



Publication Year	2020
Acceptance in OA	2021-12-17T10:55:28Z
Title	Planck 2018 results. IV. Diffuse component separation
Authors	Planck Collaboration, Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Benabed, K., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J. -M., Lasenby, A., Lattanzi, M., Lawrence, C. R., Le Jeune, M., McEwen, J. D., Salvati, L., Levrier, F., Liguori, M., Lilje, P. B., Lindholm, V., López-Caniego, M., Lubin, P. M., Ma, Y. -Z., Macías-Pérez, J. F., Maggio, G., Maino, D., Sandri, M., Meinhold, P. R., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P. G., Melchiorri, A., Mennella, A., Migliaccio, M., Miville-Deschênes, M. -A., Savelainen, M., Molinari, D., Bersanelli, M., Moneti, A., Montier, L., Morgante, G., Natoli, P., Oppizzi, F., Pagano, L., PAOLETTI, DANIELA, Partridge, B., Scott, D., Peel, M., Pettorino, V., Bielewicz, P., Piacentini, F., Polenta, G., Puget, J. -L., Rachen, J. P., Reinecke, M., Remazeilles, M., Renzi, A., Seljebotn, D. S., Rocha, G., Roudier, G., Rubiño-Martín, J. A., Bond, J. R., Ruiz-Granados, B., Sirignano, C., Spencer, L. D., Suur-Uski, A. -S., Tauber, J. A., Borrill, J., Fernandez-Cobos, R., Tavagnacco, D., Tenti, M., Thommesen, H., Toffolatti, L., Tomasi, M., Trombetti, T., Valiviita, J., Van Tent, B., Vielva, P., VILLA, FABRIZIO, Finelli, F., Bouchet, F. R., Vittorio, N., Wandelt, B. D., Wehus, I. K., ZACCHEI, Andrea, Zonca, A., Boulanger, F., Bucher, M., BURIGANA, CARLO, Calabrese, E., Forastieri, F., Cardoso, J. -F., Carron, J., Casaponsa, B., Challinor, A., Colombo, L. P. L., Combet, C., Crill, B. P., CUTTAIA, FRANCESCO, de Bernardis, P., De Rosa, A., Frailis, M., de Zotti, G., Delabrouille, J., Delouis, J. -M., Di Valentino, E., Dickinson, C., Diego, J. M., Donzelli, S., Doré, O., Ducout, A., Dupac, X., Fraisse, A. A., Efstathiou, G., Elsner, F., Enßlin, T. A., Eriksen, H. K., Falgarone, E., FRANCESCHI, ENRICO, Frolov, A., GALEOTTA, Samuele, Galli, S., Ganga, K., Martínez-González, E., Génova-Santos, R. T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K. M., Gratton, S., GRUPPUSO, ALESSANDRO, Gudmundsson, J. E., Handley, W., Hansen, F. K., Matarrese, S., Helou, G., Herranz, D., Hildebrandt, S. R., Huang, Z., Jaffe, A. H., Karakci, A., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Mauri, N., Kisner, T. S.
Publisher's version (DOI)	10.1051/0004-6361/201833881
Handle	http://hdl.handle.net/20.500.12386/31239
Journal	ASTRONOMY & ASTROPHYSICS
Volume	641

Planck 2018 results

IV. Diffuse component separation

Planck Collaboration: Y. Akrami^{14,51,53}, M. Ashdown^{60,5}, J. Aumont⁸⁵, C. Baccigalupi⁷¹, M. Ballardini^{20,37}, A. J. Banday^{85,8}, R. B. Barreiro⁵⁵, N. Bartolo^{26,56}, S. Basak⁷⁷, K. Benabed^{50,84}, M. Bersanelli^{29,41}, P. Bielewicz^{70,69,71}, J. R. Bond⁷, J. Borrill^{12,82}, F. R. Bouchet^{50,80}, F. Boulanger^{79,49,50}, M. Bucher^{2,6}, C. Burigana^{40,27,43}, E. Calabrese⁷⁵, J.-F. Cardoso⁵⁰, J. Carron²¹, B. Casaponsa⁵⁵, A. Challinor^{52,60,11}, L. P. L. Colombo²⁹, C. Combet⁶², B. P. Crill^{57,10}, F. Cuttaia³⁷, P. de Bernardis²⁸, A. de Rosa³⁷, G. de Zotti³⁸, J. Delabrouille², J.-M. Delouis^{50,84}, E. Di Valentino⁵⁸, C. Dickinson⁵⁸, J. M. Diego⁵⁵, S. Donzelli^{41,29}, O. Doré^{57,10}, A. Ducout⁶¹, X. Dupac³², G. Efstathiou^{60,52}, F. Elsner⁶⁶, T. A. Enßlin⁶⁶, H. K. Eriksen^{53,*}, E. Falgarone⁷⁹, R. Fernandez-Cobos⁵⁵, F. Finelli^{37,43}, F. Forastieri^{27,44}, M. Frailis³⁹, A. A. Fraisse²³, E. Franceschi³⁷, A. Frolov⁷⁸, S. Galeotta³⁹, S. Galli⁵⁹, K. Ganga², R. T. Génova-Santos^{54,15}, M. Gerbino⁸³, T. Ghosh^{74,9}, J. González-Nuevo¹⁶, K. M. Górski^{57,86}, S. Gratton^{60,52}, A. Gruppuso^{37,43}, J. E. Gudmundsson^{83,23}, W. Handley^{60,5}, F. K. Hansen⁵³, G. Helou¹⁰, D. Herranz⁵⁵, S. R. Hildebrandt^{57,10}, Z. Huang⁷⁶, A. H. Jaffe⁴⁸, A. Karakci⁵³, E. Keihänen²², R. Keskitalo¹², K. Kiiveri^{22,36}, J. Kim⁶⁶, T. S. Kisner⁶⁴, N. Krachmalnicoff⁷¹, M. Kunz^{13,49,3}, H. Kurki-Suonio^{22,36}, G. Lagache⁴, J.-M. Lamarre⁷⁹, A. Lasenby^{5,60}, M. Lattanzi^{27,44}, C. R. Lawrence⁵⁷, M. Le Jeune², F. Levrier⁷⁹, M. Liguori^{26,56}, P. B. Lilje⁵³, V. Lindholm^{22,36}, M. López-Caniego³², P. M. Lubin²⁴, Y.-Z. Ma^{58,73,68}, J. F. Macías-Pérez⁶², G. Maggio³⁹, D. Maino^{29,41,45}, N. Mandolesi^{37,27}, A. Mangilli⁸, A. Marcos-Caballero⁵⁵, M. Maris³⁹, P. G. Martin⁷, E. Martínez-González⁵⁵, S. Matarrese^{26,56,34}, N. Mauri⁴³, J. D. McEwen⁶⁷, P. R. Meinhold²⁴, A. Melchiorri^{28,46}, A. Mennella^{29,41}, M. Migliaccio^{31,47}, M.-A. Miville-Deschênes^{1,49}, D. Molinari^{27,37,44}, A. Moneti⁵⁰, L. Montier^{85,8}, G. Morgante³⁷, P. Natoli^{27,81,44}, F. Oppizzi²⁶, L. Pagano^{49,79}, D. Paoletti^{37,43}, B. Partridge³⁵, M. Peel^{17,58}, V. Pettorino¹, F. Piacentini²⁸, G. Polenta⁸¹, J.-L. Puget^{49,50}, J. P. Rachen¹⁸, M. Reinecke⁶⁶, M. Remazeilles⁵⁸, A. Renzi⁵⁶, G. Rocha^{57,10}, G. Roudier^{2,79,57}, J. A. Rubiño-Martín^{54,15}, B. Ruiz-Granados^{54,15}, L. Salvati⁴⁹, M. Sandri³⁷, M. Savelainen^{22,36,65}, D. Scott¹⁹, D. S. Seljebo⁵³, C. Sirignano^{26,56}, L. D. Spencer⁷⁵, A.-S. Suur-Uski^{22,36}, J. A. Tauber³³, D. Tavagnacco^{39,30}, M. Tenti⁴², H. Thommesen⁵³, L. Toffolatti^{16,37}, M. Tomasi^{29,41}, T. Trombetti^{40,44}, J. Valiviita^{22,36}, B. Van Tent⁶³, P. Vielva⁵⁵, F. Villa³⁷, N. Vittorio³¹, B. D. Wandelt^{50,84,25}, I. K. Wehus⁵³, A. Zacchei³⁹, A. Zonca⁷²

(Affiliations can be found after the references)

Received 16 July 2018 / Accepted 13 January 2019

ABSTRACT

We present full-sky maps of the cosmic microwave background (CMB) and polarized synchrotron and thermal dust emission, derived from the third set of *Planck* frequency maps. These products have significantly lower contamination from instrumental systematic effects than previous versions. The methodologies used to derive these maps follow closely those described in earlier papers, adopting four methods (Commander, NILC, SEVEM, and SMICA) to extract the CMB component, as well as three methods (Commander, GNILC, and SMICA) to extract astrophysical components. Our revised CMB temperature maps agree with corresponding products in the *Planck* 2015 delivery, whereas the polarization maps exhibit significantly lower large-scale power, reflecting the improved data processing described in companion papers; however, the noise properties of the resulting data products are complicated, and the best available end-to-end simulations exhibit relative biases with respect to the data at the few percent level. Using these maps, we are for the first time able to fit the spectral index of thermal dust independently over 3° regions. We derive a conservative estimate of the mean spectral index of polarized thermal dust emission of $\beta_d = 1.55 \pm 0.05$, where the uncertainty marginalizes both over all known systematic uncertainties and different estimation techniques. For polarized synchrotron emission, we find a mean spectral index of $\beta_s = -3.1 \pm 0.1$, consistent with previously reported measurements. We note that the current data processing does not allow for construction of unbiased single-bolometer maps, and this limits our ability to extract CO emission and correlated components. The foreground results for intensity derived in this paper therefore do not supersede corresponding *Planck* 2015 products. For polarization the new results supersede the corresponding 2015 products in all respects.

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

1. Introduction

This paper, one of a set associated with the 2018 release of data from the *Planck*¹ mission (Planck Collaboration I 2016), describes the cosmological and astrophysical component

maps derived from the full set of *Planck* observations (Planck Collaboration I 2020), and compares these to earlier versions of the corresponding products. *Planck* was launched on 14 May 2009, and observed the sky nearly without interruption for four years. The raw, time-ordered observations were released to the public in their entirety in February 2015 as part of the second *Planck* data release (PR2), together with associated frequency and component sky maps and higher-level science data products, including cosmic microwave background (CMB) power spectra and cosmological parameters. These observations represent a cornerstone of modern cosmology, and they severely constrain the history of the early Universe.

The time-ordered data selection adopted for the current (third, PR3) release is similar to that used in the second release

* Corresponding author: H. K. Eriksen,
 e-mail: h.k.k.eriksen@astro.uio.no

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).

(Planck Collaboration II 2020; Planck Collaboration III 2020); the second and third *Planck* product deliveries therefore have nearly identical scientific constraining power, as measured in terms of raw integration time and instrumental noise levels. The difference between the two releases lies in their overall levels of instrumental systematic uncertainties and calibration. A substantial fraction of the second-release papers was dedicated to identifying, quantifying, and characterizing residual uncertainties due to a wide range of instrumental effects, including effective gain variations, analogue-to-digital converter (ADC) nonlinearities, residual temporal transfer functions, and foreground bandpass leakage. Indeed, these residuals were sufficiently large to prohibit extraction of a robust polarization signal on large angular scales from the *Planck* High Frequency Instrument (HFI) observations, significantly limiting the science scope of the *Planck* polarization observations as a whole. Fortunately, as discussed extensively in Planck Collaboration III (2020), these residuals are now not only better understood and modelled, but also greatly reduced in the final dataset, particularly through the use of improved end-to-end processing techniques.

In this paper, we present updated full-sky CMB maps in both temperature and polarization, as well as new synchrotron and thermal dust emission maps in polarization, and compare these to previous versions (Planck Collaboration XII 2014; Planck Collaboration IX 2016; Planck Collaboration X 2016). In terms of temperature foreground products, we provide an update of the Generalized Needlet Internal Linear Combination (GNILC; Remazeilles et al. 2011a) thermal dust model, to be used in conjunction with the updated 2018 GNILC polarization map, but no new Commander (Eriksen et al. 2008) foreground products. The reason for this is one of necessity: as described in Planck Collaboration III (2020), the latest HFI processing exploits the full information content of each frequency in order to suppress large-scale polarization systematics, and the processing has thus been tuned to optimize the polarization solution. The cost of this choice, however, is that individual single-bolometer maps are no longer available; see Sect. 3.1.2 of Planck Collaboration III (2020) for details. Specifically, some of the single-bolometer maps only contain part of the sky signal and thus cannot be used for component separation. This, in turn, has an impact on the ability of the Commander algorithm to resolve individual foreground components in temperature. The single most important effect is on our ability to constrain CO line emission, which benefits particularly strongly from intra-frequency measurements. Because each unfiltered bolometer in principle has a different bandpass amplitude at the CO-line centre frequency of 115.27 GHz (and multiples thereof), each bolometer observes the true CO signal with different effective responses, and these differences provide a strong handle on the true intensity of the CO signal. Furthermore, both thermal dust and free-free emission correlate strongly with CO emission, and are therefore also negatively affected by the lack of single-bolometer maps. In turn, free-free emission is strongly correlated with both synchrotron and anomalous microwave emission. In summary, we believe that the *Planck* 2015 Commander-based temperature (i.e., Stokes I) foreground model represents a more accurate description of the true temperature sky than what can be extracted from the current (2018) data set. To avoid confusion, we therefore do not release the latest version publicly, although we compare the two models in Sect. 5. For the CMB component, we find that the latest processing produces results that are fully consistent with the previous incarnation, while for polarization the new results represent a major improvement, both in terms of CMB and foregrounds.

The methodologies adopted in this paper mirror those used in earlier *Planck* releases, with only minor algorithmic updates and improvements. In particular, for CMB extraction we adopt the same four component-separation implementations used in earlier releases, namely Commander, NILC, SEVEM, and SMICA, each of which was initially selected as a representative of a particular class of algorithms (blind versus non-blind methods and pixel-based versus harmonic-based methods). In combination, they represent most approaches proposed in the literature. In the current release, all four CMB methods adopt the same data selection, based only on full-frequency *Planck* maps, in order to facilitate a direct comparison of the results. As in previous releases, we strongly suggest considering all four CMB maps in any higher-level map-based CMB analysis, in order to assess robustness with respect to algorithmic choices. We also provide again cleaned CMB maps at individual frequencies constructed by SEVEM. More specifically, in this release, intensity and polarization CMB maps are produced at four different frequencies from 70 to 217 GHz. These maps are particularly useful to test, for example, the robustness of results versus the presence of foregrounds and/or systematics. In addition, one fundamentally new data product is delivered in this release, namely a CMB temperature map generated by SMICA from which Sunyaev–Zeldovich (SZ) sources have been projected out. This can be used, for instance, in lensing studies (Planck Collaboration VIII 2020).

For astrophysical component separation, which depends inherently on explicit parametric modelling, we adopt Commander as our primary computational engine, mirroring the processing adopted in the two previous *Planck* releases. However, since the last release the internal mechanics of this code have been significantly re-written. Commander now allows for analysis of data sets with different angular resolutions at each frequency, and thereby allows for production of frequency maps at the full angular resolution of the data (Seljebotn et al. 2019). In addition, we employ both GNILC and SMICA for foreground reconstruction in the new release.

The rest of the paper is organized as follows. Section 2 reviews the algorithms and methods used in the analysis, focusing primarily on updates and improvements made since the 2015 release. Section 3 describes the data selection and pre-processing steps applied to the data before analysis. Section 4 presents the *Planck* 2018 CMB maps in both temperature and polarization, and characterizes their properties in terms of residuals with respect to earlier versions, along with angular power spectra, cosmological parameters, and simple higher-order statistics. Section 5 discusses the updated polarization foreground products. Section 6 gives conclusions. The various algorithms are reviewed in Appendices A–E. A brief summary of temperature foregrounds derived from the *Planck* 2018 frequency maps is provided in Appendix F and, finally, additional CMB figures are provided in Appendices G and H.

2. Component-separation methods

Earlier publications give detailed descriptions of the four main component-separation methods used in this paper (Planck Collaboration XII 2014; Planck Collaboration IX 2016; Planck Collaboration X 2016). For some methods, notable improvements have been implemented since the last release, and these are described below. Further technical details may be found in the Appendices.

We also employ the GNILC algorithm for thermal dust extraction. This method and corresponding results are described in detail

in Remazeilles et al. (2011a), Planck Collaboration Int. XLVIII (2016), and Planck Collaboration XII (2020). A detailed comparison of the foreground products derived with Commander and GNILC is presented in the current paper.

2.1. Commander

Commander (Eriksen et al. 2004, 2008; Planck Collaboration X 2016) has undergone the most significant changes since the previous release. Commander is a Bayesian approach employing a Monte Carlo method called Gibbs sampling as its central computational engine. Within this Bayesian framework, a parametric model is fitted to the data set in question with standard posterior sampling or maximization techniques, including cosmological, astrophysical, and instrumental parameters.

We start by writing down a generic model on the form,

$$\mathbf{d}_\nu(p) = g_\nu \sum_{c=1}^{N_c} \mathbf{F}_\nu(\beta_c) \mathbf{T}(p) \mathbf{a}_c + \mathbf{n}_\nu(p). \quad (1)$$

Here $\mathbf{d}_\nu(p)$ denotes the observed data at frequency ν and pixel p . The sum runs over N_c components, each with an amplitude vector $\mathbf{a}_c$, a map projection operator $\mathbf{T}(p)$, and frequency scaling operator $\mathbf{F}_\nu(\beta_c)$ that depends on astrophysical spectral parameters β_c . The quantity g_ν denotes an overall instrumental calibration factor per frequency channel, and $\mathbf{n}_\nu(p)$ indicates instrumental noise. With this notation, the component sum runs over both astrophysical components (CMB, synchrotron, CO, thermal dust emission etc.) and possible spurious monopole and dipole terms. The projection operator $\mathbf{T}$ indicates any step required in going from a general amplitude vector (such as a pixelized sky map, a set of spherical harmonic coefficients, or a template amplitude) to a map as observed by the current detector. Thus, this matrix encodes both the choice of basis vectors (pixels, spherical harmonics, templates) and higher-level operations such as beam convolution. Given this data model, samples are drawn from the full posterior as described in Eriksen et al. (2004, 2008) and Seljebotn et al. (2019).

In previous releases the above model was fitted to the combination of *Planck* and external data using the Commander implementation described by Eriksen et al. (2008). This implementation adopted map-space pixels as its basis set for astrophysical foregrounds, for coding efficiency reasons. Although computationally fast, this approach has a significant limitation in that it requires all data sets under consideration to have the same angular resolution. Specifically, this implies that the angular resolution of the final output maps are limited to that of the lowest resolution frequency channel under consideration, which typically is 1° FWHM for the combination of *Planck*, WMAP, and Haslam 408 MHz, which formed the basis of the previous astrophysically oriented foreground analysis. Higher-resolution products could then only be derived by dropping lower-resolution channels, which in turn carried a significant cost in terms of model fidelity.

In the current release, we implement the Commander algorithm described by Seljebotn et al. (2019), which we refer to as Commander2. This approach, which models the foreground amplitude maps in terms of spherical harmonics instead of pixels, offers three important improvements over the pixel-based approach.

First, since amplitudes are modelled in harmonic space, it is computationally trivial to convolve with a separate instrumental beam transfer function at separate frequencies, so that for the first time we can solve for full-resolution signal models with multi-

resolution data sets. Commander2 is thus able to produce a foreground model at native *Planck* resolution, limited only by the effective signal-to-noise ratio of each component. The computational cost is greater; however, as shown by Seljebotn et al. (2019), this is manageable with modern computers, even for *Planck*-sized data sets.

Second, the new approach offers the option of imposing a prior on the foreground signal amplitudes in the form of an angular power spectrum. This can be used to regularize the foreground solution at small angular scales, and thereby reduce degeneracies between different components at high multipoles.

Third, the improvements allow for joint fitting of compact or unresolved objects and diffuse components. This improves the reconstruction of the diffuse components themselves, including the CMB, and also allows production of a new catalogue of compact objects. The details of this procedure are described in Appendix A.

Overall, from an algorithmic point of view the Commander2 implementation used in the current data release is more powerful than in previous releases. At the same time, there is also one important aspect of the *Planck* 2018 data release that limits our ability to perform a component separation as detailed as that in the 2015 analysis. As mentioned in Sect. 1, the *Planck* 2018 data set includes only full-frequency maps, not single-bolometer maps. For the Commander temperature analysis, this implies that a simpler foreground model must be employed than in the corresponding 2015 analysis. In the previous analysis we considered seven different physical components, namely CMB, synchrotron, free-free, spinning and thermal dust emission, a general line emission component at 95 and 100 GHz, and CO with individual components at 100, 217, and 353 GHz. Single-detector maps played a central part in constraining this rich model, in particular with respect to CO line emission. With the new and more limited data set, we instead adopt a similar model as employed in the 2013 analysis, which includes only four diffuse signal components in temperature, namely CMB, a single general low-frequency power-law component, thermal dust, and a single CO component with spatially constant line ratios between 100, 217, and 353 GHz. For polarization the model remains the same as in 2015, and includes only CMB, synchrotron, and thermal dust emission. The latter two components are as usual modelled in terms of simple power-law and modified blackbody SEDs, respectively.

The above general specification provides a basic summary of the framework used for parametric fitting. However, there are still some free choices that must be made, the two most important of which are: (1) the angular resolution of the foreground spectral indices; and (2) the spatial priors imposed on the foreground amplitudes. For the spectral indices, we are primarily driven by signal-to-noise considerations, as adopting too high resolution for such parameters leads to an undesirable increase in noise in all components. In the temperature case, we adopt a smoothing scale of $40'$ FWHM for low-frequency foregrounds, slightly larger than the 30 GHz instrumental beam. For the dust spectral index, we adopt $10'$ FWHM, which is slightly larger than the 100 GHz beam. The dust temperature is fitted at the full *Planck* resolution of $5'$ FWHM of the frequencies between 217 and 857 GHz. For polarization, we fit only a spatially-constant spectral index for synchrotron², while for thermal dust

² Note that the numerical value derived for the spectral index of polarized synchrotron emission is not directly comparable to the mean of the low-frequency component spectral index map derived in intensity, since the latter also includes free-free and spinning dust emission.

emission, we fit the dust spectral index at 3° FWHM. The dust temperature for the polarization model is fixed at the values derived in the intensity analysis, as the *Planck* 545 and 857 GHz frequency channels are unpolarized, and the *Planck* observations therefore do not constrain the thermal dust temperature in polarization.

Finally, for spatial priors, we adopt minimally informative power-spectrum priors, defined simply as flat spectra in units of $C_\ell \ell(\ell+1)/2\pi$ for all components, with an amplitude that is larger than that observed in the high signal-to-noise regime. In addition, this flat spectrum is smoothly apodized at high multipoles in order to suppress ringing around bright compact objects. For the low-frequency temperature foreground and the CO line-emission components, the apodization is performed with a Gaussian beam with a FWHM roughly matching the dominant frequency for the respective component, while for thermal dust only a mild apodization is applied in the form of an exponentially-falling cut-off between $\ell = 5000$ and 6000 . For polarization, we apodize with Gaussian smoothing kernels, as in the low-frequency foreground and CO case³. Full details regarding these choices are summarized in Appendix A.

2.2. NILC

NILC (Needlet Internal Linear Combination) is described by Basak & Delabrouille (2012, 2013). The overall goal of NILC is to extract the CMB component from multi-frequency observations while minimizing the contamination from Galactic and extragalactic foregrounds and instrumental noise. This is done by computing the linear combination of input maps that minimizes the variance in a basis spanned by a particular class of spherical wavelets called needlets (Narcowich et al. 2006). Needlets allow localized filtering in both pixel space and harmonic space. Localization in pixel space allows the weights of the linear combination to adapt to local conditions of foreground contamination and noise, whereas localization in harmonic space allows the method to favour foreground rejection on large scales and noise rejection on small scales. Needlets permit the weights to vary smoothly on large scales and rapidly on small scales, which is not possible by cutting the sky into zones prior to processing (Delabrouille et al. 2009). The NILC pipeline is applicable to scalar fields on the sphere, hence we work separately on maps of temperature and the E and B modes of polarization. The decomposition of input polarization maps into E and B is done on the full sky. At the end, the CMB Q and U maps are reconstructed from the E and B maps. Further details of the method are provided in Appendix B.

The NILC pipeline employed in the *Planck* 2018 analysis is essentially unchanged from that employed in the 2015 analysis; we therefore refer to Planck Collaboration IX (2016) and references therein for full details.

2.3. SEVEM

SEVEM (Leach et al. 2008; Fernández-Cobos et al. 2012) is an implementation of an internal template-cleaning approach in real space. It has been used in the previous *Planck* releases to produce clean CMB maps in both intensity and polarization, and has been demonstrated to provide robust results. A detailed description of the SEVEM pipeline can be found in Appendix C.

The starting point for SEVEM is a set of internal templates typically constructed as difference maps between two neighboring *Planck* channels convolved to the same resolution, ensuring that the CMB signal vanishes. These templates trace the foreground contaminants at the corresponding frequency ranges. Next, a linear combination of such templates is then subtracted from some set of CMB-dominated frequency maps, typically 70–217 GHz for *Planck*. The coefficients of the linear fit are derived by minimizing the variance of the clean map outside a given mask. A final, co-added CMB map is obtained by combining individually-cleaned frequency maps in harmonic space.

SEVEM is also able to produce cleaned CMB maps at specific channels. Individually-cleaned frequency CMB maps are useful to test the robustness of results versus the presence of foregrounds and/or systematics, for instance for isotropy and statistics estimators (Planck Collaboration XIV 2016) or the integrated Sachs-Wolfe effect stacking analysis (Planck Collaboration XXI 2016). They are also valuable to construct cross-frequency estimators, which allow one to minimize the impact of certain types of systematic effects (e.g., possible correlated noise in data splits). In addition, they can be used to search for frequency-dependent effects in the CMB itself, such as those arising from relativistic boosting (Planck Collaboration XXVII 2014) or the Sunyaev–Zeldovich effect (Sunyaev & Zeldovich 1970), although for this type of analysis the contribution of the templates (which would contain a certain level of any effect that is not constant with frequency) to the cleaned maps should be taken into account.

Since the 2015 release, we have introduced two significant improvements to the SEVEM pipeline for polarization. First, in the previous release we produced cleaned maps at three frequencies, 70, 100, and 143 GHz, and the final map was produced by combining the cleaned 100 and 143 GHz maps. However, given the improvements in the new *Planck* polarization data, we are now also able to robustly clean the 217 GHz channel map, and this is now included in the final combination. As a result, the signal-to-noise ratio of the cleaned SEVEM CMB polarization map is significantly improved with respect to the previous version. Second, in the updated pipeline, we now produce polarization maps at full resolution ($N_{\text{side}} = 2048$), whereas in the last release all polarization maps were constructed at $N_{\text{side}} = 1024$. However, recognizing the fact that the 217 GHz channel is likely to be somewhat more susceptible to large-scale systematic residuals and calibration uncertainties due its higher foreground levels than the two lower frequencies (Planck Collaboration III 2020), we introduce at the same time a relative down-weighting of the 217 GHz channel on the largest scales. In summary, these modifications yield significantly improved SEVEM polarization maps, both in terms of the combined CMB map and individually cleaned frequency maps. Regarding intensity, the SEVEM pipeline is essentially identical to that used in the previous release; however, we now also provide a cleaned 70 GHz map in intensity. In addition to the final CMB map, SEVEM therefore now provides the complete set of $\{T, Q, U\}$ CMB maps for each of the four frequency channels between 70 and 217 GHz.

2.4. SMICA

SMICA (Spectral Matching Independent Component Analysis) is described in Cardoso et al. (2008), and details regarding the actual implementation used in the following analysis (pre-processing, masking and mask correction, beam correction, binning, possible re-calibration, etc.) are provided in Appendix D.

³ The decision on whether to use a Gaussian kernel or a mild exponential high- ℓ cut-off for prior apodization is determined by the effective signal-to-noise ratio of the component in question.

SMICA synthesizes CMB $\{T, E, B\}$ maps from spherical harmonic coefficients $\hat{s}_{\ell m}$ obtained by combining the coefficients of N_{cha} frequency maps with an ℓ -dependent $N_{\text{cha}} \times 1$ vector of weights $\mathbf{w}_\ell$,

$$\hat{s}_{\ell m} = \mathbf{w}_\ell^\dagger \mathbf{x}_{\ell m} \quad \text{where} \quad \mathbf{w}_\ell = \frac{\mathbf{C}_\ell^{-1} \mathbf{a}}{\mathbf{a}^\dagger \mathbf{C}_\ell^{-1} \mathbf{a}}. \quad (2)$$

Here the $N_{\text{cha}} \times 1$ vector $\mathbf{a}$ describes the emission law of the CMB, and the $N_{\text{cha}} \times N_{\text{cha}}$ spectral covariance matrix $\mathbf{C}_\ell$ contains (estimates of) all auto- and cross-spectra of the N_{cha} input maps. On small angular scales, where a large number of harmonic coefficients are available, $\mathbf{C}_\ell$ may be accurately estimated as

$$\hat{\mathbf{C}}_\ell = \frac{1}{2\ell + 1} \sum_m \mathbf{x}_{\ell m} \mathbf{x}_{\ell m}^\dagger, \quad (3)$$

which is used “as is” in Eq. (2). On large angular scales, we resort to a parametric model $\mathbf{C}_\ell(\theta)$ of the spectral covariance matrices in order to reduce the estimation variance and mitigate the effects of chance correlation between the CMB field and the foregrounds. The model is adjusted to the data by selecting best-fit parameters θ obtained as

$$\hat{\theta} = \arg \min_{\theta} \sum_{\ell} (2\ell + 1) \left[\text{Tr}(\hat{\mathbf{C}}_\ell \mathbf{C}_\ell(\theta)^{-1}) + \log \det \mathbf{C}_\ell(\theta) \right]. \quad (4)$$

The minimization in Eq. (4) is equivalent to maximizing the joint likelihood of the N_{cha} input maps assuming that they follow a Gaussian isotropic distribution characterized by the spectra and cross-spectra collected in the spectral covariance matrices $\mathbf{C}_\ell(\theta)$. For a motivation of this likelihood, see Cardoso (2017).

The spectral model fitted by SMICA, $\mathbf{C}_\ell(\theta)$, is agnostic, as it assumes only that the foreground emission can be described by an unconstrained N_{fg} -dimensional component with a covariance matrix of the form

$$\mathbf{C}_\ell(\theta) = \begin{bmatrix} \mathbf{a} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \mathbf{C}_\ell^{\text{cmb}} & \mathbf{0} \\ \mathbf{0} & \mathbf{P}_\ell \end{bmatrix} \begin{bmatrix} \mathbf{a} & \mathbf{F} \end{bmatrix}^\dagger + \mathbf{N}_\ell. \quad (5)$$

Here the $N_{\text{cha}} \times N_{\text{fg}}$ matrix $\mathbf{F}$ represents the foreground emissivities, which are ℓ -independent, and the $N_{\text{fg}} \times N_{\text{fg}}$ matrix $\mathbf{P}_\ell$ contains the foreground auto- and cross-spectra. The diagonal matrix $\mathbf{N}_\ell$ represents the noise contribution, and θ contains whatever parameters are needed to determine the quantities $\mathbf{C}_\ell^{\text{cmb}}$, $\mathbf{a}$, $\mathbf{F}$, $\mathbf{P}_\ell$, and $\text{diag}(\mathbf{N}_\ell)$. In most cases, a SMICA fit is conducted with $\mathbf{a}$ fixed to assumed known values (i.e., assuming perfect calibration) and leaving all other parameters free. $\mathbf{P}_\ell$ is only constrained to be positive. In other words, foreground spectra (emissivities and angular spectral behaviour) and their correlations are freely fitted by SMICA.

In this release, however, we also consider two variations that include constraints on foreground emissions. The first of these is used to produce an SZ-free CMB map in intensity (see Appendix D) used in Planck Collaboration VIII (2020), and the second results in thermal dust and synchrotron maps in polarization (see Sect. 5). No attempt is made to reconstruct temperature foregrounds, since the combination of synchrotron, free-free, spinning and thermal dust, and CO emission is intrinsically much more tightly coupled and difficult to disentangle than synchrotron and thermal dust emission in polarization.

Since the last release, changes have been introduced for both intensity and polarization maps. Starting with the temperature case, the most important change in this release is the introduction of hybrid CMB rendering, merging two different CMB maps produced independently by the SMICA pipeline. The first CMB




Fig. 1. SMICA harmonic weights used to obtain the temperature X_{high} map (top panel), X_{full} map (middle panel), and polarization map (bottom panel).

map, X_{high} , is designed to describe the cleanest region of the sky and intermediate-to-small angular scales. It is obtained from all six HFI channels using a foreground dimension of $N_{\text{fg}} = 4$. The second CMB map, X_{full} , is designed to describe the full sky and all harmonic scales. It includes all nine *Planck* frequency channels using a maximal foreground dimension of $N_{\text{fg}} = 8$. The final hybrid CMB map X is then computed by merging X_{high} and X_{full} according to

$$X = \mathcal{P} X_{\text{high}} + (\mathbf{I} - \mathcal{P}) X_{\text{full}} = X_{\text{full}} + \mathcal{P}(X_{\text{high}} - X_{\text{full}}), \quad (6)$$

where $\mathcal{P}$ is a linear operator that smoothly removes large harmonic scales, and masks out an area close to the Galactic plane. Hence, in the resulting hybridized map, the multipoles of highest degree and the areas of highest Galactic latitude are provided

by X_{high} , while the remaining information is provided by X_{full} . In practice, the hybridization operator $\mathcal{P}$ is implemented by high-pass filtering in the harmonic domain (with a transfer function that smoothly transitions from 0 to 1 according to an arc-cosine function over the multipole range $50 \leq \ell \leq 150$), followed by multiplication by an apodized Galactic mask that is similar to the mask used at 100 GHz in the *Planck* 2018 likelihood (Plik) (see [Planck Collaboration V 2020](#), for details).

Hybridization of two CMB renderings has several benefits compared to using a single set of harmonic weights over all areas of the sky. First, the data suggest it: the SMICA weights are quite different if they are based on spectral statistics computed over the full sky rather than over a region with much lower foreground contamination. This is the rationale behind NILC, which extends the idea to many more than the two sky regions considered regions considered by SMICA. Second, the reason for leaving out the LFI channels in producing X_{high} except at large angular scales is that SMICA would put very small weights on those channels (this is not the case when the weights are based on statistics computed for X_{full} , as seen on Fig. 1 which shows a significant contribution from the 70 GHz channel). We could still include those channels and let SMICA automatically down-weight them, but by excluding the channels with the lowest resolution, we avoid large, “low-resolution” holes in the common point source mask, and therefore in the final CMB map. Finally, hybridization matches well the high- ℓ TT likelihood function in Plik, uses a low-foreground-contaminated fraction of the sky, does not include LFI channels, and involves only high frequency foregrounds. The spectral weights, w_ℓ , for temperature (both full-sky and high latitudes) and polarization are shown in Fig. 1.

SMICA adopts its own relative calibration between frequency channels. In 2015, this process was applied to frequency channels from 44 to 353 GHz; however, since then we have found that the uncertainty in the 44 GHz channel was larger than expected, and that the previously reported value was inaccurate (see Fig. D.5). In the new release, we adopt a more conservative approach, and limit re-calibration to 70, 100, and 217 GHz, taking the 143 GHz channel as a reference; see Appendix D.1 for further details.

For polarization, we have introduced two changes since the previous release. First, the CMB polarization maps are now generated by independently processing E and B modes, while in 2015 they were jointly fitted and filtered. Second, we run two independent SMICA fits, one targeted at CMB extraction, the other at foreground separation.

For CMB extraction, we conduct a fit using a maximal foreground dimension of $N_{\text{fg}} = 7 - 1 = 6$, which makes $[a \text{ F}]$ a square matrix. This is the largest dimension supported blindly (i.e., without any constraint on the foreground contribution) by SMICA, given the number of available polarized channels.

For foreground separation, we conduct a separate fit using a foreground model of dimension $N_{\text{fg}} = 2$, implicitly targeting synchrotron and dust emissions. The degeneracy of the SMICA foreground model (Eq. (5)) can then be fixed by requesting that synchrotron (thermal dust) emission should be negligible at 353 GHz (30 GHz); Appendix D describes the implementation details. This analysis yields, without any other prior information, the angular spectra and emissivities of both foreground components and the corresponding synchrotron and dust maps. The results are summarized in Sect. 5. Note that in 2015, a foreground model at $N_{\text{fg}} = 2$ dimensions for capturing synchrotron and thermal dust emissions was already explored, but no maps were released (although a dust comparison appeared in [Planck Collaboration X 2016](#)) because additional “foreground

dimensions” were clearly needed to accommodate the systematic errors. In 2018, we use the same dimensions as in 2015 (a SMICA fit with maximal dimension for CMB cleaning and a SMICA fit with $N_{\text{fg}} = 2$ for dust and synchrotron maps); however, contrary to 2015, the $N_{\text{fg}} = 2$ fit yields a clean CMB reconstruction, almost as clean as when using the maximal foreground dimension. For that reason, this *Planck* release includes SMICA-derived synchrotron and dust polarized maps.

2.5. GNILC

The above four methods were the standard CMB extraction algorithms in each of the three *Planck* data releases. In this release, we also consider the Generalized Needlet Internal Linear Combination (GNILC; [Remazeilles et al. 2011a](#)) method as a foreground extraction algorithm. GNILC is not designed to extract CMB information from the data⁴. GNILC is a wavelet-based component-separation method that generalizes the NILC method by exploiting not only the *spectral* information (SED) but also the *spatial* information (angular power spectra) from non-astrophysical components (cosmic infrared background, CIB, CMB, and instrumental noise) to extract clean estimates of the correlated emission from Galactic foregrounds, with reduced contamination from CIB, CMB, and noise. This additional spatial discriminator adopted by GNILC enables in particular disentanglement of emission components that suffer from spectral degeneracies, such as modified blackbody emissions like the CIB and Galactic dust. GNILC has been successfully applied to *Planck* 2015 intensity data to disentangle Galactic thermal dust emission and CIB anisotropies over the entire sky ([Planck Collaboration Int. XLVIII 2016](#)). In this paper, CMB and instrumental noise were also filtered out from the *Planck* GNILC dust intensity map by using the same strategy as for CIB removal.

In this work, we apply GNILC to the *Planck* 2018 polarization data in order to extract the Stokes parameters Q and U of Galactic thermal dust polarization, while removing the contamination from CMB polarization and instrumental noise over the entire sky. I , Q , and U dust maps have been produced in a self-consistent way by processing the seven *Planck* polarized channels (30–353 GHz). The reason for discarding the 545 and 857 GHz channels is as follows. The main characteristic of the GNILC method is to estimate the local number of independent foreground degrees of freedom over the sky and over angular scales. The estimated dimension of the foreground subspace depends on the local signal-to-noise ratio in the 9×9 (intensity) or 7×7 (polarization) observation space of the frequency-by-frequency data covariance matrix. In some parts of sky where the data are found by GNILC to be fully compatible with CIB, CMB, and noise at small angular scales, the dimension of the Galactic foreground subspace can go down to zero. The result of this is that the GNILC dust products have a variable resolution over the sky, with the local FWHM fully determined and publicly released ([Planck Collaboration Int. XLVIII 2016](#)). However, because of decorrelation effects, the local dimension of the foreground subspace found by GNILC will be larger in a 9-dimensional space of observations (30–857 GHz) than in a 7-dimensional space of observations (30–353 GHz), so that the effective local resolution of the GNILC dust products will be different over the sky for intensity and polarization. For the purpose of polarization fraction studies in the 2018 release ([Planck Collaboration XII 2020](#)), we prefer

⁴ GNILC should not be confused with the “Constrained ILC” method ([Remazeilles et al. 2011b](#)), which was designed to extract SZ-free CMB temperature anisotropies.

to have the same local resolution over the sky both for intensity and polarization, hence our choice of processing with GNILC the same data set for I , Q , and U , namely the seven *Planck* polarized channels (30–353 GHz).

Omission of the 545 and 857 GHz channels limits the ability of GNILC to clean CIB anisotropies in the *Planck* 2018 dust intensity map compared to the *Planck* 2015 dust intensity map (Planck Collaboration Int. XLVIII 2016), for which the full set of unpolarized channels (30–857 GHz) and the IRAS map were used in the component-separation pipeline. For analyses of dust intensity (e.g., dust optical depth, emissivity, and temperature), we recommend use of the *Planck* 2015 GNILC dust intensity map, which has reduced CIB contamination. Conversely, for analysis of dust polarization (e.g., polarization fraction) we recommend use of GNILC 2018 I , Q , and U maps.

3. Data selection, preprocessing, splits, and simulations

3.1. Frequency maps

The low-level data processing and mapmaking algorithms adopted for the current release are described in detail in Planck Collaboration II (2020) and Planck Collaboration III (2020). For the LFI maps at 30, 44, and 70 GHz, there are only minor changes compared to the previous release, the most important of which is a better calibration procedure that explicitly accounts for polarized foregrounds in the calibration sources. For HFI, more significant changes have been implemented, all designed to suppress instrumental systematics at various scales. These include better ADC and transfer-function corrections, and explicit bandpass corrections employing a detailed foreground model.

A particularly important problem for both LFI and HFI with respect to polarization reconstruction is bandpass mismatch between multiple detectors within a single frequency channel. The issue may be summarized as follows. In order to solve for both temperature and linear polarization in each pixel on the sky, a total of three parameters per pixel, it is necessary to include information from at least three polarization-sensitive detectors in any given mapmaking operation. The polarization signal is estimated by taking pairwise differences between the signals observed by these detectors, while accounting for the relative orientation of their polarization detector angles at any given time. However, there are other effects in addition to true sky polarization signals that may induce effective signal differences between detectors. The largest of these is different effective bandpasses. Since each detector has a slightly different frequency response function, each detector observes a slightly different foreground signal. Unless explicitly accounted for during mapmaking, these differences create a spurious polarization signal in the maps.

In the LFI mapmaking procedure, this effect is accounted for in two different ways, as described in Planck Collaboration II (2020). First, for gain estimation, an iterative scheme is established, in which a proper foreground model is derived jointly with the sky maps using Commander. Each iteration of this procedure consists of three individual steps. First, a gain model is established for each radiometer, accounting for the orbital and Solar dipoles as well as astrophysical foregrounds as estimated by Commander. Second, frequency maps are derived based on this gain model using MADAM (Keihänen et al. 2005; Planck Collaboration VI 2016), a well-established destriper. Third, these frequency maps are used by Commander to derive a new foreground model. A total of four such iterations are used to derive

the final LFI maps; however, even after these iterations there may be non-negligible large-scale residuals present in the 70 GHz sky map, as described by Planck Collaboration II (2020). To account for this, a gain correction template, based on differences between consecutive iterations, is subtracted from the final LFI 70 GHz map, with an amplitude derived from a low-resolution likelihood fit (Planck Collaboration V 2020). These procedures account for biases in the time-variable gain solutions; however, they do not remove direct temperature-to-polarization leakage from bandpass mismatch. That effect, which is stationary on the sky, is corrected through use of static templates, as described in detail in Planck Collaboration II (2016). The same procedure is applied to the LFI sky maps in the current release with an updated foreground model (Planck Collaboration II 2020).

For HFI a different but related approach is adopted. The 2015 Commander temperature model is used to explicitly adjust the effective bandpass response of all bolometers within a frequency channel, by subtracting a small fraction of each foreground signal (thermal dust, free-free, and CO emission, but not synchrotron or spinning dust emission) from the individual bolometer timestreams. These “foreground-equalized” timestreams are then combined into a single frequency map by standard destripping. Since only a spin-0 temperature signal is subtracted in this procedure, the resulting polarization maps are unbiased with respect to foreground leakage, to the extent that the foreground model is accurate. However, the resulting temperature maps will be very slightly biased, in the sense that the predicted bandpass response of a given map does not perfectly match the observed signal, and this causes complications for any method that explicitly employs such information. In the current paper, this applies to Commander and GNILC. The three remaining methods (NILC, SEVEM and SMICA) do not explicitly use such information.

An additional complication arises from the updated 2018 HFI mapmaking procedure, due to the fact that the single-bolometer maps produced by the latest processing are not reliable for component-separation purposes (Planck Collaboration III 2020). Since the CO emission lines are very narrow, their measured amplitudes are very sensitive to small variations in bandpass shape between individual detectors. In 2015, this sensitivity was exploited to extract line-emission maps at each of the affected frequencies. However, since single-bolometer maps are not available in 2018, this is no longer possible. The new processing represents a conscious choice of optimizing the polarization extraction at non-negligible expense in terms of our ability to perform high-fidelity astrophysical foreground reconstruction with temperature maps. For individual foreground components in temperature, we therefore recommend continued usage of the *Planck* 2015 data products.

To summarize the overall data selection, all diffuse component separation codes except GNILC employ all nine *Planck* frequency maps between 30 and 857 GHz in temperature, and all seven frequency maps between 30 and 353 GHz in polarization, for the 2018 analysis. GNILC uses only the seven lowest frequencies in temperature in order to match the polarization analysis. For the LFI polarization maps, we apply a set of template corrections that account for bandpass mismatch and gain corrections, as described in Planck Collaboration II (2020), while no additional corrections are applied to the HFI maps. All maps are defined by the HEALPix⁵ pixelization (Górski et al. 2005).

⁵ <http://healpix.jpl.nasa.gov>

3.2. Instrument characterization

In addition to the raw frequency maps, each method requires various degrees of knowledge about the *Planck* instrument itself. The most important characterization is the beam response of the individual frequency channels. These have been updated to reflect the latest changes in the data processing pipelines, and are described in [Planck Collaboration II \(2020\)](#) and [Planck Collaboration III \(2020\)](#). We note that in the 2015 data release, CMB polarization maps for two of the methods (*Commander* and *SEVEM*) were given at 10' FWHM, compared to 5' FWHM for the temperature maps; however, in this release all CMB maps in both polarization and temperature are provided at the maximum angular resolution of 5' FWHM.

Each CMB map must also be associated with a statistical characterization of the instrumental noise. For this purpose, we compute and analyse null maps derived from subsets of the full data set, as done in earlier releases. In the previous release, we focused on half-mission splits, yearly splits, and half-ring splits ([Planck Collaboration IX 2016](#)). In the current release, we drop the yearly split, since this behaves similarly to the half-mission split, and we replace the half-ring split with a so-called “odd-even” split, in which scanning rings from HFI are grouped according to odd or even pointing IDs. The odd-even split nullifies long-time-correlated signals, similarly to the half-ring split, but suffers less from inter-ring correlations. For LFI, we still adopt the same half-ring split as in 2015, but nevertheless refer to this split as “odd-even,” recognizing the different signal-to-noise ratios of the LFI and HFI maps. We consider this to be our best instrumental noise tracer among the splits, whereas the half-mission split represents the best instrumental systematics tracer. Simulations including either pure CMB signal or the sum of instrumental noise and residual systematics are individually propagated through each analysis pipeline, and these simulations form the basis of all subsequent goodness-of-fit tests.

As described in [Planck Collaboration III \(2020\)](#), the HFI polarization frequency maps are associated with a significant uncertainty regarding polarization efficiencies, corresponding in effect to an uncertainty in the overall calibration of the Stokes Q and U maps. Ideally, such polarization efficiencies would be perfectly accounted for during mapmaking. However, as reported by [Planck Collaboration V \(2020\)](#), a cosmological analysis of power spectra of the individual frequency maps suggests that small but notable residual calibration uncertainties may remain in a few channels. The reported best-fit correction values are $+0.7 \pm 1.0\%$ (100 GHz), $-1.7 \pm 1.0\%$ (143 GHz), and $+1.9 \pm 1.0\%$ (217 GHz). For 353 GHz, the foreground contribution is too large to allow a robust CMB-based measurement. These corrections are only marginally statistically significant, therefore we do not apply them by default in this paper. Instead, we compute results with and without the corrections, and report the difference between the two solutions as a known systematic error. For the CMB, we find that the differences due to polarization efficiency uncertainties are small, while for polarized foregrounds, we find that the inclusion of polarization efficiencies changes the spectral index of thermal dust by $\Delta\beta_d = -0.03$. See Sect. 5 for details.

3.3. Treatment of unobserved pixels

As described in [Planck Collaboration III \(2020\)](#), the HFI split maps contain a non-negligible number of unobserved pixels at the full $N_{\text{side}} = 2048$ HEALPix resolution. These are pixels that were either never seen by any bolometer at a given frequency, or for which the polarization angle coverage is too poor to support

a reliable decomposition into the three Stokes parameters. For most methods considered in this paper⁶, such unobserved pixels represent a notable algorithmic problem, and must be treated before analysis. For these methods, we simply replace all unobserved pixels in a given frequency map by the same pixels in a corresponding map downgraded to a HEALPix resolution of $N_{\text{side}} = 64$, corresponding to a pixel size of 55'. Of course, this procedure introduces correlations between neighbouring unobserved pixels, and we therefore mask all high-resolution pixels after the analysis; separate masks for each data split are provided to account for this effect. The details of how the unobserved pixel mask has been generated are described in Sect. 4.2. Finally, to account for possible leakage from unobserved to observed pixels during inter-analysis smoothing operations, we apply the same procedure to the reference simulations described below.

3.4. Comparison between 2015 and 2018 frequency maps

It is useful to compare the new 2018 frequency maps to the previous 2015 frequency maps. Structures seen in these difference maps should be expected to propagate into the corresponding CMB differences at some level. Starting with the temperature case, the left columns of Figs. 2 and 3 show the differences between each 2018 frequency map and the 2018 *Commander* CMB solution⁷. Overall, the behaviour is consistent with what has been found in earlier releases, with: an absolute foreground minimum around 70 GHz; LFI monopoles of $10\text{--}20\mu\text{K}$; increasing HFI monopoles with frequency, corresponding to the expected offset due to the cosmic infrared background (CIB), which is manually introduced into the HFI frequency maps ([Planck Collaboration III 2020](#)); and overall morphologies consistent with some combination of synchrotron, free-free, CO, and thermal and spinning dust emission.

More interesting are the second and third columns in each figure, which show the raw and the fractional differences between the 2018 and 2015 frequency maps, respectively. In the latter we have removed the best-fit offset and dipole outside a Galactic mask, defining the fractional difference, f , as

$$f = \frac{m^{2018} - m^{2015} - \Delta M - \Delta D}{m^{2015} - m^{\text{CMB}}}, \quad (7)$$

where m^{2018} is the new *Planck* 2018 frequency map, m^{2015} is the *Planck* 2015 map, ΔM and ΔD are the monopole and dipole differences between these maps, and m^{CMB} is the *Commander* 2015 CMB temperature map.

Starting with the LFI 30 GHz difference maps, two effects stand out. At high latitudes, we see broad stripes following the *Planck* scanning pattern. These are due to an improved time-varying gain calibration procedure in the 2018 analysis that takes into account astrophysical foregrounds as computed by *Commander*, in an iterative gain-estimation→mapmaking→component-separation procedure. This new iterative scheme is one of the main new features of the LFI 2018 processing pipeline ([Planck Collaboration II 2020](#)). A second effect is seen in the Galactic plane, where the 2018 amplitude is lower by about 0.2%

⁶ *Commander* behaves differently from the other codes with respect to unobserved pixels. It applies per-pixel inverse noise weighting per frequency channel, and unobserved pixels in a given channel are simply given zero weight in the parametric fits.

⁷ We remove a common estimate of the CMB signal in order to highlight the foreground and residual monopole and dipole contents of each map. Visually identical results would be obtained by adopting any of the other solutions as a reference instead of *Commander*.




Fig. 2. Comparison of 2018 and 2015 LFI temperature maps. *From left to right columns:* (1) difference between the 2018 intensity maps and the 2018 Commander CMB map; (2) difference between the 2018 and 2015 frequency maps; and (3) fractional difference between the 2018 and 2015 frequency maps. Note the different temperature scales. In the third column, ΔM and ΔD denote the relative monopole and dipole differences between the 2018 and 2015 sky maps. Rows indicate results for each of the three LFI frequency channels. All maps are smoothed to a common resolution of 1° FWHM.

compared to 2015. This is due to re-estimation of the overall absolute calibration, due to a new estimate of the Solar CMB dipole (Planck Collaboration I 2020).

Similar considerations hold for the 44 GHz channel, although with a significantly lower striping level. In fact, in this case the striping is sufficiently low to reveal a small residual dipole of about $1 \mu\text{K}$ in the raw difference map, directly showing the effect of the new Solar dipole estimate. Even smaller differences are seen in the 70 GHz channel, but in this case the iterative foreground estimation process was not used, because the foreground level of this channel near the foreground minimum is too low to allow robust foreground estimation (Planck Collaboration II 2020).

The HFI frequencies (Fig. 3) show many qualitatively similar structures, in addition to a few unique HFI-type features. First, in the 100 GHz channel we see a fairly large dipole of $2\text{--}3 \mu\text{K}$. In the new HFI processing, thermal dust emission is explicitly included in the dipole estimation model, resulting in improved consistency in the dipole estimates among the various frequency channels. As a result of this process, the best-fit 2018 dipole estimate changed by $2.4 \mu\text{K}$ relative to 2015, and this is visually apparent in the 100 GHz raw difference map. In addition, we see significant striping in the fractional difference map, with an amplitude of more than 3% of the foreground level at high latitudes. As is the case for LFI, these stripes are due to improved time-variable gain estimation, which in turn is responsible for the overall improvement in the large-scale polarization reconstruction. Of course, for this channel the absolute foreground levels are low at high Galactic latitudes, and a 3% relative difference corresponds only to $1\text{--}2 \mu\text{K}$ in absolute value. For

temperature this is small, while for polarization it is highly relevant, as we discuss below.

Qualitatively speaking, similar considerations hold for the 143 and 217 GHz channels as well. However, in these cases we see an additional effect, namely a significantly blue Galactic plane in the fractional difference map, indicating relative absolute differences of about 1% in the high signal-to-noise regime. At first sight, this may appear puzzling, since the absolute CMB calibration between the 2018 and 2015 has changed by less than 0.1% (Planck Collaboration V 2020). The explanation is the new HFI treatment of bandpass differences among individual bolometers. As described in Sect. 3, each frequency map is now generated as the sum over all bolometer timestreams within that frequency channel, each of which has been *bandpass equalized* prior to co-addition. This equalization is implemented by fitting Commander foreground templates of thermal dust, CO, and free-free emission jointly with other instrumental parameters, with the goal of minimizing inter-bolometer bandpass differences that otherwise generate spurious polarization contamination.

For component-separation purposes, this implies that the overall bandpass profile of each HFI frequency channel has changed. Furthermore, this process also leads to a complicated bandpass definition overall, in which the bandpass in principle is component dependent. While thermal dust, free-free, and CO emission are associated with bandpasses given as straight averages of the individual bolometer bandpasses (due to their inclusion in the bandpass equalization procedure), synchrotron, spinning dust, and thermal Sunyaev–Zeldovich signals are associated with inverse noise-variance-weighted bandpasses as




Fig. 3. Comparison of 2018 and 2015 HFI temperature maps, similar to Fig. 2 for LFI. *From left to right columns:* (1) difference between the 2018 intensity maps and the 2018 Commander CMB map; (2) difference between the 2018 and 2015 frequency maps; and (3) fractional difference between the 2018 and 2015 frequency maps. Note the different temperature scales. Rows indicate results for each of the six HFI frequency channels. All maps are smoothed to a common resolution of 1° FWHM. Note that the 217 and 353 GHz difference maps have been scaled by factors of $1/2$ and $1/20$, respectively, to conform numerically to the same range as the 100 and 143 GHz maps.

in earlier releases. In practice, though, we adopt the straight averaged bandpasses for all HFI channels in the current release, since the affected non-equalized components are sub-dominant

at HFI frequencies, and implementing multi-bandpass integration would require significant algorithm re-structuring. However, this is also one of the reasons why we do not release new individual




Fig. 4. Comparison of 2015 and 2018 polarization frequency maps. *From left to right columns:* (1) 2018 Stokes Q maps; (2) 2018 Stokes U maps; (3) Stokes Q difference map between 2018 and 2015; and (4) Stokes U difference map between 2018 and 2015. Note the different temperature scales. Rows indicated results for each of the seven polarized *Planck* frequency channels. All maps are smoothed to a common resolution of 1° FWHM.

synchrotron and spinning dust products in temperature in the current release.

Turning to the 353 GHz frequency channel, two additional effects are seen. First, at high latitudes one can see a weak imprint of zodiacal light emission (Planck Collaboration XIV 2014) in the fractional difference map, taking the form of a blue band along the Ecliptic plane with an amplitude of 1%. Second, we also see two deep blue bands on either side of the Galactic plane with amplitudes of 2%; these are due to changes in the 353 GHz transfer function. From such difference maps alone, it is of course impossible to conclude whether the additional residuals are due to defects in the 2015 or 2018 maps. On the other hand, such structures tend to stand out quite prominently in maps of foreground spectral indices, which in essence

measure small differentials between frequencies. Thus, through subsequent Commander-type astrophysical analyses, we find that these two 353 GHz effects are indeed present in the 2018 maps, and not in the corresponding 2015 maps. These residual effects, along with the lack of single-bolometer maps, are thus part of the cost of producing as clean polarization maps as possible, which is the primary goal of the current data release.

At 545 and 857 GHz, most of the effects are similar to those described above, with one additional effect for the 857 GHz channel, where residual sidelobe contamination dominates the high-latitude residuals, with amplitudes of 2–3% of the full foreground signal. In this case, the 2018 processing represents an absolute improvement over the 2015 processing, in the sense that the full 2018 frequency map has lower sidelobe contamination

than the corresponding 2015 frequency map. At the same time, it is worth noting that single-bolometer maps are available in the 2015 release, and the 857-2 bolometer map has significantly lower sidelobe contamination than any of the other three (Planck Collaboration X 2016). Thus, if a given scientific analysis does not require the signal-to-noise ratio of the full 857 GHz channel, the *Planck* 2015 857-2 bolometer channel may be an even better choice than the full 857 GHz 2018 frequency map. However, in the current paper, which is dedicated to the 2018 release itself, we adopt the 2018 full-frequency map in all analyses.

Figure 4 shows the corresponding plots for polarization. Here we do not subtract any CMB component (since it is small), and we also do not show fractional difference maps (since polarized foreground amplitudes can go both positive and negative). The two leftmost columns show the raw 2018 frequency maps in Stokes Q and U , and the two rightmost columns show the straight differences between the 2018 and 2015 frequency maps.

As expected, the various features seen in the polarization difference maps trace those observed in the corresponding temperature differences. For 30 and 44 GHz, the main features at high latitudes are due to bandpass mismatch and time-variable gain corrections, achieved by iterating between gain estimation, mapmaking and component separation. For 70 GHz, only very small differences are seen, since the gain estimation procedure is unchanged from 2015; however, it is important to note that a separate residual gain template has been produced for this channel, and this is applied in the scientific processing (see Planck Collaboration II 2020).

For the HFI channels, we see similar effects of improved effective gain estimation at high latitudes, as well as improved bandpass corrections at low latitudes, in particular for 100, 217, and 353 GHz, which are strongly affected by CO emission. Most strikingly, the low-latitude structures seen in the 100 GHz map are typical examples of temperature-to-polarization leakage, where the CO morphology is modulated by the scanning orientation of the *Planck* satellite. At 353 GHz, we additionally see the residual effect of transfer-function convolution near the Galactic plane in Stokes U . Thus, caution is warranted when studying polarized thermal dust emission near the Galactic plane with this frequency map.

To summarize, we observe typically (at most) 2–3% differences between the 2015 and 2018 frequency maps at high latitudes, as measured in units of foreground signal. In most cases, these differences are directly due to improvements in the updated processing (Planck Collaboration II 2020; Planck Collaboration III 2020), although with a few notable exceptions, in particular for the 353 GHz channel. It is important to note, however, that the design philosophy of the 2018 release has been to optimize the quality of the polarization products, which in some cases comes at the expense of temperature analysis. In particular, the non-availability of single-bolometer maps represents a limiting factor for astrophysical component separation in temperature. For this reason, we expect both 2015 and 2018 temperature products to be in common use in the future, depending on the needs of a particular application, whereas for polarization we strongly recommend usage of the 2018 products.

3.5. Simulations

The instrumental noise characteristics of the *Planck* observations are complex, and a simple white-noise approximation is inadequate for high-precision analyses of these data. The only realistic approach to handling both instrumental noise and residual sys-

tematics is through end-to-end simulations. As part of the *Planck* 2018 data release, we therefore provide a set of 300 independent noise-plus-systematics simulations for each frequency band and for each of the data splits described above, as well as 999 CMB-only simulations that include the effects of satellite scanning and asymmetric beams; see Planck Collaboration II (2020) and Planck Collaboration III (2020) for full details. These simulations are available through the *Planck* Legacy Archive⁸.

These simulations are propagated through each of the pipelines; we adopt the same frequency weights (mixing matrices, spectral indices etc.) as for the real data. The two main advantages of fixing the weights are, first, that the noise properties actually correspond to the real final maps; and, second, that the system becomes linear, and CMB and noise may be propagated independently through each pipeline. In the following, we will employ CMB-only, noise-only, and CMB-plus-noise combinations for various applications.

3.6. Standardization of simulations and data

Each of the four pipelines processes both the data and simulations somewhat differently with respect to harmonic space truncation ($\ell_{\max}$) and high- ℓ regularization. In order to facilitate meaningful direct comparisons between the various maps, we convolve all four data sets to a common effective resolution prior to analysis, as described below. We emphasize, however, that the released data products are provided at their native resolution, in order to allow external users to exploit the full resolution of each data set, if so desired.

For temperature, the most aggressive smoothing applied by any of the four pipelines is defined by NILC, for which the effective high- ℓ apodization kernel reads

$$B(\ell) = \begin{cases} 1, & \ell \leq \ell_{\text{peak}}, \\ \cos^2 \left[(\pi/2)(\ell - \ell_{\text{peak}})/(\ell_{\max} - \ell_{\text{peak}}) \right], & \ell > \ell_{\text{peak}}, \end{cases} \quad (8)$$

where $\ell_{\text{peak}} = 3400$ and $\ell_{\max} = 3999$. We therefore apply this kernel to each of the three other pipelines, on top of their intrinsic 5' FWHM smoothing kernels. For SMICA we additionally apply the HEALPix pixel window for $N_{\text{side}} = 2048$, which is not by default applied for this code.

For polarization, the most aggressive high- ℓ truncation is applied by SEVEM, which enforces a hard harmonic space truncation at $\ell_{\max} = 3000$. This same truncation is applied to each of the three other codes in polarization as a post-processing step.

4. CMB maps

The CMB maps and associated products obtained by the various pipelines as applied to the *Planck* 2018 data are presented in this section; astrophysical foreground results are presented in the next section. For a detailed analysis of the higher-order statistical properties of these maps, see Planck Collaboration VII (2020).

4.1. Full-mission maps and comparison with 2015 release

Figure 5 shows the final full-mission *Planck* 2018 CMB component-separated maps derived by each of the four pipelines⁹, both in intensity (left column) and polarization (middle and right

⁸ <http://pla.esac.esa.int>

⁹ The four cleaned frequency maps (from 70 to 217 GHz) provided by SEVEM are also shown in Fig. C.1.




Fig. 5. Component-separated CMB maps at 80' resolution. Columns show Stokes I , Q , and U , respectively, while rows show results derived with different component-separation methods. The Galactic plane region in the SMICA maps results from a pre-processing step (masking and diffusive inpainting of a narrow Galactic region in all frequency channels), while no masks are applied to the other maps. In this plot, monopoles and dipoles have been subtracted with parameters fitted outside a $|b| < 30^\circ$ mask.

columns). Only SMICA has been inpainted within a Galactic mask (see Appendix D). All maps are smoothed to a common resolution of 80' FWHM for visualization purposes.

At first sight, the consistency among the various pipeline maps appears to be reasonable outside the central Galactic plane, and, as expected, more so in temperature than in polarization.

In the polarization maps, however, we can identify several notable artefacts already at this stage, which prospective future users of these maps need to be aware of. The visually most striking features are of course residual foreground contamination in the Galactic plane. In particular, the alternating sign along the plane is a classic signature of temperature-to-polarization leakage, and the *Planck* data set is particularly sensitive to residual CO emission in this respect. These features are extremely difficult to suppress to the level of the CMB fluctuations during processing, and must in practice be removed by standard Galactic masking.

The second most striking feature in the polarization maps is a blue stripe in the upper right quadrant of the Stokes U map. This stripe corresponds to a few bad scanning rings that ide-

ally should have been removed by flagging during mapmaking. Unfortunately, this issue was not caught at a sufficiently early stage of the processing, and remains in the final maps. We therefore mask this stripe in the same way that we mask Galactic residuals.

Third, and somewhat less obvious, we observe broad large-scale structures in both Stokes Q and U that are aligned with the *Planck* scanning strategy. These structures are effectively due to gain-modelling uncertainties coupled to monopole and dipole leakage, and corresponding features are present in the associated simulations. In principle, therefore, these need not be removed prior to subsequent analyses, as long as the appropriate simulations are used to quantify all relevant uncertainties. In practice, however, we note that these modes are associated with significant additional systematic uncertainties, and we therefore caution against over-interpretation of the very largest scales in these maps. In particular, we warn against employing these maps for auto-correlation type analysis, unless the statistic of choice is explicitly shown to be robust against this type of systematic effect, based on end-to-end simulations.




Fig. 6. Differences between 2015 and 2018 CMB *I* maps at 80' resolution. From top to bottom panels, results are shown for Commander, NILC, SEVEM, and SMICA. Monopoles and dipoles have been subtracted with parameters fitted outside a $|b| < 30^\circ$ mask.

Figure 6 shows maps of temperature differences between each of the 2018 pipeline maps and the corresponding 2015

pipeline maps. Corresponding maps of polarization differences are not shown, since the high level of large-scale systematics in the 2015 maps renders a direct difference-map comparison non-informative. In Fig. 6, we recognize many of the features seen in the raw input frequency difference maps shown in Figs. 2 and 3, and discussed in Sect. 3.

Starting with Commander, the most striking difference is a dark blue Galactic plane residual with a clear CO-like morphology. This reflects the fact that it is more difficult for the parametric Commander pipeline to estimate CO emission from co-added frequency maps (as in the 2018 processing) than with individual bolometer maps (as in the 2015 processing). For this reason, the Commander map adopts a larger Galactic mask in the new release than in the previous one, specifically targeting CO emission; see Appendix A for further details. The second most notable feature in the Commander difference map is a $\lesssim 2 \mu\text{K}$ blue signal at intermediate latitude with a thermal dust imprint, and this is due to the changes in bandpass modelling in the high-frequency channels.

Only small differences are observed for NILC, for which very few pipeline modifications have been introduced since 2015. NILC already used full-frequency maps in the previous release. The most significant change is a large-scale quadrupolar structure at high latitudes, which directly reflects the effective gain changes at 100, 143, and 217 GHz seen in Fig. 3. Likewise, SEVEM also used full-frequency maps in 2015, and only minor pipeline modifications have been introduced, and consequently, only minor differences are observed in temperature from 2015 to 2018.

For SMICA, three qualitatively different types of differences are seen. First, the weak large-scale background pattern is similar to that observed in NILC, and simply reflects the slight changes in input data discussed above. In addition, we see significant changes in the compact sources that can be explained by the change of masking strategy described in Sect. 2. Third and finally, the strong near-Galactic-plane differences that include free-free sources (e.g., the Gum Nebula and Rho Ophiucus) are explained by the miscalibration of the 44 GHz channel in the 2015 released map (as recalled in Sect. 2). The impact of this issue is assessed in Appendix D.

Figure 7 shows all pairwise difference maps between each of the pipeline CMB maps. The structures seen in these plots correspond closely to those already discussed above. Finally, Fig. 8 shows the standard deviation evaluated from the four cleaned CMB maps, smoothed to 80' FWHM angular scales; the polarization standard deviation is here defined as $\sqrt{\text{var}Q + \text{var}U}$.

4.2. Confidence masks

From the above discussion, it is clear that significant residuals are present in the CMB maps, in particular close to the Galactic plane. Therefore, appropriate masking is required for scientific exploration of the *Planck* 2018 maps in both temperature and polarization, as in earlier releases. For this purpose, we adopt a conservative strategy similar to that of 2015, and we construct a common confidence mask for all maps, even if the various maps may have different levels of residuals.

In previous releases, a common mask was generated simply as the product of the individual confidence masks derived for each pipeline. However, the pipeline masks were established using qualitatively different criteria in each case, and a direct comparison was therefore non-trivial. In the current analysis, we adopt a more direct route, starting with the inter-pipeline standard deviation maps shown in Fig. 8. The single most striking




Fig. 7. Pairwise differences between maps from the four CMB component separation pipelines, smoothed to $80'$ resolution. Columns show Stokes I , Q , and U , respectively, while rows show results for different pipeline combinations. The lines show the regions masked in component separation. Monopoles and dipoles have been subtracted with parameters fitted outside a $|b| < 30^\circ$ mask.




Fig. 8. Standard deviation of the CMB maps between the four component-separation methods, at 80' resolution. Temperature is shown in the *top panel* and polarization in the *bottom panel*. The polarization standard deviation is defined as $\sqrt{\text{var}(Q) + \text{var}(U)}$.

feature in these maps is the Galactic plane. At high latitudes, one can additionally see point sources and a few rings in the *Planck* scanning strategy that were observed fewer times than average, resulting in coherent stripes of excess variance.

Specifically, for temperature we first threshold at $3\mu\text{K}$ the standard deviation map evaluated at 80' FWHM smoothing scale, and adopt this as our primary mask. The specific smoothing scale of 80' represents a compromise between suppressing noise while still retaining small features, while the threshold of $3\mu\text{K}$ is defined by the region at high Galactic latitude in the top panel of Fig. 8. Second, we smooth this binary mask, consisting of 0 and 1s, with a 10° FWHM Gaussian beam, and remove any pixels with a value lower than 0.5; this is to remove isolated small “islands” within the main Galactic plane. Third, we threshold at $10\mu\text{K}$ a corresponding standard deviation map evaluated at 10' FWHM smoothing scale in order to remove compact objects.

The resulting mask ensures that only pixels for which the four pipelines agree in their CMB solutions to better than $3\mu\text{K}$ in standard deviation on large scales ($10\mu\text{K}$ on small scales) are allowed in the final analysis. However, quantitative agreement among codes is only a necessary criterion; it is not sufficient. We therefore augment this mask by the absolute individual confidence masks of *Commander* and *SEVEM* (see Appendices A and C for details), by the point-source masks used for inpainting by *SEVEM* and *SMICA*, and by the processing mask employed by *SMICA*. The first two of these employ χ^2 and difference maps to define their acceptable regions, and thereby correspond to standard absolute goodness-of-fit statistics, while the latter two correspond to basic processing masks. The *SEVEM* inpainted

point-source mask is constructed from point sources detected in the 143 and 217 GHz channels, and is described in detail in Appendix C (see also Fig. C.5). We find no evidence for significant artefacts in the NILC and *SMICA* maps outside the *Commander* and *SEVEM* masks defined above, and we therefore do not apply any special measure for these maps.

We adopt a similar procedure for polarization, but with a few notable additions. First, the inter-pipeline standard deviation map evaluated at 80' FWHM is thresholded at $1\mu\text{K}$. The resulting mask is smoothed to 5° FWHM, and thresholded at a value of 0.9, effectively expanding the original mask by a few degrees in all directions. This mask is then multiplied with a corresponding mask derived by thresholding at $0.6\mu\text{K}$ the original standard deviation map at 80' FWHM, to remove sharper features. We then exclude all pixels flagged by the *Commander* and *SEVEM* confidence masks. Next, we remove the region contaminated by cosmic rays discussed in Sect. 4.1, as defined in Ecliptic coordinates following *Planck*'s scanning path. Third, as an additional guard against temperature-to-polarization leakage from CO emission, we exclude any pixels for which the CO emission at 100 GHz (see Sect. 5) is brighter than $20\mu\text{K}$, evaluated at a smoothing scale of 5° FWHM. Isolated “islands” in the main Galactic plane are then removed with the same procedure as for temperature. Finally, we also add the point-source masks used for inpainting by *SEVEM* and *SMICA*. For polarization, the *SEVEM* inpainted point-source mask is constructed from point sources detected in the 100, 143, and 217 GHz channels and is shown in Fig. C.5.

The resulting common masks are shown in the top row of Fig. 9 for temperature (left panel) and polarization (right panel). The final accepted sky fractions are $f_T = 0.780$ and $f_P = 0.782$. These sky fractions are similar to those reported in 2015, namely $f_T^{2015} = 0.77$ and $f_P^{2015} = 0.78$.

As discussed in Sect. 3.3, the half-mission and odd-even split maps contain a number of unobserved or poorly conditioned pixels. For split-map analysis, we therefore recommend additional unobserved pixel masks. These are produced by thresholding the 3×3 Stokes parameter condition number hit-count maps produced during mapmaking (*Planck Collaboration II 2020*; *Planck Collaboration III 2020*). The resulting unobserved pixel mask is further extended in a three-step iterative process in which the neighbours of each unobserved pixel have been masked. The bottom row in Fig. 9 shows the products of the temperature and polarization unobserved pixel masks for both the half-mission (left panel) and odd-even (right panel) splits.

As a final mask-related issue, we note that the *Planck* 2018 product delivery includes Wiener-filtered versions of each pipeline map, in which all high-foreground regions are replaced with a Gaussian constrained realization. For temperature, these regions are defined simply by thresholding the maximum difference between any of the four cleaned CMB maps at $100\mu\text{K}$, and additionally removing all pixels excluded by the *SMICA* processing mask. This mask is shown in Fig. 10, and excludes 2% of the sky. For polarization inpainting we conservatively adopt the common confidence mask defined above. In either case, we note that the inpainted CMB maps are primarily intended for publication and presentation purposes, rather than scientific analysis. For full scientific analysis of the high-foreground-contaminated regions, we recommend corresponding processing of end-to-end simulations. These are, however, not provided due to large data volume and processing costs, although they may be generated by applying a Wiener filter code such as *Commander* (*Seljebotn et al. 2019*) to the cleaned CMB simulations that are provided.




Fig. 9. Masks recommended for analysis of the cleaned CMB maps. *Top panel:* common confidence masks for temperature (*left*) and polarization (*right*). These masks should always be applied to any scientific analysis of the maps presented in this paper. *Bottom panel:* unobserved pixel masks for half-mission (*left*) and odd-even (*right*) data splits. These panels show the products of the individual unobserved pixel masks for temperature and polarization, whereas separate (but very similar) temperature and polarization masks are applied during analysis.




Fig. 10. Mask used for inpainting the cleaned CMB temperature maps.

4.3. Effective transfer functions

As noted in Sect. 2, all *Planck* 2018 CMB maps have a common nominal target resolution of 5' FWHM, as output by each of the respective pipelines. However, this resolution is not exact, as it does not take into account the effect of spatially-varying asymmetric beams on the sky. The nominal 5' beam kernel must therefore be corrected by an effective transfer function for each pipeline prior to any harmonic space analysis of these maps, including cosmological power spectrum and parameter estimation.

We estimate the effective transfer functions from the CMB signal-only simulations discussed in Sect. 3.5 through the following expression,

$$f_\ell = \frac{1}{b_\ell^{5'} p_\ell^{2048}} \sqrt{\left\langle \frac{C_\ell^{\text{out}}}{C_\ell^{\text{in}}} \right\rangle}, \quad (9)$$

where C_ℓ^{out} and C_ℓ^{in} denote the simulated output and input power spectra of each CMB signal realization. The former includes both instrumental beam and pixel window convolution, while the latter includes neither. The quantity $b_\ell^{5'}$ denotes the beam

transfer function of a 5' FWHM Gaussian beam, p_ℓ^{2048} is the $N_{\text{side}} = 2048$ pixel window (Górski et al. 2005), and brackets indicate an average over 50 simulations. Equation (9) is evaluated independently for temperature and both *E*- and *B*-mode polarization. Finally, each transfer function is smoothed with a third-order Savitzky–Golay filter with a window size of $\Delta\ell = 51$ to reduce residual uncertainty from the finite number of Monte Carlo simulations. The examples shown in this paper correspond to full-sky transfer functions; these functions should in principle be re-evaluated for each sky fraction used in a given analysis¹⁰.

The resulting transfer functions are shown in Fig. 11. Starting with the temperature case, we first note that the range spanned by the four curves is well within $\pm 0.5\%$, and, therefore, these effects are quite minor for all the considered multipoles. Overall, qualitatively similar behaviour is observed for the four codes, with *Commander* showing a slightly larger deviation. In particular, for *Commander*, we see that the effective residual transfer function is very close to unity up to $\ell \approx 700$, after which it starts to fall off, eventually reaching an amplitude of about 0.3% at $\ell \approx 2000$, before it begins to rise sharply. The small excess of $\lesssim 0.1\%$ around $\ell = 500$ is associated with the effective cut-off of the LFI-dominated low-frequency signal component

¹⁰ For the particular case of *SEVEM*, the evaluation of the transfer function is in principle affected by the pixels inpainted in the cleaned frequency maps. Since those pixels are all excluded in the common confidence mask, we have evaluated this function from full-sky CMB simulations without applying this inpainting. Therefore, this effective transfer function should be a good approximation for the regions passed by the common mask. However, if a transfer function is needed for a region of the sky that contains inpainted pixels, it is recommended to re-evaluate this function for that particular sky coverage, taking into account the inpainting. Although the effect is very small, it can be noticeable, especially for aggressive masks that remove only a small fraction of the Galaxy, since in the Galactic regions a relatively large number of sources are inpainted.

employed by *Commander*. These general trends are due to small mismatches between the full asymmetric beams, as implemented through pixel-space FEFBeCoP (Mitra et al. 2011) convolutions, and the azimuthally symmetric effective beam transfer functions, as implemented with QuickBeam (Hivon et al. 2017). For instance, a fall-off of 0.3% at $\ell \approx 2000$ corresponds to a mismatch of about $0.05'$ FWHM in the two models. Further, we note that *Commander* is the only code that applies per-pixel inverse variance weighting of the raw frequency maps, and as such it has a different response to small angular scales than the other codes.

Turning our attention to the E -mode transfer functions, the most striking new feature is a pattern of systematic wiggles. These are seen both in transfer functions derived from each frequency alone (shown as dotted lines) and in the component-separated maps. These wiggles are due to temperature-to-polarization leakage through the asymmetric beam shapes, and the positions of the peaks coincide with the peaks in the CMB temperature power spectrum. Note that most of the total weight below $\ell \approx 300$ is determined by the 143 GHz channel, while above $\ell \approx 300$ it is dominated by the 217 GHz channel. The 100 GHz channel does not dominate at any angular scale, due to its lower angular resolution and higher noise as compared to the 143 GHz channel.

Similar considerations apply to the B -mode transfer functions, although in this case the wiggles are largely dominated by an increasing trend caused by a wide range of both temperature-to-polarization and polarization-to-polarization leakage effects. Overall, however, the net sum of all these effects is smaller than 10% of the underlying (lensing-induced) B -mode signal up to $\ell \lesssim 1600$. We also see that the component-separated map is strongly dominated by the 217 GHz channel for $\ell \gtrsim 500$.

4.4. Noise characterization and consistency with simulations

We now characterize the statistical properties of the component-separated CMB map, and we start with a description of instrumental noise and residual systematic effects. We adopt three main measures for this purpose, each designed to highlight different aspects of the effective noise properties; these are designed for different applications.

Our first noise measure is defined in terms of the so-called odd-even half-difference (OEHD) maps, in which the full time-ordered data volume is divided according to odd and even ring numbers. This is a fine-grained time split, and as such, the OEHD map tends to cancel most systematic effects. This noise measure is thus our cleanest probe of pure instrumental (white and correlated) noise. OEHD maps are plotted in Appendix G.1 for each pipeline and for each of the three Stokes parameters. Overall, we see that these difference maps exhibit very few visually-apparent systematic effects at high latitudes, and the only significant residuals occur in the Galactic centre, where the overall signal amplitude is very large.

Our second noise measure is defined in terms of the half-mission half-difference (HMHD) maps, in which the time-ordered data are split according to long time periods, defined by years (Planck Collaboration II 2020; Planck Collaboration III 2020). This measure is thus a coarse-grained time split, and more sensitive to systematic effects that vary on long time scales, such as gain variations or sidelobe contamination. This is our preferred estimate for the combined impact of instrumental noise and systematic effects. HMHD maps are shown in Appendix G.2 for each pipeline and for each of the three Stokes parameters. In these maps, we clearly see the imprint of the Galactic plane, which is largely caused by calibration uncertainties,




Fig. 11. Effective transfer functions f_ℓ for each of the four pipeline CMB maps, after deconvolving a $5'$ FWHM Gaussian beam and $N_{\text{side}} = 2048$ HEALPix pixel window. *From top to bottom panels:* results for T , E , and B . In the two bottom panels, the dotted lines show the effective residual transfer functions for the three CMB-dominated HFI frequencies between 100 and 217 GHz, after deconvolving the azimuthally-symmetric QuickBeam-based transfer function and HEALPix pixel window in each case.

as well as more pronounced scan-aligned structures at high latitudes.

The third noise measure comprises the full-blown end-to-end simulations, in which all known systematics have been modelled to the best of our ability (see Planck Collaboration II 2020; Planck Collaboration III 2020 for full details). These simulations are generated as raw time-ordered data, and processed through each step of the analysis pipeline, including map making and component separation. Unfortunately, this process is computationally very expensive, and only a limited set of 300 realizations has been produced for the current release. However, for each realization a full set of results are produced, including full mission maps, half-mission and odd-even splits. Combined, these form the basis of most goodness-of-fit statistics presented in the following sections.

4.4.1. Power spectrum analysis

In Fig. 12 we compare the power spectra of the cleaned CMB maps with the simulations. All spectra are evaluated outside the common mask described in Sect. 4.2 using PolSpice (Chon et al. 2004). Furthermore, all spectra have been normalized




Fig. 12. Power spectrum consistency between cleaned CMB maps and end-to-end simulations. Each panel shows the fractional difference between the angular power spectrum computed from the observed data and the mean of the simulations. Rows show different polarization spectra (TT , EE , and BB), while columns show different data splits (full, HMHD, and OEHD).

relative to the mean of the simulated ensemble, and plotted in terms of the fractional deviation,

$$\eta_\ell \equiv \frac{D_\ell^{\text{data}} - \langle D_\ell^{\text{sim}} \rangle}{\langle D_\ell^{\text{sim}} \rangle}. \quad (10)$$

This function thus measures the fractional difference of the observed power spectrum from the mean of the simulations, plotted as a percentage in Fig. 12. This function is evaluated both for temperature and polarization (TT , EE , and BB), as well as for full-mission, HMHD, and OEHD data splits. For clarity, each function has been binned with $\Delta\ell = 25$ after computing the above single- ℓ quantity.

For full-mission temperature data, we find that the CMB-plus-noise simulations agree well with the data in terms of angular power up to $\ell \lesssim 750$. At higher multipoles, we see a slow increase in power up to $\ell \approx 2000$, corresponding to a positive contribution from point sources not included in the simulations. The level of point-source residuals is highest in Commander and lowest in SMICA. At high multipoles, $\ell \gtrsim 2000$, the spectra turn over. As described in Planck Collaboration III (2020), the power in the HFI simulations for the 100–217 GHz channels underesti-

mates the true noise in the real data by a few percent (with variations depending both on angular scale and frequency), and this translates into a negative bias at high multipoles in the cleaned CMB maps presented in this paper.

Similar features are seen even more clearly in the polarization EE and BB full-mission spectra, for which the signal-to-noise ratio is lower. In these cases, the simulations agree well with the data up to $\ell \lesssim 200$, after which a negative bias of a few percent is observed in the range $200 \lesssim \ell \lesssim 500$. Then, in the range $500 \lesssim \ell \lesssim 1500$ the agreement is good, before we see the same negative high- ℓ bias as in the temperature case. The same trends are even more prominent in the HMHD and OEHD spectra, which by construction are entirely noise-dominated.

In Fig. 13 we focus on the first two multipoles, and compare the observed power to the full simulated distributions in terms of cumulative distribution functions. Overall, we observe acceptable statistical agreement between the data and simulations at these largest scales, with only a few points showing extreme values of 0 or 1; however, even in these cases the observed values lie just at the edge of the simulated histogram. No large outliers are observed. Nevertheless, it is important to note that the effective noise varies greatly between the various analysis pipelines,

and it is therefore essential to compare any given data set with its corresponding simulations.

To summarize, the end-to-end simulations presented and employed in this paper exhibit power biases of several percent with respect to the true observations on intermediate and small scales, while reasonable agreement is observed on large angular scales. These biases originate from corresponding discrepancies at the level of individual frequency bands, as reported in [Planck Collaboration II \(2020\)](#) and [Planck Collaboration III \(2020\)](#). When employing these simulations for scientific analysis, it is important to verify that the statistic of choice is not sensitive to such percentage-level differences. This will usually be the case for linear or cross-correlation type analyses, but not necessarily for quadratic or auto-correlation type analyses.

4.4.2. Pixel-space variance analysis

A complementary consistency measure is given by the total variance as measured in pixel space at different pixel resolutions (see, e.g., [Monteserín et al. 2008](#); [Cruz et al. 2011](#); [Planck Collaboration XVI 2016](#)). This method normalizes the map with respect to the total variance of the signal plus noise, where the noise variance is estimated through the simulations described above, and the variance of the signal is determined as the value that gives a normalized map variance equal to unity. For the HMHD and OEHD maps, the method simplifies, since the CMB signal is cancelled through the half-difference calculation.

Recognizing the fact that *Planck* polarization maps are generally noise dominated, we apply the methods described in [Planck Collaboration VIII \(2020\)](#) and Molinari et al. (in prep.) for polarization. These methods include both auto- and cross-estimates for the variance, which is the result of the subtraction between the variance of the $\langle Q^2 + U^2 \rangle$ signal-plus-noise map and the variance of the $\langle Q_N^2 + U_N^2 \rangle$ noise estimated from the MC simulations. For both temperature and polarization, we employ the respective union masks described above. When dealing with HMHD and OEHD maps, we consider the union mask combined with the corresponding unobserved pixel mask.

In Fig. 14 we plot the percentage of simulations with a lower variance than the real data, as a function of pixel resolution. The results from this analysis are in good agreement with those found in the power spectrum analysis. In temperature we find a generally good consistency between real data and half difference noise simulations, with few exceptions. We find that only a few simulations have a lower variance for the HMHD Commander map at very large scales and for the OEHD SEVEM map at intermediate resolutions. At the maximum resolution of $N_{\text{side}} = 2048$, there is a lack of compatibility with MC simulations for both HMHD and OEHD maps, showing that at high resolution noise in temperature data is poorly described by the simulations. However, given the very high signal-to-noise ratio of the *Planck* temperature data at all resolutions, a small noise mismatch is irrelevant compared to the CMB cosmic variance.

For the signal-plus-noise data, we observe satisfactory consistency in temperature at high pixel resolutions. At lower resolutions we observe low probabilities, with p -values of about 1.0%, which are associated with the well known lack of power on large angular scales. These results are compatible with results reported in the previous release described in [Planck Collaboration XVI \(2016\)](#). We have also investigated the higher order moments, skewness and kurtosis in temperature as shown in [Planck Collaboration VII \(2020\)](#), and find good consistency with Monte Carlo simulations at all resolutions.

In polarization at high resolutions, results are not as robust, due to the noise mismatch. We observe an incompatibility for all the component-separated HMHD and OEHD maps between the MC distribution and real data at intermediate and high resolutions. This suggests that noise in the data (including systematic effects) is not fully characterized by the simulations. At lower resolutions ($N_{\text{side}} = 256$ for HMHD and 128 for OEHD), however, data are compatible with simulations, showing that the noise properties are better represented. In Fig. 15 we show the amplitude of the noise mismatch with respect to the amplitude of the expected CMB variance as a function of pixel resolution. These results give an estimation of the bias due to the noise mismatch in the extraction of the variance from signal plus noise data. The bias is very important at the highest resolution ($N_{\text{side}} = 2048$), with values of about 40–50% for all the methods. At intermediate resolutions it is of the order of few percent. At large scales the bias is not significant, since half-difference data are compatible with the MC dispersion.

In spite of the presence of a noise mismatch, we find that cross- and auto-analyses of the signal-plus-noise maps are in agreement with MC simulations. At intermediate and large scales, auto- and cross-analyses are in good agreement with each other, showing the robustness of the analysis, although with some differences among the component-separation methods due to the presence of residual foregrounds, or systematic effects, or a different impact of the noise mismatch. At high resolution the differences between auto and cross results are mainly due to the noise mismatch, whose impact is more important for the auto analysis. On the other side, the cross-analyses may be biased by a poor description of the correlated noise that we cannot investigate with the above analyses.

In [Planck Collaboration VIII \(2020\)](#) we consider more detailed analyses of this kind, including also the analysis of the SEVEM frequency maps, in a way that minimizes the impact of the correlated noise.

4.4.3. Assessing the impact of simulation noise bias

In order to understand whether these percent-level noise discrepancies in polarization are relevant for a given analysis, we strongly recommend considering the following questions while assessing the results.

(1) Which angular scales are relevant for the statistic of choice? If the statistic is sensitive only to large angular scales ($\ell \lesssim 50$), then the simulations are likely to be adequate. If not, see next question.

(2) Is the statistic of choice sensitive to signal-plus-noise or noise alone? If the former, then the simulations are likely to be adequate for $\ell \lesssim 1500$ for temperature, and $\ell \lesssim 250$ for polarization; if the latter, then see next question.

(3) Is the statistic of choice sensitive to $\lesssim 5\%$ errors in the noise model? To quantify this, we recommend applying the statistic of choice to simulations for which the noise contribution is artificially re-scaled either up or down by 5% (see Fig. 12), while the signal contribution is unchanged. If the statistic of choice is unable to distinguish between the scaled and the unscaled ensembles, then the statistic is likely robust against the uncertainties in the current simulations. If not, caution is warranted. Typically, linear, cubic, or cross-spectrum type statistics are only marginally sensitive to this type of error in the noise model, whereas quadratic and auto-spectrum type statistics are typically highly sensitive.

Clearly, no general prescription can be given for all analyses, and we therefore stress that caution is warranted when using the




Fig. 13. Power spectrum consistency between cleaned CMB maps and end-to-end simulations for $\ell = 2$ and 3 . Solid lines show cumulative distributions computed from 300 simulations, and dashed lines show the value derived from the *Planck* data. From top to bottom, the three sections show TT , EE , and BB , and within each section the top and bottom rows show distributions for $\ell = 2$ and 3 . Columns from left to right show full, HMHD, and OEHD splits. The fraction of simulations with a lower power amplitude is given in the legends of each panel for each code.




Fig. 14. Consistency between data and simulations as quantified in terms of pixel-space variance for both temperature (*left panel*) and polarization (*right panel*), and for both full mission maps (*top row*), half difference maps (*middle two rows*), and for half-mission and odd-even cross variances (*bottom two rows*). Coloured lines show results for the four different component-separation pipelines, Commander (red), SMICA (cyan), SEVEM (green), and NILC (orange), as a function of pixel resolution, N_{side} .

end-to-end simulations. That being said, they do provide the most complete description of the uncertainties in the data set currently available, and with an appropriate level of care, they should form the basis of most goodness-of-fit tests with the current data set. For several worked examples of applications of these simulations, see [Planck Collaboration VII \(2020\)](#).

4.5. Foreground template fits

Next, we consider residual foreground contamination as measured by correlation between known foreground templates and the cleaned CMB maps. Specifically, for a given cleaned temperature CMB map $\mathbf{d}$, a foreground template $\mathbf{t}$, and the common confidence mask $\mathbf{m}$ (consisting of 0's and 1's), we compute the correlation coefficient

$$r = \frac{1}{N_{\text{pix}} - 1} \sum_{i \in \mathbf{m}} \frac{\mathbf{d}_i - \langle \mathbf{d} \rangle}{\sigma_{\mathbf{d}}} \frac{\mathbf{t}_i - \langle \mathbf{t} \rangle}{\sigma_{\mathbf{t}}}, \quad (11)$$

where the sum runs over the N_{pix} pixels not excluded by the mask, $\langle \mathbf{d} \rangle = 1/N_{\text{pix}} \sum_i \mathbf{d}_i$, $\sigma_{\mathbf{d}} = \left[1/N_{\text{pix}-1} \sum_i (\mathbf{d}_i - \langle \mathbf{d} \rangle)^2 \right]^{1/2}$, and similarly for $\mathbf{t}$. All maps are smoothed to a common resolution of 80' FWHM, and pixelized at $N_{\text{side}} = 128$.

We consider four foreground templates in intensity, namely, the 408 MHz [Haslam et al. \(1982\)](#) map as processed by




Fig. 15. Amplitude of the noise mismatch in terms of expected signal amplitude, var_{th} , in percentage as a function of the pixel resolution, N_{side} for HMHD polarization data (*top plot*) and OEHD polarization data (*bottom plot*). Coloured lines show results for the four different component-separation pipelines, Commander (red), SMICA (cyan), SEVEM (green), and NILC (orange). Error bars show the amplitude of the MC dispersion at $\pm 1 \sigma$, showing that where the noise is well characterized the bias is embedded in the uncertainty in the variance extraction and hence it is not significant.

[Remazeilles et al. \(2015\)](#) for synchrotron emission, the *Planck* 2018 857 GHz map for thermal dust emission, the [Dame et al. \(2001\)](#) map for CO line emission, and the [Finkbeiner \(2003\)](#) H_{α} map for free-free emission. For polarization, we consider the difference between the WMAP 23 GHz and 33 GHz maps as a synchrotron tracer. Uncertainties are evaluated from 300 end-to-end simulations. Corresponding results were reported in [Planck Collaboration IX \(2016\)](#) for the *Planck* 2015 CMB sky maps.

The results from these calculations are summarized in Table 1. In nearly all cases, we see that the correlation coefficients are lower for the 2018 maps than the corresponding 2015 maps, and most are within the 1σ confidence limits. The only notable issue is a marginally significant polarization correlation with the WMAP synchrotron tracer, ranging in statistical significance between 2.8σ for Commander to 3.5σ for SEVEM. The absolute level of the correlation is low, however, ranging between 3 and 4%.

Table 1. Correlation coefficients between known diffuse foreground templates and each of the four component-separated CMB maps.

		Correlation coefficient			
	Data set	Commander	NILC	SEVEM	SMICA
<i>Intensity</i>					
Haslam.....	2018	-0.037 ± 0.134	-0.023 ± 0.135	-0.055 ± 0.077	-0.027 ± 0.077
	2015	-0.062 ± 0.115	-0.051 ± 0.116	-0.065 ± 0.115	-0.023 ± 0.069
H α	2018	0.000 ± 0.032	0.006 ± 0.032	0.003 ± 0.028	0.004 ± 0.028
	2015	0.010 ± 0.071	0.011 ± 0.071	0.019 ± 0.071	0.003 ± 0.057
CO.....	2018	-0.002 ± 0.019	0.000 ± 0.019	0.001 ± 0.020	0.001 ± 0.020
	2015	-0.004 ± 0.027	0.003 ± 0.027	0.003 ± 0.027	-0.007 ± 0.022
857 GHz.....	2018	-0.033 ± 0.097	-0.019 ± 0.097	-0.018 ± 0.098	-0.019 ± 0.098
	2015	-0.043 ± 0.084	-0.032 ± 0.084	-0.037 ± 0.084	-0.029 ± 0.083
<i>Polarization</i>					
WMAP K–Ka.....	2018	-0.031 ± 0.011	-0.037 ± 0.011	-0.039 ± 0.011	-0.033 ± 0.011
	2015	-0.057 ± 0.026	-0.116 ± 0.024	-0.026 ± 0.025	-0.027 ± 0.026

Notes. The intensity templates are: (1) a 408 MHz map for synchrotron emission from [Haslam et al. \(1982\)](#); (2) an H α template for free-free emission from [Finkbeiner \(2003\)](#); (3) a tracer of CO emission in the Galactic plane from [Dame et al. \(2001\)](#); and (4) the *Planck* 857 GHz map for thermal dust emission. For polarization, we only include the difference between the WMAP K (23 GHz) and Ka (33 GHz) frequency maps as a synchrotron tracer. All maps have been smoothed to a common resolution of 80' FWHM prior to the fitting process, and the correlation coefficients are evaluated outside the confidence masks described in Sect. 4.2.




Fig. 16. Comparison of half-mission temperature power spectra. The *top panel* shows the half-sum (HMHS; solid lines) and half-difference (HMHD; dashed lines) power spectra, while the *bottom panel* shows the difference between the half-sum and the best-fit *Planck* 2018 Λ CDM and half-difference spectra. The latter residual spectrum is binned with $\Delta\ell = 25$.

4.6. Power spectrum comparison

Next, we characterize the cleaned CMB maps in terms of angular power spectra. As above, we employ the PolSpice estimator for these calculations, and all spectra are evaluated outside the common mask defined in Sect. 4.2.

Figure 16 shows a comparison of power spectra evaluated from the four cleaned CMB temperature half-mission maps. In the top panel, the solid lines show spectra computed from the half-mission half-sum (HMHS) maps, and thereby contain both

signal and noise, while dashed lines show spectra computed from the HMHD, and thereby should contain only instrumental noise and systematic uncertainties. The black solid line shows the best-fit *Planck* 2018 Λ CDM model derived from the combination of the low- ℓ TT , low- ℓ EE , high- ℓ $TT + TE + EE$, and lensing likelihoods ([Planck Collaboration XIII 2016](#)). The bottom panel shows the residuals after subtracting both the best-fit Λ CDM model (as a signal tracer) and the half-difference spectrum (as a noise tracer) from each of the half-sum spectra.

Overall, we observe good agreement among the four pipelines in terms of half-sum spectra up to $\ell \lesssim 1500$. The main notable feature is a small power deficit of about $10 \mu\text{K}^2$ in NILC between $\ell = 100$ and 300 , corresponding to a relative deficit of 0.2%. At these multipoles, the *Planck* data are strongly CMB dominated, and algorithmic variations make little difference in terms of overall power. However, at higher multipoles the noise and compact source contributions become relevant, and in that regime the various approaches show slightly different behaviour, with Commander having the largest unresolved source imprint and NILC the smallest.

At low multipoles we also see differences among the codes in terms of noise. The lowest noise is achieved by Commander, which also exhibits nearly white noise with a scaling of $O(\ell^2)$. The highest low- ℓ noise – almost an order of magnitude higher than Commander – is seen for NILC for $\ell \lesssim 300$. This is not unexpected given the nature of the NILC algorithm. On large angular scales, the NILC frequency weights primarily adjust themselves to suppress foregrounds, while on small scales, they converge to inverse-noise-variance weighting. In this respect, Commander is different from the other three codes in that it explicitly uses estimates of the noise standard deviation to perform inverse-variance noise weighting per pixel. Finally, for SMICA we note that the noise decreases around $\ell \approx 100$, which corresponds to the multipole at which the three LFI frequencies are excluded at high latitudes (see Appendix D).

In Fig. 17 we present a similar comparison for the EE and BB HMHD polarization power spectra. As in Fig. 16, the solid lines includes contributions from both signal and noise. Here we see, at least at the level of visual inspection, that all four codes perform similarly in terms of polarization power




Fig. 17. Comparison of half-mission polarization power spectra. The *top panel* shows the half-sum (solid lines) and half-difference (dashed lines) power spectra, while the *bottom panel* shows the difference between the half-sum and the best-fit *Planck* 2018 Λ CDM and half-difference spectra. The latter residual spectrum is binned with $\Delta\ell = 25$.

spectrum reconstruction, both for HMHS and HMHD spectra. The only marginal outlier is NILC, which exhibits slightly higher BB HMHS and HMHD spectra at multipoles lower than $\ell \lesssim 100$. However, this excess disappears in the difference between the HMHS and HMHD spectrum (bottom panels of Fig. 17), suggesting that it is due to a somewhat higher noise level in the NILC map compared to the others, and not a signal bias.

Finally, in Fig. 18 we show an expansion of the low multipole part of the polarization power spectra without applying any multipole binning. The black solid line in the left panel shows the best-fit *Planck* 2018 Λ CDM model for which the posterior mean optical depth of re-ionization is $\tau = 0.054 \pm 0.019$ (Planck Collaboration VI 2020). The left and right panels show the EE and BB spectra, respectively. The grey curves indicate corresponding spectra computed from 300 end-to-end Commander simulations.

Starting with the BB spectrum, we note that there is an overall significant excess compared to zero. This excess, however, is reproduced in the simulations, as seen by the non-zero mean of simulations, and therefore reflects the presence of understood residual correlations in the data; see Planck Collaboration III (2020) for further discussion. Consequently, we re-emphasize the importance of comparing these data with full end-to-end simulations when subjecting them to cosmological analysis, in order to adequately capture this type of residual noise correlation. A similar noise excess is seen in the EE spectrum for $\ell \gtrsim 8$, with an amplitude similar to the BB spectrum.

The EE spectrum does not show a clear detection of the reionization peak. As mentioned, the solid black curve shown in Fig. 18 indicates a spectrum for which $\tau = 0.054$, and this amplitude is too low to be visually observed in an ℓ -by- ℓ spectrum, in particular in the presence of the noise excess mentioned above. In order to detect this peak, a full likelihood analysis is essential, as presented in Planck Collaboration V (2020) and Planck Collaboration VI (2020).

While the *Planck* 2018 best-fit value for the optical depth of reionization is $\tau = 0.054$, the value used to generate the simulations (which had to be adopted well before the final results were available for computational expense reasons) was $\tau = 0.060$. The effect of this difference can be seen in Fig. 18 as an excess of power in the simulations relative to the best-fit model at low multipoles in EE . While the difference is within 1σ , it is worth having this issue in mind when studying low- ℓ polarization effects with these maps and simulations; see Sect. 3 of Planck Collaboration VII (2020) for a quantitative analysis of this issue.

A full cosmological likelihood and parameter analysis of the *Planck* 2018 data is presented in Planck Collaboration V (2020), based on cross-spectrum techniques. In this paper, we perform a simple consistency test between the full likelihood analysis and the cleaned CMB maps presented in this paper, by fitting a CMB spectrum (parametrised by an amplitude A_{CMB} and tilt n relative to the best-fit *Planck* 2018 Λ CDM model) and an ℓ^2 point-source contribution to the difference between the HMHS and HMHD spectra. Explicitly, we adopt the signal model

$$D_\ell = A_{\text{CMB}} (\ell/\ell_0)^n f_\ell^2 D_\ell^{\Lambda\text{CDM}} + A_{\text{ps}} \ell(\ell+1)/(\ell_{\text{ps}}(\ell_{\text{ps}}+1)), \quad (12)$$

where f_ℓ is the transfer function shown in Fig. 11, $\ell_0 = 600$ is a pivot multipole for the CMB fit, and $\ell_{\text{ps}} = 500$ is a pivot multipole for the point source contribution. The quantity $D_\ell^{\Lambda\text{CDM}}$ is the best-fit *Planck* 2018 Λ CDM power spectrum. Each analysis includes multipoles between $\ell = 2$ and 1500 (for which the simulations agree well with the data; see Fig. 12), and all spectra are binned with $\Delta\ell = 20$. Uncertainties within each bin are defined as the standard deviation of the observed spectrum within the bin. The number of degrees of freedom for the χ^2 is $n_{\text{d.o.f.}} = 75$. The fit is performed with a simple Metropolis MCMC sampler.

The results from these calculations are summarized in Table 2. Overall, we find good agreement between the cleaned CMB maps and the likelihood analysis, with most Λ CDM amplitudes consistent with unity within 2σ and all tilt parameters consistent with zero within 1.5σ . All χ^2 s are also reasonable, ranging




Fig. 18. Comparison of low- ℓ half-mission polarization EE (left) and BB (right) power spectra. Each spectrum is computed as the difference between the respective half-sum and the half-difference spectrum. Grey bands show 1σ confidence regions derived from 300 end-to-end simulations as processed by Commander. Formally speaking, these can therefore only be directly compared with the red curve. Note that the value for the optical depth of reionization adopted for these simulations is $\tau = 0.060$ (see Planck Collaboration II 2020; Planck Collaboration III 2020 for details), which is larger than the best-fit *Planck* 2018 value of $\tau = 0.054$.

between 57.8 and 74.0 for 75 degrees of freedom, corresponding to probabilities-to-exceed (PTEs) ranging between 0.07 and 0.49. Finally, as already noted, Commander exhibits the largest point-source contribution, with an amplitude of $A_{\text{ps}} = 2.7 \pm 0.5 \mu\text{K}^2$ at $\ell = 500$ in temperature, while NILC and SMICA show the smallest contribution, with amplitudes of $A_{\text{ps}} = 1.9 \mu\text{K}^2$ at $\ell = 500$.

4.7. The real-space N -point correlation functions

A complementary measure of correlations and non-Gaussianity are given by real-space 2- and 3-point correlation functions. These functions are defined as the average product of N observed fields, measured in a fixed relative distance on the sky. In the case of the CMB, the fields correspond to temperature anisotropy ΔT and two Stokes parameters Q and U describing the linear polarization of the radiation in a given direction (see Planck Collaboration VII 2020 for more detail).

Because of computational limitations, we restrict our analysis to the pseudo-collapsed and equilateral configurations of the 3-point functions and low resolution CMB maps with a resolution parameter $N_{\text{side}} = 64$ and smoothed with a $160'$ FWHM Gaussian beam (see Planck Collaboration VII 2020 for a description of the procedure for downgrading and smoothing the maps). For both temperature and polarization, we employ the common masks described above, downgraded in the same way as CMB maps. Because we analyse half-difference maps, the common mask is combined with the corresponding unobserved pixel mask. The resulting 2- and 3-point correlation functions for the Commander HMHD and OEHD maps are presented in Fig. 19 (figures for the remaining component separation maps can be found in the Appendix H). We use a simple χ^2 statistic to quantify the agreement between the observed data and the full-focal-plane noise simulations (FFP10; Planck Collaboration XII 2016). Table 3 lists the significance level in terms of the fraction of simulations with a larger χ^2 value than the observed map. Corresponding analysis for full CMB maps is provided in Planck Collaboration VII (2020).

We can observe quite significant scatter in the results for the half-difference maps estimated using the different component-separation methods. This is not surprising, as different component-separation methods respond to noise and systematic effects in a different way. No statistically significant deviations between data and simulations are found at these angular scales

Table 2. Parameter fits to HM power spectra.

Pipeline	A_{CMB}	n	$A_{\text{ptsrc}} [\mu\text{K}^2]$	χ^2
<i>TT</i>				
Commander	0.997 ± 0.002	-0.003 ± 0.004	2.7 ± 0.5	61.6
NILC	0.994 ± 0.002	-0.006 ± 0.004	1.9 ± 0.6	59.3
SEVEM	0.996 ± 0.002	-0.001 ± 0.004	2.1 ± 0.6	60.3
SMICA	0.996 ± 0.002	-0.001 ± 0.002	1.9 ± 0.5	58.5
<i>EE</i>				
Commander	0.985 ± 0.007	-0.002 ± 0.009	0.35 ± 0.06	74.0
NILC	0.984 ± 0.007	-0.007 ± 0.010	0.26 ± 0.08	70.3
SEVEM	0.983 ± 0.008	-0.001 ± 0.010	0.27 ± 0.08	62.7
SMICA	0.983 ± 0.007	-0.002 ± 0.008	0.30 ± 0.08	66.7

Notes. In each case, the observed spectrum is taken as the difference between the HMHS and HMHD spectra, and the model fitted reads $D_\ell = A_{\text{CMB}} (\ell/\ell_0)^n f_\ell^2 D_\ell^{\Lambda\text{CDM}} + A_{\text{ps}} \ell(\ell+1)/(\ell_{\text{ps}}(\ell_{\text{ps}}+1))$, where f_ℓ is the transfer function shown in Fig. 11, $\ell_0 = 600$, and $\ell_{\text{ps}} = 500$. $D_\ell^{\Lambda\text{CDM}}$ is the best-fit *Planck* 2018 ΛCDM power spectrum. Each analysis includes multipoles between $\ell = 2$ and 1500. Spectra are binned with $\Delta\ell = 20$. The uncertainties within each bin are defined as the standard deviation of the observed spectrum within the bin. The number of degrees of freedom for the χ^2 is $n_{\text{dof}} = 75$.

except a few cases for the 3-point functions with deviation significance around 99%. Note, however, that the confidence regions derived from the noise simulations do vary between methods, indicating that each method results in different effective statistical properties. To avoid biases, it is therefore essential to analyse each map together with the simulations constructed specifically for that map.

4.8. Gravitational lensing

As an example of science that may be extracted from the cleaned CMB maps presented in this paper, we consider reconstruction of the gravitational lensing potential. For a complete analysis of this topic, we refer the interested reader to Planck Collaboration VIII (2020), from which the following results are reproduced.

Gravitational lensing of CMB photons by large-scale structures induces slight distortions in the statistics of the CMB. In particular, lensing deflections result in a characteristic acoustic peak smoothing signature in the angular power spectrum, and




Fig. 19. 2-point (*upper panels*), pseudo-collapsed (*middle panels*) and equilateral (*lower panels*) 3-point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* Commander HMHD (*left panels*) and OEHD (*right panels*) temperature and polarization map. The red solid lines correspond to the half-difference maps (HMHD or OEHD). The green triple-dot-dashed lines indicate the mean determined from 300 FFP10 noise simulations. The shaded dark and light grey regions indicate the corresponding 68% and 95% confidence regions, respectively.

they induce a non-zero four-point CMB correlation function. We use the methodology described in [Planck Collaboration VIII \(2020\)](#) to reconstruct the lensing power spectrum. With the sensitivity and sky coverage of *Planck*, this approach constrains the lensing deflection power spectrum to a few percent, with most of the signal coming from temperature observations at high multipoles $\ell \sim 1500$. These measurements therefore result in a strin-

gent consistency test between the various component-separation methods at small angular scales.

For masking, we employ the union of the intensity and polarization mask recommended in Sect. 4.2, combined with a Galactic mask allowing $f_{\text{sky}} = 0.70$, as well as a mask removing resolved SZ clusters with $S/N > 5$, as given by the *Planck* 2015 SZ catalogue ([Planck Collaboration XXVI 2016](#); for full details,

Table 3. Probabilities in percentages of obtaining values for the χ^2 statistic of the N -point functions for FFP10 simulations at least as large as the those obtained from the observed Commander, NILC, SEVEM, and SMICA temperature and polarization maps with resolution parameter $N_{\text{side}} = 64$.

Function	HMHD split				OEHD split			
	Comm.	NILC	SEVEM	SMICA	Comm.	NILC	SEVEM	SMICA
<i>2-point functions</i>								
TT	50.0	81.3	97.0	87.7	6.7	81.7	75.3	60.3
$Q_r Q_r$	20.7	42.3	40.3	49.3	52.7	30.3	90.0	54.7
$U_r U_r$	12.0	21.7	48.3	51.0	40.7	51.7	75.3	33.7
TQ_r	36.0	2.0	29.0	36.0	98.0	90.0	93.3	83.7
TU_r	1.3	21.3	45.0	33.7	49.0	35.7	96.7	44.7
$Q_r U_r$	47.7	51.3	76.3	68.7	76.0	78.7	77.3	96.3
<i>Pseudo-collapsed 3-point functions</i>								
TTT	32.3	29.3	22.7	50.3	50.0	28.0	20.0	90.7
$Q_r Q_r Q_r$	33.7	18.3	25.0	30.7	96.0	27.0	72.3	75.7
$U_r U_r U_r$	8.0	9.0	42.7	11.0	84.0	99.7	45.3	94.3
TTQ_r	63.7	20.7	77.0	67.7	17.0	60.3	99.0	95.3
TTU_r	82.0	59.0	42.3	74.7	94.3	18.3	63.0	26.3
$TQ_r Q_r$	23.0	2.3	33.7	7.0	99.7	94.7	50.7	78.7
$TU_r U_r$	70.0	64.3	55.7	48.3	98.3	95.3	75.3	90.7
$TQ_r U_r$	82.7	69.7	96.0	39.7	30.3	98.0	29.7	94.0
$Q_r Q_r U_r$	50.0	68.3	73.7	80.0	66.3	72.7	77.7	98.3
$Q_r U_r U_r$	64.7	90.0	89.0	69.3	45.3	27.3	52.7	79.3
<i>Equilateral 3-point functions</i>								
TTT	74.0	62.0	70.7	91.7	82.7	84.3	85.0	53.3
$Q_r Q_r Q_r$	30.3	26.0	95.7	14.3	81.7	94.7	23.3	41.0
$U_r U_r U_r$	90.7	91.0	91.0	74.0	35.3	93.3	4.7	70.0
TTQ_r	34.0	50.3	50.7	38.0	93.0	56.3	49.7	55.3
TTU_r	94.0	29.0	37.0	40.3	82.3	75.3	96.0	86.3
$TQ_r Q_r$	54.3	72.7	51.0	58.3	51.7	73.0	82.3	85.7
$TU_r U_r$	92.3	99.0	96.7	96.0	88.7	70.3	>99.7	69.0
$TQ_r U_r$	96.3	87.3	27.0	89.3	90.7	52.0	70.0	54.0
$Q_r Q_r U_r$	64.0	82.7	97.3	74.0	68.0	55.7	73.3	72.0
$Q_r U_r U_r$	58.7	69.7	68.3	39.0	38.7	26.3	41.3	29.3

Notes. Results are given for both the HMHD and OEHD data splits.

see [Planck Collaboration VIII \(2020\)](#)). We consider quadratic lensing estimates built from temperature only ($\hat{\phi}^{TT}$), as well as the full minimum variance combination ($\hat{\phi}^{MV}$). The minimum-variance estimator is derived from the full set of quadratic estimators TT, TE, TB, EE , and EB , which increases the signal-to-noise ratio with respect to TT by roughly 20%. As discussed in Sect. 4, there is slight power mismatch between data and simulation power on the scales relevant for lensing. To account for this, we add in each case additional power as an isotropic, Gaussian component either to the simulations or to the data.

Figure 20 shows the our minimum-variance lensing spectrum estimates evaluated from lensing multipoles $8 \leq L \leq 2048$. Summary amplitude statistics are listed in Table 4, both on the conservative ($8 \leq L \leq 400$) and high- L ($401 \leq L \leq 2048$) ranges. As we see, the four component-separation methods result in almost identical constraining power. No clear band-power outliers are observed in Fig. 20, and all summary statistics are consistent with each other within uncertainties. However, all four methods show a lensing power that is slightly tilted with respect to the fiducial model, with slightly less power at high multipoles. More detailed analysis and consistency tests are presented in [Planck Collaboration VIII \(2020\)](#).

4.9. Limits on primordial non-Gaussianity

The foreground-cleaned CMB maps may also be used to constrain primordial non-Gaussianity, which is often parameterized




Fig. 20. Lensing reconstruction power spectrum from the four cleaned CMB maps, including lensing multipoles $8 \leq L \leq 2048$ in the minimum variance estimator. For comparison, the black line shows the lensing potential power spectrum adopted for the FFP10 simulation suite.

Table 4. Summary of reconstructed gravitational lensing amplitudes.

Multipole range	Lensing amplitude, $\hat{A}$			
	Commander	NILC	SEVEM	SMICA
MV, $L = 8-400$	0.99 ± 0.03	0.98 ± 0.02	1.00 ± 0.03	0.98 ± 0.02
MV, $L = 401-2048$	0.87 ± 0.10	0.86 ± 0.10	0.80 ± 0.10	0.83 ± 0.10
TT , $L = 8-400$	1.00 ± 0.03	0.99 ± 0.03	0.99 ± 0.03	0.99 ± 0.03
TT , $L = 401-2048$	0.77 ± 0.10	0.78 ± 0.10	0.69 ± 0.11	0.75 ± 0.10

Notes. These amplitudes are defined relative to the Λ CDM spectrum adopted for the FFP10 simulation, which is close, but not identical, to the best-fit *Planck* 2018 Λ CDM spectrum. The first two lines show results derived with the minimum variance estimator that includes both temperature and polarization data, while the last two rows show results derived from temperature data alone.

in terms of the amplitude, f_{NL} , of quadratic corrections to the gravitational potential. This amplitude may be measured through the harmonic-space 3-point correlation function, evaluated for different triangle configurations. A detailed f_{NL} analysis applied to the current cleaned CMB maps is presented in [Planck Collaboration IX \(2020\)](#).

Table 5 summarizes some of the main results presented in that paper, specifically f_{NL} as evaluated from each map by the KSW estimator after correcting for gravitational lensing and the ISW effect. Three sets of results are provided, corresponding to constraints derived from temperature data alone, from polarization data alone, and from temperature and polarization combined. Corresponding results for the *Planck* 2015 data were presented in [Planck Collaboration VIII \(2016\)](#). However, due to the presence of systematic effects, the largest angular scales in polarization were excluded from that analysis. In contrast, the new results presented for the 2018 data set include all angular scales.

As in previous analyses with *Planck* measurements ([Planck Collaboration XXIV 2014](#); [Planck Collaboration XVII 2016](#)), no statistically significant detection of primordial non-Gaussianity is found in the *Planck* 2018 data set, even when including large angular scales in polarization. Statistically speaking, the most significant excursion from zero corresponds to a 2.4σ deviation. With six statistically independent tests (three in each

Table 5. Amplitude of primordial non-Gaussianity, f_{NL} , estimated by the KSW estimator after correcting for gravitational lensing and the ISW effect.

Type	f_{NL}			
	Commander	NILC	SEVEM	SMICA
<i>T</i>				
Local	-2 ± 6	0 ± 6	0 ± 6	-2 ± 6
Equilateral	15 ± 66	-10 ± 66	17 ± 66	14 ± 66
Orthogonal	25 ± 37	0 ± 36	24 ± 37	-15 ± 36
<i>E</i>				
Local	31 ± 29	9 ± 30	38 ± 29	47 ± 28
Equilateral	170 ± 170	39 ± 160	180 ± 170	170 ± 160
Orthogonal	-180 ± 88	-130 ± 88	-180 ± 88	-210 ± 86
<i>T + E</i>				
Local	-2 ± 5	-1 ± 5	-2 ± 5	-1 ± 5
Equilateral	-10 ± 47	-31 ± 46	-9 ± 47	-18 ± 47
Orthogonal	-13 ± 23	-24 ± 23	-15 ± 23	-37 ± 23

Notes. See Table 1 in [Planck Collaboration IX \(2020\)](#) for full details.

of temperature and polarization), this has a probability-to-exceed of about 10% by chance alone.

4.10. Analysis of end-to-end simulations

We finish this CMB-targeted analysis section with a brief discussion of end-to-end simulations, focusing on polarization extraction from the FFP10 set. For a corresponding analysis of temperature simulations, see [Planck Collaboration IX \(2016\)](#).

Unlike the simulations discussed in Sect. 4.4, which only included the CMB and instrumental noise, the simulations considered in this section also includes polarized synchrotron and thermal dust emission. These simulations are processed through each pipeline, allowing each code to estimate spectral parameters (i.e., weights for NILC, SEVEM, and SMICA, and spectral indices for Commander) directly from the simulations.

Figure 21 shows the CMB polarization reconstruction error for each of the four CMB analysis pipelines, as evaluated from the end-to-end FFP10 analysis pipeline, defined by

$$\Delta P = \sqrt{(Q_{\text{out}} - Q_{\text{in}})^2 + (U_{\text{out}} - U_{\text{in}})^2}, \quad (13)$$

where Q_{out} and U_{out} are the estimated Stokes parameters, and Q_{in} and U_{in} are the true Stokes parameters. All maps have been smoothed to 80' FWHM before computing this quantity, to reduce the impact of instrumental noise.

In these plots, one may observe generally similar behaviour between Commander and SEVEM, and between NILC and SMICA. Explicitly, NILC and SMICA result in slightly lower residuals in the Galactic plane, whereas Commander and SEVEM appear slightly less sensitive to stripes at high Galactic latitudes. As evaluated over the common polarization mask, the standard deviations of the four maps (in alphabetical order) are $0.74 \mu\text{K}$, $0.86 \mu\text{K}$, $0.74 \mu\text{K}$, and $0.75 \mu\text{K}$, respectively.

5. Polarized foregrounds

We now turn to the scientific characterization of diffuse microwave foregrounds as derived from the *Planck* 2018 polarization maps; a corresponding discussion of temperature fore-

ground products is given in Appendix F. Three different algorithms are employed in the following, namely Commander ([Eriksen et al. 2004, 2008](#); [Planck Collaboration X 2016](#); [Seljebotn et al. 2019](#)), GNILC ([Remazeilles et al. 2011a](#)), and SMICA ([Cardoso et al. 2008](#)).

5.1. Internal consistency and goodness-of-fit

Before considering astrophysical components, it is instructive to consider the internal consistency between the *Planck* 2018 polarization frequency maps. For this purpose, we employ the Commander model described in Sect. 2.1, fitting a minimal three-component signal model (CMB, synchrotron, and thermal dust emission) to the seven polarized *Planck* frequencies between 30 and 353 GHz. The synchrotron component is modelled by a single power-law with a free spectral index, β_s , in the frequency domain, while the thermal dust component is modelled as a modified blackbody with free spectral index, β_d , and temperature, T_d . In the main analyses, the synchrotron spectral index is fixed spatially to $\beta_s = -3.1$, matching the high-latitude temperature result found from the combination of *Planck* 2015, WMAP, and Haslam data ([Planck Collaboration X 2016](#)); as shown in Sect. 5.3, the *Planck* measurements by themselves have little sensitivity to the synchrotron spectral index. For thermal dust, we fix T_d at the Commander result found from the *Planck* 2018 temperature data in Appendix F; with a highest frequency of 353 GHz, the *Planck* polarization observations are insensitive to this parameter.

Additionally, we impose a spatial smoothness prior on both synchrotron and thermal dust emission to reduce noise-induced degeneracies between the various components. This takes the form of a Gaussian smoothing kernel with 40' FWHM for synchrotron emission and 10' FWHM for thermal dust emission. The widths of these priors are chosen to match the resolution at which the data have a significant signal-to-noise ratio; see Appendix A for further details.

Given this model, the top panels in Fig. 22 show residual maps of the form data minus model ($d_v - s_v$) for each *Planck* frequency map, all smoothed to a common resolution of 40'. The colour scales cover $\pm 20 \mu\text{K}$ for the LFI channels, and $\pm 5 \mu\text{K}$ for the HFI channels. Ideally, each of these maps should be consistent with instrumental noise alone, and for the three LFI channels this appears to be a reasonable approximation. The only clearly visible artefacts in these maps correspond to regions of high foreground amplitudes, which most likely are due to a low level of residual temperature-to-polarization leakage, for instance from bandpass mismatch between individual detectors. In particular, the sharp morphology of the Galactic plane residuals corresponds to the shape of temperature foregrounds, not polarization foregrounds.

In contrast, significant large-scale residuals may be seen at all four HFI frequencies, with patterns typically aligning with the *Planck* scanning strategy. Collectively, these features correspond to effective calibration uncertainties that couple the CMB dipole and foregrounds to the reconstructed CMB polarization signal. Although these residuals are significant, their amplitudes are almost an order of magnitude smaller than in the 2015 data. Moreover, the latest end-to-end simulations describe the residuals to a high level of precision ([Planck Collaboration III 2020](#)).

The bottom panel in Fig. 22 shows the reduced χ^2 per pixel, as defined by

$$\chi^2_{\text{red}}(p) = \frac{1}{\nu_{\text{d.o.f.}}} \sum_{v=1}^{N_{\text{band}}} \left(\frac{d_v - s_v(p)}{\sigma_v(p)} \right)^2. \quad (14)$$




Fig. 21. CMB polarization reconstruction error for each of the four CMB analysis pipelines, as evaluated from the end-to-end FFP10 analysis pipeline. This error is defined as $\sqrt{(Q_{\text{out}} - Q_{\text{in}})^2 + (U_{\text{out}} - U_{\text{in}})^2}$, where Q_{out} and U_{out} are the estimated Stokes parameters, and Q_{in} and U_{in} are the true Stokes parameters. Each difference map has been smoothed to $80'$ FWHM before computing the polarization amplitude, to reduce the impact of instrumental noise.

This map is summed over Stokes Q and U parameters and evaluated at $N_{\text{side}} = 1024$, corresponding to the resolution of the LFI frequency maps. The total number of degrees of freedom is therefore approximately $\nu_{\text{d.o.f.}} = 2 \cdot (3 + 4 \cdot 4) - 2 \cdot 3 = 32$, accounting for three LFI maps at $N_{\text{side}} = 1024$, four HFI maps at $N_{\text{side}} = 2048$, and three fitted component maps, each with an angular resolution comparable to the size of an $N_{\text{side}} = 1024$ pixel. The colour range corresponds to $\pm 3\sigma$ in terms of expected statistical variation for 32 degrees of freedom. Note that $\sigma_{\nu}(p)$ only accounts for white noise. The smoothness of this χ^2 map clearly suggests that the *Planck* 2018 polarization observations are dominated by instrumental white noise on intermediate and small angular scales, not by systematic effects or foreground artefacts.

5.2. Polarization amplitude

Next, we consider the polarization amplitude of synchrotron emission at 30 GHz and thermal dust emission at 353 GHz, naively defined as $P = \sqrt{Q^2 + U^2}$. As discussed by [Plaszczynski et al. \(2014\)](#), this estimator is intrinsically noise-biased; however, since we are only interested in it for comparison and consistency purposes, the noise bias is not critical for this paper. The resulting maps are shown in Figs. 23–27, as estimated by Commander, GNILC, and SMICA. For Commander, the synchrotron map is smoothed to $40'$ FWHM and the thermal dust emission map is smoothed to $5'$ FWHM. For SMICA, the corresponding smoothing scales are $40'$ and $12'$ FWHM. For GNILC the effective angular resolution varies over the sky, depending on the local signal-to-noise ratio. The Commander maps correspond to the amplitudes evaluated at monochromatic reference frequencies, while the GNILC and SMICA maps correspond to bandpass-integrated maps at 30 and 353 GHz, respectively.

Two sets of GNILC products are delivered for the *Planck* 2018 release: (i) the GNILC Stokes I , Q , and U maps of thermal dust emission at uniform $80'$ resolution, with the associated GNILC noise covariance matrix maps (II , IQ , IU , QQ , QU , and UU); and (ii) the GNILC Stokes I , Q , and U maps of thermal dust emission at variable resolution ($80'$ – $5'$) over the sky, with the associated GNILC noise-covariance-matrix maps, along with a beam FWHM map indicating the corresponding variable resolution of the dust over the sky regions. The *Planck* 2018 GNILC dust products are analysed in great detail in [Planck Collaboration XII \(2020\)](#).

Figure 28 shows a scatter plot between the Commander and GNILC thermal dust amplitudes, both evaluated for a common resolution of $80'$ FWHM. Overall, the agreement is very good, and the Pearson's correlation coefficient between the two maps is $r = 0.999$. Similar good agreement is observed between the SMICA and the Commander and GNILC maps, except for very high values of P , for which SMICA applies an inpainting mask during processing to avoid ringing. The main notable difference between the Commander and GNILC maps is an overall relative scaling of around 5%, corresponding to the fact that no colour corrections are applied to the GNILC map, and it therefore corresponds to the dust signal as observed through the *Planck* 353 GHz bandpass. This distinction between the two maps is important to bear in mind when subjecting either one to statistical analysis.

Based on these polarization amplitude maps, one can compute the corresponding polarization fraction, defined as $p = P/I$, where P is the polarization amplitude, and I is the corresponding total intensity. This quantity is useful for modelling and characterizing astrophysical emission processes, and is therefore of great interest to astrophysical theorists. However, it is also highly




Fig. 22. Top panel: Commander polarization residual maps, $d_v - s_v$, for each polarized *Planck* frequency channel. All maps are smoothed to a common resolution of 40' FWHM. Bottom panel: reduced χ^2 map for the high-resolution polarization analysis. The grayscale range corresponds to $\pm 3\sigma$ in terms of expected statistical variation.

sensitive to systematic errors in the intensity component, and in particular to the zero level, which is difficult to constrain for the *Planck* measurements. A careful analysis of the thermal dust polarization fraction derived from the *Planck* 2018 measurements, including zero level uncertainties, is provided in [Planck Collaboration XII \(2020\)](#), and we refer the interested reader to that paper for full details.

5.3. Synchrotron and thermal dust spectral indices

Next, we consider the spectral energy distributions (SEDs) for polarized synchrotron and thermal dust emission. For simplicity, we focus primarily on the effective spectral index for either process, noting that *Planck* has very limited sensitivity to estimate additional spectral parameters in polarization.

Starting with Commander, we note that the main analysis discussed above is performed with informative (delta function or Gaussian) priors on both β_s and β_d . In order to quantify the intrinsic information content and statistical strength of the *Planck* data to constrain these parameters at a more basic level, it is useful also to perform *prior-free* runs. The results from such analyses are summarized in Fig. 29, for synchrotron emission in the top panel and thermal dust emission in the bottom panel. In either case, the Gaussian prior is removed only on the component in question, not both simultaneously. In all cases, however, a broad uniform prior is imposed in order to exclude completely unphysical values. The synchrotron analysis is performed at a smoothing scale of 5° FWHM, while the thermal dust analysis is performed at a smoothing scale of 3° FWHM. This scale was determined by considering a series of scales (1, 2, 3, 5, 10 deg), and identifying the largest scale that did not result in leakage artifacts.

For synchrotron emission, we find a very broad distribution between $\beta_s = -4$ and -1.5 , with both ends being defined by the uniform prior. There is a weak preference for values between $\beta_s = -3.5$ and -3.0 , consistent with the value of $\beta_s = -3.1$ found by combining *Planck*, WMAP, and Haslam temperature data in [Planck Collaboration X \(2016\)](#), but overall, it is clear that the *Planck* polarization data by themselves do not significantly constrain the spectral index of synchrotron emission at scales smaller than 5°. For the main analysis, we therefore fix the spectral index for polarized synchrotron emission at the best-fit value derived from the 2015 temperature data, corresponding to $\beta_s = -3.1$. This value is also consistent within the uncertainties with corresponding results derived by [Kogut et al. \(2007\)](#), [Dunkley et al. \(2009\)](#), [Bennett et al. \(2013\)](#), [Fuskeland et al. \(2014\)](#), [Vidal et al. \(2015\)](#), and [Krachmalnicoff et al. \(2018\)](#).

For thermal dust emission, the situation is more informative, since the HFI data constrain thermal dust emission more strongly than the LFI data constrain synchrotron emission. Focusing for the moment on the blue curve in Fig. 29, corresponding to the nominal data set considered in this paper, we observe a clear peak centred around $\beta_d \approx 1.60$, and with a width of 0.10–0.15. The distribution exhibits heavy tails toward both steep and shallow spectral indices, which is typical for noise-dominated data; these pixels are mostly located at high Galactic latitudes, where the dust amplitude is low. Motivated by these results, we adopt a Gaussian prior for the Commander analysis of $\beta_d = 1.60 \pm 0.10$ for the main analysis, acknowledging that the standard deviation quoted above over-estimates the intrinsic scatter in the dust population because of instrumental noise. Note that the uncertainty in this prior refers to the standard deviation of the map, not the error in the mean of the central value.

As mentioned in Sect. 3, the *Planck* 2018 HFI polarization measurements are associated with small but non-negligible




Fig. 23. Commander 2018 polarized thermal dust amplitude map at 5' FWHM resolution, evaluated at a mono-chromatic reference frequency of 353 GHz.




Fig. 24. Commander 2018 polarized synchrotron amplitude map at 40' FWHM resolution, evaluated at a mono-chromatic reference frequency of 30 GHz.

uncertainties in terms of polarization efficiencies, ϵ . By default, polarization efficiency corrections are not included in the analyses presented in this paper, but instead we assess their impact by comparing results with and without these corrections. The green

curve in the bottom panel Fig. 29 shows the distribution of β_d with application of these corrections at frequencies between 100 and 217 GHz. Overall, we see that these polarization efficiencies shift the distribution by $\Delta\beta_d = -0.03$.




Fig. 25. SMICA 2018 polarized thermal dust amplitude map at 12' FWHM resolution, evaluated at 353 GHz. No colour corrections have been applied to this map.




Fig. 26. SMICA 2018 polarized synchrotron amplitude map at 40' FWHM resolution, evaluated at 30 GHz. No colour corrections have been applied to this map.

The nominal polarization-efficiency corrections described in [Planck Collaboration III \(2020\)](#) and [Planck Collaboration V \(2020\)](#) do not include any robust estimates for the 353 GHz

channel, since the CMB signal that is used to estimate these corrections is faint at this frequency. However, it is reasonable to assume that it is associated with similar uncertainties as the




Fig. 27. GNILC 2018 polarized thermal dust amplitude map evaluated at 353 GHz. The angular resolution varies over the sky, as described in Remazeilles et al. (2011a). No colour corrections have been applied to this map.




Fig. 28. P - P scatter plot between the thermal dust polarization amplitude at 353 GHz, as estimated with GNILC and Commander. Colours indicate the density of points on a logarithmic scale.

other HFI channels. In Fig. 30, we show the β_d posterior distributions resulting from changing ϵ_{353} by 1% in either direction from its nominal value. In this case, we find that a shift of ϵ_{353} by 1% translates into a change in β_d of 0.013. Combined with the uncertainties arising from the 100 to 217 GHz frequencies,

we therefore consider the total systematic uncertainty on β_d due to polarization efficiency corrections to be 0.04.

The top panel in Fig. 31 shows the spatial distribution of β_d from the prior-free analysis without polarization efficiency corrections. In this plot the statistical power of the *Planck* observations to constrain the spectral index is seen very clearly from position to position, depending on the local dust polarization amplitude. Near the Galactic plane, the data are sufficiently strong to determine the spectral index well per resolution element, while at high latitudes the measurements are fully dominated by instrumental noise. The bottom panel shows the corresponding result when applying the supporting Gaussian prior. From this figure, it is clear that the β_d distribution and prior presented above are dominated by measurements in the Galactic plane, where the signal-to-noise ratio is substantially larger than at high Galactic latitudes.

Next, we perform a blind analysis of polarization spectral indices with SMICA. This analysis is performed by running SMICA with a foreground dimension of $N_{\text{fg}} = 2$ (that is, with a two-column foreground emissivity matrix F), as defined in Eq. (5), corresponding to synchrotron and thermal dust emission. Spectral priors are imposed during the multi-frequency fit so that synchrotron emission vanishes at 353 GHz and thermal dust emission vanishes at 30 GHz.

The results from these calculations are summarized in Fig. 32 for both E -mode and B -mode polarization. Parametric best-fits are indicated by dotted lines. These are, however, only the products of post-processing the raw SMICA results by fitting a modified black-body spectrum to the measured data points with χ^2 minimization. They do not correspond to active priors as they do in the Bayesian analysis discussed above. In these particular fits, polarization efficiency corrections are applied to the 100, 143, and 217 GHz data, and colour corrections are applied in post-analysis.




Fig. 29. Distribution of spectral indices for polarized synchrotron (*top panel*) and thermal dust (*bottom panel*) emission as estimated with Commander without applying any informative Gaussian prior. The synchrotron spectral index shown in this plot is estimated with a 5° FWHM smoothing scale, and the thermal dust spectral index is estimated with a 3° FWHM smoothing scale. For the thermal dust case, results are shown both with (green curve) and without (blue curve) applying polarization efficiency corrections at 100–217 GHz. The dashed lines in this case indicate Gaussian fits to the central peak.

The best-fit spectral parameters derived in this blind manner are $\beta_s = -3.10 \pm 0.06$ and $\beta_d = 1.53 \pm 0.01$, both corresponding to full-sky averages. Furthermore, these fits provide a statistically sufficient model across the full frequency range, as indicated by the residual spectra shown in the bottom panel of Fig. 32. All residuals are within 2σ of their statistical errors.

The SMICA measurements of the polarized thermal dust spectral index are in excellent agreement with the corresponding results presented in Planck Collaboration XI (2020), based on both frequency cross-correlation power spectra at high Galactic latitudes and simple colour ratios between the 217 and 353 GHz channels at low Galactic latitudes. At the same time, β_d is lower by 0.07 or 3σ compared to the Commander results presented above. To understand the origin of these differences, it is instructive to take a closer look at the 217/353 colour ratio, which is the fastest, simplest and most transparent estimator available.

The results from this estimator may be summarized as follows. We subtract one of the cleaned CMB maps from the Planck




Fig. 30. Effect on the spectral index of polarized thermal dust emission, β_d , when changing the polarization efficiency correction at 353 GHz, ϵ_{353} . A shift of ϵ_{353} by 1% translates into a change in β_d of 0.013.

217 and 353 GHz polarization HM split maps to form two statistically independent foreground-plus-noise maps. We smooth these maps to 3° FWHM to increase the effective signal-to-noise ratio per pixel. We then compute the cross-polarization amplitude between the two halves of the split, and we finally form the CMB-corrected colour ratio between the 217 and 353 GHz maps. Given some estimate of the thermal dust temperature, this ratio may then be easily translated into estimates of the thermal dust spectral index. We adopt a constant temperature of 19.6 K in the following.

First, we consider the impact of different CMB estimates produced by each of the four analysis pipelines. With the above procedure, we find median estimates of $\beta_d = 1.57, 1.54, 1.55$, and 1.54, when subtracting the Commander, NILC, SEVEM, and SMICA CMB polarization maps, respectively. Different noise-weighting and foreground-modelling assumptions thus account for $\Delta\beta_d \approx 0.03$.

Second, the effect of polarization efficiencies has already been addressed above in the context of Commander. We find similar sensitivities to the polarization efficiencies on the 217/353 colour ratio, as the median estimates for each of the four codes when applying these corrections are $\beta_d = 1.54, 1.52, 1.52$, and 1.52, corresponding to an effective shift of $\Delta\beta_d \approx 0.02$ –0.03.

Third and finally, a small effect is due to different bandpass treatments. Specifically, in Planck Collaboration XI (2020), bandpass integration effects are taken into account by the so-called colour correction technique, in which a multiplicative correction based on some fiducial spectral parameters is applied to the nominal thermal dust SED at a given reference frequency. The same approach is adopted for the SMICA results. In contrast, Commander performs a full integral over the product of the bandpass and the SED for each set of spectral parameters. These two different approaches agree to 0.07% at 143 GHz, 0.7% at 217 GHz, and 1.3% at 353 GHz. In sum, these small differences translate into a net shift of $\Delta\beta_d = 0.015$ in terms of the thermal dust spectral index.

Recognizing the significant systematic uncertainties on the thermal dust spectral index from both modelling aspects and polarization efficiencies, we adopt a total systematic uncertainty of 0.05, defined by the above shifts added in quadrature with a statistical uncertainty of 0.02 (Planck Collaboration XI 2020). As a single point estimate, we adopt the average value of




Fig. 31. Spatial distribution of the spectral index of polarized thermal dust emission, β_d , as estimated with *Commander*, adopting a smoothing scale of 3° FWHM. *Top panel*: no Gaussian prior is applied. *Bottom panel*: a Gaussian prior of $\beta_d = 1.60 \pm 0.10$ is applied. In both cases, the spectral index of synchrotron emission is fixed to $\beta_s = -3.1$.

the colour-ratio-derived estimates without polarization efficiency corrections, for a total final estimate of $\beta_d = 1.55 \pm 0.05$. This estimate is conservative, and corresponds to marginalizing over all analysis methods and known uncertainties.

5.4. Synchrotron and thermal dust angular power spectra

Finally we consider the angular power spectra of polarized synchrotron and thermal dust emission as estimated by *Commander* and *SMICA*. We estimate the *EE* and *BB* angular cross-spectra outside the common CMB mask for the half-mission split with *XPoI* (see [Tristram et al. 2005](#); [Planck Collaboration XI 2020](#) for details). The results from these calculations are summarized in Fig. 33. *Commander* results are shown in red (for thermal dust emission) and green (for synchrotron emission); *SMICA* results are shown in orange and light green. For comparison, direct 353 GHz cross-correlation results are shown in purple, derived using the same methodology as in [Planck Collaboration XI \(2020\)](#). As in that analysis, the best-fit *Planck* 2018 Λ CDM CMB spectrum, shown as a black solid line, has been subtracted from the raw estimate. Dotted coloured lines indicate best-fit power law fits to the *Commander* spectra, as defined by

$$D_\ell = q \left(\frac{\ell}{80} \right)^\alpha. \quad (15)$$

Overall, we find excellent agreement between the *Commander*, *SMICA*, and 353 GHz results, demonstrating that the derived component maps are robust with respect to specific algorithmic details for the particular angular ranges and sky coverage considered here.




Fig. 32. *Top panel*: synchrotron and thermal dust full-sky-averaged SEDs as estimated blindly by *SMICA*. Red and blue curves indicate thermal dust and synchrotron *E* modes, respectively, and orange and cyan curves indicate corresponding *B* modes. Dotted lines indicate the best-fit spectra for a power-law fit with $\beta_s = -3.10 \pm 0.06$ for synchrotron, and a modified blackbody fit with $\beta_d = 1.53 \pm 0.01$ and $T_d = 19.6$ K for thermal dust. *Bottom panel*: residual spectral energy densities relative to best-fit models, measured in units of the data uncertainty, σ_v .

Table 6 summarizes the angular power spectra in terms of best-fit power-law models for *Commander* and *SMICA*, and in terms of the *EE/BB* ratio, all derived using the same machinery as in [Planck Collaboration XI \(2020\)](#). For thermal dust, corresponding results are also given for the direct 353 GHz cross-correlation approach. Power-spectrum amplitudes have been colour corrected to monochromatic reference frequencies of 30 and 353 GHz for synchrotron and thermal dust emission, respectively. The two masks considered in Table 6 are defined in [Planck Collaboration XI \(2020\)](#). Note, however, that only thermal dust emission results are shown for the mask with a sky fraction of $f_{\text{sky}} = 0.42$. The signal-to-noise ratio for synchrotron emission is too low to support robust power spectrum estimates in the same region.

Overall, in terms of angular power spectra for polarized thermal dust emission, we find excellent agreement over 78% of the sky between the frequency cross-correlation technique and the *Commander* and *SMICA* component-separation techniques. The only statistically significant discrepancy is seen for the spatial power-law index parameter, α , for which formally a 6σ difference is observed between *Commander* and *SMICA*. However, we note that in terms of absolute values the difference is only $\Delta\alpha = 0.06$, and no systematic uncertainties are included in these numbers. Finally, it is worth noting that the two analyses are carried out with different angular resolutions, corresponding to $5'$ and $12'$ FWHM respectively. Due to its lower resolution, the *SMICA* analysis is somewhat more sensitive to high-multipole systematics than the *Commander* and 353 GHz analyses.




Fig. 33. *EE* (top panel) and *BB* (bottom panel) power spectra for synchrotron and thermal dust as computed from the Commander, SMICA, and 353 GHz frequency maps; see [Planck Collaboration XI \(2020\)](#) for algorithmic details. All spectra are evaluated outside the common polarization mask, over 78% of the sky. Dashed lines indicate the best-fit power-law fits for the Commander case, as reported in Table 6. Error bars indicate 3σ uncertainties. All spectra have been colour corrected to monochromatic reference frequencies of 30 GHz for synchrotron and 353 GHz for thermal dust emission, respectively.

We also observe excellent agreement between the Commander and SMICA maps in terms of polarized synchrotron emission. In addition, we note that the *BB/EE* ratio measured from the *Planck* 2018 data is 0.34, which is very similar to the corresponding value of 0.36 estimated from the *Planck* 2015 data.

The thermal dust *BB* spectrum is in general lower than the *EE* spectrum. For intermediate values of ℓ , this has been interpreted as the result of statistical alignment of filamentary structure in the interstellar medium with the local direction of the Galactic magnetic field ([Planck Collaboration XI 2020](#), and references therein). Empirically, this asymmetry of *BB* relative to *EE* extends to the lowest multipoles, as seen in Fig. 33 by the decrease of *BB* below the power law and, interestingly here, no such strong decrease for *EE*. The amount of asymmetry appears to depend on both multipole and sky fraction ([Planck Collaboration XI 2020](#)). For *BB* in particular, the departures from the power law at lower multipoles show a dependence on Galactic hemisphere (see the comparison on Northern and Southern cuts of the sky in [Planck Collaboration XI 2020](#)),

Table 6. Best-fit power-law parameters to the angular power spectra of synchrotron (30 GHz) and thermal dust emission (353 GHz), evaluated with the XPol power spectrum estimator ([Tristram et al. 2005](#)) as detailed in [Planck Collaboration XI \(2020\)](#).

	$q [\mu\text{K}_{\text{CMB}}^2]$	α
<i>Thermal dust, $f_{\text{sky}} = 0.42$, $\ell = 40\text{--}600$</i>		
Commander		
<i>EE</i>	60 ± 2	-0.39 ± 0.03
<i>BB</i>	32 ± 1	-0.49 ± 0.05
<i>BB/EE</i>		0.52
SMICA		
<i>EE</i>	62 ± 2	-0.18 ± 0.04
<i>BB</i>	32 ± 1	-0.45 ± 0.05
<i>BB/EE</i>		0.48
Frequency map cross-correlation		
<i>EE</i>	59 ± 2	-0.28 ± 0.04
<i>BB</i>	32 ± 1	-0.48 ± 0.06
<i>BB/EE</i>		0.50
<i>Thermal dust, $f_{\text{sky}} = 0.78$, $\ell = 40\text{--}600$</i>		
Commander		
<i>EE</i>	323 ± 4	-0.40 ± 0.01
<i>BB</i>	199 ± 3	-0.50 ± 0.02
<i>BB/EE</i>		0.57
SMICA		
<i>EE</i>	318 ± 4	-0.34 ± 0.01
<i>BB</i>	205 ± 3	-0.55 ± 0.02
<i>BB/EE</i>		0.54
Frequency map cross-correlation		
<i>EE</i>	313 ± 4	-0.41 ± 0.01
<i>BB</i>	187 ± 3	-0.50 ± 0.02
<i>BB/EE</i>		0.57
<i>Synchrotron, $f_{\text{sky}} = 0.78$, $\ell = 4\text{--}140$</i>		
Commander		
<i>EE</i>	2.3 ± 0.1	-0.84 ± 0.05
<i>BB</i>	0.8 ± 0.1	-0.76 ± 0.09
<i>BB/EE</i>		0.34
SMICA		
<i>EE</i>	2.4 ± 0.2	-0.88 ± 0.04
<i>BB</i>	0.9 ± 0.2	-0.75 ± 0.07
<i>BB/EE</i>		0.34

Notes. Frequency cross-correlation results are derived using precisely the same methodology as in [Planck Collaboration XI \(2020\)](#), while Commander and SMICA results are derived using the same power spectrum estimation tools, but applied to the half-mission maps presented in this paper. Note that all uncertainties are statistical, and do not account for systematic or modelling uncertainties. Power spectrum amplitudes refer to monochromatic reference frequencies of 30 and 353 GHz for synchrotron and thermal dust emission, respectively.

which in turn suggests some relationship to the large-scale structure of the magnetic field. Further discussion is beyond the scope of this paper.

Based on the best-fit power spectrum and SED parameters reported above, Fig. 34 summarizes the foreground-to-CMB ratio in terms of the quantity $f(\ell, \nu) = [C_\ell^{\text{fg}}(\nu)/C_\ell^{\text{CMB}}]^{1/2}$ as a function of both frequency and angular scale. As expected, the overall picture is very similar to that presented from the *Planck* 2015 data in [Planck Collaboration X \(2016\)](#), with one small but notable exception: Because the best-fit value of the optical depth of reionization is lower in the *Planck* 2018 ΛCDM model than in the corresponding 2015 model, the relative foregrounds-to-CMB




Fig. 34. Amplitude ratio between total polarized foregrounds and CMB as a function of both multipole moment and frequency, as defined by $f(\ell, \nu) = [C_\ell^{\text{fg}}(\nu)/C_\ell^{\text{CMB}}]^{1/2}$, with parameters derived from 78% of the sky as estimated by *Commander*. The *top* and *bottom* panels show *EE* and *BB* spectra, and the black and red contours in the latter corresponds to tensor-to-scalar ratios of $r = 0.0$ and 0.05 , respectively.

ratio is higher at low *EE* multipoles, further emphasizing the importance of accurate foreground modelling for large-scale polarization CMB analysis.

6. Conclusions

In this paper we have presented cleaned CMB temperature and polarization maps derived from the *Planck* 2018 data set, as well as new polarized synchrotron and thermal dust emission maps. These maps represent a new state-of-the-art characterization of the microwave sky.

The main scientific motivation underlying the work between the *Planck* 2015 and 2018 data releases has been reduced instrumental systematics, in particular for the polarization measurements. As demonstrated in this and companion papers, the work has been successful, as the updated *Planck* frequency maps exhibit significantly lower contamination on all angular scales. For polarization, we find that the lower systematics in frequency maps translates directly into lower systematics in CMB

and foreground maps. Additionally, new end-to-end CMB-plus-noise simulations have been constructed that more accurately reproduce residual systematics observed in the real data. For full-mission data, these simulations are accurate to $\lesssim 3\%$ for $\ell \lesssim 1500$ in both temperature and polarization. On smaller scales, non-negligible biases are observed, and caution is warranted when subjecting the maps to detailed statistical analysis on scales smaller than $\ell \gtrsim 1500$.

It is important to note that the 2018 data release does not represent a globally optimal reduction of the *Planck* time-ordered data that is ideal for all purposes. In particular, the updated data set does not include single detector maps, and the new frequency maps have complicated bandpass properties ([Planck Collaboration III 2020](#)). As a result, accurate reconstruction of astrophysical temperature foreground properties is non-trivial. Thus, while the *Planck* 2018 release represents a significant step forward in our understanding of the polarized microwave sky compared to the 2015 release, the associated temperature results, for which the astrophysics are richer, do not represent a similar improvement. Indeed, for several intensity applications we anticipate that external users may find the 2015 products more useful than the corresponding 2018 products. One concrete example of this is the *Planck* astrophysical sky model as presented in [Planck Collaboration X \(2016\)](#), which includes intensity estimates of both CO line emission and thermal dust emission. The same considerations apply both to *GNILC* and *Commander*; while chronologically formally superseded by the current results, we believe that the 2015 temperature astrophysical foreground models represent more accurate approximations to the true sky than the ones presented in the 2018 data release. To avoid confusion, we therefore do not release the corresponding 2018 foreground temperature products.

Fortunately, these issues are largely unimportant for CMB reconstruction purposes. The analyses presented in this paper and in [Planck Collaboration VI \(2020\)](#) reach the same conclusion regarding the CMB temperature results, namely that the *Planck* 2018 CMB temperature data are for all practical purposes statistically consistent with the corresponding 2015 rendition. Of course, this is the direct result of the very high signal-to-noise ratio of the *Planck* measurements, in that small variations in the processing procedure make very little difference in the final maps compared to the intrinsic sample variance of the true CMB sky.

For large-scale CMB polarization at $\ell \lesssim 50$, we find that the *Planck* 2018 data are compatible with end-to-end simulations. However, it is critical to note that the observations are *not* consistent with uncorrelated white noise at any angular scales. Any statistical analysis of the *Planck* 2018 polarization data must therefore always be accompanied by a corresponding analysis of the associated end-to-end simulations. In addition, analysis of half-data split sky maps is strongly encouraged in order to probe stability with respect to both noise and residual instrumental systematics.

In addition to improving the large-scale CMB polarization map, the new data processing also results in improved astrophysical polarization results. One concrete example of this is the fact that the *Planck* 2018 data for the first time allow a pixel-by-pixel estimation of the spectral index of thermal dust emission over the full sky. Corresponding analyses based on previous data sets invariably led to clearly nonphysical results obviously driven by instrumental systematics. With this new data set, we obtain a typical spectral index of polarized thermal dust emission of $\beta_d = 1.55 \pm 0.05$, where the uncertainty accounts both for systematic uncertainties and different analysis




Fig. 35. Polarization amplitude rms as a function of frequency and astrophysical components, evaluated at a smoothing scale of $40'$ FWHM. The green band indicates polarized synchrotron emission, and the red band indicates polarized thermal dust emission. The cyan curve shows the CMB rms for a Λ CDM model with $\tau = 0.05$, and is strongly dominated by E -mode polarization. The dashed black lines indicate the sum of foregrounds evaluated over three different masks with $f_{\text{sky}} = 0.83$, 0.52 , and 0.27 . The widths of the synchrotron and thermal dust bands are defined by the largest and smallest sky coverages.

techniques. This estimate is largely consistent with comparable results derived from temperature measurements. Also, for polarized synchrotron emission, we are for the first time able to fit the spectral index pixel-by-pixel, and obtain physically meaningful values, even if the signal-to-noise ratio is low; the full-sky averaged synchrotron spectral index for polarized emission is $\beta_s = -3.1 \pm 0.1$. For thermal dust emission we find a BB/EE angular power spectrum ratio of 0.5 , largely independent of sky fraction, while for synchrotron emission we find a lower ratio of 0.34 .

In Fig. 35 we plot the rms of the polarization amplitude as a function of frequency for polarized CMB, synchrotron, and thermal dust emission, evaluated with an angular resolution of $40'$ FWHM. The CMB component is estimated from a simulation drawn from the best-fit *Planck* 2018 Λ CDM spectrum, and is dominated by E -mode polarization. The synchrotron and thermal dust emission components are based on the *Commander* sky model, by cross-correlating half-mission sky maps. The dotted lines indicate the sum of the foreground components for three different masks, defined by thresholding the total *Commander* foreground model evaluated at 70 GHz, near the foreground minimum. Three masks are shown, corresponding to 27 , 52 , and 82% of the sky. The widths of the foreground bands are defined by the two extreme masks. This figure provides a convenient summary of the properties of the polarized sky in the CMB frequencies measured by *Planck*, and it updates the corresponding polarization panel of Fig. 51 in *Planck Collaboration X* (2016).

Before concluding we briefly summarize some important points regarding limitations and recommended usage of the various *Planck* component separation products presented in this paper.

- For polarization analysis, the *Planck* 2018 data products are superior to the 2015 products in all respects, and the new maps entirely supercede the previous release.
- For CMB temperature analysis, we consider the 2015 and 2018 data products as equivalent in terms of overall data quality. Most differences between the two generations of

cleaned CMB maps are due to different processing choices, rather than fundamental data quality. For instance, for *Commander* the 2018 CMB temperature maps are more constrained by data selection issues than the 2015 maps, and as a result the new maps are more contaminated by CO emission. In contrast, for *SMICA* some minor glitches regarding inter-frequency calibration have been resolved in the 2018 maps, and the new maps are therefore somewhat more reliable. For *NILC* and *SEVEM*, only small changes are observed between the two releases. In all cases, the differences are small, typically less than $2 \mu\text{K}$ at high Galactic latitudes with a smoothing scale of $80'$ FWHM.

- For temperature foreground analysis, the 2015 release provides a number of distinct advantages compared to the 2018 release, including no pixelization issues near bright sources in the Galactic plane, more transparent bandpass definitions, and, most importantly, the availability of robust single-bolometer and detector-set maps. For these reasons, we consider the 2015 temperature foreground products from both *Commander* and *GNILC* to be more reliable than the 2018 products. For the same reason, we anticipate the 2015 temperature data set to continue to play an important role for astrophysical component-separation purposes.
- The noise properties of the *Planck* observations are complicated both in temperature and polarization, and usage of end-to-end simulations is essential to capture all uncertainties. However, even the best currently available simulations are only accurate to a few percent in power. When employing these simulations for quantitative scientific analysis, it is essential to check that the statistic of choice is not sensitive to this level of uncertainty.

With these caveats in mind, we end our discussion by recalling the original motivation and goal of the *Planck* mission, namely “...to measure the fluctuations of the CMB with an accuracy set by fundamental astrophysical limits” (*Planck Collaboration* 2005). For temperature, this goal was achieved already with the *Planck* 2015 release. With the 2018 data release, *Planck* provides a new state-of-the-art for the field also in terms of polarization.

Acknowledgements. The *Planck* Collaboration acknowledges the support of: ESA; CNES, and CNRS/INSU-IN2P3-INP (France); ASI, CNR, and INAF (Italy); NASA and DoE (USA); STFC and UKSA (UK); CSIC, MINECO, JA, and RES (Spain); Tekes, AoF, and CSC (Finland); DLR and MPG (Germany); CSA (Canada); DTU Space (Denmark); SER/SSO (Switzerland); RCN (Norway); SFI (Ireland); FCT/MCTES (Portugal); ERC and PRACE (EU). A description of the *Planck* Collaboration and a list of its members, indicating which technical or scientific activities they have been involved in, can be found at <http://www.cosmos.esa.int/web/planck/planck-collaboration>. This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement numbers 687312, 776282 and 772253.

References

- Argüeso, F., Sanz, J. L., Herranz, D., López-Caniego, M., & González-Nuevo, J. 2009, *MNRAS*, **395**, 649
- Basak, S., & Delabrouille, J. 2012, *MNRAS*, **419**, 1163
- Basak, S., & Delabrouille, J. 2013, *MNRAS*, **435**, 18
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, **208**, 20
- Cardoso, J., Martin, M., Delabrouille, J., Betoule, M., & Patanchon, G. 2008, *IEEE J. Sel. Top. Sign. Proces.*, **2**, 735, special issue on Signal Processing for Astronomical and Space
- Cardoso, J.-F. 2017, *International Conference on Latent Variable Analysis and Signal Separation* (Springer, Springer International Publishing), 403
- Chon, G., Challinor, A., Prunet, S., Hivon, E., & Szapudi, I. 2004, *MNRAS*, **350**, 914
- Chu, M., Eriksen, H. K., Knox, L., et al. 2005, *Phys. Rev. D*, **71**, 103002

- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, **115**, 1693
- Cruz, M., Vielva, P., Martínez-González, E., & Barreiro, R. B. 2011, *MNRAS*, **412**, 2383
- Dame, T. M., Hartmann, D., & Thaddeus, P. 2001, *ApJ*, **547**, 792
- Delabrouille, J., Cardoso, J.-F., & Patanchon, G. 2003, *MNRAS*, **346**, 1089
- Delabrouille, J., Cardoso, J.-F., Le Jeune, M., et al. 2009, *A&A*, **493**, 835
- Dunkley, J., Spergel, D. N., Komatsu, E., et al. 2009, *ApJ*, **701**, 1804
- Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, *ApJS*, **155**, 227
- Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, **676**, 10
- Fernández-Cobos, R., Vielva, P., Barreiro, R. B., & Martínez-González, E. 2012, *MNRAS*, **420**, 2162
- Fernández-Cobos, R., Marcos-Caballero, A., Vielva, P., Martínez-González, E., & Barreiro, R. B. 2016, *MNRAS*, **459**, 441
- Finkbeiner, D. P. 2003, *ApJS*, **146**, 407
- Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Naess, S. K. 2014, *ApJ*, **790**, 104
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Gregory, P. C., Scott, W. K., Douglas, K., & Condon, J. J. 1996, *ApJS*, **103**, 427
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, **47**, 1
- Hivon, E., Mottet, S., & Ponthieu, N. 2017, *A&A*, **598**, A25
- Jewell, J., Levin, S., & Anderson, C. H. 2004, *ApJ*, **609**, 1
- Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005, *MNRAS*, **360**, 390
- Kogut, A., Dunkley, J., Bennett, C. L., et al. 2007, *ApJ*, **665**, 355
- Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. 2018, *A&A*, **618**, A166
- Leach, S. M., Cardoso, J.-F., Baccigalupi, C., et al. 2008, *A&A*, **491**, 597
- López-Caniego, M., Herranz, D., González-Nuevo, J., et al. 2006, *MNRAS*, **370**, 2047
- Mitra, S., Rocha, G., Górski, K. M., et al. 2011, *ApJS*, **193**, 5
- Monteserín, C., Barreiro, R. B., Vielva, P., et al. 2008, *MNRAS*, **387**, 209
- Murphy, T., Sadler, E. M., Ekers, R. D., et al. 2010, *MNRAS*, **402**, 2403
- Narcowich, F. J., Petrushev, P., & Ward, J. D. 2006, *SIAM J. Math. Anal.*, **38**, 574
- Planck Collaboration. 2005, ESA publication ESA-SCI(2005)/01 [arXiv:[astro-ph/0604069](https://arxiv.org/abs/astro-ph/0604069)]
- Planck Collaboration VIII. 2014, *A&A*, **571**, A8
- Planck Collaboration IX. 2014, *A&A*, **571**, A9
- Planck Collaboration XI. 2014, *A&A*, **571**, A11
- Planck Collaboration XII. 2014, *A&A*, **571**, A12
- Planck Collaboration XIV. 2014, *A&A*, **571**, A14
- Planck Collaboration XXIV. 2014, *A&A*, **571**, A24
- Planck Collaboration XXVII. 2014, *A&A*, **571**, A27
- Planck Collaboration XXX. 2014, *A&A*, **571**, A30
- Planck Collaboration I. 2016, *A&A*, **594**, A1
- Planck Collaboration II. 2016, *A&A*, **594**, A2
- Planck Collaboration VI. 2016, *A&A*, **594**, A6
- Planck Collaboration VIII. 2016, *A&A*, **594**, A8
- Planck Collaboration IX. 2016, *A&A*, **594**, A9
- Planck Collaboration X. 2016, *A&A*, **594**, A10
- Planck Collaboration XI. 2016, *A&A*, **594**, A11
- Planck Collaboration XII. 2016, *A&A*, **594**, A12
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Planck Collaboration XIV. 2016, *A&A*, **594**, A14
- Planck Collaboration XVI. 2016, *A&A*, **594**, A16
- Planck Collaboration XVII. 2016, *A&A*, **594**, A17
- Planck Collaboration XXI. 2016, *A&A*, **594**, A21
- Planck Collaboration XXVI. 2016, *A&A*, **594**, A26
- Planck Collaboration I. 2020, *A&A*, **641**, A1
- Planck Collaboration II. 2020, *A&A*, **641**, A2
- Planck Collaboration III. 2020, *A&A*, **641**, A3
- Planck Collaboration IV. 2020, *A&A*, **641**, A4
- Planck Collaboration V. 2020, *A&A*, **641**, A5
- Planck Collaboration VI. 2020, *A&A*, **641**, A6
- Planck Collaboration VII. 2020, *A&A*, **641**, A7
- Planck Collaboration VIII. 2020, *A&A*, **641**, A8
- Planck Collaboration IX. 2020, *A&A*, **641**, A9
- Planck Collaboration X. 2020, *A&A*, **641**, A10
- Planck Collaboration XI. 2020, *A&A*, **641**, A11
- Planck Collaboration XII. 2020, *A&A*, **641**, A12
- Planck Collaboration Int. XLVIII. 2016, *A&A*, **596**, A109
- Planck Collaboration Int. L. 2017, *A&A*, **599**, A51
- Plaszczynski, S., Montier, L., Levrier, F., & Tristram, M. 2014, *MNRAS*, **439**, 4048
- Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011a, *MNRAS*, **418**, 467
- Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011b, *MNRAS*, **410**, 2481
- Remazeilles, M., Dickinson, C., Banday, A. J., Bigot-Sazy, M.-A., & Ghosh, T. 2015, *MNRAS*, **451**, 4311
- Seljebotn, D. S., Bærdal, T., Eriksen, H. K., Mardal, K. A., & Wehus, I. K. 2019, *A&A*, **627**, A98
- Sunyaev, R. A., & Zeldovich, Y. B. 1970, *ApSS*, **7**, 3
- Tassis, K., & Pavlidou, V. 2015, *MNRAS*, **451**, L90
- Tristram, M., Macías-Pérez, J. F., Renault, C., & Santos, D. 2005, *MNRAS*, **358**, 833
- Vidal, M., Dickinson, C., Davies, R. D., & Leahy, J. P. 2015, *MNRAS*, **452**, 656
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, **70**, 083511
- Wehus, I. K., Fuskeland, U., Eriksen, H. K., et al. 2017, *A&A*, **597**, A131
-
- ¹ AIM, CEA, CNRS, Université Paris-Saclay, Université Paris-Diderot, Sorbonne Paris Cité, 91191 Gif-sur-Yvette, France
 - ² APC, AstroParticule et Cosmologie, Université Paris Diderot, CNRS/IN2P3, CEA/Irfu, Observatoire de Paris, Sorbonne Paris Cité, 10 rue Alice Domon et Léonie Duquet, 75205 Paris Cedex 13, France
 - ³ African Institute for Mathematical Sciences, 6-8 Melrose Road, Muizenberg, Cape Town, South Africa
 - ⁴ Aix-Marseille Univ., CNRS, CNES, LAM, Marseille, France
 - ⁵ Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Avenue, Cambridge CB3 0HE, UK
 - ⁶ Astrophysics & Cosmology Research Unit, School of Mathematics, Statistics & Computer Science, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa
 - ⁷ CITA, University of Toronto, 60 St. George St., M5S 3H8 Toronto, ON, Canada
 - ⁸ CNRS, IRAP, 9 Av. colonel Roche, BP 44346, 31028 Toulouse Cedex 4, France
 - ⁹ Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA
 - ¹⁰ California Institute of Technology, Pasadena, CA, USA
 - ¹¹ Centre for Theoretical Cosmology, DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK
 - ¹² Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
 - ¹³ Département de Physique Théorique, Université de Genève, 24 quai E. Ansermet, 1211 Genève 4, Switzerland
 - ¹⁴ Département de Physique, École normale supérieure, PSL Research University, CNRS, 24 rue Lhomond, 75005 Paris, France
 - ¹⁵ Departamento de Astrofísica, Universidad de La Laguna (ULL), 38206 La Laguna, Tenerife, Spain
 - ¹⁶ Departamento de Física, Universidad de Oviedo, C/ Federico García Lorca, 18, Oviedo, Spain
 - ¹⁷ Departamento de Física Matemática, Instituto de Física, Universidade de São Paulo, Rua do Matão, 1371 São Paulo, Brazil
 - ¹⁸ Department of Astrophysics/IMAPP, Radboud University, PO Box 9010, 6500 GL Nijmegen, The Netherlands
 - ¹⁹ Department of Physics & Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, British Columbia, Canada
 - ²⁰ Department of Physics & Astronomy, University of the Western Cape, 7535 Cape Town, South Africa
 - ²¹ Department of Physics and Astronomy, University of Sussex, Brighton BN1 9QH, UK
 - ²² Department of Physics, Gustaf Hållströmin katu 2a, University of Helsinki, Helsinki, Finland
 - ²³ Department of Physics, Princeton University, Princeton, NJ, USA
 - ²⁴ Department of Physics, University of California, Santa Barbara, CA, USA
 - ²⁵ Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, IL, USA
 - ²⁶ Dipartimento di Fisica e Astronomia G. Galilei, Università degli Studi di Padova, via Marzolo 8, 35131 Padova, Italy
 - ²⁷ Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, Via Saragat 1, 44122 Ferrara, Italy
 - ²⁸ Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 2, Roma, Italy
 - ²⁹ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria, 16, Milano, Italy
 - ³⁰ Dipartimento di Fisica, Università degli Studi di Trieste, Via A. Valerio 2, Trieste, Italy

- ³¹ Dipartimento di Fisica, Università di Roma Tor Vergata, Via della Ricerca Scientifica, 1, Roma, Italy
- ³² European Space Agency, ESAC, Planck Science Office, Camino bajo del Castillo, s/n, Urbanización Villafranca del Castillo, Villanueva de la Cañada, Madrid, Spain
- ³³ European Space Agency, ESTEC, Keplerlaan 1, 2201, AZ Noordwijk, The Netherlands
- ³⁴ Gran Sasso Science Institute, INFN, Viale F. Crispi 7, 67100 L'Aquila, Italy
- ³⁵ Haverford College Astronomy Department, 370 Lancaster Avenue, Haverford, PA, USA
- ³⁶ Helsinki Institute of Physics, Gustaf Hållströmin katu 2, University of Helsinki, Helsinki, Finland
- ³⁷ INAF – OAS Bologna, Istituto Nazionale di Astrofisica – Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Area della Ricerca del CNR, Via Gobetti 101, 40129 Bologna, Italy
- ³⁸ INAF – Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, Padova, Italy
- ³⁹ INAF – Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, Trieste, Italy
- ⁴⁰ INAF, Istituto di Radioastronomia, Via Piero Gobetti 101, 40129 Bologna, Italy
- ⁴¹ INAF/IASF Milano, Via E. Bassini 15, Milano, Italy
- ⁴² INFN – CNAF, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ⁴³ INFN, Sezione di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ⁴⁴ INFN, Sezione di Ferrara, Via Saragat 1, 44122 Ferrara, Italy
- ⁴⁵ INFN, Sezione di Milano, Via Celoria 16, Milano, Italy
- ⁴⁶ INFN, Sezione di Roma 1, Università di Roma Sapienza, Piazzale Aldo Moro 2, 00185 Roma, Italy
- ⁴⁷ INFN, Sezione di Roma 2, Università di Roma Tor Vergata, Via della Ricerca Scientifica, 1, Roma, Italy
- ⁴⁸ Imperial College London, Astrophysics group, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK
- ⁴⁹ Institut d'Astrophysique Spatiale, CNRS, Univ. Paris-Sud, Université Paris-Saclay, Bât. 121, 91405 Orsay Cedex, France
- ⁵⁰ Institut d'Astrophysique de Paris, CNRS (UMR7095), 98 bis boulevard Arago, 75014 Paris, France
- ⁵¹ Institute Lorentz, Leiden University, PO Box 9506, 2300 RA Leiden, The Netherlands
- ⁵² Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK
- ⁵³ Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway
- ⁵⁴ Instituto de Astrofísica de Canarias, C/Vía Láctea s/n, La Laguna, Tenerife, Spain
- ⁵⁵ Instituto de Física de Cantabria (CSIC-Universidad de Cantabria), Avda. de los Castros s/n, Santander, Spain
- ⁵⁶ Istituto Nazionale di Fisica Nucleare, Sezione di Padova, Via Marzolo 8, 35131 Padova, Italy
- ⁵⁷ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA
- ⁵⁸ Jodrell Bank Centre for Astrophysics, Alan Turing Building, School of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK
- ⁵⁹ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA
- ⁶⁰ Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK
- ⁶¹ Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU, WPI), UTIAS, The University of Tokyo, 277-8583 Chiba, Japan
- ⁶² Laboratoire de Physique Subatomique et Cosmologie, Université Grenoble-Alpes, CNRS/IN2P3, 53, rue des Martyrs, 38026 Grenoble Cedex, France
- ⁶³ Laboratoire de Physique Théorique, Université Paris-Sud 11 & CNRS, Bâtiment 210, 91405 Orsay, France
- ⁶⁴ Lawrence Berkeley National Laboratory, Berkeley, CA, USA
- ⁶⁵ Low Temperature Laboratory, Department of Applied Physics, Aalto University, 00076 Aalto Espoo, Finland
- ⁶⁶ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany
- ⁶⁷ Mullard Space Science Laboratory, University College London, Surrey RH5 6NT, UK
- ⁶⁸ NAOC-UKZN Computational Astrophysics Centre (NUCAC), University of KwaZulu-Natal, 4000 Durban, South Africa
- ⁶⁹ National Centre for Nuclear Research, ul. A. Soltana 7, 05-400 Otwock, Poland
- ⁷⁰ Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, Bartycka 18, 00-716 Warsaw, Poland
- ⁷¹ SISSA, Astrophysics Sector, via Bonomea 265, 34136 Trieste, Italy
- ⁷² San Diego Supercomputer Center, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093, USA
- ⁷³ School of Chemistry and Physics, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, 4000 Durban, South Africa
- ⁷⁴ School of Physical Sciences, National Institute of Science Education and Research, HBNI, 752050 Jatni, Odisha, India
- ⁷⁵ School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff CF24 3AA, UK
- ⁷⁶ School of Physics and Astronomy, Sun Yat-sen University, 2 Daxue Rd, Tangjia, Zhuhai, PR China
- ⁷⁷ School of Physics, Indian Institute of Science Education and Research Thiruvananthapuram, Maruthamala PO, Vithura, 695551 Thiruvananthapuram, Kerala, India
- ⁷⁸ Simon Fraser University, Department of Physics, 8888 University Drive, Burnaby, BC, Canada
- ⁷⁹ Sorbonne Université, Observatoire de Paris, Université PSL, École normale supérieure, CNRS, LERMA, 75005 Paris, France
- ⁸⁰ Sorbonne Université-UPMC, UMR7095, Institut d'Astrophysique de Paris, 98bis boulevard Arago, 75014 Paris, France
- ⁸¹ Space Science Data Center – Agenzia Spaziale Italiana, Via del Politecnico snc, 00133 Roma, Italy
- ⁸² Spacem Sciences Laboratory, University of California, Berkeley, CA, USA
- ⁸³ The Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, AlbaNova, 106 91 Stockholm, Sweden
- ⁸⁴ UPMC Univ Paris 06, UMR7095, 98 bis boulevard Arago, 75014 Paris, France
- ⁸⁵ Université de Toulouse, UPS-OMP, IRAP, 31028 Toulouse Cedex 4, France
- ⁸⁶ Warsaw University Observatory, Aleje Ujazdowskie 4, 00-478 Warszawa, Poland

Appendix A: Commander

The **Commander** analysis framework as applied to previous *Planck* releases is described in detail by Eriksen et al. (2004, 2008) and Planck Collaboration X (2016). This approach implements a standard Bayesian fitting procedure based on Monte Carlo and Gibbs sampling, in which an explicit parametric model including cosmological, astrophysical, and instrumental parameters is fitted to the observations through the posterior distribution.

Due to its very general approach to CMB analysis, it is more appropriate to refer to **Commander** as a framework rather than as a specific and well-defined algorithm. For instance, the implementation that is employed for the *Planck* 2017 analysis has been re-written from scratch compared to the 2015 version, and the current version is sometimes referred to as **Commander2** (Seljebotn et al. 2019). The main difference between the old and the new implementations is their different choice of basis functions for the amplitude degrees of freedom, and their different treatment of instrumental beams. While **Commander1** adopted real-space pixels as its fundamental basis set and required uniform angular resolution across frequencies, **Commander2** adopts spherical harmonics as its fundamental basis set and supports different angular resolutions at different frequencies. As a result, the new implementation supports signal reconstruction at the full angular resolution of the *Planck* observations.

A.1. Amplitude sampling algorithm

For full algorithmic specifics regarding the new implementation, see Seljebotn et al. (2019); here we review only the main equations. First, we adopt a general signal model on the following form,

$$s_\nu(\theta) = s_\nu(\mathbf{a}_i, \beta_i, g_\nu, \mathbf{m}_\nu) \quad (\text{A.1})$$

$$= g_\nu \sum_{i=1}^{N_{\text{comp}}} \mathbf{F}_\nu^i(\beta_i) \mathbf{a}_i \quad (\text{A.2})$$

where $\mathbf{a}_i$ is an amplitude vector 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 , and $\mathbf{m}_\nu$ are monopole and dipole amplitudes. The quantity $\mathbf{F}_\nu^i(\beta_i)$ is a general projection operator that translates from the reference amplitude vector to the basis of the observed data at a given frequency. As such, it accounts for both the choice of basis functions, and for spectral effects such as the frequency dependence of the component in question and unit conversions.

As mentioned above, **Commander1** adopted pixels as its basis set for all diffuse components, requiring identical angular resolution at all frequencies. In this case, the projection operator reduces to the so-called mixing matrix, $\mathbf{F} = \mathbf{M}$, which translates signal amplitudes from a reference frequency to any other observed frequency. In contrast, **Commander2** employs different types of basis functions for different components. For diffuse components, it adopts spherical harmonics, and the projection operator therefore becomes the product of the mixing matrix, which is defined in pixel space, and a spherical harmonics transform, $\mathbf{F} = \mathbf{M}\mathbf{Y}$. For compact objects (radio sources in the current analysis), the map projection is performed through a local real-space FBeCoP template per source, $\mathbf{B}_F$, and therefore $\mathbf{F} = \mathbf{M}\mathbf{B}_F$. Finally, fixed template corrections such as monopole, dipole, or zodiacal light corrections, summarized by some overall real-space template matrix per frequency, T_ν , are implemented directly as $\mathbf{F} = \mathbf{T}$, and the fitted parameters are thus

Table A.1. Overview of spectral and spatial priors adopted in the **Commander** analysis.

Component	Prior
<i>Temperature</i>	
CMB	$T_{\text{CMB}} = 2.755\text{K}$
Low-frequency component	$\beta_{\text{lf}} = -3.1 \pm 0.5$ $A_{\text{lf}} = 10^5 \mu\text{K}_{\text{RJ}}^2$ $\theta_{\text{lf}} = 30' \text{ FWHM}$
Thermal dust emission	$\beta_{\text{d}} = 1.55 \pm 0.1$ $T_{\text{d}} = (19.5 \pm 3) \text{K}$ Cosine apodization, $5000 \leq \ell \leq 6000$
CO emission	Spatially uniform line ratios $A_{\text{CO}} = 10^4 \mu\text{K}_{\text{RJ}}^2$ $\theta_{\text{CO}} = 15' \text{ FWHM}$
Radio source component	$a_{\text{cs}} \geq 0$
<i>Polarization</i>	
CMB	$T_{\text{CMB}} = 2.755 \text{K}$
Synchrotron emission	β_{s} spatially uniform $A_{\text{s}} = 10^2 \mu\text{K}^2$ $\theta_{\text{s}} = 40' \text{ FWHM}$
Thermal dust emission	$\beta_{\text{d}} = 1.6 \pm 0.1$ $T_{\text{d, pol}} = T_{\text{d, int}}$ $A_{\text{d}} = 50 \mu\text{K}_{\text{RJ}}^2$ $\theta_{\text{d}} = 10' \text{ FWHM}$

Notes. Parameters denoted A and θ correspond to the spatial angular power spectrum prior, as defined in Eq. (A.10), where A is defined relative to the reference frequency of the component in question in units of $\mu\text{K}_{\text{RJ}}^2$.

defined directly as the template amplitude at the respective frequency.

Computationally speaking, by far the most expensive part in **Commander** is to fit for the linear amplitudes, which corresponds to sampling from the conditional distribution $P(\mathbf{a}|\mathbf{d}, \dots)$. As shown by, e.g., Jewell et al. (2004) and Wandelt et al. (2004), this can be done by solving the so-called Wiener filter equation by conjugate gradients,

$$(\mathbf{S}^{-1} + \mathbf{P}^T \mathbf{N}^{-1} \mathbf{P}) \mathbf{a} = \mathbf{P}^T \mathbf{N}^{-1} \mathbf{d} + \mathbf{P}^T \mathbf{N}^{-1/2} \omega_1 + \mathbf{S}^{-1} \omega_2. \quad (\text{A.3})$$

Here $\mathbf{S}$ is the (prior) covariance matrix of the signal amplitudes, $\mathbf{P}$ is the end-to-end projection operator from amplitude space to data space, $\mathbf{N}$ is the data noise covariance matrix, and ω_i are Gaussian random vectors with zero mean and unit variance. If the maximum posterior solution is desired rather than a sample drawn from the posterior, one simply sets ω_i to zero.

The computational expense for solving this equation depends directly on the complexity of the projection operator, $\mathbf{P}$. In most **Commander1**-type analyses, which employ a pixel basis for all components and impose no spatial priors, i.e., $\mathbf{S} = 0$, all matrix multiplications are given by diagonal matrices. In contrast, as implemented in **Commander2**, $\mathbf{P}$ involves one spherical harmonic transform per frequency channel, and the computational scaling of the left-hand side increases from $O(N_{\text{pix}})$ to $O(N_{\text{pix}}^{3/2})$. Accordingly, the associated CPU time required per sample increases from minutes to tens of hours. This additional cost, however, is very well justified by the new flexibility in terms of beam treatment, which now supports arbitrary resolution at each frequency.

By virtue of being a Gibbs-sampling procedure, *Commander* requires one sampling step for each parameter under consideration, such as spectral or calibration parameters. However, these parameters are sampled with exactly the same methods in *Commander2* as in *Commander1*, and the details will not be repeated here; see [Eriksen et al. \(2008\)](#) and [Planck Collaboration X \(2016\)](#).

A.2. *Commander 2018 signal model and priors*

As discussed in Sects. 2 and 3, the maps provided in the *Planck* 2018 release include only full frequency maps, not individual detector or detector-set maps. This has significant consequences for our ability to reconstruct some important parameters, in particular CO line emission. For this reason, we adopt a simpler signal model in the 2018 analysis than in the 2015 analysis. Explicitly, the basic model considered in the current analysis, as defined in Rayleigh–Jeans temperature units, reads¹¹

$$s_\nu = g_\nu \left[Y a_{\text{cmb}} \gamma(\nu) \right. \quad (\text{A.4})$$

$$+ Y a_{\text{lf}} \left(\frac{\nu}{\nu_{\text{lf}}} \right)^{\beta_{\text{lf}}(p)} \quad (\text{A.5})$$

$$+ Y a_{\text{d}} \left(\frac{\nu}{\nu_{\text{d}}} \right)^{\beta_{\text{d}}(p)+1} \left(\frac{e^{h\nu_{\text{d}}/kT_{\text{d}}(p)} - 1}{e^{h\nu/kT_{\text{d}}(p)} - 1} \right) \quad (\text{A.6})$$

$$+ Y a_{\text{co}} h_\nu \quad (\text{A.7})$$

$$+ \sum_{i=1}^{N_{\text{arc}}} B_{\text{F},\nu,i} a_{\text{cs},i} \left(\frac{\nu}{\nu_{\text{cs}}} \right)^{\alpha_{\text{cs}}(p)} \quad (\text{A.8})$$

$$+ \sum_{i=1}^{N_{\text{temp}}} T_{\nu} a_{\text{temp},i} \left. \right], \quad (\text{A.9})$$

where the various components correspond to, from top to bottom, CMB, low-frequency/synchrotron emission, thermal dust emission, CO line emission, compact objects, and template corrections. The full model applies only to temperature analysis, since CO line emission, point sources, and template corrections are all omitted from the polarization analysis. For temperature, we refer to the second term as a “low frequency component”, since it includes both synchrotron, free-free, and anomalous microwave emission, while for polarization we refer to it as “synchrotron”, since that is the only component that is significantly detected at low frequencies in polarization.

In the above expression, $\gamma(\nu)$ is the conversion factor between thermodynamic and Rayleigh–Jeans units, $\nu_{\text{lf}} = 30$ GHz is the reference frequency for the low-frequency component, $\nu_{\text{d}} = 857$ GHz is the thermal dust reference frequency for temperature (353 GHz for polarization), h_ν is the CO line ratio between 100 and 217 or 353 GHz, respectively, and all other quantities are defined above.

To complete the specification of a model used for Bayesian analysis, we also have to choose priors for the various parameters. Starting with the spectral parameters, we adopt the same types of priors as in previous analyses. Technically speaking, for each parameters these are given as the product of three different priors, each serving a different purpose. First, we impose a uniform prior between two hard limits for numerical reasons; this makes

it easier to precompute look-up tables for all unit conversion and bandpass integration quantities. Second, we impose a Jeffreys’ ignorance prior, which effectively normalizes posterior volume effects due to the specific choice of parametrization. Third, and by far most importantly, we adopt Gaussian informative priors with physically motivated means and standard deviations for all spectral parameters. The values of these are listed in Table A.1.

Next, we need to specify spatial priors on the amplitude degrees of freedom. With the new *Commander2* implementation – which models all diffuse components, not just the CMB, in spherical harmonic space – we are now able to impose informative spatial priors on the foregrounds through the signal covariance matrix, $\mathbf{S}$, in Eq. (A.3). In this paper, we define this matrix in harmonic space in terms of a standard angular power spectrum, $\mathcal{D}_\ell$, per component. In principle, this could be used to enforce physically motivated power spectra for each component, for instance a Λ CDM spectrum for the CMB, or a power-law spectrum for synchrotron or thermal dust emission. However, in the present analysis, we choose to be minimally constraining, and simply use this new feature to enforce smoothness of the foreground components on small scales. For all components except thermal dust intensity, we implement this by defining a reference prior spectrum given by the shape of a Gaussian smoothing kernel multiplied by an overall amplitude that is larger than the actual sky signal in the high signal-to-noise regime. Thus, the prior takes the form

$$\mathcal{D}_\ell^i = A^i e^{-\ell(\ell+1)\sigma^2}, \quad (\text{A.10})$$

where $\sigma^2 = \theta_{\text{FWHM}}^2 / (8 \ln 2)$, θ_{FWHM} is the FWHM of the desired Gaussian smoothing kernel in radians, and A^i is the uniform power spectrum amplitude. This type of prior simply acts as a smooth apodization of the high- ℓ spectra, and its main function is to prevent the ringing that would otherwise occur around objects with a sharp cutoff in harmonic space, given by some ℓ_{max} . For the special case of thermal dust emission in intensity, we employ a simple cosine apodization between $\ell = 5000$ and 6000 , in order to retain as much signal as possible. The spatial prior values adopted for the various components are summarized in Table A.1.

Finally, we need to impose priors on the zero levels and dipoles for each map. For HFI zero levels, we adopt the CIB offsets defined in Table 6 of [Planck Collaboration VIII \(2016\)](#), while for the LFI we adopt a vanishing monopole at 30 GHz. At 44 and 70 GHz, we impose no priors on the zero levels, but rather fit them freely, obtaining best-fit values of 17 and 21 μK , respectively. For the HFI channels between 100 and 545 GHz, the best-fit zero levels are 12.4 μK , 22.0 μK , 71.0 μK , 431 μK , and 0.346 MJy sr^{-1} , respectively, while the 857 GHz zero level is fixed at 0.64 MJy sr^{-1} from [Planck Collaboration VIII \(2016\)](#). For comparison, the nominal CIB offsets listed in the same reference correspond to 12 μK , 21 μK , 68 μK , 451 μK , and 0.35 MJy sr^{-1} MJy sr^{-1} .

We only fit for dipoles in the 70 and 100 GHz channels, a choice determined by inspection of the residual maps resulting from an initial analysis in which no dipoles are fitted. The best-fit *Commander*-derived amplitudes of the 70 and 100 GHz dipoles are 2.0 and 2.3 μK , respectively.

A.3. *Sampling compact objects*

A significant new feature of *Commander2* is its ability to fit compact sources with multi-resolution frequency maps, while at the same time accounting for the full asymmetric beam structure at

¹¹ For simplicity, bandpass integration and unit conversion effects are omitted from this expression. Such effects are handled as in earlier implementations, through construction of fast, splined look-up tables based on direct bandpass convolution for the relevant parameters; see [Planck Collaboration IX \(2014\)](#) for an overview of the basic equations.

each frequency. As described above, for a single source this is done through the following parametric model,

$$s_{\nu}(p) = B_{F,\nu,i} a_{cs} \left(\frac{\nu}{\nu_{cs}} \right)^{\alpha_{cs}(p)}, \quad (\text{A.11})$$

where $B_{F,\nu,i}$ is a full FEBeCoP template evaluated at the pixel closest to the point source in question, a is the source amplitude in units of mJy, and we assume a simple power-law frequency scaling with a spectral index of α_{cs} in flux density units (in the current analysis we only consider this component for radio sources, for which a power-law model is a reasonable spectrum). Thus, each source is associated with only two free parameters, the amplitude a and the spectral index α , across all frequencies. This simple two-parameter model, however, is not likely to be adequate for a full fit between 30 and 857 GHz for many sources. As a result, when fitting the free parameters, we only include frequencies between 30 and 143 GHz in the actual fit; however, the resulting parameters are used to extrapolate to the higher frequencies when fitting other parameters.

Source locations are not identified internally in *Commander*, but rather defined by external catalogues. Unfortunately, no full-sky, deep, and high-resolution catalogue of radio sources exists for the microwave frequencies, and we therefore construct a hybrid catalogue by combining four different catalogues. First, we include all sources in the AT20G catalogue for declinations below -15° , for a total of 4499 sources. By virtue of being closest to our frequency range, this catalogue is adopted as an overall reference. Thus, we compute an effective source number density per area of AT20G sources, and adopt this as a threshold density. This threshold is then applied to the GB6 catalogue, including all sources above a flux density defined by requiring that the area number density is the same as for AT20G. This results in 5814 GB6 sources. Next, for sky regions not covered by either AT20G or GB6, we employ the same algorithm to the NVSS catalogue, resulting in 1527 NVSS sources. Finally, we also include all sources found in the PCCS2 catalogue, except for excluding duplicates in the already considered catalogues; this results in 352 unique sources. Thus, at this stage, the full sky has been populated by sources with a nearly uniform number density, for a total of 12 192 sources.

It is important to note that the catalogue positions defined above are only used as candidates for source positions. Including a non-existing source will not bias any other parameter, since its relevant parameters are fitted jointly with all other parameters; the only detrimental effect of including too many sources is a slight increase in the overall noise level.

Figure A.1 shows an enlargement of the final *Commander* compact source map for 30 and 100 GHz, generated as described above. The plot shows a $10^\circ \times 10^\circ$ region centred on the South Galactic Pole, for which the *Planck* scanning strategy provides relatively poor cross-linking. As a result, the asymmetric properties of the 30 GHz beams are clearly visible. Another feature seen in these plots is the large number of overlapping sources in the 30 GHz map. If we included only this single frequency while fitting the spectral properties of the sources, there would be significant degeneracies between such overlapping sources. However, when we include higher frequencies, for which the beams are smaller and neighboring sources overlap less, these degeneracies are effectively broken.

A.4. Confidence masks

For *Commander*, we establish the following prescription for defining a temperature confidence mask. First, the base tempera-




Fig. A.1. Enlargement of the compact source map fitted with *Commander* using real-space spatial FEBeCoP templates and a power-law spectral model. Shown here is a $10^\circ \times 10^\circ$ region centred on the South Galactic Pole (SGP), and the *top* and *bottom* panels showing the effective point source maps at 30 and 100 GHz, respectively. Note the significantly asymmetric beam structures in the 30 GHz map.

ture mask is defined by smoothing the *Commander* χ^2 map with a $30'$ FWHM Gaussian beam, suppressing instrumental noise fluctuations, and then thresholding the smoothed map at a value of 50, which corresponds to a roughly 4σ confidence level at high Galactic latitudes. This mask removes any pixel for which the *Commander* model obviously breaks down in terms of total χ^2 . However, based on frequency residual maps, one does observe residuals corresponding to specific components that are not easily picked up by the total χ^2 . To capture these, we augment the base mask with three specifically targeted masks. First, we




Fig. A.2. Commander masks in temperature (*top panel*) and polarization (*bottom panel*).

remove any pixels brighter than 10 mK, to eliminate particularly bright radio sources. Second, we exclude by hand the Virgo and Coma clusters and the Crab Nebula. Third, noting that CO emission represents a particularly difficult problem with the current data set, we smooth the Commander 2018 CO emission map shown in Fig. F.1 with a 30' FWHM beam, and exclude any pixels for which the CO amplitude is larger than $50 \mu\text{K}_{\text{RJ}}$ at 100 GHz.

The polarization confidence mask is constructed in a similar way, with a few specific modifications. First, a base mask is produced by smoothing and thresholding the χ^2 map shown in Fig. 22. Second, we remove all pixels for which the polarized thermal dust amplitude smoothed to 3° FWHM is larger than $20 \mu\text{K}_{\text{RJ}}$ at 353 GHz (see Fig. 23); this mask excludes pixels for which large values are observed in the frequency residual maps shown in Fig. 22, but which are not robustly picked up by the χ^2 values. Third, we remove all pixels corresponding to the cosmic ray contaminated ring discussed in Sect. 4.1. Fourth, we remove particularly bright point sources based on the PCCS2 source catalogue. The resulting masks for both temperature and polarization are shown in Fig. A.2.

A.5. Comparison between low- ℓ likelihood and full-resolution Commander maps

As described in Planck Collaboration V (2020), the Commander algorithm is used to generate the low- ℓ temperature sky map that feeds the Planck 2018 CMB likelihood, as it was in previous Planck releases. The set-up adopted for that analysis is, however, somewhat different than the one adopted for the main analysis presented in this paper, primarily due to the different angular scales in question. Specifically, since the likelihood map is only used for large angular scales, covering primarily only $\ell \leq 30$, the full analysis is carried out with Commander1, and all input frequency maps are smoothed to a common angular resolution of 40' FWHM, similar to the Planck 2015 processing.




Fig. A.3. *Top panel:* Commander CMB temperature map used for the Planck low- ℓ temperature likelihood analysis, smoothed to 60' FWHM resolution. The grey region indicates the mask adopted for the likelihood analysis, which retains 86% of the sky. *Bottom panel:* difference between the low- ℓ likelihood and full-resolution Commander CMB temperature maps, smoothed to 60' FWHM resolution.

Finally, the Commander1 low- ℓ analysis internally estimates the CMB power spectrum as one of the parameters in the Bayesian parametric model, and the corresponding Gaussian constrained realization samples (Eriksen et al. 2008) provide the necessary inputs for the Blackwell-Rao likelihood estimator employed by the Planck temperature-only likelihood (Chu et al. 2005). For the combined Planck temperature and polarization likelihood, which is map-based rather than power-spectrum-based, a single constrained-realization sample is adopted as the low- ℓ likelihood temperature component. We have verified that the choice of the particular sample used has no significant effect on the actual power spectrum outside the analysis mask.

The top panel of Fig. A.3 shows the low- ℓ likelihood temperature map with the corresponding likelihood mask marked in grey. The bottom panel shows the difference map with respect to the full-resolution Commander map, after the latter is smoothed to 60' FWHM resolution. Over most of the sky, the absolute difference between the two maps is $\lesssim 2 \mu\text{K}$, increasing to $5 \mu\text{K}$ near the Galactic plane. A few bright spots exhibit differences at the $10 \mu\text{K}$ level. These differences are dominated by thermal dust emission (see, e.g., Fig. F.1), and are in effect due to the different angular resolutions adopted for the two analyses. Since the likelihood analysis is performed at an angular resolution of 40' FWHM, the thermal dust spectral index and temperature are also estimated at an angular resolution of 40' FWHM. In contrast, the high-resolution analysis estimates the thermal dust spectral index at 10' FWHM, and the corresponding temperature




Fig. A.4. *Top panel:* low- ℓ temperature power spectra derived from the low- ℓ likelihood Commander map (red) and from the full-resolution Commander map (blue), both evaluated over the low- ℓ likelihood mask shown in Fig. A.3. *Bottom panel:* difference between the two spectra shown in the top panel.

at $5'$ resolution. As a consequence, the assumed spectral priors have a relatively larger impact in the high-resolution analysis than in the low- ℓ analysis.

Nevertheless, these differences are small in terms of absolute numbers, and have a negligible impact on the derived angular power spectrum. This is explicitly shown in Fig. A.4, in which the top panel shows the individual spectra computed from each of the two maps, and the bottom panel shows their difference. Overall, the absolute differences are smaller multipole-by-multipole than $10 \mu\text{K}^2$, corresponding to $\lesssim 1\%$ in absolute power and $\lesssim 0.05 \sigma$ in terms of cosmic variance. There is also no overall trend in the difference spectrum that might pull systematically on cosmological parameters. The two maps are statistically equivalent in terms of temperature power spectra.

Appendix B: Needlet Internal Linear Combination

The goal of NILC is to estimate the CMB from multi-frequency observations while minimizing the contamination from Galactic and extragalactic foregrounds, and instrumental noise. The method makes a linear combination of the data from the input maps with minimum variance on a frame of spherical wavelets called needlets (Narcowich et al. 2006). Due to their unique properties, needlets enable localized filtering in both pixel space and harmonic space. Localization in pixel space allows the weights of the linear combination to adapt to local conditions of foreground contamination and noise, whereas localization in harmonic space allows the method to favour foreground rejection on large scales and noise rejection on small scales. Needlets permit the weights to vary smoothly on large scales and rapidly on small scales, which is not possible by cutting the sky into zones prior to processing (Delabrouille et al. 2009).

The NILC pipeline (Basak & Delabrouille 2012, 2013) is applicable to scalar fields on the sphere, hence we work separately

Table B.1. Needlet bands used in the NILC analysis.

Band	$\ell_{\min}$	ℓ_{peak}	$\ell_{\max}$	N_{side}
$j = 1 \dots \dots \dots$	0	0	100	64
2 $\dots \dots \dots$	0	100	200	128
3 $\dots \dots \dots$	100	200	300	256
4 $\dots \dots \dots$	200	300	400	256
5 $\dots \dots \dots$	300	400	600	512
6 $\dots \dots \dots$	400	600	800	512
7 $\dots \dots \dots$	600	800	1000	512
8 $\dots \dots \dots$	800	1000	1400	1024
9 $\dots \dots \dots$	1000	1400	1800	1024
10 $\dots \dots \dots$	1400	1800	2200	2048
11 $\dots \dots \dots$	1800	2200	2800	2048
12 $\dots \dots \dots$	2200	2800	3400	2048
13 $\dots \dots \dots$	2800	3400	4000	2048




Fig. B.1. Needlet bands used in the analysis. The solid black line shows the normalization of the needlet bands, i.e., the total filter applied to the original map after needlet decomposition and synthesis of the output map from needlet coefficients.

on maps of temperature and the E and B modes of polarization. The decomposition of input polarization maps into E and B is performed on the full sky. At the end of the processing, the CMB Q and U maps are reconstructed from the E and B maps.

Prior to applying NILC, all of the input maps are convolved or deconvolved in harmonic space to a common resolution corresponding to a Gaussian beam of $5'$ FWHM. Each map is then decomposed into a set of needlet coefficients. For each scale j , needlet coefficients of a given map are stored in the form of a single HEALPix map. The filters h_l^j used to compute filtered maps are given by

$$h_l^j = \begin{cases} \cos \left[\left(\frac{\ell_{\text{peak}}^j - \ell}{\ell_{\text{peak}}^j - \ell_{\min}^j} \right) \frac{\pi}{2} \right] & \text{for } \ell_{\min}^j \leq \ell < \ell_{\text{peak}}^j, \\ 1 & \text{for } \ell = \ell_{\text{peak}}^j, \\ \cos \left[\left(\frac{\ell - \ell_{\text{peak}}^j}{\ell_{\max}^j - \ell_{\text{peak}}^j} \right) \frac{\pi}{2} \right] & \text{for } \ell_{\text{peak}}^j < \ell \leq \ell_{\max}^j. \end{cases}$$

For each scale j , the filter has compact support between the multipoles $\ell_{\min}^j$ and $\ell_{\max}^j$ with a peak at ℓ_{peak}^j (see Table B.1 and

Fig. B.1). The needlet coefficients are computed from these filtered maps on HEALPix pixels with N_{side} equal to the smallest power of 2 larger than $\ell_{\text{max}}^j/2$.

Due to the deconvolution of sky maps to an effective smoothing scale of $5'$ FWHM, noise levels in lower-resolution frequency maps are boosted. To limit this effect, we include only those multipoles for which the ratio between the corresponding beam transfer function and that of a corresponding $5'$ Gaussian is larger than 0.01 for each frequency map. This cut is made separately for each needlet scale, such that only those frequency maps containing valid harmonic content that spans the entire bandwidth of a given needlet scale contribute to that particular scale.

In order to improve the measurement of CMB temperature anisotropy near the Galactic plane, we have used a very small preprocessing mask with a sky fraction of 99.8%. The procedure to generate the preprocessing mask is as follows. First we implement the NILC pipeline on the full-mission data sets. Then we identify the pixels where the CMB is more than $500 \mu\text{K}_{\text{CMB}}$, and assign a value of 0 to all the pixels that are within $6'$ and a value of 1 to other pixels. We implement this preprocessing mask on the sky maps in the next run of the NILC pipeline. Prior to implementing the pipeline on the sky maps, the mask regions are filled using the inpainting procedure adopted by the Planck Sky Model; see Planck Collaboration XIII (in prep.) for details.

Estimates of the covariance matrices of needlet coefficients for each scale are computed by smoothing all possible products of needlet coefficients with Gaussian beams. In this way, an estimate of needlet covariances at each point is obtained as a local, weighted average of needlet coefficient products. The FWHMs of the Gaussian windows used for the analysis are chosen to support the computation of statistics; 4225 samples or more samples are averaged. Choosing a smaller FWHM results in excessive error in the covariance estimates, and hence excessive bias. Choosing a larger FWHM results in less localization, and hence some loss of efficiency of the needlet approach.

A patch of angular radius θ and area $2\pi(1 - \cos(\theta))$ contains $N/(4\pi) \times 2\pi(1 - \cos(\theta))$ modes. If we wish to have M modes in that patch, we simply solve for the corresponding θ . We chose $\text{FWHM} = 2 \times \theta$ for the Gaussian beam that we use to smooth the covariance matrix. Hence in order to determine the covariance matrix at a particular point, we have given more weight to those pixels that are close to that point, and less weight to those pixels that are far away. However, this strategy is not optimal for the largest scales.

Figure B.2 shows that the 70, 100, 143, and 217 GHz channels have contributed most to the final reconstruction of the NILC CMB map. However, other channels are also important because these channels are tracers of the foreground signals, and help us to find optimal weights for the final solution.

Calibration errors are a serious issue for precise measurement of the CMB, as they conspire with the ILC filter to cancel out the CMB. This effect is particularly strong in the high signal-to-noise ratio regime, on large scales in particular. We investigate the impact of this calibration bias by redoing the analysis with slightly modified calibration coefficients, and computing the difference between the CMB spectra estimated in the two cases. We adopt the calibration coefficients determined by SMICA in Appendix D. Figure B.3 shows the impact of calibration on the angular power spectrum of the CMB temperature. Comparison of angular power spectra for two data splits shows that the impact is less than 0.5%.




Fig. B.2. Full-sky average of needlet weights for different frequency channels and needlet bands. *From top to bottom panels:* results for temperature, E modes, and B modes.

Appendix C: SEVEM

The SEVEM method (Leach et al. 2008; Fernández-Cobos et al. 2012) produces cleaned CMB maps at different frequencies by




Fig. B.3. Difference of angular power spectra obtained with and without considering the correction to calibration coefficients estimated by SMICA.

subtracting a linear combination of templates constructed internally from the data. In particular, the templates are typically obtained as the subtraction of two close *Planck* frequency channels, filtered to the same resolution to remove the CMB signal. The cleaning is achieved simply by subtracting a linear combination of the templates $t_j(\mathbf{x})$ from the data, with coefficients α_j obtained by minimizing the variance outside a given mask:

$$T_c(\mathbf{x}, \nu) = d(\mathbf{x}, \nu) - \sum_{j=1}^{n_t} \alpha_j t_j(\mathbf{x}). \quad (\text{C.1})$$

Here n_t is the number of templates used, while $T_c(\mathbf{x}, \nu)$ and $d(\mathbf{x}, \nu)$ correspond to the cleaned and raw maps at frequency ν , respectively. The same expression applies for T , Q , or U . Note that we estimate the linear coefficients α_j independently for Q and U maps, following what was done for the previous release¹².

The cleaned frequency maps are then combined in harmonic space, taking into account the noise level, resolution, and (optionally) an estimate of the foreground or systematic residuals of each cleaned channel, to produce a final CMB map at the required resolution.

C.1. Implementation for temperature

For temperature, we have followed the same procedure as for the *Planck* 2015 release (see [Planck Collaboration IX 2016](#) for further details). As before, we clean the 100, 143, and 217 GHz frequency channels with four templates, three of them constructed as the difference between two nearby *Planck* channels (30–44, 44–70, 545–353), such that the first channel is convolved with the beam of the second one and vice versa, and a fourth template given by the 857 GHz channel, convolved with the 545 GHz beam. The cleaned frequency maps have the same resolution as the corresponding original raw data map. To reduce the contamination from point sources in the templates,

we follow the same approach as in the previous release. First, point sources are detected in each frequency map using the Mexican-Hat-Wavelet algorithm ([López-Cañiego et al. 2006](#); [Planck Collaboration XXVI 2016](#)). The upper part of Table C.1 gives the number of point sources detected in intensity and polarization for all the *Planck* frequency channels over the full-sky, at Galactic latitudes $|b| > 20^\circ$. We then inpaint the holes corresponding to the positions of those point sources in the frequency maps (at their original resolution) involved in the construction of the templates. Note that the size of the hole depends on the amplitude of the detected source and the resolution of the considered channel. The filling is done with a simple diffusive inpainting scheme, which replaces one pixel with the mean value of the neighbouring pixels in an iterative way. To avoid possible inconsistencies when performing the subtraction of two maps to construct a template, the diffusive inpainting is performed for all of the sources detected in both channels. For instance, when constructing the (30–44) GHz template, all sources detected at 30 and 44 GHz are inpainted in the two frequency maps before subtraction¹³.

In addition, for this release, we also provide a cleaned CMB map for the 70 GHz channel. This map is constructed at its original resolution and $N_{\text{side}} = 1024$, and has been cleaned with two templates, one constructed as the 30 GHz channel (convolved with the 44 GHz beam) minus the 44 GHz one (convolved with the 30 GHz beam), and a second template obtained as the difference between the 353 and 143 channels, constructed at a resolution equal to that of the 70 GHz channel. This second template has been chosen to trace the emission of the thermal dust, but avoiding, as far as possible, the CO contamination (which is mostly present in the 100 and 217 GHz maps). Point source emission in the templates has also been reduced with the inpainting mechanism already described.

The coefficients of the linear combination used for cleaning the frequency maps are given in Table C.2. They have been calculated by minimizing the variance of the corresponding cleaned map outside a mask that excludes the brightest 1% of the sky and all the point sources detected in intensity. The cleaning procedure introduces a certain level of correlation between the 100, 143, and 217 GHz cleaned frequency maps, due to the use of the same templates, but one frequency map is not used to clean the others. The cleaned 70 GHz channel is, however, more correlated, since it is part of one of the templates used to clean the higher frequency channels. Moreover, the 143 GHz map is also used to clean the 70 GHz channel. Therefore, one should bear in mind these correlations when carrying out analyses with the cleaned single frequency maps. A possible way to reduce the correlations introduced by the cleaning process would be, when possible, to use pairs of cleaned frequency maps constructed with different splits, although this would be at the expense of decreasing the signal-to-noise ratio (e.g., to work with the cleaned 143 GHz even-ring and with the 217 GHz odd-ring maps, since the templates are constructed with the corresponding split).

Following the same approach as in the previous release, after the frequency maps are cleaned, they are inpainted, in a first step, at the positions of the point sources identified in the corresponding raw maps. In a second step, the Mexican-Hat-Wavelet algorithm is again run on the cleaned frequency maps, and the newly detected sources (see lower part of Table C.1) are further

¹² In principle, it would be desirable to estimate the linear coefficients taking into account the spinorial character of the Q and U components, since this allows us to keep the physical coherence of the foreground residuals, following, for instance, the method proposed by [Fernández-Cobos et al. \(2016\)](#). However, in practice, this does not seem to have any significant impact on the results from *Planck* data and, therefore, for simplicity, we work with independent coefficients for Q and U maps.

¹³ Note that if a map is used to construct more than one template, several inpainted versions of that map will be constructed in the appropriate way in order to match the pair.

Table C.1. Number of detected sources in intensity and polarization

Map	30 GHz	44 GHz	70 GHz	100 GHz	143 GHz	217 GHz	353 GHz	545 GHz	857 GHz
<i>Raw</i>									
<i>T</i> (full-sky)	1593	923	1307	2162	3479	4955	5794	8145	11876
<i>T</i> ($ b > 20^\circ$)	977	470	648	809	1093	1289	1588	2898	6117
<i>P</i> (full-sky)	195	64	74	237	349	632	1075		
<i>P</i> ($ b > 20^\circ$)	65	19	15	56	63	87	134		
<i>Cleaned</i>									
<i>T</i> (full-sky)	...	...	420	1475	2117	3675	...	...	...
<i>T</i> ($ b > 20^\circ$)	...	...	37	93	230	553	...	...	...
<i>P</i> (full-sky)	...	...	10	48	73	199	...	...	...
<i>P</i> ($ b > 20^\circ$)	...	...	1	1	4	16	...	...	...

Notes. The upper part of the table refers to sources detected in the raw frequency maps, while the lower part gives the number of point sources detected in the cleaned SEVEM frequency maps after inpainting the originally detected sources. A list with the positions of all the sources and the corresponding masks used in the SEVEM pipeline are available in the *Planck* Legacy Archive.

Table C.2. Linear coefficients, α_j , of the templates used to clean individual frequency maps with SEVEM for temperature.

Template	Coefficients α_j			
	70GHz	100GHz	143GHz	217GHz
30–44	1.68×10^{-1}	-9.15×10^{-2}	3.47×10^{-3}	-1.57×10^{-1}
44–70	...	4.19×10^{-1}	1.97×10^{-1}	4.22×10^{-1}
353–143	6.68×10^{-3}	...	...	...
545–353	...	4.21×10^{-3}	6.32×10^{-3}	1.72×10^{-2}
857	...	-3.23×10^{-5}	-5.04×10^{-5}	-1.04×10^{-4}

Notes. The 353–143 template has been produced at the same resolution as the 70 GHz frequency channel, while the 857 GHz map has been convolved with the 545 GHz beam. The rest of the templates are constructed such that the first map in the subtraction is convolved with the beam of the second map and vice versa.

inpainting. The combined area inpainted outside the SEVEM confidence mask for the 143 and 217 GHz channels (those used to construct the final CMB map) corresponds to around 0.4% of the sky, while it is fully covered by the common confidence mask. Note that the same inpainting strategy is applied to the simulations processed through the SEVEM pipeline, to make sure that any possible effect introduced by this procedure is statistically taken into account. Finally, the monopole and dipole are removed from the full-sky cleaned maps (note that this is different from the previous release, where monopole and dipole were removed outside the SEVEM confidence mask). The cleaned intensity maps for the 70, 100, 143, and 217 GHz channels are shown in Fig. C.1.

The final SEVEM CMB map is constructed by combining the cleaned 143 and 217 GHz maps in harmonic space¹⁴. In particular, the maps are weighted at each multipole, taking into account the noise level and resolution of the maps, as well as a rough estimation of the expected foreground residuals. This estimation has been updated with respect to the previous release by using the FFP8 simulations. The total weights are shown in Fig. C.2. The resolution of the combined map corresponds to a Gaussian beam of 5' FWHM and HEALPix resolution $N_{\text{side}} = 2048$, with

maximum multipole $\ell_{\text{max}} = 4000$. A monopole and a dipole are also removed from the full-sky map.

C.2. Implementation for polarization

A similar procedure is applied independently to the frequency maps of the Stokes Q and U parameters to obtain cleaned polarization CMB maps, which are afterwards combined in harmonic space to produce the final SEVEM maps. Given the narrower frequency coverage available for polarization, a different choice of templates needs to be defined in this case. In the previous release, only two cleaned channels (100 and 143 GHz) were combined to produce the final polarization map. However, due to the significant improvement of the *Planck* data in polarization for the current release, we are now able to clean the 217 GHz channel and to include this map in the final combination. This produces a significant improvement in the signal-to-noise ratio of the cleaned SEVEM CMB polarization maps with respect to the previous version. In addition, in the updated pipeline, the cleaned maps are produced at full resolution ($N_{\text{side}} = 2048$ instead of $N_{\text{side}} = 1024$ for the 100, 143, and 217 GHz channels, as well as the combined map). As for the previous release, a cleaned 70 GHz map is also provided at its native resolution. To reduce point source contamination, inpainting similar to that in the previous release is performed.

The first step of the pipeline is to inpaint the positions of the sources detected in those channels that will be used to construct templates. These point sources are detected using a non-blind approach, among the intensity candidates, using the

¹⁴ In principle one could also include the cleaned 70 and 100 GHz maps in the combined solution. However, given the lower resolution and higher noise level of these channels, the improvement in the signal-to-noise ratio of the combined map is modest. Taking into account also that the addition of these channels could potentially introduce contamination from low-frequency foregrounds or CO emission, we decided to combine only the 143 and 217 GHz cleaned channels in the final map.




Fig. C.1. Cleaned single-frequency CMB maps from the SEVEM pipeline. The cleaned maps in intensity (*left column*) are given at their original resolution, while the polarization maps (*Q* and *U*, *middle* and *right columns*) have been smoothed with a Gaussian beam of 80' FWHM resolution for better visualization. Rows show results for different frequencies (70, 100, 143, and 217 GHz).

filtered fusion technique (Argüeso et al. 2009). The upper part of Table C.1 shows the number of sources detected in polarization in each of the frequency channels. The size of the holes to be inpainted takes into account both the amplitude of the source and the beam of the channel. As in the intensity case, when performing the subtraction of two maps to construct a template, the diffuse inpainting is performed for all of the sources detected in both channels. Note that the inpainting is always done at the native resolution of the channel and independently for *Q* and *U* maps.

Once the maps have been inpainted, each template is constructed as the subtraction of two frequency channels processed to a common resolution. Given the smaller number of channels in polarization, the maps to be cleaned are also used to construct templates. In this sense, the cleaned maps at different frequencies are, in general, less independent than in the intensity case (the exception is the 70 GHz channel, which is not used as part of the templates for polarization). Six templates (one of them at two

different resolutions) are generated to produce cleaned maps at 70, 100, 143, and 217 GHz. In particular, to trace the synchrotron emission, the (30–44) GHz template is constructed, where the 30 GHz map is smoothed with the 44 GHz beam and vice versa. To trace the thermal dust, templates are produced at (217–143), (217–100), and (143–100) with 1° resolution (this smoothing is included in order to increase the signal-to-noise ratio of the template), and at (353–217) and (353–143) with 10' resolution. In addition, this last template is also constructed at the resolution of the 70 GHz beam, in order to clean that channel. The produced templates are then subtracted from the (non-inpainted) raw data at their native resolution. Table C.3 shows the list of templates used to clean each map, as well as the corresponding coefficients of the linear combination. These coefficients have been obtained by minimizing the variance of each cleaned map outside a mask excluding the brightest 3% of the sky and all the point sources detected in polarization. Note that to clean the 100 and 143 GHz maps, the same combination of templates as in the




Fig. C.2. Weights used to combine the cleaned single-frequency maps into the final SEVEM CMB map for temperature, corresponding to 143 GHz (blue line) and to 217 GHz (green line). The weights do not sum to unity because they include the effect of deconvolving the beams of the frequency maps and convolving with the 5' Gaussian beam of the final map.




Fig. C.3. Weights used to combine the cleaned single-frequency maps into the final SEVEM CMB maps for polarization. The different lines correspond to 100 (red), 143 (blue), and 217 GHz (green) channels. The weights do not sum to unity because they include the effect of deconvolution by the beams of the frequency maps and convolving with the 5' Gaussian beam of the final map.

previous release is used, although now templates and cleaned maps are produced at $N_{\text{side}} = 2048$.

Once the frequency maps are cleaned, inpainting at the position of the point sources detected at each of those channels is carried out. Moreover, once these cleaned inpainted maps are ready, the non-blind point source detection algorithm is run again on them and additional point sources detected (see lower part of Table C.1). These newly identified sources are also inpainted. This second iteration of the algorithm was performed for intensity in the previous release but not for polarization; in this version it is done for both cases. The joint area inpainted outside the SEVEM mask in the three cleaned channels used to produce the combined maps corresponds to around a 0.04% of the sky, and is fully covered by the common confidence mask. As for intensity, exactly the same inpainting procedure is applied to the simulations processed through the SEVEM pipeline, to account for any possible effects introduced by this step. The cleaned Q and U maps for the 70, 100, 143, and 217 GHz channels are shown in Fig. C.1. The maps have been smoothed with a Gaussian beam with 80' FWHM resolution to allow for better visualization.

The last step is to combine the cleaned single-frequency maps in order to produce the final Q and U cleaned CMB maps. This is done by combining in harmonic space the cleaned 100, 143, and 217 GHz maps. The weights take into account the noise of each channel and its resolution. In addition, recognizing the fact that the 217 GHz channel is likely to be somewhat more susceptible to large-scale systematic residuals than the other two channels, we also introduce a relative down-weighting of the 217 GHz channel on the largest scales. This can be seen in Fig. C.3, where the harmonic weights are given for the 100 (red), 143 (blue), and 217 GHz (green) channels. The same weights are applied for E and B . The resolution of the combined map corresponds to a Gaussian beam of FWHM 5' and HEALPix resolution $N_{\text{side}} = 2048$, with a maximum multipole $\ell_{\text{max}} = 3000$. We consider a lower ℓ_{max} for polarization than for intensity due to the lower signal-to-noise ratio of the polarization data.



C.3. Masks

In temperature, the SEVEM confidence mask is generated following a similar procedure to that of the previous release. Specifically, we define the mask by thresholding maps constructed as the difference between two different CMB reconstructions. As in 2015, we construct these differences at $N_{\text{side}} = 256$, with resolution given by a Gaussian beam with $\text{FWHM} = 30'$. In particular, three combinations are considered: the cleaned (217–143) GHz and (143–100) GHz maps and the difference between two cleaned, combined CMB maps, whose linear coefficients have been obtained by minimizing the variance outside two different masks. From each of these maps, one mask is constructed by removing the brightest pixels (and its direct neighbours) down to a certain threshold. The three masks are multiplied to produce the final confidence mask, which is then smoothed with a Gaussian beam of 1° to avoid sharp edges and upgraded to full resolution. The thresholds that define the masks are chosen by looking at the amplitude of the extrema and the dispersion of the cleaned 100, 143, and 217 GHz channels and the combined map after applying the considered mask, trying to find a compromise between reducing the values of these quantities while keeping a reasonable sky fraction. In particular, thresholds removing between 8 and 10% of the sky were found to be adequate for the differences considered. In addition, a small region near the Galactic plane with a relatively high contamination, but that was not captured with these values of the thresholds, was manually masked by applying a circle of 0.3 radius. This removed around 350 additional pixels, without modification of the thresholds, which would lead to a larger reduction of the area allowed by the mask. The final SEVEM confidence mask in intensity leaves a suitable sky fraction of 83.8%, and is shown in the top panel of Fig. C.4.

In polarization, given the lower signal-to-noise ratio of the reconstructed CMB maps, a different approach from that of intensity needs to be considered to identify the reliable regions of

Table C.3. Linear coefficients α_j for each of the templates used to clean individual frequency maps with SEVEM for polarization.

Template	Coefficients α_j							
	70GHz Q	70GHz U	100GHz Q	100GHz U	143GHz Q	143GHz U	217GHz Q	217GHz U
30–44	2.72×10^{-2}	3.53×10^{-2}	0.96×10^{-2}	1.29×10^{-2}	3.43×10^{-3}	6.81×10^{-3}	1.21×10^{-2}	1.79×10^{-2}
143–100	...	...	...	...	...	...	8.63×10^{-1}	7.19×10^{-1}
217–100	...	...	...	...	1.52×10^{-1}	1.47×10^{-1}	...	...
217–143	...	...	9.38×10^{-2}	8.27×10^{-2}	...	...	...	...
353–143	1.13×10^{-2}	0.98×10^{-2}	...	...	...	...	1.17×10^{-1}	1.13×10^{-1}
353–217	...	...	1.32×10^{-2}	1.30×10^{-2}	2.83×10^{-2}	2.72×10^{-2}	...	...

**Fig. C.4.** SEVEM masks in temperature (*top panel*) and polarization (*bottom panel*).**Fig. C.5.** SEVEM masks in temperature (*top panel*) and polarization (*bottom panel*) for inpainted point sources.

the sky. Several aspects of the approach to construct the polarization confidence mask have been modified with respect to the previous release, and the new method is described here in detail. In particular, we have defined the confidence mask as the product of two individual masks: one specific mask based on the achieved CMB reconstruction, and a second one customized to avoid the regions more contaminated by thermal dust.

For the specific mask, the first step is to downgrade the CMB reconstructed maps (Q and U) to a resolution equivalent to a Gaussian beam with $FWHM = 90'$ and $N_{\text{side}} = 128$. From these maps, we estimate locally the rms of P (i.e., $\sqrt{Q^2 + U^2}$) at each position by calculating the rms of the pixels included in a circle with a given radius centred on the considered pixel. We then estimate the expected rms of P for a map containing only CMB and noise. For the noise, this is obtained by estimating this quantity locally for the odd-even half-difference map, processed through the SEVEM pipeline, at the resolution being considered, using the same procedure as for the cleaned maps. For the CMB, we simply obtained the rms of P , averaging over simulations. Since the CMB and noise are independent, their rms values are added quadratically. The ratio between the rms of the cleaned maps over that expected for a CMB-plus-noise map is then constructed. Pixels with larger ratios are expected to be more con-

taminated; the specific mask is defined by those pixels above a given threshold. This mask is then smoothed with a Gaussian beam of $FWHM = 90'$ to avoid sharp boundaries, and upgraded to $N_{\text{side}} = 2048$. We explored several values for the radius of the circle (to locally estimate the rms) and for the amplitude of the threshold, finding that a value of 15 pixels (at $N_{\text{side}} = 128$) for the radius and a threshold of 1.5 produced good results.

To construct the dust mask, we use the raw 353 GHz channel, smoothed at a resolution of $90'$ and $N_{\text{side}} = 128$. The rms of P is obtained at each pixel as explained above, and a fixed fraction of pixels with the largest rms values is included in the mask. This mask is again smoothed with a Gaussian beam of $90'$ and upgraded to $N_{\text{side}} = 2048$. To construct this mask, we have chosen a radius for estimating the rms of four pixels and excluded 15% of the sky. Finally, the specific mask and the dust mask are multiplied together, passing 80.3% of the sky. The SEVEM confidence mask in polarization is shown in the bottom panel of Fig. C.4.

We conclude with some additional comments about the best way to deal with inpainted pixels. The most conservative approach is to explicitly exclude all of the inpainted areas from the analysis. This implies the inclusion of a large number of holes in the confidence mask, which can be damaging

for certain analyses, especially those performed in harmonic space. The diffusive inpainting strategy considered above seems to effectively reduce the emission from detected point sources while, at the same time, not introducing evident artefacts in the cleaned maps (recall that we are filling small holes, corresponding to scales where the background is usually smooth). Therefore, we have only masked those inpainted pixels which are directly excluded by the general algorithm used to construct the confidence mask. For intensity, this leaves a joint inpainted area outside the SEVEM mask in the two cleaned channels (143 and 217 GHz) used to construct the final CMB map of around 0.4% of the sky. For polarization, the corresponding joint area (from the cleaned 100, 143, and 217 GHz channels) also covers around 0.04% of the sky. Moreover, for both intensity and polarization, the exact same procedure is applied to the simulations processed through the SEVEM pipeline, to ensure that any unexpected spurious effects are statistically taken into account. We believe that this is a good way to proceed in order to find a compromise between reducing point-source contamination in the cleaned maps and providing a well-behaved confidence mask for CMB analysis. This is the same approach used in the previous release. Nonetheless, for certain types of analysis, as for example the local study of compact objects, it may be necessary to discard, or at least to be aware of, the inpainted regions. For these cases, we also provide masks of the pixels inpainted in each of the cleaned frequencies, as well as the joint mask for those channels that are used to construct the final CMB maps. The joint masks of inpainted pixels are given in Fig. C.5 for intensity (top) and polarization (bottom). Note that additional inpainting is also performed during the template construction, but those positions are not included in these masks since those pixels are not directly inpainted on the cleaned maps. Finally, we point out that the masks for inpainted point sources given in Fig. C.5 have been included in the confidence common masks (Fig. 9), to reduce possible point source contamination in all the CMB maps. If it is desired to carry out an analysis of the SEVEM CMB maps without explicitly including point source holes in the mask, the SEVEM confidence masks should be considered.

Appendix D: Spectral Matching Independent Component Analysis (SMICA)

The general operation of SMICA (Spectral Matching Independent Component Analysis; Delabrouille et al. 2003; Cardoso et al. 2008) and the main changes with respect to the 2015 release are summarized in Sect. 2.4. In this appendix, we provide additional implementation details. There are several masking and pre-processing operations whose specifics vary depending on the target map (CMB or foregrounds, temperature or polarization), but the general methodology is the same, following these steps:

- (1) Preprocessing of the input maps by point source subtraction and masking/inpainting. This step also includes additional masking (Galactic plane, etc.).
- (2) Estimation of the spectral statistics $\widehat{C}_\ell$ via Eq. (3) from the spherical harmonic coefficients computed from the pre-processed maps, possibly with some additional masking to remove particularly bright objects.
- (3) Fitting of a SMICA model to beam-corrected $\widehat{C}_\ell$, from which the SMICA harmonic weights w_ℓ are computed.
- (4) Computation of the spherical harmonic coefficients from the preprocessed maps and linear combination as per Eq. (2), to synthesize a map with a specified effective Gaussian beam.
- (5) Determination of a “confidence mask”.

The specifics of the production of each SMICA map are given below.

D.1. Temperature analysis

The SMICA 2018 temperature map is a hybrid of two complementary CMB renderings, namely X_{high} , which includes only HFI observations, and is specialized for high Galactic latitudes, and intermediate and small angular scales, and X_{full} , which includes all *Planck* channels, and provides us with additional content. The final SMICA temperature map is then constructed as a weighted sum of these two maps, following Eq. (6). The two sky areas to be hybridized are defined by a smooth mask shown at Fig. D.2. In polarization, we do not resort to such a hybrid scheme.

Recalibration. As in previous releases, a preliminary SMICA fit (calibration run) is conducted, with calibration coefficients left unconstrained at 100 and 217 GHz. This fit involves only HFI channels, is limited to the first peak ($30 \leq \ell \leq 300$), and involves spectral matrices estimated over a clean part of the sky. It yields relative calibration coefficients 1.0004 at 100 GHz and 1.0005 at 217 GHz. These values are consistent with the results reported in Planck Collaboration III (2020) and Planck Collaboration V (2020).

Preprocessing. The input maps are preprocessed for point sources as follows. In the maps from 30 GHz to 353 GHz, we try to fit and subtract the strongest point sources detected at the 5σ level in the PCCS2 catalogue (Planck Collaboration XXVI 2016). Any point source with an unsatisfactory fit is left “as is” in the map. In a second step, in each of the input maps from 44 GHz to 353 GHz, we mask all the point sources detected at more than 50σ (unless they have already been subtracted in the previous step). The masked areas at all frequencies are then combined to form a common point-source mask. In addition to that point-source mask, we include a small mask, hereafter “the Galactic mask”, blocking the Galactic plane, plus a small number of selected regions (such as the LMC). The resulting “preprocessing mask” is shown in Fig. D.1. In order to minimize leakage in the subsequent computation of spherical harmonic coefficients, the masked areas under this common mask are filled in by a simple diffusive inpainting procedure.

Spectral statistics. The computation of the spherical harmonic coefficients entering in the spectral statistics $\widehat{C}_\ell$ differs between X_{high} and X_{full} . For X_{high} , we apply an apodized version of the transition mask, while for X_{full} , we use the full sky. In both cases, we use the preprocessed maps with additional masking of bright objects or regions. For X_{full} , which involves all *Planck* frequency channels, the point source mask is augmented with all the sources detected at more than 50σ at frequencies 30 GHz, 545 GHz, and 857 GHz, and the new holes are again filled in by diffusive inpainting. We also mask part of Galactic region using an apodized version of the Galactic mask. For X_{high} , which involves only HFI channels, we mask all the point sources detected at 5σ at frequencies 100 GHz, 143 GHz, and 217 GHz, even if already subtracted. The resulting holes are then apodized over $30'$.

Spectral fits. For producing the X_{high} map, SMICA processing is conducted, fitting the spectral covariance matrices $\widehat{C}_\ell$ over the multipole range $25 \leq \ell \leq 1000$. For this fit, the calibration is kept fixed at the values found in the calibration run. The free parameters are the (binned) CMB spectrum C_ℓ^{cmb} , the positive matrices P_ℓ , and the $6 \times N_{\text{fg}}$ foreground emissivity matrix F .




Fig. D.1. SMICA pre-processing masks. *Left panel:* for intensity analysis, covering $f_{\text{sky}} = 98.5\%$. *Right panel:* for polarization analysis, covering $f_{\text{sky}} = 97.3\%$.




Fig. D.2. SMICA transition mask used to combine the X_{high} and the X_{full} CMB renderings.

For producing the X_{full} map, a first run is devoted to estimating the foreground emissivity matrix F , and a recalibration factor for the 70 GHz channel (this factor is found to be 1.0019). This fit is conducted over the multipole range $2 \leq \ell \leq 150$. In a second run, we fit (binned versions of) C_{ℓ}^{cmb} and P_{ℓ} over the multipole range $10 \leq \ell \leq 1000$, keeping fixed the calibration (vector $\mathbf{a}$) and the foreground emissivity matrix F .

Map synthesis. The SMICA fits produce parametric estimates of C_{ℓ} , from which spectral weights $\mathbf{w}_{\ell}$ are readily obtained. They are shown in Fig. 1 for X_{high} (top panel) and X_{full} (middle panel). Those weights are applied to spherical harmonic coefficients computed from the preprocessed input maps. The spatial transition weights used to hybridize X_{high} and X_{full} are shown in Fig. D.2.

Confidence mask. The confidence mask combines a point source mask and a Galactic mask determined by a procedure similar to the one used for the 2015 release. It is documented in the Explanatory Supplement.

Inpainting. Final inpainting of the CMB maps is no longer performed in the SMICA pipeline, but is carried out through a procedure common to all methods, as described in Sect. 4.2.

SZ-free CMB map. A CMB map free of SZ contamination is produced by a simple adaptation of Eq. (2) as follows. That expression yields weights $\mathbf{w}_{\ell}$, which, at each multipole ℓ , minimize the output power while enforcing unit gain towards the CMB signal. In other words, it is the minimizer of $\mathbf{w}_{\ell}^{\dagger} C_{\ell} \mathbf{w}_{\ell}$ subject to $\mathbf{w}_{\ell}^{\dagger} \mathbf{a} = 1$. One can solve the same problem with the additional constraint that the weights should also cancel the SZ

signal, that is, enforcing the additional constraint $\mathbf{w}_{\ell}^{\dagger} \mathbf{b} = 0$ where $\mathbf{b}$ denotes the SZ emission law. The minimizer of $\mathbf{w}_{\ell}^{\dagger} C_{\ell} \mathbf{w}_{\ell}$ subject to $\mathbf{w}_{\ell}^{\dagger} \mathbf{a} = 1$ and $\mathbf{w}_{\ell}^{\dagger} \mathbf{b} = 0$ is easily found in closed form (see Remazeilles et al. 2011b) as

$$\mathbf{w}_{\ell} = C_{\ell}^{-1} G (G^{\dagger} C_{\ell}^{-1} G)^{-1} \mathbf{c} \quad (\text{D.1})$$

where $G = [\mathbf{a} \ \mathbf{b}]$ and $\mathbf{c} = [1 \ 0]^{\dagger}$.

Figure D.3 shows an enlargement of the difference between the SMICA CMB maps derived with and without SZ projection. Figure D.4 compares the angular power spectra of these two maps. Both versions of the SMICA CMB maps are considered in the lensing study (Planck Collaboration VIII 2020).

Changes with respect to the 2015 release. Figure 6 shows, for all pipelines, the differences in CMB temperature maps from 2015 to 2018. In the SMICA case, the difference could have three origins: changes in the input data, changes in the SMICA pipeline, and changes in recalibration procedure. We show here that the difference is mostly due to recalibration by producing a CMB map, referred to as the “2015b map”, obtained from the 2015 data by running the 2015 pipeline with the sole exception that, as for the 2018 release, the 30 GHz and 44 GHz channels are *not* recalibrated. Figure D.5 shows the differences from the 2015 map to this 2015b map (top panel) and from this 2015b map to the 2018 map, while the bottom panel of Fig. 6 shows the difference from the 2015 map to the 2018 map. These three pairwise comparisons make it clear that, in temperature, most of the differences between 2015 and 2018 should be attributed to changes in calibration, rather than to changes in the SMICA pipeline.

D.2. Polarization analysis

CMB reconstruction. We now turn to the construction of SMICA polarization maps, and start with the CMB map. First, a significant modification to the SMICA 2018 pipeline is the fact that E and B modes are now processed independently; in contrast, the 2015 analysis fitted and filtered these modes jointly. SMICA uses all seven *Planck* polarized channels. When producing either E -mode or B -mode CMB maps, the foreground emission is taken to have maximal dimension: $N_{\text{fg}} = 7 - 1 = 6$.

The input maps are preprocessed as follows. First, in each of the input frequency maps, all point sources detected at the 5σ level are masked and the holes are filled by diffusive inpainting. Second, the bright pixels (with amplitude ten times larger than the standard deviation of the map) are similarly masked and inpainted. Finally, a small Galactic mask – obtained by thresholding a combination of the 30 GHz and 353 GHz maps – is applied. The resulting mask, shown in Fig. D.1, covers 97% of




Fig. D.3. Difference between the SMICA CMB map and its SZ-free version. The patch shown is $20^\circ \times 20^\circ$ centered on $(l, b) = (46.3, 53^\circ)$.




Fig. D.4. Angular spectra for the CMB (blue lines), the CMB SZ-free version (green lines), and their difference spectra (red lines), computed on the SMICA confidence mask. Half-mission cross-spectra (solid line) and half-mission difference spectra (dashed line) are shown to assess the signal and noise differences between the two CMB maps.

the sky. In the 2015 release, the same processing mask was used for polarization and intensity.

As for temperature, we proceed in two steps. A first SMICA fit is performed to estimate the foreground emissivity matrix F over the range $5 \leq \ell \leq 150$. A second SMICA fit is then performed in the range $2 \leq \ell \leq 1000$ over parameters C_ℓ^{cmb} and P_ℓ , while F is kept fixed at the value found in the first run. The right panel of Fig. 1 shows the resulting harmonic weights. These result from spectral statistics $\hat{C}_\ell$ computed from the preprocessed maps without additional masking, unlike in the temperature case.




Fig. D.5. CMB difference maps in temperature at $80'$ resolution. *Top panel:* difference between the SMICA 2015 released map and the 2015b map (without recalibration of the 44 GHz channel). *Bottom panel:* difference between the 2015b and the 2018 map.

Confidence mask. A SMICA polarization confidence mask has been produced and released, but appears not to be conservative enough. For that reason, we recommend using the common confidence mask to analyse SMICA polarized CMB maps.

D.3. Polarized foreground reconstruction.

The results presented in Sect. 5.3, regarding the polarized dust and synchrotron emission, are based on a dedicated, blind SMICA fit with a foreground emissivity matrix F composed only of $N_{\text{fg}} = 2$ columns. The total foreground contribution to a spectral covariance matrix C_ℓ being $FP_\ell F^\dagger$, a blind fit can only determine the factors F and P_ℓ up to multiplication by an invertible 2×2 matrix T . Indeed, for any such matrix T , one can define $\tilde{P}_\ell = TP_\ell T^\dagger$ and $\tilde{F} = FT^{-1}$ and see that the transformed pair $(\tilde{F}, \tilde{P}_\ell)$ contributes as much as the original pair (F, P_ℓ) to the spectral covariance matrix, since, by construction $FP_\ell F^\dagger = \tilde{F}\tilde{P}_\ell \tilde{F}^\dagger$. Therefore the likelihood is insensitive to the value of T . Since a blind fit is (by definition) conducted without constraining either F nor P_ℓ , the matrix T cannot be determined from the data without imposing extra constraints. This degeneracy could be fixed by constraining P_ℓ to be diagonal, but this would be equivalent to fitting a (wrong) model of uncorrelated synchrotron and dust emissions. We choose instead to fix the degeneracy as follows. We conduct a blind SMICA fit and, in a post processing step, we select (without affecting the quality of the SMICA fit) a matrix T given by

$$T = \begin{bmatrix} F_{30,1} & F_{30,2} \\ F_{353,1} & F_{353,2} \end{bmatrix},$$

so that the first row of $\tilde{\mathbf{F}} = \mathbf{F}\mathbf{T}^{-1}$ becomes $[0, 1]$ and its last row becomes $[1, 0]$. In other words, we fix the indeterminacy in the blind fit of a two-template foreground model by assuming that the entire foreground signal at 30 GHz is only synchrotron, and that the entire foreground signal at 353 GHz is only thermal dust. We checked that performing a second fit where the synchrotron contribution at 353 GHz is not zero but an extrapolated value (and similarly for dust at 30 GHz), has no significant effect on fitted values and, unsurprisingly, that it does not affect either of the reconstructed maps.

Maps of polarized dust and synchrotron emission are synthesized from harmonic coefficients computed over the full sky, except for point sources detected at 5σ , which are masked and inpainted. This is carried out independently for each input map. The Q and U maps are synthesized with an effective Gaussian beam of 3° (FWHM) for synchrotron and $12'$ for dust. The SEDs of dust and synchrotron emission shown in Fig. 32 are determined from a dedicated SMICA fit based on spectral covariance matrices computed from about 70% of the sky.

Appendix E: GNILC

The formalism of GNILC has been described in detail in [Remazeilles et al. \(2011a\)](#) and [Planck Collaboration Int. XLVIII \(2016\)](#). The main characteristics can be summarized as follows: (i) a GNILC map at given frequency is a weighted linear combination (ILC) of the *Planck* frequency maps having minimum variance; (ii) GNILC performs localized analysis in both harmonic space and pixel space via needlet (spherical wavelet) decomposition ([Narcowich et al. 2006](#)), and as such it adapts component separation to local conditions of contamination both over the sky and over angular scale; and (iii) GNILC uses not only spectral information, but also spatial information (angular power spectra) of the non-Galactic components (CIB, CMB, and noise) in order to disentangle the Galactic signal from the CIB, CMB, and noise contamination. Therefore, GNILC is a blind, model-independent, data-driven component-separation method, in the sense that there is no prior assumption/parametrization of the Galactic foreground properties.

There are, however, a few differences in the GNILC processing steps between intensity and polarization. For intensity, the processing is identical to that of [Planck Collaboration Int. XLVIII \(2016\)](#), i.e., the prior information is both spectral and spatial, and consists of the *Planck* best-fit CMB temperature power spectrum, $C_\ell^{\Lambda\text{CDM}}$ ([Planck Collaboration XI 2016](#)), the *Planck* CIB best-fit auto/cross power spectra across frequency pairs, $C_\ell^{\text{CIB}}(\nu_1, \nu_2)$ ([Planck Collaboration XXX 2014](#)), and the *Planck* noise power spectra across frequencies, $C_\ell^{\text{noise}}(\nu)$. For polarization, prior information is only spectral for the CMB, consisting of the CMB SED¹⁵, while the noise prior is still spectral and spatial, comprising the *Planck* noise power spectra at each frequency. In practice, *Planck* noise power spectra are derived from the half-difference of the first and second halves of each stable pointing period (“rings”) of *Planck*, in which the sky emission cancels out and leaves an estimate of the full-survey noise.

From those prior power spectra, we simulate Gaussian realizations of the CMB map, $y_\nu^{\text{CMB}}(p)$, the correlated CIB maps, $y_\nu^{\text{CIB}}(p)$, and the noise maps, $y_\nu^{\text{noise}}(p)$, where ν denotes the frequencies and p the pixels. The simulated total “nuisance” map is

defined as

$$y_\nu(p) \equiv g_\nu y_\nu^{\text{CMB}}(p) + y_\nu^{\text{CIB}}(p) + y_\nu^{\text{noise}}(p) \quad (\text{E.1})$$

for intensity, where g_ν is the derivative of a blackbody with respect to temperature, and

$$y_\nu(p) \equiv y_\nu^{\text{noise}}(p) \quad (\text{E.2})$$

for polarization, since the CIB is assumed to be unpolarized and we have no spatial information on the CMB polarization.

We perform a needlet decomposition of both the simulated nuisance maps and the *Planck* frequency maps. We thus define ten needlet windows, $\{h_\ell^{(j)}\}_{1 \leq j \leq 10}$, as Gaussian bandpass filters in harmonic space to perform component separation on different ranges of multipoles independently¹⁶. The spherical harmonic coefficients of the simulated maps, $y_\nu(p)$, are bandpass-filtered as $h_\ell^{(j)} a_{\ell m}(\nu)$. The inverse spherical harmonic transform of the filtered coefficients produces ten needlet maps, $y_\nu^{(j)}(p)$ (one for each needlet scale), for each frequency. Each needlet map, $y_\nu^{(j)}(p)$, contains temperature fluctuations at the specific range of angular scales probed by the associated needlet window, with statistical properties determined by the prior power spectra at these scales.

For each needlet scale (j), we compute the covariance matrix of the nuisance map (noise for polarization; CMB plus CIB plus noise for intensity) in each pixel p for all pairs of frequencies a and b as:

$$[\mathbf{R}_n^{(j)}(p)]_{ab} = \sum_{p' \in \mathcal{D}^{(j)}(p)} y_a^{(j)}(p') y_b^{(j)}(p'), \quad (\text{E.3})$$

where in practice the pixel domain, $\mathcal{D}^{(j)}(p)$, is defined by the convolution in real space of the product of needlet maps, $y_a^{(j)}(p) y_b^{(j)}(p)$, with a Gaussian kernel whose the width is a function of the needlet scale considered. Note that the prior covariance matrix of the nuisance map, $\mathbf{R}_n(p)$, is blind about the particular realization of CMB, CIB, and noise that is found in the observed *Planck* data.

Similarly, the data (*Planck* frequency maps), $d_\nu(p)$, are decomposed onto the same needlet frame, and the frequency-frequency covariance matrix of the data is computed in each pixel for each needlet scale as:

$$[\widehat{\mathbf{R}}_d^{(j)}(p)]_{ab} = \sum_{p' \in \mathcal{D}^{(j)}(p)} d_a^{(j)}(p') d_b^{(j)}(p'). \quad (\text{E.4})$$

As described in [Planck Collaboration Int. XLVIII \(2016\)](#), the prior power spectra are thus used to obtain a model of the frequency-frequency covariance matrix, $\mathbf{R}_n$, of the nuisance contribution (CIB, CMB, and noise) to the total data covariance matrix, $\widehat{\mathbf{R}}_d$ (9×9 matrices for intensity, 7×7 for polarization). The signal-to-noise ratio, where signal stands for Galactic emission, is obtained via the matrix $\mathbf{R}_n^{-1/2} \widehat{\mathbf{R}}_d \mathbf{R}_n^{-1/2}$, which is estimated locally over the sky and over different ranges of angular scales via needlet decomposition of the maps. The eigenstructure of the matrix $\mathbf{R}_n^{-1/2} \widehat{\mathbf{R}}_d \mathbf{R}_n^{-1/2}$ allows us to discriminate those eigenvalues that are close to unity (therefore corresponding to nuisance) from

¹⁵ Given that the amplitude of the CMB B -mode power spectrum is unknown, we cannot use spatial information on the CMB as a prior when performing GNILC on polarization data.

¹⁶ The needlet windows satisfy the relation $\sum_{j=1}^{10} (h_\ell^{(j)})^2 = 1$ to ensure the conservation of the total power when synthesizing all the needlet maps to reconstruct the complete map.

those that correspond to the contribution of Galactic emission¹⁷. This allows us to estimate the local dimension, m , of the Galactic signal subspace over the sky and over scales, i.e., the finite number of independent (not physical) components¹⁸ onto which the correlated Galactic emission can be decomposed. We note $\mathbf{U}_S$ the matrix collecting the selected subset of m eigenvectors of the matrix $\mathbf{R}_n^{-1/2} \widehat{\mathbf{R}}_d \mathbf{R}_n^{-1/2}$ that form an orthogonal basis of the Galactic signal subspace.

The data¹⁹ are then projected onto the identified Galactic signal subspace, and an m -dimensional ILC is performed on the projected data in order to further minimize any part of the nuisance that did not project orthogonally to the Galactic subspace:

$$\widehat{s}_v^{\text{dust}(j)}(p) = \sum_{v'} \mathbf{W}_{vv'}^{(j)}(p) d_{v'}^{(j)}(p). \quad (\text{E.5})$$

The matrix of GNILC weights can be written in compact form as (Remazeilles et al. 2011a):

$$\mathbf{W} = \mathbf{F}^t \widehat{\mathbf{R}}_d^{-1} \mathbf{F}^{-1} \mathbf{F}^t \widehat{\mathbf{R}}_d^{-1}, \quad (\text{E.6})$$

with the estimated foreground mixing matrix, $\mathbf{F}$, given by

$$\mathbf{F} = \mathbf{R}_n^{1/2} \mathbf{U}_S. \quad (\text{E.7})$$

For polarization, where there is no prior on the CMB power spectra, the ILC is replaced by a constrained ILC (Remazeilles et al. 2011b), for which the vector of weights in frequency is constrained to be orthogonal to the CMB SED. In practice, this is done through a Gram-Schmidt orthogonalization of the set of eigenvectors collected in matrix $\mathbf{U}_S$ with respect to the CMB SED vector g_v . This constraint ensures that the GNILC weights (Eq. (E.6)) project out any CMB polarization signal in the reconstructed dust polarization map.

The GNILC filters (Eq. (E.6)) are invariant if $\mathbf{F}$ is replaced by $\mathbf{F}\mathbf{T}$ for any invertible matrix $\mathbf{T}$. Therefore, the true foreground mixing matrix does not need to be known by GNILC; the only useful information is a set of independent components onto which the correlated Galactic emission can be decomposed.

The estimated needlet maps of dust emission (Eq. (E.5)) are then synthesised to reconstruct the complete GNILC dust maps, as follows. The spherical harmonic coefficients, $\widehat{a}_{\ell m}^{(j)}(v)$, of the needlet dust maps, $\widehat{s}_v^{\text{dust}(j)}(p)$, are again bandpass-filtered by the needlet windows as $h_\ell^{(j)} \widehat{a}_{\ell m}^{(j)}(v)$. The filtered coefficients are then inverse-spherical-harmonic transformed into maps, and coadded across needlet scales to form the complete GNILC dust map, accounting for all the angular scales.

GNILC has many advantages over template subtraction, parametric methods, or smoothing procedures. First, it is a one-shot component-separation method that does not rely on subtraction of any template, such as a CMB template map, coming from another component-separation process. This prevents the propagation of CMB foreground residuals (e.g., dust and CIB residuals in the CMB map) to the reconstructed Galactic map.

The second advantage is related to noise filtering in *Planck* polarization maps, where GNILC performs better than a simple smoothing. Given that GNILC is a minimum-variance linear combination of frequency maps, the overall noise level in

the GNILC maps will always be lower than the noise level in smoothed *Planck* maps at the same frequency and equal resolution:

$$\frac{1}{\sigma_{\text{GNILC}}^2(353 \text{ GHz})} = \frac{1}{\sigma_{\text{Planck}}^2(30 \text{ GHz})} + \dots + \frac{1}{\sigma_{\text{Planck}}^2(353 \text{ GHz})}, \quad (\text{E.8})$$

where $\sigma_{\text{GNILC}}(353 \text{ GHz})$ is the noise rms in the GNILC 353 GHz map, and $\sigma_{\text{Planck}}(353 \text{ GHz})$ is the noise rms in the *Planck* 353 GHz map. Moreover, a simple smoothing of the *Planck* 353 GHz Q and U maps will mitigate CMB E and B modes but not cancel them on large scales, and there is no reliable CMB B -mode template to be subtracted. Conversely, GNILC is an orthogonal projection to the flat CMB SED, and therefore cancels out any CMB E - and B -mode polarization at all angular scales.

Third, GNILC filtering is performed locally over the sky and over scales via wavelet decomposition. This enables optimization of the component-separation process given local variations of contamination over the sky and over scales.

Finally, the GNILC method is blind, since it does not rely on any assumption about Galactic foregrounds. Most important, GNILC allows for outputting Galactic foreground maps at all frequencies, e.g., at 100–143 GHz, without relying on the extrapolation of high-frequency templates with arbitrary emission laws. This is particularly useful in the context of decorrelation effects and searches for primordial B modes (Tassis & Pavlidou 2015; Planck Collaboration Int. L 2017), where we can no longer rely on simple emission laws to extrapolate dust foregrounds to CMB frequencies.

Appendix F: Intensity foregrounds

In this appendix, we review the temperature foreground products derived by Commander and GNILC from the *Planck* 2018 frequency maps. As discussed in Sect. 3 and elsewhere, these results are not intended for scientific analysis, but are included here for reference and completeness purposes.

F.1. Commander analysis

We start our discussion with a review of the Commander intensity analysis. For a summary of the methodology and model definitions used in this work, see Sect. 2.1 and Appendix A. In short we fit a parametric five-component model to the *Planck* 2018 data by maximizing the standard Bayesian posterior. The 2018 model includes the following components: (1) CMB; (2) a single power-law foreground model with a free spectral index per pixel to describe the sum of low-frequency foregrounds (synchrotron, free-free, and anomalous microwave emission); (3) a modified blackbody with a free emissivity and temperature to describe thermal dust; (4) a line-emission component at 100, 217, and 353 GHz, with fixed line ratios between channels to describe CO emission; and (5) a catalogue of 12 192 known point source positions, each source being fitted with a free flux density and spectral index.

We first consider the parameters of the derived astrophysical model in intensity, starting with the point source component, which represents one of the most novel aspects of the Commander 2018 model compared to previous versions.

Starting with the amplitude maps, the most notable difference with earlier results is caused by the explicit inclusion of a radio point source component in the latest model. Each object in this component is associated with an overall flux density and spectral index across all frequencies, while the spatial projection into each frequency component is performed through a full

¹⁷ In practice, the distinction between the two sets of eigenvalues is performed via the Akaike Information Criterion, which prevents the method from overfitting the foreground subspace.

¹⁸ Those independent components are related to the subset of eigenvectors, or principal components, of the matrix $\mathbf{R}_n^{-1/2} \widehat{\mathbf{R}}_d \mathbf{R}_n^{-1/2}$ for which the associated eigenvalues depart from unity.

¹⁹ Needlet coefficients of *Planck* frequency maps.

FEBeCoP calculation, accounting for the asymmetric beam profile in the respective frequency channel. Only frequencies up to and including 143 GHz are included when fitting the flux densities and spectral indices, to avoid biases from modelling errors at high frequencies. However, the resulting model is also extrapolated to higher frequencies when fitting other components. Infrared and sub-mm sources are not explicitly modelled in this approach, since they are well described for the *Planck* frequencies within the diffuse thermal dust component, which has 5' FWHM resolution.

As described in Appendix A, the total catalogue used in this work represents a combination of four separate source catalogs, three of which (AT20G, GB6, and NVSS; [Murphy et al. 2010](#); [Gregory et al. 1996](#); [Condon et al. 1998](#)) are selected to cover disjoint regions of the sky, and the fourth (PCSS2; [Planck Collaboration XXVI 2016](#)) includes microwave sources that are not detected by any of the former three. In Table F.1, we provide summary statistics for the fits produced in the current analysis, broken down according to reference catalogue. From left to right, columns show: (1) catalogue name; (2) catalogue reference frequency; (3) total number of sources used in our combined catalogue; (4) number of sources statistically detected by Commander in the *Planck* 2018 data; (5) average flux density recalibration factor relative to the reference catalogue (no colour corrections are applied); and (6) Pearson's r correlation coefficient evaluated between the reference catalogue and Commander-estimated flux densities.

Several interesting features may be seen in Table F.1. Starting with the PCSS2 sources ([Planck Collaboration XXVI 2016](#)), the correlation between the Commander and PCSS2 flux densities at 30 GHz is very high, with a Pearson's correlation coefficient of 0.99. However, the best-fit relative amplitude between the two catalogues is $a = 0.867$. Part of this is due to the fact that the Commander flux densities are intrinsically colour corrected, and therefore correspond to a monochromatic reference frequency of 30 GHz, whereas the PCSS2 values correspond to flux densities directly observed in the 30 GHz map without colour correction. Considering that the effective frequency of the 30 GHz channel for a flat-spectrum source with a spectral energy distribution proportional to ν^{-2} is 28.4 GHz, the difference in amplitude is expected to be roughly $(28.4/30)^2 \approx 0.90$. In addition, the Commander analysis takes into account the full asymmetry of the *Planck* beams, and also exploits all frequencies between 30 and 143 GHz in the fit, while the PCSS2 catalogue only considers a symmetric Gaussian beam model, and employs the LFI 30 GHz observations alone.

The Commander fits exhibit a slightly lower correlation coefficient relative to the AT20G source catalogue at 20 GHz, with a numerical value of $r = 0.74$ and a detection rate of 91%. However, the flux-density calibration is very good, with a relative normalization factor of $a = 0.977$. At 4.85 GHz, the correlation with the GB6 catalogue flux densities is again very slightly weaker at $r = 0.69$, and this time the detection rate is 59%, with a relative normalization of $a = 0.56$. Finally, this general trend of weakening correlations becomes even stronger at lower frequencies, with the NVSS catalogue at 1.4 GHz only having a correlation coefficient of $r = 0.10$ and a relative normalization of $a = 0.163$. However, the detection rate remains fairly high, at 72%. NVSS and Commander thus agree on the existence of the set of sources, but disagree significantly on their amplitudes. This is, of course, not unexpected, when extrapolating all the way from 1.4 GHz to 30–143 GHz. The point source component as evaluated for the 30 GHz channel is plotted in the top right panel of Fig. F.1.

Next, we consider the amplitude parameter maps of the diffuse foreground components, as shown in Fig. F.1. Starting with the top left panel, this figure shows the joint low-frequency foreground component, which includes synchrotron, free-free, and anomalous microwave emission as evaluated at 30 GHz and smoothed to a resolution of 40' FWHM²⁰. A similar low-frequency foreground map was presented in [Planck Collaboration XII \(2014\)](#), derived from the *Planck* 2013 data, and the most visually striking difference between these two maps is the absence of small-scale compact objects in the updated map. This is of course due to the fact that these sources are explicitly fitted out in the new model. The resulting source amplitude map at 30 GHz is shown in the top right panel.

The bottom left panel of Fig. F.1 shows the thermal dust amplitude map evaluated at 857 GHz and smoothed to 10' FWHM. Visually speaking, this map is nearly identical to the corresponding 2015 map, since the thermal dust component is strongly dominated by the 545 and 857 GHz HFI frequency maps, and these have only changed by one or two percent since the last release (see Fig. 3).

At a strictly visual level, the same holds true for the CO component, shown in the bottom-right panel of Fig. F.1. However, in this case the reconstruction quality of the new map is notably worse than in the corresponding 2015 map, as shown in the top panel of Fig. F.2. This figure shows scatter plots between the Dame et al. CO survey map ([Dame et al. 2001](#)) and the Commander CO 2015 (cyan dots) and 2018 (grey dots) maps. Two effects are notable. First, we note that the slopes are different between the two maps, corresponding simply to the different overall normalization conventions adopted for the two maps. In particular, for the 2015 analysis we employed conversion factors between μK_{CMB} and $\text{K}_{\text{RJ}} \text{ km s}^{-1}$ derived directly from the *Planck* bandpasses measured on the ground ([Planck Collaboration IX 2014](#)). This is significantly more complicated with the single-CO line model employed in the current analysis, and with the 2018 co-added frequency maps. The scale of the current CO amplitude map is therefore instead directly set by regressing against the Dame et al. map, and the resulting scatter plot therefore by definition has a slope of unity.

More important than this choice of normalization, however, is the width and shape of the two scatter plots. Specifically, while the 2015 scatter plot exhibits a very tight overall correlation and no visually notable outliers, the 2018 scatter plot is broader overall and exhibits several outliers in the Commander map. The reasons for this weaker correlation have already been discussed in Sect. 3 and [Planck Collaboration III \(2020\)](#), and can be summarized as being due to the lack of single-bolometer HFI maps and inaccuracies in the CO template corrections used during mapmaking. As described in Appendix A, the Commander CO map is used as a tracer for CO emission in the Commander confidence mask.

Finally, we consider the spectral parameters for various components, shown in the left column of Fig. F.3 for the low-frequency and thermal dust components. These can be compared to similar maps presented in the 2013 and 2015 *Planck* releases ([Planck Collaboration XII 2014](#); [Planck Collaboration X 2016](#)). Starting with the low-frequency spectral-index map, the two most notable changes with respect to the corresponding 2013 products are different priors on spectral index ($\beta_{\text{lf}} =$

²⁰ Although all components are formally estimated without internal smoothing during the Commander analysis, the resulting maps are completely noise dominated on small scales. In practice, each component map therefore needs to be smoothed to the resolution corresponding to the most relevant frequency map for visualization purposes.




Fig. F.1. Commander foreground amplitude maps, derived from the *Planck* 2018 data set in intensity. *Top-left panel:* combined low-frequency foreground map at 40' FWHM resolution, evaluated at 30 GHz, and accounts for synchrotron, free-free, and anomalous microwave emission. *Top-right panel:* derived radio point source map, as observed in the 30 GHz frequency channel. *Bottom-left panel:* thermal dust emission at 10' FWHM resolution, evaluated at 857 GHz. Neither the CIB nor high-frequency point sources are fitted explicitly in the Commander 2018 temperature model, and these are therefore in effect included in this thermal dust emission map. *Bottom-right panel:* CO line-emission map, evaluated for the 100 GHz channel.

Table F.1. Summary of Commander point-source fits.

Catalog	ν_{ref} [GHz]	N_{tot}	N_{det}	f_{det}	a	Pearson's r	Reference
AT20G	20	4499	4096	0.91	0.977	0.74	Murphy et al. (2010)
GB6	4.85	5814	3415	0.59	0.560	0.69	Gregory et al. (1996)
NVSS	1.4	1527	1094	0.72	0.163	0.10	Condon et al. (1998)
PCCS2	28.5	352	313	0.89	0.867	0.99	Planck Collaboration XXVI (2016)

Notes. Each row corresponds to one reference catalogue, as described in the text. Columns indicate, from left to right: (1) catalogue name; (2) catalogue reference frequency; (3) total number of catalogue sources selected for the current analysis; (4) number of statistically detected sources in the current analysis; (5) detection rate; (6) relative average normalization factor between Commander-derived and reference flux densities; (7) Pearson's r correlation coefficient between Commander-derived and reference flux densities; and (8) reference publication.

-2.9 ± 0.3 in 2013 versus $\beta_{\text{lf}} = -3.1 \pm 0.5$ in 2018), resulting in a darker map at high latitudes, and an overall higher signal-to-noise ratio resulting from the inclusion of four-years of LFI observations in these new maps, as opposed to only 14 months in 2013, resulting in larger areas being data-driven. Otherwise, the two maps are largely consistent.

Relatively speaking, larger changes are seen for the thermal dust spectral parameters when compared to the 2015 model presented in Planck Collaboration X (2016). Starting with the emissivity or spectral index, β_{d} , one can see bright CO-like structures in the 2018 version, for instance near the Fan region at $(l, b) = (140^\circ, 10^\circ)$; this indicates a stronger degeneracy between CO and thermal dust in the 2018 map than in the 2015 map, and results most likely from the lack of single-bolometer maps

in the 2018 analysis. Similarly, one can see a strong dark region extending from the North to the South Ecliptic Pole in the new map. This feature is well-known in *Planck* mapmaking, and arises from bandpass mismatch between different bolometers used to create a single map. Although the most recent map-making process makes a great effort to suppress this effect (Planck Collaboration III 2020), the lack of single-bolometer and detector-set maps carries a significant price for subsequent component separation: while it was possible to remove single bolometers for which this effect was particularly pronounced in 2015 (see Fig. 2 of Planck Collaboration X 2016), only full frequency maps are available in the 2018 analysis.

At high latitudes, the most notable effect is a brighter overall distribution of small-scale fluctuations, which correspond




Fig. F.2. T – T scatter plots between the Dame et al. (2001) $J = 1 \rightarrow 0$ map and the Commander 2015 (blue dots) and 2018 (grey dots) CO maps. Note that the 2018 map has been directly calibrated to the Dame et al. map, and is therefore expected to have unity slope by construction, while the 2015 map was calibrated using the *Planck* bandpasses; this difference explains the overall shift in slopes. The lower level of scatter around the best-fit slope in the 2015 map is due to including single-bolometer and detector-set maps, as opposed to the 2018 map, which exclusively uses co-added frequency maps.

to small-scale cosmic infrared background (CIB) fluctuations. When interpreting these fluctuations, however, it is important to recall that the two-parameter β – T modified blackbody model exhibits a strong degeneracy between the spectral index and temperature in the low signal-to-noise regime. The fluctuations seen in the 2018 β map were thus also present in the 2015 rendition, but in that case were seen in the temperature map. The main reason for the apparent shift is the choice of thermal dust temperature prior, or, to be more precise, the angular resolution at which it is fitted. In 2015 the thermal dust temperature was fitted at $40'$ FWHM, while in the updated analysis it is fitted at $5'$ FWHM. As a result, the 2015 temperature map had higher effective signal-to-noise per resolution element, and therefore less dependence on the prior and more structure at high latitudes. In contrast, the 2018 temperature map has less signal-to-noise per resolution element, stronger prior dependency, and accordingly also shows less structure at high latitudes, as the temperature is driven to the prior mean, and fluctuations are instead captured in the spectral index map. In general, we caution against over-interpreting the individual parameters of the modified blackbody model in the low signal-to-noise regime, since small changes in the input can lead to relatively large variations in parameter values. In contrast, the resulting SED arising from the parameters is robust.

For completeness, we note that the best-fit CO line ratio between 100 and 217 GHz (353 GHz) is $h_{2018} = 0.58$ ($h_{353} = 0.20$), as estimated by Commander from the *Planck* 2018 data set. For comparison, the corresponding 2013 values for these two parameters were $h_{2018} = 0.595$ and $h_{353} = 0.295$. However, for the reasons discussed above, we do not attach physical significance to the lower value found in the new data set, but rather recommend




Fig. F.3. Commander 2018 foreground spectral parameters. *From top to bottom rows:* low-frequency spectral index at a $40'$ FWHM smoothing scale, thermal dust spectral index at $10'$ FWHM, and thermal dust temperature at $5'$ FWHM, respectively.

continued usage of the previous values when using *Planck* results for astrophysical analysis and forecasts.

Before concluding our discussion, we emphasize that while we do not consider the Commander 2018 intensity foreground analysis to be as robust as the corresponding 2015 analysis, this has only a very small effect on the corresponding CMB reconstruction after accounting explicitly for CO emission in the Commander confidence mask (see Appendix A.4). As far as CMB reconstruction is concerned, the only important factor is whether the sum of the apparent foregrounds may be modelled within the parameter space of the Bayesian model; whether or not those best-fit values represents the physically true sky is irrelevant. This is of course also precisely why blind CMB reconstruction methods, such as NILC, SEVEM, and SMICA, perform very well. Nevertheless, the fact that the Commander 2018 intensity products appear reasonable, and that shortcomings are understood, is reassuring.




Fig. F.4. *Top panel:* GNILC thermal dust intensity map at 353 GHz with spatially varying angular resolution. *Bottom panel:* T – T scatter plot between the thermal dust intensity Commander, both smoothed to a common angular resolution of 80' FWHM. An offset of 421 μ K has been subtracted from the GNILC map in both panels (see main text for details).

F.2. Thermal dust intensity maps and their zero levels

Finally, we compare the thermal dust intensity maps derived with Commander and GNILC. Specifically, the top panel of Fig. F.4 shows the GNILC thermal dust intensity map evaluated at 353 GHz, and the bottom panel shows a scatter plot between the Commander and GNILC estimates, where the Commander model has been integrated over the 353 GHz channel bandpass. Overall, we observe good agreement between the two estimates.

The behaviour at low intensities is particularly interesting because it is sensitive to how the zero level of each map has been set. By construction, the frequency maps delivered by the HFI DPC and used for component separation have a Galactic zero level consistent with an intensity of the dust foreground at high Galactic latitudes proportional to the column density of the ISM traced by the 21 cm emission of H I at low column densities. In the case of GNILC, the processing does not adjust the monopoles contained in the input maps, the largest of which is the CIB monopole. Therefore, the zero levels of the resulting GNILC dust maps need to be adjusted prior to Galactic

applications. This has been accomplished here, just as in Planck Collaboration Int. XLVIII (2016), by correlation with the H I map at high latitude, following the methodology set out in Planck Collaboration VIII (2014) and Planck Collaboration XI (2014). At 353 GHz, 421 μ K is subtracted. In the case of Commander, the zero level at each frequency is solved for explicitly within the component separation processing, with priors set equal to the value of the CIB monopole (see Appendix A.2). The Commander offset found at 353 GHz is 431 μ K, separate from the thermal dust emission model. Given these zero level adjustments, the agreement at low intensities is satisfactory.

Especially for applications at low intensity, it is critical to appreciate that there are significant uncertainties in the zero levels of the Commander thermal dust intensity maps derived from the Planck 2015 and 2018 frequency maps, as discussed in Sect. 6.1.1 of Planck Collaboration X (2016), and of GNILC, as discussed in Sect. 2.2 of Planck Collaboration XII (2020), including the possibility of dust associated with ionized gas. These uncertainties need to be evaluated and then propagated in any subsequent analyses using these thermal dust maps, in particular when estimating modified blackbody parameters or the polarization fraction. Ideally, the uncertainties can be reduced through improved methods of zero level determination, such as exploitation of correlations with external data sets, including H I and optical extinction (e.g., Planck Collaboration Int. XLVIII 2016), or via spatial spectral variations (Wehus et al. 2017).

Appendix G: Extra CMB plots

In this Appendix, we present supporting plots relevant for the CMB discussion. These complement and elucidate the analyses and results presented in the main text, and are useful for reference purposes.

First, Figs. G.1 and G.2 show odd-even and half-mission half-difference maps, and as such, they represent our preferred tracers of noise and instrumental systematics, respectively. The former exhibit very few large-scale correlated features, whereas the latter show clear signatures of both the Planck scanning strategy at high latitudes and Galactic contamination through calibration and leakage effects at low latitudes.

Next, Fig. G.3 shows a $20^\circ \times 20^\circ$ zoom-in of the four cleaned CMB maps, centered on the North Ecliptic Pole. The polarization pattern expected from a typical E -mode signal (“+”-type in Stokes Q , and “ $\times$ ”-type in Stokes U) is clearly visible.

Figures G.4 and G.5 show enlargements of the odd-even and half-mission half-difference maps for the same region. In these maps, notable qualitative differences between the four CMB maps are observed, perhaps the most striking of which is the effect of different point source treatments adopted by the four pipelines. For instance, in the half-mission splits one can clearly see bright source residuals in the temperature maps for Commander, NILC, and SMICA, but not for SEVEM. These are due to changes in the amplitude of point sources between both periods of observations, which show up when subtracting the half-mission splits. SEVEM does not present these residuals because it explicitly inpaints known sources positions in each split, and therefore it reduces significantly this contaminant emission in the half-mission data before constructing the half-difference maps. In the case of the SMICA polarization maps, one can also see outlines of the processing mask adopted for inpainting in that case.

Another type of qualitative difference is seen between Commander on the one side, and the other three codes on the




Fig. G.1. Odd-even half-difference CMB maps at 80' resolution. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. The common mask is marked in red.

other side. *Commander* accounts explicitly for spatial variations in instrumental sensitivity at each frequency during Wiener filtering, which corresponds to evaluating an exact inverse-noise-variance weighting pixel-by-pixel in the different channels. This procedure produces somewhat more uniform effective residual maps than the other three codes.

Next, Figs. G.6–G.9 show a single Gaussian-constrained realization evaluated for each of the cleaned CMB maps, with the inpainting mask shown in Fig. 10 applied. The temperature maps are shown at 80' FWHM resolution, and the polarization maps are shown at 80' FWHM resolution. These maps are primarily intended for presentation purposes, rather than scientific analysis, since their noise properties are complicated. If similar constrained realizations are required for quantitative analysis, we recommend users to employ a Gibbs sampler, for instance as implemented in *Commander*, to produce an ensemble of such realizations, which then collectively may be used to propagate uncertainties.

Finally, for illustration, Fig. G.10 shows one of the first half-mission noise simulations at 80' FWHM resolution propagated through the four component separation methods. The simulation

contains both instrumental noise and residual systematic effects. Some residual systematics can be seen in the Galactic plane, and are especially apparent in the SEVEM intensity map. These residuals come mainly from the 545 GHz simulated map, which seems to have larger systematics than the other channels. This explains why this structure is not visible in polarization, and also why the greatest effect is on SEVEM, which gives greater weight to this channel than the other methods. Since the maps have been smoothed to 80', the residuals extend beyond their original locations. Nevertheless, the amplitude of these residuals is relatively small in absolute values, and moreover, they are mostly contained within the common confidence mask (marked in red), even without considering an extended version of the mask that should take into account the additional smoothing of the map. Therefore, we do not expect these residuals to affect significantly the analysis carried out with the simulations. The NILC intensity map has higher noise at this resolution than the other pipelines, consistent with what has been seen in previous figures. Figure G.11 shows the same plot for one even-ring, split-noise simulation, from which similar conclusions can be derived.




Fig. G.2. Half-mission half-difference CMB maps at 80' resolution. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. The common mask is marked in red.




Fig. G.3. CMB maps smoothed to a common resolution of $10'$ FWHM. The patch shown is $20^\circ \times 20^\circ$ centred on the North Ecliptic Pole, $(l, b) = (96.38, 29.81)$. Columns show Stokes I , Q , and U , while rows show results derived with different component separation methods. The common mask is marked in red.




Fig. G.4. Odd-even half-difference CMB maps smoothed to a common resolution of 10' FWHM. The patch shown is $20^\circ \times 20^\circ$ centred on the North Ecliptic Pole, $(l, b) = (96.38, 29.81)$. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. The common mask is marked in red.




Fig. G.5. Half-mission half-difference CMB maps at $20'$ resolution. The patch shown is $20^\circ \times 20^\circ$ centred on the North Ecliptic Pole, $(l, b) = (96.38, 29.81)$. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. The common mask is marked in red.




Fig. G.6. Commander constrained-realization CMB maps. The masked regions shown in Fig. 10 have been replaced with a Gaussian-constrained realization. *From top to bottom panels:* Stokes parameters I , Q , and U . The temperature map is shown at $5'$ FWHM angular resolution, while the polarization maps are shown at $80'$ FWHM angular resolution.




Fig. G.7. NILC constrained-realization CMB maps. The masked regions shown in Fig. 10 has been replaced with a Gaussian-constrained realization. *From top to bottom panels:* Stokes parameters I , Q , and U . The temperature map is shown at $5'$ FWHM angular resolution, while the polarization maps are shown at $80'$ FWHM angular resolution.




Fig. G.8. SEVEM constrained-realization CMB maps. The masked regions shown in Fig. 10 has been replaced with a Gaussian-constrained realization. *From top to bottom panels:* Stokes parameters I , Q , and U . The temperature map is shown at 5' FWHM angular resolution, while the polarization maps are shown at 80' FWHM angular resolution.




Fig. G.9. SMICA constrained-realization CMB maps. The masked regions shown in Fig. 10 have been replaced with a Gaussian-constrained realization. *From top to bottom panels:* Stokes parameters I , Q , and U . The temperature map is shown at $5'$ FWHM angular resolution, while the polarization maps are shown at $80'$ FWHM angular resolution.




Fig. G.10. First half-mission split-noise simulation maps at $80'$ resolution. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. Monopoles and dipoles have been subtracted from the intensity maps, with parameters estimated outside a $|b| < 30^\circ$ Galactic cut.




Fig. G.11. Even ring split-noise simulation maps at $80'$ resolution. Columns show Stokes I , Q , and U , while rows show results derived with different component-separation methods. Monopoles and dipoles have been subtracted from the intensity maps, with parameters estimated outside a $|b| < 30^\circ$ Galactic cut.

Appendix H: N-point functions

Here we present 2-point and 3-point correlation functions for the HMHD and OEHD maps. These complement analyses and

figures presented in the main text (Sect. 4.7). Figures H.1–H.3 show the correlation functions for half-differences of the NILC, SEVEM, and SMICA maps, respectively.




Fig. H.1. 2-point (*upper panels*), pseudo-collapsed (*middle panels*), and equilateral (*lower panels*) 3-point correlation functions determined from the $N_{\text{side}} = 64$ Planck NILC HMHD (*left panels*) and OEHD (*right panels*) temperature and polarization map. The red solid line corresponds to the half-difference maps (HMHD or OEHD). The green triple-dot-dashed line indicates the mean determined from 300 FFP10 noise simulations. The shaded dark and light grey regions indicate the corresponding 68% and 95% confidence regions, respectively.




Fig. H.2. 2-point (*upper panels*), pseudo-collapsed (*middle panels*), and equilateral (*lower panels*) 3-point correlation functions determined from the $N_{\text{side}} = 64$ *Planck* SEVEM HMHD (*left panels*) and OEHD (*right panels*) temperature and polarization map. The red solid line corresponds to the half-difference maps (HMHD or OEHD). The green triple-dot-dashed line indicates the mean determined from 300 FFP10 noise simulations. The shaded dark and light grey regions indicate the corresponding 68% and 95% confidence regions, respectively.




Fig. H.3. 2-point (*upper panels*), pseudo-collapsed (*middle panels*), and equilateral (*lower panels*) 3-point correlation functions determined from the $N_{\text{side}} = 64$ Planck SMICA HMHD (*left panels*) and OEHD (*right panels*) temperature and polarization map. The red solid line corresponds to the half-difference maps (HMHD or OEHD). The green triple-dot-dashed line indicates the mean determined from 300 FFP10 noise simulations. The shaded dark and light grey regions indicate the corresponding 68% and 95% confidence regions, respectively.