Hierarchical Bayesian Detection Algorithm for Early-Universe Relics in the Cosmic Microwave Background

Stephen M. Feeney,[1, *] Matthew C. Johnson,[2, 3, †] Jason D. McEwen,[1, ‡]
Daniel J. Mortlock,[4, 5, §] and Hiranya V. Peiris[1, ¶]

[1]*Department of Physics and Astronomy, University College London, London WC1E 6BT, U.K.*
[2]*Department of Physics and Astronomy, York University, Toronto, Ontario, Canada M3J 1P3*
[3]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada N2L 2Y5*
[4]*Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, U.K.*
[5]*Department of Mathematics, Imperial College London, London SW7 2AZ, U.K.*
(Dated: August 20, 2013)

A number of theoretically well-motivated additions to the standard cosmological model predict weak signatures in the form of spatially localized sources embedded in the cosmic microwave background (CMB) fluctuations. We present a hierarchical Bayesian statistical formalism and a complete data analysis pipeline for testing such scenarios. We derive an accurate approximation to the full posterior probability distribution over the parameters defining any theory that predicts sources embedded in the CMB, and perform an extensive set of tests in order to establish its validity. The approximation is implemented using a modular algorithm, designed to avoid *a posteriori* selection effects, which combines a candidate-detection stage with a full Bayesian model-selection and parameter-estimation analysis. We apply this pipeline to theories that predict cosmic textures and bubble collisions, extending previous analyses by using: (1) adaptive-resolution techniques, allowing us to probe features of arbitrary size, and (2) optimal filters, which provide the best possible sensitivity for detecting candidate signatures. We conclude that the WMAP 7-year data do not favor the addition of either cosmic textures or bubble collisions to ΛCDM, and place robust constraints on the predicted number of such sources. The expected numbers of bubble collisions and cosmic textures on the CMB sky within our detection thresholds are constrained to be fewer than 4.0 and 5.2 at 95% confidence, respectively.

---

[*] stephen.feeney.09@ucl.ac.uk
[†] mjohnson@perimeterinstitute.ca
[‡] jason.mcewen@ucl.ac.uk
[§] mortlock@ic.ac.uk
[¶] h.peiris@ucl.ac.uk

I. INTRODUCTION

The cosmic microwave background (CMB) radiation provides our best picture of the primordial universe, and therefore the best set of observations available to confront theories of the early universe with data. The angular power spectrum of the CMB, together with complementary datasets (e.g., from large-scale structure and supernova surveys), has established the standard model of cosmology, a spatially flat universe dominated by a cosmological constant and cold dark matter (ΛCDM). However, many theories of high-energy physics predict that there should be deviations from the isotropic and purely Gaussian density fluctuations predicted by ΛCDM. In this paper, we are concerned with the question of how to optimally test theories that predict spatially-localized sources embedded in the CMB. We present a statistical formalism and a set of approximations that are implemented in a full analysis pipeline to construct the posterior probability distribution over the parameters describing a class of theories. We implement a two-step algorithm in which we first locate the most promising candidate signatures, and then use the information about the number, location, and properties of the candidate sources to construct an approximation to the full posterior probability distribution.

To illustrate the application of this pipeline, we focus on two signatures that are predicted by theories with spontaneous symmetry breaking giving rise to phase transitions in the early universe: cosmic textures and cosmic bubble collisions. Cosmic textures are a type of topological defect produced when a non-Abelian global symmetry is broken [1]. Textures are not stable, but instead undergo collapse as they come within the expanding cosmological horizon, eventually unwinding into scalar radiation [1–5]. CMB photons passing through a collapsing texture will be red-shifted, while those passing through an exploding texture will be blue-shifted, giving rise to a set of features in the CMB [5]. Cosmic bubble collisions are predicted by theories of eternal inflation, where our observable universe is postulated to be embedded inside one bubble among many, formed during a first-order phase transition out of an inflating false vacuum (for a review of eternal inflation see, e.g., Refs. [6, 7]). The density perturbations induced by collisions between our bubble and others can lead to localized features in the CMB, providing an observable signature of the dynamics of eternal inflation [8].

In previous work, we presented the first constraints on theories giving rise to cosmic bubble collisions [9, 10] and the first full-sky constraints on cosmic textures [11]. The present paper focuses on:

- Generalizing the statistical formalism and approximation scheme used in Refs. [9–11];

- Implementing an adaptive-resolution analysis, allowing us to overcome the limitations in Refs. [9, 10] on the size of candidate bubble collisions;

- Including and refining the candidate detection scheme using optimal filters presented in Ref. [12];

- Performing a complete suite of tests of the formalism, approximations, and analysis pipeline;

- Performing a new analysis of the posterior probability distribution for bubble collisions and cosmic textures that includes new candidates from the optimal filtering step in combination with the upgraded adaptive-resolution analysis pipeline.

The paper is structured as follows. In Sec. II, we outline the formalism and approximations we use. In Sec. III, we describe the theoretical predictions for cosmic textures and bubble collisions. The algorithm used to calculate the approximated posterior is described in Secs. IV-V and tested in Secs. VI and VII. A null test of the pipeline is carried out in Sec. VIII before the pipeline is applied to CMB data from the Wilkinson Microwave Anisotropy Probe (WMAP) [13] in Sec. IX. The results of this analysis are compared with previous analyses in Sec. X, and our conclusions are summarised in Sec. XI.

II. HIERARCHICAL BAYESIAN SOURCE DETECTION FORMALISM

A. The theory

The observed fluctuations in the CMB can be modeled as the realization of a random field on the sphere, which, under the assumption of isotropy and Gaussianity, is completely characterized by its angular power spectrum. A number of extensions of this model predict various populations of distinct sources embedded in the background random field. This includes astrophysical sources such as clusters of galaxies (which affect the CMB through the Sunyaev-Zel'dovich effect [14]), and primordial sources such as cosmic textures and cosmic bubble collisions. We

restrict our attention to cases where the temperature anisotropies can be described as

$$\frac{\Delta T}{T}(\theta, \phi) = \delta(\theta, \phi) + n(\theta, \phi) + \sum_{i=1}^{N_s} t_i(\theta, \phi). \tag{1}$$

Here, $\delta(\theta, \phi)$ is a realization of the background random field, $n(\theta, \phi)$ describes the instrumental noise as well as residual foregrounds, and $t_i(\theta, \phi)$ are templates for the temperature anisotropies laid down by each of N_s distinct sources. We assume that the sources under consideration are non-overlapping. All terms other than the instrumental noise are assumed to include the effects of a finite instrumental beam. Such a theory can be described by

- **Model parameters:** This includes the parameters describing the background random field and the source templates.[1] The background random field is described by the parameters of ΛCDM, which we denote by the vector $\mathbf{m}_{\Lambda\text{CDM}}$. These parameters include: the fraction of energy density in baryons, $\Omega_b h^2$; cold dark matter, $\Omega_{\text{CDM}} h^2$; and dark energy, Ω_Λ; the scalar spectral index, n_s; the primordial scalar amplitude, A_s; and the optical depth to reionization, τ. Modeling the instrument gives a characterization of the expected noise properties. No model of the Galactic foreground residuals is available for the dataset considered in this analysis, and we therefore resort to null tests of simulations including foreground residuals in order to determine their effects (although a model of the foreground residuals could, in principle, be included in the formalism).

  It is convenient to treat the extension hypothesis as a hierarchical Bayesian model (e.g. Ref. [15]) in which the population-level parameters are considered separately from the lower-level parameters describing the individual sources. The parameters describing the templates are hence divided into two categories: global parameters, $\mathbf{m}_0$, which describe the source population as a whole; and local parameters, $\mathbf{m}_i$, characterizing individual sources. Any model will possess at least one global parameter – $\bar{N}_s$, the expected total number of detectable sources – in addition to any properties common to all templates. Further, any model will possess at least one set of local parameters: $\{\theta_i, \phi_i\}$, the central position of the i^{th} template. Other properties that can differ from template to template (e.g., size) are also classified as local parameters. Global template parameters, in addition to the parameters of ΛCDM, can be thought of as labeling different theories, characterizing the background cosmology and the type of source. Local parameters characterize the properties of sources in the context of a specific theory.

- **Theoretical priors:** An important component of the theory is the prior probability distribution over the model parameters. In principle, a complete theory of cosmology would provide an explanation for the observed properties of the population of sources *and* the background random field. In general such a full theory is not available. To make progress, we will assume that there are no correlations between the properties of the background field and the properties of the sources, rendering the prior separable. In the context of a specific theory of sources, the prior over local parameters can be fully determined. The priors over the parameters of ΛCDM and the global template parameters are somewhat less certain in the absence of a theoretical construction in which different models can be compared.[2] A reasonable assumption is therefore to use an uninformative prior, which assigns equal probability to all possibilities. The use of uninformative priors requires care to be taken when defining the prior range, and will be discussed in later sections.

- **Model statistics:** Given a set of model parameters, it is necessary to understand how particular realizations of the temperature anisotropies are determined. For the sources, this is most efficiently encoded in the theoretical prior over local model parameters of the templates. For the background random field and instrumental noise, this is most efficiently encoded in the two-point correlation function. Under ΛCDM, for perfect data, the correlation in the temperature between two positions on the sphere is given by

$$C_{ij} \equiv C(\theta_{ij}) = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell(\mathbf{m}_{\Lambda\text{CDM}}) P_\ell(\cos(\theta_{ij})), \tag{2}$$

  where θ_{ij} is the angular distance between two points on the sphere labeled by i and j, and $C_\ell(\mathbf{m}_{\Lambda\text{CDM}})$ is the angular power spectrum, which is dependent on the choice of parameters $\mathbf{m}_{\Lambda\text{CDM}}$. The characterization of the instrumental noise and beam depends on the experiment in question, and will be described in greater detail below for the WMAP experiment.

---

[1] The formalism could also be extended to allow for marginalization over any imperfectly known experimental parameters. For simplicity, we assume that the parameters of the WMAP experiment are perfectly known.

[2] Of course, the best example is the eternally inflating multiverse, in which regions with diverse physical properties are sampled. Defining the theoretical prior in this case is difficult due to the infinite number of regions that must be compared; this is the "Measure Problem" of eternal inflation (see Refs. [16, 17] for recent reviews).

B. The full posterior

Having fully specified the theory, we can now ask how to test it. Our goal is to construct the posterior probability distribution over the global source parameters, given a dataset $\mathbf{d}$ consisting of pixelized temperature measurements covering a solid angle $\Omega_{\mathrm{obs}} = 4\pi f_{\mathrm{sky}}$ of the sky (and, optionally, any statistics derived from them). Bayes' theorem gives the posterior as

$$\Pr(\mathbf{m}_0|\mathbf{d}, f_{\mathrm{sky}}) = \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}})}{\Pr(\mathbf{d}|f_{\mathrm{sky}})}, \tag{3}$$

where $\Pr(\mathbf{m}_0)$ is the prior distribution on the global parameters, $\mathbf{m}_0$, $\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}})$ is the likelihood of getting the observed data, and $\Pr(\mathbf{d}|f_{\mathrm{sky}})$ ensures that the posterior is normalized. The posterior can also be used to derive summary statistics, such as confidence intervals on the global model parameters. The quantity $\bar{N}_{\mathrm{s}}$ is always included in the set of global model parameters $\mathbf{m}_0$, and ΛCDM is specified by $\bar{N}_{\mathrm{s}} = 0$. Therefore, we can also perform model selection by comparing the posterior probability of a model for which $\bar{N}_{\mathrm{s}} = 0$ and one which admits $\bar{N}_{\mathrm{s}} > 0$.

The model likelihood $\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}})$ is obtained by marginalizing over:

1. The parameters of ΛCDM;

2. The actual number of sources present on the observable sky (given the expected number of sources, $\bar{N}_{\mathrm{s}}$, the actual number is a realization of a Poisson-like process of mean $f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}$);

3. The local template parameters.

Unpacking the model likelihood, we therefore have:

$$\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}}) = \sum_{N_{\mathrm{s}}=0}^{\infty} \frac{(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^{N_{\mathrm{s}}} e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}}{N_{\mathrm{s}}!} \tag{4}$$
$$\times \int d\mathbf{m}_{\Lambda\mathrm{CDM}} \Pr(\mathbf{m}_{\Lambda\mathrm{CDM}}) \int d\mathbf{m}_1 \ldots d\mathbf{m}_{N_{\mathrm{s}}} \Pr(\mathbf{m}_1, \ldots, \mathbf{m}_{N_{\mathrm{s}}}) \Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots, \mathbf{m}_{N_{\mathrm{s}}}),$$

where $\Pr(\mathbf{m}_{\Lambda\mathrm{CDM}})$ is the prior over the parameters of ΛCDM, $\Pr(\mathbf{m}_1, \ldots \mathbf{m}_{N_{\mathrm{s}}})$ is the prior over the local model parameters for N_{s} sources, and $\Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots \mathbf{m}_{N_{\mathrm{s}}})$ is the likelihood. For measurements of the CMB temperature anisotropies under the assumption of ΛCDM as the background random field, the likelihood is written as

$$\Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots, \mathbf{m}_{N_{\mathrm{s}}}) = \frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-(\mathbf{d}-\sum_{i=1}^{N_{\mathrm{s}}} \mathbf{t}(\mathbf{m}_i))\mathsf{C}^{-1}(\mathbf{d}-\sum_{i=1}^{N_{\mathrm{s}}} \mathbf{t}(\mathbf{m}_i))^{\mathrm{T}}/2}, \tag{5}$$

where C is the pixel-pixel covariance matrix including a ΛCDM CMB signal and instrumental noise.

One must evaluate Eq. 5 in order to construct the full posterior, Eq. 3. Evaluating the full posterior exactly is impossible, as it requires marginalizing over a formally infinite dimensional parameter space (since one must consider realizations with an arbitrary number of templates). Even if the parameter space were finite, the enormous size of modern CMB datasets, such as those produced by the WMAP or Planck [18] experiments, makes inverting the (non-diagonal) covariance matrix prohibitively expensive. Nevertheless, it is possible to apply a controlled and testable sequence of approximations to estimate the full posterior, as we now describe.

C. Approximation to the posterior

In order to proceed, we must construct a suitable approximation to the model likelihood. Let us first deal with the cosmological parameters $\mathbf{m}_{\Lambda\mathrm{CDM}}$. We assume that our inferences on $\mathbf{m}_i$ do not vary significantly for the range of ΛCDM parameters allowed by current cosmological data. Under this assumption, we can proceed as if $\Pr(\mathbf{m}_{\Lambda\mathrm{CDM}}) = \delta(\mathbf{m}_{\Lambda\mathrm{CDM}} - \bar{\mathbf{m}}_{\Lambda\mathrm{CDM}})$, where $\bar{\mathbf{m}}_{\Lambda\mathrm{CDM}}$ is the best-fit concordance cosmological model. Performing the integral over $\mathbf{m}_{\Lambda\mathrm{CDM}}$ we obtain the approximation to the model likelihood

$$\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}}) \simeq \sum_{N_{\mathrm{s}}=0}^{\infty} \frac{(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^{N_{\mathrm{s}}} e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}}{N_{\mathrm{s}}!} \Pr(\mathbf{d}|N_{\mathrm{s}}), \tag{6}$$

where we define the helpful short-hand

$$\Pr(\mathbf{d}|N_{\mathrm{s}}) \equiv \int d\mathbf{m}_1 \ldots d\mathbf{m}_{N_{\mathrm{s}}} \Pr(\mathbf{m}_1, \ldots, \mathbf{m}_{N_{\mathrm{s}}}) \Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}, \bar{\mathbf{m}}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots, \mathbf{m}_{N_{\mathrm{s}}}). \tag{7}$$
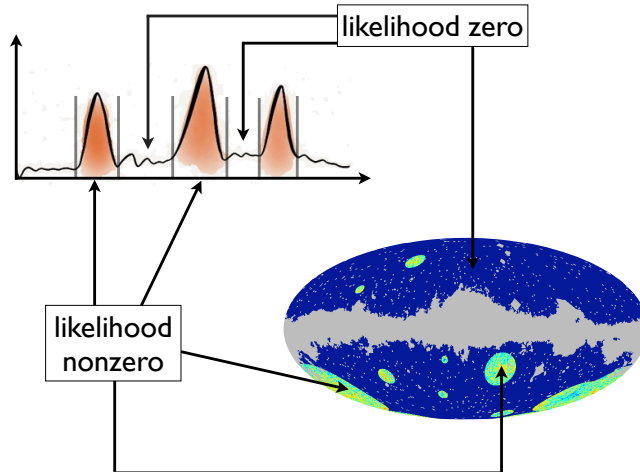
FIG. 1. A schematic depicting the approximation scheme we employ. By locating a set of well-separated candidate sources, it is possible to determine in which regions of parameter space the likelihood function is appreciably different from zero. Eq. 7 can be approximated by integrating over only those regions where the likelihood is large. This collapses the sum in Eq. 6 to a finite number of terms. Finally, we neglect correlations of the random Gaussian background CMB between pixels inside and outside each blob.

1. Likelihood Derivation for One Candidate Source

In a similar spirit, if we knew something about the dependence of the likelihood on the local model parameters, it would be possible to approximate the remaining integrals. This is depicted schematically in Fig. 1. To see how this works in detail, let us begin with a particular example. Imagine that there is a region of the sky which has been judged by some independent method to be a good candidate source. We can segment the data into a "blob", containing the candidate source, and the rest of the sky. The details of the size and shape of the blob are treated in abstract here, and will depend on the particular theory of the sources being tested. We can now evaluate the sum over N_s in Eq. 6 term by term. The likelihood in the first term, for $N_\mathrm{s} = 0$, is simply given by

$$\Pr(\mathbf{d}|N_\mathrm{s} = 0) = \frac{1}{(2\pi)^{N_\mathrm{pix}/2}|\mathsf{C}|}e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^\mathrm{T}/2}, \tag{8}$$

which is the likelihood for the null hypothesis, i.e., no sources. Here, and in what follows, C is evaluated at the best fit values $\bar{\mathbf{m}}_{\Lambda\mathrm{CDM}}$. Moving on to the $N_\mathrm{s} = 1$ term, we must evaluate the integral over $\mathbf{m}_1$. Recall that the local model parameters always include the location at which the template is centered and, if relevant, its size. We can therefore separate the integral over $\mathbf{m}_1$ into the region inside the blob containing our candidate source, which we will refer to as region b, and the rest of the sky, which we will refer to as region $\bar{b}$:

$$\Pr(\mathbf{d}|N_\mathrm{s} = 1) = \int_b \mathrm{d}\mathbf{m}_1 \Pr(\mathbf{m}_1)\frac{1}{(2\pi)^{N_\mathrm{pix}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^\mathrm{T}/2}$$
$$+ \int_{\bar{b}} \mathrm{d}\mathbf{m}_1 \Pr(\mathbf{m}_1)\frac{1}{(2\pi)^{N_\mathrm{pix}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^\mathrm{T}/2}. \tag{9}$$

If there are no sources in the region $\bar{b}$, then we can approximate the likelihood by integrating over region b alone. The accuracy of this approximation depends on our ability to locate the candidate source; however, it will always provide a *lower* (i.e., conservative) bound on the likelihood (since we are integrating a positive-definite function). We therefore have

$$\Pr(\mathbf{d}|N_\mathrm{s} = 1) \simeq \int_b \mathrm{d}\mathbf{m}_1 \Pr(\mathbf{m}_1)\frac{1}{(2\pi)^{N_\mathrm{pix}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^\mathrm{T}/2}. \tag{10}$$

While we have reduced the parameter space over which we must integrate, this expression is still numerically intractable for large datasets, since we must invert the large, generally non-diagonal covariance matrix. To make progress, we

must make a few further approximations. Expanding the exponential in the likelihood, we have

$$\frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2} = \frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-(d_i-t_i)C_{ij}^{-1}(d_j-t_j)/2} \times e^{-d_\alpha C_{\alpha\beta}^{-1}d_\beta/2} \times e^{-(d_i-t_i)C_{i\alpha}^{-1}d_\alpha}, \quad (11)$$

where the indices i and j refer to pixels in region b while the indices α and β correspond to pixels in region $\bar{b}$. We have used the fact that the template vanishes in region $\bar{b}$. Re-arranging, we obtain

$$\frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2} = \frac{e^{-(d_i-t_i)C_{ij}^{-1}(d_j-t_j)/2}}{e^{-d_iC_{ij}^{-1}d_j/2}} \times \left(\frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2}\right) \times \left(e^{t_iC_{i\alpha}^{-1}d_\alpha}\right) \quad (12)$$

$$\simeq \frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)-1}\mathbf{d}^{\mathrm{T}}/2}} \times \left(\frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2}\right) \times \left(e^{t_iC_{i\alpha}^{-1}d_\alpha}\right) \quad (13)$$

$$\simeq \frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)-1}\mathbf{d}^{\mathrm{T}}/2}} \times \Pr(\mathbf{d}|N_{\mathrm{s}}=0). \quad (14)$$

In these expressions, $\mathsf{C}^{(bb)}$ is the covariance matrix constructed using only the data in region b. We have made two approximations in deriving this final expression. First, we have neglected correlations between the template and the data in region $\bar{b}$, which is equivalent to assuming

$$e^{2t_iC_{i\alpha}^{-1}d_\alpha^{\mathrm{T}}} \simeq 1. \quad (15)$$

This is justified in the limit where the inverse covariance falls off sufficiently fast with angular distance. Our second approximation was to assume that we can make the replacement

$$(d_i - t_i)C_{ij}^{-1}(d_j - t_j) \rightarrow (\mathbf{d} - \mathbf{t})\mathsf{C}^{(bb)-1}(\mathbf{d} - \mathbf{t})^{\mathrm{T}}/2, \quad (16)$$

which is justified to the extent that the subgroup of the inverse of the full covariance matrix corresponding to pixels in region b can be approximated as the inverse of a covariance matrix defined only using pixels in region b. For a diagonal covariance, this is exact. For the non-diagonal covariance matrix of ΛCDM it is only approximate. We comment on the validity of these approximations under ΛCDM in Section VII.

Finally, performing the integral over $\mathbf{m}_1$ in Eq. 10, we obtain

$$\Pr(\mathbf{d}|N_{\mathrm{s}}=1) \simeq \Pr(\mathbf{d}|N_{\mathrm{s}}=0)\,\rho_b(\mathbf{m}_0), \quad (17)$$

where the patch-based evidence ratio ρ_b is given by

$$\rho_b(\mathbf{m}_0) \equiv \frac{\int_b d\mathbf{m}_1 \Pr(\mathbf{m}_1)e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)-1}\mathbf{d}^{\mathrm{T}}/2}}. \quad (18)$$

This is a measure of how much better the theory with a template at fixed $\mathbf{m}_0$ fits the patch of data than the theory with only the random field.

Now, we evaluate the $N_{\mathrm{s}} = 2$ term. Again, we approximate the full integral as the integral over region b alone. This yields

$$\Pr(\mathbf{d}|N_{\mathrm{s}}=2) = \int_b \int_b d\mathbf{m}_1 d\mathbf{m}_2 \Pr(\mathbf{m}_1)\Pr(\mathbf{m}_2)\frac{1}{(2\pi)^{N_{\mathrm{pix}}/2}|\mathsf{C}|}e^{-(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))^{\mathrm{T}}/2}. \quad (19)$$

Making the same approximation about the covariance matrix as above, we have

$$\Pr(\mathbf{d}|N_{\mathrm{s}}=2) \simeq \Pr(\mathbf{d}|N_{\mathrm{s}}=0)\frac{\int_b \int_b d\mathbf{m}_1 d\mathbf{m}_2 \Pr(\mathbf{m}_1)\Pr(\mathbf{m}_2)e^{-(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))\mathsf{C}^{(bb)-1}(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)-1}\mathbf{d}^{\mathrm{T}}/2}}. \quad (20)$$

If there is in fact only a single source in region b then the addition of another template will not increase the likelihood: in effect, we would be trying to fit a single feature with a template possessing *twice* the number of parameters ($\mathbf{m}_1$ and $\mathbf{m}_2$). As a concrete example, assume that the source location is the only local model parameter. If the source can be located anywhere on the sky with equal probability, the theory prior is simply

$$\Pr(\mathbf{m}_i) = \frac{1}{4\pi}. \quad (21)$$

In the case where the likelihood function is roughly equal for all positions inside b, with the solid angle contained in b given by Ω_b, and there is no improvement from adding a second template, the relative size of the $N_{\rm s} = 1$ and $N_{\rm s} = 2$ terms can be estimated as

$$\frac{\Pr(\mathbf{d}|N_{\rm s} = 1)}{\Pr(\mathbf{d}|N_{\rm s} = 2)} \simeq \frac{4\pi}{\Omega_b}. \tag{22}$$

Assuming the blob does not cover the entire sky, this is always larger than one. Subsequent terms in the $N_{\rm s}$ expansion will be penalized by higher powers of this ratio. While this is a toy model, this property is expected to hold generally.

More generally, if there are multiple sources in a single blob, increasing the number of templates will improve the fit. If the sources are well-separated, however, this will be a very unlikely occurrence, and terms with $N_{\rm s} > 1$ can therefore be neglected. We operate under this assumption, providing justification in the context of the specific sources considered in Sec. III. Relaxation of this assumption would require the inclusion of all terms in the sum in Eq. 6.

For a single blob, we can therefore approximate the full-sky posterior Eq. 3 as

$$\Pr(\mathbf{m}_0|\mathbf{d}, f_{\rm sky}) \simeq \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|N_{\rm s} = 0)\, e^{-f_{\rm sky}\bar{N}_{\rm s}}\left(1 + (f_{\rm sky}\bar{N}_{\rm s})\,\rho_b(\mathbf{m}_0)\right)}{\Pr(\mathbf{d}|f_{\rm sky})}, \tag{23}$$

where

$$\Pr(\mathbf{d}|f_{\rm sky}) \equiv \int \mathrm{d}\mathbf{m}_0 \Pr(\mathbf{m}_0) e^{-f_{\rm sky}\bar{N}_{\rm s}}\left(1 + (f_{\rm sky}\bar{N}_{\rm s})\,\rho_b(\mathbf{m}_0)\right) \tag{24}$$

is the evidence which ensures $\Pr(\mathbf{m}_0|\mathbf{d}, f_{\rm sky})$ is normalized to unity. Recall that $\bar{N}_{\rm s}$ is included in the vector of parameters $\mathbf{m}_0$.

2. Likelihood Derivation for an Arbitrary Number of Candidate Sources

We now move on to discuss the general case, where there are $N_{\rm b}$ blobs, labelled $b_1 \ldots b_{N_{\rm s}}$, identified as containing a candidate source. We assume that the blobs in question do not overlap. For $N_{\rm r} = N_{\rm b} + 1$ regions on the sky, making the approximation that the template likelihood is small when evaluated outside a blob and neglecting correlations between blobs, we obtain for the general case

$$\Pr(\mathbf{d}|N_{\rm s}) = \begin{cases} 0, & \text{if } N_{\rm s} > N_{\rm b}, \\ \Pr(\mathbf{d}|N_{\rm s} = 0)\sum_{b_1,b_2,\ldots,b_{N_{\rm s}}=1}^{N_{\rm b}} \Delta^{b_1 b_2 \ldots b_{N_{\rm s}}} \prod_{s=1}^{N_{\rm s}} \rho_{b_s}(\mathbf{m}_0), & \text{if } N_{\rm s} \leq N_{\rm b}. \end{cases} \tag{25}$$

The quantity $\Delta^{b_1 b_2 \ldots b_{N_{\rm s}}}$ is one when all indices take distinct values and zero otherwise: the sum hence generates all permutations of $N_{\rm s}$ sources located in $N_{\rm b}$ blobs, assuming no more than one source per blob. If there are fewer blobs on the sky than proposed sources, then the likelihood is very small: this would involve fitting more than one template within a single blob, and incurring the penalisation previously discussed. If there are at least as many blobs as proposed sources, then the likelihood takes the form of a sum that includes every possible association of the $N_{\rm s}$ sources with the $N_{\rm b}$ blobs, provided that no two sources are matched to the same blob. For example, the approximation to the $N_{\rm s} = 1$ term in the presence of multiple blobs is

$$\Pr(\mathbf{d}|N_{\rm s} = 1) = \Pr(\mathbf{d}|N_{\rm s} = 0)\left[\rho_1 + \rho_2 \ldots + \rho_{N_{\rm b}}\right], \tag{26}$$

which is simply an approximation to the integral over positions including only the contributions where the likelihood is large. The approximation to the $N_{\rm s} = 2$ term is

$$\Pr(\mathbf{d}|N_{\rm s} = 2) = \Pr(\mathbf{d}|N_{\rm s} = 0)\left[2\rho_1\rho_2 + 2\rho_1\rho_3 + 2\rho_2\rho_3 + \ldots\right], \tag{27}$$

where the product results from the double integral over positions and again we only consider contributions from regions where the likelihood is large. The factor of two multiplying the blob evidences derives from the fact that there are two ways of placing two sources in two separate blobs. Note that terms for $N_{\rm s} \geq 2$ would be altered if our assumption of well-separated blobs was violated.

Substituting Eq. 25 into Eq. 6, the expression for the approximation to the full posterior is given by

$$\Pr(\mathbf{m}_0|\mathbf{d}, f_{\rm sky}) \simeq \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|N_{\rm s} = 0)}{\Pr(\mathbf{d}|f_{\rm sky})}e^{-f_{\rm sky}\bar{N}_{\rm s}}\sum_{N_{\rm s}=0}^{N_{\rm b}}\frac{(f_{\rm sky}\bar{N}_{\rm s})^{N_{\rm s}}}{N_{\rm s}!}\sum_{b_1,b_2,\ldots,b_{N_{\rm s}}=1}^{N_{\rm b}}\Delta^{b_1 b_2 \ldots b_{N_{\rm s}}}\prod_{s=1}^{N_{\rm s}}\rho_{b_s}(\mathbf{m}_0). \tag{28}$$

Eq. 28 is the main result of this calculation, from which all following results are derived. In the limit of a single isolated observation, Eq. 28 reproduces the Bayesian source detection formalism developed in Refs. [19, 20].

III. SOURCES

In this paper, we consider two theories that give rise to localized sources in the CMB: cosmic bubble collisions in the eternal inflation scenario and cosmic textures. For bubble collisions, the only global parameter is $\bar{N}_s$, and the final result of the analysis is a one-dimensional posterior probability distribution. The first analysis of cosmic bubble collisions using a variant of the approximation scheme outlined in the previous section was presented in Refs. [9, 10], where, in addition to the location, three local model parameters (size, edge discontinuity, and amplitude) were included. Cosmic textures have two global parameters: $\bar{N}_s$ and a measure of the symmetry breaking scale ϵ, and therefore the final product is a two-dimensional posterior probability distribution. An analysis of textures, which also used a variant of the approximation scheme outlined above, was presented in Ref. [11]. In this study a model of textures with one local parameter (size) in addition to the position was considered. Previous work on testing for the signature of textures in the CMB was presented in Refs. [21–25]. Below, we outline our models of these two types of sources, including the prior probability distribution on the local model parameters. For bubble collisions, we update the model assumptions of Refs. [9, 10] in light of improved theoretical understanding, and remove the edge discontinuity parameter from our analysis.

A. Bubble collisions

For an overview of the theory of eternal inflation and the observable effects of bubble collisions, we refer the reader to the reviews [6, 26]. For a detailed discussion of the expected signature of bubble collisions in the CMB, we refer the reader to Refs. [9, 27–30]; here, we provide only a brief overview.

Based on the symmetry of the bubble collision spacetime, the existence of a causal boundary splitting the bubble interior into regions affected and not affected by a collision event, and the fact that a bubble collision is a pre-inflationary relic, the most general template for the temperature fluctuation caused by a single bubble collision is given by [9, 10, 27]

$$t(\theta, \phi) = \left(\frac{z_{\text{crit}} - z_0 \cos\theta_{\text{crit}}}{1 - \cos\theta_{\text{crit}}} + \frac{z_0 - z_{\text{crit}}}{1 - \cos\theta_{\text{crit}}} \cos\theta \right) \Theta(\theta_{\text{crit}} - \theta), \tag{29}$$

where θ_{crit} is the angular scale of the source, corresponding to the causal boundary of the collision event, z_0 and z_{crit} are the amplitudes at the center and edge of the template, Θ is the Heaviside step function, and we have centered the template on the Galactic north pole. In the limit of small amplitude, this is an additive contribution to the CMB temperature anisotropies as in Eq. 1. Theoretical work [29, 30] which appeared subsequent to the previous analysis suggests that there is no discontinuity in temperature at the causal boundary, and we therefore restrict our attention to $z_{\text{crit}} = 0$.

The bubble collision model contains only one global parameter, the expected number of detectable sources $\bar{N}_s$. This is partially a function of the properties of the potential sourcing inflation, and as such is impossible to predict without a model for the potential. In the context of an inflationary landscape, $\bar{N}_s$ can be considered as a continuous parameter with some prior distribution reflecting the typical vacua produced, but without a measure for the landscape it is difficult to estimate even an order of magnitude for this number. It may even be quite likely that $\bar{N}_s \gg 1$, in which case the approximation of looking for widely separated sources is not valid. Ref. [31] has derived the observational signatures in this case (see their Fig. 7), and has used WMAP measurements of the low-ℓ CMB temperature power spectrum to constrain a combination of the number of observable collisions, their average amplitude and the current energy density in curvature. We shall concentrate on the small-$\bar{N}_s$ limit and, in the absence of a detailed theory, assume that $\bar{N}_s$ can take values of $\mathcal{O}(1)$, and set $\bar{N}_s$ to be uniform in the range $0 \leq \bar{N}_s \leq 10$. As we shall see, this parameter is constrained by data, and the precise choice of upper limit has no effect on the analysis. In this model, ΛCDM corresponds to $\bar{N}_s = 0$.

The local parameters are the collision signature's central amplitude, z_0, size, θ_{crit}, and location, $\{\theta_0, \phi_0\}$. The modulations are equally likely to be hot or cold and are isotropically distributed across the sky. Theory does not fix the expected amplitude of the collisions, so we assume that the amplitude is uniform in the range $-10^{-4} \leq z_0 \leq 10^{-4}$ (as stronger collisions would have been obvious in previous CMB data). Neglecting the back-reaction of the collision on the geometry of the bubble interior, the distribution of source sizes is proportional to $\sin\theta_{\text{crit}}$ [8, 26, 32, 33]. Further assuming no correlation between the various local parameters, the final *normalized* prior on the local parameters is

$$\Pr(\mathbf{m}_1) = \Pr(z_0)\Pr(\theta_0, \phi_0)\Pr(\theta_{\text{crit}}) = \frac{1}{2 \times 10^{-4}} \left(\frac{\sin\theta_0}{4\pi} \right) \left(\frac{\sin\theta_{\text{crit}}}{\cos\theta_{\text{crit}}^{\min} - \cos\theta_{\text{crit}}^{\max}} \right), \tag{30}$$

where $0 \leq \theta_0 \leq \pi$, $0 \leq \phi_0 < 2\pi$, and the extrema of the size distribution, $\{\theta_{\text{crit}}^{\min}, \theta_{\text{crit}}^{\max}\}$ are chosen such that the bubble collisions are detectable.[3] The lower limit, $\theta_{\text{crit}}^{\min} = 2°$, stems from the fact that the CMB contains considerable power on the degree-scale, greatly increasing the difficulty of detection. The observable signature of a bubble collision with θ_{crit} larger than the upper limit, $\theta_{\text{crit}}^{\max} = 90°$, would be indistinguishable from the signature of a collision of size $180° - \theta_{\text{crit}}$.

B. Cosmic textures

For a detailed discussion of the production, evolution, and observational signature of cosmic textures we refer the reader to the original literature [1–5, 24]. In brief, CMB photons passing through a collapsing or exploding texture will be red- or blue-shifted, producing an azimuthally symmetric feature on the CMB sky of angular size θ_{crit} with temperature profile of the form

$$t(\theta, \phi) = \frac{(-1)^p \epsilon}{\sqrt{1 + 4\left(\frac{\theta}{\theta_{\text{crit}}}\right)^2}}, \quad \theta \leq \theta_*; \quad t(\theta, \phi) = \frac{(-1)^p \epsilon}{2} \exp\left(-\frac{(\theta^2 - \theta_*^2)}{2\theta_{\text{crit}}^2}\right), \quad \theta > \theta_*; \tag{31}$$

where the amplitude $\epsilon = 8\pi^2 G\eta^2$ depends on the scale of symmetry breaking, η, θ_{crit} is the size of the texture, $\theta_* = \sqrt{3}\theta_{\text{crit}}/2$, $p = \{0, 1\}$ and the template is centered at the Galactic north pole.

Assuming that all textures are produced in a single symmetry-breaking phase transition, the texture model has two global parameters: the dimensionless symmetry-breaking scale ϵ and the expected number of detectable sources $\bar{N}_{\text{s}}$. Texture unwinding events produce features that all have the same amplitude on the CMB sky. The total expected number of unwinding events depends on the particular model giving rise to textures. Simulations [24] of SU(2) textures indicate that we can expect to have causal access to ~ 7 textures with $\theta_{\text{crit}} > 2°$ in the CMB; the precise number of *detectable* unwinding events further depends on our particular realization of the background CMB, the dominant source of noise in the analysis. We adopt a uniform prior on $\bar{N}_{\text{s}}$ between $0 \leq \bar{N}_s \leq 10$ to encode our ignorance of the precise theory and the effect of our CMB realization on detectability. Again, we shall see that this parameter is constrained by data, and thus the choice of upper limit has no effect on the analysis. Requiring that the symmetry-breaking scale for textures is below the scale of cosmological inflation (bounded to be lower than $\sim 10^{16}$ GeV by the absence of observed B-mode polarization), we place an upper bound on ϵ of 10^{-4}. We assume a uniform prior on ϵ down to 2.5×10^{-5}, which corresponds to the estimated detection limit of our pipeline (details of which can be found in Sec. IV). Under the assumption of a uniform prior, the posterior is simply proportional to the likelihood; results for a different prior can be obtained easily by re-weighting the posterior.

The local parameters for textures are the size, θ_{crit}, location, $\{\theta_0, \phi_0\}$ and p, which specifies whether the texture is hot or cold. Textures are expected to be isotropically distributed over the sky, with a distribution of sizes proportional to $\theta_{\text{crit}}^{-3}$ (see Ref. [24] for a derivation), and so the normalized prior on the local texture parameters is

$$\Pr(\mathbf{m}_1) = \frac{\sin\theta_0}{4\pi\theta_{\text{crit}}^3} \left(\frac{1}{(2°)^2} - \frac{1}{(50°)^2}\right)^{-1}, \tag{32}$$

where $0 \leq \theta_0 \leq \pi$, $0 \leq \phi_0 < 2\pi$, and we take $2° \leq \theta_{\text{crit}} \leq 50°$. The lower limit on θ_{crit} results from the large power on degree scales in the CMB; the upper limit stems from the fact that templates with $\theta_{\text{crit}} > 50°$ are large enough to cover the whole sky and overlap themselves, rendering Eq. 31 invalid. To marginalize over the different signs, we can simply sum the likelihoods evaluated at the same magnitude of ϵ, but where the template takes opposite sign.

We now describe our implementation of the formalism derived in Sec. II. There are two main steps in the algorithm – candidate-source detection and patch-based evidence calculation – in which the non-zero regions of the likelihood function are first estimated and then evaluated. In the first step, optimal filters matched to the signals of interest are used to identity the positions and sizes of the most likely sources in the dataset. We describe this procedure in Sec. IV below. In the second step, the patch-based evidence ratios, Eq. 18, are calculated for each candidate using the nested sampler MultiNest [34, 35], before combining to form the posterior, Eq. 28, on the global model parameters. We describe this part of the calculation in Sec. V.

---

[3] This restriction to "detectable" signatures comes about as the overall number of collisions on the sky, $\bar{N}_{\text{s}}$, is the primary parameter of interest, as opposed to, for example, the parameters describing any one putative signature. Once this approach is taken it is only possible to place data-driven constraints on the number of collisions with parameters that would make them plausibly detectable: it is impossible to conclude anything about the number of, say, small-θ_{crit} bubble collisions that would have been completely smeared out by the WMAP beam.
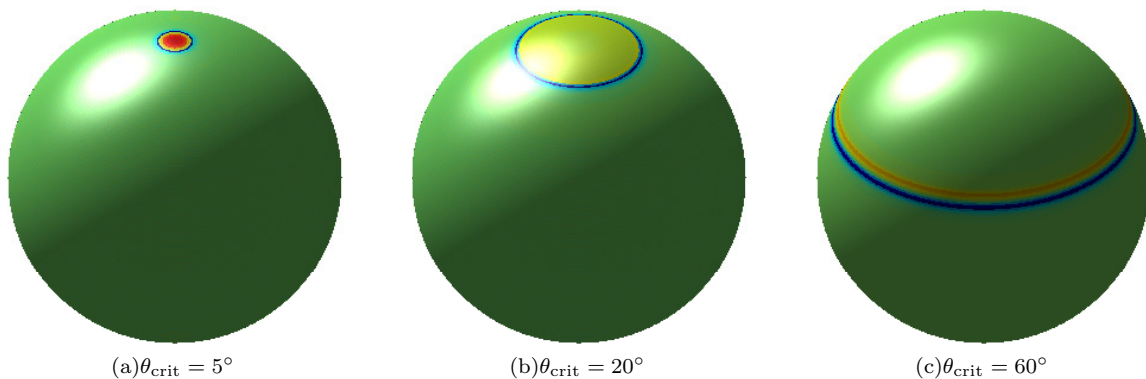
(a)$\theta_{\mathrm{crit}} = 5°$ (b)$\theta_{\mathrm{crit}} = 20°$ (c)$\theta_{\mathrm{crit}} = 60°$

FIG. 2. Matched filters optimized to bubble collision signatures of varying size embedded in a ΛCDM CMB background.

IV.   CANDIDATE DETECTION WITH OPTIMAL FILTERS

In order to effectively approximate the full posterior distribution describing the population of candidate sources, it is necessary to first locate the regions that provide significant contributions to the likelihood. We follow the approach of Ref. [12], using optimal filters defined on the sphere that are matched to the source profile of either bubble collision or texture signatures. First, we construct optimal matched filters for the purpose of detecting candidate sources embedded in full-sky WMAP 7-year data and assess their performance. Second, we briefly describe the optimal-filter-based algorithm for detecting sources of unknown and differing sizes, highlighting differences between the bubble collision and texture cases. Third, we calibrate the algorithm on an end-to-end simulation of WMAP observations, before assessing its sensitivity.

A.   Optimal bubble collision and cosmic texture filters

We construct two sets of matched filters: one set that enhances the contributions of bubble collision signatures and one set that enhances the contributions of texture signatures. The matched filters are constructed to enhance the source profile in a specified stochastic background. A stochastic background of CMB fluctuations is assumed, characterized by the lensed ΛCDM power spectrum that best fits the WMAP 7-year data, baryon acoustic oscillations and supernovae observations [36] (hereafter referred to as the lensed WMAP7+BAO+H_0 power spectrum). The bubble collision and texture source profiles for which we search are relatively large-scale; thus we consider spherical harmonics up to the band-limit $\ell_{\max} = 256$ only. Since we eventually apply these filters to W-band WMAP observations, we assume observations are made in the presence of a Gaussian beam of full-width-half-maximum FWHM = 13.2 arcminutes and isotropic white noise of $N_\ell = 0.02 \ \mu K^2$. The optimal matched filters are constructed in harmonic space: thus the assumption of an azimuthally-symmetric beam profile and isotropic noise simplifies their construction and application considerably, while remaining highly accurate for the relatively low band-limit considered. Once the source profile and stochastic background are defined, the filters are constructed on the sphere as outlined by Ref. [37].

In Fig. 2 and Fig. 3 we show the matched filters recovered for the bubble collision and texture profiles, respectively, for a range of source sizes. Notice that the bubble collision filters on smaller scales contain a central broad hot region to enhance the main bubble collision contribution, surrounded by hot and cold rings to enhance the transition from the collision to the background. On larger scales, however, the matched filters contain only the hot and cold rings that enhance the transition. Since the CMB has more power on large scales, the matched filters on large scales do not respond to the large-scale features of the bubble collision signature but rather the transition region near the location where the template goes to zero. The texture source profile has a smooth, Gaussian transition to the background, and consequently the matched filters recovered for textures contain only a central broad region, without any strong contribution from the perimeter of the profile.

The matched filters constructed are optimal in the sense that no other filter can yield a greater enhancement of the signal-to-noise ratio (SNR) of the filtered field. It is possible to calculate analytically the SNR of the filtered field for various filter types, as derived for example by Ref. [12]. In Fig. 4 we plot the SNR computed for bubble collision and texture profiles for the unfiltered field, for the optimal filters constructed here, and for needlets [38, 39], which have been used previously to detect candidate sources [9–11]. Note that the lack of a sharp transition in the texture template means that the SNR for textures are lower than those of bubble collisions. Nevertheless, it is clear that matched filters yield the highest SNR, in accordance with expectations.
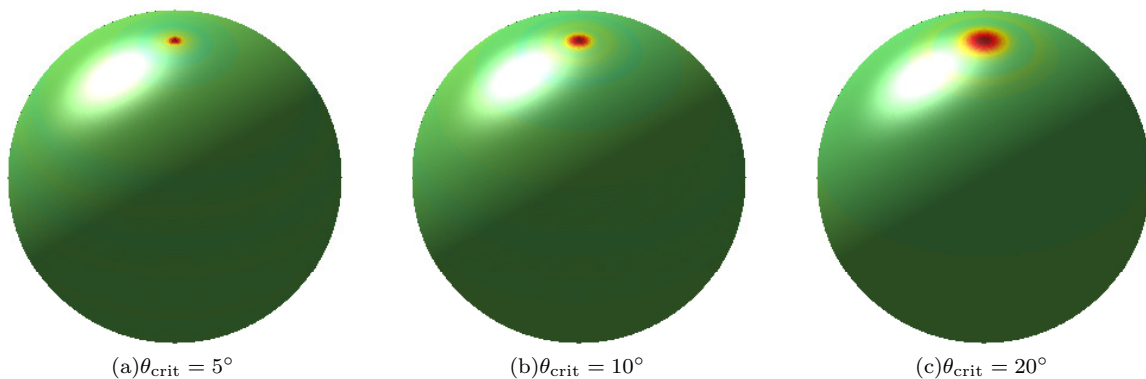
(a)$\theta_{\mathrm{crit}} = 5°$ (b)$\theta_{\mathrm{crit}} = 10°$ (c)$\theta_{\mathrm{crit}} = 20°$

FIG. 3. Matched filters optimized to cosmic texture signatures of varying size embedded in a ΛCDM CMB background.

B. Candidate object detection algorithm

Although we have constructed optimal filters for a range of source sizes, we have not yet addressed the problem of detecting sources of unknown and differing sizes. We adopt the algorithm described in detail by Ref. [12] for this purpose, which we review here briefly. First, matched filters are constructed for a grid of source sizes $R \in \{\theta_{\mathrm{crit}}^k\}_{k=1}^{N_{\theta_{\mathrm{crit}}}}$. All filters are then applied to the full-sky observed data by convolving the matched filter kernel with the observed data, which may be computed efficiently in harmonic space (see e.g. Ref. [12]). Significance maps are then computed by normalizing the filtered field to the mean and standard deviation of filtered fields computed from realizations of the background process (i.e. CMB fluctuations and instrumental noise) in the absence of sources. The significance maps are then thresholded (the calibration of threshold levels is discussed below), before potential candidate sources are found from the localized peaks of the thresholded significance maps. Potential candidate sources are eliminated if a stronger source is found on adjacent scales, where the set of scales adjacent to scale R is defined by the set $\{R_{\mathrm{adj}} \in \{\theta_{\mathrm{crit}}^k\}_{k=1}^{N_{\theta_{\mathrm{crit}}}} : |R_{\mathrm{adj}} - R| \leq \theta_{\mathrm{adj}}\}$, i.e. where the distance between R and R_{adj} is less than the parameter θ_{adj}. Once candidate sources are detected, the parameters of the source size, location and amplitude are estimated from the corresponding filter scale, peak position of the thresholded significance map and amplitude of the filtered field, respectively.

The construction of optimal filters is implemented in the `S2FIL` code [37] (which in turn relies on the codes `S2` [40] and `HEALPix` [41]), while the `COMB` code [37] has been used to simulate bubble collision signatures embedded in a CMB background.[4] The candidate object detection algorithm described here is implemented in a modified version of `S2FIL` that will soon be made publicly available.

There is no guarantee that the peak in the filtered field across scales will coincide with the size of the unknown source. Nevertheless, for bubble collision signatures embedded in the CMB this has indeed been found to be the case [12]. For texture profiles, however, we have found this phenomenon to hold below scales $\theta_{\mathrm{crit}} \sim 10°$ only. Through numerical experiments we found that the difference in the behavior of the filtered field between bubble collisions and textures on large scales is due to the absence of a well-defined transition region from the source to the background in the texture profile. For large texture sizes, not only is there no peak in the filtered field at the scale of the unknown source, but the SNR of the filtered field does not drop off rapidly when applied to nearby scales (Fig. 4), as is the case for bubble collisions. We trivially modify the candidate detection algorithm described above to account for this behavior. For textures sizes below $\theta_{\mathrm{crit}} = 10°$ we look across adjacent scales as usual to find the most significant potential candidate source, whereas for sizes above $\theta_{\mathrm{crit}} = 10°$ we do not (by a judicious choice of the adjacency parameter θ_{adj} there is in fact no need to modify the algorithm, as described below).

Although the candidate detection algorithm considers a grid of candidate scales R, it is overwhelmingly probable that the signal for any given source peaks at scales between the samples of the grid. It is thus important to examine how sensitive the matched filters are to small errors in the source size. In Fig. 4 (b) and (d) we plot SNR curves for matched filters constructed on the grid of candidate scales. For bubble collision profiles, a sharp degradation in the SNR away from the scale used to construct each filter is clearly apparent. For texture profiles the degradation is much less pronounced, especially at large scales (as discussed above). Provided that the θ_{crit} grid is sampled sufficiently densely, the matched filters for both bubble collision and texture profiles remain effective and are superior to needlets.

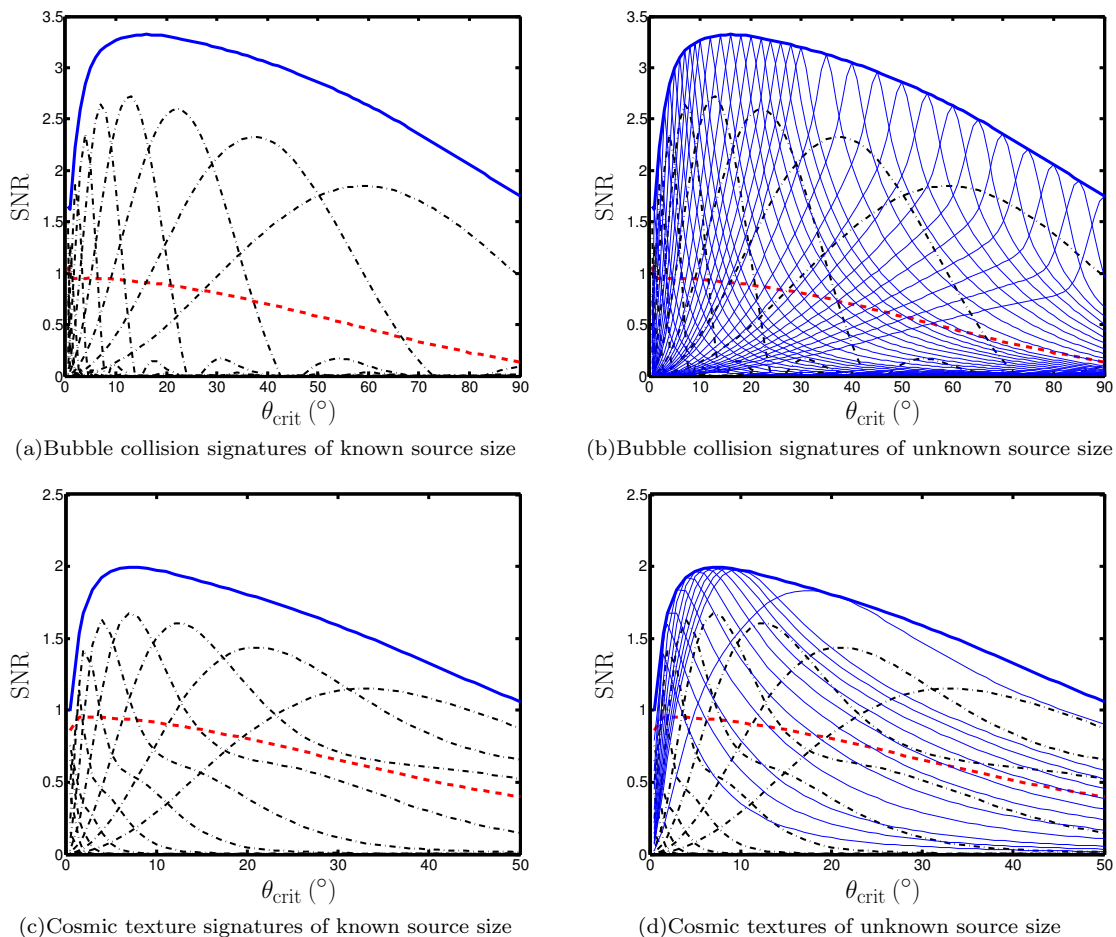
FIG. 4. SNRs of bubble collision (top row) and texture (bottom row) signatures of varying size with amplitude 100 μK embedded in a ΛCDM CMB background. SNR curves are plotted for matched filters (solid blue curve), needlets [38] with scaling parameter $B = 1.8$ for a range of needlet scales j (dot-dashed black curves) and for the unfiltered field (dashed red curve). In panels (b) and (d) SNR curves for the matched filters constructed at a given scale and applied at all other scales are also shown (thin solid blue curves). The scale for which the filters are constructed may be read off the plot from the intersection of the heavy and light solid blue curves. Provided the $\theta_{\rm crit}$ grid is sampled sufficiently densely, the matched filters remain superior to needlets.

C. Candidate object detection calibrated to WMAP

We define the parameters of the optimal-filter-based candidate object detection algorithm here and calibrate the threshold levels for WMAP observations. Throughout the calibration we apply the WMAP KQ85$yr7$ mask [42], since this is the mask adopted when analysing WMAP data subsequently. We select the less-conservative KQ85 mask so as to reduce the variance in reconstructing large-scale information masked by the Galactic sky cut [43]. This choice has the additional advantage of revealing more of the sky and hence increasing the number of candidates. The Bayesian stage of the pipeline will assess the overall evidence for the source model in each candidate region: if any candidate is found to contribute evidence in favor of the source model, then it will be examined closely for frequency-dependence signifying potential foreground contamination.

For bubble collisions we consider the grid of scales set out in Ref. [12], as defined in Table I (left). The $\theta_{\rm crit}$ prior range is smaller for textures than for bubble collisions – the texture profile extends well beyond $\theta_{\rm crit}$, covering the full sky for $\theta_{\rm crit} \gtrsim 50°$ – and hence for textures we consider a smaller grid of scales, also defined in Table I (right). Since we found the matched filters for textures to be sensitive to a large range of nearby scales, we nevertheless remain sensitive to the full prior range of sizes. The SNR curves for the matched filters constructed for these scales for bubble collisions and textures are shown in Fig. 4 (b) and (d). These grids of scales are thus sufficiently sampled to ensure that the matched filters remain effective for scales between the samples of the grid. We set the adjacency parameter to $\theta_{\rm adj} = 5°$ for both bubble collisions and textures. For textures this ensures that we look across scales for sizes

$\theta_{\rm crit}(°)$	N_{σ_R}	$\theta_{\rm crit}(°)$	N_{σ_R}
1	4.25	1	4.25
1.5	4.25	1.5	4.00
2	4.25	2	3.75
3	4.00	3	3.50
4	4.00	4	3.50
5	4.00	5	3.25
6	4.00	6	3.00
7	4.00	7	3.00
8	4.00	8	3.00
9	3.75	9	2.75
10	3.75	10	2.75
12	3.75	20	2.50
14	3.75		
16	3.75		
18	3.50		
20	3.50		
22	3.50		
24	3.50		
26	3.50		
28	3.25		
30	3.25		
35	3.25		
40	3.25		
45	3.25		
50	3.25		
55	3.25		
60	3.25		
65	3.25		
70	3.25		
75	3.25		
80	3.00		
85	3.00		
90	3.00		

TABLE I. The $\theta_{\rm crit}$ grid and threshold levels N_{σ_R} adopted for the optimal-filter-based candidate source detection algorithm for bubble collisions (left) and textures (right). Threshold levels are calibrated to the WMAP end-to-end simulation to allow at most three false detections on each scale.

below $\theta_{\rm crit} = 10°$ but not above, whereas for bubble collisions we always look across adjacent scales.

We use 3,000 Gaussian CMB simulations to calculate the mean and standard deviation of the filtered field at each scale in the absence of sources, in order to compute significance maps. For these simulations, and for the WMAP data analysed subsequently, we perform Wiener filtering to recover spherical harmonic coefficients with $\ell \leq 10$ from masked CMB maps [43], where we adopt a Gaussian prior for the harmonic coefficients specified by the lensed WMAP7+BAO+H_0 power spectrum. Note that this differs from the maximum likelihood reconstruction [44] of harmonic coefficients performed by Refs. [9–11] and the cut-sky estimation performed by Ref. [12]. This can alter the spherical harmonic coefficients recovered on large scales, and thus the detected candidate sources, in a non-negligible manner. However, Wiener filtering should give the most reliable reconstruction of the large-scale harmonic modes [43].

Finally, we calibrate the threshold levels N_{σ_R} applied to the significance maps for each filter scale from a realistic WMAP simulation that does not contain embedded sources. The thresholds are chosen to allow a manageable

number of false detections while remaining sensitive to weak sources. For this calibration we use a complete end-to-end simulation of the WMAP experiment provided by the WMAP Science Team [42]. The temperature maps in this simulation are produced from a simulated time-ordered data stream, which is processed using the same algorithm as the actual data. The data for each frequency band are obtained separately using simulated diffuse Galactic foregrounds and CMB fluctuations, and include realistic noise, smearing from finite integration time, finite beam size, and other instrumental effects. We use the foreground-reduced W-band simulation for calibration. The threshold levels N_{σ_R} are selected to allow at most three false detections on each scale on this simulated map (recall that detections on one scale can be eliminated by stronger detections made on adjacent scales). These threshold levels are less conservative that those set by Refs. [9–12], in order to increase the number of false detections passed by the candidate source detection stage of the analysis pipeline and hence improve its sensitivity. The calibrated threshold levels for both bubble collisions and textures are shown in Table I. Once candidate sources are detected by the optimal-filter-based detection algorithm, we discard those objects that are significantly masked. For the WMAP end-to-end simulation, we make 12 false detections of candidate bubble collisions and 4 false detections of candidate textures.

In the analysis of the WMAP end-to-end simulation (and in the analysis of the WMAP data considered subsequently), some of the bubble collision candidates that we detect differ from those found with optimal filters previously [12]. This is expected since we now use Wiener filtering to recover spherical harmonic coefficients, have included a more accurate model of the WMAP noise in the optimal filter construction (noise was neglected previously), and have reduced the threshold levels in order to increase the sensitivity of the entire pipeline. Further, the thresholding-based nature of the candidate source detection algorithm means that small differences in the filtered field can have an impact of the final candidates detected if regions move below or above the threshold. In the cases where candidates disappeared, we nevertheless found peaks in the filtered field; these were simply no longer above the threshold or were eliminated by stronger detections on nearby scales or positions. Given these differences are due to improvements made to the pipeline, the results given here are to be preferred to those presented previously.

Following the candidate source-detection stage of the analysis pipeline, we pass to the Bayesian stage an estimate of the domain of parameter space over which the likelihood is expected to be non-negligible. These regions of parameter space are estimated from each of the candidate sources detected. For the size of each source, the relevant region is determined first by finding the range of nearby filter scales for which significance maps exceed their threshold level at the source position. This range of scales is extended to the next smallest and largest filter scales (or the edge of the prior if encountered) to yield an estimate of the range of scales over which the likelihood is non-negligible. For example, for a bubble collision candidate found to be significant by the $40°$ and $45°$ optimal filters, we would estimate the likelihood to be non-negligible over the range $35 - 50°$. To estimate the integration limits of the central positions for each source, we first find all pixels within a radius r of the source position estimated by the optimal filters, where r is 25% of the maximum source size estimated from the previous step. The extrema $\{\theta_0^{\min}, \theta_0^{\max}, \phi_0^{\min}, \phi_0^{\max}\}$ of these pixels are found, and the source positions are then sampled from the region defined by $\theta_0^{\min} \leq \theta_0 \leq \theta_0^{\max}$, $\phi_0^{\min} \leq \phi_0 \leq \phi_0^{\max}$. Tests of these assumptions are included in the suite of pipeline tests detailed in later sections.

We conclude this section by assessing the level to which the optimal-filter-based candidate detection algorithm is sensitive for each source type. In previous studies, simulations were performed for this purpose [9, 10, 12]. Here we instead take a probabilistic approach based on the analytic SNRs of the filters computed previously (see Fig. 4). This allows us to probe the source size-amplitude parameter space at higher resolution and accuracy than would be achievable with modest numbers of simulations (to reach an equivalent resolution and accuracy through simulations would be extremely computationally demanding). In Fig. 5 we plot the sensitivity of the matched filters constructed for bubble collisions and cosmic textures. These plots are produced as follows. For each scale θ_{crit} we compute the source amplitude (z_0 for bubbles; ϵ for textures) that would be required to ensure that the SNR reaches the threshold specified in Table I. This level defines the 50% completeness curves shown in Fig. 5 since, in the presence of noise, we expect half of the sources with this amplitude to be detected and half to be missed. Similarly, we compute approximate completeness curves for one-, two- and three-standard-deviation differences from the 50% completeness curve (note that the probabilities quoted on each curve are computed assuming a Gaussian distribution of the filtered field at the source position). For the 50% completeness curve, the bubble collision matched filters are sensitive to $z_0 \sim 10^{-4.4}$, while the cosmic texture matched filters are sensitive to $\epsilon \sim 10^{-4.2}$. Note that the sensitivities computed in this manner are similar to those computed previously through simulations [9, 10, 12], both in terms of the sensitivity levels obtained and the shape of the sensitivity regions. Further, we see that optimal matched filters are ~ 1.7 times more sensitive than needlets for detecting bubble collision signatures, as found previously [9, 10, 12].

V. ADAPTIVE-RESOLUTION EVIDENCE CALCULATION

Modern CMB experiments map the sky with extremely high resolution: the beam of the Planck experiment in the main CMB bands is expected to be $\sim 5'$ [18], resulting in maps pixelated on the arcminute scale. While this
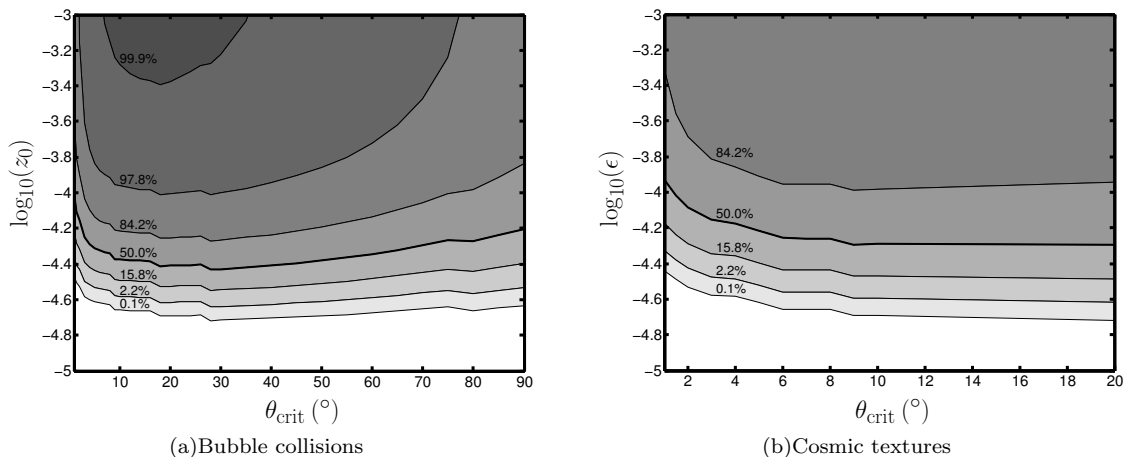
FIG. 5. Sensitivity of the optimal-filter-based candidate detection algorithm, with completeness curves plotted for one, two and three standard deviations from the 50% completeness curve. The completeness curves are computed in the following manner. For each scale θ_{crit} we compute the source amplitude (z_0 for bubbles; ϵ for textures) that would be required to ensure that the SNR reaches the threshold specified in Table I. This level defines the 50% completeness curve, since we expect half of the sources with this amplitude to fall below the curve and half to fall above. Similarly, we compute approximate completeness curves for one-, two- and three-standard-deviation differences from the 50% completeness curve. The probabilities quoted on each completeness curve are computed assuming a Gaussian distribution of the filtered field at the source position.

is necessary for pinning down the secondary CMB anisotropies at small scales, it means that calculating pixel-space covariance matrices becomes extremely memory-intensive. We illustrate this point in Fig. 6, which shows the memory needed to calculate covariance matrices from $1°$ to $180°$ in radius at `HEALPix` resolutions ranging from $N_{side} = 8$ to $N_{side} = 2048$ (i.e., Planck resolution).[5] It is clear that the memory costs, which to a good approximation rise as angular radius to the fourth power, make processing even relatively small patches prohibitive at full Planck resolution.

In previous work [9, 10], we chose to truncate both our patches and the integration limits of θ_{crit} to the maximum radius invertible with our memory constraints. While this allowed us to at least partially process almost all features at full WMAP resolution, it meant that we were unable to probe the large-θ_{crit} region of parameter space for which the prior for bubble collisions is highest. If we are to do so, it is clear that the input maps must be processed at degraded resolution: the larger the patch, the lower the resolution at which it can be processed. The maximum memory accessible per core in this analysis is ~ 90 GB, which means that the full-sky covariance matrix can be inverted at $N_{side} = 64$: this is therefore the minimum resolution at which any feature will be processed. The degradation process will now be described in detail, along with the suite of tests performed to assess its performance.

A. Processing maps

The candidate detection stage returns estimates for the size and position of features of interest within a map, defining the set of patches to be considered in the evidence calculation. The memory cost of computing each covariance matrix is derived from the number of unmasked pixels within the patch: if this is greater than the memory available, the patch must be processed at a degraded resolution.

In the `HEALPix` pixelization scheme, each step down in resolution reduces the pixel count, N_{pix}, by a factor of four. As the covariance matrices are N_{pix}^2 in size, an estimate of the number of steps down required is therefore given by

$$n_{deg} = \text{ceiling}\left(\frac{\log\left(m_{est}/m_{max}\right)}{\log 16}\right), \tag{33}$$

where m_{est} and m_{max} are the estimate of the memory required and the memory available. The corresponding estimate of the finest resolution at which the patch can be processed is then

$$N_{side,deg} = \frac{N_{side,full}}{2^{n_{deg}}}. \tag{34}$$

---

[5] The quantity plotted corresponds to a total number of matrix elements equal to $\sim 1.5 \times N_{pix}^2$. Our algorithm calculates the full N_{pix}^2 covariance matrix in order to make use of the `LAPACK` inversion routines [45], then compresses the inverted matrix to a 1-D array containing its upper triangle to reduce memory costs while sampling.
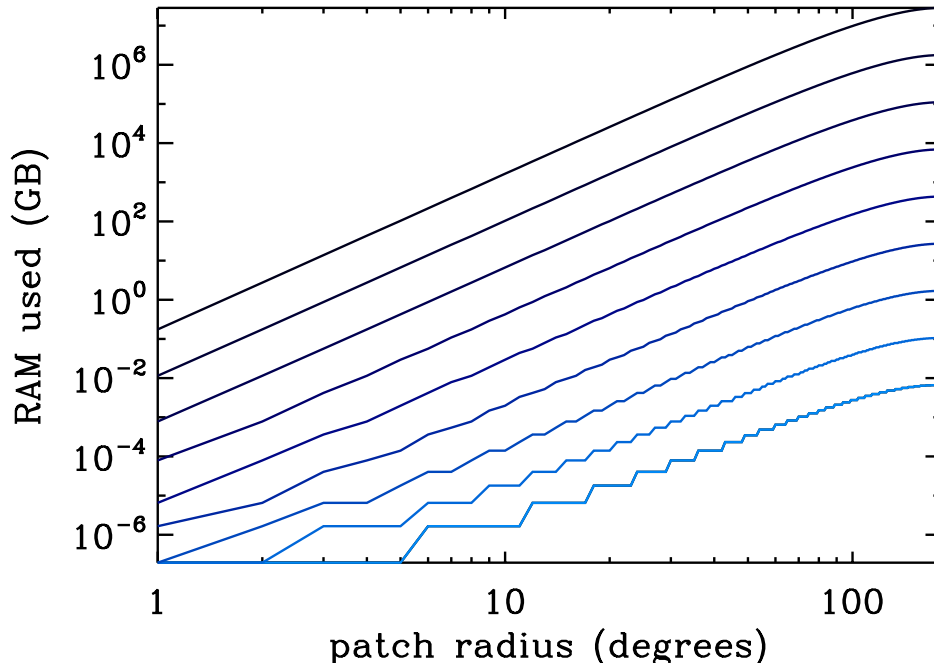
FIG. 6. The memory needed to calculate the covariance matrix for patches of a given size at `HEALPix` resolutions ranging from $N_{\mathrm{side}} = 8$ (lightest-blue, lowest curve) to $N_{\mathrm{side}} = 2048$ (highest, darkest curve). The effects of pixel size are visible at small patch radii and low resolutions.

This estimate is tested by re-counting the number of unmasked pixels in the patch at the new resolution. An $N_{\mathrm{side,deg}}$ mask is created by averaging within degraded-resolution pixels: any $N_{\mathrm{side,deg}}$ pixel which is more than half masked at the input resolution is considered to be masked. A precise calculation of the memory cost of the degraded covariance matrix is then made: if this is below the memory threshold, as expected, the degraded resolution is accepted and the algorithm proceeds; if not, the resolution is decreased once more.

Once the required resolution has been determined, the input CMB temperature map can be degraded. As the CMB is a smooth field, it is not sufficient to simply average within $N_{\mathrm{side,deg}}$ pixels: doing so will introduce large pixelization effects, which will act as an extra noise term unaccounted for in the pixel-pixel covariance matrix. This can be avoided by smoothing such that the input map is smooth on the degraded-pixel scale prior to reducing the resolution. This is equivalent to introducing a band-limit, ℓ_{max}, in harmonic space. Choosing the band-limit is a balance. If the smoothing scale is set too large (i.e., ℓ_{max} is too low), too much information will be discarded with each degrade and performance will suffer. If the smoothing scale is set too small (i.e., smaller than an $N_{\mathrm{side,deg}}$ pixel), the smoothed maps will contain pixelization artefacts.

The choice of ℓ_{max} is somewhat arbitrary, but experimentation shows that the degradation is stable if the harmonic-space Gaussian smoothing kernel is 1% of its maximum at $\ell_{\mathrm{max}} = 2N_{\mathrm{side,deg}}$. This defines a smoothing scale at each resolution: the FWHM, f_{deg}, of the pixel-space kernel is given by

$$f_{\mathrm{deg}} = \sqrt{\frac{8 \log 2 \log 100}{\ell_{\mathrm{max}}(\ell_{\mathrm{max}} + 1)}}. \tag{35}$$

Assuming the $12N_{\mathrm{side}}^2$ pixels in a `HEALPix` map of given resolution are flat and square (a safe assumption at high-resolution), the smoothing scale is approximately 2.5 times the size of a pixel, and the maps are clearly smooth on the pixel-scale. The smoothing kernel sizes used for the three degraded resolutions considered in this work are listed in Table II. Note that for speed and simplicity the smoothing is carried out on the full-sky, rather than within a patch: the time taken to smooth in pixel-space scales poorly with patch size.

$N_{\text{side,deg}}$	N_{pix}	Pixel Scale (arcmin)	Smoothing FWHM
256	786432	13.7	33.9
128	196608	27.5	67.7
64	49152	55.0	135.2

TABLE II. The full-widths at half-maximum of the Gaussian kernels used to smooth input maps prior to degradation. Also tabulated are the pixel count at each resolution, and an approximate pixel scale, derived assuming each pixel is square.
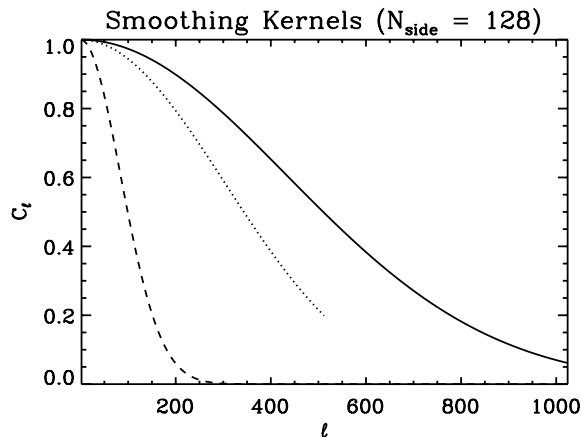

FIG. 7. Window functions of the various smoothing kernels appearing in the adaptive-resolution analysis. Shown are the WMAP7 W-band beam (approximated as a Gaussian, solid line), the Gaussian smoothing applied before degradation to $N_{\text{side,deg}} = 128$ (dashed), and the pixel window function at this resolution (dotted). Note that the pixel window function is only defined up to $\ell = 4N_{\text{side,deg}}$, the maximum multipole allowed by the `HEALPix` software.

B. Calculating the degraded evidence

Care must be taken when calculating the covariance matrix for use in the likelihood: the covariance matrix must include as faithful as possible a representation of the components of the data. We must therefore capture every important feature of both the CMB "signal" (the CMB is, in fact, the dominant noise in the analysis) and instrumental noise measured in the input map, as well as the effects of the degradation process. It is helpful to break up the full covariance matrix, C, into its CMB, S, and noise, N, constituents, as they are morphologically different.

1. CMB covariance

The CMB, as a correlated random field on the sphere, is most simply defined in harmonic space by its power spectrum, C_ℓ. At full resolution, the CMB power is smoothed by the instrumental beam, which we approximate in this analysis with a Gaussian of FWHM f_{WMAP}. In patches that are processed at reduced resolution, the CMB signal is also smoothed by the anti-aliasing beam (another Gaussian, of FWHM f_{deg}) and further by the pixel window function of the degraded resolution map. This final effect, shown for the relevant beams in Fig. 7, is small but non-negligible. Ignoring the degraded pixel window function means that the covariance contains an overestimate of the CMB power – in our analysis, a noise term – and log-evidences can be underestimated by as much as 1 when degrading to $N_{\text{side,deg}} = 64$.

Taking all of these effects into account, the CMB covariance between two pixels i and j is

$$S_{ij} = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell P_\ell(\cos\theta_{ij}) B^2_{\ell,\text{WMAP}} \tag{36}$$

if the patch is processed at full resolution, and

$$S_{ij} = \sum_\ell \frac{2\ell+1}{4\pi} C_\ell P_\ell(\cos\theta_{ij}) B^2_{\ell,\mathrm{WMAP}} B^2_{\ell,\mathrm{deg}} W^2_{\ell,\mathrm{deg}} \tag{37}$$

if the patch must be processed at degraded resolution. Here, $P_\ell(x)$ are the Legendre polynomials, $W_{\ell,\mathrm{deg}}$ is the $N_{\mathrm{side,deg}}$ pixel window function, and the B_ℓ are the WMAP and anti-aliasing Gaussian beams, represented in harmonic space as

$$B_\ell = \exp\left(\frac{-\ell(\ell+1)f^2}{16\log 2}\right), \tag{38}$$

where f is the relevant FWHM.

2. Instrumental noise covariance

At full resolution, the WMAP noise is uncorrelated in pixel space (see, e.g., Ref. [46]). The noise covariance matrix is therefore diagonal, and can be written as

$$N_{ij} = \frac{\sigma^2_{\mathrm{WMAP}}}{N_{\mathrm{obs},i}} \delta_{ij}, \tag{39}$$

where $\sigma_{\mathrm{WMAP}} = 6.549$ mK is the RMS noise of the W-band detectors, $N_{\mathrm{obs},i}$ is the number of times pixel i has been observed, and δ_{ij} is the Kronecker delta.

The anti-aliasing smoothing applied as part of the degradation process induces correlations between the noise measured in each pixel. Coupled with the variations in sky coverage represented by the N_{obs} values, this makes the exact pixel-space covariance more difficult to write down. Progress can be made by separating the full-resolution noise model into isotropic white noise, calculated from the power spectrum in harmonic space, which is modulated by a map encoding the variations in sky coverage, which belongs entirely in pixel space. The effects of the smoothing and degradation can then be applied individually to each component, and recombined in the covariance matrix.

Uncorrelated pixel noise is identical to the noise generated by white (i.e., flat) noise in harmonic space. To produce the correct RMS noise, the amplitude of the white noise power spectrum, N_ℓ, should be set such that

$$N_\ell = \frac{\sigma^2_{\mathrm{WMAP}}\Omega_{\mathrm{pix}}}{\bar{N}_{\mathrm{obs}}}, \tag{40}$$

where Ω_{pix} is the pixel area, and we have incorporated the mean number of observations per pixel, $\bar{N}_{\mathrm{obs}}$, into the definition. After smoothing and degradation, the noise power spectrum is no longer flat, having been multiplied by both the anti-aliasing beam and the degraded pixel window function (both squared).

Having absorbed the mean number of pixel observations into the isotropic noise, to include the effects of varying sky coverage we need only consider relative changes in the number of observations in each pixel. These are captured by generating a map of $\sqrt{\bar{N}_{\mathrm{obs}}/N_{\mathrm{obs},i}}$. After degradation, this map will have been smoothed and degraded exactly as the input temperature map; its mean at any resolution is one. The noise covariance matrix at degraded resolution is then

$$N_{ij} = \frac{\bar{N}_{\mathrm{obs}}}{\sqrt{N_{\mathrm{obs},i}N_{\mathrm{obs},j}}} \sum_\ell \frac{2\ell+1}{4\pi} N_\ell P_\ell(\cos\theta_{ij}) B^2_{\ell,\mathrm{deg}} W^2_{\ell,\mathrm{deg}}, \tag{41}$$

where the separation into components residing in pixel and harmonic space is clear. Note that this expression does not include the instrumental beam, as this is *detector* noise.

3. Full covariance

The sum over ℓ in the covariance matrix calculations is strictly a sum to infinity, but is approximated with a sum from $0 \le \ell \le \ell_{\mathrm{max,cov}}$, truncated at a point where all significant contributions have been included. At full resolution, the data are noise-dominated at high-ℓ (see Fig. 8), so the covariance is not particularly sensitive to the precise value
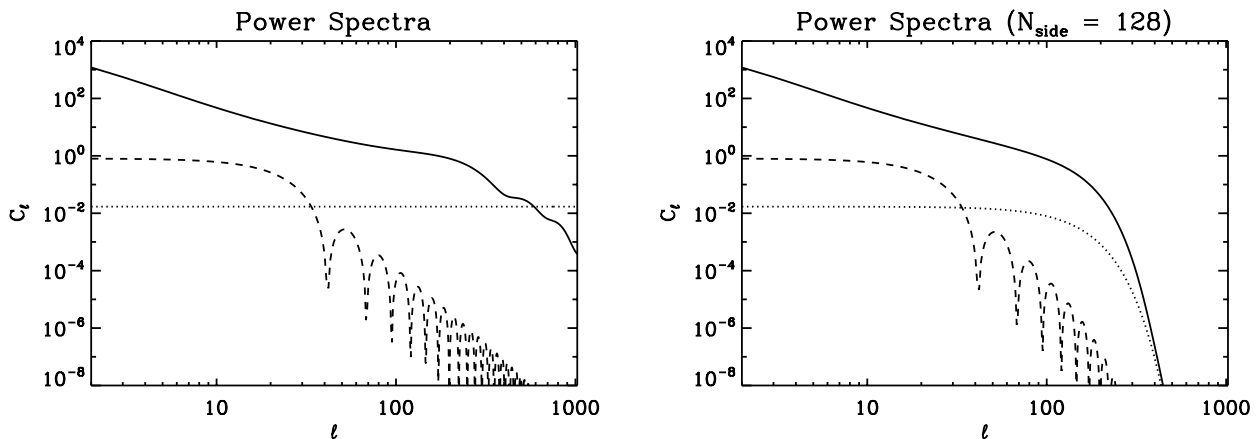
FIG. 8. Power spectra of the WMAP7 best-fit CMB signal (solid lines), WMAP7 noise (dotted) and a single simulated bubble collision (dashed) with amplitude $z_0 = 5 \times 10^{-5}$ and angular radius $\theta_{\rm crit} = 7°$. The power spectra are plotted at full WMAP resolution, $N_{\rm side} = 512$ (left), and after smoothing and degradation to $N_{\rm side,deg} = 128$ (right).

of $\ell_{\rm max,cov}$, provided it is larger than ~ 1000. We use $\ell_{\rm max,cov} = 1024$.[6] After degradation, the CMB and noise power spectra have been additionally damped by the extra smoothing and pixel window function. These effects combine to reduce the map power to $\sim 0.7\%$ of its maximum at $\ell_{\rm max} = 2N_{\rm side,deg}$, but the kernel's long tail means that there is still information at higher-ℓ. If this is ignored, the maps contain more information than is accounted for in the covariance matrix. The signal-to-noise ratio of any relic present is therefore artificially boosted, and the evidence is over-estimated: truncating at $2N_{\rm side,deg}$ yields over-estimates of ~ 5 in the log-evidence at $N_{\rm side,deg} = 64$. Tests reveal that evidence ratios are stable provided multipoles $\ell \gtrsim 3N_{\rm side,deg}$ are included in the covariance matrix. We choose to be conservative and push this to $\ell_{\rm max,cov} = 4N_{\rm side,deg}$.

The two components of the covariance are added to form the full covariance matrix.

$$\mathsf{C} = \mathsf{S} + \mathsf{N}. \tag{42}$$

This matrix must be inverted for use in the likelihood function; the inversion is carried out using the Cholesky decomposition implemented in the `LAPACK` software package. The full-resolution covariance matrix has a prominent diagonal due to the uncorrelated pixel noise, and hence is readily invertable. However, at lower resolutions the additional smoothing can make the covariance matrix nearly singular. Diagonal regularizing noise is therefore added at the level of $2\,\mu{\rm K}^2$ to all degraded-resolution covariance matrices to aid inversion: we have checked that this does not affect the results of the likelihood calculation.

C. Calculation of patch-based evidence ratios with MultiNest

To form the full posterior (Eq. 28), we must calculate the patch-based evidence ratios (Eq. 18) for each blob. Our pipeline uses the `MultiNest` sampler [34, 35], which outputs the normalized posterior and the evidence for the data in each patch under a specific model. This is not precisely the information required in Eq. 18, which integrates the product of the likelihood and the full-sky prior. However, we can use Bayes' theorem in the patch to convert the output posterior into a likelihood according to

$$\Pr_b(\mathbf{m}_0, \mathbf{m}_1 | \mathbf{d}) = \frac{\Pr(\mathbf{m}_0)\Pr_b(\mathbf{m}_1)\Pr_b(\mathbf{d}|\mathbf{m}_0, \mathbf{m}_1)}{\Pr_b(\mathbf{d})}. \tag{43}$$

Here, the subscript b indicates that the probabilities are formed using only the data for a single blob. In particular, we identify

$$\Pr_b(\mathbf{d}|\mathbf{m}_0, \mathbf{m}_1) \propto e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\rm T}/2} \tag{44}$$

---

[6] The WMAP7 observations [36] extend the power spectrum measurements to $\ell = 1200$, but the CMB signal-to-noise ratio for the W-band is below one for $\ell \gtrsim 600$. Setting $\ell_{\rm max,cov} = 1024$ ensures the CMB contribution is characterized well into the noise-dominated regime.

as the quantity necessary for the patch-based evidence ratio. The normalization assumed in Eq. 43 is

$$\int_b \mathrm{d}\mathbf{m}_0 \mathrm{d}\mathbf{m}_1 \mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}_b(\mathbf{m}_1) = 1, \tag{45}$$

which implies that $\mathrm{Pr}_b(\mathbf{m}_1)$ is related to the prior $\mathrm{Pr}(\mathbf{m}_1)$ in Eq. 18 by an overall normalization:

$$F_b = \frac{\mathrm{Pr}(\mathbf{m}_1)}{\mathrm{Pr}_b(\mathbf{m}_1)} = 1 - \int_{\bar{b}} \mathrm{d}\mathbf{m}_0 \mathrm{d}\mathbf{m}_1 \mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}(\mathbf{m}_1). \tag{46}$$

We also have

$$\mathrm{Pr}_b(\mathbf{d}) = \int_b \mathrm{d}\mathbf{m}_0 \mathrm{d}\mathbf{m}_1 \mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}_b(\mathbf{m}_1)\mathrm{Pr}_b(\mathbf{d}|\mathbf{m}_0,\mathbf{m}_1). \tag{47}$$

Using these expressions, we can solve for the patch-based evidence in terms of known quantities:

$$\rho_b(\mathbf{m}_0) = F_b \times \frac{\mathrm{Pr}_b(\mathbf{d})}{\mathrm{Pr}_b(\mathbf{d}|0,0)} \times \frac{1}{\mathrm{Pr}(\mathbf{m}_0)} \int_b \mathrm{d}\mathbf{m}_1 \mathrm{Pr}_b(\mathbf{m}_0,\mathbf{m}_1|\mathbf{d}), \tag{48}$$

where $\mathrm{Pr}_b(\mathbf{d}|0,0)$ is the likelihood in the patch with no sources.

All steps of the algorithm are parallelised wherever possible using OpenMP and MPI to take advantage of both shared- and distributed-memory clusters.

VI. ADAPTIVE-RESOLUTION TESTS

A. Stability of degraded evidence values

The adaptive-resolution analysis pipeline was tested thoroughly to estimate the effect the degrading has on the calculated evidence values. Naïvely, we would expect the evidence not to change (on average) for resolutions at which the feature is well-sampled, i.e., for which the pixel scale is much smaller than the feature itself. Once the resolution has decreased enough for the feature – and hence the template used to fit it – to appear pixelated, the evidence should begin to drop off sharply. In harmonic space this is equivalent to requiring that enough modes are left intact by the pre-degradation smoothing that the template power spectrum can be discerned.

The ideal test in this situation is to create a simulation containing a sufficiently large template to be un-pixelated at the lowest resolution at which a feature can be processed: in this case, $N_{\mathrm{side}} = 64$. This feature should then be processed at all resolutions considered, from the highest (WMAP-resolution) to the lowest, to determine how the evidence behaves. Unfortunately, the very nature of the problem makes this is impossible: such templates would require enormous covariance matrices at full resolution.

Progress can be made by breaking the test into parts. Simulations containing templates – strong bubble collision signatures, in this case – are generated on a range of angular scales. Each simulation is processed twice: first at the highest resolution possible, then again at one resolution lower. This indicates how the evidence changes when a patch is not processed at its "ideal" resolution, which could occur if the maximum angular size, θ_{crit}, is overestimated. The lower-resolution evidence values are calculated with four times fewer pixels: if the evidences returned by the two runs do not differ greatly, we can be confident that small reductions in resolution do not affect the evidence values returned.

Three maps are generated containing single bubble collision signatures of angular scale θ_{crit} equal to $7°$, $15°$ and $30°$, each of which contains approximately the same number of pixels at $N_{\mathrm{side}} = 512$, 256 and 128, respectively. In each case the signatures have the same amplitude, $z_0 = 8 \times 10^{-5}$, and position, $(\theta_0, \phi_0) = (45°, 45°)$; the maps also contain the same CMB and noise realizations, and are plotted in Fig. 9. The evidence is calculated once at the highest resolution possible – N_{side} of 512, 256 and 128 for the $7°$, $15°$ and $30°$ collisions, respectively – and once again after degrading one step in resolution – i.e., for N_{side} of 256, 128 and 64. In each pair of tests the same ranges in size and position are sampled, and no mask is used. The only small difference comes in the pixels included in the patch, as the lower-resolution patches sample a slightly larger region than the high-resolution patches.[7]

----

[7] When defining the patches using the `HEALPix` query_disc subroutine, the "inclusive" option is set to include all pixels which fall even partly within the radius we sample.
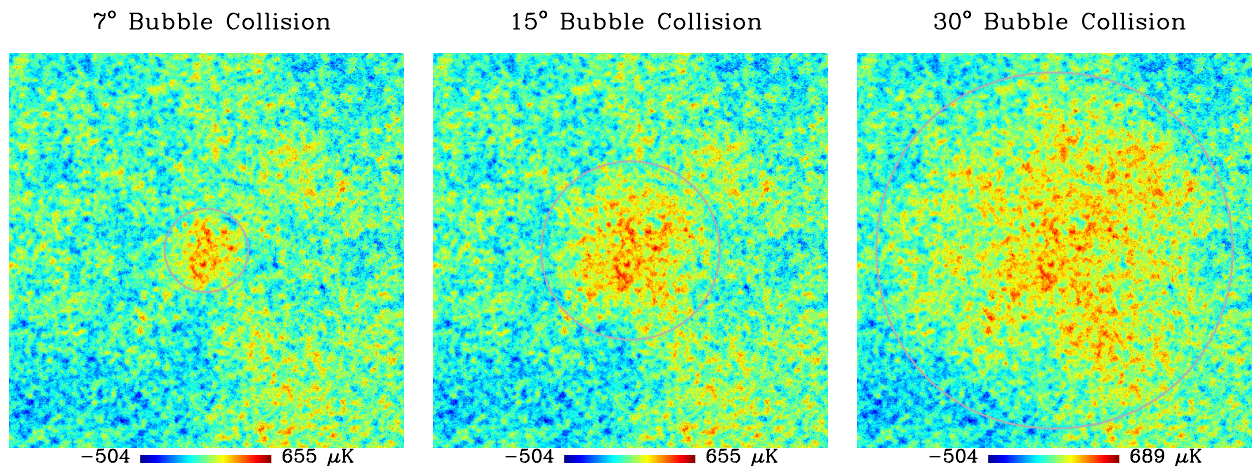
FIG. 9. The maps used to test the adaptive-resolution analysis. Three maps are generated, each containing a single bubble collision of radius $\theta_{\mathrm{crit}} = 7°$, $15°$ or $30°$. The maps contain identical CMB and noise realizations, and in each case the bubble collision is placed at $(\theta_0, \phi_0) = (45°, 45°)$ with amplitude $z_0 = 8.0 \times 10^{-5}$. The $7°$ and $30°$ collision maps are also used to test the effects of neglecting correlations with data outside the patch and the restriction of template locations to the regions highlighted by the candidate detection stage. All plots are shown at the same scale, and are $67°$ on a side.

Bubble θ_{crit} (°)	$N_{\mathrm{side,deg}}$	N_{pix}	$\log \rho$
7	512	41618	12.6 ± 0.3
7	256	10544	12.8 ± 0.3
15	256	40715	14.2 ± 0.3
15	128	10314	14.0 ± 0.3
30	128	49894	13.2 ± 0.3
30	64	12619	13.1 ± 0.3

TABLE III. Tests of the stability of the degraded evidence values. Three maps, each containing a small, medium or large simulated bubble collision, are used to examine how the evidence ratio changes when a patch is degraded from $N_{\mathrm{side}} = 512$ to 256, 256 to 128 and 128 to 64.

The results of the tests are shown in Table III. In each case, the evidence calculated at high resolution matches the evidence at low resolution to within `MultiNest` precision. We conclude that the adaptive-resolution analysis produces stable evidence ratios for the resolutions considered in this work. Note that the evidence does not always decrease: in fact it can increase. This is because the realization noise (i.e., the combined CMB and noise signal) is different after degradation: the bubble collision signature is compared to larger-scale modes at lower $N_{\mathrm{side,deg}}$. Further, note that the log-evidence errors tabulated are the typical random variations due to sampling, estimated by repeatedly testing an individual patch with initial conditions set from different random seeds. `MultiNest` also provides a statistical estimate of the error in an evidence calculation, derived from the relative entropy of the samples (see Refs. [34, 47]). We find these estimates (typically ~ 0.1 in log-evidence) to be subdominant to the variations due to sampling.

B. Robustness to smoothing-induced contamination

The simulations used to test the adaptive-resolution pipeline contain only a CMB realization, a noise realization and a bubble collision template; no mask was used. Real CMB datasets also contain foregrounds (or foreground residuals after component separation), the worst of which are masked. Due to the need to smooth prior to degradation, there is a potential for contaminants to leak from behind the mask.[8] While it is possible to mitigate this effect by extending the

---

[8] The smoothing procedure is fastest in harmonic space, where it is a multiplication rather than a convolution. Inclusion of the mask in this procedure is complex, and it is simplest to perform the smoothing on full-sky data.

mask (and potentially using a smoothing kernel that is localized in pixel space) [43], it is highly undesirable to discard hard-won data. The likely scale and amplitude of any smoothing-induced contamination is therefore investigated to determine whether the mask should be extended.

The masks recommended by the WMAP team [42] comprise two components: a Galactic cut (of varying conservatism) and a point-source mask. The point-source mask is created from a range of external and internal catalogs (as listed in Ref. [48]), and is updated with each data release. The point-sources, which appear in the data as approximately Gaussian peaks with the same FWHM $(0.22°)$ as the instrumental beam, are masked by excising a region of radius $\sim 0.6°$ centered on each point-source. A small number of the strongest sources are more aggressively masked, out to a $\sim 1.2°$ radius. The cut made to remove the extended emission of the Milky Way is much larger, and forms an irregular band $\sim 20° - 40°$ in width centered on the Galactic Plane. For clarity, the effects of the two components of the mask are investigated individually.

We can estimate the effects of point-source smoothing using a very simple test. A W-band point-source is simulated by placing a normalized $1\,\mu$K delta-function at a position taken from the WMAP point-source catalog [42], then convolving it with a Gaussian of $0.22°$ FWHM. Since smoothing is a linear process, this can then be scaled to investigate the effects of sources of different amplitudes. This map is then smoothed and degraded to each scale used in the analysis (i.e., $N_{\rm side} = \{256, 128, 64\}$), and masked using the degraded point-source mask at each resolution. The resulting maps, plotted in Fig. 10, can then be scaled to mimic a source of a given temperature. The plots show that, at all resolutions considered, the maximum contamination injected into a single pixel is a few thousandths of the point-source's amplitude. Assuming that such sources have amplitudes of $100 - 1000\,\mu$K [49], these results suggest that our degradation technique induces contaminants of at most $1 - 2\,\mu$K into fewer than 10 pixels. This level is completely subdominant to the CMB signal, and so should not affect the analysis. We therefore need not extend the point-source mask when smoothing and degrading.

An estimate of the contaminants leaked from the Galactic cut can be obtained in a similar fashion. In place of the simulated point source, we must make an estimate of the Galactic foreground residuals. Modeling this precisely is difficult, so, following Ref. [50], 1% of the difference between the V-band signal and the WMAP7 Internal Linear Combination (ILC) map [42] is used for illustrative purposes. This combination is indicative of the morphology and amplitude of residuals within the ILC: there are contaminants of around 50 times the amplitude visible to the eye in the WMAP foreground-reduced maps. As with the simulated point-source, we take this map, smooth and degrade it to $N_{\rm side}$ values of 256, 128 and 64, and then mask using only the Galactic portion of the 7-year KQ85 mask. The resultant maps are plotted in Fig. 11, along with the input. The extra smoothing creates a strip, a few pixels wide, of contamination around the Galactic mask, typically at a level of $0.3 - 0.4\,\mu$K. Scaling this up by a factor of 50 yields contaminants of $\sim 20\,\mu$K. Although this is an order of magnitude higher than that created by point-sources, it is extremely localized, and does not mimic any of the target signals. We conclude that, as with the point-source mask, there is no need to extend the Galactic cut when degrading. Our adaptive-resolution algorithm should be robust to smoothing-induced contamination, a hypothesis that will be further tested at a later point by processing a null simulation containing realistic foreground residuals.

VII. TESTING THE FORMALISM'S APPROXIMATIONS

The formalism set out in Section II includes three main approximations: that the parameters of ΛCDM can be set to the best-fit measured values, that the likelihood need only be integrated over ranges of $\{\theta_0, \phi_0\}$ (and, indeed, $\theta_{\rm crit}$) corresponding to patches containing candidate sources, and that correlations between data inside and outside these patches can be neglected. The error made in the latter two approximations will depend on the nature of the source templates. In particular, the accuracy of the second approximation will depend on how well the candidate regions cover the underlying sources, and how the likelihood falls off as a function of the location and size of the feature. It is also important to note that the approximations will most likely improve as a function of the signal-to-noise for the source: if the likelihood is not peaked in any region of parameter space, we are not justified in choosing to integrate only over particular regions.

We begin by discussing the assumption that the parameters of ΛCDM can be set to the best-fit measured values. The only place these parameters enter the evidence calculation is through the pixel-pixel covariance matrix used to formulate the likelihood function. To get an idea for how this might vary with the parameters of ΛCDM, in Fig. 12 we show the pixel-pixel covariance as a function of angle for three cosmologies: the fiducial best-fit WMAP7+BAO+H_0 cosmology and two cosmological models where the dark energy density, Ω_Λ, is two standard deviations higher or lower than the best-fit value (chosen because varying Ω_Λ significantly changes power on the largest scales, which are most relevant for the signatures considered). The difference at the level of the covariance is very small, and the difference in the inverse covariance appearing in Eq. 5 is correspondingly small. We can write the inverse covariance as

$$\mathsf{C}(\mathbf{m}_{\Lambda{\rm CDM}})^{-1} = \mathsf{C}(\bar{\mathbf{m}}_{\Lambda{\rm CDM}})^{-1} + \Delta\mathsf{C}(\mathbf{m}_{\Lambda{\rm CDM}})^{-1}, \tag{49}$$
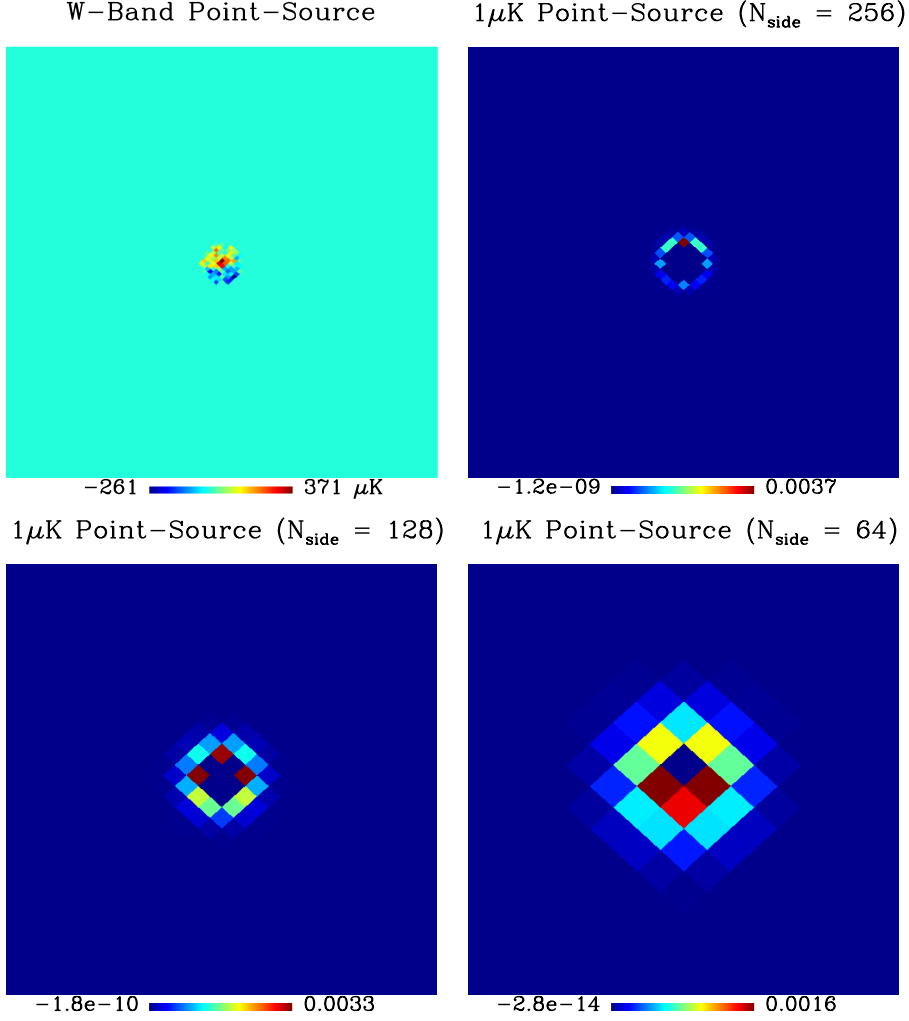
FIG. 10. A typical WMAP W-band point-source contaminant (top-left, plotted using the inverted point-source mask), and the effects of smoothing a unit-amplitude simulated point-source and degrading to $N_{\text{side,deg}} = 256$, 128 and 64. All low-resolution plots are shown with a degraded mask applied, which limits any residuals to $\mathcal{O}(\mu\text{K})$. All plots are 12.5° on a side.

where $\Delta\mathsf{C}(\mathbf{m}_{\Lambda\text{CDM}})^{-1}$ is the small difference in the inverse covariance resulting from varying cosmological parameters. Neglecting the effect of changing the cosmological parameters amounts to requiring

$$\mathbf{d}\Delta\mathsf{C}(\mathbf{m}_{\Lambda\text{CDM}})^{-1}\mathbf{d}^{\text{T}} \gg \mathbf{d}\Delta\mathsf{C}(\mathbf{m}_{\Lambda\text{CDM}})^{-1}\mathbf{t} \gg \mathbf{t}\Delta\mathsf{C}(\mathbf{m}_{\Lambda\text{CDM}})^{-1}\mathbf{t}, \tag{50}$$

which is natural for small-amplitude templates that have compact support on the sky. In this case, the $N_{\text{s}} = 0$ and $N_{\text{s}} = 1$ terms in the expansion Eq. 4 would take the form

$$\Pr(\mathbf{d}|N_{\text{s}} = 0) = \frac{1}{(2\pi)^{N_{\text{pix}}/2}|\mathsf{C}|}e^{-\mathbf{d}\mathsf{C}(\bar{\mathbf{m}}_{\Lambda\text{CDM}})^{-1}\mathbf{d}^{\text{T}}/2} \int d\mathbf{m}_{\Lambda\text{CDM}}\Pr(\mathbf{m}_{\Lambda\text{CDM}})e^{-\mathbf{d}\Delta\mathsf{C}(\mathbf{m}_{\Lambda\text{CDM}})^{-1}\mathbf{d}^{\text{T}}/2} \tag{51}$$

$$\Pr(\mathbf{d}|N_{\text{s}} = 1) \simeq \Pr(\mathbf{d}|N_{\text{s}} = 0)\rho_b, \tag{52}$$

with the patch-based evidence ratio ρ_b defined as before (i.e. with the covariance defined for fixed cosmological parameters and confined to pixels in the blob). Higher order terms would follow suit. Therefore, in this approximation all of the dependence on cosmological parameters would be absorbed into an overall renormalization of $\Pr(\mathbf{d}|N_{\text{s}} = 0)$, and would not affect our final answer for the posterior of the global model parameters.

To explore the latter two assumptions, consider the single-source contribution to the posterior Eq. 10. Neglecting correlations between regions inside and outside the patch (see Eq. 15) corresponds to the assumption that

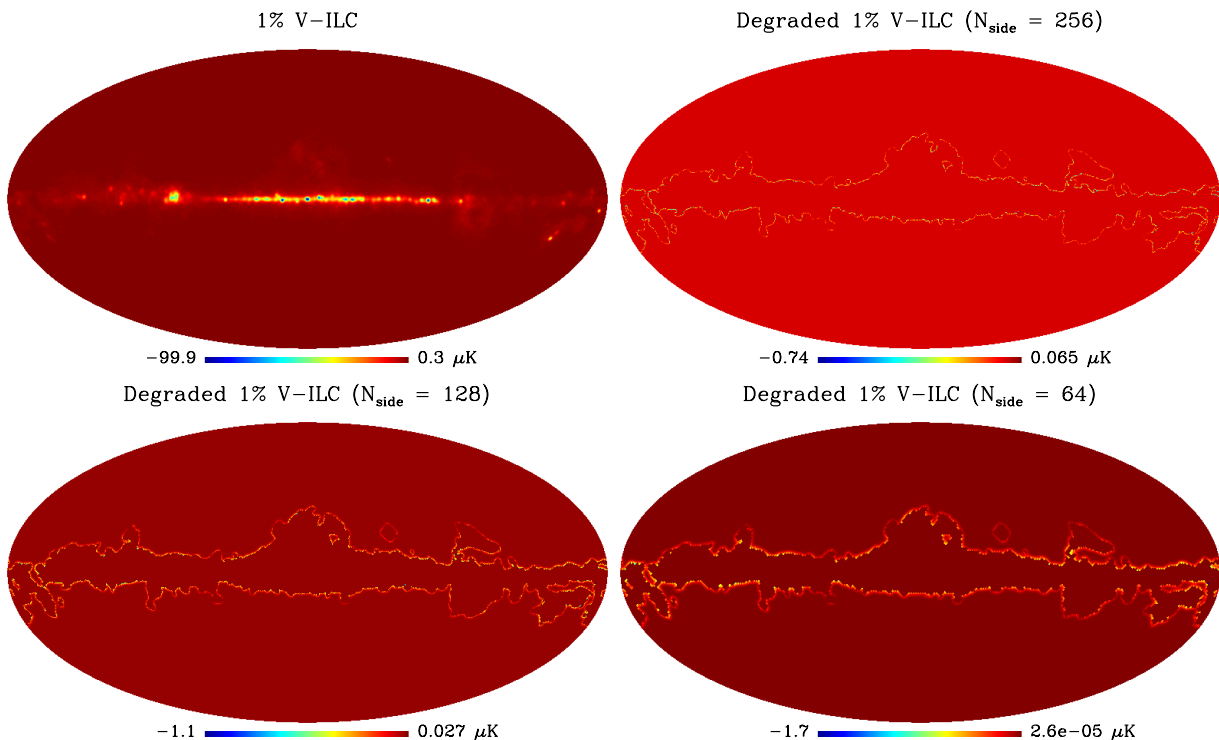$$\mathbf{t}\mathsf{C}^{-1}\mathbf{d}_b^{\text{T}} \ll 1, \tag{53}$$

FIG. 11. One percent of the difference between the WMAP V-band and ILC maps provides an estimate of the Galactic contamination present in WMAP data products. This map is plotted here (top-left), alongside the residuals produced by smoothing and degrading it to $N_{\text{side,deg}} = 256$, 128 and 64. All low-resolution plots are shown with a degraded KQ85 Galaxy-only mask applied.

where the template has support in a blob b, and the data consist of pixels in region $\bar{b}$ outside the blob. In Fig. 13, we plot the inverse covariance between several positions and the rest of the sky (e.g., a set of rows of the inverse covariance matrix) in ΛCDM using the best-fit WMAP 7-year cosmological parameters, keeping only the first 50 multipole moments. It can be seen that the inverse covariance is only significant within a disk of radius $\sim 15°$ around each of the template pixels. Therefore, we need not retain all of the pixels on the sky. In Fig. 14, we depict the case where the likelihood is peaked for templates inside a region well-contained within the blob (shaded disk). The inverse covariance will be significant within a disk (dashed circle) of $\sim 15°$ around each pixel where the template is non-zero (black dot). Our approximation neglects correlations between the template and the pixels contained within the dashed circle, but outside the blob (i.e., in region $\bar{b}$). The exponential will clearly yield a decreasing correction to the integral as the size of the blob is increased, becoming vanishingly small when the radius of the blob is $\sim 15°$ larger than the size of templates near the maximum of the likelihood. For terms of higher order in N_s, the correction is slightly more complicated, but for blobs separated by a distance greater than $15°$, the argument is the same. Another assumption we have made is encoded in Eq. 16: that the pixels in region $\bar{b}$ do not contribute to the inverse of the covariance inside region b. Further, if the actual source is well-contained within the blob, the likelihood will presumably peak in a region well-contained within the blob (which is the assumption behind our first approximation of integrating over region b alone).

The reasoning set out above provides qualitative support for the approximations that make this analysis feasible. To determine quantitatively how good these approximations are, we have performed three numerical tests. The first two tests are designed to determine whether correlations with pixels outside the blobs can indeed be discounted; the third determines the effects of restricting the position integral to within our candidate blobs.

A. Tests of neglected correlations

The ideal test of the effect of neglected correlations would be to perform evidence integrals for simulated collisions of varying sizes using covariance matrices ranging from patch-sized to full-sky. Unfortunately, memory restrictions mean we can only hold the full-sky covariance matrix in memory for `HEALPix` resolutions smaller than $N_{\text{side}} = 64$. At
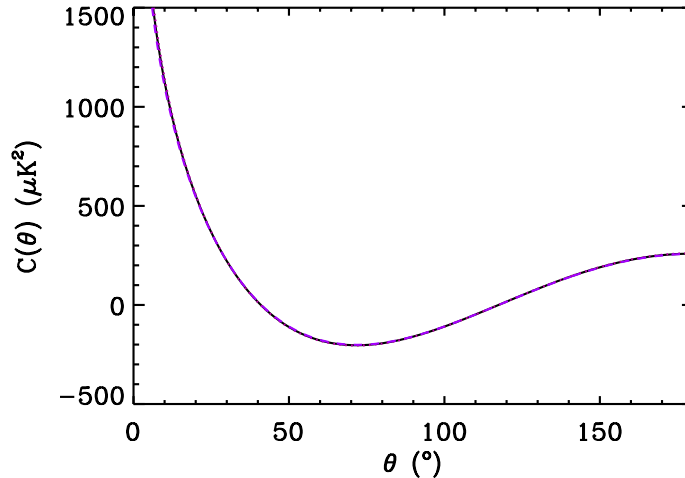
FIG. 12. A plot of the angular correlation function of the CMB for three cosmological models: the best-fit cosmological parameters from WMAP 7-year CMB, BAO and H_0 [36] (solid), a cosmological model where the dark energy density Ω_Λ is two standard deviations too high (dashed), and a cosmological model where Ω_Λ is two standard deviations too low (dotted). The three curves are nearly indistinguishable.
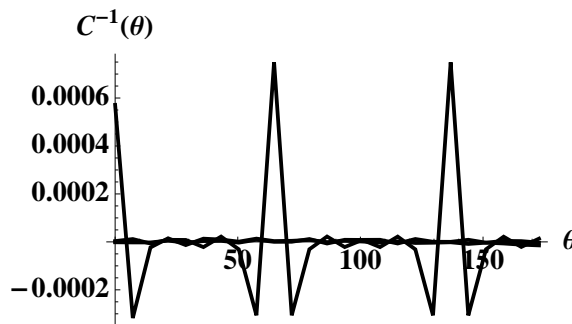

FIG. 13. The inverse of the pixel-space ΛCDM correlation function between $\theta = 0$ and all θ, $\theta = 72°$ and all θ, and $\theta = 144°$ and all θ (left to right). It can be seen that C^{-1} is largest in magnitude over a $\sim 30°$ window around the pixel being correlated.

this resolution each pixel is $\sim 1°$ across, so only large collisions are faithfully represented. As with the tests of the adaptive resolution analysis, we therefore split the test up, performing two tests: one checking the effect of neglected correlations on the largest scales, and the other at smaller scales.

The first test uses the map simulated for the largest-scale degradation tests, containing a $30°$ bubble collision template placed at $(\theta_0, \phi_0) = (45.0°, 45.0°)$ with amplitude $z_0 = 8 \times 10^{-5}$ (as plotted in Fig. 9). The simulation is passed to the candidate detection stage, and the patch corresponding to the bubble collision is singled out. The evidence ratio is then calculated twice at $N_{\mathrm{side,deg}} = 64$, first using the covariance matrix calculated for the patch, and second using the covariance matrix calculated on the full-sky. For clarity, the test is carried out without masking, and the integration limits on the template parameters are kept constant for both runs. The difference between evidence ratios returned will indicate the scale of any error induced by neglecting correlations at the largest scales.

The second test reuses the smallest map considered in the degradation tests (again, plotted in Fig. 9), containing a $7°$ bubble collision template with the same amplitude and position as in test one. As in the first test, the candidate detection algorithm is applied, and the feature containing the template is extracted. The evidence is calculated first using the standard patch size $15.4°$ in radius, then using progressively larger patches of sky until the patch is $30°$ in radius. At this point, Fig. 13 implies that the covariance matrix contains all pixels significantly correlated with those in the candidate collision region. The difference between evidence ratios will therefore indicate the errors associated with neglecting correlations on smaller scales. As with the first test, the integration limits used in each case are the same, as is the resolution ($N_{\mathrm{side,deg}} = 256$) at which the calculation is performed, and no mask is used.

The results of the two tests are presented in Tables IV and V. In each case, increasing the patch size does not
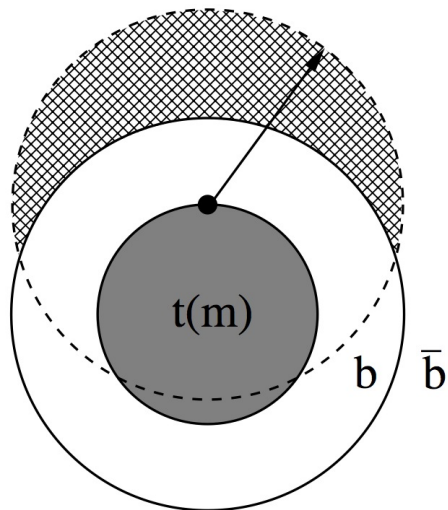
FIG. 14. The case where the source is well-contained within the blob. There is a clearly peaked likelihood for templates contained within the shaded disk. For the pixel denoted by the black dot, the inverse covariance is significant only within the dashed circle. Our approximation neglects correlations with pixels in the hatched region. This approximation is worst for pixels on the edge of the template.

$\theta_{\mathrm{patch}}(^\circ)$	N_{pix}	$\log \rho$
60.1	12619	13.1 ± 0.3
180.0	49149	13.3 ± 0.3

TABLE IV. The evidence ratios obtained when a patch covariance matrix is used versus the full-sky covariance matrix. Note that the full-sky covariance matrix does not quite cover the entire sky: three pixels are left out. This is a consequence of the patch-based nature of the algorithm, and does not affect the conclusions.

change the evidence value obtained beyond `MultiNest` precision. This supports that the assumption that correlations outside of the patch can be neglected, and indicates that doing so does not add a significant source of systematic error, given the scale of the variations induced by the nested sampler.

B.   Test of localization of likelihood peaks

The third assumption test is designed to assess the claim that the likelihood is peaked in position space, and that the evidence values obtained are unchanged when only the peaks are considered; this amounts to changing the limits of integration in Eq. 18. This test uses the same 7°-bubble collision map as in the second test. This time, the patch radius is held constant at 30°, but the set of central positions sampled is incrementally increased until the template

$\theta_{\mathrm{patch}}(^\circ)$	N_{pix}	$\log \rho$
15.4	14670	13.0 ± 0.3
20.0	24167	13.1 ± 0.3
25.0	37388	13.2 ± 0.3
30.0	53338	13.0 ± 0.3

TABLE V. The evidence ratios obtained when the size of the patch covariance matrix is incrementally increased until all correlations are included.

Center Range $/\theta_{\mathrm{patch}}$	Center Range (°)	$\log \rho$
10%	1.4	13.0 ± 0.3
25%	3.5	13.0 ± 0.3
50%	7.0	13.0 ± 0.3
75%	10.5	13.1 ± 0.3
100%	14.0	12.8 ± 0.3

TABLE VI. The evidence ratios obtained when the range of collision centers sampled is incrementally increased until the collision can be placed anywhere within the patch.

can be centered anywhere within the entire patch. The nested sampler therefore has access to greater portions of the collision environs – by the fourth run it can sample central positions placing the templates entirely outside the simulated collision region – and can provide an estimate of how much evidence is discarded by restricting the template position to lie within the blob.

The results of the test are reported in Table VI. The evidence ratio is stable, indicating that sampling from a larger range of central positions does not affect the outcome. This implies that the likelihood is indeed well-localized, supporting the assumption that the evidence integration need only be carried out over a restricted range of positions. The likelihood is plotted as a function of the two angular coordinates in Fig. 15, and indeed is very strongly peaked about the position highlighted by the candidate detection stage. For comparison, we also plot a collision with the same θ_{crit} but with an amplitude of $z_0 = 2.5 \times 10^{-5}$, which is near the lower-limit of our sensitivity (from Fig. 5 only 1% of such sources would be flagged by the candidate-detection algorithm). Even in this case, the likelihood is strongly peaked near the centre of the collision.

VIII. NULL TEST: ANALYSIS OF WMAP END-TO-END SIMULATION

The WMAP data contain a number of components that cannot be included in the pixel-space covariance matrix; in particular, the foregrounds are not known precisely, and so their subtraction leaves behind unknown, highly anisotropic residuals. It is therefore important to apply the Bayesian analysis pipeline on a null dataset containing estimated foreground residuals – and any other potential systematic effects – to determine whether they generate false-positive results. The full end-to-end simulation of the WMAP experiment used to set the optimal-filter detection thresholds is perfectly suited to the task, as it is created from simulated time-ordered data, including foreground contamination, and is processed by exactly the same pipeline as the WMAP observations.

The raw optimal filter analysis of the WMAP 7-year W-band end-to-end simulation (with the KQ85 mask applied) generates 19 bubble collision candidates and 10 texture candidates. Any candidates which are heavily masked are discarded, as are candidates whose estimated range of sizes has no overlap with the relevant prior on θ_{crit}. Finally, any candidates which clearly correspond to a single feature are merged, leaving a set of 12 and 4 bubble collision and texture candidates, respectively. The sizes and locations of these candidates as estimated by the optimal filters are tabulated in Table VII.

The patch evidence ratios obtained by applying the adaptive-resolution Bayesian evidence calculation to these candidates are also presented in Table VII. As expected, no candidate exhibits strong evidence in favor of either the bubble collision or texture model: the maximum evidence ratios are $e^{-5.4}$ and $e^{-4.0}$, respectively. Merging these results produces the posteriors on the global model parameters as plotted in Fig. 16. The posteriors for both models are clearly peaked at $\bar{N}_{\mathrm{s}} = 0$, confirming our prior knowledge that there are no bubble collisions or textures in the end-to-end simulation of the WMAP W-band data. The null test therefore indicates that we should not expect un-modeled foreground residuals and unknown systematics to generate false positives in the WMAP data. Further, the maximum evidence ratios obtained provide indicators of the level of response foregrounds and systematics can produce: any features found in the data should exceed these values in order to be considered interesting.

IX. ANALYSIS OF WMAP 7-YEAR DATA

Our analysis of the WMAP 7-year W-band foreground-reduced temperature map (with the KQ85 mask applied) produces a total of 32 bubble collision candidates and 33 texture candidates. The candidates' locations are visually
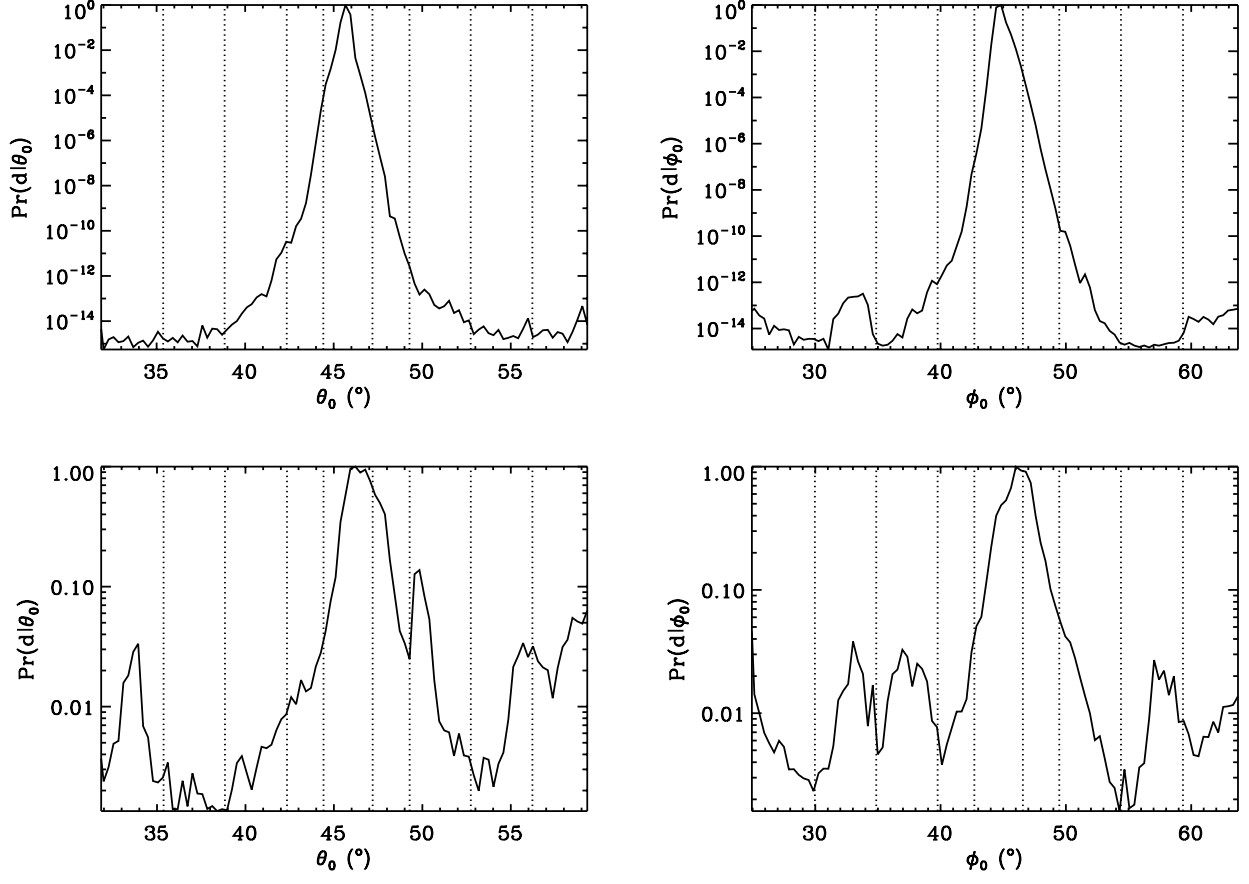
FIG. 15. The likelihood as a function of co-latitude θ_0 (left) and longitude ϕ_0 (right), plotted on a logarithmic scale. The top row is for a collision with $z_0 = 8 \times 10^{-5}$ (as used to calculate the localization-test evidence ratios); the bottom row is for a collision with $z_0 = 2.5 \times 10^{-5}$ (shown for comparison). Overplotted as dashed lines are the integration limits used in each of the five runs testing how the evidence changes as the central positions are sampled from larger regions. These limits correspond to sampling central positions from 10%, 25%, 50%, 75% and 100% of the patch by angular radius. The likelihood is strongly peaked in both angular coordinates.

ID	θ_{crit} range (°)	$\theta_0(°)$	$\phi_0(°)$	$\log \rho$		ID	θ_{crit} range (°)	$\theta_0(°)$	$\phi_0(°)$	$\log \rho$
1	2-4	63.1	142.0	-9.68 ± 0.3		1	2-5	50.2	221.5	-4.41 ± 0.3
2	3-5	50.2	221.5	-10.87 ± 0.3		2	4-9	45.8	153.1	-6.01 ± 0.3
3	6-8	126.4	214.8	-8.75 ± 0.3		3	4-15	143.2	3.2	-3.99 ± 0.3
4	8-12	43.8	152.7	-8.57 ± 0.3		4	7-50	38.7	94.8	-6.10 ± 0.3
5	10-14	19.1	326.3	-9.21 ± 0.3						
6	14-18	139.7	11.3	-7.43 ± 0.3						
7	20-26	137.0	42.7	-6.74 ± 0.3						
8	30-40	56.8	154.0	-6.01 ± 0.3						
9	35-45	118.6	232.4	-5.83 ± 0.3						
10	60-70	102.0	326.6	-6.43 ± 0.3						
11	80-90	86.1	305.9	-6.64 ± 0.3						
12	85-90	53.9	19.7	-5.43 ± 0.3						

TABLE VII. The size ranges and locations of the final 12 candidate bubble collisions (left) and 4 candidate textures (right) found in the end-to-end simulation of the WMAP experiment, along with their patch evidence ratios. The angular positions are related to Galactic longitude, l, and latitude, b, by the transformations $l = \phi$ and $b = 90° - \theta$.
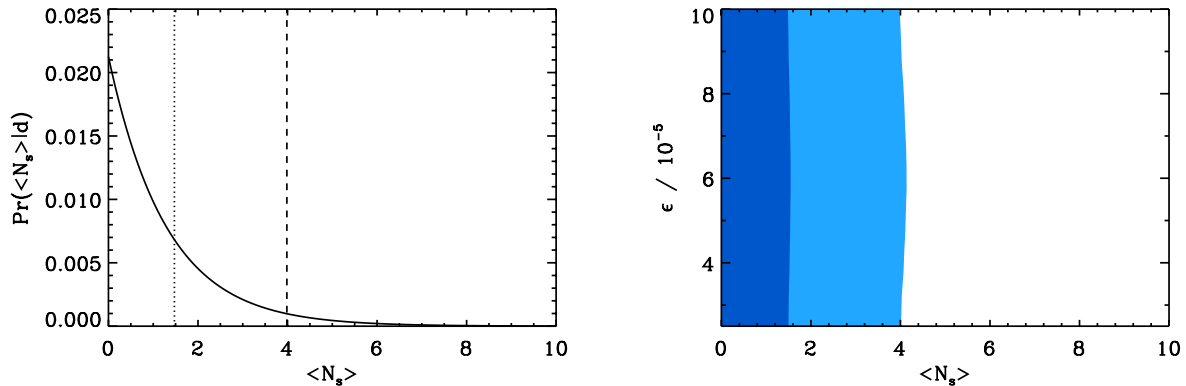
FIG. 16. The posterior probabilities of the global parameters of the bubble collision (left) and texture (right) models, given the end-to-end simulation of the WMAP 7-year W-band. The posterior is plotted as a function of one parameter, $\bar{N}_s$, for bubble collisions, and two parameters, $\bar{N}_s$ and ϵ, for textures. The regions containing 68% and 95% of the posterior probability are indicated by the dotted and dashed lines in the bubble collision plot, and as dark and light regions in the texture plot. Both posteriors are strongly peaked at zero sources.

ID	θ_{crit} range (°)	$\theta_0(°)$	$\phi_0(°)$	$\log\rho$		ID	θ_{crit} range (°)	$\theta_0(°)$	$\phi_0(°)$	$\log\rho$
1	2-6	169.0	187.5	-8.82 ± 0.3		1	1-3	114.6	22.1	-6.08 ± 0.3
2	3-5	167.2	268.7	-7.79 ± 0.3		2	1.5-5	168.7	184.4	-3.12 ± 0.3
3	4-14	147.4	207.1	-5.80 ± 0.3		3	1.5-4	166.8	268.8	-4.58 ± 0.3
4	4-8	123.2	321.3	-7.17 ± 0.3		4	1.5-3	72.4	150.1	-5.57 ± 0.3
5	6-8	62.7	220.4	-9.33 ± 0.3		5	2-15	123.5	321.0	-3.31 ± 0.3
6	14-22	136.6	176.5	-5.96 ± 0.3		6	2-7	123.2	79.5	-4.73 ± 0.3
7	12-20	118.0	212.0	-7.39 ± 0.3		7	2-4	128.3	93.5	-5.56 ± 0.3
8	20-24	75.8	168.8	-9.46 ± 0.3		8	2-15	147.4	210.2	-2.06 ± 0.3
9	28-35	85.8	166.3	-7.04 ± 0.3		9	5-10	69.3	202.1	-8.04 ± 0.3
10	22-40	126.8	220.1	-5.67 ± 0.3		10	8-15	119.3	155.7	-7.71 ± 0.3
11	80-90	69.6	61.9	-6.66 ± 0.3		11	4-15	135.4	172.1	-5.48 ± 0.3
						12	7-50	120.7	220.4	-5.84 ± 0.3

TABLE VIII. The size ranges and locations of the final 11 candidate bubble collisions (left) and 12 candidate textures (right), along with their patch evidence ratios. The size ranges tabulated are those derived from the optimal filter analysis; where they extend beyond the relevant priors they are truncated before the evidence calculation. The angular positions are related to Galactic longitude, l, and latitude, b, by the transformations $l = \phi$ and $b = 90° - \theta$. Note that, although bubble collision candidates 8 and 9 overlap considerably, their positions and size ranges are sufficiently discrepant to justify processing individually.

inspected, and those which are mostly obscured by the mask are discarded; candidates found to have no overlap with the relevant priors on θ_{crit} are likewise cut. Any candidates which are obviously coincident are merged at this point. The remaining candidates are then required to also be significant in an optimal filter analysis of the WMAP 7-year V-band foreground-reduced temperature map. This simple check requires that each feature is interesting across a range of frequencies, indicating that it is not due to foregrounds. This final cut leaves a set of 11 and 12 bubble collision and texture candidates respectively. The most probable sizes and locations of these candidates are tabulated in Table VIII and plotted in Fig. 17.

Applying the adaptive-resolution evidence calculation to the candidates produces the patch evidence ratios also reported in Table VIII. No single candidate is strong enough to claim a detection on its own. However, as demonstrated in Refs. [9, 11], it is possible for a number of weak candidates to favor the addition of relics to ΛCDM even if their individual evidence ratios are less than one: only by combining the results obtained for all candidates can the overall predictive power of the underlying model be revealed. The posteriors on the global parameters of the bubble collision
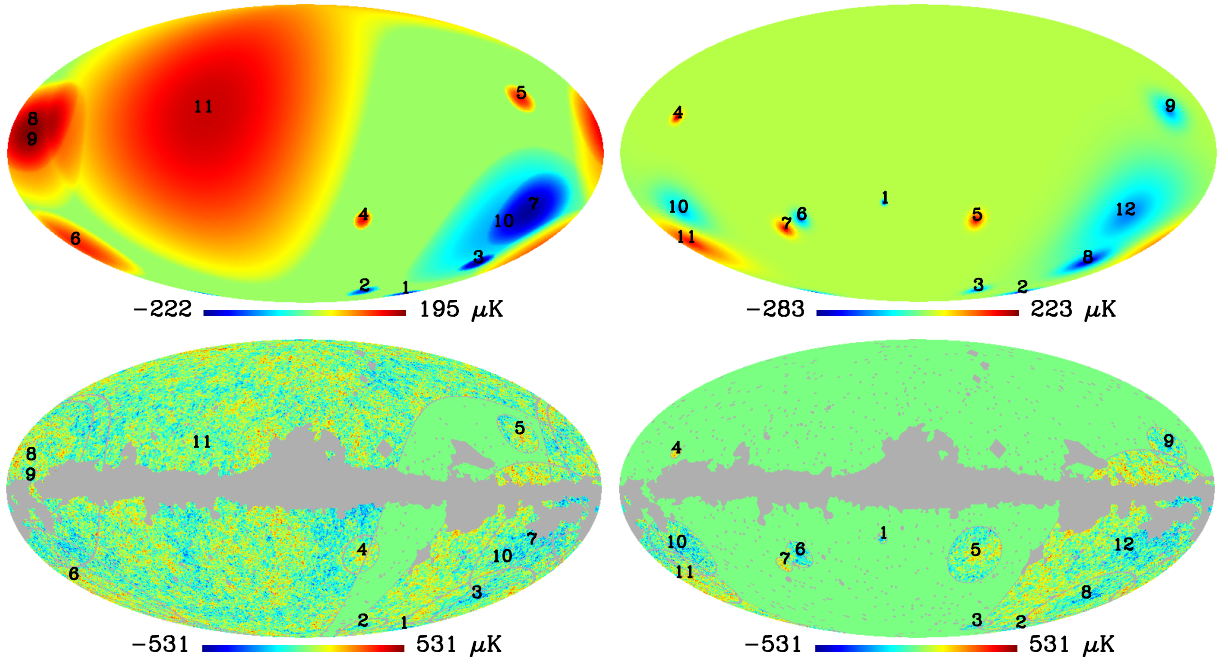
FIG. 17. Top: the estimated sizes and amplitudes of the bubble collision (left) and texture (right) candidates located in the WMAP 7-year data by the optimal filters. Bottom: the patches of WMAP 7-year data passed to the Bayesian evidence calculation for each of these candidates. The bubble collision plot shows all of the data involved in the evidence calculation for each candidate; for clarity, the texture plot only shows the core region of each patch. Note that the plots of the optimal filter candidates (top) contain only the estimated contributions due to additional sources, and the temperature ranges therefore differ from the plots of the WMAP data (bottom).
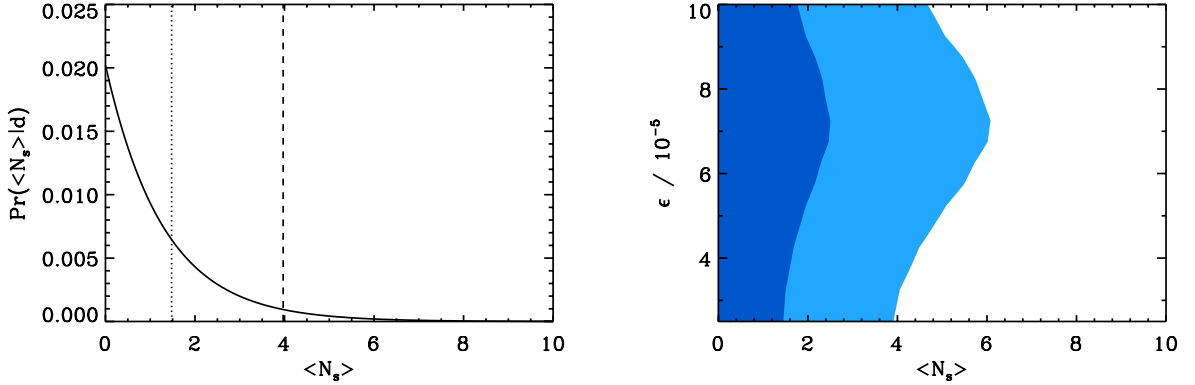

FIG. 18. The posterior probabilities of the global parameters of the bubble collision (left) and texture (right) models, given the WMAP 7-year data. The posterior is plotted as a function of one parameter, $\bar{N}_s$, for bubble collisions, and two parameters, $\bar{N}_s$ and ϵ, for textures. The most probable regions containing 68% and 95% of the posterior probability are indicated by the dotted and dashed lines in the bubble collision plot, and as dark and light regions in the texture plot. Both posteriors are strongly peaked at $\bar{N}_s = 0$.

and texture models, derived by combining the results from the candidates, are plotted in Fig. 18: both posteriors are peaked at zero sources. The texture model's dimensionless scale of symmetry breaking is constrained to be $2.6 \times 10^{-5} \leq \epsilon \leq 1.0 \times 10^{-4}$ (at 95% confidence), which, as the prior is defined only within the range $2.5 \times 10^{-5} \leq \epsilon \leq 1.0 \times 10^{-4}$, indicates that the WMAP data do not provide any interesting constraint on this parameter.

The WMAP 7-year data do not favor the addition of either bubble collisions or textures to ΛCDM. As none of the candidates exhibits significant evidence for the addition of sources to ΛCDM, we do not check the candidates for foreground residuals.

source	new ID	old ID	expected $\Delta \log \rho$	observed $\Delta \log \rho$
bubble	3	2	-2.1	-1.2
bubble	4	3	-4.1	-3.1
bubble	1	7	-4.0	-3.4
bubble	2	10	-3.8	-4.0
texture	2	6	-0.1	-0.5
texture	3	8	-0.1	-0.4
texture	8	9	-0.1	-0.2
texture	5	10	-0.1	-0.2

TABLE IX. The expected and observed differences in evidence ratio found for the four best bubble collision and texture candidates between the previous incarnation of the pipeline and the present analysis. Included in the expected change are the new form of the θ_{crit} prior for the bubble collisions, the removal of the edge from the bubble collision template, and the change in the mask. Note that the evidence ratios and IDs of the texture candidates were not previously published.

X. DISCUSSION

In Refs [9, 10] and [11], searches for bubble collisions and textures using earlier versions of the Bayesian source detection pipeline were published. Each previous analysis shares a number of candidate features in common with the current analysis, allowing consistency checks to be carried out between versions of the pipeline. Comparing results between versions is non-trivial, and must take into account each change made to the algorithm. In particular:

1. The prior on the bubble collision size has changed from uniform in the range 2-11.25° to being proportional to $\sin \theta_{\mathrm{crit}}$ in the range $2 - 90°$. *Ceteris paribus*, this will reduce evidence ratios previously reported for bubble collision candidates, particularly those at small scales.

2. The bubble collision template previously allowed for a discontinuity at the template boundary with amplitude z_{crit}. This parameter is now set to zero due to updated theoretical results [29, 30], and the bubble collision model considered in this analysis is consequently nested within the model considered previously. The effects of removing the edge can be determined exactly using the Savage-Dickey Density Ratio [51]: the change in evidence will be the ratio of the posterior and prior probabilities of the edge amplitude, evaluated at $z_{\mathrm{crit}} = 0$ using the results of the previous analysis, i.e.,

$$\Delta \log \rho = \log \frac{\Pr(z_{\mathrm{crit}}|\mathbf{d}, \mathrm{old})}{\Pr(z_{\mathrm{crit}}, \mathrm{old})} \bigg|_{z_{\mathrm{crit}}=0}.$$

   As there was little evidence to support the edge parameter in the earlier analysis, the ratio of posterior to prior at $z_{\mathrm{crit}} = 0$ is typically ~ 10, and the new evidence ratios are boosted accordingly.

3. The new analysis replaces the WMAP 7-year KQ75 mask used in all prior analyses with the KQ85 mask, revealing $\sim 8\%$ more of the sky. The change in the fraction of the sky available to the algorithm increases the prior volume on observable source positions by $\sim 8\%$ as well, and the log evidence ratios hence decrease by a similar amount.

4. The candidate detection method has changed from needlets to optimal filters, and the ranges of size and position deemed significant therefore also change. The integration limits for the patch evidences can differ by small amounts if the new ranges either reveal or truncate regions of non-zero likelihood.

In addition, since the first bubble collision analysis we have used an increased number of `MultiNest` live points and tighter tolerance and efficiency settings. The current settings were chosen to ensure accurate calculation of the evidence; however, tests show that there is no significant difference in our results due to the new settings.

Table IX shows the expected and observed changes in evidence ratio obtained for the four best bubble collision and texture candidates processed in the previous analyses. In the majority of cases the differences between the observed and expected changes in evidence ratios are consistent to `MultiNest` precision, but there are two bubble collision cases where the new results show a significant difference.

The first case is the Cold Spot, candidate 3. The memory restrictions of the previous analysis [9, 10] required that the θ_{crit} range sampled be truncated (compare Tables VII in Ref. [9] and VIII in the current analysis). The new adaptive-resolution algorithm allows the full range of θ_{crit} estimated by the candidate detection stage to be sampled, revealing an additional peak in the posterior and boosting the evidence accordingly. The second case is candidate 4. The patches of data used to calculate the evidence for this candidate contain a $\sim 4° \times 3°$ region masked by the KQ75 mask but not by the KQ85 mask. The improvement in evidence ratio most likely derives from uncovering these extra pixels, which produce an extra hot contribution to an already hot feature.

The only differences between the current texture analysis and that of Ref. [11] are the candidate detection algorithm and the mask employed. The candidates that contributed most significantly to the posterior in the previous analysis are all still detected by the optimal filters, and the changes in evidence observed are consistent within `MultiNest` precision, given the new mask and small differences in integration limits from changing the candidate detection stage. Indeed, the posteriors produced by the two analyses (Fig. 2 of Ref. [11] and the right-hand plot of Fig. 18) are almost indistinguishable by eye.

XI. CONCLUSIONS

We have presented a hierarchical Bayesian algorithm for the detection of spatially-localized sources in high-resolution CMB datasets. The algorithm uses the posterior over the global parameters describing the population of sources to determine whether their presence is warranted by the data and prior theoretical knowledge. To cope with the volume of data available, a conservative approximation to the posterior is calculated by selecting the most likely candidate sources using optimal filters and assuming that the remaining data do not contribute to the likelihood. Candidates are processed at the highest data resolution possible, given their size and the computing power available. The effects of the approximations and adaptive-resolution analysis have been quantified using a suite of tests, and are found to be comparable to the typical variance in sampling from the un-approximated posterior.

As a demonstration, the pipeline has been applied to search for evidence of bubble collisions and cosmic textures in the WMAP 7-year data. This work removes the size restriction imposed by memory constraints on the previous bubble collision analysis [9, 10], as well as optimising the candidate detection stages of previous bubble collision and texture analyses [9–11]. The WMAP data do not favor the addition of either bubble collisions or cosmic textures to the ΛCDM model: even though such sources provide higher-likelihood fits, they are not sufficiently predictive to overcome the extra model complexity. In the context of these models our results also place limits on the average numbers of bubble collisions and textures within our detection thresholds on the CMB sky, which are constrained to be fewer than 4.0 and 5.2 at 95% confidence, respectively. The WMAP data do not place any significant constraint on the dimensionless scale of symmetry breaking for textures, ϵ.

The Planck satellite [52] will soon release temperature data with a factor of 2-3 improvement in resolution and ~ 10 in pixel noise over WMAP. Further, it will extract essentially all of the information from the temperature power spectrum, providing a near-ideal characterization of the dominant source of noise in the analysis. These facts strongly motivate performing the texture and bubble collision analyses on the Planck data when they become available. In addition, high-quality CMB polarization data are being gathered by experiments such as Planck, ACTPol [53], SPTPol [54] and Spider [55]. Textures do not induce a polarization signal [22], but bubble collisions are expected to create characteristic imprints [28], complementary to those in the temperature data. Extension of the hierarchical Bayesian analysis pipeline to process polarization data, either in isolation or by cross-correlating with temperature maps, therefore represents a promising avenue for future tests of these models. Further, such an upgraded pipeline could be readily applied to other localized signatures in the CMB, such as the Sunyaev-Zel'dovich effect produced by clusters of galaxies.

---

[1] N. Turok, Phys. Rev. Lett. **63**, 2625 (1989).
[2] N. Turok and D. N. Spergel, Phys. Rev. Lett. **66**, 3093 (1991).
[3] D. N. Spergel, N. Turok, W. H. Press, and B. S. Ryden, Phys. Rev. **D43**, 1038 (1991).
[4] U.-L. Pen, D. N. Spergel, and N. Turok, Phys. Rev. **D49**, 692 (1994).
[5] N. Turok and D. Spergel, Phys. Rev. Lett. **64**, 2736 (1990).
[6] A. Aguirre, in *Beyond the Big Bang* (Springer, 2008).
[7] A. H. Guth, J.Phys. **A40**, 6811 (2007), arXiv:hep-th/0702178 [HEP-TH].
[8] A. Aguirre, M. C. Johnson, and A. Shomer, Phys. Rev. **D76**, 063509 (2007).
[9] S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris, Phys. Rev. **D84**, 043507 (2011), arXiv:1012.3667 [astro-ph.CO].
[10] S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris, Phys. Rev. Lett. **107**, 071301 (2011), arXiv:1012.1995 [astro-ph.CO].
[11] S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris, Phys.Rev.Lett. **108**, 241301 (2012), arXiv:1203.1928 [astro-ph.CO].
[12] J. D. McEwen, S. M. Feeney, M. C. Johnson, and H. V. Peiris, Phys. Rev. **D85**, 103502 (2012), arXiv:1202.2861 [astro-ph.CO].
[13] C. L. Bennett *et al.* (WMAP), Astrophys. J. **583**, 1 (2003), arXiv:astro-ph/0301158.
[14] R. A. Sunyaev and Y. B. Zeldovich, Comments on Astrophysics and Space Physics **4**, 173 (1972).
[15] T. J. Loredo, ArXiv e-prints (2012), arXiv:1208.3036 [astro-ph.IM].
[16] B. Freivogel, Class.Quant.Grav. **28**, 204007 (2011), arXiv:1105.0244 [hep-th].
[17] M. P. Salem, JCAP **1201**, 021 (2012), arXiv:1108.0040 [hep-th].
[18] Planck Collaboration, P. A. R. Ade, *et al.*, Astron. Astrophys. **536**, A1 (2011), arXiv:1101.2022 [astro-ph.IM].
[19] M. P. Hobson and C. McLachlan, Mon. Not. R. Astron. Soc. **338**, 765 (2003), arXiv:astro-ph/0204457.
[20] M. P. Hobson, G. Rocha, and R. S. Savage, "Bayesian source extraction," in *Bayesian Methods in Cosmology*, edited by Hobson, M. P., Jaffe, A. H., Liddle, A. R., Mukeherjee, P., & Parkinson, D. (2010) pp. 167–+.
[21] M. Cruz, E. Martínez-Gonzalez, P. Vielva, J. M. Diego, M. Hobson, and N. Turok, Mon. Not. R. Astron. Soc. **390**, 913 (2008), arXiv:0804.2904.
[22] P. Vielva, E. Martinez-Gonzalez, M. Cruz, R. B. Barreiro, and M. Tucci, Mon. Not. R. Astron. Soc. **410**, 33 (2011), arXiv:1002.4029 [astro-ph.CO].
[23] M. Cruz, E. Martinez-Gonzalez, P. Vielva, and L. Cayon, Mon. Not. R. Astron. Soc. **356**, 29 (2005), arXiv:astro-ph/0405341.
[24] M. Cruz, N. Turok, P. Vielva, E. Martinez-Gonzalez, and M. Hobson, Science **318**, 1612 (2007), arXiv:0710.5737 [astro-ph].
[25] M. Cruz, E. Martinez-Gonzalez, and P. Vielva, in *Highlights of Spanish Astrophysics V* (2010) arXiv:0901.1986 [astro-ph].
[26] A. Aguirre and M. C. Johnson, Rept.Prog.Phys. **74**, 074901 (2011), arXiv:0908.4105 [hep-th].
[27] S. Chang, M. Kleban, and T. S. Levi, JCAP **0904**, 025 (2009), arXiv:0810.5128 [hep-th].
[28] B. Czech, M. Kleban, K. Larjo, T. S. Levi, and K. Sigurdson, JCAP **1012**, 023 (2010), arXiv:1006.0832 [astro-ph.CO].
[29] R. Gobbetti and M. Kleban, JCAP **1205**, 025 (2012), arXiv:1201.6380 [hep-th].
[30] M. Kleban, T. S. Levi, and K. Sigurdson, (2011), arXiv:1109.3473 [astro-ph.CO].
[31] J. Kozaczuk and A. Aguirre, (2012), arXiv:1206.5038 [hep-th].
[32] A. Aguirre and M. C. Johnson, Phys. Rev. **D77**, 123536 (2008).
[33] B. Freivogel, M. Kleban, A. Nicolis, and K. Sigurdson, JCAP **0908**, 036 (2009).
[34] F. Feroz and M. P. Hobson, Mon. Not. R. Astron. Soc. **384**, 449 (2008), arXiv:0704.3704 [astro-ph].
[35] F. Feroz, M. P. Hobson, and M. Bridges, Mon. Not. R. Astron. Soc. **398**, 1601 (2009), arXiv:0809.3437.
[36] D. Larson, J. Dunkley, G. Hinshaw, E. Komatsu, M. R. Nolta, *et al.*, Astrophys. J. Suppl. Ser. **192**, 16 (2011), arXiv:1001.4635 [astro-ph.CO].
[37] J. D. McEwen, M. P. Hobson, and A. N. Lasenby, IEEE Trans. Sig. Proc. **56**, 3813 (2008), astro-ph/0612688.
[38] D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkyacharian, P. Natoli, D. Picard, and N. Vittorio, Mon. Not. R. Astron. Soc. **383**, 539 (2008).
[39] S. Scodeller, O. Rudjord, F. Hansen, D. Marinucci, D. Geller, *et al.*, Astrophys. J. **733**, 121 (2011), arXiv:1004.5576 [astro-ph.CO].
[40] J. D. McEwen, M. P. Hobson, D. J. Mortlock, and A. N. Lasenby, IEEE Trans. Sig. Proc. **55**, 520 (2007), astro-ph/0506308.
[41] K. M. Gorski *et al.*, Astrophys. J. **622**, 759 (2005), arXiv:astro-ph/0409513.
[42] B. Gold, N. Odegard, J. L. Weiland, R. S. Hill, A. Kogut, *et al.*, Astrophys. J. Suppl. Ser. **192**, 15 (2011), arXiv:1001.4555 [astro-ph.GA].
[43] S. M. Feeney, H. V. Peiris, and A. Pontzen, Phys. Rev. **D84**, 103002 (2011), arXiv:1107.5466 [astro-ph.CO].
[44] A. de Oliveira-Costa and M. Tegmark, Phys. Rev. **D74**, 023005 (2006), arXiv:astro-ph/0603369.
[45] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics,

Philadelphia, PA, 1999).

[46] G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, S. S. Meyer, C. Barnes, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright, Astrophys. J. Suppl. Ser. **148**, 135 (2003), arXiv:astro-ph/0302217.

[47] J. Skilling, in *AIP Conference Proceedings of the 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, American Institute of Physics Conference Series, Vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint (2004) pp. 395–405.

[48] C. Bennett *et al.* (WMAP Collaboration), Astrophys.J.Suppl. **148**, 97 (2003), arXiv:astro-ph/0302208 [astro-ph].

[49] E. Wright *et al.* (WMAP Collaboration), Astrophys.J.Suppl. **180**, 283 (2009), arXiv:0803.0577 [astro-ph].

[50] A. Pontzen and H. V. Peiris, Phys. Rev. **D81**, 103008 (2010), arXiv:1004.2706 [astro-ph.CO].

[51] J. Dickey, Ann. Mathemat. Statist. **42**, 204 (1971).

[52] J. A. Tauber, N. Mandolesi, J. Puget, *et al.*, Astron. Astrophys. **520** (2010), 10.1051/0004-6361/200912983.

[53] M. D. Niemack, P. A. R. Ade, J. Aguirre, F. Barrientos, J. A. Beall, J. R. Bond, J. Britton, H. M. Cho, S. Das, M. J. Devlin, S. Dicker, J. Dunkley, R. Dünner, J. W. Fowler, A. Hajian, M. Halpern, M. Hasselfield, G. C. Hilton, M. Hilton, J. Hubmayr, J. P. Hughes, L. Infante, K. D. Irwin, N. Jarosik, J. Klein, A. Kosowsky, T. A. Marriage, J. McMahon, F. Menanteau, K. Moodley, J. P. Nibarger, M. R. Nolta, L. A. Page, B. Partridge, E. D. Reese, J. Sievers, D. N. Spergel, S. T. Staggs, R. Thornton, C. Tucker, E. Wollack, and K. W. Yoon, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7741 (2010) arXiv:1006.5049 [astro-ph.IM].

[54] J. J. McMahon, K. A. Aird, B. A. Benson, L. E. Bleem, J. Britton, J. E. Carlstrom, C. L. Chang, H. S. Cho, T. de Haan, T. M. Crawford, A. T. Crites, A. Datesman, M. A. Dobbs, W. Everett, N. W. Halverson, G. P. Holder, W. L. Holzapfel, D. Hrubes, K. D. Irwin, M. Joy, R. Keisler, T. M. Lanting, A. T. Lee, E. M. Leitch, A. Loehr, M. Lueker, J. Mehl, S. S. Meyer, J. J. Mohr, T. E. Montroy, M. D. Niemack, C. C. Ngeow, V. Novosad, S. Padin, T. Plagge, C. Pryke, C. Reichardt, J. E. Ruhl, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, B. Stadler, A. A. Stark, Z. Staniszewski, K. Vanderlinde, J. D. Vieira, G. Wang, R. Williamson, V. Yefremenko, K. W. Yoon, O. Zhan, and A. Zenteno, in *American Institute of Physics Conference Series*, American Institute of Physics Conference Series, Vol. 1185, edited by B. Young, B. Cabrera, and A. Miller (2009) pp. 511–514.

[55] B. P. Crill, P. A. R. Ade, E. S. Battistelli, S. Benton, R. Bihary, J. J. Bock, J. R. Bond, J. Brevik, S. Bryan, C. R. Contaldi, O. Doré, M. Farhang, L. Fissel, S. R. Golwala, M. Halpern, G. Hilton, W. Holmes, V. V. Hristov, K. Irwin, W. C. Jones, C. L. Kuo, A. E. Lange, C. Lawrie, C. J. MacTavish, T. G. Martin, P. Mason, T. E. Montroy, C. B. Netterfield, E. Pascale, D. Riley, J. E. Ruhl, M. C. Runyan, A. Trangsrud, C. Tucker, A. Turner, M. Viero, and D. Wiebe, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7010 (2008) arXiv:0807.1548.