

A framework for testing isotropy with the cosmic microwave background

Daniela Saadeh,¹★ Stephen M. Feeney,² Andrew Pontzen,¹ Hiranya V. Peiris¹
and Jason D. McEwen³

¹Department of Physics and Astronomy, University College London, London WC1E 6BT, UK

²Department of Physics, Blackett Laboratory, Imperial College London, Prince Consort Road, London SW7 2AZ, UK

³Mullard Space Science Laboratory (MSSL), University College London, Surrey RH5 6NT, UK

Accepted 2016 July 14. Received 2016 July 1; in original form 2016 April 14

ABSTRACT

We present a new framework for testing the isotropy of the Universe using cosmic microwave background data, building on the nested-sampling `ANICOSMO` code. Uniquely, we are able to constrain the scalar, vector and tensor degrees of freedom alike; previous studies only considered the vector mode (linked to vorticity). We employ Bianchi type VII_h cosmologies to model the anisotropic Universe, from which other types may be obtained by taking suitable limits. In a separate development, we improve the statistical analysis by including the effect of Bianchi power in the high- ℓ , as well as the low- ℓ , likelihood. To understand the effect of all these changes, we apply our new techniques to *Wilkinson Microwave Anisotropy Probe* data. We find no evidence for anisotropy, constraining shear in the vector mode to $(\sigma_V/H)_0 < 1.7 \times 10^{-10}$ (95 per cent confidence level). For the first time, we place limits on the tensor mode; unlike other modes, the tensor shear can grow from a near-isotropic early Universe. The limit on this type of shear is $(\sigma_{T,\text{reg}}/H)_0 < 2.4 \times 10^{-7}$ (95 per cent confidence level).

Key words: gravitation – cosmic background radiation – cosmology: miscellaneous – early Universe.

1 INTRODUCTION

The standard cosmological model assumes that space is homogeneous and isotropic on large scales. Observational data, particularly measurements of the cosmic microwave background (CMB), allow this assumption to be tested quantitatively. When homogeneity and isotropy are assumed from the outset, the cosmological solutions to Einstein’s field equations are described by Friedmann–Lemaître–Robertson–Walker (FLRW) metrics. By relaxing the requirement for isotropy, one is instead led to Bianchi metrics (Bianchi 1898; Ellis & MacCallum 1969). The basis for testing isotropy is therefore to consider which of these backgrounds better describes our Universe. In the limit that departure from isotropy is small, the observed fluctuations in the CMB are well approximated by the sum of a deterministic Bianchi template and the stochastic contribution from the inflationary Λ cold dark matter (Λ CDM) cosmological model.

Testing isotropy has received considerable attention since *Wilkinson Microwave Anisotropy Probe* (WMAP; Bennett et al. 2003) full-sky maps became available. Jaffe et al. (2005) found a correlation between WMAP temperature data and a pattern induced by the Bianchi VII_h model; employing the new background also

improved the fit to the temperature power spectrum (Jaffe et al. 2005; McEwen et al. 2006; Bridges et al. 2007), suggesting that the Universe indeed departs from isotropy. However, the same authors pointed out that the best-fitting Bianchi template is characterized by cosmological parameters (for example a large negative curvature) that are inconsistent with other available observations.

McEwen et al. (2013) subsequently introduced `ANICOSMO`, a tool for robust statistical analysis of the effects of an anisotropic background on the CMB. This code has been employed for a number of studies of the type of Bianchi models considered by Jaffe et al. (2005), most recently using *Planck* data (Planck Collaboration XXVI 2014; Planck Collaboration XVIII 2015). In these analyses, in which the parameters for the background and stochastic components were required to be mutually consistent, no preference was found for anisotropy. However, the tests only took into account two out of a total five degrees of freedom of the Bianchi anisotropy, and thus did not allow an upper limit to be placed on anisotropy in general.

Pontzen & Challinor (2011) presented a systematic linearization, from which the most general anisotropies that respect homogeneity can be expressed as a set of non-interacting modes on an isotropic background. To date, no statistical analysis is available for the additional degrees of freedom highlighted by this analysis, so that a true test of universal isotropy is lacking. Some of the previously unconstrained modes are expected to be the most compatible with probes

*E-mail: daniela.saadeh.13@ucl.ac.uk

such as nucleosynthesis and CMB polarization, both of which are particularly sensitive to early-time anisotropy (Pontzen 2009).

In this work, we create a framework for testing all modes simultaneously and validate it in the VII_h case. This requires a number of changes to the previous approaches. We use the *WMAP* 9-yr (Bennett et al. 2013) data set previously considered by McEwen et al. (2013) so that we can isolate the impact of these changes on the posterior distributions and the Bayesian evidence ratios with respect to Λ CDM. The *Planck* analyses (Planck Collaboration XXVI 2014; Planck Collaboration XVIII 2015) did not include a high- ℓ likelihood, and hence do not provide a full setting for comparison. In a future paper, we will extend the framework further to apply the new methods to *Planck* data, including polarization which is expected to be highly constraining (Pontzen & Challinor 2007).

The paper is structured as follows. We introduce our method in Sections 2 and 3, where we respectively discuss how we model anisotropy and the statistical analysis we perform to look for its signatures. In Section 4, we present results from our analysis of *WMAP* 9-yr data, and conclude in Section 5.

2 IMPLEMENTING ANISOTROPY

In this section, we describe the framework underlying *ABSOLVE* (Anisotropic Boltzmann Solver), a Boltzmann-hierarchy code we have developed to compute the deterministic CMB temperature and polarization perturbations induced by a Bianchi background. We first discuss how anisotropy is modelled within the code in Section 2.1, before presenting the computational details of the calculation in Section 2.2.

2.1 The anisotropy degrees of freedom

ABSOLVE can compute the anisotropy-induced perturbations for Bianchi types¹ I, V, VII_0 and VII_h . These cosmologies are sufficient to take into account all the open and flat Bianchi models that are close to isotropy, and therefore compatible with observations (Pontzen & Challinor 2011). The closed Bianchi type IX only induces power at $\ell = 2$, making it difficult to constrain: it is therefore currently not included.² In this paper, we focus on the Bianchi type VII_h , together with its flat limit VII_0 .

Anisotropy is characterized quantitatively by means of the shear tensor $\sigma_{\mu\nu}$, which describes the deformation experienced by fluid elements in the Universe during anisotropic expansion; the shear scalar σ is defined by $\sigma^2 \equiv \sigma_{\mu\nu}\sigma^{\mu\nu}/2$. To model the evolution of shear, we follow Pontzen & Challinor (2011) in decomposing the anisotropic expansion into a set of five non-interacting modes on an underlying isotropic (FLRW) background. Each mode corresponds to a degree of freedom in the shear tensor. The deterministic perturbations induced in the CMB transform like scalars (1 dof), vectors (2 dof) or tensors (2 dof) under rotations around a preferred axis of the Bianchi symmetry, and are labelled accordingly throughout this work. In particular, we write the magnitude of the shear associated with each component as σ_S , σ_V and σ_T , respectively, by considering their evolution independently; when all modes are combined one may show that the total shear obeys $\sigma^2 = \sigma_S^2 + \sigma_V^2 + \sigma_T^2$. Currently,

published constraints on departures from isotropy use the approach of Barrow, Juszkiewicz & Sonoda (1985) and consequently consider only σ_V , i.e. the two Bianchi VII_h vector modes (e.g. Jaffe et al. 2005; McEwen et al. 2013; Planck Collaboration XXVI 2014; Planck Collaboration XVIII 2015).

The evolution of anisotropy is dictated by the Einstein equations. We assume the contents of the Universe can be described as a sum of perfect fluids corresponding to radiation, matter and dark energy. Under these assumptions, scalar and vector modes exhibit a fast decay ($\sigma_V, \sigma_S \propto a^{-3}$, where a is the scalefactor), linking small levels of anisotropy today to larger levels at recombination. Extrapolating backwards to the big bang, eventually the shear becomes comparable to the Hubble parameter, at which point the linear decomposition of Pontzen & Challinor (2011) ceases to apply. We refer to such behaviour in this paper as ‘irregular’, since it cannot be immediately reconciled with a near-isotropic early-Universe implied by the inflationary scenario. The steep decay also gives rise to a high degree of polarization in the CMB (Pontzen & Challinor 2007).

Solutions for the tensor modes can behave in this irregular way; but there is also a growing solution (Collins & Hawking 1973) allowing observable anisotropy to emerge from a near-isotropic early Universe. We term this behaviour ‘regular’, since it can more easily be fitted into the modern cosmological paradigm, although fine-tuning of inflation is still required for shear to reach an observable amplitude at the present day. Only the tensor degrees of freedom, which have not been tested previously, can exhibit this behaviour. Scalar and vector modes also possess a second solution, but it can be removed by a coordinate transformation and therefore has no physical effect.

The maps in Fig. 1 show, from upper left to lower right, CMB anisotropies imprinted by scalar, irregular tensor, vector and regular tensor modes.³ The inset panels show the polarization. For the case of the vector and regular tensor modes, the Bianchi power spectra are plotted underneath as solid and dashed curves, respectively. Black, red and blue lines show temperature, E -mode and B -mode polarization power spectra, respectively.

The magnitude of the shear in these cases has been chosen to produce a similar rms temperature anisotropy amplitude of approximately 75 μ K. For the case of the vector modes a present-day shear (normalized to the isotropic Hubble expansion rate H_0 to form a dimensionless quantity) of $(\sigma_V/H_0)_0 \simeq 10^{-9}$ is sufficient. However, for the regular tensor modes, this amplitude must be considerably higher, $(\sigma_{T,\text{reg}}/H_0)_0 \simeq 5 \times 10^{-6}$, because the steep scaling with redshift is absent. We can therefore immediately anticipate that constraints on present-day anisotropy in regular tensor modes will be considerably weaker than the corresponding constraints for the vector modes.

Having decomposed the shear into $(\sigma_S/H_0)_0$, $(\sigma_V/H_0)_0$ and $(\sigma_T/H_0)_0$, there are a number of further degrees of freedom to be considered. First, the specific initial conditions for the tensor mode are placed into either the decaying, irregular solution (in which case we will refer to $\sigma_{T,\text{irr}}$) or the growing, regular solution ($\sigma_{T,\text{reg}}$). Additionally, the Bianchi morphology is set by the matter and dark energy densities today $\{\Omega_m, \Omega_\Lambda\}$ and the parameter x , which regulates the scale over which the shear principal axes rotate in the vector and tensor modes (see Collins & Hawking 1973; Barrow et al. 1985 for a formal definition). The orientation relative to the Galaxy is determined by three Euler angles, labelled α , β and γ .

¹ Bianchi ‘types’ are subclasses into which Bianchi cosmologies are divided, grouping the different, inequivalent, ways for a 3-space to be homogeneous. Accessible reviews on Bianchi models may be found in Ellis, Maartens & MacCallum (2012) and Wainwright & Ellis (1997); also see Pontzen & Challinor (2011) for a different approach.

² Bianchi I also only induces power in the quadrupole, but it arises as the natural flat limit of Bianchi V.

³ For animations of the Bianchi pattern for varying cosmological parameters, see <http://zenodo.org/record/48654>.

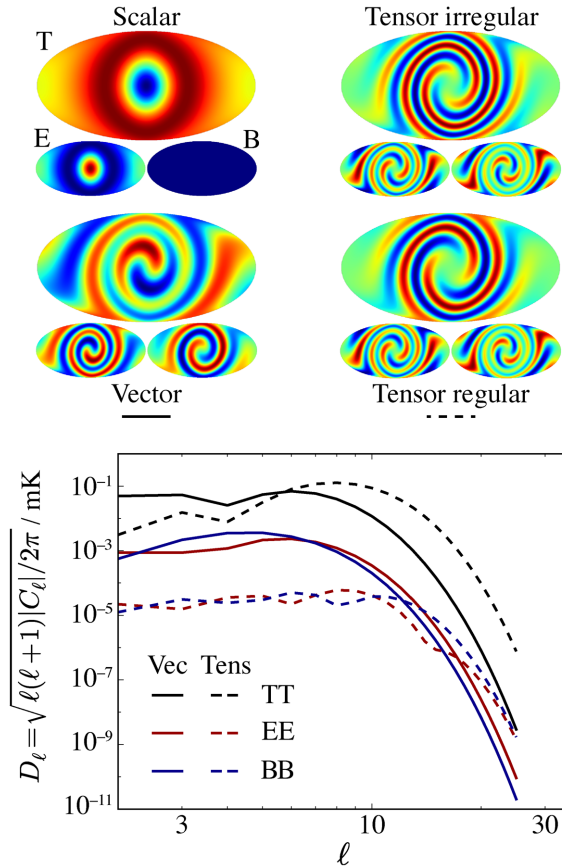


Figure 1. Maps: example scalar, vector, regular and irregular tensor patterns induced in the CMB temperature (upper panels) and polarization (lower panels, E - and B -mode to the left and right) for $(\Omega_m, \Omega_\Lambda, x) = (0.27, 0.7, 0.62)$. These maps were produced with the `ABSOLVE` code developed for this analysis. Plot: vector and regular tensor power spectra for $(\Omega_m, \Omega_\Lambda, x) = (0.27, 0.7, 0.62)$ and $(\sigma_V/H)_0 = 1 \times 10^{-9}$ (solid line), $(\sigma_{T,\text{reg}}/H)_0 = 5 \times 10^{-6}$ (dashed line): significantly smaller values of the vector amplitude today lead to comparable temperature signals, and larger polarization signals, compared to that of the tensors.

There is then a further freedom to rotate the tensor shear contribution relative to that of the vectors by an angle γ_{VT} ; however, we will not consider vector and tensor modes simultaneously in this work, and this freedom therefore does not enter. Finally, Bianchi vector and tensor modes have a handedness, so for these models parity must also be specified. For reference, all the Bianchi parameters employed in this work are summarized in Table 1.

2.2 Computational aspects

The deterministic contribution to the CMB is obtained by computing the Boltzmann equation for photons in the context of a given Bianchi background: the complete set of equations can be found in Pontzen & Challinor (2007), except for the shear evolution, derived by Pontzen & Challinor (2011). Our code `ABSOLVE` is written in `PYTHON` and `CYTHON`. Run times vary considerably across the parameter space, but typically take a few seconds on one 2.6 GHz core.

The code starts by computing the recombination history using `RECFAST 1.3`⁴ (Seager, Sasselov & Scott 1999, 2000). The Boltzmann

Table 1. Summary of the parameters defining the Bianchi morphology, amplitude and orientation. The amplitudes are expressed in terms of the shear scalar normalized to the Hubble parameter. In addition to the parameters tabulated below, there is the discrete choice of handedness.

Morphology	
Ω_m	Matter density
Ω_Λ	Dark energy density
x	Rotation scale of shear principal axes
Signal amplitude	
$(\sigma_S/H)_0$	Amplitude of scalar modes
$(\sigma_V/H)_0$	Amplitude of vector modes
$(\sigma_{T,\text{reg}}/H)_0$	Amplitude of regular tensor modes
$(\sigma_{T,\text{irr}}/H)_0$	Amplitude of irregular tensor modes
Orientation	
α, β, γ	Euler angles defining the orientation of the Bianchi pattern

integration starts at redshift $z_{\text{start}} = 1500$ with zero initial Bianchi power; the power quickly builds from the shear-induced temperature quadrupole and scatters into the E -mode polarization quadrupole. Anisotropies are subsequently advected to smaller scales and partially converted into B -mode polarization due to free-streaming effects after recombination. We have checked for one example model⁵ that the Bianchi pattern is unchanged if integration is started significantly earlier, at $z_{\text{start}} = 1800$ or $z_{\text{start}} = 2000$. This reflects the structure of the Boltzmann equation, where anisotropy is built by the shear tensor but damped by Thomson scattering through a viscous-friction term proportional to anisotropy. A limit equilibrium exists in this setting that is reached for any initial redshift sufficiently above that of recombination. A possible future improvement to `ABSOLVE` would be to include a more refined treatment of z_{start} that fixes its value based on the model parameters (in particular, taking into account the sensitivity of recombination to the baryon and dark matter physical densities).

The maximum multipole used to characterize the Bianchi pattern, ℓ_{max} , must be chosen carefully to avoid missing small-scale power. We implemented a test that compares the power around a trial ℓ_{max} with the power at the quadrupole; if this test fails, integration is repeated up to a much larger multipole. Further details of this procedure and its importance are given in Section 3.4; typical values are $\ell_{\text{max}} = 200$ as an initial guess and $\ell_{\text{max}} = 1000$ subsequently.

We must also avoid edge effects that can propagate errors from the hierarchy truncation. We found empirically that we needed to extend the calculated hierarchy to $\ell_{\text{trunc}} = \ell_{\text{max}} + 50$ to obtain converged results at ℓ_{max} . To prevent instabilities from developing, we additionally apply a Fermi–Dirac damping of the form

$$D(\ell) = \frac{1}{\exp[\ell - (\ell_{\text{trunc}} - 10)] + 1} \quad (1)$$

to the hierarchy,⁶ which damps the power at $\ell_{\text{trunc}} - 10 = \ell_{\text{max}} + 40$ to prevent advected power reflecting at the unphysical boundary.

⁵ $\{\Omega_b h^2 = 0.022, \Omega_c h^2 = 0.11, x = 0.5, \Omega_m = 0.27, \Omega_\Lambda = 0.7, (\sigma_S/H)_0 = 1 \times 10^{-9}, (\sigma_V/H)_0 = 1 \times 10^{-9}, (\sigma_{T,\text{reg}}/H)_0 = 1 \times 10^{-6}, (\sigma_{T,\text{irr}}/H)_0 = 1 \times 10^{-7}, \gamma_{VT} = 0\}$

⁶ Technically, we implement this by multiplying the advection coefficients in equations (48), (52) and (53) of Pontzen & Challinor (2007) by $D(\ell)$. The chosen form for $D(\ell)$ is heuristically motivated, rather than following from any analytic approximation: we verified that it leads to convergence in all the examined cases.

⁴ <http://www.astro.ubc.ca/people/scott/recfast.html>

Once our code has produced output in harmonic space, we use `HEALPIX`⁷ (Górski et al. 2005) to perform the required Euler rotations and to produce maps when needed.

3 STATISTICAL FRAMEWORK

In this section, we introduce the statistical framework required for constraining all possible modes of background anisotropy using CMB data. We start with an overview in Section 3.1 before discussing the choice of priors in Section 3.2. Section 3.3 presents an illustrative analysis as a guide to the interpretation of our results. Finally, in Section 3.4 we show how the new analysis improves constraints on the background anisotropy by employing information from smaller scales (higher ℓ) than considered in previous analyses.

3.1 Analysis overview

In the models under study, stochastic Λ CDM fluctuations are superimposed on a Bianchi background. In this setting, the observed CMB data, \mathbf{d} , are assumed to consist of a stochastic, Gaussian Λ CDM component, \mathbf{s} , a deterministic Bianchi component, \mathbf{b} , and Gaussian instrumental noise, \mathbf{n} :

$$\mathbf{d} = \mathbf{s}(\Theta_{\Lambda\text{CDM}}) + \mathbf{b}(\Theta_{\text{B}}) + \mathbf{n}. \quad (2)$$

Here, $\Theta_{\Lambda\text{CDM}} = \{\Omega_{\text{b}}h^2, \Omega_{\text{c}}h^2, \Omega_{\Lambda}, \Omega_{\text{K}}, n_{\text{s}}, A_{\text{s}}, \tau\}$ is a vector of parameters from the standard Λ CDM framework: baryon and dark matter physical densities $\Omega_{\text{b}}h^2$ and $\Omega_{\text{c}}h^2$ (where $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$), dark energy and curvature densities Ω_{Λ} and Ω_{K} , scalar spectral index and power amplitude n_{s} and A_{s} , and optical depth to reionization τ . Θ_{B} is a vector of the Bianchi parameters summarized in Table 1 (see Section 2.1). The partial overlap between Θ_{B} and $\Theta_{\Lambda\text{CDM}}$ is discussed in more detail in Section 3.2.

The likelihood function $P(\mathbf{d} | (\Theta_{\text{B}}, \Theta_{\Lambda\text{CDM}}), \mathbf{M}) = \mathcal{L}(\Theta_{\text{B}}, \Theta_{\Lambda\text{CDM}})$ takes the form of a Gaussian with mean \mathbf{b} set by the deterministic Bianchi template and covariance matrix \mathbf{C} set by the stochastic Λ CDM component and instrumental noise properties

$$\mathcal{L}(\Theta_{\text{B}}, \Theta_{\Lambda\text{CDM}}) \propto \frac{1}{\sqrt{|\mathbf{C}(\Theta_{\Lambda\text{CDM}})|}} \exp[-\chi^2(\Theta_{\Lambda\text{CDM}}, \Theta_{\text{B}})/2], \quad (3)$$

where

$$\chi^2(\Theta_{\Lambda\text{CDM}}, \Theta_{\text{B}}) = [\mathbf{d} - \mathbf{b}(\Theta_{\text{B}})]^{\dagger} \mathbf{C}(\Theta_{\Lambda\text{CDM}})^{-1} [\mathbf{d} - \mathbf{b}(\Theta_{\text{B}})]. \quad (4)$$

The data and Bianchi template may be expressed in either pixel or harmonic space.

Limits on anisotropy can be considered as either a parameter-estimation or a model-comparison problem. In the former case, the product of the prior and likelihood gives the posterior probability density for different shear amplitudes. In the latter case, one must compute and compare each model's Bayesian evidence. The Bayesian evidence is defined as the probability of obtaining the data given a model \mathbf{M} , integrating over all of \mathbf{M} 's parameters Θ , i.e. the marginal likelihood

$$E = P(\mathbf{d} | \mathbf{M}) = \int d\Theta P(\mathbf{d} | \Theta, \mathbf{M}) P(\Theta | \mathbf{M}). \quad (5)$$

Here, $P(\Theta | \mathbf{M})$ is the a priori probability for Θ and $P(\mathbf{d} | \Theta, \mathbf{M})$ is the likelihood function. The evidence can be used to estimate the relative probability of two models \mathbf{M}_1 and \mathbf{M}_2 , given available data

$$\frac{P(\mathbf{M}_1 | \mathbf{d})}{P(\mathbf{M}_2 | \mathbf{d})} = \frac{P(\mathbf{d} | \mathbf{M}_1) P(\mathbf{M}_1)}{P(\mathbf{d} | \mathbf{M}_2) P(\mathbf{M}_2)} = \frac{E_1 P(\mathbf{M}_1)}{E_2 P(\mathbf{M}_2)}. \quad (6)$$

If, as assumed in this work, no a priori knowledge is available that favours one model over the other (i.e. $P(\mathbf{M}_1) = P(\mathbf{M}_2)$), then the probability ratio equals the evidence ratio: $P(\mathbf{M}_1 | \mathbf{d})/P(\mathbf{M}_2 | \mathbf{d}) = E_1/E_2$. The log-Bayes factor $\ln(E_1/E_2)$ will be used in Section 4 to assess the degree by which one model is favoured over the other.

In order to evaluate the probability distributions described above, we have integrated `ABSOLVE` into the `ANICOSMO` package (McEwen et al. 2013). `ANICOSMO` uses the nested sampler `MULTINEST`⁸ (Feroz & Hobson 2008; Feroz, Hobson & Bridges 2009; Feroz et al. 2013) to explore the parameter space of each model and hence efficiently calculate its Bayesian evidence. At each sampled point, the theoretical mean of the CMB data is evaluated using `ABSOLVE`, and the covariance (set by the stochastic Λ CDM fluctuations) is calculated using power spectra produced by `CAMB` (Lewis, Challinor & Lasenby 2000).⁹

When applied to masked data, the likelihood (3) is difficult to evaluate at high ℓ since the mask becomes problematic in harmonic space while the covariance becomes problematic in pixel space. It is therefore necessary to adopt an approximate power-spectrum-based likelihood at high ℓ (typically $\ell > 32$; e.g. Hinshaw et al. 2007; Page et al. 2007). In previous versions of `ANICOSMO`, including those used by the Planck Collaboration (Planck Collaboration XXVI 2014; Planck Collaboration XVIII 2015), only the low- ℓ part of the likelihood has been modified to take into account the Bianchi template. This results in a good estimate of the overall likelihood provided that the power in the Bianchi component is negligible compared to the stochastic power in the high- ℓ modes. However, Bianchi models have two physical scales (a spiralling and curvature scale controlled by x and Ω_{K} , respectively); when either of these is sufficiently small relative to the horizon, high- ℓ power can be significant; see Section 3.4.

For this reason, we complement the low- ℓ likelihood (3) with an approximate modification to the high- ℓ likelihood using the summed contributions of the Bianchi and Λ CDM power spectra. For fluctuations around an anisotropic Bianchi background, the power spectrum does not provide lossless data compression, but in the limit that the Bianchi signal is subdominant relative to the Λ CDM component, the approximation is valid. See Appendix A for further details.

In test runs on simulated maps (see Section 3.3), full-sky information is available at all multipoles. In these cases, we use the likelihood function in equation (3) up to $\ell = 400$, without an additional high- ℓ component. The absence of a high- ℓ likelihood constraining the temperature power spectrum in the damping tail makes the recovered constraints on test Λ CDM parameters less stringent. This is acceptable, since we only need to verify our ability to recover the Bianchi template parameters.

3.2 Models

As described in Section 3.1, parameters related to density appear in both Θ_{B} and $\Theta_{\Lambda\text{CDM}}$. A physically meaningful analysis must have

⁸ <http://ccpforge.cse.rl.ac.uk/gf/project/multinest/>

⁹ <http://camb.info/>

⁷ <http://healpix.sourceforge.net>

Table 2. Models tested in this work. The ^(d) superscript, where present, refers to the phenomenological ‘decoupled’ models; physical ‘coupled’ models are indicated with ^(c). S, V, T_{reg} and T_{irr} refer, respectively, to Bianchi scalar, vector, regular and irregular tensor modes.

Notation	Model
Λ CDM	Pure Λ CDM, no Bianchi component
S ^(c)	Bianchi VII _h /VII ₀ scalar modes, ‘coupled’ model
S ^(d)	Bianchi VII _h /VII ₀ scalar modes, ‘decoupled’ model
V ^(c)	Bianchi VII _h /VII ₀ vector modes, ‘coupled’ model
V ^(d)	Bianchi VII _h /VII ₀ vector modes, ‘decoupled’ model
T _{reg} ^(c)	Bianchi VII _h /VII ₀ regular tensor modes, ‘coupled’ model
T _{reg} ^(d)	Bianchi VII _h /VII ₀ regular tensor modes, ‘decoupled’ model
T _{irr} ^(c)	Bianchi VII _h /VII ₀ irregular tensor modes, ‘coupled’ model
T _{irr} ^(d)	Bianchi VII _h /VII ₀ irregular tensor modes, ‘decoupled’ model

Table 3. Priors for all the model parameters, along with the search to which they are applied. Where no indication is given, the stated prior is applied in all cases: B^(c), B^(d) and pure Λ CDM. See also Section 3.2.

Parameter	Prior range	Prior type	Models
$\Omega_b h^2$	[0.005, 0.05]	Uniform	
$\Omega_c h^2$	[0.05, 0.3]	Uniform	
Ω_Λ	[0, 0.99]	Uniform	
Ω_K	[10 ⁻⁵ , 0.5]	Uniform	B ^(c)
Ω_K	0	Fixed	B ^(d) and Λ CDM
n_s	[0.9, 1.05]	Uniform	
A_s	[1, 5] × 10 ⁻⁹	log-uniform	
τ	[0.082, 0.092]	Uniform	
Ω_m	[0, 0.99]	Uniform	B ^(d)
Ω_Λ	[0, 0.99]	Uniform	B ^(d)
x	–	–	S
x	[0.05, 2]	Uniform	V, T _{reg} and T _{irr}
$(\sigma_S/H)_0$	[–10 ⁻⁸ , 10 ⁻⁸]	Uniform	S
$(\sigma_V/H)_0$	[10 ⁻¹² , 10 ⁻⁸]	log-uniform	V
$(\sigma_{T,\text{reg}}/H)_0$	[10 ⁻¹² , 10 ⁻⁴]	log-uniform	T _{reg}
$(\sigma_{T,\text{irr}}/H)_0$	[10 ⁻¹² , 10 ⁻⁴]	log-uniform	T _{irr}
α	[0°, 360°]	Uniform	B ^(c) and B ^(d)
β	[0°, 180°]	Sine-uniform	B ^(c) and B ^(d)
γ	–	–	S
γ	[0°, 360°]	Uniform	V
γ	[0°, 180°]	Uniform	T _{reg} and T _{irr}

self-consistent matter and dark energy densities Ω_m and Ω_Λ when calculating the background and stochastic perturbation contributions: such analyses are referred to as Bianchi ‘coupled’ runs in the following, or B^(c). Our upper limits on anisotropy will be derived in these settings. However, to connect with the early analyses that found evidence in favour of Bianchi cosmologies, we also test phenomenological models in which the two components are allowed to vary independently: these models are labelled as ‘decoupled’, or B^(d).

Table 2 lists the models considered in this work. We use a Bianchi VII_h model and allow the curvature to approach zero such that the VII₀ models are also naturally included. We separately test for the scalar (S), vector (V), regular (T_{reg}) and irregular (T_{irr}) tensor degrees of freedom in turn; the priors that we adopt for the Bianchi and Λ CDM parameters are listed in Table 3. For $(\sigma_V/H)_0$, $(\sigma_{T,\text{reg}}/H)_0$ and $(\sigma_{T,\text{irr}}/H)_0$, we adopt log-uniform priors so as to avoid setting a preferred scale for these parameters. This choice is not possible for $(\sigma_S/H)_0$, which can take negative values in our parametrization: we adopt a uniform prior in this case.




Figure 2. Comparison of the prior on $(\sigma_V/H)_0$ employed in this work (log-uniform) with that implied by McEwen et al. (2013) in which a uniform prior is taken over $(\omega/H)_0$, x and Ω_K . The transformation between the spaces is described by equation (7). The approximate range is comparable, but the prior of McEwen et al. (2013) places considerable added emphasis on shear values around $(\sigma_V/H)_0 \simeq 10^{-9}$.

The expanded set of modes requires us to use a general parametrization that uses shear, rather than vorticity, to control the amplitude of anisotropy. Consequently our priors cannot be made identical to those in previous work. The link between expansion-normalized shear $(\sigma_V/H)_0$ and vorticity $(\omega/H)_0$ is provided by the time-space component of the Einstein equations, which gives (Barrow et al. 1985)

$$\left(\frac{\omega}{H}\right)_0 = \frac{\sqrt{1+x^2\Omega_K}\sqrt{1+9x^2\Omega_K}}{6x^2(1-\Omega_K)} \left(\frac{\sigma_V}{H}\right)_0. \quad (7)$$

In McEwen et al. (2013), the prior on $(\omega/H)_0$ is uniform over the range [0, 10⁻⁹]. However, the appearance of parameters x and Ω_K in the relationship (7) shows that the implied prior on $(\sigma_V/H)_0$ is not uniform. By sampling points from the prior of McEwen et al. (2013) and applying the mapping (7), we established that the marginalized prior on $\log(\sigma_V/H)_0$ in previous work is peaked around $(\sigma_V/H)_0 \simeq 10^{-9}$; see Fig. 2. By contrast we assign equal a priori probability to all the scales in the range [10⁻¹², 10⁻⁸]. As an example of the Bianchi signal strengths that are consequently included in the search, the lower and upper limits of our prior range correspond to rms temperature fluctuations of, respectively, 0.05 and 500 μ K at $x = 0.3$.

The allowed ranges for $(\sigma_{T,\text{reg}}/H)_0$ and $(\sigma_{T,\text{irr}}/H)_0$ need to be significantly wider than that on $(\sigma_S/H)_0$ and $(\sigma_V/H)_0$ because similar values of S or V and T_{reg} or T_{irr} at recombination correspond to considerably larger values for the tensor shear amplitude today, especially in the case of T_{reg}: for the tensor cases, we allow values as large as 10⁻⁴. Scalar modes possess rotational symmetry around a preferred axis, which makes them insensitive to the parameters x and γ ; they are consequently held fixed in the S runs. Similarly, because of the spin-2 tensor symmetry, the γ angle is only required to vary in the range [0°, 180°] in T runs, whereas it takes the full range [0°, 360°] in V cases. The prior range for the optical depth to reionization, τ , was chosen to match that of McEwen et al. (2013) to ease comparison.




Figure 3. An illustration of the strong geometric degeneracy in the Bianchi parameter space. The triangle plot shows the recovered posterior distribution for a mock map containing a stochastic Λ CDM and underlying deterministic Bianchi pattern with $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$, $x = 0.7$ and $(\sigma_v/H)_0 = 0.7 \times 10^{-9}$. These input parameters are indicated by grey lines. The parameters are recovered but are strongly broadened by the geometric degeneracy arising from the projection of the Bianchi spiral on to the surface of last scattering. The two inset maps correspond to models near the ends of the recovered degeneracy as indicated by the dot (upper map) and triangle (lower map). The overall orientation of the pattern is constrained as seen in the sharply peaked marginalized posterior distributions of the Euler angles in the bottom row.

3.3 Illustrative analysis

We applied the analysis outlined in Section 3.1 to several simulated maps of the CMB sky containing a Bianchi signal to check that the input parameters were recovered correctly. The mocks mimic stochastic Λ CDM fluctuations on top of a Bianchi background and were generated as follows: we computed the temperature power spectrum given a set of cosmological parameters through `CAMB` and obtained a realization of the corresponding Gaussian random field; we then added the resulting fluctuations linearly to a map containing a Bianchi template. The deterministic Bianchi contribution was calculated using the code from Pontzen (2009) as a blind test of the `ABSOLVE` implementation.¹⁰ We applied a Gaussian beam with full width at half-maximum (FWHM) of 1° to the maps and assumed instrumental noise to be negligible on the relevant scales. When applying our modified version of `ANICOSMO` to these mocks, we employed a likelihood function of the form in equation (3) up to $\ell = 400$, with no supplementary high- ℓ likelihood.

Fig. 3 shows the recovered constraints in an example decoupled test run on a mock CMB map containing a Bianchi vector

$V^{(d)}$ signal. The recovered posterior is consistent with the input parameters, which are, respectively, $\Omega_m = 0.31$, $\Omega_\Lambda = 0.69$, $x = 0.7$ and $(\sigma_v/H)_0 = 0.7 \times 10^{-9}$. However, there is a strong degeneracy in the $\{\Omega_m, \Omega_\Lambda, x, \sigma/H\}$ dimensions; this reflects how the angular scale of the recovered pattern is approximately set by the Bianchi spiralling scale projected on to the last scattering surface. The inset temperature maps show how the input signal is mimicked extremely well along this degeneracy line. As a consequence, the marginalized posteriors on these parameters are broad. The orientation of the pattern, defined by the Euler angles, is recovered well, with sharply peaked Gaussians around the input values, with the exception of γ which is biased by 2.3σ . We verified that the cause of this bias was a slight net rotation between the Pontzen (2009) maps which form the basis of the test and the `ABSOLVE` maps for the same input values. Given the more careful numerical choices described above, we believe the `ABSOLVE` results to be more robust.

Improved constraints on the Bianchi parameters can be obtained by employing complementary information to break the degeneracy: this is the case in $B^{(c)}$ models, where Λ CDM fluctuations strongly limit the range of allowed matter and dark energy densities $\{\Omega_m, \Omega_\Lambda\}$, thereby also tightening the constraints on x and σ/H . Taking into account the CMB polarization in addition to temperature would also partially break the model degeneracy.

Our framework was additionally tested against mock $S^{(d)}$, $V^{(c)}$, $V^{(d)}$, $T_{\text{reg}}^{(c)}$ and $T_{\text{irr}}^{(d)}$ maps.

3.4 The importance of small scales

As discussed in Section 3.1, one novel feature of our improved framework is that it includes the effect of the Bianchi temperature fluctuations in the high- ℓ part of the likelihood. While some Bianchi models induce power that decays rapidly with ℓ and is negligible for $\ell > 32$, there is a large part of our parameter space in which the high- ℓ corrections are significant. In particular, models with low values of x or large negative curvatures have anisotropic features on strongly sub-horizon scales which project on to high ℓ s.

Fig. 4 illustrates a case where neglecting high- ℓ information is inappropriate. The split map shows how the intense alternating cold-and-hot spirals (lower-right portion) are lost when $\ell > 32$ is ignored (upper-left portion). This is reflected in the Bianchi power spectrum D_ℓ which peaks at $\ell \simeq 70$. Without the high- ℓ information, models such as this will be inaccurately characterized.

Losing this information can equally cause false negatives and false positives. In the case of a true Bianchi universe characterized by a tight spiral, the statistical search will underestimate the likelihood around the correct parameters. Conversely, in the case of a pure Λ CDM universe, the statistical search will overestimate the likelihood of a tightly spiralling feature. In particular, the final results of existing analyses must therefore spuriously favour low values of x and overestimate the upper limit on $(\sigma/H)_0$. Our inclusion of the high- ℓ information will produce tighter and more robust limits on anisotropy.

For practical purposes, we still need to apply a truncation to the Boltzmann hierarchy; since as $x \rightarrow 0$ the power is advected to arbitrarily high ℓ , this generates a lower limit on the values of x we can meaningfully consider. We define ℓ^* to be the minimum multipole at which the average power over the range

¹⁰ The Pontzen (2009) code is several times slower than `ABSOLVE` and its design decisions concerning timestepping and high- ℓ truncation are not suited to a statistical search.




Figure 4. The effects of neglecting the small scales in the treatment of the Bianchi background. An ABSOLVE temperature map for the model $\{\Omega_m = 0.2, \Omega_\Lambda = 0, x = 0.2\}$ was produced twice; first, using multipoles only to $\ell_{\max} = 32$ (upper-left portion of map) and secondly, using multipoles up to $\ell_{\max} = 200$ (lower-right portion of map). Most of the defining features in the Bianchi pattern are lost with the $\ell = 32$ cut. The power spectra of the model are illustrated, with the $\ell = 32$ cut highlighted with a vertical line, reinforcing how significant information is discarded if higher multipoles are not considered.

$\{\ell^* - 10 < \ell \leq \ell^*\}$ is significantly smaller than the power at the quadrupole according to the criterion¹¹

$$\frac{1}{10} \sum_{\ell=\ell^*-10}^{\ell^*} C_\ell < \frac{1}{100} C_{\ell=2}. \quad (8)$$

As discussed in Section 2.2, this test is performed after each integration. If ℓ^* exceeds $\ell_{\max} = \ell_{\text{trunc}} - 50$, the entire integration is repeated with a higher value of ℓ_{\max} up to 1000. Higher values of ℓ_{\max} become extremely slow to evaluate so we designed our priors to avoid models where ℓ^* exceeds 1000; specifically, we exclude regions with $x < 0.05$. This final threshold was chosen after calculating ℓ^* across a grid of 27 000 ($30 \times 30 \times 30$) models spanning the $\{\Omega_m, \Omega_\Lambda, x\}$ unit cube at regular intervals. By this approach, we also verified that ℓ^* rises towards small x and large negative curvatures, as expected. We calculated that the implicit $\ell = 32$ cut that has been applied previously mischaracterizes 16 per cent of the total cube. Therefore, we expect significant changes to posteriors when our high- ℓ treatment is included.

¹¹ We consider the average $\frac{1}{10} \sum_{\ell=\ell^*-10}^{\ell^*} C_\ell$, as opposed to the power C_{ℓ^*} at a single multipole, to reduce the impact of localized dips in the power spectrum spuriously satisfying the constraint in equation (8) at lower multipoles.

Table 4. Constraints on the Bianchi VII_h/VII₀ anisotropy (95 per cent confidence level).

Mode	Parity	95 per cent confidence level
S ^(c)	–	$-3.4 \times 10^{-10} < (\sigma_S/H)_0 < 3.8 \times 10^{-10}$
V ^(c)	Left	$(\sigma_V/H)_0 < 1.7 \times 10^{-10}$
	Right	$(\sigma_V/H)_0 < 1.6 \times 10^{-10}$
T _{reg} ^(c)	Left	$(\sigma_{T,\text{reg}}/H)_0 < 2.4 \times 10^{-7}$
	Right	$(\sigma_{T,\text{reg}}/H)_0 < 2.2 \times 10^{-7}$
T _{irr} ^(c)	Left	$(\sigma_{T,\text{irr}}/H)_0 < 2.4 \times 10^{-9}$
	Right	$(\sigma_{T,\text{irr}}/H)_0 < 2.1 \times 10^{-9}$

4 APPLICATION TO WMAP TEMPERATURE DATA

As a demonstration of the framework developed above, we analyse the WMAP 9-yr temperature data using a combination of the internal linear combination (ILC) map for large scales (Bennett et al. 2003) ($\ell \leq 32$) with the TT high- ℓ likelihood¹² (Bennett et al. 2013) for small scales ($\ell > 32$). The combination is chosen because it allows access to the full sky for the low- ℓ modes while avoiding the complex noise properties of the ILC on small scales (Hinshaw et al. 2007). An earlier version of the ILC was the basis for finding a correlation with a Bianchi VII_h template (Jaffe et al. 2005).

Following the approach described in Section 3.1, the low- ℓ likelihood is specified by equation (3). For the ILC, we employ a Gaussian beam with FWHM of 1° and assume that instrumental noise and residual foreground contamination is negligible at $\ell \leq 32$. The calculation is performed in harmonic space and no masking is applied. The WMAP high- ℓ likelihood code models noise and beams internally.

All the results discussed in this section were obtained by applying the priors listed in Table 3. The complete posterior distributions and triangle plots are available from <https://zenodo.org/record/48653>.

This section is structured as follows. In Section 4.1, we present the constraints we recover for the different anisotropy modes. In Section 4.2, we compare the Bayesian evidence for Bianchi models and Λ CDM. In Section 4.3, we discuss how our prior choices impact on the calculations compared to previous work.

4.1 Constraints on anisotropy

We tested the full anisotropy freedom of the Bianchi VII_h/VII₀ expansion using the WMAP 9-yr data. As described above, our analysis considers the vector (vorticity) modes that have been studied previously, as well as new degrees of freedom that have not previously been included. Table 4 summarizes the constraints we recover for the amplitudes of scalar, vector, regular and irregular tensor modes, as obtained when searching for B^(c) models.

For left-parity V^(c) modes, we obtain $(\sigma_V/H)_0 < 1.7 \times 10^{-10}$ (95 per cent confidence level). Ω_m and Ω_Λ peak around concordance values, driven by the stochastic component. In Fig. 5 (left-hand panel), we show how the Bianchi degree of freedom affecting the morphology, x , is largely unconstrained, with only the tightest spirals being marginally disfavoured. This suggests no overall preference by the data for specific Bianchi patterns; accordingly, the Euler angles are also unconstrained (see right-hand panel of Fig. 5 for an example). Similar results hold for the right-parity V^(c) runs.

¹² http://lambda.gsfc.nasa.gov/product/map/dr5/likelihood_get.cfm




Figure 5. Posterior distribution for the Bianchi x parameter (left-hand panel) and one of the Euler angles α (right-hand panel) for $V^{(c)}$, $T_{\text{reg}}^{(c)}$ and $T_{\text{irr}}^{(c)}$ in left-parity models (right-parity models give similar results). In all cases the data do not strongly prefer any particular values, which reflect the lack of evidence for anisotropy in the data. A comparable analysis of the $V^{(c)}$ case (McEwen et al. 2013) resulted in posteriors rising towards small x (tight spirals); the difference can be traced to our improved analysis including high- ℓ information.

In runs on $V^{(d)}$, the Bianchi sub-case for which early evidence was found in the *WMAP* ILC map, we see a small preference in the parameters controlling the morphology (weakened by the degeneracy) and a sharper preference in the orientation, but, unlike in previous work, we only recover upper limits on the shear amplitude even in the presence of this considerable extra freedom. In Section 4.3, we discuss how the difference in our result can be traced to the choice of priors.

For the case of V modes, it is possible to compute the universal vorticity $(\omega/H)_0$ through equation (7). This allows constraints to be placed on $(\omega/H)_0$ from x , Ω_K and $(\sigma_V/H)_0$; we obtain $(\omega/H)_0 < 1.6 \times 10^{-10}$ (95 per cent confidence level). However, it must be stressed that this limit depends not only on the amplitude of the Bianchi signal, but also on the priors for x and Ω_K .

The three remaining degrees of freedom are constrained for the first time. Table 4 shows that the $T_{\text{reg}}^{(c)}$ mode (which we argued in Section 2.1 to be the best-motivated scenario from the standpoint of accommodating residual anisotropy within the standard cosmological paradigm) is constrained at a level three orders of magnitude weaker than $V^{(c)}$ modes. This results from the regular behaviour which allows for relatively high levels of late-time anisotropy even when the early universe was near-isotropic (see Section 2.1). For the other tensor solution $T_{\text{irr}}^{(c)}$, we obtain limits that are more closely comparable to the $V^{(c)}$ case. In all tensor cases, as in the vector cases, we find that Ω_m and Ω_Λ remain sharply peaked around concordance values and the Euler angles and Bianchi x parameter display no strong preferences (Fig. 5); in other words, the data do not support the existence of anisotropy in these modes.

The $S^{(c)}$ modes are constrained to $-3.4 \times 10^{-10} < (\sigma_S/H)_0 < 3.8 \times 10^{-10}$ (95 per cent confidence level). Due to the scalar symmetry, the parameter x has no effect and the orientation of the scalar pattern only requires two Euler angles to be defined, as it is rotationally invariant around a preferred axis; the data do not prefer any particular values for these two angles. Once again, concordance values are recovered for Ω_m and Ω_Λ . The upper limit on $S^{(c)}$ is slightly less stringent than that on $V^{(c)}$ only because of the different prior shape on $(\sigma_S/H)_0$ and $(\sigma_V/H)_0$.

In summary, for all modes considered, the marginalized posterior distributions peak around the concordance values for the matter and dark energy densities, with little difference between the S , V , T_{reg} and T_{irr} runs. The data do not display any significant preference

Table 5. Log-Bayes factor for different Bianchi+ Λ CDM models with respect to standard flat Λ CDM (positive/negative values favour/disfavour the addition of a Bianchi component).

Mode	Parity	$B^{(c)}$ models	$B^{(d)}$ models
S	–	-6.3 ± 0.2	-2.0 ± 0.2
V	Left	-3.4 ± 0.2	-0.1 ± 0.2
	Right	-3.3 ± 0.2	0.0 ± 0.2
T_{reg}	Left	-3.0 ± 0.2	0.2 ± 0.2
	Right	-3.3 ± 0.2	0.1 ± 0.2
T_{irr}	Left	-3.5 ± 0.2	-0.1 ± 0.2
	Right	-3.6 ± 0.2	0.1 ± 0.2

in the remaining parameters that control Bianchi morphology or orientation. In the specific case of the x parameter, a distinct preference for low values has been found in previous work considering the same data (McEwen et al. 2013), but is absent in our analysis (Fig. 5). This is a consequence of our refined treatment of small scales in the background modelling, which results in a more accurate assessment of the relative probability of Bianchi models with tightly wound spirals (Section 3.4).

4.2 Model comparison

Table 5 shows the log evidence ratios of all examined models with respect to Λ CDM. The self-consistent $B^{(c)}$ models are all strongly disfavoured compared to standard flat Λ CDM. In McEwen et al. (2013), however, the Bianchi hypothesis had comparable evidence to Λ CDM. This difference results from a combination of the improvements introduced in our method (particularly the treatment of the small scales, Section 3.4) and the choice of the prior on $(\sigma_V/H)_0$ (Fig. 2; Section 3.2). The updated prior choice also affects the evidence ratio for the $V^{(d)}$ decoupled model, removing the preference found in McEwen et al. (2013) for the left parity over the right one.

S models stand out as they display a significantly smaller log-Bayes factor than the other degrees of freedom; however, this is a consequence of the uniform prior that we have to adopt for $(\sigma_S/H)_0$ (see Section 3.2). The smallest shear amplitudes are favoured by the data but are given less weight by the uniform (rather than log-uniform) prior, so the evidence values are pushed down. To verify that this effect accounts for the apparent disfavouring of S models, we calculated the log evidence ratio for $V^{(c)}$ and $V^{(d)}$ with a uniform prior on $(\sigma_V/H)_0$ $[0, 10^{-8}]$. The values become, respectively, -6.0 ± 0.2 and -1.7 ± 0.2 for the left parity, confirming that the uniform prior accounts for the down-weighting.

4.3 The effects of prior choices in searches for Bianchi signatures

Fig. 6 shows the posterior distributions recovered for the Bianchi parameters $(\sigma_V/H)_0$ and α in searches for V modes in the *WMAP* ILC map, assuming the following prior choices:

- (i) $V^{(c)}$, log-uniform prior on $(\sigma_V/H)_0$ (solid black line);
- (ii) $V^{(d)}$, log-uniform prior on $(\sigma_V/H)_0$ (dotted green line);
- (iii) $V^{(d)}$, uniform prior on $(\sigma_V/H)_0$ (dashed blue line).

The parameter $(\sigma_V/H)_0$ controls the amplitude of the Bianchi component while α partially controls its orientation. The remaining Euler angles β , γ and the spiral parameter x exhibit similar behaviour to α .

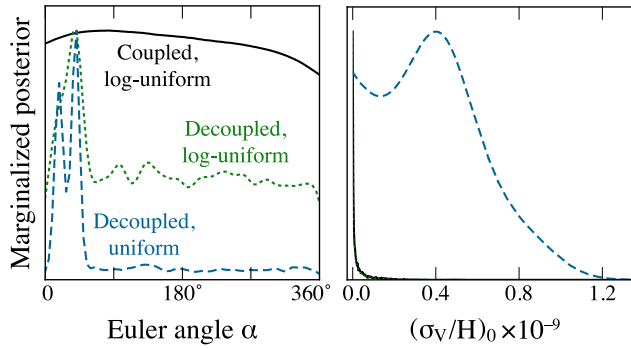


Figure 6. The strong prior-dependence of Bianchi template correlations is illustrated using posterior distributions on Euler angle α and shear $(\sigma_V/H)_0$ for (black solid lines) coupled model with a log-uniform prior on $(\sigma_V/H)_0$, (green dotted lines) decoupled model with a uniform prior on $(\sigma_V/H)_0$ and (blue dashed lines) decoupled model with a uniform prior on $(\sigma_V/H)_0$. The last of these hints at the presence of anisotropy, since the position of the anisotropic features (from α and the other Euler angles) begin to be constrained and significant shear appears to be permitted. However, this prior is the least motivated physically and statistically.

In the log-uniform coupled case (solid black line), we recover the previously quoted upper limit on shear. On the linear scale of Fig. 6, the posterior is very sharply peaked towards zero shear (right-hand panel) and there is no preference for orientation (left-hand panel). By decoupling the parameters (dotted green lines), we find that a preference for a particular orientation begins to emerge, in agreement with McEwen et al. (2013). However, we additionally find that non-zero shear at the amplitude corresponding to the Jaffe et al. (2005) template is only permitted once we also switch to a uniform prior on $(\sigma_V/H)_0$ (dashed blue line). The strong prior-dependence of the analysis shows that even in the case that the parameters are allowed to decouple the data do not robustly support the addition of a Bianchi signal.

5 CONCLUSIONS

We have presented a new framework to search for departures from background isotropy in CMB data. Our approach extends the ANICOSMO tool in two ways. First, we implemented a new Boltzmann hierarchy solution (ABSOLVE) to calculate CMB temperature and polarization patterns arising from anisotropy in the background; this allows a much wider variety of solutions to be probed and in future will allow for the inclusion of CMB polarization data in a coherent way. The wider variety of modes in the new solutions requires us to adopt shear, rather than vorticity, as the primary parameter. Secondly, we improved the statistical approach; in particular, we include the previously neglected effects of the anisotropic background on modes at $\ell > 32$.

As a test of the new approach and to compare with previous results, we searched for departures from isotropy in WMAP temperature data. In doing so, we included three hitherto unconstrained anisotropy modes, including the regular tensor solution which (unlike other modes) is compatible with exiting inflation in a highly isotropic state. Our setup focuses on Bianchi type VII_h and is easily applicable to other types since these are found by allowing the VII_h parameters to limit to boundary values; in this work, we specifically included VII₀ models which are obtained as $\Omega_K \rightarrow 0$.

We find no evidence for anisotropic expansion from WMAP data and place upper limits on present-day shear, as reported in Table 4.

Our constraints on vector modes (linked to vorticity) are tighter than those presented in prior work by a factor of 5, which we showed to be due to a combination of different priors and our improved treatment of small-scale power. Scalar modes are constrained for the first time at a similar level. However, the first constraint on tensor shear – and in particular the regular solution to the tensor anisotropy – is much weaker than the constraint we are able to obtain on the other modes. We showed that this difference arises from the different dynamical nature of the solutions.

In the near future, we plan to extend this framework to allow for analysis of CMB polarization data in addition to temperature. This is expected to further tighten limits on the anisotropy of the Universe (Pontzen & Challinor 2007).

ACKNOWLEDGEMENTS

DS thanks Sabino Matarrese, Alessandro Renzi and Boris Leistedt for helpful discussions. DS is supported by the Perren Fund and the IMPACT fund and was partially supported by the RAS Small Grant Scheme. SMF is supported by the Science and Technology Facilities Council in the UK. AP is supported by the Royal Society. HVP was partially supported by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 306478-CosmicDawn. JDM was partially supported by the Engineering and Physical Sciences Research Council (grant number EP/M011852/1). We acknowledge use of the following publicly available codes: RECFAST (Seager et al. 1999, 2000); HEALPIX/HEALPY (Górski et al. 2005); CAMB (Lewis et al. 2000); MULTINEST (Feroz & Hobson 2008; Feroz et al. 2009); s^2 ¹³ (McEwen et al. 2007). We acknowledge use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

REFERENCES

- Barrow J. D., Juszkiewicz R., Sonoda D. H., 1985, MNRAS, 213, 917
 Bennett C. L. et al., 2003, ApJS, 148, 1
 Bennett C. L. et al., 2013, ApJS, 208, 20
 Bianchi L., 1898, Mem. Mat. Fis. Soc. Ital. Sci., Ser. Terza XI, 11, 267
 Bridges M., McEwen J. D., Lasenby A. N., Hobson M. P., 2007, MNRAS, 377, 1473
 Collins C. B., Hawking S. W., 1973, ApJ, 180, 317
 Ellis G. F. R., MacCallum M. A. H., 1969, Commun. Math. Phys., 12, 108
 Ellis G. F. R., Maartens R., MacCallum M. A. H., 2012, Relativistic Cosmology. Cambridge Univ. Press, Cambridge
 Feroz F., Hobson M. P., 2008, MNRAS, 384, 449
 Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601
 Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2013, preprint (arXiv:1306.2144)
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
 Hinshaw G. et al., 2007, ApJS, 170, 288
 Jaffe T. R., Banday A. J., Eriksen H. K., Górski K. M., Hansen F. K., 2005, ApJ, 629, L1
 Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473
 McEwen J. D., Hobson M. P., Lasenby A. N., Mortlock D. J., 2006, MNRAS, 369, 1858
 McEwen J. D., Hobson M. P., Mortlock D. J., Lasenby A. N., 2007, IEEE Trans. Signal Process., 55, 520
 McEwen J. D., Josset T., Feeney S. M., Peiris H. V., Lasenby A. N., 2013, MNRAS, 436, 3680

¹³ <http://www.jasonmcewen.org/codes.html>

Page L. et al., 2007, ApJS, 170, 335
 Planck Collaboration XXVI, 2014, A&A, 571, A26
 Planck Collaboration XVIII, 2015, A&A, preprint (arXiv:1502.01593)
 Pontzen A., 2009, Phys. Rev. D, 79, 103518
 Pontzen A., Challinor A., 2007, MNRAS, 380, 1387
 Pontzen A., Challinor A., 2011, Class. Quantum Gravity, 28, 185007
 Seager S., Sasselov D. D., Scott D., 1999, ApJ, 523, L1
 Seager S., Sasselov D. D., Scott D., 2000, ApJS, 128, 407
 Wainwright J., Ellis G. F. R., eds, 1997, Dynamical Systems in Cosmology.
 Cambridge Univ. Press, Cambridge, p. 35

APPENDIX A: APPROXIMATIONS IN THE HIGH- ℓ LIKELIHOOD

In this Appendix, we discuss how we have incorporated a correction for the background Bianchi power in high- ℓ likelihoods. For motivation and an overview, see Section 3. In brief, at $\ell > 32$ we incorporate the Bianchi power by calculating the sky-averaged equivalent power spectrum C_ℓ^{B} and adding it to the isotropic stochastic power $C_\ell^{\Lambda\text{CDM}}$ before passing to a standard high- ℓ likelihood.

To understand why this produces a reasonable approximation to the true likelihood, let the measured CMB fractional temperature be expanded in complex spherical harmonics as

$$\frac{\Delta T}{T}(\vartheta, \varphi) = \sum_{\ell=0}^{\ell_{\text{max}}} d_{\ell m} Y_{\ell m}(\vartheta, \varphi), \quad (\text{A1})$$

and similarly for the Bianchi template

$$\left(\frac{\Delta T}{T}\right)^{\text{(B)}}(\vartheta, \varphi) = \sum_{\ell=0}^{\ell_{\text{max}}} b_{\ell m} Y_{\ell m}(\vartheta, \varphi). \quad (\text{A2})$$

We furthermore define the data and Bianchi power spectra \hat{C}_ℓ and C_ℓ^{B} as

$$\hat{C}_\ell \equiv \frac{\sum_{m=-\ell}^{\ell} d_{\ell m}^* d_{\ell m}}{2\ell + 1}, \quad (\text{A3})$$

$$C_\ell^{\text{(B)}} \equiv \frac{\sum_{m=-\ell}^{\ell} b_{\ell m}^* b_{\ell m}}{2\ell + 1}, \quad (\text{A4})$$

note that, unlike in the ΛCDM case, the Bianchi power spectrum does not provide lossless data compression (which is to say that the true likelihood cannot be expressed purely in terms of $C_\ell^{\text{(B)}}$).

The probability of obtaining the data set $\mathbf{d} = \{d_{\ell m}, 2 \leq \ell \leq \ell_{\text{max}}\}$ given the model, \mathbf{M} , is

$$P(\mathbf{d}|\mathbf{b}, C_\ell^{\Lambda\text{CDM}}, \mathbf{M}) \equiv \prod_{\ell=2}^{\ell_{\text{max}}} \mathcal{L}_{\text{exact}}^{(\ell)}, \quad (\text{A5})$$

with

$$\mathcal{L}_{\text{exact}}^{(\ell)} \equiv \frac{2^\ell}{(2\pi C_\ell^{\Lambda\text{CDM}})^{\frac{2\ell+1}{2}}} \times \exp \left\{ -\frac{(d_{\ell 0} - b_{\ell 0})^2 + 2 \sum_{m=1}^{\ell} |d_{\ell m} - b_{\ell m}|^2}{2C_\ell^{\Lambda\text{CDM}}} \right\}, \quad (\text{A6})$$

which is the explicit form of equation (3) for complex spherical-harmonic coefficients.

A typical high- ℓ likelihood code assumes a pure ΛCDM sky, which in the limit of a full-sky approximation may be written as

$$P(\mathbf{d}|C_\ell^{\Lambda\text{CDM}}, \Lambda\text{CDM}) \equiv \prod_{\ell=2}^{\ell_{\text{max}}} \mathcal{L}_{\Lambda\text{CDM}}^{(\ell)}, \quad (\text{A7})$$

with

$$\mathcal{L}_{\Lambda\text{CDM}}^{(\ell)} \equiv \frac{2^\ell}{(2\pi C_\ell)^{\frac{2\ell+1}{2}}} \exp \left\{ -\frac{2\ell + 1}{2} \frac{\hat{C}_\ell}{C_\ell} \right\} \quad (\text{A8})$$

and where C_ℓ is the input theoretical power spectrum. Our treatment of high multipoles amounts to setting $C_\ell = C_\ell^{\text{(B)}} + C_\ell^{\Lambda\text{CDM}}$. We regard the resulting likelihood as an approximation to equation (3); it may be written as

$$P(\mathbf{d}|C_\ell^{\text{(B)}}, C_\ell^{\Lambda\text{CDM}}, \mathbf{M}) \equiv \prod_{\ell=2}^{\ell_{\text{max}}} \mathcal{L}_{\text{approx}}^{(\ell)}, \quad (\text{A9})$$

with

$$\mathcal{L}_{\text{approx}}^{(\ell)} \equiv \frac{2^\ell}{\left[2\pi \left(C_\ell^{\text{(B)}} + C_\ell^{\Lambda\text{CDM}}\right)\right]^{\frac{2\ell+1}{2}}} \times \exp \left\{ -\frac{2\ell + 1}{2} \frac{\hat{C}_\ell}{C_\ell^{\text{(B)}} + C_\ell^{\Lambda\text{CDM}}} \right\}. \quad (\text{A10})$$

To assess the validity of the approximation, we consider the variance of the quantity

$$\epsilon^{(\ell)} \equiv \frac{\mathcal{L}_{\text{exact}}^{(\ell)} - \mathcal{L}_{\text{approx}}^{(\ell)}}{\mathcal{L}_{\text{exact}}^{(\ell)}}, \quad (\text{A11})$$

which has mean zero by the normalization of the likelihoods. If $C_\ell^{\text{(B)}} > C_\ell^{\Lambda\text{CDM}}$, then $\text{var}(\epsilon^{(\ell)})$ can be shown to be infinite (i.e. the approximation can be arbitrarily wrong). However, this is generally not the case, since the Bianchi C_ℓ s tend to be much smaller than the ΛCDM C_ℓ s. For $C_\ell^{\text{(B)}} < C_\ell^{\Lambda\text{CDM}}$, we have

$$\text{var}(\epsilon^{(\ell)}) = \left\langle \left(\frac{\mathcal{L}_{\text{approx}}^{(\ell)}}{\mathcal{L}_{\text{exact}}^{(\ell)}} \right)^2 \right\rangle - 1 = \left[\frac{(C_\ell^{\Lambda\text{CDM}})^2}{(C_\ell^{\Lambda\text{CDM}})^2 - (C_\ell^{\text{(B)}})^2} \right]^{\frac{2\ell+1}{2}} \times \exp \left\{ \frac{(2\ell + 1)C_\ell^{\text{(B)}}}{C_\ell^{\Lambda\text{CDM}} - C_\ell^{\text{(B)}}} \right\} - 1, \quad (\text{A12})$$

where the angular brackets indicate averaging over all the realizations.

In the limit $C_\ell^{\text{(B)}} \ll C_\ell^{\Lambda\text{CDM}}$, we obtain

$$\text{var}(\epsilon^{(\ell)}) \approx (2\ell + 1) \frac{C_\ell^{\text{(B)}}}{C_\ell^{\Lambda\text{CDM}}}, \quad (\text{A13})$$

which is first order in $C_\ell^{\text{(B)}}/C_\ell^{\Lambda\text{CDM}}$. Models such that $C_\ell^{\text{(B)}} > C_\ell^{\Lambda\text{CDM}}$ for some ℓ are excluded during the sampling as they fall outside the range within which this approximation is valid – they constitute, however, a small fraction of the number of explored models.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.