (7)$$

where the number N_{modes} of available modes is:

$$N_{\text{modes}} = f_{\text{sky}} \sum_{\ell=\ell_{\text{min}}}^{\ell_{\text{max}}} (2\ell + 1) \quad (8)$$

where f_{sky} denotes the fraction of sky coverage. The average noise power spectrum, \bar{N}_ℓ , is obtained from the noise power spectra of the different channels

$$\bar{N}_\ell = \left(\sum_{\nu=1}^{N_{\text{chan}}} \frac{B_{\nu\ell}^2}{N_{\nu\ell}} \right)^{-1}, \quad (9)$$

$$N_{\nu\ell} = \frac{4\pi\sigma_{\text{hit}}^2}{n_{\text{pix}}^2} \sum_p \frac{1}{n_{\text{hit}}(p)}, \quad (10)$$

where $B_{\nu\ell}$ is the beam window function for channel ν , and using the calculated values of $N_{\nu\ell}$ given in Table 1. This theoretical limit Eq. (7) is used below to assess the impact of foregrounds on power spectrum estimation, taking the 70 to 217GHz channels and assuming the noise levels from Table 1 together with an $f_{\text{sky}} = 0.8$, and is displayed in the lower panel of Figure 7.

Ideally, the figure of merit given by Eq. (6) should be much less than one in the cosmic-variance limited regime (i.e. for $\ell \leq 500$ according to Figure 6). Significant deviations from zero at low ℓ and over ± 1 at high ℓ are indications of significant departures from optimality.

Focusing first on the range $\ell < 20$ we can discern the best performance from Commander, which models the spatial variation of the foreground spectral indices, thus improving the subtraction of foregrounds on large scales. On the range $20 < \ell < 500$, SEVEM, specifically designed for an estimate of the CMB power spectrum, performs best among the methods tested on this challenge. Beyond $\ell = 500$ we see the best performance from SEVEM and SMICA.

5. Results for other components

5.1. Point sources

Additional efforts have been directed towards producing a catalogue of unresolved galaxies, a catalogue of SZ clusters and maps of thermal SZ effect and of Galactic components.

The detection of point sources is both an objective of PLANCK component separation (for the production of the PLANCK early release compact source catalogue (ERCSC) and of the final point source catalogue), and also a necessity for CMB science, to evaluate and subtract the contamination of CMB maps and power spectra by this population of astrophysical objects (Wright et al. 2008; González-Nuevo et al. 2008).

The two methods have advantages and drawbacks. In principle, the matched filter is the optimal linear filter. However, it often suffers from inaccurate estimation of the required correlation matrix of the contaminants, and from the difficulty to adapt the filter to the local contamination conditions. On the data from the present challenge, this resulted in excessive contamination of the point source catalogue by small scale dust emission. The Mexican hat wavelet is not optimal, and has proved to be less effective than the matched filter in most of the channels, but is less vulnerable to dust.

On our data, the best catalogues are obtained with the MF method described above between 30 and 353 GHz, and with the MHW at 545 and 857 GHz. Table 3 summarises the PS detection achieved by these methods.

It should be noted that the five sigma detection limit, for all channels, is somewhat above what would be expected from (unfiltered) noise alone (by a factor 1.33 for the best case, 44 GHz, to 4.8 for the worst case, 857 GHz). This is essentially due to the impact of other foregrounds and of the CMB, as well as with the confusion with other sources. In particular, this effect is more evident at 545 and 857 GHz, due to high dust contamination but also to the confusion with the highly correlated population of SCUBA sources (Granato et al. 2004; Negrello et al. 2004), which constitute a contaminant whose impact on point source detection was until now somewhat underestimated.

The number of detections for each frequency channel in Table 3 has been compared to the predictions made by López-Caniego et al. (2006), properly rescaled for our sky coverage. In general, there is a good agreement between the predictions and the results of this exercise, except for the 857 GHz channel, where the number of detections is roughly half the predicted one. Again, the difference may be due to the confusion of correlated infrared sources, that are now present in the PSM but were not considered by López-Caniego et al. (2006).

5.2. SZ effect

The recovery of an SZ map from the challenge data is illustrated in Figure 8. The recovered full sky SZ is obtained by Wiener-filtering in harmonic space the needlet ILC map of the SZ y parameter. Wiener filtering enhances the visibility of SZ clusters. We clearly identify by eye the brightest clusters in the map.

One of the main results of this study is the recovery of around 2300 clusters. This is significantly lower than the performance one could expect if the main limitation was the nominal PLANCK noise, and if most detectable clusters were unresolved. Many of the recovered clusters are in fact resolved, and thus emit on scales where the contamination from CMB is not neg-

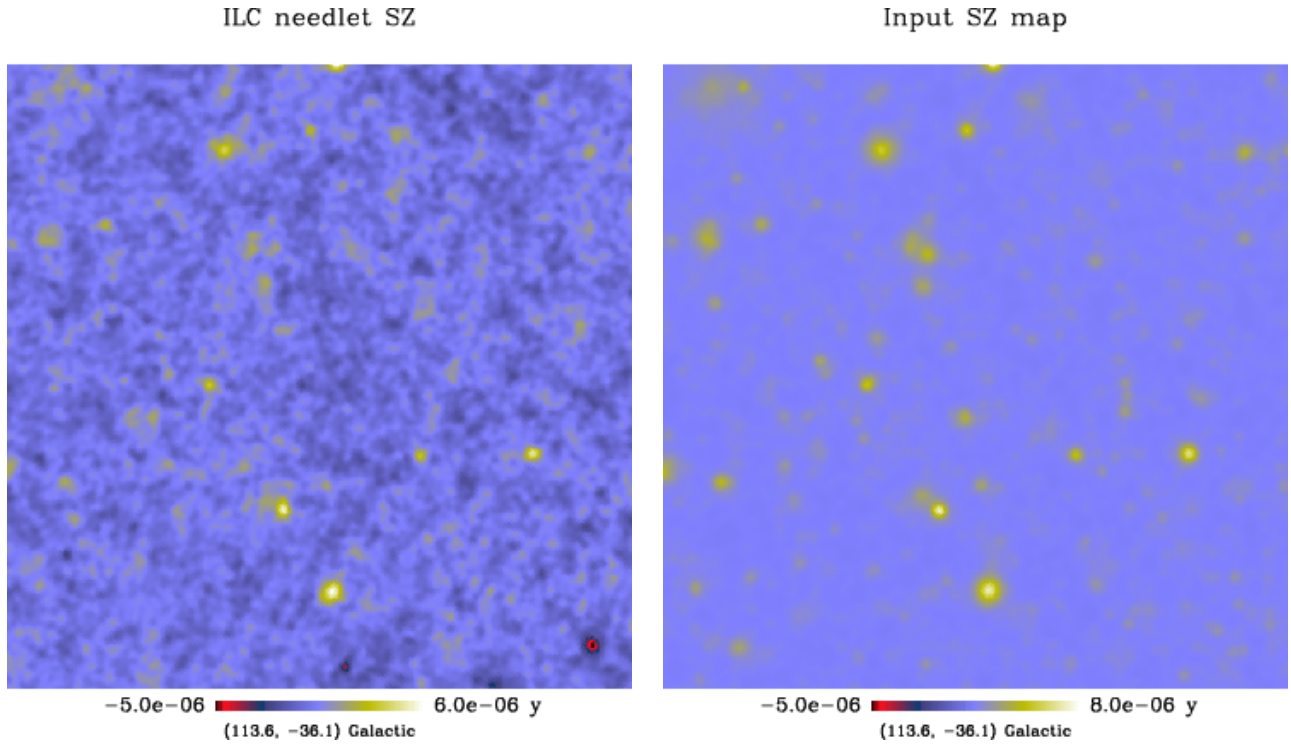


Fig. 8. Patch of the recovered ILC needlet SZ map and input SZ map. For easier comparison of the two maps ($12.5 \text{ deg} \times 12.5 \text{ deg}$), the input SZ map has been filtered to the same resolution as the output.

Table 4. Performance of the SZ cluster detection methods

The table gives the absolute number of detection for $|b| > 20^\circ$. The ILC + SExtractor method gives the highest completeness, at 91% reliability. The ILC needlet + Matched Filter (MF) performs as well as the Matched Multifilter (MMF). Some different implementation of the MMF leads to slightly different results but performances are comparable.

Method	Detections	False	Reliability
ILC needlet + SExt.	2564	225	91%
ILC needlet + MF	1804	179	90%
MMF Saclay	1803	178	90%
MMF IFCA	1535	144	91%

ligible. Small scale Galactic emissions and the background of extragalactic sources, now included in the simulations, further complicate the detection. Further study is necessary to find the exact origin of the lack of performance, and improve detection methods accordingly.

Actual detection performances, limited to 67% of the sky at Galactic latitudes above 20 degrees, are shown in Table 4. The ILC + SExtractor method gives the best result. The ILC+MF approach performs as well as the matched multifilter here. The two implementations of the MMF perform similarly. The difference in the number of detection achieved (about 13.5%), however, suggests that implementation details are important for this task.

Using the detected cluster catalogue obtained with MMF, we have produced a mask of the detected SZ clusters. For each of the 1625 clusters we mask a region of radius given by the corresponding input cluster size given by the virial radius, i.e. ten times the core radius here.

5.3. Galactic components

For the Challenge, a number of methods are applied for separating Galactic components. Table 2 lists which Galactic components were separated by the different methods. Five groups have attempted to separate a high frequency dust-like component. Four groups have attempted separation of synchrotron and free-free at low frequencies.

We compared the reconstructed component maps with their counterpart input maps, both in terms of the absolute residual error and in terms of the relative error. Both these measures are computed after removing the best fit monopole and dipole from the residual error map (fitted when excluding a region $\pm 30^\circ$ in Galactic latitude). We then defined a figure of merit $f_{20\%}$, which corresponds to the fraction of sky where the foreground amplitude is reconstructed with a relative error of less than 20%. The main results can be summarised as follows: The dust component was the best reconstructed component with $f_{20\%} \approx 0.7$ for all methods. The relative error typically becomes largest at the higher galactic latitudes where the dust emission is faintest. Synchrotron is reconstructed with $f_{20\%} \approx 0.3-0.5$. Here Commander achieved the best results at 3° resolution, but with noticeable errors along the galactic ridge where, in our simulations, the synchrotron spectral index flattens off. Free-free emission is detected and identified in regions such as the Gum Nebula, Orion A and B, and the Ophiucus. However, the reconstruction of the free-free emission at low Galactic latitudes needs improving. On the other hand, the *total* Galactic emission (free-free plus synchrotron) at low-frequencies is better reconstructed, with $f_{20\%} \approx 0.5-0.8$, with the best results from Commander.

In Figure 9, we show for illustration the recovered total Galactic emission at 23GHz from Commander, the dust emission at 143GHz from FastICA and, for comparison, the recovered CMB from SMICA on the same patch.

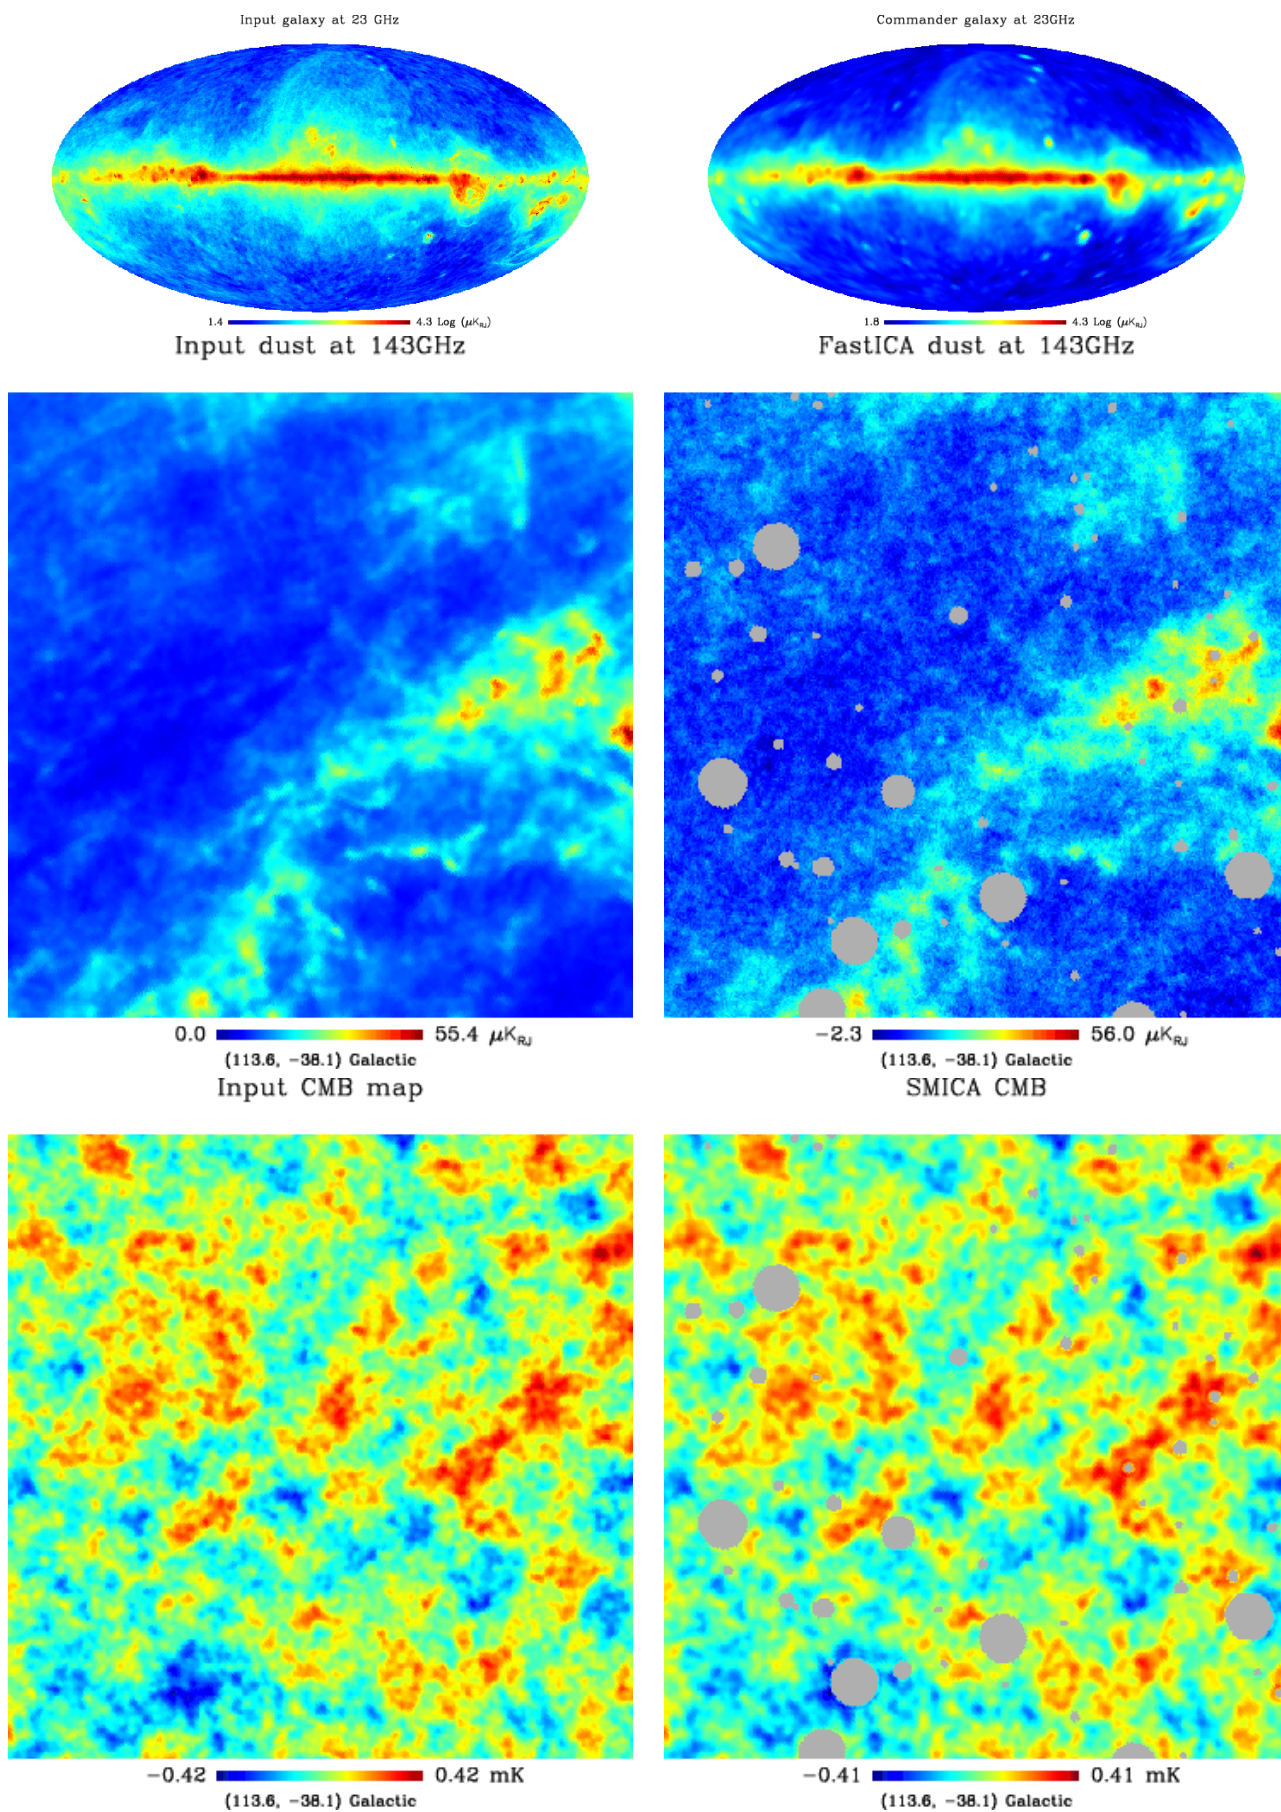


Fig. 9. Example input and recovered total galaxy emission at 23GHz, dust at 143GHz and CMB components.

6. Summary and conclusions

In this paper we have described a CMB component separation Challenge based on realistic simulations of the PLANCK satellite mission. The simulated data were based on a development version of the PLANCK Sky Model, and included the foreground emission from a three component Galactic model of free-free, synchrotron and dust, as well as radio and infra-red sources, the infra-red background, the SZ effect and PLANCK-like inhomogeneous noise. We caution that the simulations, while complex, still relied on some simplifying assumptions. Thus, there is no guarantee that the priors and data models that yielded the best separation on simulations will work equally well on real data. In other words, the amplitudes of the residuals shown in Figures 6 and 7, which indicate a difference of performance on the particular simulations used here, are not necessarily a very reliable measure of the quality of the methods or their performance on real data.

As a combined set of tools, component separation methods developed and tested in this work offer very different ways to address the component separation problem. Comparable performance between different tools, when achieved, provides confidence in the conclusions of this work against some of the simplifications used in the model used for simulating sky emission.

We found that the recovered CMB maps were clean on large scales, in the sense that the RMS of the residual contamination was much less than the cosmic variance: at best the RMS residual of the cleaned CMB maps was of the order of $2\mu\text{K}$ across the sky on a smoothing scale of $45'$, with a spectral distribution described by $A = 0.015 \times \ell^{-0.7} \mu\text{K}^2$ and $A = 0.02 \times \ell^{-0.9} \mu\text{K}^2$ at high and low Galactic latitudes respectively. The effectiveness of the foreground removal is illustrated by the comparison of the input foreground power spectra of Figure 3 with the residuals shown in Figure 6. The two panels of the latter figure show that, with few exceptions, the methods manage to clean the low Galactic latitude Zone 2, where the foreground contamination is high, almost at the same level as they do for the high Galactic latitude Zone 1, where the CMB dominates at the frequencies near the foreground minimum. The amplitude of the power spectrum of residuals is, on the largest scales, four orders of magnitude lower than that of the input Galaxy power spectrum at the foreground minimum. This means that the CMB map could be recovered, at least by some methods, over the whole sky except for a sky cut, at the 5% percent level (see Fig. 4). The CMB power spectrum was accurately recovered up to the sixth peak.

As detailed by Table 2, the outputs of the methods were diverse. While all have produced a CMB map, only a subset of them were used to obtain maps of individual diffuse Galactic emissions. Five (Commander, CCA, FastICA, FastMEM, SMICA) reconstructed thermal dust emission maps at high frequencies, and another five (Commander, CCA, FastMEM, GMCA, SMICA) yielded a map of the low-frequency Galactic emissions (synchrotron and free-free).

It is not surprising that the dust component was more easily reconstructed because it is mapped over a larger frequency range, and benefits from observations at high frequencies where it dominates over all other emissions, except the IR background at high Galactic latitudes. Moreover at high frequencies the noise level is lower and the angular resolution is better. Low frequency Galactic foregrounds suffer from more confusion, with a mixture of several components observed in only few channels, at lower resolution.

The relative errors of the reconstructed foreground maps are larger at high Galactic latitudes where the foregrounds are

fainter. We have defined a figure of merit $f_{20\%}$, which corresponds to the fraction of sky where the dust amplitude has been reconstructed with a relative error of less than 20%. For most methods, $f_{20\%} \simeq 70\%$ in the case of dust, while for the radio emission $f_{20\%} \simeq 50\%$, increasing to 80% if component separation is performed at a relatively low resolution of 3° . Clearly, there is ample room for, and need of improvement in this area.

The flux limits for extragalactic point source detection are minimum at 143 and 217 GHz, where they reach $\simeq 100$ mJy. About 1000 radio sources and about 2600 far-IR sources are detected over about 67% of the sky ($|b| > 20^\circ$). Over the same region of the sky, the best methods recover about 2300 clusters.

Work in progress includes an upgrade of the sky model, including an anomalous emission component and polarization. Also, we are in the process of integrating point source and SZ extraction algorithms together with those doing the separation of diffuse components in a single component separation pipeline. This is expected, on one side, to decrease the contamination of CMB maps on small angular scales, where point and compact sources (including SZ effects) dominate and, on the other side, to achieve a more efficient point and compact source extraction.

Acknowledgements. The work reported in this paper was carried out by Working Group 2 of the PLANCK Collaboration. PLANCK is a mission of the European Space Agency. The Italian group were supported in part by ASI (contract Planck LFI Phase E2 Activity). The US PLANCK Project is supported by the NASA Science Mission Directorate. The ADAMIS team at APC has been partly supported by the Astro-Map and Cosmostat ACI grants of the French ministry of research, for the development of innovative CMB data analysis methods. The IFCA team acknowledges partial financial support from the Spanish project AYA2007-68058-C03-02 This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. DH, JLS and ES would like to acknowledge partial financial support from joint CNR-CSIC research project 2006-IT-0037. RS acknowledges support of the EC Marie Curie IR Grant (mirg-ct-2006-036614). CB was partly supported by the NASA LTSA Grant NNG04CG90G. SML would like to thank the kind hospitality of the APC, Paris for a research visit where part of this work was completed. We thank Carlo Burigana for a careful reading of the manuscript. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science. The results in this paper have been derived using the HEALPix package (Górski et al. 2005).

References

- Ashdown, M. A. J., Baccigalupi, C., Balbi, A., et al. 2007, *A&A*, 467, 761
 Bedini, L., Herranz, D., Salerno, E., et al. 2005, *EURASIP Journal on Applied Signal Processing*, 2005, No. 15, 2400
 Bedini, L. & Salerno, E. 2007, *Lecture Notes in Artificial Intelligence*, 4694, 9
 Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003a, *ApJS*, 148, 1
 Bennett, C. L., Hill, R. S., Hinshaw, G., et al. 2003b, *ApJS*, 148, 97
 Bertin, E. & Arnouts, S. 1996, *A&AS*, 117, 393
 Bobin, J., Starck, J.-L., & Fadili, J. 2007, *IEEE Transactions on Image Processing*, In Press
 Bonaldi, A., Bedini, L., Salerno, E., Baccigalupi, C., & de Zotti, G. 2006, *MNRAS*, 373, 271
 Bonaldi, A., Ricciardi, S., Leach, S., et al. 2007, *MNRAS*, 382, 1791
 Bouchet, F. R. & Gispert, R. 1999, *New Astron.*, 4, 443
 Cardoso, J.-F., Martin, M., Delabrouille, J., Betoule, M., & Patanchon, G. 2008, *ArXiv e-prints*, 803
 Colafrancesco, S., Mazzotta, P., Rephaeli, Y., & Vittorio, N. 1997, *ApJ*, 479, 1
 Davies, R. D., Dickinson, C., Banday, A. J., et al. 2006, *MNRAS*, 370, 1125
 de Zotti, G., Ricci, R., Mesa, D., et al. 2005, *A&A*, 431, 893
 Delabrouille, J., Cardoso, J.-F., & Patanchon, G. 2003, *MNRAS*, 346, 1089
 Dickinson, C., Davies, R. D., & Davis, R. J. 2003, *MNRAS*, 341, 369
 Dunkley, J., Komatsu, E., Nolta, M. R., et al. 2008, *ArXiv e-prints*, 803
 Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, 676, 10
 Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999a, *ApJ*, 524, 867
 Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999b, *ApJ*, 524, 867
 Gawiser, E., Finkbeiner, D., Jaffe, A., et al. 1998
 Gold, B., Bennett, C. L., Hill, R. S., et al. 2008, *ArXiv e-prints*, 803
 González-Nuevo, J., Argüeso, F., López-Cañiego, M., et al. 2006, *MNRAS*, 369, 1603

González-Nuevo, J., Massardi, M., Argüeso, F., et al. 2008, MNRAS, 384, 711
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
 Granato, G. L., De Zotti, G., Silva, L., Bressan, A., & Danese, L. 2004, ApJ, 600, 580
 Hansen, F. K., Banday, A. J., Eriksen, H. K., Gorski, K. M., & Lilje, P. B. 2006, Astrophys. J., 648, 784
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
 Herranz, D., Sanz, J. L., Hobson, M. P., et al. 2002, MNRAS, 336, 1057
 Hinshaw, G., Nolta, M. R., Bennett, C. L., et al. 2007, ApJS, 170, 288
 Hinshaw, G., Spergel, D. N., Verde, L., et al. 2003, ApJS, 148, 135
 Kinshaw, G., Weiland, J. L., Hill, R. S., et al. 2003, ArXiv e-prints, 803
 Hobson, M. P., Barreiro, R. B., Toffolatti, L., et al. 1999, MNRAS, 306, 232
 Hobson, M. P., Jones, A. W., Lasenby, A. N., & Bouchet, F. R. 1998, MNRAS, 300, 1
 Huffenberger, K. M., Eriksen, H. K., Hansen, F. K., Banday, A. J., & Gorski, K. M. 2007
 Hyvarinen, A. 1999, IEEE Signal Processing Lett. 6, 6, 145
 Jones, W. C., Ade, P. A. R., Bock, J. J., et al. 2006, ApJ, 647, 823
 Komatsu, E., Kogut, A., Nolta, M. R., et al. 2003, ApJS, 148, 119
 López-Cañiego, M., González-Nuevo, J., Herranz, D., et al. 2007, ApJS, 170, 108
 López-Cañiego, M., Herranz, D., González-Nuevo, J., et al. 2006, MNRAS, 370, 2047
 Maino, D., Donzelli, S., Banday, A. J., Stivoli, F., & Baccigalupi, C. 2007, MNRAS, 374, 1207
 Maino, D., Farusi, A., Baccigalupi, C., et al. 2002, MNRAS, 334, 53
 Martínez-González, E., Diego, J. M., Vielva, P., & Silk, J. 2003, Mon. Not. Roy. Astron. Soc., 345, 1101
 Masi, S., Ade, P. A. R., Bock, J. J., et al. 2006, A&A, 458, 687
 Melin, J.-B., Bartlett, J. G., & Delabrouille, J. 2006, A&A, 459, 341
 Miville-Deschenes, M., Ysard, N., Lavabre, A., et al. 2008, ArXiv e-prints, 802
 Miville-Deschenes, M.-A., Lagache, G., Boulanger, F., & Puget, J.-L. 2007, A&A, 469, 595
 Negrello, M., Magliocchetti, M., Moscardini, L., et al. 2004, MNRAS, 352, 493
 Poutanen, T., de Gasperis, G., Hivon, E., et al. 2006, Astron. Astrophys., 449, 1311
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in FORTRAN. The art of scientific computing (Cambridge: University Press, —c1992, 2nd ed.)
 Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Enßlin, T. A. 2006, A&A, 445, 373
 Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525
 Serjeant, S. & Harrison, D. 2005, MNRAS, 356, 192
 Stolyarov, V., Hobson, M. P., Ashdown, M. A. J., & Lasenby, A. N. 2002, Mon. Not. Roy. Astron. Soc., 336, 97
 Stolyarov, V., Hobson, M. P., Lasenby, A. N., & Barreiro, R. B. 2005, Mon. Not. Roy. Astron. Soc., 357, 145
 Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. 2000, ApJ, 530, 133
 The Planck Collaboration. 2005, ESA-SCI(2005), ArXiv Astrophysics e-prints
 Wright, E. L., Chen, X., Odgaard, N., et al. 2008, ArXiv e-prints, 803

Appendix A: Description of methods

A.1. Commander

‘Commander’ is an implementation of the CMB and foregrounds Gibbs sampler most recently described by Eriksen et al. (2008); This algorithm maps out the joint CMB-foreground probability distribution, or ‘posterior’, by sampling. The target posterior may be written in terms of the likelihood and prior using Bayes’ theorem,

$$\Pr(s, C_\ell, \theta_{\text{fg}} | \mathbf{d}) = \mathcal{L}(\mathbf{d} | s, \theta_{\text{fg}}) \Pr(s | C_\ell) \Pr(C_\ell) \Pr(\theta_{\text{fg}}) \quad (\text{A.1})$$

Here θ_{fg} is the collection of all parameters required to describe the non-cosmological foregrounds. Since the noise is assumed to be Gaussian, the likelihood is simply given by the χ^2 .

In the current analysis, the foregrounds are modeled by a sum of synchrotron, free-free and thermal dust emission, and free monopole and dipoles at each frequency band. The thermal dust component is approximated by a single-component modified blackbody with a fixed dust temperature $T_d = 21\text{K}$. Thus,

the total foreground model reads

$$\begin{aligned} s_{\text{fg}}(\nu, p) = & A_s(p)g(\nu) \left(\frac{\nu}{\nu_s}\right)^{\beta_s(p)} + A_f(p)g(\nu) \left(\frac{\nu}{\nu_f}\right)^{-2.15} \\ & + A_d(p)g(\nu) \frac{e^{h\nu_d/kT_d} - 1}{e^{h\nu/kT_d} - 1} \left(\frac{\nu}{\nu_d}\right)^{\beta_d(p)+1} \\ & + m_\nu^0 + \sum_{i=1}^3 m_\nu^i(\hat{\mathbf{n}}(p) \cdot \hat{\mathbf{e}}_i), \end{aligned} \quad (\text{A.2})$$

where $g(\nu)$ is the conversion factor between antenna and thermodynamic temperatures, and $\hat{\mathbf{n}}$ is the unit vector of pixel p . The free parameters are thus the foreground amplitudes, A_s , A_f and A_d , and spectral indices, β_s and β_d , for each pixel, and the overall monopole, m_ν^0 , and dipole amplitudes, m_ν^i , for each band. For priors, we adopt the product of the Jeffreys’ ignorance prior and an informative Gaussian prior ($\beta_s = -3 \pm 0.3$ for synchrotron and $\beta_d = 1.5 \pm 0.3$ for dust) for the spectral indices, while no constraints are imposed on the amplitudes.

This posterior is mapped out with using the Gibbs (conditional) sampling technique, and the basic output is a set of samples drawn from the posterior. From these samples we provide marginal posterior mean and RMS component maps, as well as the marginal CMB power spectrum posterior.

The code assumes identical beams at all frequencies, and it is therefore necessary to smooth the data to a common resolution, limiting the analysis to large angular scales. For this particular data set, we have chosen a common resolution of 3° FWHM, with $54'$ pixels (Healpix $N_{\text{side}}=64$) and with $\ell_{\text{max}} = 150$. For more details on the degradation process, see Eriksen et al. (2008). At this resolution, the CPU time for producing one sample is around one wall-clock minute. A total of 5400 samples were produced over four independent Markov chains, of which the first 2400 were rejected due to burn-in. Twelve frequency bands (covering frequencies between 23 and 353 GHz) were included, for a total cost of around 1000 CPU hours.

The main advantage of this approach is simply that it provides us with the exact joint CMB and foreground posterior for very general foreground models. From this joint posterior, it is trivial to obtain the exact marginal CMB power spectrum and sky signal posteriors. Second, since any parametric foreground model may be included in the analysis, the method is very general and flexible. It also provides posterior maps for individual components, and is therefore a true component separation method, and not only a foreground removal tool.

Currently, the main disadvantage of the approach is the assumption of identical beam profiles at each frequency. This strictly limits the analysis to the lowest resolution of a particular data set. However, this is a limitation of the current implementation, and not of the method as such. Work is currently ongoing to extend the foreground sampler to multi-resolution experiments.

A.2. Correlated Component Analysis (CCA)

CCA (Bedini et al. 2005) is a semi-blind approach that relies on the second-order statistics of the data to estimate the mixing matrix on sub-patches of the sky. CCA assumes the data model given by Eq. 4, and makes no assumptions about the independence or lack of correlations between pairs of radiation sources. The method exploits the spatial structure of the individual source maps and adopts commonly accepted models for source frequency scalings in order to reduce the number of free parameters to be estimated.

The spatial structures of the maps are accounted for through the covariance matrices at different shifts. From the data model adopted, the data covariance matrices at shifts (τ, ψ) are given by

$$\begin{aligned} \mathbf{C}_d(\tau, \psi) &= \langle [\mathbf{d}(\theta, \phi) - \mu_d] [\mathbf{d}(\theta + \tau, \phi + \psi) - \mu_d]^t \rangle \\ &= \mathbf{A}\mathbf{C}_s(\tau, \psi)\mathbf{A}^t + \mathbf{C}_n(\tau, \psi). \end{aligned} \quad (\text{A.3})$$

where μ_d is the mean data vector, and (θ, ϕ) is the generic pixel index pair. The matrices $\mathbf{C}_d(\tau, \psi)$ can be estimated from the data, and the noise covariance matrices $\mathbf{C}_n(\tau, \psi)$ are derived from the map-making noise estimations. From Eq. (A.3), we can estimate the mixing matrix and free parameters of the source covariance matrices by matching the known quantities to the unknowns, that is by minimizing the following function for \mathbf{A} and $\mathbf{C}_s(\tau, \psi)$

$$\sum_{\tau, \psi} \|\mathbf{A}\mathbf{C}_s(\tau, \psi)\mathbf{A}^t - [\mathbf{C}_d(\tau, \psi) - \mathbf{C}_n(\tau, \psi)]\|, \quad (\text{A.4})$$

where the Frobenius norm is used and the summation is taken over the set of shift pairs (τ, ψ) for which data covariances are non-zero. Given an estimate of \mathbf{C}_s and \mathbf{C}_n , Eq. 4 can be inverted and component maps obtained via the standard inversion techniques of Wiener filtering or generalized least square inversion. For the Challenge, harmonic space Wiener filtering was applied, using a mixing matrix obtained by averaging the mixing matrices of different patches. More details on the method can be found in Bonaldi et al. (2006, 2007).

CCA can treat the variability of the spectral properties of each component with the direction of observation by working on sufficiently small sky patches, which must however be large enough to have sufficient constraining power; typically the number of pixels per patch must be around 10^5 . To obtain a continuous distribution of the free parameters of the mixing matrix, CCA is applied to a large number of partially overlapping patches.

A drawback of the present version of CCA is common to many pixel-domain approaches to separation: the data must be smoothed to a common resolution. A Fourier-domain implementation of CCA (Bedini & Salerno 2007) would be able to cope with this problem. Alternatively for the pixel-domain version, the mixing matrix could be estimated from the smoothed maps and then used to separate the sources using the full resolution data.

A.3. Generalised morphological component analysis (GMCA)

GMCA (Bobin et al. 2007) is a blind source separation method devised for separating sources from instantaneous linear mixtures. GMCA disentangles between the sources assuming that their spatial morphology is well represented in a fixed waveform dictionary (such as wavelets). In this context, the waveform dictionary leads to a so-called sparse representation: the components are well defined from only a few samples in the waveform dictionary. Sparsity enhances the diversity between the components thus improving the separation quality.

GMCA is adapted to account for physical prior knowledge in the scope of CMB component separation: simple physical priors are used to model Galactic foregrounds (dust, free-free, synchrotron). Further work will be devoted to improve the performances of the GMCA method by handling more efficiently foreground components.

A.4. Independent component analysis (FastICA)

Independent Component Analysis is an approach to component separation, looking for the components which maximize some measure of the statistical independence (Hyvarinen 1999). The FastICA algorithm presented here exploits the fact that non-Gaussianity is usually a convenient and robust measure of the statistical independence and therefore it searches for linear combinations \mathbf{y} of the input multi-frequency data, which maximize some measure of the non-Gaussianity. In the specific implementation of the idea, employed here, the non-Gaussianity is quantified by the *neg-entropy*. Denoting by $H(\mathbf{y}) = -\int p(\mathbf{y}) \log p(\mathbf{y}) d\mathbf{y}$ the entropy associated with the distribution p , we define the neg-entropy as,

$$\text{neg-entropy}(\mathbf{y}) = H(\mathbf{y}_G) - H(\mathbf{y}), \quad (\text{A.5})$$

where \mathbf{y}_G is a Gaussian variable with the same covariance matrix as \mathbf{y} . The search for the maxima of the neg-entropy is usually aided by enhancing the role of the higher order moments of \mathbf{y} , which is achieved by means of a non-linear mapping. In the present implementation, the FastICA finds the extrema of the neg-entropy approximation given by $|E[g(\mathbf{y})] - E[g(\mathbf{y}_G)]|^2$, where E means the average over the pixels, and g represents the non-linear mapping of the data, which may be a power law in the simplest case. The algorithm is straightforwardly implemented in real space, and requires the same angular resolution for all channels. Note that for an experiment like PLANCK where the resolution varies with frequency, this requires smoothing the input data to the lowest resolution before processing. The use of an efficient minimization procedure, with a required number of floating point operations scaling linearly with the size of the data set, makes the computational requirements essentially dominated by memory, needed to allocate and quickly access the multi-frequency data.

The algorithm has been tested so far as a CMB cleaning procedure, because the hypothesis of statistical independence is expected to be verified at least between CMB and diffuse foregrounds. It produced results on real (BEAST, COBE, WMAP) and simulated total intensity data, as well as on polarization simulations, on patches as well as all sky (see Maino et al. (2007) and references therein). The performance is made possible by two contingencies, i.e. the validity of the assumption of statistical independence for CMB and foregrounds, as well as the high resolution of the present CMB observations, which provides enough of statistical realizations (pixels) for the method to decompose the data into the independent components.

A.5. Harmonic-space maximum entropy method (FastMEM)

The Maximum Entropy Method (MEM) can be used to separate the CMB signal from astrophysical foregrounds including Galactic synchrotron, dust and free-free emission as well as SZ effects. The particular implementation of MEM used here works in the spherical harmonic domain. The separation is done mode-by-mode allowing one to split a huge optimisation problem into a number of smaller problems. The solution can thus be obtained faster, giving this implementation its name: FastMEM. This approach is described by Hobson et al. (1998, 1999) for Fourier modes on flat patches of the sky and by Stolyarov et al. (2002, 2005) for the full-sky case.

If we have a model (or hypothesis) H in which the measured data \mathbf{d} is a function of an underlying signal \mathbf{s} , then Bayes' theorem tells us that the posterior probability $\Pr(\mathbf{s}|\mathbf{d}, H)$ is the product of the likelihood $\Pr(\mathbf{d}|\mathbf{s}, H)$ and the prior probability $\Pr(\mathbf{s}, H)$,

divided by the evidence $\Pr(\mathbf{d}, H)$,

$$\Pr(\mathbf{s}|\mathbf{d}, H) = \frac{\Pr(\mathbf{d}|\mathbf{s}, H) \Pr(\mathbf{s}|H)}{\Pr(\mathbf{d}|H)}. \quad (\text{A.6})$$

The objective here is to maximise the posterior probability of the signal given the data. Since the evidence in Bayes' theorem is merely a normalisation constant we maximise the product of the likelihood and the prior

$$\Pr(\mathbf{s}|\mathbf{d}, H) \propto \Pr(\mathbf{d}|\mathbf{s}, H) \Pr(\mathbf{s}, H). \quad (\text{A.7})$$

We assume that the instrumental noise in each frequency channel is Gaussian-distributed, so that the log-likelihood has a form of a χ^2 misfit statistic. We make the assumption that the noise is uncorrelated between spherical harmonic modes. We also assume that the beams are azimuthally symmetric, so that they are fully described by the beam transfer function B_ℓ in harmonic space. For mode (ℓ, m) , the log-likelihood is

$$\chi^2(\mathbf{s}_{\ell m}) = (\mathbf{d}_{\ell m} - B_\ell \mathbf{A} \mathbf{s}_{\ell m})^T N_{\ell m}^{-1} (\mathbf{d}_{\ell m} - B_\ell \mathbf{A} \mathbf{s}_{\ell m}) \quad (\text{A.8})$$

where \mathbf{A} is the fixed frequency conversion matrix which describes how the components are mixed to form the data, and $N_{\ell m}^{-1}$ is the inverse noise covariance matrix for this mode. If the instrumental noise is uncorrelated between channels, then this matrix is diagonal. However, unresolved point sources can be modeled as a correlated noise component.

The prior can be Gaussian, and in this case we recover the Wiener filter with the well-known analytical solution for the signal s . However, the astrophysical components have strongly non-Gaussian distribution, especially in the Galactic plane. Therefore Hobson et al. (1998) suggested that an entropic prior be used instead. In this case, maximising the posterior is equivalent to the minimising the following functional for each spherical harmonic mode

$$\Phi_{\text{MEM}}(\mathbf{s}_{\ell m}) = \chi^2(\mathbf{s}_{\ell m}) - \alpha S(\mathbf{s}_{\ell m}) \quad (\text{A.9})$$

where $S(\mathbf{s})$ is the entropic term, and α is the regularisation parameter. The minimisation can be done numerically using one of a number of algorithms (Press et al. 1992).

FastMEM is a non-blind method, so the spectral behaviour of the components must be known in advance. Since \mathbf{A} is fixed, the spectral properties of the components must be the same everywhere on the sky. However, small variations in the spectral properties, for example, dust temperature, synchrotron spectral index or SZ cluster electron temperature, can be accounted for by introducing additional components. These additional components correspond to terms in the Taylor expansion of the frequency spectrum with respect to the relevant parameter.

The initial priors on the components are quite flexible and they can be updated by iterating the component separation, especially if the signal-to-noise is high enough.

It is not necessary for all of the input maps to be at the same resolution since FastMEM solves for the most probable solution for unsmoothed signal, deconvolving and denoising maps simultaneously. It is flexible enough to include any datasets with known window function and noise properties. A mask can easily be applied to the input data (the same mask for all frequency channels) and this does not cause problems with the separation.

Since FastMEM uses priors on the signals, the solution for the signals is biased. This is especially evident if the signal-to-noise ratio is low. It is possible to de-bias the power spectrum statistically, knowing the priors and the FastMEM separation errors per mode. However, one can not de-bias the recovered maps

since the errors are quadratic and de-biasing will introduce phase errors in the harmonics.

No information about the input components was used in the separation, and the prior power spectra were based solely on the physical properties of the components and templates available in the literature. The prior on the CMB component was set using the best-fit theoretical spectrum, instead of a WMAP-constrained realisation. This has a significant effect at low multipoles.

A.6. Spectral estimation via expectation-maximization (SEVEM)

SEVEM (Martínez-González et al. 2003) tries to recover only the CMB signal, treating the rest of the emissions as a generalised noise. As a first step, the cosmological frequency maps, 100, 143 and 217 GHz, are foreground cleaned using an internal template fitting technique. Four templates are obtained from the difference of two consecutive frequency channels, which are smoothed down to the same resolution if necessary, to avoid the presence of CMB signal in the templates. In particular, we construct maps of (30-44), (44-70), (545-353) and (857-545) differences. The central frequency channels are then cleaned by subtracting a linear combination of these templates. The coefficients of this combination are obtained minimising the variance of the final clean map outside the considered mask. The second step consists on estimating the power spectrum of the CMB from the three cleaned maps using the method (based on the Expectation-Maximization algorithm) described in Martínez-González et al. (2003), which has been adapted to deal with spherical data. Using simulations of CMB plus noise, processed in the same way as the Challenge data, we obtain the bias and statistical error of the estimated power spectrum and construct an unbiased version of the C_ℓ 's of the CMB. This unbiased power spectrum is used to recover the CMB map from the three clean channels through Wiener filter in harmonic space. Finally, we estimate the noise per pixel of the reconstructed map using CMB plus noise simulations.

One of the advantages of SEVEM is that it does not need any external data set or need to make any assumptions about the frequency dependence or the power spectra of the foregrounds, other than the fact that they are the dominant contribution at the lowest and highest frequency channels. This makes the method very robust and, therefore, it is expected to perform well for real PLANCK data. Moreover, SEVEM provides a good recovery of the power spectrum up to relatively high ℓ and a small error in the CMB map reconstruction. In addition the method is very fast, which allows one to characterise the errors of the CMB power spectrum and map using simulations. The cleaning of the data takes around 20 minutes, while the estimation of the power spectrum and map requires around 15 and 30 minutes respectively. In fact, the whole process described, including producing simulations to estimate the bias and errors, takes around 30 hours in one single CPU. Regarding weak points, the method reconstructs only the CMB and does not try to recover any other component of the microwave sky although it could be generalised to reconstruct simultaneously the both the CMB and the thermal SZ effect. Also, the reconstructed CMB map is not full-sky, since the method does not aim to remove the strong contamination at the centre of the Galactic plane or at the point source positions. In any case, the masked region excluded for the analysis is relatively small.

A.7. Spectral matching independent component analysis (SMICA)

The principle of SMICA can be summarised in three steps: 1) Compute spectral statistics 2) Fit a component-based model to them 3) Use the result to implement a Wiener filter in harmonic space. More specifically, an idealized operation goes as follows. Denote $\mathbf{d}(\xi)$ the column vector whose i -th entry contains the observation in direction ξ for the i -th channel and denote $\mathbf{d}_{\ell m}$ the vector of same size (the number of frequency channels) in harmonic space. This is modeled as the superposition of C components $\mathbf{d}_{\ell m} = \sum_{c=1}^C \mathbf{d}_{\ell m}^c$. In Step 1), we compute spectral matrices $\widehat{\mathbf{C}}_{\ell} = \frac{1}{2\ell+1} \sum_m \mathbf{d}_{\ell m} \mathbf{d}_{\ell m}^T$. In Step 2) we model the ensemble-averaged spectral matrix $\mathbf{C}_{\ell} = \langle \widehat{\mathbf{C}}_{\ell} \rangle$ as the superposition of C uncorrelated components: $\mathbf{C}_{\ell} = \sum_{c=1}^C \mathbf{C}_{\ell}^c$ and, for each component, we postulate a parametric model, that is, we let the matrix set $\{\mathbf{C}_{\ell}^c\}_{\ell=0}^{\ell_{\max}}$ be a function of a parameter vector θ^c . This parameterization embodies our prior knowledge about a given component. For instance, for the CMB component, we take $[\mathbf{C}_{\ell}^{\text{cmb}}]_{ij} = e_i e_j c_{\ell}$ where e_i is the known CMB emission coefficient for channel i and c_{ℓ} is the unknown angular power spectrum at frequency ℓ . The parameter vector for CMB would then be $\theta^{\text{cmb}} = \{c_{\ell}\}_{\ell=0}^{\ell_{\max}}$. All unknown parameters for all components are then estimated by fitting the model to the spectral statistics, i.e. by solving $\min_{\theta^1, \dots, \theta^C} \sum_{\ell=0}^{\ell_{\max}} (2\ell+1) K[\widehat{\mathbf{C}}_{\ell} | \sum_{c=0}^C \mathbf{C}_{\ell}^c(\theta^c)]$ where $K[\mathbf{C}_1 | \mathbf{C}_2]$ is a measure of mismatch between two covariance matrices \mathbf{C}_1 and \mathbf{C}_2 . The resulting values $\hat{\theta}^1, \dots, \hat{\theta}^C$ provide estimates $\mathbf{C}_{\ell}^c(\hat{\theta}^c)$ of \mathbf{C}_{ℓ}^c . The Wiener filter estimate of $\mathbf{d}_{\ell m}^c$ can be expressed as $\hat{\mathbf{d}}_{\ell m}^c = \mathbf{C}_{\ell}^c \mathbf{C}_{\ell}^{-1} \mathbf{d}_{\ell m}$. In practice, we use the fitted spectral matrices estimated at the previous step: component c is estimated as

$$\hat{\mathbf{d}}_{\ell m}^c = \mathbf{C}_{\ell}^c(\hat{\theta}^c) \mathbf{C}_{\ell}(\hat{\theta})^{-1} \mathbf{d}_{\ell m} \quad (\text{A.10})$$

and the maps of each component in each channel are finally computed by inverse spherical harmonic transforms.

For processing the current data set, we have used a model containing four components: the CMB, the SZ component, a 4-dimensional Galactic component and a noise component.

The actual processing includes several modifications with respect to this outline: a) beam correction applied to each spectral matrix $\widehat{\mathbf{C}}_{\ell}$; b) spectral binning by which the (beam corrected) spectral matrices are averaged over bins of increasing lengths; c) localization implemented via aopdized masks, by which the SMICA process is conducted independently over two different sky zones.

Strengths: a) No prior information used regarding Galactic emission b) Accurate recovery of CMB via Wiener filter c) Relatively fast algorithm d) Built-in goodness of fit.

Weaknesses: a) The results reported here do not account for the contribution of point sources for which a convenient model is lacking; b) Localization in two zones is probably too crude c) No separation of Galactic components.

A.8. Wavelet-based high-resolution Fitting of Internal Templates (WI-FIT)

WI-FIT (Hansen et al. 2006) is based on fitting and subtraction of internal templates. Regular (external) template fitting uses external templates of Galactic components based on observations at frequencies different from the ones used to study the CMB. These templates are fitted to CMB data, the best fit coefficients for each component are found and the templates are subtracted from the map using these coefficients in order to obtain a clean

CMB map. WI-FIT differs from this procedure in two respects: (1) It does not rely on external observations of the galaxy but forms templates by taking the difference of CMB maps at different channels. The CMB temperature is equal at different frequencies whereas the Galactic components are not. For this reason, the difference maps contain only a sum of Galactic components. A set of templates are constructed from difference maps based on different combinations of channels. (2) The fitting of the templates are done in wavelet space where the uncertainty on the foreground coefficients is much lower than a similar pixel based approach (in the pixel based approach, no pixel-pixel correlations are taken into account since the correlation matrix will become too large for PLANCK-like data sets. In the wavelet based approach, a large part of these correlations are taken into account in scale-scale covariance matrices).

For calibration purposes, a set of 500 simulated CMB maps need to be produced and the full wavelet fitting procedure applied to all maps. This is where most CPU time goes. For PLANCK resolution maps, around 1 Gb of memory is necessary to apply WI-FIT and a total of around 400 CPU hours are required.

The strength of WI-FIT is that it relies on very few assumptions about the Galactic components. WI-FIT does however assume that the spectral indices do not vary strongly from pixel to pixel within the frequency range used in the analysis. If this assumption is wrong then WI-FIT leaves residuals in the areas where there are strongly varying spectral indices.

Another advantage of WI-FIT is that it is easy to apply and is completely linear, i.e. the resulting map is a linear combination of frequency channels with well known noise and beam properties. This will in general result in increased noise variance in the cleaned map. In order to avoid this, we smooth the internal templates in order to make the noise at small scales negligible and at the same time not make significant changes to the shape of the diffuse foregrounds. If the diffuse foregrounds turn out to be important at small scales $l > 300$, the smoothing of the internal templates will significantly reduce the ability of WI-FIT to perform foreground cleaning at these scales. Tests on the WMAP data have shown that diffuse foregrounds do not seem to play an important role at such small scales. This is valid for the frequency range observed by WMAP (i.e. at LFI-frequencies), similar tests will need to be made for the PLANCK HFI data.

Finally, WI-FIT does not do anything to the point sources, which need to be masked.