Planck 2015 results

DOI:
10.1051/0004-6361/201525967

Citation for published version (APA):
Adam, R., Ade, P. A. R., Aghanim, N., Alves, M. I. R., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A. J., Barreiro, R. B., Bartlett, J. G., Bartolo, N., Battaner, E., Benabed, K., Benoît, A., Benoit-Lévy, A., Bernard, J. P., Bersanelli, M., Bielewicz, P., ... Zonca, A. (2016). Planck 2015 results: X. Diffuse component separation: Foreground maps. *Astronomy and Astrophysics*, *594*, [A10]. https://doi.org/10.1051/0004-6361/201525967

Published in:
Astronomy and Astrophysics

A&A 594, A10 (2016)
DOI: 10.1051/0004-6361/201525967

Planck 2015 results

Astronomy & Astrophysics
Special feature

Planck 2015 results

X. Diffuse component separation: Foreground maps

Planck Collaboration: R. Adam[82], P. A. R. Ade[97], N. Aghanim[65], M. I. R. Alves[107,11,65], M. Arnaud[80], M. Ashdown[76,6], J. Aumont[65], C. Baccigalupi[95], A. J. Banday[107,11], R. B. Barreiro[71], J. G. Bartlett[1,73], N. Bartolo[33,72], E. Battaner[109,110], K. Benabed[66,106], A. Benoît[63], A. Benoit-Lévy[27,66,106], J.-P. Bernard[107,11], M. Bersanelli[36,53], P. Bielewicz[90,11,95], J. J. Bock[73,13], A. Bonaldi[74], L. Bonavera[71], J. R. Bond[10], J. Borrill[16,101], F. R. Bouchet[66,99], F. Boulanger[65], M. Bucher[1], C. Burigana[52,34,54], R. C. Butler[52], E. Calabrese[103], J.-F. Cardoso[81,1,66], A. Catalano[82,79], A. Challinor[68,76,14], A. Chamballu[80,18,65], R.-R. Chary[62], H. C. Chiang[30,7], P. R. Christensen[91,39], D. L. Clements[61], S. Colombi[66,106], L. P. L. Colombo[26,73], C. Combet[82], F. Couchot[78], A. Coulais[79], B. P. Crill[73,13], A. Curto[71,6,76], F. Cuttaia[52], L. Danese[95], R. D. Davies[74], R. J. Davis[74], P. de Bernardis[35], A. de Rosa[52], G. de Zotti[49,95], J. Delabrouille[1], F.-X. Désert[59], C. Dickinson[74], J. M. Diego[71], H. Dole[65,64], S. Donzelli[53], O. Doré[73,13], M. Douspis[65], A. Ducout[66,61], X. Dupac[42], G. Efstathiou[68], F. Elsner[27,66,106], T. A. Enßlin[86], H. K. Eriksen[69], E. Falgarone[79], J. Fergusson[14], F. Finelli[52,54], O. Forni[107,11], M. Frailis[51], A. A. Fraisse[30], E. Franceschi[52], A. Frejsel[91], S. Galeotta[51], S. Galli[75], K. Ganga[1], T. Ghosh[65], M. Giard[107,11], Y. Giraud-Héraud[1], E. Gjerløw[69], J. González-Nuevo[22,71], K. M. Górski[73,112], S. Gratton[76,68], A. Gregorio[37,51,58], A. Gruppuso[52], J. E. Gudmundsson[104,93,30], F. K. Hansen[69], D. Hanson[88,73,10], D. L. Harrison[68,76], G. Helou[13], S. Henrot-Versillé[78], C. Hernández-Monteagudo[15,86], D. Herranz[71], S. R. Hildebrandt[73,13], E. Hivon[66,106], M. Hobson[6], W. A. Holmes[73], A. Hornstrup[19], W. Hovest[86], K. M. Huffenberger[28], G. Hurier[65], A. H. Jaffe[61], T. R. Jaffe[107,11], W. C. Jones[30], M. Juvela[29], E. Keihänen[29], R. Keskitalo[16], T. S. Kisner[84], R. Kneissl[41,8], J. Knoche[86], M. Kunz[20,65,3], H. Kurki-Suonio[29,47], G. Lagache[5,65], A. Lähteenmäki[2,47], J.-M. Lamarre[79], A. Lasenby[6,76], M. Lattanzi[34,55], C. R. Lawrence[73], M. Le Jeune[1], J. P. Leahy[74], R. Leonardi[9], J. Lesgourgues[67,105], F. Levrier[79], M. Liguori[33,72], P. B. Lilje[69], M. Linden-Vørnle[19], M. López-Caniego[42,71], P. M. Lubin[31], J. F. Macías-Pérez[82], G. Maggio[51], D. Maino[36,53], N. Mandolesi[52,34], A. Mangilli[65,78], M. Maris[51], D. J. Marshall[80], P. G. Martin[10], E. Martínez-González[71], S. Masi[35], S. Matarrese[33,72,44], P. McGehee[62], P. R. Meinhold[31], A. Melchiorri[35,56], L. Mendes[42], A. Mennella[36,53], M. Migliaccio[68,76], S. Mitra[60,73], M.-A. Miville-Deschênes[65,10], A. Moneti[66], L. Montier[107,11], G. Morgante[52], D. Mortlock[61], A. Moss[98], D. Munshi[97], J. A. Murphy[89], P. Naselsky[92,40], F. Nati[30], P. Natoli[34,4,55], C. B. Netterfield[23], H. U. Nørgaard-Nielsen[19], F. Noviello[74], D. Novikov[85], I. Novikov[91,85], E. Orlando[111], C. A. Oxborrow[19], F. Paci[95], L. Pagano[35,56], F. Pajot[65], R. Paladini[62], D. Paoletti[52,54], B. Partridge[46], F. Pasian[51], G. Patanchon[1], T. J. Pearson[13,62], O. Perdereau[78], L. Perotto[82], F. Perrotta[95], V. Pettorino[45], F. Piacentini[35], M. Piat[1], E. Pierpaoli[26], D. Pietrobon[73], S. Plaszczynski[78], E. Pointecouteau[107,11], G. Polenta[4,50], G. W. Pratt[80], G. Prézeau[13,73], S. Prunet[66,106], J.-L. Puget[65], J. P. Rachen[24,86], W. T. Reach[108], R. Rebolo[70,17,21], M. Reinecke[86], M. Remazeilles[74,65,1], C. Renault[82], A. Renzi[38,57], I. Ristorcelli[107,11], G. Rocha[73,13], C. Rosset[1], M. Rossetti[36,53], G. Roudier[1,79,73], J. A. Rubiño-Martín[70,21], B. Rusholme[62], M. Sandri[52], D. Santos[82], M. Savelainen[29,47], G. Savini[94], D. Scott[25], M. D. Seiffert[73,13], E. P. S. Shellard[14], L. D. Spencer[97], V. Stolyarov[6,102,77], R. Stompor[1], A. W. Strong[87], R. Sudiwala[97], R. Sunyaev[86,100], D. Sutton[68,76], A.-S. Suur-Uski[29,47], J.-F. Sygnet[66], J. A. Tauber[43], L. Terenzi[96,52], L. Toffolatti[22,71,52], M. Tomasi[36,53], M. Tristram[78], M. Tucci[20], J. Tuovinen[12], G. Umana[48], L. Valenziano[52], J. Valiviita[29,47], F. Van Tent[83], P. Vielva[71], F. Villa[52], L. A. Wade[73], B. D. Wandelt[66,106,32], I. K. Wehus[73,69,⋆], A. Wilkinson[74], D. Yvon[18], A. Zacchei[51], and A. Zonca[31]

(Affiliations can be found after the references)

Received 26 February 2015 / Accepted 22 February 2016

ABSTRACT

Planck has mapped the microwave sky in temperature over nine frequency bands between 30 and 857 GHz and in polarization over seven frequency bands between 30 and 353 GHz in polarization. In this paper we consider the problem of diffuse astrophysical component separation, and process these maps within a Bayesian framework to derive an internally consistent set of full-sky astrophysical component maps. Component separation dedicated to cosmic microwave background (CMB) reconstruction is described in a companion paper. For the temperature analysis, we combine the *Planck* observations with the 9-yr *Wilkinson* Microwave Anisotropy Probe (WMAP) sky maps and the Haslam et al. 408 MHz map, to derive a joint model of CMB, synchrotron, free-free, spinning dust, CO, line emission in the 94 and 100 GHz channels, and thermal dust emission. Full-sky maps are provided for each component, with an angular resolution varying between 7.′5 and 1°. Global parameters (monopoles, dipoles, relative calibration, and bandpass errors) are fitted jointly with the sky model, and best-fit values are tabulated. For polarization, the model includes CMB, synchrotron, and thermal dust emission. These models provide excellent fits to the observed data, with rms temperature residuals smaller than $4\,\mu$K over 93% of the sky for all *Planck* frequencies up to 353 GHz, and fractional errors smaller than 1% in the remaining 7% of the sky. The main limitations of the temperature model at the lower frequencies are internal degeneracies among the spinning dust, free-free, and synchrotron components; additional observations from external low-frequency experiments will be essential to break these degeneracies. The main limitations of the temperature model at the higher frequencies are uncertainties in the 545 and 857 GHz calibration and zero-points. For polarization, the main outstanding issues are instrumental systematics in the 100–353 GHz bands on large angular scales in the form of temperature-to-polarization leakage, uncertainties in the analogue-to-digital conversion, and corrections for the very long time constant of the bolometer detectors, all of which are expected to improve in the near future.

Key words. ISM: general – cosmology: observations – polarization – cosmic background radiation – diffuse radiation – Galaxy: general

⋆ Corresponding author: I. K. Wehus; i.k.wehus@astro.uio.no

1. Introduction

This paper, one of a set associated with the 2015 release of data from the *Planck*[1] mission (Planck Collaboration I 2016), presents a coherent astrophysical model of the microwave sky in both temperature and polarization, as derived from the most recent *Planck* observations. For temperature, the analysis also incorporates the 9-yr WMAP observations (Bennett et al. 2013) and a 408 MHz survey (Haslam et al. 1982), allowing the separation of synchrotron, free-free, and spinning dust emission.

In March 2013, the *Planck* Consortium released its first temperature measurements of the microwave sky, summarized in terms of nine frequency maps between 30 and 857 GHz (Planck Collaboration I 2014). The richness of these data has enabled great progress in our understanding of the astrophysical composition of the microwave sky. The current *Planck* data release presents additionally high-sensitivity, full-sky maps of the polarized microwave sky, offering a fresh view on both cosmological and astrophysical phenomena.

With increased data volume and quality comes both greater scientific potential and more stringent requirements on model complexity and sophistication. The current *Planck* data release is more ambitious than the 2013 release in terms of component separation efforts, accounting for more astrophysical effects and components. In this round, three related papers summarize the *Planck* 2015 component separation products and approaches. First, cosmic microwave background (CMB) reconstruction and extraction are discussed in Planck Collaboration IX (2016). Second, this paper presents the diffuse astrophysical foreground products derived from the 2015 *Planck* observations, both in temperature and polarization. Third, Planck Collaboration XXV (2016) discusses the scientific interpretation of the new low-frequency *Planck* foreground products.

The main goal of the current paper is to establish a single, internally coherent and global parametric model of the microwave sky, simultaneously accounting for all significant diffuse astrophysical components and relevant instrumental effects using the Bayesian `Commander` analysis framework (Eriksen et al. 2004, 2006, 2008). As such, our discussion does not focus on any single emission component, but rather emphasize the global picture. In the 2013 data release, the same framework was applied to the *Planck* temperature measurements for frequencies between 30 and 353 GHz, considering only angular scales larger than 40′ full-width half-maximum (FWHM). This resulted in low-resolution CMB, CO, and thermal dust emission maps, as well as a single low-frequency foreground component combining contributions from synchrotron, free-free, and spinning dust emission (Planck Collaboration XII 2014). Here we extend that analysis in multiple directions. First, instead of 15.5 months of temperature data, the new analysis includes the full *Planck* mission data, 50 months of Low Frequency Instrument (LFI) and 29 months of High Frequency Instrument (HFI) data, in both temperature and polarization. Second, we now also include the 9-yr WMAP observations between 23 and 94 GHz and a 408 MHz survey map, providing enough frequency constraints to decompose the low-frequency foregrounds into separate synchrotron, free-free, and spinning dust components. Third, we now include the *Planck* 545 and 857 GHz frequency bands, allowing us to constrain the thermal dust temperature and emissivity index with greater precision, thereby reducing degeneracies between CMB, CO, and free-free emission. At the same time, we find that the calibration and bandpass measurements of these two channels represent two of the most important sources of systematic uncertainty in the analysis. Fourth, the present analysis implements a multi-resolution strategy to provide component maps at high angular resolution. Specifically, the CMB is recovered with angular resolution 5′ FWHM (Planck Collaboration IX 2016), thermal dust emission and CO $J = 2 \rightarrow 1$ lines are recovered at 7.′5 FWHM, and synchrotron, free-free, and spinning dust are recovered at 1° FWHM. The resulting parameter fits define the *Planck* 2015 baseline astrophysical model in temperature and polarization. We emphasize, however, that these models are not unique, but instead represent minimal physically well-motivated models that are able to reproduce the current data.

As in the 2013 data release, the CMB solutions derived, using this Bayesian approach, form the basis of the *Planck* 2015 CMB temperature likelihood on large angular scales. This is described in detail in Planck Collaboration XI (2016), which also presents a detailed characterization of the low-multipole CMB angular power spectrum. The low-frequency astrophysical model presented here is used as input for the temperature-to-polarization bandpass mismatch corrections for the LFI polarization maps (Planck Collaboration II 2016).

The paper is organized as follows. Section 2 gives an overview of the computational framework implemented in the `Commander` code. Section 3 describes the data selection and processing. Section 4, gives an overview of the relevant astrophysical components and systematic effects. Sections 5 and 6 give the main temperature and polarization products. We summarize in Sect. 7.

[1] *Planck* (`http://www.esa.int/Planck`) is a project of the European Space Agency (ESA) with instruments provided by two scientific consortia funded by ESA member states and led by Principal Investigators from France and Italy, telescope reflectors provided through a collaboration between ESA and a scientific consortium led and funded by Denmark, and additional contributions from NASA (USA).

2. Algorithms

2.1. Data, posterior distribution and priors

Most of the results derived in this paper are established within a standard Bayesian analysis framework, as implemented in the `Commander` code, in which an explicit parametric model, $\boldsymbol{s}(\theta)$, is fitted to a set of observations, $\boldsymbol{d}$, either by maximizing or mapping out the corresponding posterior distribution,

$$P(\theta|\boldsymbol{d}) = \frac{P(\boldsymbol{d}|\theta)P(\theta)}{P(\boldsymbol{d})} \propto \mathcal{L}(\theta)P(\theta). \tag{1}$$

Here θ denotes some general set of free parameters in the model, $\mathcal{L}(\theta) = P(\boldsymbol{d}|\theta)$ is the likelihood, and $P(\theta)$ denotes a set of priors on θ. The evidence, $P(\boldsymbol{d})$, is a constant with respect to the parameter set, and is neglected in the following.

The data are defined by a set of pixelized frequency-channel sky maps, $\boldsymbol{d} = \{\boldsymbol{d}_\nu\}$, comprising the three Stokes parameters I, Q and U. In this paper, however, we analyse temperature and polarization separately; therefore the data vector comprises either I or $\{Q, U\}$.

We start by assuming that the data at a given frequency ν may be described as a linear sum of signal $\boldsymbol{s}_\nu$ and noise $\boldsymbol{n}_\nu$,

$$\boldsymbol{d}_\nu = \boldsymbol{s}_\nu + \boldsymbol{n}_\nu, \tag{2}$$

where $\boldsymbol{n}_\nu$ is assumed to be Gaussian-distributed with a known covariance matrix N_ν. For the signal, we adopt the following parametric expression:

$$\boldsymbol{s}_\nu(\theta) = s_\nu(\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu) \tag{3}$$

$$= g_\nu \sum_{i=1}^{N_{\mathrm{comp}}} \mathsf{F}^i_\nu(\beta_i, \Delta_\nu)\boldsymbol{a}_i + \mathsf{T}_\nu \boldsymbol{m}_\nu, \tag{4}$$

where $\boldsymbol{a}_i$ is an amplitude map for component i at a given reference frequency, β_i is a general set of spectral parameters for the same component, g_ν is a multiplicative calibration factor for frequency ν, Δ_ν is a linear shift in the bandpass central frequency, and $\boldsymbol{m}_\nu$ is a set of template correction amplitudes, such as monopole, dipole, or zodiacal light corrections for temperature, or calibration leakage templates for polarization. The corresponding spatially fixed templates are organized column-wise in a template matrix T_ν. The mixing matrix, $\mathsf{F}^i_\nu(\beta_i, \Delta_\nu)$, accounts for the effect of spectral changes as a function of frequency for component i, parametrized by β_i, as well as bandpass integration effects and unit conversions. For numerical stability, all internal calculations are performed in units of brightness temperature, and $\boldsymbol{a}$ is therefore naturally defined in the same units at some specified reference frequency.

The posterior distribution takes the usual form,

$$\begin{aligned} P(\theta|\boldsymbol{d}_\nu) &= P(\boldsymbol{d}_\nu|\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu, C_\ell) P(\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu, C_\ell) \\ &= \mathcal{L}(\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu) P(\boldsymbol{a}_i) P(\beta_i) P(\boldsymbol{a}_\mathrm{cmb}|C_\ell), \end{aligned} \quad (5)$$

where we have included the CMB power spectrum, C_ℓ, and also implicitly adopted uniform priors on $\boldsymbol{g}_\nu$, Δ_ν, $\boldsymbol{m}_\nu$, and C_ℓ. Because the noise is assumed to be Gaussian and independent between frequency channels, the likelihood reads

$$\mathcal{L}(\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu) \propto \exp\left(-\frac{1}{2}\sum_\nu [\boldsymbol{d}_\nu - \boldsymbol{s}_\nu(\theta)]^\mathrm{T} \mathsf{N}^{-1} [\boldsymbol{d}_\nu - \boldsymbol{s}_\nu(\theta)]\right) \quad (6)$$

Likewise, we further assume the CMB signal to be Gaussian distributed with a covariance matrix, $\mathsf{S}(C_\ell)$, given by the power spectrum, and the corresponding CMB prior factor therefore reads

$$P(\boldsymbol{a}_\mathrm{cmb}|C_\ell) = \frac{e^{-\frac{1}{2}\boldsymbol{a}^\mathrm{T}_\mathrm{cmb}\mathsf{S}^{-1}(C_\ell)\boldsymbol{a}_\mathrm{cmb}}}{\sqrt{|\mathsf{S}(C_\ell)|}}. \quad (7)$$

The only undefined factors in the posterior are the amplitude and spectral parameter priors, $P(\boldsymbol{a}^i)$ and $P(\beta^i)$. These represent the most difficult problem to handle from a conceptual point of view, since the prior is to some extent a matter of personal preference. However, we adopt the following general practices in this paper. First, for low-resolution analyses that include fitting of template amplitudes (e.g., monopoles and dipoles), we always impose a strict positivity prior, i.e., $\boldsymbol{a}_i > 0$, on all signal amplitudes except the CMB. Without such a prior, there are large degeneracies between the zero-points of the amplitude maps and the individual template amplitudes. Second, for the high angular resolution analysis, we fix the template amplitudes at the low-resolution values and disable the positivity prior, in order to avoid noise bias. Third, to further break degeneracies, we adopt fiducial values for the monopole, dipole, and calibration factors for a few selected channels, effectively imposing a set of external priors from CMB dipole measurements and H I cross-correlation (Planck Collaboration VIII 2014) to anchor the full solution. Fourth, for the spectral parameters, we adopt Gaussian priors with means and variances informed by the high signal-to-noise values observed in the Galactic plane, which for all practical purposes are independent of the adopted priors. Intuitively, we demand that a map of the spectral parameter in question should not be much different in the data-dominated and the prior-dominated regions of the sky. Fourth, one of the components in the temperature model is free-free emission, which has two free parameters, namely the effective emission measure, EM, and the electron temperature, T_e. The latter of these is very poorly constrained with the current data set except in the central Galactic plane, and we therefore adopt a smoothness prior on this paper to increase the effective signal-to-noise ratio, demanding that is must be smooth on 2° FWHM scales. This in turn has a large computational cost by making the overall foreground parameter estimation process non-local, and T_e is therefore only varied in fast maximum-likelihood searches, not in expensive sampling analyses. Its effect on other parameters is, however, minimal, precisely because of its low signal-to-noise ratio. Finally, in addition to these informative priors, we adopt a Jeffreys prior for the spectral parameters in order to suppress prior volume effects (Jeffreys 1946; Eriksen et al. 2008; Dunkley et al. 2009).

2.2. Gibbs sampling and posterior maximization

As described above, the posterior distribution contains many millions of free (both non-Gaussian and strongly correlated) parameters for *Planck* – ranging from 11 million in the following low-resolution analysis to 200 million in the corresponding high-resolution analysis – and mapping out this distribution poses a significant computational problem. Indeed, no direct sampling algorithm exists for the full distribution, and the only computationally efficient solution currently known is that of Gibbs sampling, a well-known textbook algorithm in modern statistical analysis (e.g., Gelman et al. 2003). The underlying idea of this method is that samples from a complicated multivariate distribution may be drawn by iteratively sampling over the corresponding conditional distributions, which usually have much simpler, and often analytic, sampling algorithms. This framework was originally introduced to the CMB analysis field by Jewell et al. (2004) and Wandelt et al. (2004), and subsequently developed into a fully functional computer code called `Commander` by Eriksen et al. (2004, 2008).

For the problem in question in this paper, this algorithm may be schematically translated into an explicit set of sampling steps through the following Gibbs chain:

$$\boldsymbol{a}_i \leftarrow P(\boldsymbol{a}_i|\beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu, C_\ell) \quad (8)$$

$$\beta_i \leftarrow P(\beta_i|\boldsymbol{a}_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu, C_\ell) \quad (9)$$

$$g_\nu \leftarrow P(g_\nu|\boldsymbol{a}_i, \beta_i, \boldsymbol{m}_\nu, \Delta_\nu, C_\ell) \quad (10)$$

$$\boldsymbol{m}_\nu \leftarrow P(m_\nu|\boldsymbol{a}_i, \beta_i, g_\nu, \Delta_\nu, C_\ell) \quad (11)$$

$$\Delta_\nu \leftarrow P(\Delta_\nu|\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, C_\ell) \quad (12)$$

$$C_\ell \leftarrow P(C_\ell|\boldsymbol{a}_i, \beta_i, g_\nu, \boldsymbol{m}_\nu, \Delta_\nu). \quad (13)$$

Here "←" denotes drawing a sample from the distribution on the right-hand side. After some burn-in period, the theory of Gibbs sampling guarantees that the joint set of parameters is indeed drawn from the correct joint distribution. For a full description of the various steps in the algorithm, see Eriksen et al. (2008).

While no fully functional alternatives to Gibbs sampling have been established for this full joint distribution to date, Gibbs sampling alone by no means solves all computational problems. In particular, this algorithm is notorious for its slow convergence for nearly degenerate parameters, since it by construction only moves through parameter space parallel to coordinate axes. For this reason, we implement an additional posterior maximization phase, in which we search directly for the posterior maximum point rather than attempt to sample from the full distribution. The resulting solution may then serve either as a final product in its own right, by virtue of being a maximum-posterior estimate, or as the starting position for a regular Gibbs sampling analysis. The crucial point, though, is that special-purpose nonlinear search algorithms can be used in this phase, moving in arbitrary directions through parameter space, and individual optimization combinations may be introduced to jointly probe directions with particularly strong degeneracies. Perhaps the single most important example in this respect is the parameter combination between component amplitudes, detector calibrations, and bandpass uncertainties, $\{\boldsymbol{a}_i, g_\nu, \Delta_\nu\}$, all three of which essentially correspond to scaling parameters. However, since both g_ν and $\boldsymbol{a}_i$ are conditionally linear parameters, and only Δ_ν is truly nonlinear, it is possible to solve analytically for g_ν or $\boldsymbol{a}_i$, *conditioning* on any given fixed value of Δ_ν. Consequently, one can set up a nonlinear Powell-type search (Press et al. 2002) for Δ_ν, in which the optimal values of either g_ν or $\boldsymbol{a}_i$ are quickly computed at each iteration in the search. A second example is the electron temperature discussed above, for which non-local optimization is feasible, whereas a full-blown sampling algorithm is too expensive to converge robustly. In this situation, fixing the parameter at its maximum-posterior value is vastly preferable compared to adding an unconverged degree of freedom in the full sampler.

Even with this optimization phase, however, there is always an inherent danger of the algorithm being trapped in a local posterior maximum. Indeed, with a distribution involving millions of highly correlated parameters, it is exceedingly difficult to prove that the derived solution is the true global posterior maximum. As a partial solution to this problem, we initialize the search using different starting positions, and carefully monitor the convergence properties of the chains.

3. Data selection and processing

The primary data used in this paper are the 2015 *Planck* temperature and polarization sky maps (Planck Collaboration VI 2016; Planck Collaboration VIII 2016). For the temperature analysis we additionally include the 9-yr WMAP observations[2] (Bennett et al. 2013) and a full-sky 408 MHz survey map (Haslam et al. 1982), with the goal of individually resolving synchrotron, free-free, and spinning dust emission. For WMAP, we adopt the beam-symmetrized frequency maps for the foreground-dominated *K*- and *Ka*-bands, to mitigate beam artifacts around compact sources, but we use the standard maps for the CMB-dominated *Q*-, *V*-, and *W*-bands, because of their more accurate noise description. At the lowest frequency, we adopt the destriped version of the 408 MHz survey map recently published by Remazeilles et al. (2015).

In order to maximize our leverage with respect to bandpass measurement uncertainties and line emission mechanisms, we employ individual detector and detector set ("ds") maps for all *Planck* frequencies between 70 and 857 GHz, and differencing assembly ("DA") maps for WMAP. However, the polarization analysis employs frequency maps in order to maximize signal-to-noise ratio and to minimize correlated noise from destriper mapmaking uncertainties. Intensity-to-polarization leakage from bandpass mismatch between detectors is suppressed through the use of precomputed leakage templates (Planck Collaboration II 2016; Planck Collaboration VIII 2016). For LFI, these templates are based on a preliminary version of the foreground products presented in this paper. The full set of clean channels used in this analysis is summarized in Table 1 in terms of centre frequencies, resolution, and noise levels. In total, 32 individual detector and detector set maps[3] and frequency maps are included in the temperature analysis, and seven frequency maps in the polarization analysis.

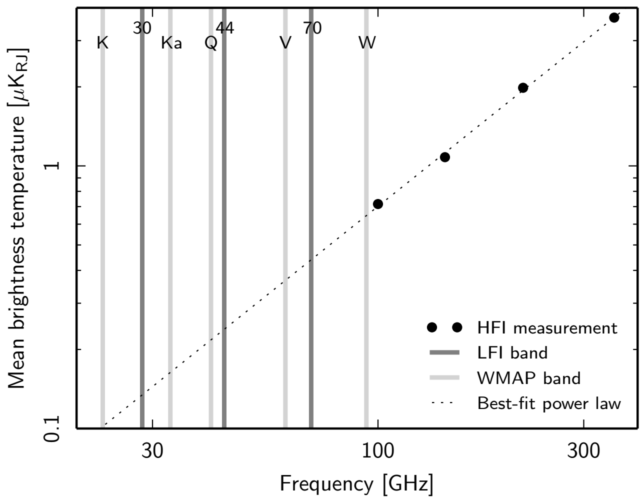

Fig. 1. Zodiacal light extrapolation from HFI to LFI and WMAP frequency channels in terms of full-sky mean brightness temperature. The dotted line shows the power-law fit to the HFI observations between 100 and 353 GHz, $s(\nu) = 0.70\,\mu\mathrm{K_{RJ}}\,(\nu/100\,\mathrm{GHz})^{1.31}$, and the vertical grey lines indicate the central frequencies of the LFI and WMAP frequency bands.

For the *Planck* HFI channels, a model of zodiacal light emission is subtracted from the time-ordered data prior to mapmaking (Planck Collaboration VIII 2016). In addition, in this paper we apply a small correction to the low-frequency LFI and WMAP channels by scaling the effective HFI 100 GHz zodiacal light correction map (i.e., uncorrected minus corrected map) to each frequency according to a power law fitted to frequencies between 100 and 353 GHz (Planck Collaboration XIV 2014), as illustrated in Fig. 1; the actual template amplitudes relative to the 100 GHz correction map (in thermodynamic units) are listed in Table 2. Although the magnitude of this correction is small, with a maximum amplitude of 2 μK in the 70 GHz map, applying no correction at all below 100 GHz results in a visually noticeable bias in the derived CO $J = 1 \rightarrow 0$ map at high Galactic latitudes, in the characteristic form of the zodiacal light. Extending the zodiacal light model to low frequencies efficiently eliminates this structure.

In our 2013 release, colour corrections and unit conversions for all *Planck* channels were based on individual bandpass profiles as measured on the ground before launch (Planck Collaboration V 2014; Planck Collaboration IX 2014). However, as discussed in detail in Sects. 2, 4.3, and 5, during the component separation process we find that systematic uncertainties in the nominal bandpasses induce significant residuals between data and model, and it is necessary to fit for these bandpass uncertainties in order to obtain statistically acceptable fits. For WMAP, we adopt the nominal bandpasses[4] for the first DA within each frequency band, and fit for the remaining DA bandpasses within each frequency. For the 408 MHz survey, we adopt a delta function response at the nominal frequency. The unit conversion factors between thermodynamic and brightness

[2] http://lambda.gsfc.nasa.gov

[3] For uniformity, we refer to the 70 GHz horn pair maps as "detector set" maps in this paper, with the {ds1, ds2, ds3} maps corresponding to horns {18+23, 19+22, 20+21}, respectively.

[4] As described by Bennett et al. (2013), the WMAP bandpasses evolved during the 9 yr of WMAP observations, resulting in slightly lower effective full-mission frequencies as compared to the nominal bandpasses. We correct for these small shifts in the present analysis by shifting the bandpasses accordingly.