3-2019

Taller in the Saddle: Constraining CMB Physics Using Saddle Points

Jow L. Dylan

Dagoberto Contreras

Douglas Scott

Emory F. Bunn
University of Richmond, ebunn@richmond.edu

Recommended Citation

Taller in the saddle: constraining CMB physics using saddle points

Dylan L. Jow,[a,b] **Dagoberto Contreras,**[a,c,d] **Douglas Scott**[a] **and Emory F. Bunn**[e]

[a]Department of Physics and Astronomy, University of British Columbia,
  6224 Agricultural Road, Vancouver, BC, Canada
[b]Department of Physics, University of Toronto,
  ON M5S 1A7, Canada
[c]Department of Physics and Astronomy, York University,
  ON M3J 1P3, Canada
[d]Perimeter Institute,
  31 Caroline St N, Ontario, Canada
[e]Physics Department, University of Richmond,
  Richmond, VA 23173, U.S.A.

  E-mail: djow@physics.utoronto.ca, dago@yorku.ca, dscott@phas.ubc.ca,
  ebunn@richmond.edu

Abstract. The statistics of extremal points in the cosmic microwave background (CMB) temperature (hot and cold spots) have been well explored in the literature, and have been used to constrain models of the early Universe. Here, we extend the study of critical points in the CMB to the set that remains after removing extrema, namely the saddle points. We perform stacks of temperature and polarization about temperature saddle points in simulations of the CMB, as well as in data from the *Planck* satellite. We then compute the theoretical profile of saddle-point stacks, given the underlying power spectra of the CMB. As an example of the utility of such stacks, we constrain models of cosmic birefringence, and compare the constraining power of the saddle points with that of extremal points. We find that, in the specific example of birefringence, we can place tighter constraints using saddle points in our analysis than using extrema. In fact, we find our saddle-point analysis yields close to optimal constraints, as seen by comparing to a power spectrum analysis. We, therefore, suggest that stacking on saddle points may, in general, be a useful way of testing for non-standard physics effects that change the CMB power spectra.

Contents

1   Introduction

Essential to the process of extracting cosmological information from the cosmic microwave background (CMB) radiation is a full characterization of its statistics. The temperature fluctuations in the CMB are consistent with an isotropic Gaussian random field, with some minor exceptions (e.g. due to gravitational lensing, and the integrated Sachs-Wolfe effect) [1]. Thus, to appreciate the CMB we must understand the statistics of a 2D Gaussian random field on a 2-sphere (i.e. the sky).

The study of 1D Gaussian random fields originated in the study of telecommunications [2, 3]. This approach was adapted for cosmological applications in the seminal 1986 paper by Bardeen, Bond, Kaiser & Szalay, in which the perturbations of the evolving early Universe were modelled as a 3D Gaussian field [4]. These results were particularized to the 2D case of Gaussian fields on a sphere by Bond & Efstathiou in 1987 [5]. The focus of these papers was the statistical properties of extremal points, i.e. local maxima and minima in Gaussian random fields, which has led to a fruitful avenue of study of the statistics of collapsing peaks in large-scale structure. The average profiles of the temperature and polarization fields around temperature peaks were calculated in the flat-sky limit by the *WMAP* team [6], and a more general formalism for calculating these profiles, valid at large scales (i.e. with no flat-sky approximation) and allowing for biased peak eccentricity, was developed by Marcos-Caballero et al. [7]. Stacks of temperature and polarization patches around

Figure 1. Un-oriented stacks of T, Q, U, $Q_{\rm r}$, and $U_{\rm r}$ around temperature maxima for the *Planck* SMICA CMB map. The stacks are in units of μK, and the x- and y-axes are in units of degrees. See appendix A for more details.

temperature peaks in CMB data can be compared to these calculations to test cosmological models. Importantly, the theoretical predictions for the profiles of these stacks rely on the assumptions of Gaussianity and statistical isotropy, and, therefore, they can be used as general tests for these assumptions. Using such stacks, the *WMAP* [6] and *Planck* [1] teams have found complete consistency with the standard model. Moreover, the bullseye-like plots produced by both the *WMAP* and *Planck* teams have been frequently used as a simple real-space way to visualize the oscillatory physics of the model [1, 6, 8, 9]. In order to contrast with our new stacks, we show the standard averages around temperature extrema in figure 1. The fact that these stacks match the predictions has been used to constrain non-standard physics such as cosmic birefringence [10].

As well as maxima and minima, there is a third type of critical point in 2D: saddle points. Though the formalism developed in the literature is valid for general critical points,

the focus has been almost entirely on extrema, as opposed to saddle points. However, since polarization is sourced by quadrupole anisotropies in temperature [11], temperature saddle points, which take the form of quadrupoles on the sky, are natural objects to study. Moreover, scalar, vector, and tensor perturbations source specific modes of quadrupole anisotropies [11], and so saddle points may be of interest in studying the characteristics of different sources of perturbation. Some of the general statistical properties of saddle points in Gaussian random fields have already been studied (see e.g. refs. [12–14]); however, the use of stacks to visualize and constrain physics has been primarily focused on extrema.

In section 3 of this paper, we present stacks around temperature saddle points for simulated CMB maps containing purely scalar, vector, or tensor perturbations, as well as real *Planck* data. In section 4, we compute the expected profiles of temperature and polarization about temperature saddle points in the CMB in the flat-sky limit. For completeness, we also present the result computed for the spherical sky, extending the formalism of earlier papers, particularly ref. [7]. We end with an example of the application of the use of stacked saddle-point plots for constraining exotic physics, in section 5. This enables us to explore how the constraining power of saddle points compares with that of extrema.

2 Data and simulations

In this work, we use the 2018 full-mission and half-mission *Planck* data [15–17], specifically adopting the `SMICA` component-separation [18] procedure to obtain maps of the Stokes parameters T, Q, and U. The maps are provided at a resolution of `HEALPix`[1] $N_{\text{side}} = 2048$, to which we apply a $10'$ beam. Note that the *Planck* satellite has a $5'$ beam; however, for the purposes of stacking around extrema and saddle points, applying a $10'$ beam produces stacks with less noise, but does not wash out any of the relevant features. We also use simulations of the CMB, with power spectra generated using `CAMB` [19] and the map manipulation tools of `HEALPix` [20]. We generate three sets of simulations for purely scalar, vector, and tensor perturbations, described in section 3. For all simulations, we input a fiducial ΛCDM model into `CAMB`, except for simplicity we set the number of massive neutrinos to 0 (since massive neutrino models are incompatible with vector modes in `CAMB`). We generate simulated maps at a resolution of $N_{\text{side}} = 1024$, also smoothed with a $10'$ beam to match the real data. The maps are analysed using `HEALPy`, the Python wrapper for `HEALPix` [20].

3 Stacks on temperature saddle points

3.1 Oriented stacking of temperature and polarization

We are interested in the average profiles of temperature and polarization around temperature saddle points on the sky. For maps of the CMB, we can compute the mean profile around saddle points by averaging over (i.e. stacking) small patches of the sky centred around each saddle point. Saddle points are defined as critical points that are neither maxima nor minima. That is, they are points with vanishing first derivatives, and with det $H \leq 0$, where H is the matrix of second derivatives, the Hessian. In practice, to find the location of temperature saddle points on the sky, we compute the magnitude of the gradient of the temperature, $|\nabla T|^2 = \partial_\theta^2 T + \frac{1}{\sin^2 \theta} \partial_\phi^2 T$, in the usual spherical coordinates. Critical points coincide with zeros in $|\nabla T|^2$, for which we use minima as proxies, where minima in the gradient are identified

---

[1]See http://healpix.sourceforge.net for more details.

by comparing the value of each pixel with the value of its eight nearest neighbours in the HEALPix scheme. We consider minima of $|\nabla T|^2$ with Hessians of negative determinant to be saddle points. Note that this procedure of identifying saddle points will, in general, include points that are not genuinely saddle points due to pixelization issues and the fact that not all minima of $|\nabla T|^2$ are zeros. However, tests show that these additional points give a negligible contribution to the total number of points selected.

Since we are interested in the quadrupolar profile of the saddle points, we need to perform oriented stacking, similar to what is described in section 8 of ref. [1]. That is, before stacking the patches around each saddle point, we orient them so that the major axes of the saddle points are aligned. This can be achieved by rotating each patch into a local basis where the Hessian of the temperature field at the saddle point is diagonal. For a general field f on a sphere, the Hessian, in the usual spherical coordinate basis $(\mathbf{e}_\theta, \mathbf{e}_\phi)$, is given by [e.g. 21]

$$H[f] = \begin{bmatrix} \partial_{\theta\theta} f & \frac{1}{\sin\theta}\partial_{\phi\theta} f - \frac{\cos\theta}{\sin^2\theta}\partial_\phi f \\ \frac{1}{\sin\theta}\partial_{\phi\theta} f - \frac{\cos\theta}{\sin^2\theta}\partial_\phi f & \frac{1}{\sin^2\theta}\partial_{\phi\phi} f + \frac{\cos\theta}{\sin\theta}\partial_\theta f \end{bmatrix}. \tag{3.1}$$

Thus, when stacking on saddle points, we orient the patches so that $H[T]$ is diagonal at the centre. To ensure that the patches are consistently oriented, we always select a right-handed basis, where the x-axis is chosen so that $(H_{xx} - H_{yy}) \geq 0$. In order to compute the derivatives of a full-sky map, we use the HEALPy function alm2map_der1, which computes the derivatives from the spherical harmonic transform of a given map. Higher derivatives are computed using successive applications of this function.

We stack both the temperature and polarization fields on the locations of temperature saddle points. Polarization of the light from the CMB is measured using the Stokes parameters Q and U, where positive and negative Q describe the degree of linear polarization in the $\mathbf{e}_\theta$ and $\mathbf{e}_\phi$ directions, respectively, and positive and negative U describe the polarization in the $(\mathbf{e}_\theta \pm \mathbf{e}_\phi)/\sqrt{2}$ directions, respectively. There is a third Stokes parameter, V, which describes the degree of circular polarization; however, since CMB polarization is generated through Thomson scattering, which does not produce circular polarization, V is not sourced and hence is typically ignored [11]. The polarization pseudo-vector can be computed from the Stokes parameters by

$$P = \sqrt{Q^2 + U^2}, \tag{3.2a}$$

$$\psi_{\mathrm{pol}} = \frac{1}{2}\arctan\frac{U}{Q}, \tag{3.2b}$$

where P is the magnitude of the polarization, and ψ_{pol} is the angle of the polarization pseudo-vector measured from $\mathbf{e}_\theta$ [1]. Polarization is a spin-2 quantity and transforms with rotations of the local basis according to

$$Q' + iU' = (Q + iU)e^{-2i\beta}, \tag{3.3}$$

where β is the rotation angle. Thus, when stacking we both project $(Q + iU)$ onto the new basis and multiply by a phase to obtain the oriented patch. The stacks are given in flat-sky coordinates

$$x = \varpi\cos\psi, \quad y = \varpi\sin\psi, \tag{3.4}$$

where ϖ is the angular distance from the origin, and ψ is the angle measured from the x-axis. Additionally, we can obtain the so-called radial Stokes parameters Q_{r} and U_{r} from

$$Q_{\mathrm{r}} + iU_{\mathrm{r}} = (Q + iU)e^{-2i\psi}. \tag{3.5}$$

Here, Q_{r} describes the radial component of the polarization and is positive or negative for radial or tangential polarization, respectively. As with Q and U, U_{r} describes the polarization pattern of Q_{r} rotated by 45° [22].

It is often convenient to expand polarization using the spin-2 spherical harmonics, $_{\pm 2}Y_\ell^m$, as [11, 22]

$$(Q \pm iU)(\theta, \phi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^{\pm 2} {}_{\pm 2}Y_\ell^m(\theta, \phi). \tag{3.6}$$

We use the expansion coefficients to define

$$a_{\ell m}^{E} \equiv \frac{1}{2}\big(a_{\ell m}^{2} + a_{\ell m}^{-2}\big), \tag{3.7a}$$

$$a_{\ell m}^{B} \equiv \frac{-i}{2}\big(a_{\ell m}^{2} - a_{\ell m}^{-2}\big). \tag{3.7b}$$

These can be used as coefficients in a spin-0 spherical harmonic expansion to obtain rotationally invariant quantities E and B, which describe the curl-less and divergence-less parts of the polarization pattern, respectively [11].

3.2 Saddle-point stacks for scalar, vector, and tensor maps

Anisotropies in the CMB are sourced by gravitational redshift induced by perturbations, $h_{\mu\nu}$, to the metric,

$$g_{\mu\nu} = a^2(\eta_{\mu\nu} + h_{\mu\nu}), \tag{3.8}$$

where $\eta_{\mu\nu}$ is the background Minkowski metric [23]. Any general symmetric tensor, such as the metric perturbation, can be separated into three mutually independent components that transform as scalars, vectors, and tensors, respectively, under transformations of the Poincaré group [11, 24]. Thus, CMB anisotropies are sourced by three distinct modes of perturbation. As discussed in section 2, we simulate maps of the CMB for purely scalar, vector, and tensor perturbations, separately. We then compute the temperature saddle-point stacks for each of these simulated maps. Figure 2 shows a side-by-side comparison of the stacks of T, Q, U, Q_{r}, and U_{r} around temperature saddle points for each of the scalar, vector, and tensor maps. Immediately we see that the temperature stacks are simply quadrupoles about the centre, as expected. Since tensor perturbations manifest themselves on a larger scale than scalar and vector perturbations [11], the features in the tensor stacks extend to larger scales. For this reason, while the scalar and vector stacks are shown for $4° \times 4°$ patches of the sky, the tensor stacks are shown for $20° \times 20°$ patches.

In order to extract the salient information contained in figure 2, we compute the Fourier-Bessel expansion of the stacks. Any 2D function, $f(r, \phi)$, within a disk of radius a with Dirichlet boundary conditions (i.e. $f(a, \phi) = 0$) can be expanded according to

$$f(r, \phi) = \sum_{n=1}^{\infty} \sum_{m=-\infty}^{\infty} P_{nm} \Psi_{nm}(r, \phi), \tag{3.9a}$$

where

$$\Psi_{nm} = \frac{1}{\sqrt{2\pi N_{nm}}} J_m\Big(\frac{x_{mn}}{s}r\Big)e^{im\phi}, \tag{3.9b}$$

and

$$N_{nm} = \frac{s^2}{2} J_{m+1}^2(x_{mn}). \tag{3.9c}$$

Figure 2. Oriented stacks of T, Q, U, Q_r, and U_r (rows) on temperature saddle points for simulated scalar, vector, and tensor maps (columns), in units of μK. The x- and y-axes label the flat-sky coordinates, given in degrees. Note that the tensor stacks are $20° \times 20°$, while the scalar and vector stacks are $4° \times 4°$.

Here, J_m is the Bessel function of the first kind of order m, and x_{mn} is the nth zero of J_m. The Ψ_{nm} form a complete orthonormal basis on the disk $r \leq s$, and we can compute the expansion coefficients as

$$P_{nm} = \int_0^s \int_0^{2\pi} f(r,\phi)\Psi^*(r,\phi)r\,dr\,d\phi. \tag{3.10}$$

The Fourier-Bessel expansion is the flat analogue to the spherical-harmonic transformation on a sphere. Note that this expansion implicitly assumes Dirichlet boundary conditions at

some radius $r = s$. In the case of CMB saddle-point stacks, this generally holds, since both the temperature and polarization stacks vanish at infinity, and in all cases we compute the stacks for a radius sufficiently large that the stacks are effectively zero at the boundary. For the scalar and vector stacks, this radius is $s = 2°$, while it is $s = 10°$ for the tensor stacks. Also note that when computing the Fourier-Bessel expansion of the stacks we treat the independent Stokes parameters as single complex quantities. From the coefficients, we compute their averages over n and m ($\langle |P_{mn}|^2 \rangle_n$ and $\langle |P_{mn}|^2 \rangle_m$, respectively) to determine the relative presence of various modes.

Figure 3 shows the averaged expansion coefficients for the stacks of the simulated scalar map. As we discuss further in section 4, the T and $(Q_r + iU_r)$ stacks contain only $m = \pm 2$ modes; i.e. they are purely quadrupolar. When transforming from $(Q_r + iU_r)$ to $(Q + iU)$, we multiply by $e^{2i\psi}$. Thus, the $m = \pm 2$ modes in $(Q_r + iU_r)$ are shifted up by 2 when going to $(Q + iU)$, yielding the observed $m = 0$ and $m = 4$ modes. We can then decompose $(Q + iU)$ into two distinct types of mode. Since the Fourier component for $m = 0$ is real, when splitting $(Q + iU)$ into real and imaginary parts to obtain Q and U, we will find that Q is a combination of $m = 0$ and 4 modes, while U is purely $m = 4$. Thus, the T, U, Q_r, and U_r stacks are all purely one mode, whereas Q also has a monopole contribution. We have checked that this holds true for both the vector and tensor maps as well.

We can separate out this monopole component of Q to analyse it separately from the $m = 4$ mode. Figure 4 shows the $m = 4$ component of Q compared with the full U stack for the scalar maps. We find that $Q_{m=4}$ is simply a $22.5°$ rotation of U. That this should be the case becomes apparent when we consider the fact that the Fourier mode $(Q + iU)_{m=4} = Q_{m=4} + iU$ (since U only has an $m = 4$ mode to begin with). Thus $Q_{m=4}$ and U are simply the real and imaginary parts, respectively, of a single $m = 4$ Fourier mode. In other words, $Q_{m=4}$ is modulated by $\cos(4\phi)$ and $U_{m=4}$ is modulated by $\sin(4\phi)$, and are, therefore, simply $\frac{\pi}{8}$ ($22.5°$) rotations of each other. The stacked polarization is, then, fully specified by U and the monopole contribution to Q. Since $m = 0$ modes are completely radially symmetric, it is straightforward to compute the radial profile from the azimuthal average of $Q_{m=0}$. We do this for the simulated scalar, vector, and tensor maps and show the results in figure 5. As demonstrated in section 4, the profile of the polarization stacked on temperature saddle points is determined by the two-point correlation functions, i.e. the TE and TB power spectra. From this, the features of the radial profiles in figure 5 can be qualitatively understood in a similar way to the extrema stacks, as described in section 8 of ref. [1].

So far in this paper, we have only plotted stacks for simulated data. Figure 6 shows the results of temperature saddle-point stacks on real CMB data, namely the *Planck* 2018 SMICA CMB maps. As expected, the results are consistent with the simulated scalar CMB stacks. We can also apply the same stacking procedure to other kinds of CMB maps. In particular, since the E-modes define a scalar map with the same Gaussian statistics, as with temperature, it is natural to also stack on saddle points in E. Following the above procedure, we stack E, Q, and U on E saddle points. Figure 7 shows the result of these stacks. Note the similarity of these results to the temperature stacks for the simulated vector map, with the only major difference being the scale of the features in the stacks, which is because of the different power spectrum.

Here we have shown the results of stacks performed specifically on the SMICA component-separated map; however, we have performed the same procedure on other *Planck* component-separated maps (Commander, NILC, and SEVEM [18]) and have found consistent results.
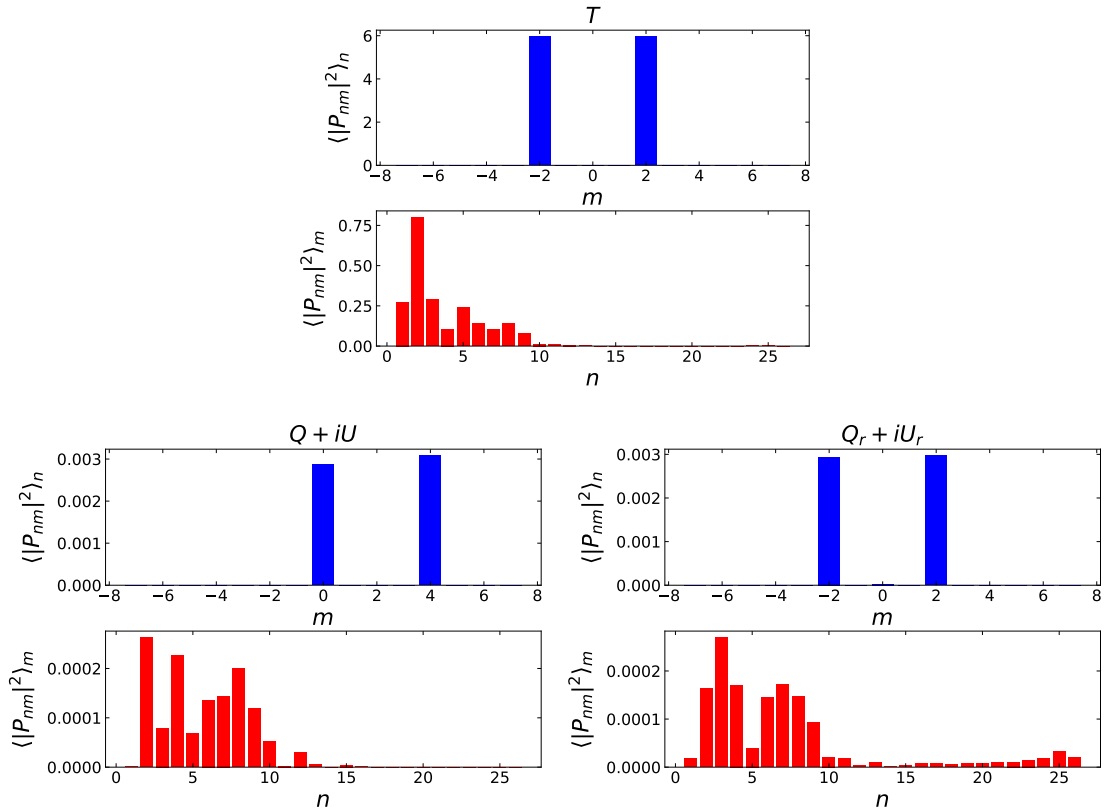
Figure 3. $\langle |P_{mn}|^2 \rangle_n$ and $\langle |P_{mn}|^2 \rangle_m$ computed for the T, $(Q+iU)$, and (Q_r+iU_r) stacks around temperature saddle points for the simulated scalar CMB. When computing the Fourier-Bessel expansion of the stacks we treat the independent Stokes parameters as single complex quantities. As expected, the T and (Q_r+iU_r) stacks have only a quadrupole ($m=\pm2$) moment. Transforming (Q_r+iU_r) to $(Q+iU)$, by multiplying by $e^{2i\psi}$, shifts the modes up by 2, yielding an $m=0$ and $m=4$ combination in $(Q+iU)$.

4 Saddle-point statistics

Now we turn to a mathematical analysis of what we have plotted in the previous section. The statistics of extremal points as a subset of critical points in a Gaussian random field have been well-characterized in the literature [e.g. refs. 4–7]. Here we generalize some of the results to saddle points in order to calculate the profiles of temperature and polarization patterns around saddle points in the CMB. Specifically, we consider critical points in the CMB temperature, $T(\theta, \phi)$, on the sky, but the following analysis holds for any scalar field on a sphere, such as E or B modes.

4.1 Critical points in the flat-sky approximation

Critical points are completely characterized by their derivatives up to and including second order. That is, the statistics of critical points are determined by a total of six degrees of freedom. In order to compute any statistical properties of critical points, we need to know the probability distribution function of these degrees of freedom. Ref. [7] presents a general formalism for the statistics of critical points on the spherical sky; however, here we work in the flat-sky approximation for the sake of simplicity.

Figure 4. Comparison between the full U stack and the $m = 4$ part of the Q stack on temperature saddle points for the simulated scalar maps. Note that by removing the $m = 0$ mode, the Q and U stacks become identical, but rotated by $22.5°$. We find that this also holds true for both the vector and tensor stacks.

For a field, $T(x, y)$, on the flat sky, the peak degrees of freedom are given by

$$\mathbf{X} = \begin{pmatrix} T \\ \partial_x T \\ \partial_y T \\ \partial_{xx}^2 T \\ \partial_{yy}^2 T \\ \partial_{xy}^2 T \end{pmatrix}. \tag{4.1}$$

For a Gaussian random field, these degrees of freedom follow a multivariate Gaussian distribution [4, 7]. In the case of the CMB anisotropy maps, the means of all of these quantities have been subtracted out, so that all of the statistical information lies in the covariance matrix $\langle \mathbf{X}\mathbf{X}^\dagger \rangle$.

To compute the covariance matrix it is convenient to work in Fourier space, defining the Fourier transform $\widetilde{T}$ to be

$$T(\boldsymbol{\theta}) = \int d^2\boldsymbol{\ell} \, \widetilde{T}(\boldsymbol{\ell}) e^{i\boldsymbol{\ell}\cdot\boldsymbol{\theta}}, \tag{4.2}$$

where

$$\langle \widetilde{T}(\boldsymbol{\ell})\widetilde{T}^*(\boldsymbol{\ell'}) \rangle = P_{TT}(\ell)\delta^2(\boldsymbol{\ell} - \boldsymbol{\ell'}). \tag{4.3}$$

Here, $\boldsymbol{\ell} = (\ell_x, \ell_y) = (\ell\cos\psi, \ell\sin\psi)$, where ψ is the angle measured from the x-axis. The entries of the covariance matrix are then simply integrals over the power spectrum $P_{TT}(\ell)$ with different factors of ℓ in the integrand for the order of the derivative. For example, $\langle T, \partial_{xx}^2 T \rangle = \int \ell_x^2 P_{TT}(\ell)d^2\ell = \pi \int P_{TT}(\ell)\ell^3 d\ell$. The rest of the entries are similarly easy to

Figure 5. Radial profiles computed from the azimuthal average of the $m = 0$ modes for the Q stacks on temperature saddle points for the simulated scalar, vector, and tensor maps, in black. The dashed blue lines are the theoretically predicted profiles in each case, obtained from eqs. (4.16)–(4.17), as described in section 4. Note that eq. (4.17) yields the profiles for Q_r and U_r, as opposed to Q and U. Since Q_r and U_r are undefined at the origin, this yields the observed departure between the stacks and predicted profiles as $r \to 0$.

compute, and we have

$$\langle \mathbf{X}\mathbf{X}^\dagger \rangle = \sigma^2 \begin{pmatrix} 1 & 0 & 0 & q^{-2} & q^{-2} & 0 \\ 0 & q^{-2} & 0 & 0 & 0 & 0 \\ 0 & 0 & q^{-2} & 0 & 0 & 0 \\ q^{-2} & 0 & 0 & 3a^{-4} & a^{-4} & 0 \\ q^{-2} & 0 & 0 & a^{-4} & 3a^{-4} & 0 \\ 0 & 0 & 0 & 0 & 0 & a^{-4} \end{pmatrix}, \tag{4.4}$$

where

$$\sigma^2 \equiv 2\pi \int d\ell\, P_{TT}(\ell)\ell, \tag{4.5a}$$

$$\frac{\sigma^2}{q^2} \equiv \pi \int d\ell\, P_{TT}(\ell)\ell^3, \tag{4.5b}$$

$$\frac{\sigma^2}{a^4} \equiv \frac{\pi}{4} \int d\ell\, P_{TT}(\ell)\ell^5. \tag{4.5c}$$
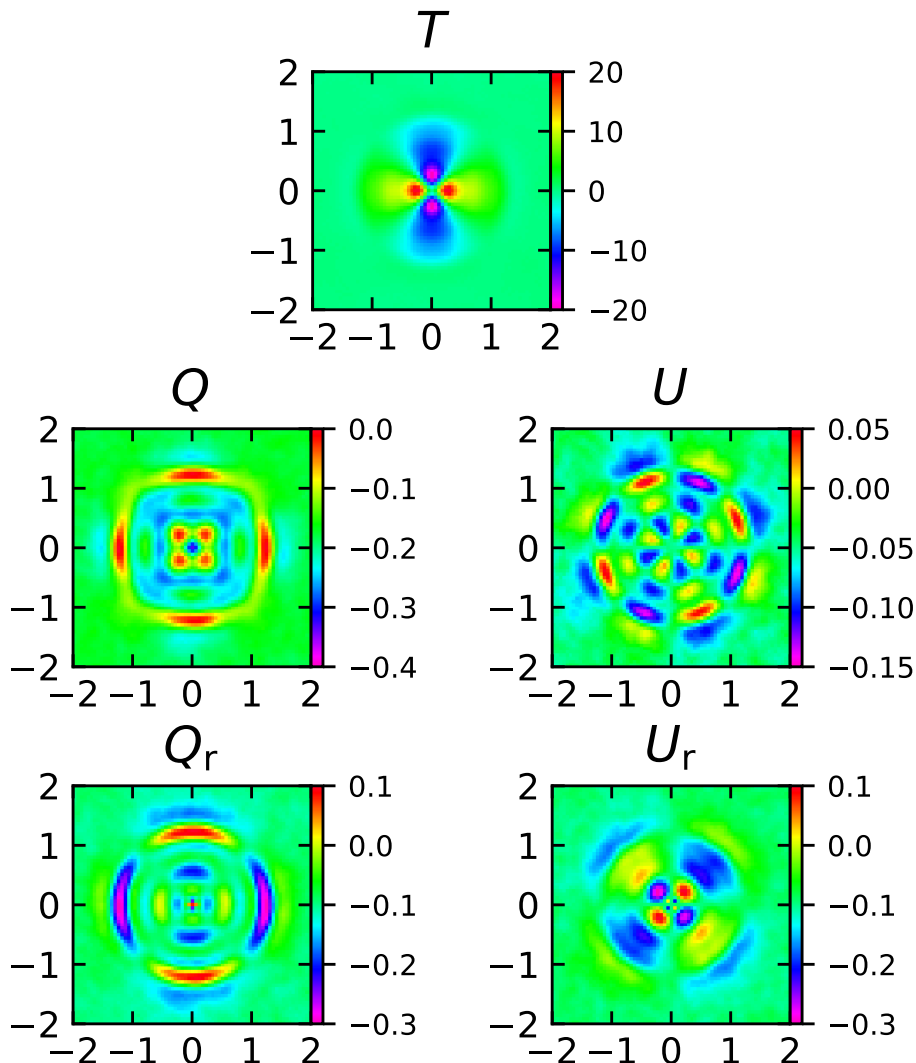
Figure 6. Oriented stacks of T, Q, U, Q_r, and U_r around temperature saddle points for the *Planck* SMICA CMB map. The stacks are in units of μK, and the x- and y-axes are in degrees. The stacks for the real data are visually hard to distinguish from the simulated scalar CMB stacks.

With this covariance matrix we can, in principle, compute any statistical property of critical points. The statistics of saddle points can then be obtained by imposing the constraint $\det H < 0$.

4.2 Saddle-point profiles

Here we compute the average profiles of T, E, and B around temperature saddle points. We will work in Fourier space, and as with temperature, we define the Fourier transforms and power spectra of E and B to be

$$E(\boldsymbol{\theta}) = \int d^2\boldsymbol{\ell}\, \widetilde{E}(\boldsymbol{\ell}) e^{i\boldsymbol{\ell}\cdot\boldsymbol{\theta}}, \tag{4.6a}$$

$$B(\boldsymbol{\theta}) = \int d^2\boldsymbol{\ell}\, \widetilde{B}(\boldsymbol{\ell}) e^{i\boldsymbol{\ell}\cdot\boldsymbol{\theta}}, \tag{4.6b}$$
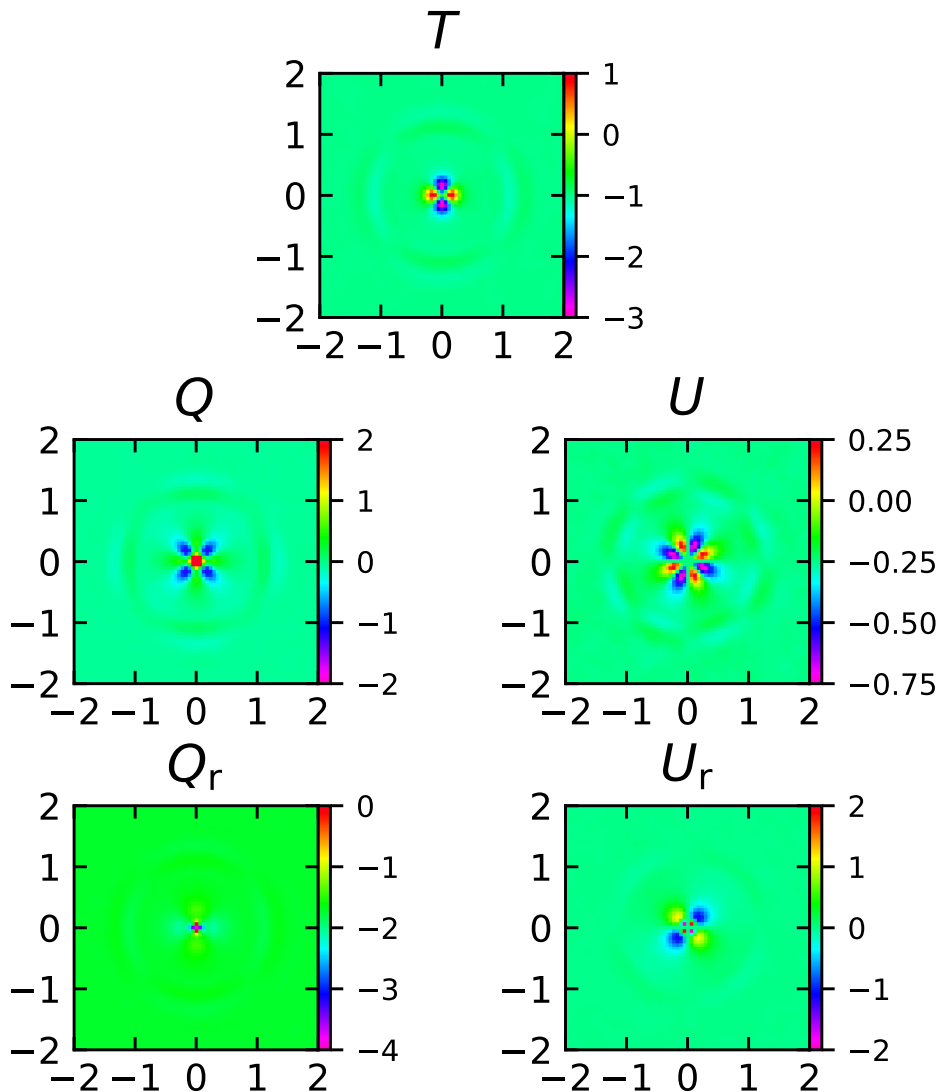
Figure 7. Oriented stacks of E, Q, U, $Q_{\rm r}$, and $U_{\rm r}$ around E saddle points for the *Planck* SMICA CMB map. The stacks are in units of μK, and the x- and y-axes are in units of degrees. Note the similarity with the temperature stacks for the simulated vector map shown in figure 2.

where

$$\langle \widetilde{E}(\boldsymbol{\ell})\widetilde{E}^*(\boldsymbol{\ell}')\rangle = P_{EE}(\ell)\delta^2(\boldsymbol{\ell}-\boldsymbol{\ell}'), \tag{4.7a}$$

$$\langle \widetilde{B}(\boldsymbol{\ell})\widetilde{B}^*(\boldsymbol{\ell}')\rangle = P_{BB}(\ell)\delta^2(\boldsymbol{\ell}-\boldsymbol{\ell}'), \tag{4.7b}$$

$$\langle \widetilde{T}(\boldsymbol{\ell})\widetilde{E}^*(\boldsymbol{\ell}')\rangle = P_{TE}(\ell)\delta^2(\boldsymbol{\ell}-\boldsymbol{\ell}'), \tag{4.7c}$$

$$\langle \widetilde{T}(\boldsymbol{\ell})\widetilde{B}^*(\boldsymbol{\ell}')\rangle = P_{TB}(\ell)\delta^2(\boldsymbol{\ell}-\boldsymbol{\ell}'). \tag{4.7d}$$

We consider the case when T has a saddle point in a small neighbourhood of the origin. As we explain in appendix B, it is necessary to phrase the condition in this way, rather than imposing the condition that there be a saddle point *at* the origin, because the latter has probability zero and leads to conditional probabilities that are undefined. We assume that the saddle point is oriented with positive eigenvector along the x-axis, so that the Hessian

has eigenvalues $\lambda_+ > 0$ and $\lambda_- < 0$. Thus, we have a set of constraints

$$\mathbf{g} \equiv \begin{pmatrix} \partial_x T \\ \partial_y T \\ \partial_{xx}^2 T \\ \partial_{yy}^2 T \\ \partial_{xy}^2 T \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \lambda_+ \\ \lambda_- \\ 0 \end{pmatrix}. \tag{4.8}$$

Since we want to average all saddle points, we will eventually want to marginalize over all λ_+ and λ_- values.

We want to know the mean values of T, E, and B subject to these constraints. In Fourier space, we are interested in

$$\overline{T}(\boldsymbol{\ell}) \equiv \langle \widetilde{T}(\boldsymbol{\ell}) | \mathbf{g} \rangle, \tag{4.9}$$

and similarly for $\overline{E}$ and $\overline{B}$. That is, we want the conditional expectation of T, E, and B given $\mathbf{g}$.

All of the quantities we are interested in are drawn from multivariate Gaussian distributions. In general, for a Gaussian random vector $\mathbf{h}$, subject to constraint $\mathbf{g}$, the expectation is given by

$$\langle \mathbf{h} | \mathbf{g} \rangle = \langle \mathbf{h} \mathbf{g}^\dagger \rangle \langle \mathbf{g} \mathbf{g}^\dagger \rangle^{-1} \mathbf{g}. \tag{4.10}$$

In our case, $\mathbf{h}$ consists of the Fourier modes $\widetilde{T}(\boldsymbol{\ell})$, $\widetilde{E}(\boldsymbol{\ell})$, $\widetilde{B}(\boldsymbol{\ell})$.

Thus, we need to compute the covariance matrix $\langle \mathbf{g} \mathbf{g}^\dagger \rangle$, but this is just the covariance matrix $\langle \mathbf{X} \mathbf{X}^\dagger \rangle$, given by eq. (4.4), restricted to the lower-right 5×5 sub-matrix. It follows that

$$\langle \mathbf{g} \mathbf{g}^\dagger \rangle^{-1} \mathbf{g} = \begin{pmatrix} 0 \\ 0 \\ \frac{a^4}{8\sigma^2}(3\lambda_+ - \lambda_-) \\ \frac{a^4}{8\sigma^2}(3\lambda_- - \lambda_+) \\ 0 \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 0 \\ \alpha_+ \\ \alpha_- \\ 0 \end{pmatrix}. \tag{4.11}$$

Now, to obtain the expectation of any element of $\mathbf{h}$, we need to multiply the corresponding row of $\langle \mathbf{h} \mathbf{g}^\dagger \rangle$ by the vector expressed in eq. (4.11). This gives us

$$\overline{Z}(\boldsymbol{\ell}) = \alpha_+ \langle \widetilde{Z}(\boldsymbol{\ell}) \partial^2 T_{xx}(0) \rangle + \alpha_- \langle \widetilde{Z}(\boldsymbol{\ell}) \partial^2 T_{yy}(0) \rangle, \tag{4.12}$$

where Z is T, E, or B. We can easily compute these last expectations, since, for instance,

$$\partial_{xx}^2 T(0) = - \int d^2 \boldsymbol{\ell}' \widetilde{T}^*(\boldsymbol{\ell}') \ell_x'^2, \tag{4.13}$$

so

$$\langle \widetilde{T}(\boldsymbol{\ell}) \partial^2 T_{xx}(0) \rangle = - \int d^2 \boldsymbol{\ell}' \ell_x'^2 \langle \widetilde{T}(\boldsymbol{\ell}) \widetilde{T}^*(\boldsymbol{\ell}') \rangle = -l_x^2 P_{TT}(\ell). \tag{4.14}$$

Thus, in the end we obtain

$$\overline{T}(\boldsymbol{\ell}) = -(\alpha_+ \ell_x^2 + \alpha_- \ell_y^2) P_{TT}(\ell), \tag{4.15a}$$

$$\overline{E}(\boldsymbol{\ell}) = -(\alpha_+ \ell_x^2 + \alpha_- \ell_y^2) P_{TE}(\ell). \tag{4.15b}$$

$$\overline{B}(\boldsymbol{\ell}) = -(\alpha_+ \ell_x^2 + \alpha_- \ell_y^2) P_{TB}(\ell). \tag{4.15c}$$

– 13 –

We now want to marginalize over λ_+ and λ_-. Since the results are linear in these quantities, this amounts to simply replacing them with their average values. Appendix B contains a calculation of these means, leading to the result $\lambda_\pm = \pm 1.027\sigma/a^2$. Using this, we obtain

$$\overline{T}(\boldsymbol{\ell}) = -0.513\frac{a^2}{\sigma}(\ell_x^2 - \ell_y^2)P_{TT}(\ell), \tag{4.16a}$$

$$\overline{E}(\boldsymbol{\ell}) = -0.513\frac{a^2}{\sigma}(\ell_x^2 - \ell_y^2)P_{TE}(\ell), \tag{4.16b}$$

$$\overline{B}(\boldsymbol{\ell}) = -0.513\frac{a^2}{\sigma}(\ell_x^2 - \ell_y^2)P_{TB}(\ell). \tag{4.16c}$$

This fully specifies the temperature and polarization profiles around temperature saddle points. As anticipated, the result gives the form of purely quadrupolar patterns in the temperature and polarization. We can then obtain the Stokes parameters using [6]

$$Q_{\rm r}(\theta) = \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2}\big\{\widetilde{E}(\boldsymbol{\ell})\cos[2(\phi-\psi)] + \widetilde{B}(\boldsymbol{\ell})\sin[2(\phi-\psi)]\big\}e^{i\boldsymbol{\ell}\cdot\theta}, \tag{4.17a}$$

$$U_{\rm r}(\theta) = \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2}\big\{\widetilde{E}(\boldsymbol{\ell})\sin[2(\phi-\psi)] - \widetilde{B}(\boldsymbol{\ell})\cos[2(\phi-\psi)]\big\}e^{i\boldsymbol{\ell}\cdot\theta}. \tag{4.17b}$$

The same calculations can also be done for the profiles of saddle points in E. Since E and T are both scalar Gaussian random fields, it is simply a matter of replacing every occurrence of "T" in the results for temperature saddle points with "E".

Note that we have derived the expected profiles conditioning on there being a saddle point within some small neighbourhood of the origin. We could have done the calculation with the condition that the critical point be located exactly at the origin — indeed, this may seem like the most natural thing to do. However, such a calculation would lead to a different numerical prefactor in eqs. (4.16). We compare the two approaches explicitly in appendix B. We show that conditioning on having a critical point at the origin is incorrect, because it involves conditioning on an event of probability zero. We note that the latter approach leads to mathematical inconsistencies. For example, one can naively compute the number of extrema and saddle points by conditioning on being at a critical point, and then compute the relative probabilities of that point being an extremum or a saddle point. This calculation gives the result that saddle points make up $1/\sqrt{3} \approx 58\%$ of all critical points, whereas a well-known theorem from Morse theory, and effectively going back to Euler and Poincaré, states that for analytic functions in the plane the numbers of saddle points and extrema are equal, and that they differ by exactly 2 points on the sphere [25].

4.3 Saddle-point profiles in the curved sky

The previous calculation gives the saddle-point profiles in the flat-sky approximation, which is more than sufficient here, since we are considering stacks for at most $20° \times 20°$ patches of the sky. However, ref. [7] presents a general formalism for critical point statistics that can be used to compute the saddle-point profiles on the sphere, which we will briefly describe here for completeness. In this formalism, the Hessian is decomposed into two parts, namely a part with non-vanishing trace and a traceless part. That is, in a local basis around the critical point, where the basis vectors $\mathbf{e}_x$ and $\mathbf{e}_y$ correspond to the usual spherical coordinate

Figure 8. Comparison between the T, Q_r, and U_r oriented stacks on temperature saddle points for the simulated scalar map (left column), and their profiles computed according to the model expressed by eq. (4.16) (middle column). The residuals were then computed (right column). Note that the residual panels have a different contrast range than the signal panels. We carried out the same procedure for the vector and tensor maps and found similar consistency.

basis vectors $\mathbf{e}_\theta$ and $\mathbf{e}_\phi$, we can write the Hessian as

$$H[T] = \begin{bmatrix} \partial_x^2 T & \partial_x \partial_y T \\ \partial_x \partial_y T & \partial_y^2 T \end{bmatrix} = \frac{1}{2} \begin{bmatrix} \nabla^2 T & 0 \\ 0 & \nabla^2 T \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \mathrm{Re}(\bar{\eth}^*)^2 T & -\mathrm{Im}(\bar{\eth}^*)^2 T \\ -\mathrm{Im}(\bar{\eth}^*)^2 T & -\mathrm{Re}(\bar{\eth}^*)^2 T \end{bmatrix}. \quad (4.18)$$

Here the derivatives $\bar{\eth}$ and $\bar{\eth}^*$ are the spin ladder operators on the sphere, which are proportional to the covariant derivatives on the sphere in the helicity basis (i.e. $\mathbf{e}_\pm = (\mathbf{e}_\theta \pm \mathbf{e}_\phi)/\sqrt{2}$). In the local, flat coordinates the ladder operator can be written as $\bar{\eth}^* = -\partial_x + i\partial_y$. The Hessian of a critical point contains information about the curvature around that point. In separating out the traceless part of the Hessian, we have decomposed the Hessian into one part describing the magnitude of the curvature, characterized by the Laplacian $\nabla^2 T$, and another describing the eccentricity of the curvature, characterized by the derivative $(\bar{\eth}^*)^2 T$.

In this way, we can define four critical-point parameters, corresponding to the value of the field at the point, the first derivatives, and the magnitude and eccentricity of the local curvature. That is, one can define [7]

$$\nu \equiv \frac{T}{\sigma_\nu}, \qquad \kappa \equiv -\frac{\nabla^2 T}{\sigma_\kappa}, \qquad \eta \equiv \frac{\bar{\eth}^* T}{\sigma_\eta}, \qquad \epsilon \equiv \frac{(\bar{\eth}^*)^2 T}{\sigma_\epsilon}, \quad (4.19)$$

where we have divided by the square roots of the variances, $\sigma_\nu^2 = \langle T^2 \rangle$, $\sigma_\eta^2 = \langle |\bar{\eth}^* T|^2 \rangle$, $\sigma_\kappa^2 = \langle |\nabla^2 T|^2 \rangle$, and $\sigma_\epsilon^2 = \langle |(\bar{\eth}^*)^2 T|^2 \rangle$, so that the parameters are normalized to unit variance.

Since η and ϵ are complex numbers, there are once again six degrees of freedom characterizing the statistics of critical points. For a Gaussian random field, the probability distribution over the variables ν, κ, η and ϵ is Gaussian, with covariance matrix [4, 7]

$$S = \begin{pmatrix} 1 & \rho & 0 & 0 \\ \rho & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{4.20}$$

where $\rho = \sigma_\eta^2/\sigma_\nu\sigma_\kappa$.

In order to compute the saddle-point profiles, we can expand fields on the sphere according to

$$f(\theta, \phi) = \sum_{m=-\infty}^{\infty} f_m(\theta)e^{im\phi}, \tag{4.21a}$$

$$f_m(\theta) = \frac{1}{2\pi} \int d\phi f(\theta, \phi)e^{-im\phi}. \tag{4.21b}$$

Note that if f is a real field, then we have the property that $f_m = f_{-m}^*$. We then need to compute the expected radial profiles $\langle T_m(\theta) \rangle$, $\langle Q_{\mathrm{r},m}(\theta) \rangle$, and $\langle U_{\mathrm{r},m}(\theta) \rangle$, where Q_r and U_r are the radial Stokes parameters for a saddle point at $\theta = 0$. Since critical points are completely characterized by their value, curvature, and eccentricity (together with the fact that they have vanishing first derivatives), the only non-vanishing parts are the $m = 0$ and $m = 2$ modes [7], which depend on the spherical-harmonic power spectra C_ℓ^{TT}, C_ℓ^{TE} and C^{TB}, and the constraints on the peak variables ν, η, κ, and ϵ. In particular, the critical-point profiles depend on the mean values of the peak parameters subject to these constraints [7]. Thus, in order to compute the expected saddle-point profiles, we need to determine the biases b_ν, b_κ, and b_ϵ, given by

$$\begin{pmatrix} b_\nu\sigma_\nu \\ b_\kappa\sigma_\kappa \end{pmatrix} = \Sigma^{-1} \begin{pmatrix} \langle\nu\rangle \\ \langle\kappa\rangle \end{pmatrix}, \tag{4.22a}$$

$$b_\epsilon = \frac{\langle\epsilon\rangle}{\sigma_\epsilon}, \tag{4.22b}$$

where Σ is the covariance matrix between ν and κ. We compute these expectations subject to the constraint that there is a critical point in a small neighbourhood around the origin, and that $|\epsilon| \geq \sqrt{a}|\kappa|$, where $a \equiv \sigma_\kappa^2/\sigma_\epsilon^2$, which selects for saddle points. We stress again that simply conditioning on $\eta = 0$ at the origin leads to inconsistent results. Instead, we impose the constraint that $\eta = 0$ for some point in a neighbourhood around the origin, which amounts to multiplying the integrand of the expectation integral by the determinant of the Hessian, which, in terms of the peak variables, is given by $\det(H) = (1/4)[\sigma_\kappa^2\kappa^2 - \sigma_\epsilon^2(\Re(\epsilon))^2 - \sigma_\epsilon^2(\Im(\epsilon))^2]$. Intuitively, we might expect to find that b_ν and b_κ vanish for saddle points, since there is no reason for there to be a preference for saddle points to be either positive or negative in value or in curvature. We numerically verify that, indeed, b_ν and b_κ are much less than b_ϵ for saddle points, so that only the $m = 2$ mode effectively contributes to the saddle-point profiles. Note, also, that in general the bias b_ϵ is a complex number; however, we can remove the imaginary part by rotating to the basis defined by the principle axes of the saddle point, since $\Im(\epsilon) \propto \partial_x\partial_y T$, which vanishes along the principle axes. In this paper, we always compute oriented saddle point stacks so that b_ϵ is real.

The $m = 2$ radial profiles are computed to be [7]

$$\langle T_2(\theta) \rangle = b_\epsilon \sum_{\ell=0}^{\infty} \frac{2\ell + 1}{4\pi} (W_\ell^T)^2 C_\ell^{TT} \, P_\ell^2(\cos\theta), \tag{4.23a}$$

$$\langle Q_{r2}(\theta) \rangle = -2b_\epsilon \sum_{\ell=0}^{\infty} \frac{2\ell + 1}{4\pi} \sqrt{\frac{(\ell-2)!}{(\ell+2)!}} W_\ell^T W_\ell^P \left[C_\ell^{TE} \, P_\ell^+(\cos\theta) + iC_\ell^{TB} \, P_\ell^-(\cos\theta) \right], \tag{4.23b}$$

$$\langle U_{r2}(\theta) \rangle = 2ib_\epsilon \sum_{\ell=0}^{\infty} \frac{2\ell + 1}{4\pi} \sqrt{\frac{(\ell-2)!}{(\ell+2)!}} W_\ell^T W_\ell^P \left[C_\ell^{TE} \, P_\ell^-(\cos\theta) + iC_\ell^{TB} \, P_\ell^+(\cos\theta) \right]. \tag{4.23c}$$

Here, the P_ℓ^m are the associated Legendre polynomials, and W_ℓ^T and W_ℓ^P are the spherical harmonic transformations of the smoothing functions (i.e. the beam) applied to the temperature and polarization maps. We have also written

$$
\begin{aligned}
P_\ell^+(x) &= -\left[\frac{\ell-4}{1-x^2} + \frac{1}{2}\ell(\ell-1) \right] P_\ell^2(x) + (\ell+2)\frac{x}{1-x^2}P_{\ell-1}^2(x), \\
P_\ell^-(x) &= -2\left[(\ell-1)\frac{x}{1-x^2}P_\ell^2(x) - (\ell+2)\frac{1}{1-x^2}P_{\ell-1}^2(x) \right].
\end{aligned}
\tag{4.24}
$$

Thus, the saddle-point profiles depend only on the eccentricity bias, and the power spectra C_ℓ^{TT} and C_ℓ^{TE}, since, in general, we have $C_\ell^{TB} = 0$. Equation 4.23 can be shown to reduce to eq. (4.16) in the flat-sky limit (i.e. for $\ell \gg 1$, $\theta \ll 1$, and $\ell\theta \sim 1$) using the fact that $\epsilon = (\lambda_+ - \lambda_-)/\sigma_\epsilon$, $\int_\alpha^{2\pi+\alpha} d\psi \cos[2(\phi-\psi)]e^{ix\cos(\phi-\psi)} = -2\pi J_2(l\theta)$, and the substitutions for going to the flat-sky approximation given in appendix C of ref. [7].

5 Constraints on cosmic birefringence

Equipped with the theoretical calculations for the saddle-point profiles, as well as the extremum profiles given in refs. [6] and [7], we can, in principle, use stacks of the data to constrain any model that affects the CMB power spectra, particularly models that alter the polarization spectra. It is clear, in the case of extrema, that when using hot spots one can add the constraining power of cold spots as independent information [6]. However, two obvious questions arise when extending this to saddle points: how do the constraints from saddle-point stacks compare to hot- and cold-spot stacks? and to what extent are the saddle-point constraints independent? Here we use saddle-point stacks to test cosmic birefringence as an example of their utility in constraining physical models. We compare the performance of saddle-point stacks with a similar test using temperature hot and cold spots previously performed in ref. [10]. We stress that, for the purposes of this paper, birefringence is just chosen as an explicit example of physics that could be constrained using stacks. We expect the lessons learned in this section to apply generally.

Models of parity violation in the electromagnetic sector of the Standard Model of particle physics predict a rotation of the plane of polarization of photons as they travel in vacuum [26]. Within such models, the polarization of photons emitted from the last-scattering surface will rotate as the photons travel through the Universe towards our detectors. The amount of rotation, α, is the cosmic birefringence angle, which can, in principle, vary with direction [27]. Here, we only consider the case in which α is the same in all directions. The effect of a nonzero α results in a mixing of the E- and B-modes of the CMB polarization, in addition to

the generation of TB and EB correlations that would otherwise vanish. Explicitly, for a constant α, we can write the observed power spectra as [26]

$$C_\ell'^{TT} = C_\ell^{TT}, \tag{5.1a}$$

$$C_\ell'^{EE} = C_\ell^{EE} \cos^2(2\alpha) + C_\ell^{BB} \sin^2(2\alpha), \tag{5.1b}$$

$$C_\ell'^{BB} = C_\ell^{EE} \sin^2(2\alpha) + C_\ell^{BB} \cos^2(2\alpha), \tag{5.1c}$$

$$C_\ell'^{TE} = C_\ell^{TE} \cos(2\alpha), \tag{5.1d}$$

$$C_\ell'^{TB} = C_\ell^{TE} \sin(2\alpha), \tag{5.1e}$$

$$C_\ell'^{EB} = \frac{1}{2}(C_\ell^{EE} - C_\ell^{BB}) \sin(4\alpha), \tag{5.1f}$$

Here, the C_ℓ' are the observed spectra, while the unprimed C_ℓs are the spectra for a universe without birefringence. The modulation of the TE correlation and generation of the TB correlation naturally results in different profiles for the polarization stacks around temperature saddle-points compared to the case of $\alpha = 0$, as seen clearly in eq. (4.23). Thus, we can potentially use the saddle-point stacks to compute likelihoods for constraining the value of α.

In order to test the constraining power of the saddle-point stacks compared to the extrema stacks, we generate birefringent maps with a known α from the simulated scalar power spectra generated by CAMB. It is simple to generate birefringent maps, since the new polarization is simply a rotation of the old polarization by α, which can be computed from eq. (3.3). An estimate of the signal-to-noise ratio, S/N, for the angle α can be obtained from the power spectra. For small α, and in the case where $C_\ell^{BB} = 0$, the signal manifests predominantly in the generation of TB correlations (as well as EB correlations; however, for simplicity we restrict our attention to TB in the following analysis). Then the signal-to-noise is just the signal-to-noise of the TB correlations, which is given by [28]

$$\frac{S}{N} = \sum_{\ell=2}^{\infty} \frac{(2\ell+1)C_\ell'^{TB}}{\widetilde{C}_\ell^{BB}\widetilde{C}_\ell^{TT} + (\widetilde{C}_\ell^{TB})^2}, \tag{5.2}$$

where $\widetilde{C}_\ell$ indicates the power spectra in the case of no birefringence, but with noise. Note that if the birefringent maps are generated using power spectra without noise, then the signal-to-noise ratio is infinite in the case of vanishing BB correlations. For current experiments, the signal-to-noise is most sensitive to noise in the polarization, and remains relatively unaffected by noise in the temperature. Therefore we choose to add noise only to the polarization. The noisy power spectra are then given by

$$\widetilde{C}_\ell^{EE} = C_\ell^{EE} + \frac{n^2}{(W_\ell^P)^2}, \tag{5.3a}$$

$$\widetilde{C}_\ell^{BB} = C_\ell^{BB} + \frac{n^2}{(W_\ell^P)^2}, \tag{5.3b}$$

where W_ℓ^P is the polarization beam function, and n is a constant parameterizing the amount of noise. Since we are interested in a comparison of the performance of the saddle-point and extrema stacks, the choice of n is arbitrary. Here we choose $n = 0.04\,\mu\text{K}$.

In order to recover the value of α using saddle-point stacks, we need to compute the residuals from the stacks performed on the (simulated) data, and the predicted profile of the

stacks computed from the power spectra. Using a uniform prior on α (for the range $(-\frac{\pi}{2}, \frac{\pi}{2}]$), we have

$$P(\alpha|d) \propto P(d|\alpha), \tag{5.4}$$

with

$$P(d|\alpha) = \frac{1}{\sqrt{2\pi|\Sigma|}} \exp -\frac{1}{2}(d - (Q_{\mathrm{r}}, U_{\mathrm{r}})(\alpha))^{\mathsf{T}} \Sigma^{-1}(d - (Q_{\mathrm{r}}, U_{\mathrm{r}})(\alpha)). \tag{5.5}$$

Here, d is the data, i.e. the stacked images of Q_{r} and U_{r}, and $(Q_{\mathrm{r}}, U_{\mathrm{r}})(\alpha)$ are the theoretical profiles computed as a function of α given by eqs. (4.16)–(4.17) combined with eq. (5.1). The quantity Σ is the covariance matrix of the saddle-point stacks and is computed explicitly for general stacks in ref. [7]. Here, we follow ref. [26] and estimate the covariance matrix by weighting the rms of the noise in polarization with the inverse of the total number of pixels used in each re-gridded pixel in the stacked image. We assume that the covariance is diagonal in pixel space, which, for large enough pixel bins, is a reasonable approximation.

We identify critical points in the CMB maps as points in the map of $|\nabla T|^2$ that are smaller than their eight nearest-neighbour pixels, as discussed in section 3. Saddle points are then critical points with $\det(H) < 0$ and extrema are critical points with $\det(H) > 0$. Hot spots are distinguished from cold spots by the sign of the Laplacian, $\nabla^2 T$. This method of identifying hot and cold spots differs from the method used, for example, in refs. [1, 26], where hot spots are identified as pixels in the map that have a greater value than their nearest-neighbour pixels, and similarly for cold spots. We, however, choose to find the extrema in a manner that is consistent with our method of finding saddle points, so that, for example, we find an equal number of extrema and saddle points, consistent with Morse theory [25]. Using the nearest neighbour-method for identifying extrema returns 180,000 extrema, whereas the critical point method returns 274,000 extrema, which is equal to the number of saddle points found. This is an approximately 50% increase in the number of extrema identified. Since not all minima are, in fact, zeros of the gradient, the critical point method will generally include points that are not true extrema. It is reasonable to ask, then, whether these additional points contain any information or are simply adding noise. In the present example of testing birefringence, we performed the analysis with both the critical-point and nearest-neighbour methods, and found that both methods recovered the same signal-to-noise ratio. Thus, while it is difficult to argue that the additional points are comprised of mostly genuine extrema, both sets of points contain the same information. Here we choose to use the critical point method for consistency with the saddle-point analysis.

For both the extrema and saddle points, we consider all points weighted uniformly, without proof that this is an optimal weighting for either case. The posterior on α is approximately Gaussian, so we only need two numbers to describe the recovered signal, namely, the mean, μ, and the standard deviation, σ. To compare the constraining power of the saddle points and the extrema, we compare the signal-to-noise ratio, μ/σ, of both methods.

We find that for $\alpha = 0.2°$, the extrema recover a signal of $\alpha = 0.190° \pm 0.030°$, while the saddle points recover $\alpha = 0.190° \pm 0.023°$. The signal-to-noise ratio is, therefore, (rounding the results) 6 for the extrema, and 8 for the saddle points. Thus, in this specific example, the saddle-point stacks contain more constraining power than the hot and cold spots combined. One might imagine that, ideally, one would be able to combine the information from the saddle points and extrema to obtain even higher signal-to-noise. However, it is not immediately clear how the saddle points and extrema stacks covary, and thus how to combine the two. In the best case scenario, the saddle points and extrema would simply be independent,

so that the noise effectively scales as the inverse of the square of the total number of saddle points and extrema combined. To gain some insight into whether this is the case, we can compare the signal-to-noise ratios of the stacks with the signal-to-noise estimated from the power spectra, given by eq. (5.2). For $\alpha = 0.2°$, the power spectra give an estimated signal-to-noise of 9. Under the assumption of Gaussianity, the power spectra contain all of the physical information in the CMB. Thus, we can take the estimate of the signal-to-noise from the power spectrum to be essentially optimal. With this assumption, we can see that the saddle points return a nearly optimal constraint. Thus, it must be that the information in the extrema and saddle points is at least partly redundant, and cannot be naively combined.

We have also tested that we obtain the same general trends for other values of α. For example, for $\alpha = 0.5$, the signal-to-noise for extrema is 16, whereas it is 23 for all saddle points. The signal-to-noise ratio estimated from the power spectra is also 23, and so in this case the saddle points achieve equivalent constraining power as the power spectra. Saddle-point stacks, therefore, appear to perform better than stacks on extrema when constraining non-standard physics. It is important to note that this is not a rigorously proven mathematical result; we have only shown that the saddle points identified by the critical-point algorithm constrain birefringence close to optimally and contain more information than the extrema identified by either the critical-point of nearest-neighbour algorithm. Since we do not claim that any of these algorithms identify saddle points or extrema perfectly, we are unable to compare the theoretical constraining power for these sets of points. However, we have empirically shown that our saddle point analysis outperforms the extrema analysis commonly used in the literature. Of course, such stacks do not outperform power-spectrum analysis (at least for Gaussian statistics), but there are occasions where there may be advantages to using a real-space approach. For one thing, stacks may be less susceptible to some kinds of systematic effects, and hence stacks (particularly on saddle points) provide a valuable alternative approach for testing any physics that might alter the power spectra.

6 Conclusions

Although stacks around hot spots and cold spots in CMB maps have been extensively studied, little attention has been given to saddle points. We have performed stacks of saddle points on simulated microwave skies, generated from purely scalar, vector, and tensor perturbations, as well as for real *Planck* data. The features in these stacks are visually striking, and can also be used to characterize some of the physics underlying the CMB correlation functions. Additionally, such saddle-point stacks can be used to constrain non-standard physics. To that end, we have computed the theoretically predicted profiles of the stacks, both for the spherical sky, and in the flat-sky approximation. As an example of using the stacks to test physics, we simulated birefringent skies and used saddle-point and extremal-point stacks to constrain the birefringent angle. Comparing the constraining power, we found that points identified as saddle points by the critical-point algorithm contain more information than points identified as extrema using either the critical-point or nearest-neighbour algorithms. Therefore, the addition of saddle points in tests of new physics, along with extremal points, has the potential to improve physical constraints. Indeed, the saddle-point stacks achieve a signal-to-noise ratio on the birefringence angle approaching the estimated signal-to-noise from the power spectrum, suggesting that saddle-point stacks alone may provide a close-to-optimal estimator of the birefringence angle. We also performed preliminary tests comparing saddle points to extrema for constraining the presence of tensor modes in simulated CMB

maps and found, again, that saddle points have more constraining power; however, further work should be done to characterize the statistics of using stacks to constrain tensor modes, and other changes in the physics of the power spectrum. Although we have not investigated a wide range of examples, and have not provided a strict proof of any general claim, it appears to be the case in practice that saddle-point stacks provide a more sensitive test than stacks on extrema. Whether or not this is always the case, we certainly recommend that future tests of physics using stacks should also include saddle-point stacks, in addition to extrema, as a potentially stronger constraint.

Acknowledgments

A Un-oriented stacks on hot and cold spots

For the sake of comparison with the saddle-point stacks, in figure 1 we show un-oriented stacks of temperature and polarization (given in terms of the Stokes parameters Q and U) on temperature maxima (defined as pixels with temperature greater than the eight nearest neighbours) above a threshold $\nu = 0$, for the `SMICA` data. The stacks are un-oriented in the sense that the patches around each hot spot are added together with random orientations with respect to one another. We also show the so-called radial Stokes parameters $Q_{\rm r}$ and $U_{\rm r}$, defined in section 3.1. We stress that the rotation of the Stokes parameters to obtain the radial Stokes parameters parameters is unrelated to the rotations performed to obtain oriented stacks. Stacks similar to those of figure 1 can be found elsewhere, for example, in refs. [6] and [1].

In contrast to the stacks found in refs. [6] and [1], which identify hot spots and cold spots as pixels whose value is greater or less than, respectively, their nearest-neighbour pixels, in the rest of this paper we first identify critical points as minima (with respect to their nearest neighbours) and then extrema as those critical points with positive $\det(H)$. Hot spots are, then, those pixels with $\nabla^2 T < 0$ and cold spots are those with $\nabla^2 T > 0$. This method of finding extrema is consistent with our method of finding saddle points.

B The ratio of saddle points to critical points

In this appendix, we calculate the probability that a given critical point is a saddle point. This appears to be a straightforward computation of a conditional probability, but as we will see care must be taken in specifying the problem correctly. We work in the flat-sky approximation throughout (but note that it is easy to extend the results to the full sky).

We will begin by deliberately performing an incorrect calculation, and then show how it must be corrected. We conclude with a calculation of the mean Hessian eigenvalues for saddle points, which is needed in section 4.2.

B.1 Incorrect calculation

Let T be a realisation of a Gaussian random process on the plane, following the notation of section 4.1. We calculate the conditional probability that the origin is a saddle point, given that it is a critical point. The latter condition requires that $\partial_x T(0) = \partial_y T(0) = 0$, while the former requires that the determinant of the Hessian be negative. We parameterize the Hessian matrix elements in terms of the quantities

$$h_0 = \left(\frac{a^2}{\sigma}\right)\{\partial_{xx}T(0) + \partial_{yy}T(0)\}, \tag{B.1}$$

$$h_+ = \left(\frac{a^2}{\sigma}\right)\{\partial_{xx}T(0) - \partial_{yy}T(0)\}, \tag{B.2}$$

$$h_\times = \left(\frac{a^2}{\sigma}\right)2\partial_{xy}T(0). \tag{B.3}$$

Using the covariance matrix elements in eq. (4.4), it is straightforward to show that these are independent Gaussian random variables with mean zero and variances 2, 1, and 1, respectively.

The determinant of the Hessian is

$$\det(\mathbf{H}) = \frac{\sigma^2}{4a^4}\left\{h_0^2 - (h_+^2 + h_\times^2)\right\}, \tag{B.4}$$

so the conditional probability we wish to compute is

$$p_{\text{saddle}} \equiv P\left(h_0^2 < h_+^2 + h_\times^2 \mid \partial_x T(0) = \partial_y T(0) = 0\right). \tag{B.5}$$

The second derivatives of T are all independent of the first derivatives, as seen in eq. (4.4) (or by a parity argument). So we can ignore the condition in this expression and compute the probability by integrating over the Gaussian distribution of $h_0, h_+, h_\times$:

$$p_{\text{saddle}} = \int_{h_0^2 < h_+^2 + h_\times^2} G(h_0; 0, 2)G(h_+; 0, 1)G(h_\times; 0, 1)\, dh_0\, dh_+\, dh_\times = \frac{1}{\sqrt{3}}. \tag{B.6}$$

Here $G(x; \mu, \sigma^2)$ is the Gaussian distribution with mean μ and variance σ^2.

We conclude that $1/\sqrt{3} = 0.577$ of all critical points are saddle points. This result is found in the literature, for example in ref. [13], although it does not affect the main results of their paper.

We can see that the above conclusion is incorrect by comparing with a well-known result from Morse theory, which says that all functions on the sphere (subject to very mild hypotheses) have precisely two more extrema than saddle points. In general, the number of saddle points exceeds the number of extrema in an analytic function on a manifold by the Euler characteristic of the manifold [25]. Suppose that we have a Gaussian random process on the sphere, whose power spectrum contains mostly small-scale power, so that the flat-sky approximation is good. Then both the number of extrema and the number of saddle points are large, and the fraction of all critical points that are saddle points must be $1/2$, not $1/\sqrt{3}$.

B.2 Correcting the problem

The error in the above calculation arises from the definition of p_{saddle} in equation (B.5). The condition $\partial_x T(0) = \partial_y T(0) = 0$ has probability zero, so the conditional probability is undefined.

To solve this problem, we must replace that condition with the condition $\partial_x T(\mathbf{r}) = \partial_y T(\mathbf{r}) = 0$ for some $\mathbf{r}$ within a given small neighbourhood $\mathcal{N}$ of the origin. In such a neighbourhood, we can approximate

$$\nabla T(\boldsymbol{\theta}) = \nabla T(0) + \mathbf{H}\boldsymbol{\theta}, \tag{B.7}$$

where $\mathbf{H}$ is the Hessian matrix at the origin. There will be a critical point in the neighbourhood as long as $-\nabla T(0)$ lies within the neighbourhood $\mathbf{H}\mathcal{N}$. The area of the latter is $|\det(\mathbf{H})|$ times the original area, and hence the probability of finding a critical point in a neighbourhood, given $\mathbf{H}$, is also proportional to $|\det(\mathbf{H})|$. By Bayes's theorem, therefore, the probability density of $h_0, h_+, h_\times$, given that there is a critical point in a small neighbourhood of the origin, is

$$f(\mathbf{h} \,|\, \text{nearby critical point}) \propto f(\mathbf{h})|\det(\mathbf{H})|, \tag{B.8}$$

where $f(\mathbf{h})$ is the Gaussian distribution on $\mathbf{h} = (h_0, h_+, h_\times)$.

The fraction of all critical points that are saddle points then becomes

$$p_{\text{saddle}} = \frac{\int_{h_0^2 < h_+^2 + h_\times^2} G(h_0; 0, 2) G(h_+; 0, 1) G(h_\times; 0, 1)|\det(\mathbf{H})|\, dh_0\, dh_+\, dh_\times}{\int G(h_0; 0, 2) G(h_+; 0, 1) G(h_\times; 0, 1)|\det(\mathbf{H})|\, dh_0\, dh_+\, dh_\times}. \tag{B.9}$$

This expression evaluates to $1/2$, in agreement with Morse theory.

B.3 Mean eigenvalues

In section 4.1, we need the mean values of the Hessian eigenvalues, taken over all saddle points. We can find this by integrating over the probability density (B.8). The Hessian eigenvalues are

$$\lambda_\pm = \frac{\sigma}{2a^2}\left(h_0 \pm \sqrt{h_+^2 + h_\times^2}\right). \tag{B.10}$$

So the mean over all saddle points is

$$\bar{\lambda}_\pm = \frac{\int_{h_0^2 < h_+^2 + h_\times^2} \lambda_\pm G(h_0; 0, 2) G(h_+; 0, 1) G(h_\times; 0, 1)|\det(\mathbf{H})|\, dh_0\, dh_+\, dh_\times}{\int_{h_0^2 < h_+^2 + h_\times^2} G(h_0; 0, 2) G(h_+; 0, 1) G(h_\times; 0, 1)|\det(\mathbf{H})|\, dh_0\, dh_+\, dh_\times}, \tag{B.11}$$

which evaluates to

$$\bar{\lambda}_\pm = \pm\frac{\sigma}{a^2}\left(\frac{10 + 3\sqrt{2}\cot^{-1}\sqrt{2}}{4\sqrt{3}\pi}\right) = 1.027\frac{\sigma}{a^2}. \tag{B.12}$$

References

[1] PLANCK collaboration, *Planck 2015 results. XVI. Isotropy and statistics of the CMB*, *Astron. Astrophys.* **594** (2016) A16 [arXiv:1506.07135] [INSPIRE].

[2] S.O. Rice, *Mathematical analysis of random noise*, *Bell Syst. Techn. J.* **23** (1944) 282.

[3] S. . Rice, *Mathematical analysis of random noise*, *Bell Syst. Techn. J.* **24** (1945) 46,

[4] J.M. Bardeen et al., *The statistics of peaks of Gaussian random fields*, *Aastrophys. J.* **304** (1986) 15.

[5] J.R. Bond and G. Efstathiou, *The statistics of cosmic background radiation fluctuations*, *Mon. Not. Roy. Astron. Soc.* **226** (1987) 655.

[6] WMAP collaboration, *Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: cosmological interpretation*, *Astrophys. J. Suppl.* **192** (2011) 18 [arXiv:1001.4538] [INSPIRE].

[7] A. Marcos-Caballero, R. Fernández-Cobos, E. Martínez-González and P. Vielva, *The shape of CMB temperature and polarization peaks on the sphere*, *JCAP* **04** (2016) 058 [arXiv:1512.07412] [INSPIRE].

[8] PLANCK collaboration, *Planck 2013 results. I. Overview of products and scientific results*, *Astron. Astrophys.* **571** (2014) A1 [arXiv:1303.5062] [INSPIRE].

[9] WMAP collaboration, *Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: cosmological parameter results*, *Astrophys. J. Suppl.* **208** (2013) 19 [arXiv:1212.5226] [INSPIRE].

[10] PLANCK collaboration, *Planck intermediate results. XLIX. Parity-violation constraints from polarization data*, *Astron. Astrophys.* **596** (2016) A110 [arXiv:1605.08633] [INSPIRE].

[11] W. Hu and M.J. White, *A CMB polarization primer*, *New Astron.* **2** (1997) 323 [astro-ph/9706147] [INSPIRE].

[12] P.D. Naselsky and D.I. Novikov, *General statistical properties of the CMB polarization field*, *Astrophys. J.* **507** (1998) 31 [astro-ph/9801285] [INSPIRE].

[13] O. Dore, S. Colombi and F.R. Bouchet, *Probing CMB non-Gaussianity using local curvature*, *Mon. Not. Roy. Astron. Soc.* **344** (2003) 905 [astro-ph/0202135] [INSPIRE].

[14] D. Huterer and T. Vachaspati, *Distribution of singularities in the cosmic microwave background polarization*, *Phys. Rev.* **D 72** (2005) 043004 [astro-ph/0405474] [INSPIRE].

[15] PLANCK collaboration, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, arXiv:1807.06205 [INSPIRE].

[16] PLANCK collaboration, *Planck 2018 results. II. Low frequency instrument data processing*, arXiv:1807.06206 [INSPIRE].

[17] PLANCK collaboration, *Planck 2018 results. III. High frequency instrument data processing and frequency maps*, arXiv:1807.06207 [INSPIRE].

[18] PLANCK collaboration, *Planck 2018 results. IV. Diffuse component separation*, arXiv:1807.06208 [INSPIRE].

[19] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [astro-ph/9911177] [INSPIRE].

[20] K.M. Gorski et al., *HEALPix — A framework for high resolution discretization and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [astro-ph/0409513] [INSPIRE].

[21] M. Masi, *On compressive radial tidal forces*, *Am. J. Phys.* **75** (2007) 116 [arXiv:0705.3747] [INSPIRE].

[22] M. Kamionkowski, A. Kosowsky and A. Stebbins, *Statistics of cosmic microwave background polarization*, *Phys. Rev.* **D 55** (1997) 7368 [astro-ph/9611125] [INSPIRE].

[23] S. Dodelson, *Modern cosmology*, Academic Press, U.S.A. (2003).

[24] L.F. Abbott and R.K. Schaefer, *A general, gauge-invariant analysis of the cosmic microwave anisotropy*, *Astrophys. J.* **308** ((1986) 546.

[25] J. Arango and O. Perdomo, *Morse theory for analytic functions on surfaces*, *J. Geom.* **84** (2006) 13.

[26] S.M. Carroll, G.B. Field and R. Jackiw, *Limits on a Lorentz- and parity-violating modification of electrodynamics*, *Phys. Rev.* **D 41** (1990) 1231.

[27] D. Contreras, P. Boubel and D. Scott, *Constraints on direction-dependent cosmic birefringence from Planck polarization data*, *JCAP* **12** (2017) 046 [arXiv:1705.06387] [INSPIRE].

[28] D. Scott, D. Contreras, A. Narimani and Y.-Z. Ma, *The information content of cosmic microwave background anisotropies*, *JCAP* **06** (2016) 046 [arXiv:1603.03550] [INSPIRE].