Novel Algorithms for Early-Universe Cosmology

Stephen Martin Feeney

Submitted for the degree of Doctor of Philosophy

Department of Physics and Astronomy

University College London

September 2012

I, *Stephen Martin Feeney*, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Fluctuations in the cosmic microwave background (CMB), the radiation left over from the Big Bang, contain information which has been pivotal in establishing the current cosmological model. CMB data can also be used to test theoretically well-motivated additions to the model, including pre-inflationary relics (signatures of bubble collisions arising in eternal inflation) and topological defects that form after inflation (cosmic strings and textures). These relics typically leave sub-dominant, spatially localised signals, hidden in the "noise" of the primary CMB, the instrumental noise, foreground residuals and other systematics.

Standard approaches for searching for such signals involve focusing on statistical anomalies, which carry the danger of extreme *a posteriori* biases. The self-consistent approach to this problem is Bayesian model comparison; however, the full implementation of this approach is computationally intractable with current CMB datasets, and will only become more difficult with data from the next generation of CMB experiments. I will describe a powerful modular algorithm, capable of coping with the volume of data, which combines a candidate-detection stage (using wavelets or optimal filters) with a full Bayesian parameter-estimation and model-selection stage performed in pixel space within the candidate regions. The algorithm is designed to fully account for the "look-elsewhere" effect, and its use of blind analysis techniques further enhances its robustness to unknown systematics. Finally, I will present the results of applying the algorithm to hunt for the signatures of bubble collisions and cosmic textures in the seven-year data from the Wilkinson Microwave Anisotropy Probe.

Acknowledgements

I owe an enormous debt of gratitude to random chance, which paired me up with the best supervisor, Hiranya Peiris, I could ever hope for. For sharing the way you think and work, the opportunities you have granted me, the people I have met, the incredible science to which I have been exposed, for the inspiration, the help and the occasional shove, I will never be able to thank you enough. Thank you also to my amazing collaborators, Matt Johnson, Daniel Mortlock, Andrew Pontzen and Jason McEwen (in order of appearance), without whom this thesis would be so much blank paper: it has been an honour and a privilege to share the pages of the Physical Review with you. Thank you to Ofer Lahav and Filipe Abdalla, my official and unofficial second supervisors, for all of their help and, in Filipe's case, "stimulating" discussions. Thanks also to John Deacon and Dugan Witherick for all of the cores and support you have granted me. To Ingo Waldmann and all of the other UCL PhD students and post-docs, thank you for your support, hilarity and general insanity. My deepest gratitude and all of my love to my mum, dad and sister for launching me onto the path I now follow, and to the rest of my family, both old and new, and friends for filling my life with happiness. And finally, thank you to my beautiful wife Jenny: my best friend, my strength and my inspiration, whom I love more than words can describe.

Dedicated to Dr Jennifer Anne Feeney,
who got there first.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Universe

Cosmology – the study of the large-scale structure and evolution of the Universe – is a topic that has fascinated mankind for millennia. Initially bounded only by imagination, through confrontation with carefully taken data cosmologists have transformed philosophy into science. In the last twenty years, the quality of cosmological datasets has improved dramatically, painting a firm but broad-brush picture of the last 14 billion or so years of the Universe's history. There is, however, a lot to learn about almost every single brush-stroke. With the quality of modern datasets we can map out the Universe's history in unprecedented detail.

1.1.1 Composition

The modern concordance cosmological model, ΛCDM, describes a homogeneous and isotropic universe that is spatially flat. The contents of the Universe are listed below.

1. Baryonic matter (often simply "baryons"): matter comprising mostly baryons by mass. Baryonic matter includes all nuclei and electrons (even though electrons are leptons), and is predominantly primordial hydrogen formed early in the Universe's history.

2. Currently small, but historically extremely important, relic components of electromagnetic radiation (discussed in detail in Section 1.2) and neutrinos, known as the cosmic microwave and neutrino backgrounds (CMB and CνB, respectively).

3. Cold dark matter (CDM): pressureless, non-relativistic and *non-baryonic* matter that does not interact via the electromagnetic force. Though not yet observed directly, its existence

is required to explain a host of phenomena on Galactic (e.g. Oort (1932); Rubin and Ford (1970)) to cosmological (e.g. Zwicky (1933); Davis et al. (1982)) scales. The strongest evidence for the non-baryonic nature of dark matter comes from measurements of the CMB, which reveal that the amount of matter required to form the structures observed today is far greater than the amount of baryons known to exist (Komatsu et al., 2011). A full discussion of the effects of dark matter on the CMB is provided in Sec. 1.2.3. The precise nature of dark matter is the subject of intense ongoing investigation (see e.g. Angle et al. (2008); Bernabei et al. (2008); Aalseth et al. (2011)).

4. A dark energy component (Λ), currently indistinguishable from a cosmological constant, responsible for the recent acceleration of the Universe's expansion (Riess et al., 1998; Perlmutter et al., 1999). If dark energy is indeed a cosmological constant, its value is lower than expected by some 120 orders of magnitude, a discrepancy which could indicate that the observed Universe is but one of many (Weinberg, 1989).

The dark energy is currently the dominant component, making up $\sim 73\%$ of the Universe's energy density. Dark matter accounts for a further $\sim 23\%$, with normal baryonic matter making up only $\sim 5\%$ (Komatsu et al., 2011). While the model explains the abundance of light elements (Coc et al., 2004), CMB (Komatsu et al., 2011), galaxy power spectra (Reid et al., 2010), supernovae redshift-distance relationship (Suzuki et al., 2012) and more,[1] it presents almost as many questions as it answers, not least the provenance of the two dominant components for which it is named. Although the currently accepted composition is fairly well-defined, there are strong indications from both theory (e.g. Kibble (1976)) and experiment (e.g. Fukuda et al. (1998)) that further additions should be made. This thesis is primarily concerned with additions to the ΛCDM model.

1.1.2 Expansion

The Universe is expanding, an observation first made when Hubble noted a linear relationship between the distance and recessional velocity of "extra-galactic nebulae", i.e. local galaxies (Hubble, 1929). Assuming the cosmological principle – that the Universe has no preferred position or direction and is therefore homogeneous and isotropic – the Universe is described by the Friedmann-Robertson-Walker (FRW) metric (Friedman, 1922)

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t) \left[\frac{\mathrm{d}r^2}{1 - kr^2} + r^2(\mathrm{d}\theta^2 + \sin^2\theta \mathrm{d}\phi^2) \right], \tag{1.1}$$

---

[1]Neither the list of physical phenomena nor references provided here are exhaustive, but are rather meant to provide illustrative recent examples of the experimental success of the model.

where $\{r, \theta, \phi\}$ are spherical polar coordinates, $a(t)$ is a time-dependent scale factor, k is a measure of the curvature of the space-time and I have used natural units ($c = 1$).

Assuming General Relativity is the correct theory of gravity, the evolution of the scale factor in a universe containing a perfect fluid with pressure p and density ρ is governed by the Friedmann equations

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G \rho}{3} - \frac{k}{a^2} + \frac{\Lambda}{3} \tag{1.2}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(3p + \rho) + \frac{\Lambda}{3}, \tag{1.3}$$

where G is Newton's constant, Λ is the cosmological constant, and dots denote derivatives with respect to time. Solving the Friedmann equations requires an equation of state for the perfect fluid (or, equivalently, a relationship between the energy density and the scale factor). For cold dark matter, which is pressureless, this is simply $p_{\mathrm{mat}} = 0$; photons and other relativistic particles obey $p_{\mathrm{rad}} = \frac{1}{3}\rho_{\mathrm{rad}}$. These imply, and are implied by, the fact that the densities of cold dark matter and photons scale as $\rho_{\mathrm{mat}} \propto a^{-3}$ and $\rho_{\mathrm{rad}} \propto a^{-4}$, respectively.

The critical density, ρ_{crit}, is the total energy density of a flat universe, where $k = 0$. Defining the energy density in the cosmological constant as

$$\rho_\Lambda = \frac{\Lambda}{8\pi G} \tag{1.4}$$

allows the absorption of the cosmological constant term in Equation 1.2 into the energy density. The critical density is then

$$\rho_{\mathrm{crit}} = \frac{3H^2}{8\pi G}, \tag{1.5}$$

where $H = \frac{\dot{a}}{a}$ is the Hubble parameter.

The Hubble parameter is a measure of the Universe's expansion rate, and defines typical length- and time-scales for the Universe at a given epoch. Recalling that $c = 1$, the *comoving*[2] Hubble radius is given by $(aH)^{-1}$, and is the comoving distance at which the recession velocity is the speed of light. The Hubble radius can be thought of as defining the maximum distance separating particles in causal contact at a given epoch. Another important length scale is the comoving particle horizon

$$\eta = \int_0^t \frac{\mathrm{d}t'}{a(t')}, \tag{1.6}$$

the maximum comoving distance traveled by a photon since the beginning of the Universe. This

---

[2] A comoving distance is a physical *proper* distance divided by the scale factor to remove the expansion of the Universe. The comoving distance between two objects moving only due to the Universe's expansion – i.e. with zero peculiar velocity – is therefore constant.

is the maximum distance separating particles that have *ever* been in causal contact. In universes containing only radiation and matter, the Hubble and horizon scales are approximately equal (indeed, exactly equal for a universe containing only radiation); this is not the case for universes containing only dark energy. An excellent discussion of these, and other, distance scales can be found in Davis and Lineweaver (2004).

1.1.3 History

The history of a ΛCDM universe begins with the universe in an extremely hot and dense state. The Universe is initially dominated by a scalar field known as the inflaton, an exotic form of matter which exerts negative pressure. Driven by the inflaton, the Universe undergoes a period of exponential expansion known as inflation (Guth, 1981) (discussed in more detail in Section 1.4), increasing the scale factor by at least e^{60} and leaving the currently-observable portion of the Universe flat, homogeneous and isotropic.

At the end of inflation, the inflaton is assumed to have decayed into the particles of the Standard Model in a process known as reheating (see Allahverdi et al. (2010) for a review). Quantum fluctuations in the inflaton are transferred to the decay products as a distinctive spectrum of perturbations on a smooth background. The electromagnetic radiation produced by reheating dominates the energy budget of the Universe at this stage, decelerating the expansion such that the scale factor grows only as $a \propto t^{1/2}$, and allowing only logarithmic growth of sub-horizon matter perturbations. The Universe is filled with a plasma of photons and exotic particles, which gradually become less exotic as they cool and combine into nucleons. Between around two and twenty minutes after the onset of inflation, when the temperature of the plasma drops below $\sim 10^9$K, the protons and neutrons combine to produce deuterium, helium and lithium nuclei in characteristic ratios.

As the energy density of radiation decreases faster than that of matter, the radiation becomes steadily less influential and the dark matter comes to dominate the the Universe. The expansion continues to slow, albeit at a reduced deceleration ($a \propto t^{2/3}$), but critically matter perturbations now grow linearly with the scale factor. Once the temperature reaches $\sim$ 3000 K, some 380,000 years after inflation, it becomes energetically favourable for free electrons to combine with protons in an event known as recombination, and the Universe rapidly becomes neutral. At this point the mean-free path of the photons grows larger than the Hubble scale, and the photons free-stream, forming the CMB: our earliest picture of the Universe. A more in-depth description of the CMB is provided in Section 1.2.

The epoch of matter domination sees the formation of recognisable structures, from the first

stars or quasars responsible for reionisation of the intergalactic medium – the precise identity of the sources responsible being a particularly active field of current research – to galaxies and the filamentary structures of galaxy clusters and voids identified in large-scale surveys. After approximately 10 billion years, the dark energy overtakes matter as the dominant component in the Universe, and the Universe's expansion starts to accelerate. We find ourselves observing the Universe around 14 billion years after inflation, at a particularly unique – and, indeed, poignant – epoch, at the point where the complexity created by the collapse of ever-larger structures starts to unravel.

1.2 The Cosmic Microwave Background

1.2.1 Discovery and characterisation

As discussed in Section 1.1.3, the cosmic microwave background (CMB) is the radiation released after free electrons and protons combined when the Universe cooled below $\sim$ 3000 K. Now gravitationally redshifted into the microwave regime, it was first measured by Penzias and Wilson (1965) and identified by Dicke et al. (1965). The COBE satellite measured the spectrum of the CMB to be a near-perfect black-body at 2.725 K (Mather et al., 1994), implying that the entire CMB was once in thermal equilibrium (more on this later), and discovered fluctuations in its temperature of order 1 in 100,000 (Smoot et al., 1992). These anisotropies are (primarily) primordial, and therefore offer an early glimpse of the perturbations that went on to form structures we see today. The characteristics of the anisotropies depend intricately on the contents of the Universe and the source of the primordial perturbations, and yield an extraordinary amount of cosmological information. The anisotropies have accordingly been studied in great detail by experiments such as Saskatoon (Netterfield et al., 1997), BOOMERanG (de Bernardis et al., 2000) and MAXIMA (Hanany et al., 2000), and latterly WMAP (Larson et al., 2011), ACT (Das et al., 2011), SPT (Keisler et al., 2011) and Planck (Planck Collaboration et al., 2011).[3] Their power spectrum, plotted in Figure 1.1, forms the cornerstone of the ΛCDM model, and, indeed, the modern view of cosmology as a precision science; a discussion of its major features is presented in the following section.

---

[3]Among many others! See `http://lambda.gsfc.nasa.gov/links/experimental_sites.cfm` for a complete listing.
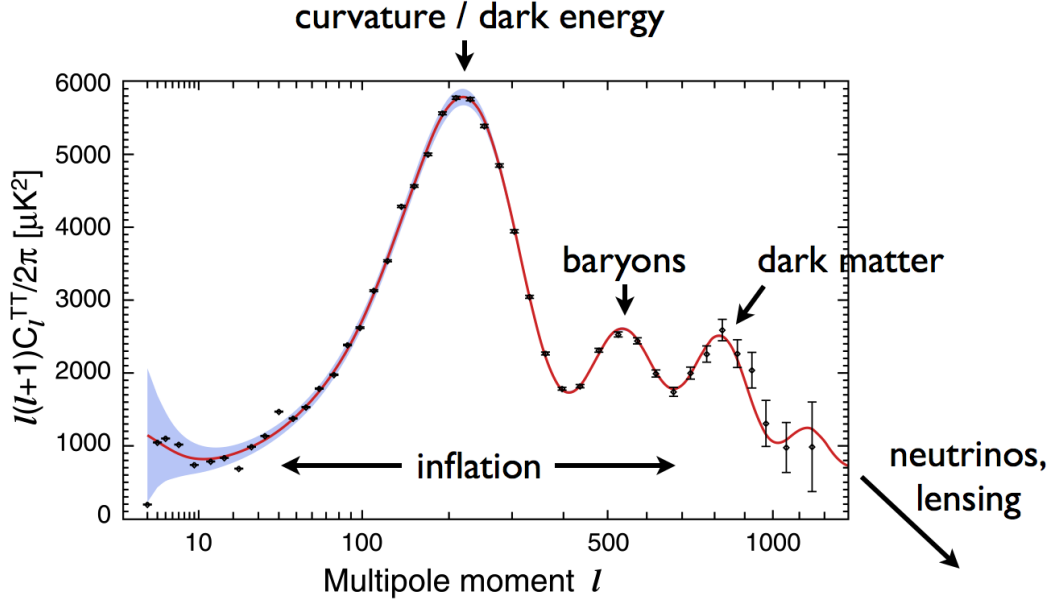
Figure 1.1: The CMB power spectrum as measured by seven years of WMAP observations. The best-fit ΛCDM power spectrum is plotted as a red line, and cosmic variance – the sample variance inherent in estimating the power spectrum from a limited number of modes – is indicated by the shaded region. Original (annotation-free) plot credit: Larson et al. (2011).

1.2.2 Anisotropies

Inflation (of which much more will be said in Section 1.4) creates a universe containing a nearly scale-invariant set of adiabatic perturbations of radiation, baryons and dark matter.[4] At early times, the photons are so energetic that the baryons are entirely ionised, and the photons and baryons are tightly coupled by Compton and Thomson scattering processes. The dark matter and radiation interact more indirectly, via the gravitational potentials set up by the primordial density fluctuations, which are dominated at early times by the radiation, and at later times by the dark matter. The evolution of the perturbations in each of the Universe's components are therefore coupled until the temperature is low enough for hydrogen atoms to form and the CMB is released.

The evolution of the perturbations depends greatly on scale. The largest scales measured in the CMB power spectrum were super-horizon until just before the CMB decoupled, and therefore have had very little time in causal contact to evolve. These modes are therefore largely unchanged from the primordial fluctuations, although late-time effects from dark energy also affect this region of the power spectrum (Rees and Sciama, 1968).

Smaller scale modes, on the other hand, are sub-horizon for long enough to evolve through causal processes. The evolution is governed by two competing effects: the radiation pressure

---

[4]Scale-invariance here means that the full three-dimensional power spectrum of the perturbations is given by $P(k) \propto k^{n_s}$, where k is the wavenumber and $n_s = 1$ for exact scale-invariance. Adiabaticity means that all species were created with the same number density at reheating.

of the photons tending to equilibrate variations in the density field, and gravity causing the dark matter and baryons to cluster in potential wells. The two competing forces set up acoustic oscillations in the photon-baryon plasma, which in turn produce features in the power spectrum. Modes that have completed a half-integer number of oscillations at recombination are maximally rarefied or compressed when the CMB photons decouple. Structure on these scales – the harmonics of the horizon size at recombination – is therefore enhanced relative to intermediate scales, creating acoustic peaks in the power spectrum (see Figure 1.1). The finite speed of sound in the photon-baryon plasma means that the photons and baryons are not perfectly coupled, and the photons are therefore able to travel a small distance between collisions. This serves to smooth out structures smaller than a characteristic diffusion length-scale (Silk, 1968), and as a result the smallest-scale modes in the CMB power spectrum are heavily damped.

1.2.3 Applications

The intricate dependence on the CMB anisotropies on the contents of the Universe makes the CMB an ideal dataset for characterising the cosmological model. The primary application of CMB data is therefore the use of the power spectrum to constrain models of the Universe. The features of the power spectrum constraining particular facets of the cosmological model are indicated in Figure 1.1. The overall dependence of the power spectrum on scale indicates the shape of the primordial power spectrum, and strongly favours the nearly scale-invariant spectrum produced by inflation. The positions of the acoustic oscillations constrain a combination of the curvature and the amount of dark energy, which, when combined with an independent measure of the Hubble constant, indicates that the Universe is flat. A measurement of the first two acoustic peaks allows the baryonic fraction of the Universe to be determined, as an increased baryon fraction causes the photon-baryon plasma to compress further than it rarefies, weighting odd peaks over even. Observations of smaller-scale peaks allow the dark matter fraction to be constrained: modes which begin to oscillate when the Universe is still radiation-dominated, and which are therefore small, do so in decaying potentials, driving up their amplitudes. Finally, measurements of the smallest scales – the so-called damping tail – constrain the amount of relativistic matter (in particular, the neutrino fraction) present at recombination.

As the physics of the primary CMB anisotropies is so well-understood, the CMB is also an excellent dataset in which to search for small deviations from Gaussianity imparted both primordially and in the intervening ~ 14 billion years. These effects include the primordial non-Gaussianities inherent in all but the simplest models of inflation, the Integrated Sachs-Wolfe effect due to late-time dark energy domination, gravitational lensing of the CMB by clumps of

dark matter, and, potentially, signatures of as-yet undiscovered physical processes occurring at very early times (as discussed in Sections 1.4.3 and 1.5). As these effects induce non-Gaussian features into the CMB signal, the power spectrum no longer completely characterises the CMB data. Such studies therefore consider either higher-order correlations of the data, such as the bispectrum or trispectrum, or maps of the CMB sky, either in isolation or in cross-correlation with other datasets. Map-based approaches are the main focus of the thesis, and will be discussed in great detail in Chapters 2-5.

1.3 Matter Perturbations

As the CMB power spectrum describes the perturbations to the primordial photons, so the matter power spectrum describes the perturbations to the cold dark matter and baryons which evolve into galaxies, clusters and filamentary large-scale structure. The form of the matter power spectrum is derived from linear perturbation theory, and shows that the behaviour of matter perturbations is heavily scale-dependent. The theoretical ΛCDM matter power spectrum is plotted in Figure 1.2 as a function of wavenumber k, and its form is described below.

Let us first consider super-horizon modes: that is, matter perturbations with wavelengths larger than the horizon at a given epoch. These modes can not evolve through causal processes. Linear perturbation theory tells us that their amplitudes grow linearly with conformal time[5] during the radiation- and matter-dominated epochs, but remain constant once dark energy begins to dominate. This picture changes, however, when the modes become smaller than (or "enter") the horizon. The smallest-scale modes, which enter the horizon before matter-radiation equality, evolve in a Universe dominated by radiation. As discussed in previous sections, radiation pressure causes the photon plasma to oscillate, the associated potential wells to decay, and therefore the growth of matter perturbations to slow to logarithmic with conformal time. Modes which enter the horizon *after* matter-radiation equality[6], however, behave very differently: as the *pressureless* cold dark matter is now dominant, there is no pressure support, and the perturbations are free to grow linearly with conformal time. This means that the power in modes which enter the horizon before matter-radiation equality is suppressed (by a factor of $\sim k^{-4}$) relative to those which enter the horizon afterwards. When this is evolution is superimposed upon a scale-invariant (i.e. $P(k) \propto k$) initial power spectrum, a characteristic turnover is seen at the wavenumber corresponding to the horizon scale at matter-radiation equality.

---

[5]Conformal time is defined to be $\eta = \int_0^t \frac{dt'}{a(t')}$, i.e. exactly the same as the particle horizon: the maximum distance traveled by a photon corresponds to a particularly useful time variable.

[6]But before dark energy begins to dominate: the rapid expansion associated with dark energy domination stops the growth of matter perturbations on all scales.

Figure 1.2: The best-fit theoretical ΛCDM matter power spectrum (red line), and measurements of it from a variety of sources, including galaxy redshifts from the Sloan Digital Sky Survey (SDSS) (Tegmark et al., 2004). This plot is reproduced from Tegmark et al. (2004), which contains a full description of the other datasets featured.

Though it is, of course, impossible to probe dark matter using electromagnetic radiation, by assuming that galaxies are biased tracers of the underlying dark matter distribution one can derive the full matter power spectrum using galaxy surveys. The results of such a survey, namely the Sloan Digital Sky Survey (SDSS), are superimposed upon the theoretical power spectrum shown in Figure 1.2. Such observations are critical to our understanding of the Universe. The matter power spectrum is an extremely rich source of cosmological information, and can be used to constrain the Universe's composition, initial conditions (i.e. the primordial power spectrum), and the nature of its more-exotic constituents (dark matter, dark energy and neutrinos, for example) through their effects on the growth of structure.

1.4 Inflation

1.4.1 Motivation

The phase of inflation invoked in the early Universe is motivated by a number of "fine-tuning" problems with the observed CMB, namely the horizon, flatness, and multipole problems. From the very first observations of the CMB (Penzias and Wilson, 1965; Dicke et al., 1965), it was clear that the CMB is uniform in temperature across the sky, implying that the entire observable Universe was in thermal equilibrium at last scattering. This is at odds with expectations from simple calculations assuming that the Universe contains only matter and radiation, which show that the size of the particle horizon – the maximum separation between two points in causal contact at the start of the Universe – is $\sim 1°$ across on the CMB sky. This is the *horizon problem*: any two points in the CMB separated by more than a degree should have $\mathcal{O}(1)$ variations in temperature, but clearly do not.

The source of the second, *flatness* problem, is the observation that the current Universe is so close to flat, or equivalently that the density today is so close to the critical density. The critical density is an unstable point, so if there is any small deviation from the critical density at the Universe's inception it grows with time. Looking back, this means that the Universe must have been even flatter in the past, so much so as to imply fine-tuning. The observed near-flatness of the Universe also uncovers another problem, the *monopole* problem, as phase transitions in the early Universe (of which more will be said in Section 1.5) are expected to have produced relics, known as monopoles, which over-close the Universe.

1.4.2 Implementation

Alan Guth (Guth, 1981) observed that a phase of exponential expansion has the potential to fix each of these problems. Guth proposed that the early Universe contained a scalar field, the inflaton, trapped in a false vacuum, dominated by the associated potential energy. This constant potential energy density produced an exponential expansion rate, greatly increasing the horizon size, flattening the Universe and diluting away any relic particles. Provided the scale factor increased by a factor $\gtrsim e^{60}$, the fine-tuning problems were solved. The Universe then stopped inflating by undergoing a phase transition via bubble nucleation, whereby bubbles of the true vacuum instantaneously appear within the false vacuum. A simple example of a scalar-field potential in which this scenario, termed "old inflation", arises is shown in the left-hand plot of Figure 1.3.

This original proposal suffers from the so-called "graceful exit" problem, as the Universe

Figure 1.3: Examples of inflationary potentials. In the left-hand plot, the inflaton, depicted by a blue ball, is trapped within a false vacuum and must tunnel through a potential energy barrier in order to reach the true vacuum. In the right-hand plot, the inflaton rolls slowly down a plateau before oscillating about the potential minimum.

is not able to stop inflating and reheat homogeneously (Guth and Weinberg, 1983). The only method of thermalizing the Universe is for bubbles to collide, releasing the energy in their walls. For this to occur homogeneously requires the percolation of many such bubble collisions, and therefore a high bubble nucleation rate, but inflation then ends too quickly to solve the horizon and flatness problems. The graceful exit problem can be solved by presuming that the inflaton is not trapped in a false vacuum, but is instead rolling slowly down a shallow potential towards a minimum (Linde, 1982; Albrecht and Steinhardt, 1982). An example potential for this so-called "new inflation" is shown in the right-hand plot of Figure 1.3. The potential energy is still roughly constant, and so the Universe inflates, stopping only when the kinetic energy becomes comparable to the potential energy. Once the inflaton reaches the bottom of the potential, it undergoes oscillations which couple to the particles of the Standard Model, reheating the Universe.

Part of the lasting appeal of inflation is that, in addition to solving the fine-tuning problems highlighted above, it also provides a natural method of generating the primordial perturbations that grow into the structures seen today. As the Universe undergoes such an enormous expansion, quantum fluctuations in the inflaton field are blown up to super-horizon scales (Mukhanov and Chibisov, 1981) with a nearly scale-invariant power spectrum. Over 20 years after inflation was proposed, this prediction was confirmed by the CMB power spectrum obtained by the WMAP satellite (Spergel et al., 2003; Komatsu et al., 2003; Peiris et al., 2003).

1.4.3 Eternal inflation and the Multiverse

Inflation is an empirically successful theory: it predicts Gaussian, nearly scale-invariant fluctuations, and it *can* solve the fine-tuning problems, but this success is purely phenomenological. Theories of higher-energy physics implementing inflation – and, in particular, the form of the potential sourcing inflation – are needed to make probabilistic statements about its likely duration, and hence decide whether this success is deserved. If inflation embedded in a physical framework is found to be extremely unlikely, then instead of solving the fine-tuning problems discussed earlier, it simply moves them to an earlier epoch in the Universe's history.

String theory is arguably the leading candidate for quantum gravity: a "theory of everything", describing the Universe on the smallest and largest scales, and in both the weak and strong gravitational limits. String theory predicts that spacetime consists of at least 10 dimensions, of which all but four are compactified. There are an extremely large number – at least 10^{500} – of ways of doing so, each corresponding to a different false vacuum in a vast landscape potential (Susskind, 2003). In string theory, we therefore expect inflationary potentials resembling those originally proposed by Guth (1981): picture the left-hand side of Figure 1.3 with $\sim 10^{500}$ more vacua of greatly differing magnitudes.[7]

The question therefore arises as to how inflation can occur in the string theory landscape. If the inflaton simply tunnels from one false vacuum to the next, then we are stuck with the same problems that thwarted the original, "old" inflation. Progress can be made provided the requirement for percolation of the bubbles is dropped. If, instead, the inflaton tunnels onto a region of the potential where it can slowly roll, each tunnelling event produces a bubble in which slow-roll inflation, reheating, and standard cosmological evolution can then occur.[8] In this picture of "eternal inflation", the bulk of space inflates forever, stopping only within individual bubble universes. Each bubble nucleates with a finite radius before expanding at the speed of light (Coleman and De Luccia, 1980), and to observers within contains an infinite FRW metric, which is initially open, but can be flat after slow-roll.[9] The collection of these bubbles is known as the Multiverse.

The infinite metric foliating the interior of each bubble means that it is impossible for observers to see out of their bubble into the Multiverse. However, it is possible for information

---

[7]Note that multiple vacua appear generically in high-energy theories, including supersymmetry and supergravity. The following reasoning should therefore not be considered as only relating to string theory.

[8]Guth and Weinberg (1983) originally argued against the Universe residing in a single such bubble, as it could not contain enough information and reheat to the right temperature. These issues are addressed by the secondary phase of slow-roll inflation.

[9]The duration of slow-roll inflation that naturally follows bubble nucleation is the subject of ongoing study. Calculations based on single bubble nucleations from *ad hoc* potentials indicate that sufficient slow-roll to satisfy current bounds on curvature occurs in 90% of cases in which structures form on Galactic scales (Freivogel et al., 2006). Equivalent results for colliding bubbles nucleating from realistic potentials are needed to make robust statements about this probability.

about collisions between bubbles to penetrate the bubble interior (Aguirre et al., 2007). To see this, consider Figure 1.4, which plots the space-time associated with two colliding bubbles. When the bubbles collide, energy passes into the interior of each bubble in the direction of the collision. This energy perturbs the interior of the observation bubble, affecting all observers to the future of the collision. Because of the differences between the metrics describing the inflating bulk and the bubble interior, observers see the collision as taking place in their past, before the phase of slow-roll inflation creating their ΛCDM universe. In particular, the observer indicated by the black dot sees a last-scattering surface that is partly pristine and partly perturbed by the collision: its CMB sky would contain a localised, circular, smooth temperature modulation (Chang et al., 2009). Fascinatingly, it is therefore possible to observationally test the existence of the Multiverse using CMB data. Chapters 3 and 5 describe searches for bubble collisions in the WMAP 7-year data, including deeper discussion of the observational signatures and the probability of detection.

1.5 Topological defects

Topological defects are formed when a medium undergoes a phase transition in which certain symmetries are broken. Defects form because the medium can configure itself in a number of different ways after the symmetry is broken. If two regions are separated by large distances there is no reason to expect they should end up in the same configuration after the phase transition. The interfaces between such regions – long-lived, localised, high-energy states of the medium – are known as topological defects. A terrestrial example of a symmetry-breaking phase transition is the cooling of a ferromagnetic material. Above the Curie temperature, the magnetic field in such a material is randomly ordered at all positions. Below the Curie temperature, the magnetic field is aligned locally, in regions known as domains, but not globally. Topological defects, called "domain walls", are observed where the domains meet.

Topological defects are of cosmological interest because symmetry-breaking phase transitions are predicted to take place as the Universe cools (Kibble, 1976), producing, for example, domains with differing values of the Higgs field. As topological defects are long-lived regions where the medium is trapped in a high-energy state, and potential energy gravitates in General Relativity, we should generically expect topological defects to affect the evolution of the Universe. Some defects, such as monopoles (point defects) and domain walls analogous to those formed in ferromagnets, do so catastrophically, being produced in such numbers as to cause the Universe to collapse before any structure has formed ("over-closing" the Universe (Preskill, 1979)). Others, such as cosmic strings and textures (stable linear defects and dynamical, three-dimensional

Figure 1.4: Space-time diagram of two bubbles colliding. The axes plotted indicate a time and space coordinate in the eternally inflating bulk; for simplicity, the bubbles are taken to nucleate at the same time ($t = 0$). Each bubble nucleates with a finite radius before expanding. The future lightcone of the nucleation (black dashed line) forms the bubble's "Big Bang", where the scale factor tends to zero; slow-roll inflation takes place shortly after this. After nucleation, the bubble walls rapidly accelerate to the speed of light, which corresponds to slopes of 45° in this diagram. When the bubble walls collide, a domain wall (in red) forms between the bubbles, and energy is transferred into the bubble interiors affecting all regions to the future of the collision (blue dashed line). This perturbs the surfaces on which the inflaton field is constant – the natural time coordinates describing the bubble interior – which are shown as thin black lines in the unperturbed region of the bubble, and as green lines in the perturbed region. If the thick black / green line corresponds to the last-scattering surface, the highlighted observer's CMB will contain a disk-like region that has been perturbed by the collision. Original plot credit: Aguirre and Johnson (2009).

"knots" (Turok, 1989)), are more benign, and can be produced with only perturbative effects on the rest of the Universe. While inflation dilutes away any defects formed prior to its onset (Guth, 1981), any defects formed in post-inflationary phase transitions could have observable consequences. As the energy scale of inflation is far, far beyond that achievable in the laboratory, detection of a defect could hint at processes occurring in regimes otherwise completely inaccessible to laboratory experiment.

The CMB provides a particularly useful signal in which to detect defects, as it back-lights all structures between the observer and last-scattering surface. CMB photons interact with defects to leave distinctive deviations from the CMB's primordial Gaussian random field. Cosmic strings, for example, produce linear discontinuities in the temperature field, as photons passing the string on different sides experience a differential redshift (Kaiser and Stebbins, 1984). Textures are dynamical, collapsing and eventually exploding, and so CMB photons passing through

textures experience time-varying potentials. Textures therefore leave a distribution of hot and cold modulations on the CMB sky (Turok and Spergel, 1990). While we can rule out the production (after inflation, at least) of monopoles and domain walls by our presence alone, the search for cosmic textures and strings is very much alive. Chapters 4 and 5 describe searches for textures in the WMAP 7-year data, including greater detail on the texture production mechanism and the observational signatures.

1.6 Bayesian model selection

The bubble collisions and topological defects discussed in Sections 1.4 and 1.5 are theoretically well-motivated additions to the standard cosmological model. The physical signatures are understood (at least to first order), allowing likelihoods to be written down, and the prior probability distributions of the model parameters are likewise set out. We can therefore use model selection to make probabilistic statements about the degree to which each model is favoured by the latest cosmological data and theoretical knowledge.

Model selection problems require the comparison of posterior probability distributions, and are inherently a question of inference. Cox (1946) showed that Bayesian methods are the only self-consistent framework for such calculations. Bayes' theorem (Bayes and Price, 1763) states that the posterior probability of a model M (selected from a complete set of N models) being true, given a set of data $\mathbf{d}$, is

$$\Pr(M|\mathbf{d}) = \frac{\Pr(M)\Pr(\mathbf{d}|M)}{\sum_{i=1}^{N}\Pr(M_i)\Pr(\mathbf{d}|M_i)}, \tag{1.7}$$

which allows the comparison between two such models to be written as

$$\frac{\Pr(M_1|\mathbf{d})}{\Pr(M_2|\mathbf{d})} = \frac{\Pr(M_1)}{\Pr(M_2)}\frac{\Pr(\mathbf{d}|M_1)}{\Pr(\mathbf{d}|M_2)}. \tag{1.8}$$

The prior probabilities of each model, $\Pr(M)$, are typically chosen to be equal, indicating no *a priori* reason to prefer one model over another. The decision-making power in Bayesian probability theory therefore lies with the evidence, $\Pr(\mathbf{d}|M)$, which is calculated by marginalising the model likelihood over its parameters, $\mathbf{m}$,

$$\Pr(\mathbf{d}|M) = \int d\mathbf{m}\,\Pr(\mathbf{m}|M)\Pr(\mathbf{d}|\mathbf{m}, M). \tag{1.9}$$

Equation 1.9 clearly shows the core concept of Bayesian probability theory: namely the interplay between the data, in the form of the likelihood, $\Pr(\mathbf{d}|\mathbf{m}, M)$, and existing knowledge,

in the form of the parameter priors, $\Pr(\mathbf{m}|M)$. Consequently, Bayesian model selection naturally incorporates Occam's Razor. The appearance of the prior in the evidence calculation means that complex models, using more parameters, are automatically down-weighted by the volume of the extra parameter space, and so must explain the data much better than a simpler model in order to produce the same evidence. Likewise, prescriptive models, where the likelihood is non-zero over a large fraction of the prior volume, will naturally be favoured over models in which a parameter is poorly constrained by theory, and the likelihood is predominantly zero.

The use of Bayesian model selection in cosmology is growing, and has recently been applied to test for extensions to ΛCDM such as curvature (Trotta (2007): a particularly pedagogic paper), and to differentiate between models of inflation (Easther and Peiris, 2012; Mortonson et al., 2011; Norena et al., 2012). Chapters 3, 4 and 5 describe applications of Bayesian model selection to determine whether the standard cosmological model should be augmented with bubble collisions or textures.

1.7 Thesis outline

The Milky Way emits radiation across the electromagnetic spectrum, including the microwave regime. This dominates the CMB radiation in the direction of the Galactic plane, and such regions are therefore often masked in CMB studies. Masking by definition removes information from the data, but it is possible to recover information on the largest scales using reconstruction techniques. In **Chapter 2**, I discuss the performance of the most commonly used reconstruction, highlighting an issue that has recently surfaced in the literature. An improved version of the reconstruction is proposed, to be applied in later chapters.

Chapters 3 to 5 describe the development of a Bayesian algorithm designed to search for the signatures of cosmic textures and bubble collisions in the CMB. The core aim of the algorithm is the calculation of the posterior probability of the average number of signatures present, per CMB sky, given the WMAP 7-year data. Due to the volume of data available, the full calculation is computationally infeasible, and so the algorithm uses candidate-detection techniques – such as spherical needlets (Marinucci et al., 2008; Pietrobon et al., 2008; Scodeller et al., 2011) and optimal filters (Schäfer et al., 2006; McEwen et al., 2008) – to determine the portions of data which contribute to the posterior. Only these regions then need be sampled in order to generate a conservative approximation to the posterior.

The first version of this algorithm, presented in **Chapter 3**, is restricted (for computational reasons) to considering patches of data up to $11°$ in radius. The pipeline is applied to search for bubble collisions, using spherical needlets to detect the regions of interest. In **Chapter 4**,

the restriction on patch size is removed by processing patches at reduced resolution if required, and the resulting algorithm is used to search for cosmic textures. Finally, in **Chapter 5**, the candidate-detection stage is optimised by replacing the spherical needlets with optimal filters based on the signatures of interest. The resulting algorithm is then applied to search for both texture and collision signatures.

Each chapter consists of a paper written by myself and collaborators. I was critically involved in every aspect of these papers, contributing to the design of the formalism and pipeline, coding the algorithm (aside from the optimal filters), performing the data analysis and writing significant portions of the manuscripts. The content of each paper is reproduced as it appears in print; formatting changes have been made to match thesis requirements. The titles, co-authors and publication details of the papers comprising each chapter are listed below.

An additional second-author paper is included as **Appendix B**, as it contains an in-depth treatment of the optimal filters applied in Chapter 5. This work was published as Jason D. McEwen, Stephen M. Feeney, Matthew C. Johnson, Hiranya V. Peiris, 2012, Phys. Rev. D, volume 85, article 103502. I contributed to the design of the algorithm and the study of its performance, including the influence of noise and beam approximations.

The thesis contains references to Feeney et al. (2011c), Feeney et al. (2011a), Feeney et al. (2012) and McEwen et al. (2012), which should be interpreted as references to Chapters 2, 3 and 4 and Appendix B, respectively.

Chapter 2: *Avoiding bias in reconstructing the largest observable scales from partial-sky data*
This work was published as Stephen M. Feeney, Hiranya V. Peiris and Andrew Pontzen, 2011, Phys. Rev. D, volume 84, article 103002, and was carried out in collaboration with the named co-authors.

Chapter 3: *First observational tests of eternal inflation: analysis methods and WMAP 7-year results*
This work was published as Stephen M. Feeney, Matthew C. Johnson, Daniel J. Mortlock, and Hiranya V. Peiris, 2011, Phys. Rev. D, volume 84, article 043507, and was carried out in collaboration with the named co-authors. A summary was also published as Stephen M. Feeney, Matthew C. Johnson, Daniel J. Mortlock, and Hiranya V. Peiris, 2011, Phys. Rev. Lett. volume 107, article 071301.

Chapter 4: *Robust constraint on cosmic textures from the cosmic microwave background*
This work was published as Stephen M. Feeney, Matthew C. Johnson, Daniel J. Mortlock, and Hiranya V. Peiris, 2012, Phys. Rev. Lett., volume 108, article 241301, and was carried

out in collaboration with the named co-authors.

Chapter 5: *Hierarchical Bayesian detection algorithm for early-universe relics in the cosmic microwave background*

This work was carried out in collaboration with Matthew C. Johnson, Jason D. McEwen, Daniel J. Mortlock, and Hiranya V. Peiris, and is in the final stages of preparation for submission to a journal.

Chapter 2

Avoiding bias in reconstructing the largest observable scales from partial-sky data

2.1 Abstract

Obscuration due to Galactic emission complicates the extraction of information from cosmological surveys, and requires some combination of the (typically imperfect) modeling and subtraction of foregrounds, or the removal of part of the sky. This particularly affects the extraction of information from the largest observable scales. Maximum-likelihood estimators for reconstructing the full-sky spherical harmonic coefficients from partial-sky maps have recently been shown to be susceptible to contamination from within the sky cut, arising due to the necessity to band-limit the data by smoothing prior to reconstruction. Using the WMAP 7-year data, we investigate modified implementations of such estimators which are robust to the leakage of contaminants from within masked regions. We provide a measure, based on the expected amplitude of residual foregrounds, for selecting the most appropriate estimator for the task at hand. We explain why the related quadratic maximum-likelihood estimator of the angular power spectrum does not suffer from smoothing-induced bias.

2.2 Introduction

It is unavoidable that we observe the Universe through the galaxy we inhabit. The foreground contamination injected by the Milky Way into full-sky cosmological data-sets must be modeled

and removed, or the regions most conspicuously contaminated must be excised. Where no precise model of the foregrounds is available, cutting the sky is the most robust option, with the regrettable consequence that part of the signal is discarded along with the contamination. This includes information on the largest scales, which are valuable for a variety of reasons, including measurement of the integrated Sachs-Wolfe effect (Sachs and Wolfe, 1967) and constraining primordial non-Gaussianity using tracers of large-scale structure (Slosar et al., 2008).

It is impossible to uniquely recover the cosmological signal discarded in the sky cut. However, by writing down the likelihood for the region of the sky in which one trusts the data, it is possible to reconstruct an estimate of the signal at large scales which maximizes the likelihood of the residual noise (de Oliveira-Costa and Tegmark, 2006). An alternative reconstruction scheme maximizes the *posterior* probability (Wiener, 1964; Bunn et al., 1994; Zaroubi et al., 1995; Tegmark, 1997a; Bielewicz et al., 2004) of measuring the underlying cosmological signal given the available data and a prior theoretical expectation on the signal.

The reconstructions estimate the large-scale (low-ℓ) spherical harmonic coefficients, $a_{\ell m}$, by treating the signal at small scales as noise and only considering data external to the sky cut. If, as with the cosmic microwave background (CMB), the field to be reconstructed is not band-limited, the proliferation of small-scale signal makes the reconstruction noisy to the point of being useless. Input maps are therefore smoothed – necessarily prior to cutting the sky – to truncate the signal and remove sources of confusion below a chosen angular scale (Efstathiou et al., 2010). However, smoothing leaks contamination from the masked region into the trusted data (Aurich and Lustig, 2011; Copi et al., 2011), and the reconstructed spherical harmonic coefficients, $\hat{a}_{\ell m}$, are biased. In this work we explore the causes and expected magnitudes of this bias, and discuss how it can be mitigated.

2.3 Maximum-likelihood reconstruction

We begin with a description of the standard implementation of maximum-likelihood CMB $a_{\ell m}$ reconstruction. The first step of the reconstruction process is to band-limit the temperature field by smoothing, typically with a Gaussian kernel of width $10°$ FWHM. As this removes information on the smallest scales, the map resolution can be downgraded to reduce computation time. The $\hat{a}_{\ell m}$s in the range $2 \leq \ell \leq \ell_{\mathrm{max, rec}}$ (represented for ease as the $n_{a_{\ell m}}$-element vector $\hat{\mathbf{a}}$, where $n_{a_{\ell m}} = (\ell_{\mathrm{max, rec}} - 1)(\ell_{\mathrm{max, rec}} + 3)$) are then reconstructed from the n_{pix} unmasked pixel temperatures, $\mathbf{x}$, using (de Oliveira-Costa and Tegmark, 2006)

$$\hat{\mathbf{a}} = \mathbf{W}\mathbf{x}. \tag{2.1}$$

The reconstructed spherical harmonic coefficients maximize the likelihood of the residual noise, given the available data, if the reconstruction matrix, $\mathbf{W}$, is

$$\mathbf{W} = [\mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1} \mathbf{Y}^t \mathbf{C}^{-1}. \tag{2.2}$$

Here, $\mathbf{Y}$ are the spherical harmonics calculated at each unmasked pixel[1], and $\mathbf{C}$ is the pixel-space noise covariance matrix

$$C_{ij} = R_{ij} + \sum_{\ell=\ell_{\mathrm{max,\,rec}}+1}^{\ell_{\mathrm{max}}} \frac{2\ell+1}{4\pi} \bar{C}_\ell P_\ell(\hat{\mathbf{r}}_i \cdot \hat{\mathbf{r}}_j), \tag{2.3}$$

where $\mathbf{R}$ is uncorrelated, low-amplitude regularizing noise added to prevent $\mathbf{C}$ from becoming singular, $\bar{C}_\ell$ is the smoothed theory CMB angular power spectrum, and P_ℓ are the Legendre polynomials at unmasked pixels i, j. The sum over the multipoles $\ell_{\mathrm{max,\,rec}} < \ell \leq \ell_{\mathrm{max}}$ ensures that the small-angular-scale CMB multipoles we do not wish to reconstruct are treated as noise. As stated above the CMB power must be artificially truncated to restrict the number of ambiguous modes accessible to the reconstruction. The smoothing kernel is deconvolved from the $\hat{a}_{\ell m}$s after reconstruction by dividing the $\hat{a}_{\ell m}$s by the kernel's spherical harmonic transform.

If a foreground signal b_i is now introduced, so that $\mathbf{x} = \mathbf{Y}\mathbf{a} + \mathbf{b} + \mathbf{n}$ and $\mathbf{a}$ is the CMB signal uncorrelated with $\mathbf{b}$ and the noise $\mathbf{n}$, the mean and variance of the reconstruction error $\epsilon_{\ell m} = \hat{a}_{\ell m} - a_{\ell m}$ are

$$\langle \boldsymbol{\epsilon} \rangle = \mathbf{W}\mathbf{b} \tag{2.4}$$

and

$$\langle \boldsymbol{\epsilon}\boldsymbol{\epsilon}^t \rangle - \langle \boldsymbol{\epsilon} \rangle \langle \boldsymbol{\epsilon}^t \rangle = \mathbf{W}\mathbf{C}\mathbf{W}^t, \tag{2.5}$$

respectively.

Throughout this work, we reconstruct the spherical harmonic coefficients up to $\ell_{\mathrm{max,\,rec}} = 10$. The noise covariance matrix includes CMB power in the range $\ell_{\mathrm{max,\,rec}} < \ell \leq \ell_{\mathrm{max}} = 32$ unless explicitly stated; this value is chosen such that modes with $\ell > \ell_{\mathrm{max}}$ are suppressed to $\mathcal{O}(\text{few }\%)$ by the smoothing. The WMAP 5-year best-fit C_ℓs (Nolta et al., 2009) are chosen for the theory CMB angular power spectrum[2]. Input maps are smoothed at HEALPix (Górski et al., 2005) resolution $N_{\mathrm{side}} = 512$ before being downgraded to $N_{\mathrm{side}} = 16$ to retain the information required for the reconstruction while minimizing the number of pixels included in the noise covariance matrix. Diagonal regularizing noise $\mathbf{R}$ is added at the level of $2\,\mu\mathrm{K}^2$ to allow the inversion of

---

[1]Without loss of generality, the reconstruction matrix in this work is formed from the *real* spherical harmonics.
[2]Our results are not sensitive to the small differences between different WMAP releases in the best-fit cosmology.

the noise covariance matrix despite the presence of some null modes (which are irrelevant to the reconstruction).

2.4   Smoothing-induced bias

This section outlines how a bias arises from smoothing-induced contamination of the unmasked pixels (Aurich and Lustig, 2011; Copi et al., 2011). For clarity, we illustrate the smoothing-induced bias in the maximum-likelihood reconstruction (and, later, our proposed solutions) with results plotted in both harmonic and pixel space.

It is beyond the scope of this work to estimate accurate foreground residuals resulting from different component separation methods; instead, we choose some residuals for the purpose of illustration. Following Pontzen and Peiris (2010), the residual foregrounds are taken to be 1% of the difference between the WMAP 7-year Internal Linear Combination (ILC) (Gold et al., 2011) and V-band (Jarosik et al., 2011) temperature maps. The resulting map is *indicative* of the extent and amplitude of the residual foregrounds in the WMAP 7-year ILC, the data-set for which the smoothing-induced bias was first described (Efstathiou et al., 2010; Aurich and Lustig, 2011; Copi et al., 2011). It is important to note that a considerably higher level of contamination is present in the foreground-reduced maps for each individual WMAP frequency band provided by the WMAP team: contaminants of ~ 50 times those used here are visible in these maps. The residual foregrounds are restricted to the pixels within the sky cut, which in this work masks only the Galaxy and not individual point-sources. These degree-scale point-source cuts are, unsurprisingly, not found to significantly bias the large-scale reconstructed $\hat{a}_{\ell m}$s, and so for clarity Galaxy-only contaminants and masks are considered.

The addition of simulated residuals and the spherical harmonic transform are linear, so the smoothing-induced bias is given by the reconstructed $\hat{a}_{\ell m}$s of the simulated residual map. We take the "standard" $10°$-FWHM-Gaussian-smoothed maximum-likelihood reconstruction, using data outside the Galaxy-only part of the WMAP 7-year KQ85 mask, as our fiducial maximum-likelihood estimator (hereafter the "Gaussian ML" $\hat{a}_{\ell m}$s). The Gaussian ML $\hat{a}_{\ell m}$s generated from the simulated residual foregrounds are plotted (deep-blue solid line) in Fig. 2.1, along with the full-sky $a_{\ell m}$s (dotted line) for comparison. The smoothing-induced leakage from within the sky cut is clear to see: the reconstruction picks up about half of the power in the foreground residuals, even though the simulated residuals are entirely confined to the sky cut.

The largest peaks in the bias affect the $\hat{a}_{\ell m}$s satisfying $\ell = 2n, m = 0$ for integer n (in Fig. 2.1, HEALPix index $\ell^2 + \ell + m = \{6, 20, 42 \ldots\}$) (de Oliveira-Costa and Tegmark, 2006), and are positive for odd n, negative for even n. The pattern of these peaks can be explained by

Figure 2.1: The spherical harmonic coefficients of the simulated foreground residuals, calculated using the Gaussian ML reconstruction (deep-blue solid line), the full-sky data (dotted line) and the 10° Top-Hat ML reconstruction (light-blue solid line on x-axis). It is clear that the Gaussian ML reconstruction leaks around half of the information from within the sky cut; this can be counteracted by smoothing with a top-hat kernel and using an extended mask. The HEALPix index $\ell^2 + \ell + m$ maps each ℓ, m combination to a unique index into the array of $a_{\ell m}$s.

examining the reconstruction of the simulated Galactic residuals, plotted in Fig. 2.2, which are coldest along the Galactic plane. Smoothing these residuals reduces the pixel temperature values approximately symmetrically around the Galactic mask, and therefore pollutes the azimuthal modes which are also symmetric about the equator. The bias is positive for modes which have minima at the equator, and negative for those with maxima. There are also secondary peaks at $\ell = 2n + 1, m = 1$ (in Fig. 2.1, $\ell^2 + \ell + m = \{13, 31, 57 \ldots\}$), which again are positive for odd n, negative for even n. These modes pick out the concentration of reconstructed foreground power in the Galactic centre.

The leakage of information from within the sky cut can also be demonstrated in pixel-space. Taking the WMAP 7-year ILC map, the full-sky $a_{\ell m}$s are extracted, and the Gaussian ML reconstruction is performed. The spherical harmonic coefficients recovered in each case are then used to reconstruct the input ILC map using only $2 \leq \ell \leq 10$, as plotted in Fig. 2.3. The maps formed from the full-sky-$a_{\ell m}$s (top-left) and reconstructed from Gaussian-ML-$\hat{a}_{\ell m}$s (top-right) are almost identical, even in the Galactic plane, confirming that the reconstruction has access to information well inside the sky cut.

Figure 2.2: Clockwise, from top left: simulated Galactic foregrounds, smoothed by a 10°-FWHM Gaussian and masked with the Galaxy-only KQ85 mask; the $\ell \leq 10$ Gaussian ML $\hat{a}_{\ell m}$s reconstructed from the simulated Galactic foregrounds; the real spherical harmonics Y_{51} and Y_{40}. The simulated foregrounds yield a negative bias in the Y_{40} mode, and a positive bias in the Y_{51} mode.

Although the bias illustrated in Fig. 2.1 looks problematic, its amplitude is at least reduced over using the contaminated full sky, and so the simple Gaussian-ML procedure may yet turn out to be useful. To compare with other possible approaches, we first need to discuss its standard deviation – i.e. the scatter induced by the $\ell > \ell_{\max,\,\mathrm{rec}}$ modes – which is calculated using Eq. 2.5 and plotted (as the deep-blue narrowest band) in Fig. 2.4. This is a few μK at most. For a given ℓ, the modes that are reconstructed with the least precision are those with $|m| = \ell$, with the $m = \ell$ modes typically the worst. This confirms the observation in Pontzen and Peiris (2010) that the sky cut removes the most information from modes with power concentrated towards the equator, and particularly those with extrema at $\phi = 0°$, where the mask is at its widest (Copi et al., 2011). The mask is plotted for reference in Fig. 2.5, along with examples of the affected modes. We see that, typically, three modes per ℓ will have increased bias or variance, but for most modes both the mean reconstruction error and its variance will be small.

In order to compare estimators, we must first quantify their performance over the range of multipoles considered. The performance measure

$$
\begin{aligned}
z_{\ell m} &= \langle \epsilon_{\ell m}^2 \rangle \\
&= \mathrm{mean}(\epsilon_{\ell m})^2 + \mathrm{var}(\epsilon_{\ell m})
\end{aligned}
\tag{2.6}
$$

48
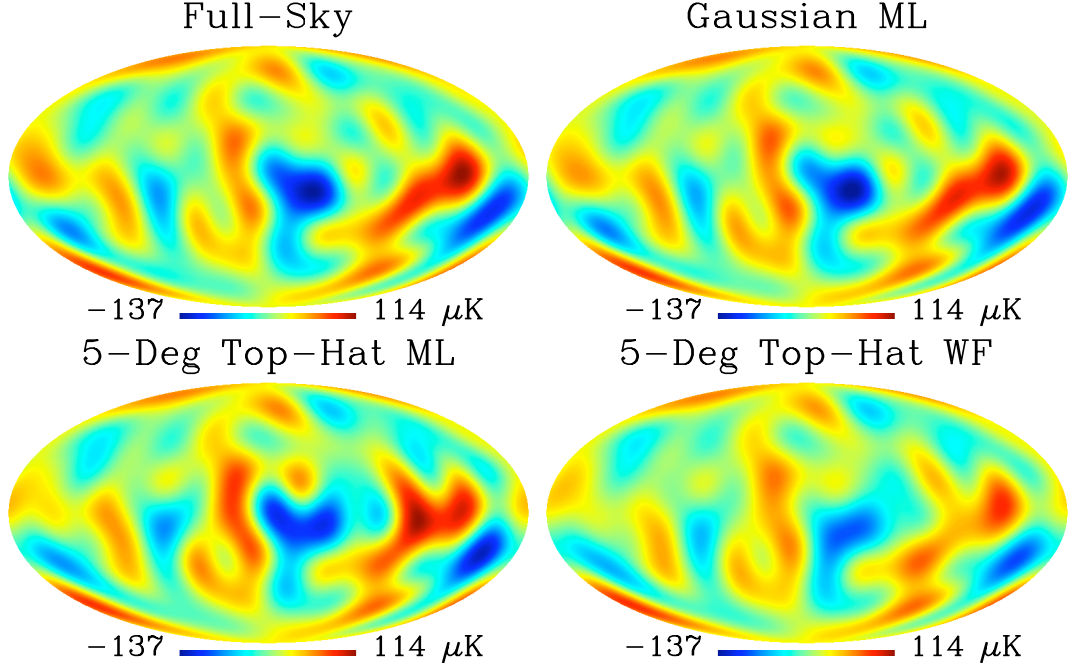
Figure 2.3: The 7-year ILC $2 \le \ell \le 10$ modes, plotted using (clockwise from top-left) full-sky $a_{\ell m}$s, Gaussian ML $\hat{a}_{\ell m}$s, 5° Top-Hat WF $\hat{a}_{\ell m}$s and 5° Top-Hat ML $\hat{a}_{\ell m}$s.

provides the expected size of the reconstruction error $\epsilon_{\ell m}$ for each mode: summing over all modes

$$Z = \sum_{\ell,\, m} z_{\ell m} \tag{2.7}$$

therefore yields a complete measure of each estimator's performance. Any alternative estimator which removes the smoothing-induced bias should be preferred only if its Z value is lower than that of the Gaussian ML reconstruction, and, indeed, the contaminated full-sky $a_{\ell m}$s. In fact, using the 1% ILC-V foreground residuals, $Z \simeq 33\,\mu\mathrm{K}^2$ using the full-sky $a_{\ell m}$s, compared to $\sim$ $265\,\mu\mathrm{K}^2$ for the Gaussian ML estimator. If the residual foregrounds employed in this work are an accurate reflection of those present in the WMAP 7-year ILC map, then the contaminated full-sky $a_{\ell m}$s provide a better estimate of the cosmological signal than the Gaussian ML reconstruction. The second form of Eq. 2.6 shows that both bias (see Fig. 2.1) and variance (see Fig. 2.4) in the reconstruction increase the value of Z. The variance term is independent of the contamination, while the bias scales linearly with the contamination. Therefore, if the amplitude of residual foregrounds in the ILC map is higher than in our illustrative example, the reduction in bias due to the use of the Gaussian ML $\hat{a}_{\ell m}$s will eventually overcome the variance introduced by the reconstruction. For residual levels $3 - 4$ times higher than those used here, the Gaussian ML reconstruction should be used instead of the contaminated full-sky $a_{\ell m}$s. However, as we have seen, the Gaussian ML estimator (as implemented thus far) does not eliminate the bias due to

Figure 2.4: The one-standard-deviation ranges of the reconstruction error $\epsilon_{\ell m}$ for the Gaussian ML $\hat{a}_{\ell m}$s (narrowest, deep-blue band), 5° Top-Hat ML $\hat{a}_{\ell m}$s (widest, light-blue band), and Top-Hat WF $\hat{a}_{\ell m}$s (intermediate, mid-blue band).

smoothing-related leakage of contaminants from within the masked region.

While our simulated foreground residuals are simply meant to be indicative, we nevertheless expect that the smoothing bias is mainly sensitive to the *amplitude* of the residuals, and not their precise morphology. This can be seen in Fig. 2.6, where we have modeled the residuals as a simple bar in the Galactic plane, while rescaling the amplitude to match our 1%(ILC-V) model. This highly simplified model is able to capture most of the features of the bias in harmonic space, as seen in the lower panels of this Figure.

2.5 Eliminating the bias

At this stage, we are presented with something of a conundrum: smoothing is essential to the reconstruction process, but it is exactly this smoothing that is biasing the results. The simplest solution to this issue is to remove the areas of the sky that are within one smoothing scale of the main Galactic sky cut. However, the smoothing kernel typically used in the standard reconstruction algorithm is a Gaussian, with support across the full sky in pixel space, and the set of contaminated pixels is hence poorly defined. This problem can be solved by using a kernel with finite pixel-space support – for simplicity we choose a top-hat – as all contaminated pixels
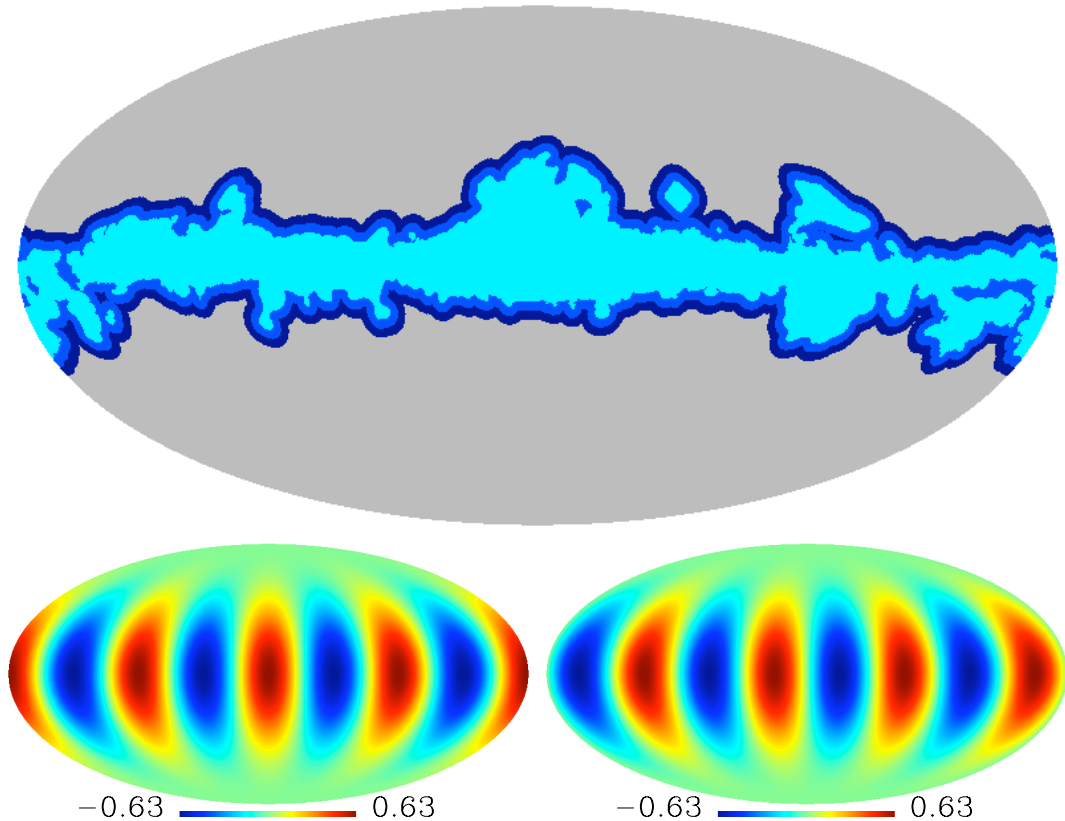
Figure 2.5: Top: the WMAP 7-year Galaxy-only KQ85 mask (light-blue central region) extended by 2.5° (mid-blue) and 5° (dark-blue). Bottom: the real spherical harmonics $Y_{4\,4}$ and $Y_{4\,-4}$. The concentration of $Y_{\ell\,\pm\ell}$ mode power towards the equator results in increased estimator variance in those modes.

fall within a kernel's radius of the mask.

Naïvely selecting the diameter of the top-hat smoothing kernel to be 10°, we reconstruct the simulated foreground residuals using the Galaxy-only KQ85 mask extended by 5° – hereafter the "10° Top-Hat ML" reconstruction. The results are plotted in light blue (pale line along x-axis) in Fig. 2.1: the smoothing-induced bias has been eliminated. However, the measure of reconstruction quality has deteriorated dramatically to $Z \simeq 11\,252\,\mu\mathrm{K}^2$, significantly worse than the Gaussian ML reconstruction. There are two reasons for this increase in Z, which is now sourced entirely by increased variance in the reconstruction. Firstly, the top-hat smoothing kernel has support over a greater range of multipoles than the Gaussian kernel, and so more ambiguous modes contribute to the covariance matrix (in this case, and for all further kernels, we increase ℓ_{max} to 1024 to capture all relevant modes, even though at 10° the smoothing kernel is effectively band-limited at $\ell \sim 100$).[3] The second factor is that the reconstruction-error variance increases very rapidly with the area of the sky that is masked (de Oliveira-Costa and Tegmark, 2006). This suggests the use of narrower smoothing kernels, although this necessarily increases

---

[3]Truncating the kernel by selecting a lower ℓ_{max} leads to $\mathcal{O}(10\%)$ variations in the results.

Figure 2.6: Top: the residual foreground map employed in this work (left), and a simple model comprising a 3°-wide $-20\,\mu$K equatorial band. Bottom: the full-sky $a_{\ell m}$s (dotted) and reconstructed Gaussian ML $\hat{a}_{\ell m}$s (solid) of the residual foreground maps. The simple band model captures most of the features of the smoothing-induced bias injected by the more-complex residuals.

the power of the high-ℓ noise.

The interplay between the variance injected by decreased smoothing and increased masking is shown in Fig. 2.7. Here, the total reconstruction-error variance is plotted for top-hat smoothing kernels of diameter $3° - 10°$, and hence mask extensions of $1.5° - 5°$. The reduction in variance due to minimizing the sky cut dominates the added noise from narrower smoothing. We therefore select the width of the smoothing kernel to be as small as possible, given the resolution of the degraded map. At $N_{\mathrm{side}} = 16$, the pixels are $\sim 3°$ across, so to avoid injecting bias through pixelization (Copi et al., 2011) (which would not be captured by Eq. 2.5) we choose our optimal kernel diameter to be just larger: $5°$. Hereafter, we refer to this reconstruction – using the $5°$ top-hat smoothing kernel and Galaxy-only KQ85 mask extended by $2.5°$ – as the "$5°$ Top-Hat ML" reconstruction.

The reconstruction-error variance calculated for the $5°$ Top-Hat reconstruction is plotted as the light-blue outermost region of Fig. 2.4. Even using the minimum possible mask extension the reconstruction-error variance is still an order of magnitude larger than that of the Gaussian ML estimator. In terms of the measure of reconstruction quality, the biased Gaussian ML reconstruction ($Z \sim 265\,\mu$K^2) should be strongly preferred to the $5°$ Top-Hat ML case ($Z \sim 8466\,\mu$K^2) for residual levels comparable to those used in this work. As the Top-Hat ML reconstruction is

Figure 2.7: The impact of mask extension and smoothing-kernel diameter on the quality of the bias-free maximum-likelihood reconstruction. Reconstruction-error variances are calculated using top-hat smoothing kernels of varying diameter, and extending the KQ85 mask by one kernel radius each time. Smoothing at the lowest-possible scale will produce the most faithful reconstruction. Note that the curve is not smooth as both the extent and *shape* of the mask change as it is extended.

unbiased, this value of Z is fixed (for a given sky cut). Thus, only if the residuals are greater than $\sim 25\%$ ILC$-$V will the 5° Top-Hat ML reconstruction outperform the Gaussian ML reconstruction. Note that the quality of the reconstruction could be improved further if it was performed at higher resolution, as smaller smoothing kernels could be used. This will necessarily have to be traded off against the increased computational requirements.

2.6 Reducing the variance

The increase in variance encountered when using extended sky cuts is far beyond that expected due to the reduction in pixel count: reducing f_{sky} from 81% to 74% should, assuming uncorrelated pixels for simplicity, increase the variance by only $\sim 10\%$. The dominant issue is that the maximum-likelihood reconstruction allows the temperature field in the masked region of the sky to have infinite variance. For small sky cuts (and small $\ell_{\mathrm{max, rec}}$) this is fine: one cannot "hide" large-scale power within the cut, and so the variance on the large-scale $\hat{a}_{\ell m}$s is low. Extending the KQ85 mask not only increases its overall width, but also closes a number of small gaps that

allow the estimator limited access to the poorly-constrained equatorial modes (see Fig. 2.5). The estimator is therefore free to fill the cut with significant low-ℓ power (compare Fig. 2.3 top-right and bottom-left), and the estimator variance rises rapidly.

The variance of the reconstruction error can be reduced by enforcing a prior on the power within the Galactic cut using Wiener-filtering (Method 5 in Tegmark (1997a); see also Bielewicz et al. (2004); Wiener (1964); Bunn et al. (1994); Zaroubi et al. (1995)). The Wiener-filtered $\hat{a}_{\ell m}$s then maximize the posterior probability of reconstructing the underlying $a_{\ell m}$s, given the trusted data and theoretical power spectrum. In practice, this is achieved by adding a theoretical (inverse) covariance matrix for the multipoles of interest to the reconstruction matrix

$$\mathbf{W} = [\mathbf{S}^{-1} + \mathbf{Y}^t\mathbf{C}^{-1}\mathbf{Y}]^{-1}\mathbf{Y}^t\mathbf{C}^{-1}. \tag{2.8}$$

$\mathbf{S}$ here is a diagonal $n_{a_{\ell m}} \times n_{a_{\ell m}}$ matrix with elements equal to the smoothed theory power spectrum $\bar{C}_\ell$. The theory prior restricts this "Top-Hat WF" reconstruction from filling the sky cut with arbitrary power (see Fig. 2.3 bottom-right), and reduces the variance of the estimator's reconstruction error accordingly (see the mid-blue intermediate region in Fig. 2.4). The analytic expectation for the reconstruction-error variance is

$$\langle \boldsymbol{\epsilon}\boldsymbol{\epsilon}^t \rangle - \langle \boldsymbol{\epsilon} \rangle\langle \boldsymbol{\epsilon}^t \rangle = \mathbf{S}\left[\mathbf{S} + \mathbf{N}\right]^{-1}\mathbf{N}, \tag{2.9}$$

where $\mathbf{N}$ is the noise power spectrum, defined to be $\mathbf{N} = [\mathbf{Y}^t\mathbf{C}^{-1}\mathbf{Y}]^{-1}$ (*i.e.* the variance of the Top-Hat ML reconstruction error). By adding a theory prior to the power within the sky cut, and hence requiring finite power in that region, the Wiener-filtered reconstruction tends to produce $\hat{a}_{\ell m}$s that are closer to zero than the maximum-likelihood case. While this could be seen as biasing the $\hat{a}_{\ell m}$s toward lower values[4], it can also be interpreted as being conservative, and applying the prior belief that the information within the mask is similar to the trusted information outside the mask. In other words, we should be happy to trade off a small multiplicative bias against a significant reduction in variance.

This is automatically encapsulated in the measure of reconstruction quality Z for the 5° Top-Hat WF estimator, which has improved to $\sim 1521\,\mu\mathrm{K}^2$. However this is still worse than that of the Gaussian ML estimator. The bias arising in the Top-Hat WF reconstruction is not from smoothing but from a prior, so Z is fixed for a given mask, and always lower than that of the corresponding Top-Hat ML reconstruction. For contamination levels of $\gtrsim 10\%$ ILC$-$V (such as those found in the foreground-reduced maps for the individual WMAP frequency bands), the

---

[4]Assuming for clarity zero noise, the Wiener-filtered reconstruction yields $\hat{\mathbf{a}} = \mathbf{S}[\mathbf{S} + \mathbf{N}]^{-1}\mathbf{a}$, *i.e.* a multiplicative bias. Note that the ensemble average $\langle \hat{\mathbf{a}} \rangle = \langle \mathbf{a} \rangle = 0$.

5° Top-Hat WF reconstruction therefore represents the most reliable estimator considered in this work.

We do not know the precise magnitude or morphology of the residual foregrounds in the WMAP data. We can, however, examine the Gaussian ML and Top-Hat WF reconstructions of the ILC by eye to determine if there is an obvious impact due to residual foregrounds. The $\hat{a}_{\ell m}$s generated from these two reconstructions are plotted in Fig. 2.8, along with the *estimator* standard deviation ($\sqrt{\langle |\hat{a}_{\ell m}|^2 \rangle - \langle \hat{a}_{\ell m} \rangle^2}$). Comparing the two plots, we see that there is very little difference between the $\hat{a}_{\ell m}$s returned in each case. Further, no modes look statistically anomalous at the 3σ-level, even those that we expect to be contaminated from the simple residuals model used here.

2.7 Relation of $a_{\ell m}$ reconstruction to the QML estimator for the C_ℓs

We have so far discussed estimating the full-sky $a_{\ell m}$s from cut-sky data, which is equivalent to reconstructing the smoothed temperature field. However, in the context under which the smoothing-induced bias was revealed (Efstathiou et al., 2010; Aurich and Lustig, 2011; Copi et al., 2011) it is in fact only the angular power spectrum C_ℓ of the temperature field which is required.

A popular method for estimating the full-sky angular power spectrum is to adopt the quadratic maximum-likelihood estimator as first derived in Tegmark (1997b). It has been noted (see e.g. Section 3 of Efstathiou et al. (2010) for a complete discussion) that the QML estimator can be formed using the maximum-likelihood $\hat{a}_{\ell m}$s. On the surface, the QML estimates (henceforth denoted $\hat{C}_\ell$) may therefore seem to be susceptible to similarly problematic contamination from a smoothing stage.

However, this is not the case: in fact the $\hat{C}_\ell$s are far more robust to the content of the cut because the smoothing can be conducted on vastly smaller scales (e.g. $1°$ in Pontzen and Peiris (2010)). Note that Copi et al. (2011) miss this point, because they consider only two extreme cases: (i) smoothing at $10°$ and (ii) failing to smooth. They therefore reach the erroneous conclusion that the QML estimator is susceptible to contamination from within the mask. We explicitly verified that the pipeline used by Pontzen and Peiris (2010) is independent of any contamination placed fully inside the mask.

The above paragraphs at first appear to be contradictory, since they simultaneously claim (a) that the QML power spectrum estimates can be formed out of the ML temperature field

reconstruction; and (b) that the QML power spectrum estimates can still be constructed from maps smoothed on degree scales (whereas the $\hat{a}_{\ell m}$s will necessarily become noisy for sufficiently high $\ell_{\max, \mathrm{rec}}$). However, this is not a true contradiction because the QML estimates are not formed directly from the noisy $\hat{a}_{\ell m}$s, but rather through an expression (Eq. 23 of Efstathiou et al. (2010)) which specifically downweights poorly constrained modes. It is this cautious treatment of ambiguous modes which makes power spectrum estimation, as opposed to $a_{\ell m}$ reconstruction, so well-behaved, irrespective of the shape of the smoothing kernel employed.

2.8   Discussion

Maximum-likelihood estimators, $\hat{a}_{\ell m}$, are often used to reconstruct the large-scale spherical harmonic coefficients, $a_{\ell m}$, from partial-sky data. The technique relies on smoothing to restrict the amount of small-scale noise accessible to the reconstruction, but smoothing has been shown to contaminate "clean" pixels with residual foregrounds from within the sky cut. In this work, we have examined the impact of this smoothing-induced bias on the maximum-likelihood reconstruction. We have shown that it is possible to mitigate the bias by removing the contaminated regions, but these are only well-defined if smoothing is performed using a kernel with finite support on the sky. This precludes the use of the commonly used Gaussian kernel. Cutting a larger portion of the sky greatly increases the variance of the reconstruction, but it is possible to counteract this effect by enforcing a prior on the reconstructed coefficients using a Wiener filter. We have therefore proposed an estimator – using top-hat smoothing, extended masks and a Wiener-filtered reconstruction – which does not suffer from smoothing-induced bias. By considering the expectation of the square of the reconstruction error, $Z = \sum_{\ell, m} \langle (\hat{a}_{\ell m} - a_{\ell m})^2 \rangle$, we have compared the performance of the maximum-likelihood and Wiener-filtered estimators in the presence of simulated CMB foreground residuals.

The reconstruction performance measure Z scales with the estimators' bias and variance, which in turn are governed by the amplitude of contamination and the size of the sky cut, respectively. The fiducial maximum-likelihood reconstruction is performed using relatively small sky cuts, but is susceptible to contamination through smoothing-induced bias; the finite-smoothing Wiener-filtered reconstruction does not suffer from smoothing-induced bias, but makes use of extended masks. Increasing the level of contamination therefore increases Z for the maximum-likelihood reconstruction only, which suggests that there is a level of contamination above which one should switch from the maximum-likelihood to the Wiener-filtered reconstruction.

Given an estimate of the morphology and amplitude of the contaminants within the cut sky, one can predict which modes will be biased and by how much, and hence determine the threshold

at which one should swap estimators. We find that this threshold is relatively insensitive to the precise morphology of foreground residuals at large scales, and is mainly governed by their amplitude. Calculating Z for the two estimators in the presence of estimated foreground residuals, we determine this threshold to be ~ 10 times the amplitude of the foreground residuals used in this work. Assuming that the ILC contains similar levels of contamination to those used here, we therefore recommend the use of either the contaminated full-sky $a_{\ell m}$s or the fiducial maximum-likelihood $\hat{a}_{\ell m}$s when handling this data-set. However, when using foreground-reduced maps for individual WMAP frequencies, which contain much greater foreground residuals, the Wiener-filtered reconstruction will provide the best estimate of the large-scale underlying CMB signal. Note that, as the Wiener-filtered $\hat{a}_{\ell m}$s are a maximum-*posterior* solution, care must be taken if the reconstruction output is being used for further model-selection steps. The reconstruction techniques are, however, most commonly used to test the null hypothesis, in which case the prior employed in this work is completely appropriate.

For problems requiring only a power spectrum (as opposed to the full temperature field) the issues described in this paper are essentially irrelevant because the smoothing can be conducted on vastly smaller scales, the resulting range of poorly constrained modes being automatically downweighted.

2.9 Acknowledgements

Gaussian ML ILC Reconstruction

Top−Hat WF ILC Reconstruction

Figure 2.8: The reconstructed WMAP 7-year ILC $\hat{a}_{\ell m}$s, calculated using the Gaussian ML reconstruction (top) and the 5° Top-Hat WF reconstruction (bottom). The shaded areas represent the estimator standard deviations. The modes that are most contaminated by the simulated foregrounds in the Gaussian ML reconstruction are indicated, along with their expected sign, by dashed ($Y_{\mathrm{even}\,0}$) and dash-dotted ($Y_{\mathrm{odd}\,1}$) lines.

Chapter 3

First observational tests of eternal inflation: analysis methods and WMAP 7-year results

3.1 Abstract

In the picture of eternal inflation, our observable universe resides inside a single bubble nucleated from an inflating false vacuum. Many of the theories giving rise to eternal inflation predict that we have causal access to collisions with other bubble universes, providing an opportunity to confront these theories with observation. We present the results from the first observational search for the effects of bubble collisions, using cosmic microwave background data from the WMAP satellite. Our search targets a generic set of properties associated with a bubble collision spacetime, which we describe in detail. We use a modular algorithm that is designed to avoid *a posteriori* selection effects, automatically picking out the most promising signals, performing a search for causal boundaries, and conducting a full Bayesian parameter estimation and model selection analysis. We outline each component of this algorithm, describing its response to simulated CMB skies with and without bubble collisions. Comparing the results for simulated bubble collisions to the results from an analysis of the WMAP 7-year data, we rule out bubble collisions over a range of parameter space. Our model selection results based on WMAP 7-year data do not warrant augmenting ΛCDM with bubble collisions. Data from the *Planck* satellite can be used to more definitively test the bubble collision hypothesis.

3.2   Introduction

Observations of the cosmic microwave background (CMB) radiation have ushered in a new era of precision cosmology. Full-sky temperature maps produced by the *Wilkinson Microwave Anisotropy Probe* (WMAP) (Bennett et al., 2003a) have confirmed with high precision that the observed temperature fluctuations are consistent with a nearly Gaussian and scale invariant primordial power spectrum, as predicted by inflation. The recently launched *Planck* satellite (Tauber et al., 2010) has a resolution three times better than that of WMAP, with an order of magnitude greater sensitivity, and significantly wider frequency coverage (allowing for far more robust foreground removal, and therefore reduced systematics). These high quality data sets allow for the possibility of observing deviations from the standard inflationary paradigm, some of which could have drastic consequences for our understanding of the universe and its origins.

Perhaps the largest gap in our description of the early universe lies in an understanding of its initial conditions. One possibility, motivated by the proliferation of vacua in compactifications of string theory (known as the string theory landscape (Susskind, 2003)), is that our observable universe is only a tiny piece of a vast multiverse, the majority of which is still inflating. This picture of *eternal inflation* (for a review, see, e.g. Aguirre (2008)) arises when the rate at which local regions exit an inflating phase is outpaced by the accelerated expansion of the inflating background. Eternal inflation is a fairly generic consequence of any theory containing positive vacuum energy and multiple vacua, highlighting the importance of understanding how this scenario might be confronted with observational tests.

The first attempts to embed our cosmology inside an eternally inflating universe led to "open inflation" (Bucher et al., 1995; Gott, 1982); see Garcia-Bellido (1997) for a review. In this scenario, a scalar field (or set of scalar fields) has a potential with a high energy metastable minimum that drives the eternally inflating phase. Transitions out of this vacuum proceed via the Coleman-de Luccia (CDL) instanton (Coleman, 1977; Coleman and De Luccia, 1980), resulting in expanding bubbles inside which the scalar field rests on an inflationary plateau. The symmetries of the CDL instanton ensure that there is a very nearly homogeneous and isotropic open universe inside the bubble; inflation, reheating, and standard cosmological evolution follow.

In any given bubble, the future light cone of the nucleation event forms the "Big Bang" (where the scale factor vanishes) of an open FRW universe. The eternally inflating phase outside our bubble can therefore be thought of as a pre-Big Bang epoch, and one might expect inflation to erase any of the scant observational evidence of our parent vacuum. In single bubble open inflation, various anomalies are induced in the CMB temperature power spectrum (see Garcia-Bellido (1997) and references therein), but unfortunately, the size of these effects decreases with

the present energy density in curvature (related to the number of inflationary e-folds), rendering them negligible at all but the lowest multipoles where cosmic variance dominates. However, our bubble does not evolve in isolation. There are other nucleation events from the false vacuum, containing a phase that might be identical to ours, or perhaps very different. If one of these secondary nucleation events occurs close enough to our bubble wall, then a collision inevitably results. In fact, since our bubble grows to reach infinite size, there are an infinite number of collisions (Guth and Weinberg, 1981; Guth and Weinberg, 1983; Gott and Statler, 1984; Garriga et al., 2007) (a finite subset of which are causally accessible to any one observer). This raises the possibility that if such collisions are both survivable and only small perturbations on top of standard cosmology, they might leave observable signatures of eternal inflation (Aguirre et al., 2007); it is these signatures which our analysis targets.

If we are to detect such bubble collisions, their predicted signatures must be consistent with our observed cosmology, but sufficiently distinct to be differentiated from other possible signals in the CMB. In addition, the theory must predict that we expect to have causal access to bubble collisions. While these criteria are not met in every model of eternal inflation, recent work (Guth and Weinberg, 1981; Guth and Weinberg, 1983; Hawking et al., 1982; Chao, 1983; Gott and Statler, 1984; Garriga et al., 2007; Aguirre et al., 2007; Aguirre and Johnson, 2008; Aguirre et al., 2009; Chang et al., 2008, 2009; Dahlen, 2010; Freivogel et al., 2009; Easther et al., 2009; Larjo and Levi, 2010; Zhang and Piao, 2010; Czech et al., 2010) (for a review, see Aguirre and Johnson (2009)) has established that bubble collisions could in some theories be both expected and detectable. Bubble collisions produce a fairly characteristic set of inhomogeneities in the very early universe, which are processed into temperature anisotropies in the CMB. From the spherical symmetry of the colliding bubbles, the collision possess azimuthal symmetry, and by causality must be confined to a disc on the sky. The CMB temperature and its derivatives need not be constant across the causal boundary. Therefore, the signals we are searching for are localized, and because they are primordial, consist of a long-wavelength modulation of the standard inflationary density fluctuations inside the affected region (Chang et al., 2009). The amplitude and angular scale of the signal is dependent upon the underlying model and kinematics of the collision.

These general features suggest a set of strategies for data analysis. The localization of the collision implies that wavelet analysis could be a sensitive tool for picking out both the location and angular scale of a candidate signal. The causal boundary, across which the temperature and its derivatives need not be constant, suggests the use of edge detection algorithms similar to those used in searches for cosmic strings (Kaiser and Stebbins, 1984; Lo and Wright, 2005;

Danos and Brandenberger, 2010; Amsel et al., 2008). Finally, the prediction that the temperature modulation induced by the collision is rather long-wavelength yields a sufficiently generic template to perform a full Bayesian parameter estimation and model selection analysis.

In this paper, we describe a modular analysis algorithm designed to look for the signatures of eternal inflation, and apply it to the WMAP 7-year data (Jarosik et al., 2011). This algorithm can easily be adapted to test any model that predicts a population of spatially localized sources in addition to the standard fluctuations predicted by ΛCDM. A summary of our results was presented in Feeney et al. (2011b); in this paper we describe our analysis in detail. Currently available full sky CMB data are rather limited in their sensitivity to the signatures of bubble collisions listed above; the main current limitation is the low resolution. Therefore, we apply our algorithm to current data mainly as a validation exercise; to exploit its full power would require future high resolution data, e.g., from *Planck*.

The individual steps of our analysis pipeline are calibrated using realistic simulations of the WMAP experiment with and without bubble collisions. The calibrated pipeline applied to data is fully automated, identifying the candidate signals and processing them without any human intervention. This removes any *a posteriori* choices from our analysis, which must be carefully avoided in any analysis of a large data-set such as the WMAP 7-year data (Bennett et al., 2011).

The plan of the paper is as follows. In Sec. 3.3, we review some of the background on bubble collisions in eternal inflation, and outline the predicted observable signatures. Our analysis pipeline is summarized in Sec. 3.4. We describe some properties of the WMAP experiment and our simulations in Sec. 3.5, and detail our analysis tools in Sec. 3.6. Sec. 3.7 summarizes the results of our analysis of the WMAP 7-year data, and we conclude in Sec. 3.8.

3.3 The observable effects of bubble collisions

The simplest model of eternal inflation involves a single scalar field in four dimensions, with a double-well potential. In many models (as long as the average curvature of the potential between the minima is small compared to the Planck scale), the Coleman-de Luccia (CDL) instanton (Coleman, 1977; Coleman and De Luccia, 1980) mediates a transition from the false (higher energy) to the true (lower energy) vacuum. This tunneling event corresponds to the appearance of an expanding bubble of the true vacuum embedded in the false. As long as the probability that a bubble nucleates in each horizon volume of the false vacuum during a Hubble time is less than one (so that the background expansion of the false vacuum on average prevents bubbles from merging), the phase transition never completes and inflation is eternal (Guth and Weinberg, 1981) (see Sekino et al. (2010) for a modern treatment of the percolation problem

in eternal inflation). The O(4)-invariance of the instanton guarantees that the bubble interior possesses SO(3,1) symmetry, and therefore contains an infinite open Friedman Robertson Walker (FRW) universe. Although homogeneity is ensured by the symmetries of the instanton, if the interior of a bubble is to resemble our own universe, a second epoch of inflation inside the bubble is necessary to dilute the negative curvature and provide the correct spectrum of primordial density perturbations to seed structure. Models of this type are known in the literature as open inflation, and have been explored in detail (see Garcia-Bellido (1997)).

The signatures of single-bubble open inflation include negative curvature and modifications to the power spectrum. These modifications are most important at large angular scales (see Garcia-Bellido (1997) and references therein) where cosmic variance is dominant, and would be very difficult to detect. Since curvature alone would not be a very distinguishing prediction, we do not consider these signals further.

A less ambiguous signature of eternal inflation would be the visible remnants of collisions between bubbles. Although the bubbles formed during eternal inflation do not percolate, there are many (in fact, an infinite number of) collisions. These collisions lead to inhomogeneities in the inner-bubble cosmology, perhaps leaving observable signatures in the CMB. Assessing the observational consequences of bubble collisions in an eternally inflating universe has been an active area of research (Garriga et al., 2007; Aguirre et al., 2007; Aguirre and Johnson, 2008; Aguirre et al., 2009; Chang et al., 2008, 2009; Dahlen, 2010; Freivogel et al., 2009; Easther et al., 2009; Larjo and Levi, 2010; Zhang and Piao, 2010; Czech et al., 2010) (for a review, see Aguirre and Johnson (2009)). These studies have established that a number of criteria necessary for the observation of bubble collisions (Aguirre et al., 2007) can be satisfied, at least in some models:

Compatibility: In order to satisfy this criterion, there must be a bubble we can collide with that only minimally disturbs the homogeneity of the observable portion of the surface of last scattering. Such collisions do seem to exist, as evidenced by thin-wall junction condition analysis (Aguirre and Johnson, 2008; Chang et al., 2008) as well as numerical simulations (Aguirre et al., 2009) and a study of the inflaton field in a background thin-wall collision geometry (Chang et al., 2009).

Probability: We should expect to have a collision in our causal past. The number of collisions in our past is $\bar{N} = \lambda V_4^F$, where λ is the bubble nucleation probability per unit four-volume and V_4^F is the four-volume outside the bubble to which we have causal access. The expected number (assuming the original open FRW foliation) is formally infinite (Aguirre et al., 2007); however, collisions that contribute to this divergence only produce very long wavelength fluctuations at last scattering, and so would not be observable (Aguirre et al., 2009; Freivogel

et al., 2009) (this is similar to the infrared divergence found in models of slow-roll inflation). The average number of collisions that affect the *observable* portion of the surface of last scattering is finite (Freivogel et al., 2009; Aguirre and Johnson, 2009), and is given by

$$\bar{N} \simeq \frac{16\pi}{3} \lambda H_F^{-4} \left(\frac{H_F}{H_I} \right)^2 \Omega_k^{1/2} \,, \tag{3.1}$$

where H_F is the false vacuum Hubble constant, H_I is the inflationary Hubble constant inside our bubble, and Ω_k is the current component of energy in curvature. For the expected number of observable collisions to be one or larger, the separation of scales between H_F and H_I must be large enough to compensate for the low probability λ (which is exponentially suppressed because this is a tunneling process) and the observational constraint $\Omega_k \lesssim .0084$ (Komatsu et al., 2011). Given a particular scalar potential underlying eternal inflation, $\bar{N}$ for each possible type of collision is fixed. However, in a theory with a complicated potential landscape for the scalar field(s), it makes sense to think of $\bar{N}$ as a continuous parameter with some prior probability distribution[1]. Without detailed knowledge of the theory underlying eternal inflation and an associated measure, it is difficult to assess how likely it is to have $\bar{N} > 1$, but see Freivogel et al. (2009) and Aguirre and Johnson (2009) for some speculative comments. There is also an exponential pressure from the nucleation rates towards $\bar{N} \gg 1$ or $\bar{N} \ll 1$. In the following, we assume $\bar{N}$ can be order one.

Observability: Since the effects of a collision must pass through the entire inner-bubble cosmology, they can be thought of as perturbations of the Big Bang in an FRW cosmology. As such, they are stretched by inflation, and we expect the strength of most signatures to scale with (some power of) Ω_k. We therefore must require that there are not too many more e-folds than required to satisfy the observational bound on curvature. Given a field theory model, the number of e-folds of inflation inside the bubble is uniquely determined by the instanton. However, if we consider a landscape of scalar potentials, then it is necessary to find a measure over the number of inflationary e-folds (or equivalently Ω_k). For some work in this direction, see e.g. Freivogel et al. (2006) and De Simone and Salem (2010).

Much remains to be learned about the full spectrum of possible outcomes of a bubble collision and the exact details of the associated observational signatures. Nevertheless, all potentially observable bubble collisions involving two bubbles share a sufficiently general set of properties to allow for a meaningful observational search even in the absence of a detailed model. In summary, we expect all such observable bubble collisions to possess:

---

[1]We return to this point in Sec. 3.6.3 when discussing the Bayesian framework for testing bubble collision models.

- Azimuthal symmetry: A collision leaves an imprint on the CMB sky that has azimuthal symmetry. This is a consequence of the SO(2,1) symmetry of the spacetime describing the collision of two vacuum bubbles (Chao, 1983; Garriga et al., 2007; Aguirre et al., 2007).

- A causal boundary: The surface of last scattering can only be affected inside the future light cone of a collision event. The intersection of our past light cone, the future light cone of a collision, and the surface of last scattering is a ring. This is the causal boundary of the collision on the CMB sky. The temperature and its derivatives need not be continuous across this boundary. Neglecting the backreaction of the collision on the geometry of the bubble interior, the distribution of ring sizes was found in Freivogel et al. (2009) to be

$$\frac{d\bar{N}}{d\theta_{\text{crit}}} \sim 4\pi\lambda H_F^{-4} \left(\frac{H_F}{H_I}\right)^2 \Omega_k^{1/2} \sin\left(\theta_{\text{crit}}\right), \tag{3.2}$$

where θ_{crit} is the angular radius measured from the center of the disc to the causal boundary and the other quantities are as defined in Eq. 3.1.

The form of Eq. 3.2 can be simply broken down into its constituent parts. The $H_F^{-4}\left(\frac{H_F}{H_I}\right)^2$ factor describes the volume in the inflating bulk – a shell around the observation bubble – in which future bubble nucleations will result in collisions; multiplying by the rate of nucleations per unit four-volume, λ, gives the expected number of collisions. For the collisions to then be *observable*, we then require two criteria to be satisfied. Firstly, there must not be too much slow-roll inflation, post-bubble-nucleation, that the effects of the collision are completely smoothed away: the number of observable collisions therefore increases with the energy density in curvature, Ω_k. Secondly, the collision must intersect the observer's last-scattering surface: a simple geometric calculation (Freivogel et al., 2009) shows that the majority of such collisions are half-sky in scale (collisions covering all or none of the last-scattering surface are by definition not detectable), and that the number depends on $\sin\theta_{\text{crit}}$.

- An overall modulation of the background fluctuations: We assume that the temperature fluctuations, including the effects of the collision, at a location on the sky $\hat{\mathbf{n}}$ can be written as (Gordon et al., 2005; Chang et al., 2009)

$$\frac{\delta T(\hat{\mathbf{n}})}{T_0} = (1 + f(\hat{\mathbf{n}}))(1 + \delta(\hat{\mathbf{n}})) - 1, \tag{3.3}$$

where $f(\hat{\mathbf{n}})$ is the modulation induced by the collision and $\delta(\hat{\mathbf{n}})$ are the temperature fluctuations induced by modes set down during inflation. This is motivated by the observation

that the main effect of the bubble collision is to slightly advance or retard the inflaton inside our bubble. The modulation is multiplicative under the assumption that the normal inflationary density fluctuations simply "paint" the perturbed surface of last scattering and have identical statistical properties in both the regions affected and unaffected by the collision.

- Long-wavelength modulation: A collision is a pre-inflationary relic. The effects of a collision inside the causal boundary are stretched by inflation, and so we can expect that the relevant fluctuations are large-scale. As we describe below, this implies that the temperature modulation due to a collision centered on the north pole ($\theta = 0$) has the form

$$f(\hat{\mathbf{n}}) = (c_0 + c_1 \cos \theta + \mathcal{O}(\cos^2 \theta))\Theta(\theta_{\mathrm{crit}} - \theta) , \tag{3.4}$$

where the c_i are constants related to the properties of the collision, θ is the angle measured from the center of the affected disc, and $\Theta(\theta_{\mathrm{crit}} - \theta)$ is a step function at the causal boundary θ_{crit}. Truncating the sum at $\mathcal{O}(\cos \theta)$, the constants c_0 and c_1 can be expressed in terms of a central amplitude z_0 and edge discontinuity z_{crit}:

$$c_0 = \frac{z_{\mathrm{crit}} - z_0 \cos \theta_{\mathrm{crit}}}{1 - \cos \theta_{\mathrm{crit}}}, \quad c_1 = \frac{z_0 - z_{\mathrm{crit}}}{1 - \cos \theta_{\mathrm{crit}}} , \tag{3.5}$$

as shown in Fig. 3.1. Allowing the collision to be centered on an arbitrary location $\{\theta_0, \phi_0\}$ on the celestial sphere, the induced temperature modulation can be expressed as a function of five parameters: $\{z_0, z_{\mathrm{crit}}, \theta_{\mathrm{crit}}, \theta_0, \phi_0\}$. A modulation of this form was first derived in Chang et al. (2009), where it was obtained from the observed modulation of a field representing the inflaton inside our bubble, numerically evolved in a background thin-wall bubble collision geometry. These authors did not predict the existence of a temperature discontinuity z_{crit}. While further work is needed to better predict the precise form of the template, in our analysis we allow bubble collisions to produce modulations with and without discontinuities.

In Fig. 3.2, we show the Poincare-disc representation[2] (see Aguirre et al. (2007) for the details of this construction) of the surface of last scattering inside our parent bubble. The collision affects the shaded portion of this surface. The observed CMB is formed at the intersection of our past light cone (dashed circle) with the surface of last scattering, which in this case includes regions both affected and unaffected by the collision. From the underlying azimuthal symmetry,

---

[2] A Poincare disc is a mapping of a hyperbolic geometry – i.e. a negatively curved space, such as a bubble interior – onto a unit disc. Points at infinity are mapped onto the boundary of the disc; geodesics are circular arcs (or diameters) that intersect the disc boundary at right angles.

Figure 3.1: The radial temperature modulation Eq. 3.4 induced by a bubble collision centered on the the north pole ($\theta = 0$).

the collision appears as a disc on the observer's CMB sky. Zooming in on the neighborhood of our past light cone (inset), we can treat the universe as being flat. In addition, because we have causal access to much less than one curvature radius at last scattering (again from the observational bound on Ω_k), the collision has an approximate planar symmetry.

The collision introduces pre-inflationary inhomogeneities into our bubble. The exact nature of these inhomogeneities depends on the specific model underlying the formation of our bubble and the subsequent epoch of slow-roll inflation, as well as the specifics of the collision. In dramatic cases, the collision ends slow-roll inflation everywhere within its future light cone (Aguirre et al., 2009), induces the transition to another vacuum state (Easther et al., 2009; Giblin et al., 2010; Johnson and Yang, 2010), or produces a post-collision domain wall that eats into our bubble interior (Aguirre and Johnson, 2008; Chang et al., 2008). These scenarios are obviously in conflict with observation, and we do not consider them further. In mild cases, which will be our focus in the remainder of this paper, collisions satisfy the "compatibility" criterion defined above: the observable portion of the surface of last scattering is only minimally disturbed by the collision. Thin-wall analysis (Aguirre and Johnson, 2008) and numerical simulations (Aguirre et al., 2009; Chang et al., 2009) indicate that it is indeed possible to find situations where the effects of a collision are compatible with our observed cosmology.

The disturbance caused by a collision is a pre-inflationary relic and thus is stretched by the period of inflation inside the bubble. From the current bound on curvature (Komatsu et al., 2011), we can infer that our past light cone encompasses less than one horizon volume at the onset of inflation. This implies that the initial disturbances caused by a collision, which is smeared out on the scale of the inflationary horizon after a few e-folds of inflation, has a wavelength today that is larger than the current horizon size. Together with the planar symmetry of the collision

Figure 3.2: A Poincare-disc representation of the surface of last scattering inside our parent bubble, with one dimension suppressed. The future light cone of the collision at this time is denoted by the dark red line, with the shaded region representing the portions of the surface of last scattering that are to the future of the collision. Our past light cone at last scattering is represented by the dashed circle. From the present bounds on curvature, the size of our past light cone must be much smaller than one curvature radius. Zooming in on the portion of the surface of last scattering that we have causal access to (inset), the universe is very close to flat, and the region affected by the collision has approximate planar symmetry. The region affected by the collision appears as a disc of angular radius θ_{crit} on the CMB sky.

at last-scattering (by convention along the y-z plane), this implies that we can Taylor-expand the Newtonian potential (see Czech et al. (2010) for a translation between the Newtonian potential and the originally postulated temperature modulation presented in Chang et al. (2009)) about the causal boundary of the collision at $x = x_{\mathrm{crit}}$ as

$$\Phi_{\mathrm{coll}} = \Phi(a) \left(\bar{c}_0 + \bar{c}_1(x - x_{\mathrm{crit}}) + \mathcal{O}((x - x_{\mathrm{crit}})^2) \right) \Theta(x - x_{\mathrm{crit}}), \tag{3.6}$$

where $\Phi(a)$ encodes the evolution of the potential with scale factor a and the $\bar{c}_i$ are model-dependent constants. [3]

There are contributions to the observed temperature modulation from the Sachs-Wolfe effect, the integrated Sachs-Wolfe effect, and a Doppler effect (coming from the induced bulk peculiar velocity $\mathbf{v}$ of the fluid in the region affected by the collision):

$$\frac{\delta T}{T} \simeq \frac{\Phi_{\mathrm{coll}}(a_{\mathrm{ls}})}{3} + 2 \int_{a_{\mathrm{ls}}}^{1} da \frac{d\Phi_{\mathrm{coll}}}{da} + \left(\mathbf{v} \cdot \hat{\mathbf{n}} + \mathcal{O}(v^2) \right), \tag{3.7}$$

where a_{ls} is the scale factor at last scattering, $a = 1$ today, and

$$\mathbf{v} \propto \nabla \Phi_{\mathrm{coll}} + a \frac{d}{da} \nabla \Phi_{\mathrm{coll}}. \tag{3.8}$$

To leading order in $\mathbf{v}$, the temperature induced by the collision is linear in Φ_{coll} and its derivatives. Therefore, since $x = x_{\mathrm{ls}} \cos\theta$ (where x_{ls} is the comoving distance out to which we can see on the surface of last scattering), the temperature fluctuations induced by a collision are generally of the form Eq. 3.4. Further, even if the Newtonian potential is continuous across $x = x_{\mathrm{crit}}$, the resulting temperature fluctuations need not be continuous across the causal boundary at θ_{crit}. This discontinuity arises from the ISW and Doppler contributions to the observed temperature fluctuation. Effects that we have neglected, including the finite thickness of the surface of last scattering and uncertainties about how the perturbations caused by a bubble collision propagate through our bubble interior, are encapsulated by the higher order terms in Eq. 3.4. These effects could smear out the causal boundary enclosing the collision on sub-degree scales. These corrections could be incorporated into our analysis as theoretical understanding improves.

Given a specific model for the scalar fields making up the bubbles and driving eternal inflation, the kinematics of a particular collision, and our position inside our bubble, it is in principle

---

[3] We are modeling the collision as a collection of super-modes truncated at the causal boundary, and our treatment is therefore very similar to the so-called "tilted universe" scenario (Turner, 1991; Erickcek et al., 2008). The important distinction in the case of bubble collisions is that the perturbation vanishes at the causal boundary x_{crit}. Because the collision entered our past light cone only relatively recently, we are still comoving with respect to the undisturbed FRW foliation, and the cancellation of the dipolar temperature modulation seen in Turner (1991), Erickcek et al. (2008) and Zibin and Scott (2008) does not occur.

Figure 3.3: On the left, we show a bubble collision template with $\{z_0 = 5.0 \times 10^{-5}, z_{\mathrm{crit}} = -5.0 \times 10^{-5}, \theta_{\mathrm{crit}} = 10.0°, \theta_0 = 57.7°, \phi_0 = 99.2°\}$. On the right we add simulated background fluctuations, smoothing, and instrumental noise.

possible to determine the free parameters in Eq. 3.4. Treating the colliding bubbles in the thin-wall approximation, some measure of the strength of a collision can be specified in terms of the vacuum energies in the bubbles, wall tensions, and kinematics as in Aguirre and Johnson (2008). The kinematics will induce a probability distribution for the free parameters in Eq. 3.4. However, an accurate treatment requires a calculation of the back-reaction of the collision on the behaviour of the inflaton inside our bubble. Preliminary work in this direction has been done (Chang et al., 2009; Aguirre et al., 2009; Easther et al., 2009; Giblin et al., 2010), providing a handful of examples. However, a systematic investigation has not yet been performed. This is distinct from the case where an ensemble of field theory models is considered, representing the string theory landscape. In this case, the fundamental parameters governing the structure of the colliding bubbles (wall tensions and vacuum energies) and the properties of the inner-bubble cosmology (including the number of inflationary e-folds etc) are drawn from some probability distribution. This again will induce a probability distribution for the free parameters in Eq. 3.4, whose nature is presently poorly understood.

What would a bubble collision embedded in a CMB temperature map look like? In Fig. 3.3 we show a large-amplitude collision with and without background CMB fluctuations. In the following sections, we apply the various stages of our analysis pipeline to this example to illustrate the algorithm. We make extensive use of such simulations in calibrating our analysis pipeline, and the details of their construction are presented in Sec. 3.5. Although there could conceivably be many overlapping collisions, the predicted observational signatures of this scenario have yet to be explored, and we focus on simulations of distinct individual bubble collisions. Again, as theoretical understanding improves, our analysis could be extended to include the possibility of overlapping collisions.

What would the detection, or absence, of a bubble collision tell us about the underlying theory of eternal inflation? To examine what the answer to this question might be, let us make

some further assumptions about the temperature modulations caused by a bubble collision. First, assume that the potential induced by the collision (Eq. 3.6) is composed mostly of a single long-wavelength mode of physical wavenumber k. Second, assume that the Sachs-Wolfe effect is the dominant contribution to the observed temperature modulation. Under these assumptions, the amplitude of an observed temperature modulation is:

$$z_0 \simeq \frac{2}{3} \frac{k}{H_0} \Phi(a_{\mathrm{ls}}) \left(1 - \cos\theta_{\mathrm{crit}}\right), \tag{3.9}$$

where a_{ls} is the scale factor at last scattering. If the initial wavelength of the disturbance was of order one inflationary Hubble length $k \sim H_I$ (since any fine-structure in the collision would be smeared within the first few e-folds of inflation), then $\Phi(a_{\mathrm{ls}}) = \Phi(a = 0)$, and the physical size of such a mode at last scattering is given by

$$k \simeq \Omega_k^{1/2} H_0. \tag{3.10}$$

In this case, we have

$$z_0 \simeq \Phi(0)\Omega_k^{1/2} \left(1 - \cos\theta_{\mathrm{crit}}\right). \tag{3.11}$$

If a bubble collision is detected, and a similar set of assumptions is proven correct in a specific model, the measured values of z_0 and θ_{crit} allow one to infer the value of Ω_k. [4] In the absence of a detected collision, Eq. 3.11 can be turned into a bound on a combination of $\Omega_k^{1/2}$ and $\Phi(0)$:

$$\Omega_k^{1/2}\Phi(0) < \left[z_0/\left(1 - \cos\theta_{\mathrm{crit}}\right)\right]_{\mathrm{observational\ upper\ bound}} \tag{3.12}$$

This analysis can be recognized as an example of the Grishchuk-Zel'dovich effect in an open universe (Garcia-Bellido et al., 1995; Turner, 1991).

Determining the detailed properties of the theory underlying eternal inflation through the observation of bubble collisions is likely to be a messy business. However, any model will predict an expectation value for the number of observable bubble collisions $\bar{N}$, making this a very useful phenomenological parameter. Any constraints on $\bar{N}$ from data will also yield interesting information about our parent vacuum through Eq. 3.1. The most naive application of such a

---

[4]The values of Ω_k that one might be able to infer are conceivably below both the observational bound $\Omega_k \leq .0084$ and the theoretical observational bound $\Omega_k \leq 10^{-5}$. For example, assuming $z_0 \sim 10^{-5}$ and $\Phi(0) \sim 1$ (since the collision involves a relatively large release of energy), if a collision were observed at large angular scale (where $\cos\theta_{\mathrm{crit}} \sim 0$), we can infer that $\Omega_k \sim 10^{-10}$. This implies that a collision is in principle observable even when curvature is not. We thank Lam Hui for elucidating this point.

constraint, where we have evidence that $\bar{N} > 1$ or $\bar{N} < 1$, would yield the inequalities

$$\lambda H_F^{-2} < \frac{3H_I^2}{16\pi\Omega_k^{1/2}} \text{ (no detected collisions)}, \quad \lambda H_F^{-2} > \frac{3H_I^2}{16\pi\Omega_k^{1/2}} \text{ (collision detection)}, \quad (3.13)$$

These bounds would be most useful if we detect Ω_k and/or B-mode polarization (the amplitude of which can be related to H_I) in future data. In the most optimistic scenario [5], if primordial B-mode polarization is detected by the *Planck* satellite, we can infer that $H_I \sim 10^{11} - 10^{13}$ GeV. Further, if curvature is detected at the level $\Omega_k \sim 10^{-3}$, then in Eq. 3.13, λH_F^{-2} would be bounded from above or below by $\sim 10^{26} \, \text{GeV}^2$. The condition for eternal inflation is $\lambda H_F^{-4} < 1$. Any application of Eq. 3.13 must be consistent with this inequality. For example, assuming a Planckian false vacuum energy ($H_F \sim 10^{19} \, \text{GeV}$), the nucleation probability λH_F^{-4} could be bounded from above or below by $\sim 10^{-4}$, remaining consistent with the condition for eternal inflation.

3.4   Summary of the analysis pipeline

Before providing a detailed description of our analysis pipeline, we motivate and summarize its various components. Eternal inflation can arise from a wide range of inflationary potentials, each producing a different expected number of detectable collisions on the CMB sky, $\bar{N}_s$. We will therefore use $\bar{N}_s$ as a continuous parameter that characterizes particular models of eternal inflation. The standard cosmological model is given by the special case in which $\bar{N}_s = 0$. Our primary goal is to determine, given the WMAP 7-year data, what constraints can be placed on $\bar{N}_s$ and whether models predicting $\bar{N}_s > 0$ should be be preferred over models predicting $\bar{N}_s = 0$.

The optimal approach to achieving this goal would be to construct the full posterior for $\bar{N}_s$ from Bayes' theorem given full-resolution CMB data on the whole sky. Unfortunately, this would require inverting the full-sky full-resolution CMB covariance matrix as well as integrating the bubble-collision likelihood over a many-dimensional parameter space. These tasks are computationally intractable. However, taking advantage of the fact that bubble collisions produce discrete localized effects on the CMB sky, it is possible to approximate the full-sky Bayesian analysis by a patch-wise analysis if the most promising candidate signatures can be identified in advance. The implementation of such an approximation scheme requires two assumptions. First, we assume that the likelihood of models predicting $\bar{N}_s > 0$ is peaked in the regions of the

---

[5]In models where a collision is expected to be in our past, there might be good reason to expect a correlation between observed B-modes and the observation of a bubble collision (Aguirre et al., 2009). This is because large-field models of inflation, which generically predict a larger value for the tensor to scalar ratio, are much more robust in the presence of a bubble collision. In models of small-field inflation, a bubble collision can end inflation everywhere in its future light cone, implying that collisions in such models are not compatible with our observed cosmology.

sky containing the candidate collisions, and that the integral over the likelihood can therefore be estimated by concentrating on these regions, which make the largest contribution to the full integral. Second, we assume that these regions are separated widely enough to be uncorrelated with each other, so much smaller local covariance matrices can be used. These assumptions allow the results of a small number of localized (and therefore computationally-feasible) Bayesian model selection tests to be combined into estimates of the required full-sky statistics. Put simply, our algorithm implements a conservative approximation to the required numerical integral. A complete treatment of the full-sky analysis and the assumptions on which it is formed can be found in Appendix A.1. In addition, once a set of candidates have been identified, it is possible to apply further tests of the data in parallel.

The full-sky approximation necessitates the development of an algorithm that identifies the most promising regions of the CMB sky and then processes them individually. Upon segmenting the full data set, it is important to avoid biasing oneself with *a posteriori* selection effects (Bennett et al., 2011), and it is therefore critical to minimize human intervention in choosing what portions of the sky to analyze. Thus our analysis pipeline is fully automated, tested and calibrated on realistic simulations of the data, and frozen before being applied to the real data. The final pipeline contains no algorithmic choices tunable via human intervention. As discussed in Appendix A.1, missing a bubble collision candidate which makes a significant contribution to the full integral leads to a *conservative* bias towards models predicting $\bar{N}_{\mathrm{s}} = 0$. This alleviates the worry that selection effects might lead to a spurious detection.

Our analysis pipeline consists of a candidate identification step, followed by two parallel verification procedures:

- **Blob detection**: To begin, we attempt to locate the most promising candidate signals using wavelets. Wavelet analysis is a compromise between working purely in position or harmonic space, and therefore yields information both about the location and angular scale of particular features in the temperature map. Specifically, we employ standard (Marinucci et al., 2008; Pietrobon et al., 2006; Pietrobon et al., 2008; Baldi et al., 2006; Guilloux et al., 2007) and Mexican (Scodeller et al., 2011) spherical needlets, two classes of wavelets defined on the sphere. The statistics of the needlet representation of a purely Gaussian CMB temperature map (expected in the absence of a bubble collision at large scales where WMAP is cosmic-variance-dominated), combined with simulations of a bubble-free masked CMB sky, can be used to quantify the significance of various features. A set of significance thresholds are then defined to ensure a manageable number of "false detections" in the end-to-end simulation of the WMAP experiment (see Sec. 3.5.1). Regions of the sky passing

these thresholds are sewn into "blobs," whose size and location is determined by the needlet responses, and passed onto the next stages of the pipeline.

- **Edge detection**: Once a set of candidate signatures is found, we look for circular edges across which the temperature is discontinuous. As discussed above, such causal edges are expected to be a generic signature of bubble collisions. We use the Canny algorithm (Canny, 1986), adding an adaptation of the Circular Hough Transform (CHT) (Kimme et al., 1975), to focus our search on circular edges. The algorithm consists of identifying the most likely centre for a noisy circular edge. The significance of this response is calibrated from a detailed analysis of bubble collision simulations including cosmic variance, spatially-varying WMAP instrumental noise, and smoothing due to the instrumental beam. We verify that this step produces no false detections in the WMAP end-to-end simulation.

- **Bayesian parameter estimation and model selection**: The regions highlighted by the blob detection step can be used to construct an approximation to the full-sky posterior probability distribution for $\bar{N}_{\mathrm{s}}$ using the methods outlined in Appendix A.1. We first perform a pixel-based evaluation of the likelihood and Bayesian evidence in each blob for bubble collision templates of the form given in Eq. 3.4, sampling the parameter space using the nested sampler Multinest (Feroz et al., 2009). The likelihood analysis includes cosmic variance, spatially varying WMAP instrumental noise, and the smoothing due to the instrumental beam. Combining the evidences from each blob we obtain the posterior probability distribution for $\bar{N}_{\mathrm{s}}$, which is used to derive constraints on $\bar{N}_{\mathrm{s}}$ and perform model selection to determine if a theory with $\bar{N}_{\mathrm{s}} \neq 0$ is preferred over a theory with no predicted collisions. The significance of a detection is again calibrated from an analysis of simulated collisions and an end-to-end collision-free simulation of the experiment.

The most important output of our pipeline is the approximation to the full-sky posterior probability distribution for $\bar{N}_{\mathrm{s}}$. This allows us to derive marginalized constraints on $\bar{N}_{\mathrm{s}}$, and perform model selection between theories with $\bar{N}_{\mathrm{s}} = 0$ and $\bar{N}_{\mathrm{s}} \neq 0$. In addition, for each blob identified by the first set of the pipeline, we obtain a set of marginalized posterior constraints on the model parameters $\{z_0, z_{\mathrm{crit}}, \theta_{\mathrm{crit}}, \theta_0, \phi_0\}$, a maximum needlet significance, CHT score, and a local Bayesian evidence ratio with respect to the no-bubble-collision model.

3.5 Simulations

Our analysis pipeline is general, but each step must be calibrated using simulations of the particular data-set under consideration, in this case the WMAP 7-year data release (Jarosik

et al., 2011). WMAP has measured the intensity and polarization of the microwave sky in five frequency bands. The resolution of the instrument in each band is limited by the detectors' beams, and is highest at $0.22°$ in the 94 GHz W band. We perform our analysis on the foreground-subtracted W-band WMAP temperature map, as this combines the highest resolution full-sky data currently available with the least foreground contamination. To minimize the effects of the residual foregrounds we cut the sky with the conservative KQ75 mask, leaving 70.6% of the sky unmasked.

We carry out extensive simulations to quantify the thresholds at which areas of the sky are passed from one step to another. To find the best approximation to the full-sky Bayesian analysis, we process as much of the sky as is computationally feasible.

To determine the response of our pipeline to bubble collisions over the range of possible parameters, we generate simulations containing a variety of bubble collisions plus CMB, realistic noise and Gaussian beam smoothing. However, we also wish to ensure that we have a method to guard against systematic effects (e.g., foreground residuals and any map-making artifacts that may be present) that we do not have capability to simulate. These effects might lead to false detections in the "blob detection" stage, or critically, the edge detection and Bayesian analysis stages. It is impossible to claim a detection without first ensuring that there are no such false detections due to systematics.

3.5.1   WMAP end-to-end simulation

A realistic simulation of a WMAP-quality data-set that does not contain a bubble collision is an important tool for calibrating and quantifying the expected false detection rate of our analysis pipeline when applied to data which may include systematics (such as foreground residuals) that are not captured in our simulations or likelihood function. For this purpose we use a complete end-to-end simulation of the WMAP experiment provided by the WMAP Science Team [6]. 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 is obtained separately from simulated sources including diffuse Galactic foregrounds, CMB fluctuations, realistic noise, smearing from finite integration time, finite beam size, and other instrumental effects. In our analysis, we utilize the foreground-reduced W-band simulation.

---

[6]`http://lambda.gsfc.nasa.gov/product/map/dr4/sim_maps_info.cfm`

3.5.2 Simulated bubble collisions

The temperature fluctuations observed in the CMB, including the effects of a bubble collision (originally found in Gordon et al. (2005) and Chang et al. (2009)), can be written as

$$\delta T(\hat{\mathbf{n}}) = [T_0'(1 + f(\hat{\mathbf{n}}))(1 + \delta(\hat{\mathbf{n}})) - T_0]_{\text{smoothed}} + \delta T_{\text{noise}}(\hat{\mathbf{n}}), \qquad (3.14)$$

where $\hat{\mathbf{n}} = \{\theta, \phi\}$ is the position on the sky[7], T_0 is the average temperature of the map including the modulation, T_0' is the average temperature without the modulation, δT_{noise} is the contribution from instrumental noise, $f(\hat{\mathbf{n}})$ and $\delta(\hat{\mathbf{n}})$ are defined as in Eq. 3.3. The quantities in the brackets are smoothed with a Gaussian beam of $0.22°$ (approximating the beam size of the WMAP experiment in the W band). We use the WMAP best-fit 7-year power spectrum (Larson et al., 2011) in the multipole range $2 \leq \ell \leq 1024$ to generate fluctuation maps $\delta T_{\text{syn}}(\hat{\mathbf{n}}) = T_0'\delta(\hat{\mathbf{n}})$ at the full WMAP resolution of $N_{\text{side}} = 512$ (with 3,145,728 pixels). The noise term δT_{noise} is generated from WMAP 7-year noise variances at the same resolution. Since the templates we consider add a relatively small temperature excess/deficit in one location, the features do not cause the power spectrum to deviate from that measured by WMAP (Chang et al., 2009). Additionally, we can replace $T_0 \approx T_0'$, which gives

$$\delta T(\hat{\mathbf{n}}) = [(1 + f(\hat{\mathbf{n}}))(T_0 + \delta T_{\text{syn}}(\hat{\mathbf{n}})) - T_0]_{\text{smoothed}} + \delta T_{\text{noise}}(\hat{\mathbf{n}}). \qquad (3.15)$$

We consider collisions with $\theta_{\text{crit}} = 5°, 10°, 25°$ and choose centers in a high-noise region ($\theta_0 = 56.6°, \phi_0 = 193.0°$) and a low-noise region ($\theta_0 = 57.7°, \phi_0 = 99.2°$) of the sky that remain significantly outside of the main body of the WMAP KQ75 7-year mask. The regions of the sky affected by $5°$ and $10°$ collisions are over-plotted in Fig. 3.4 on a masked map of the instrumental noise variance. For each θ_{crit} and location, we generate 35 simulated collisions with parameter values logarithmically spaced in the ranges $10^{-6} \leq z_0 \leq 10^{-4}$ and $-10^{-4} \leq z_{\text{crit}} \leq -10^{-6}$. The response of our pipeline depends only on the absolute value of z_0 and z_{crit}, so the choice of sign for z_0 and z_{crit} is arbitrary. We repeat this for three realizations of the background CMB fluctuations, yielding a total of 210 simulated sky maps for each of the three collision sizes.

3.6 Analysis tools

We now describe in detail the analysis tools which make up our pipeline and how they are calibrated with simulations before being applied to the data. Readers wishing to skip these

---

[7]These angular positions can be expressed in terms of Galactic coordinates through longitude $l = \phi$ and latitude $b = 90° - \theta$.

Figure 3.4: The two locations chosen for our simulated bubble collisions are over-plotted on the WMAP 7-year noise variances with the KQ75 7-year mask applied. The regions encompassed by the 5° and 10° simulated collisions are shaded black and grey respectively. Bubbles are centered in an unmasked high-noise $\{\theta_0 = 56.6°, \phi_0 = 193.0°\}$ (left) and low-noise $\{\theta_0 = 57.7°, \phi_0 = 99.2°\}$ region (right).

details may wish to study Figs. 3.11 and 3.17 and turn to Sec. 3.6.4 for a summary of the outputs of the pipeline at each stage, and the conditions under which a detection can be claimed.

3.6.1 Needlets

Wavelet analysis is a powerful tool for identifying features localized on the sky, the type of signal expected for bubble collisions. There exist families of wavelets that are defined on the sphere known as standard (Marinucci et al., 2008; Pietrobon et al., 2006; Pietrobon et al., 2008; Baldi et al., 2006; Guilloux et al., 2007) and Mexican (Scodeller et al., 2011) spherical needlets. These functions form what is known as a "tight frame," allowing for a well-defined forward and reverse needlet transform. As in other forms of wavelet analysis, decomposing the temperature map into a sum over such functions yields information both about the location and angular scale of specific features. For a purely Gaussian temperature field, the statistical properties of the needlet transform can be straightforwardly related to the power spectrum, allowing a quantification of the significance of a possible detection. In addition, the spatial localization properties of the standard and Mexican needlets make it possible to avoid many of the problems associated with working on a cut sky. In this section we outline the properties of needlet transforms and analyze their utility in searching for bubble collisions.

Definition of the spherical needlet transform

The needlet transform is defined as:

$$T(\hat{\mathbf{n}}) = \sum_{j,k} \beta_{jk} \psi_{jk}(\hat{\mathbf{n}}), \qquad (3.16)$$

where $\hat{\mathbf{n}}$ denotes a direction $\{\theta, \phi\}$ on the sky, β_{jk} are constant needlet coefficients, and $\psi_{jk}(\hat{\mathbf{n}})$ are the needlet functions. The members of this family of functions are labeled by the index k of the pixel at which they are centered, and their "frequency" j, which is related to the spatial extent of the needlet profile in real space. The sum in the needlet transform is over all pixels k, and all frequencies $j = 0, 1, 2, \ldots, \infty$. For fixed j, there is one needlet coefficient β_{jk} for each pixel k, allowing us to represent the needlet coefficients at fixed j as a map on the pixelated sky. The needlet functions are defined in terms of spherical harmonics $Y_{\ell m}(\hat{\mathbf{n}})$ as

$$\psi_{jk}(\hat{\mathbf{n}}) = \sqrt{\lambda_{jk}} \sum_{\ell} b(\ell, B, j) \sum_{m=-\ell}^{\ell} Y_{\ell m}^*(\hat{\mathbf{n}}) Y_{\ell m}(\hat{\mathbf{n}}_k). \qquad (3.17)$$

Here, λ_{jk} are the cubature weights, which are related to the area of each pixel. In the equal-area HEALPix pixelization (Górski et al., 2005) we employ, all cubature weights are equal to $\lambda_{jk} = 4\pi/N_{\text{pix}}$, where N_{pix} is the number of pixels, and we absorb this constant into the needlet coefficients β_{jk}. The function $b(\ell, B, j)$ acts as a filter in harmonic space, where B is a constant bandwidth parameter. It is chosen such that the family of functions $\psi_{jk}(\hat{\mathbf{n}})$ form a tight frame (see e.g. Marinucci et al. (2008)), which guarantees the existence of an inverse needlet transform given by

$$\beta_{jk} = \int T(\hat{\mathbf{n}}) \psi_{jk}(\hat{\mathbf{n}}) d\Omega. \qquad (3.18)$$

There are a number of possible choices for the function $b(\ell, B, j)$, which distinguish the standard and Mexican needlets. A description of the explicit form of the function $b(\ell, B, j)$ can be found for standard needlets in Marinucci et al. (2008) and Mexican needlets in Scodeller et al. (2011). We plot b as a function of the multipole moments ℓ in Fig. 3.5. For standard needlets, b only has support for values of ℓ between $B^{j-1} < \ell < B^{j+1}$. The bandwidth parameter B controls the width of each window function in harmonic space. Mexican needlets have support over all ℓ at each frequency j, and again have a bandwidth parameter B which controls the localization properties of the functions in harmonic space.

In Fig. 3.6 we plot the wavelet functions in pixel space. As is to be expected, increasing the width of the function $b(\ell, B, j)$ in harmonic space corresponds to improved localization in pixel space. In the limit of large j, there is an extremely small overlap of the needlet functions at

Figure 3.5: The filter function $b(\ell, B, j)$ for standard (left) and Mexican (right) needlets with $B = 2.5$ for $j = 0, 1, 2, 3$ (solid, dashed, dot-dashed, dotted).

Figure 3.6: Standard needlets in pixel space. On the left, we show standard needlets ψ_{jk} with $B = 2.5$ for $j = 1, 2, 3$ (dot-dashed, dashed, solid) at fixed k as a function of the polar angle θ. On the right, we show standard needlets ψ_{jk} for fixed $j = 3$ at three pixels k (dashed, solid, dot-dashed) as a function of the polar angle θ (note: since we are projecting onto θ, the needlets appear asymmetric).

nearby pixels. The compact support of $b(\ell, B, j)$ in harmonic space for standard needlets leads to slightly poorer localization in pixel space than is enjoyed by the Mexican needlets.

If we decompose the temperature map into spherical harmonics:

$$T(\hat{\mathbf{n}}) = \sum_{\ell, m} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \tag{3.19}$$

then Eq. 3.18, together with the inverse transform

$$a_{\ell m} = \int T(\hat{\mathbf{n}}) Y_{\ell m}^* d\Omega \tag{3.20}$$

leads to:

$$\beta_{jk} = \sqrt{\lambda_{jk}} \sum_{\ell} b(\ell, B, j) \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}_k), \tag{3.21}$$

where $a_{\ell m}$ are the spherical harmonic coefficients. This formula allows us to transform directly from the $a_{\ell m}$s to the spherical needlet coefficients β_{jk}. In our analysis, the needlet transforms

Figure 3.7: The spherical harmonic transform (connected dots) of a $\theta_{\rm crit} = 5°$ (left) and $\theta_{\rm crit} = 25°$ (right) bare collision template centered on the north pole on top of the filter function $b(\ell, B, j)$ for standard needlets with $B = 2.5$ for $j = 0, 1, 2, 3$ (solid, dashed, dot-dashed, dotted). The overlap of the filter function $b(\ell, B, j)$ with the spherical harmonic transform of the bubble template (see Eq. 3.21) determines for which value of j the needlet transform yields the largest signal. In these examples, the 5° collision has the largest needlet response at $j = 3$ and the 25° collision at $j = 2$. The needlet response as a function of angle for the 25° collision is plotted in Fig. 3.8.

are accelerated by generating the $a_{\ell m}$s at full WMAP resolution ($N_{\rm side} = 512$) but limiting the reconstruction multipoles to $2 \leq \ell \leq 256$, and needlet positions k to the pixels at $N_{\rm side} = 128$. This retains the resolution required to reconstruct features from half-sky to half-degree scales, encompassing the range of all detectable collisions.

Needlet response to the bubble collision templates

We now quantify the sensitivity of the needlet transform to the presence of a collision. In the absence of the background Gaussian fluctuations, we perform the needlet transform on a set of bare collision templates. As an illustration, in Fig. 3.7 we plot the spherical harmonic coefficients as a function of ℓ (all coefficients for $m \neq 0$ vanish by symmetry if we center the template on the north pole) for 25° and 5° collision templates, overlaid on top of the rescaled filter function $b(\ell, B, j)$ for spherical needlets. The spherical harmonic coefficients for the collision templates peak at a value of ℓ related to the angular scale of the causal boundary. Therefore, the needlet coefficients are largest at a frequency j that is directly related to the angular scale of the collision. This can be seen in Fig. 3.7, where the 5° collision has a maximum response at $j = 3$ and the 25° collision has a maximum response at $j = 2$.

In Fig. 3.8, we plot the needlet coefficients for the 25° template at a variety of polar angles, for $0 \leq j \leq 3$. The needlet coefficients are largest in the center of the region affected by the collision (here, the template is centered on $\theta = 0$), and at a frequency j correlated with the angular scale of the collision ($j = 2$). As expected, the needlet response is sensitive to both the location and angular scale of the collision.

By studying the needlet response to a variety of bare collision templates given by Eq. 3.4, we

Figure 3.8: $B = 2.5$ for standard needlets and a $25°$ collision. $j = 0$ (circles), $j = 1$ (squares), $j = 2$ (diamonds), and $j = 3$ (triangles).

can find optimal values of the bandwidth parameter B for each needlet type. Larger-bandwidth needlets produce stronger signals, but also respond to a greater range of bubble sizes. We have found that the values $B = 1.8$ and $B = 2.5$ for the standard and $B = 1.4$ for the Mexican needlets are a good compromise between signal strength and angular localization of response. In our analysis, we use this suite of three needlet transforms, ensuring that we are sensitive to temperature modulations with a variety of profiles, and allowing us to cross-check any candidate signals.

As an important step in our analysis pipeline, we build lookup tables containing the possible range of bubble collision scales, $\theta_{\mathrm{crit,min}} \leq \theta_{\mathrm{crit}} \leq \theta_{\mathrm{crit,max}}$, to which each needlet type and frequency is sensitive. We first generate a set of 100 templates at each integer θ_{crit} between $1°$ and $89°$ by randomly sampling z_{crit} between $-5 \times 10^{-5} \leq z_{\mathrm{crit}} \leq 5 \times 10^{-5}$ with $z_0 = 5 \times 10^{-5}$ for $z_{\mathrm{crit}} > 0$, and $z_0 = z_{\mathrm{crit}} + 5 \times 10^{-5}$ for $z_{\mathrm{crit}} < 0$. This creates a set of templates with uniform total amplitude (i.e., constant $f(\theta, \phi) - f(\theta_{\mathrm{crit}}, \phi)$) but a variety of profiles at each angular scale. Next, we calculate the needlet coefficients for each of the three members of our needlet suite, and record the frequency generating the maximum central needlet response for each template. The range in θ_{crit} recorded at each frequency is used to generate the lookup tables in Table 3.1.

Needlet coefficients on a cut sky

The CMB is completely dominated by foreground emission in the region of the Galactic plane, and is also affected by bright point-sources. These issues are typically handled by applying a mask which covers the Galactic plane and known point-sources. The needlet transform can be applied directly on the masked temperature maps, and because the needlet functions are localized in pixel space, needlet coefficients far from the mask for sufficiently high frequency j are not significantly affected. These high-frequency needlets are mainly composed of high-ℓ spherical harmonics, and so cut-sky $a_{\ell m}$s can safely be used to calculate the needlet coefficients through Eq. 3.21. Unfortunately, the low-frequency needlets are quite sensitive to the presence

j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$	j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$	j	$\theta_{\text{crit,min}}$	$\theta_{\text{crit,max}}$
0	60°	90°	1	56°	90°	0	86°	90°
1	33°	71°	2	28°	64°	1	78°	90°
2	12°	32°	3	17°	38°	2	55°	90°
3	5°	14°	4	10°	21°	3	36°	71°
4	2°	5°	5	6°	12°	4	27°	48°
5	1°	2°	6	3°	7°	5	19°	37°
			7	2°	4°	6	14°	27°
			8	1°	2°	7	10°	20°
						8	8°	16°
						9	5°	12°
						10	4°	8°
						11	3°	6°
						12	2°	4°
						13	1°	2°

Table 3.1: Angular scale lookup tables for standard needlets with $B = 2.5$ (left), standard needlets with $B = 1.8$ (center), and Mexican needlets with $B = 1.4$ (right). For a needlet frequency j, the needlet transform is sensitive to bubble collisions on scales $\theta_{\text{crit,min}} \leq \theta_{\text{crit}} \leq \theta_{\text{crit,max}}$. No results are shown for the standard needlets with $B = 1.8$, $j = 0$ as they have no support over the range of angular scales considered.

of the mask. To partially mitigate this sensitivity, we calculate the optimal unbiased maximum-likelihood estimators of the $a_{\ell m}$s (de Oliveira-Costa and Tegmark, 2006) at low ℓ.

Such maximum-likelihood estimators are computationally very expensive, and we must balance accuracy against limited computational resources. Another minor complication arises from the smoothing that is necessary to band-limit the data when performing the maximum-likelihood reconstruction algorithm of de Oliveira-Costa and Tegmark (2006): information leaks from inside the mask. Comparing the reconstruction on masked and unmasked simulated temperature maps using 10°-FWHM Gaussian smoothing, we have determined that a reasonably small bias is obtained when maximum-likelihood $a_{\ell m}$s are used for $\ell < 10$.

This set of hybrid $a_{\ell m}$s – maximum-likelihood reconstructed $a_{\ell m}$s for $\ell \leq 10$ and cut-sky $a_{\ell m}$s for $\ell > 10$ – is used in Eq. 3.21 to calculate the needlet coefficients in the analysis that follows.

Statistical properties of needlet coefficients

For a Gaussian CMB without sky cuts, the statistical properties of the spherical harmonic coefficients (Marinucci et al., 2008) are

$$\langle a_{\ell m} \rangle = 0, \quad \langle |a_{\ell m}|^2 \rangle = C_\ell. \tag{3.22}$$

These are related to the statistical properties of the β_{jk} in a straightforward way by

$$\langle \beta_{jk} \rangle = 0, \quad \langle \beta_{jk}^2 \rangle = \sum_{\ell} b(\ell, B, j) \frac{2\ell + 1}{4\pi} C_{\ell}, \tag{3.23}$$

which are identical at each pixel k. Thus, in a full-sky analysis, comparison with the Gaussian variance yields a measure of how likely it is to find a particular needlet coefficient in a purely Gaussian realization of the CMB sky.

In the presence of foregrounds, however, it is necessary to work on a cut sky, which introduces a j- and k-dependent bias. Following Pietrobon et al. (2008), we determine the significance of a needlet coefficient on the cut sky by [8]

$$S_{jk} = \frac{|\beta_{jk} - \langle \beta_{jk} \rangle_{\text{gauss,cut}}|}{\sqrt{\langle \beta_{jk}^2 \rangle_{\text{gauss,cut}}}}, \tag{3.25}$$

where the average, $\langle \beta_{jk} \rangle_{\text{gauss,cut}}$, and variance, $\langle \beta_{jk}^2 \rangle_{\text{gauss,cut}}$, are calculated at each pixel from the needlet coefficients of 3000 collision-free Gaussian CMB realisations with the WMAP 7-year KQ75 sky cut applied. Simulating only cosmic variance is sufficient here because the measurements made by WMAP are cosmic-variance-limited on the scales of interest, $\theta_{\text{crit}} \gtrsim 5°$.

Maps of the needlet variances obtained from simulations are shown in Fig. 3.9 for an example with low (top) and high needlet (bottom) frequency. On the left are the needlet variances calculated without a mask, which agree at the 5% level with the expected variances from Eq. 3.23. On the right are the masked needlet variances, which are clearly biased within a certain distance from the mask in both cases. Variances in the low-frequency example are affected predominantly by the Galactic cut. Variances in the high-frequency example are affected in a much smaller region of the Galactic cut (reflecting the increased spatial localization of needlets at high frequencies), but are much more significantly affected by the point source masks.

To illustrate the expected response to a bubble collision, in Fig. 3.10 we show the temperature map of our illustrative example of a simulated bubble collision on the cut sky, and the significances of its needlet coefficients calculated from Eq. 3.25 at $j = 3$ using standard needlets with $B = 2.5$. The location of the collision is clearly highlighted in the map of needlet coefficients. The significance of the needlet coefficients in pixels in the center of the collision form a global maximum on the entire map.

---

[8]One can also define composite significances, involving needlet coefficients at multiple frequencies. An example is

$$S_{jj'} = \frac{|\beta_{jk}\beta_{j'k} - \langle \beta_{jk}\beta_{j'k} \rangle|}{\sqrt{\langle (\beta_{jk}\beta_{j'k})^2 \rangle}}. \tag{3.24}$$

We have evaluated this statistic on a variety of collision templates modulating Gaussian realizations of the background CMB fluctuations, and found that it returns about half the significance given by Eq. 3.25.
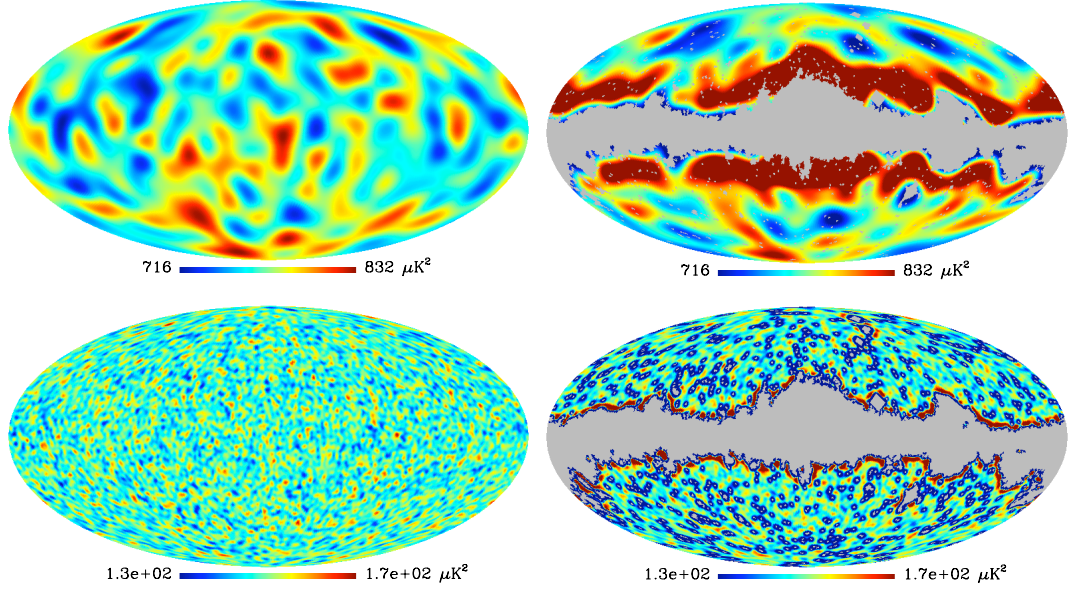
Figure 3.9: Needlet-coefficient variance maps for standard needlets with $B = 2.5$, $j = 2$ (top row) and Mexican needlets with $B = 1.4$, $j = 11$, generated using an ensemble of 3000 Gaussian CMB realizations. The map-averaged variances of the full-sky maps (left column) agree well with expectations from theory (Eq. 3.23 predicts 774 μK^2 and 150 μK^2 respectively), as do regions of sky sufficiently distant from the 7-year KQ75 mask (right column). The low-frequency needlets are affected predominantly by the Galactic cut, whereas the high-frequency needlets are affected by point-source masking.

In order to identify a set of needlet coefficients with a particular feature, we sew regions with 5 or more pixels whose needlet coefficients exceed a frequency-dependent threshold into "blobs" (we discuss in more detail below how these thresholds are set). The core of each blob contains all adjacent pixels that pass the significance threshold. This core is then extended by first finding the average position $\hat{\mathbf{n}}_0$ of the pixels in a blob and modeling it as a disc of radius $\Delta\theta/2$ (where $\Delta\theta$ is the maximum separation between any two pixels in the blob) centred on $\hat{\mathbf{n}}_0$. The blob is then extended to a radius of $1.1 \times (\theta_{\mathrm{crit,max}} + \Delta\theta/2)$, where $\theta_{\mathrm{crit,max}}$ is found from Table 3.1 (which is dependent upon the needlet type and frequency at which a feature is found) to ensure we include all related pixels. All pixels not contained in a blob are masked, and this new temperature map is passed to subsequent steps in the analysis pipeline. Eliminating irrelevant pixels allows us to drastically reduce the computational effort needed in the subsequent analysis.

Analysis of the WMAP end-to-end simulation

As there are many independent needlet coefficients over the sky, it is inevitable that highly significant features will be detected in even a purely Gaussian CMB temperature map. In addition, residual foregrounds and instrumental artifacts may give rise to features which are mis-classified as candidate bubble collisions by the pipeline. To understand how these effects
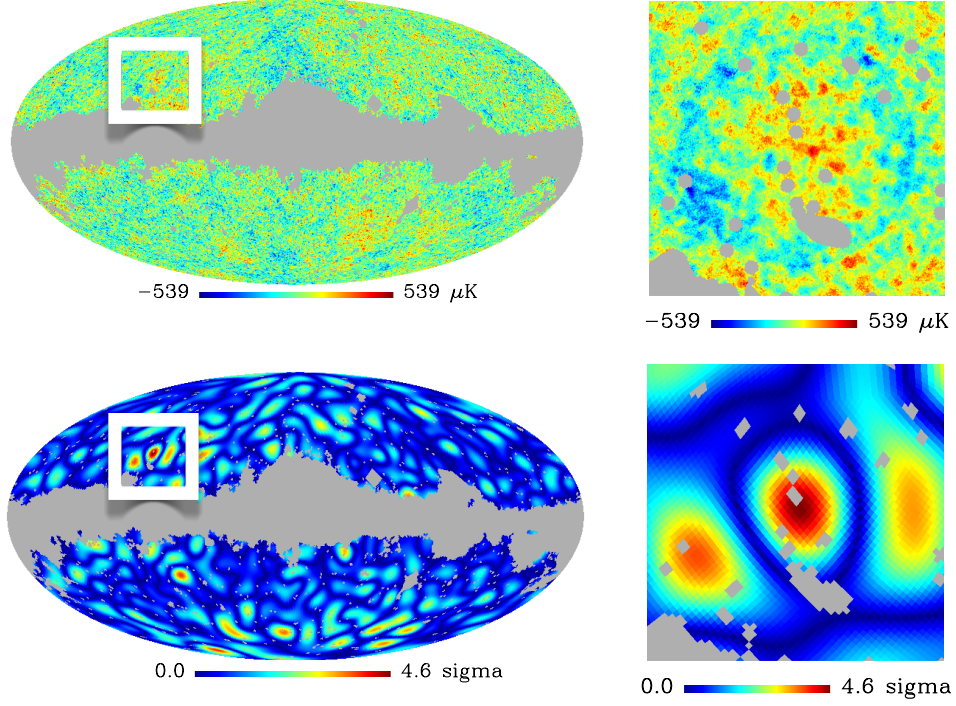
Figure 3.10: The temperature (top) and needlet coefficient significance (bottom) maps for a simulated bubble collision with $z_0 = 5 \times 10^{-5}$, $z_{\mathrm{crit}} = -5 \times 10^{-5}$, $\theta_{\mathrm{crit}} = 10°$ on the CMB sky with the WMAP 7-year KQ75 mask applied. We show the map of needlet coefficients which gives the largest needlet response, in this case $j = 3$ for standard needlets with $B = 2.5$. The right-hand panels provide close-ups of the collision region.

might contribute to the needlet significances, we ran the suite of needlet transforms on the end-to-end simulation of the WMAP experiment (described in Section 3.5) with the 7-year KQ75 mask applied. As an illustration of our results, in Table 3.2 we give the number of blobs of varying significance found in the masked end-to-end simulation using standard needlets with $B = 2.5$. At increasingly high frequency, for which there are more independent needlet coefficients, more and more blobs are found that pass relatively large significance thresholds.

We use the results of Table 3.2 (and similar tables for other members of the needlet suite) to define a set of needlet frequency-dependent significance thresholds that allow a manageable number of false-positives, while retaining sensitivity to a fairly large range of collision template parameters. The significance thresholds we use in our analysis are shown in Table 3.3. There are a total of 17 blobs in the masked end-to-end simulation that pass these thresholds. Comparing their locations on the sky, we can associate these blobs with 13 features (if a feature is picked up by multiple needlet types or frequencies, it can have multiple blobs associated with it). For three of these features, the set of pixels that pass the needlet threshold intersect the Galactic cut. We associate these with a response to the mask, and do not consider them further. For the

j	$S = 3$	3.25	3.5	3.75	4
0	0	0	0	0	0
1	0	0	0	0	0
2	0	0	0	0	0
3	10	4	2	1	0
4	23	10	4	0	0

Table 3.2: The number of blobs found in the masked WMAP end-to-end simulation above significances ranging from $3 \leq S \leq 4$ for standard needlets with $B = 2.5$.

j	S_{min}	N_{blobs}	j	S_{min}	N_{blobs}	j	S_{min}	N_{blobs}
0	3	0	1	3	0	0	3	0
1	3	0	2	3	0	1	3	0
2	3.5	0	3	3.25	0	2	3	0
3	3.5	1	4	3.25	1	3	3	0
4	3.75	0	5	3.25	3	4	3	0
			6	3.5	1	5	3	0
			7	3.5	3	6	3.5	0
						7	3.5	0
						8	3.5	0
						9	3.5	0
						10	3.5	1
						11	3.5	1
						12	3.75	1

Table 3.3: Sensitivity thresholds S_{min} and the number of significant blobs detected in the end-to-end simulations N_{blobs} for standard needlets with $B = 2.5$ (left), standard needlets with $B = 1.8$ (center), and Mexican needlets with $B = 1.4$ (right). Only blobs that do not intersect the galactic cut are reported. No results are shown for the standard needlets with $B = 1.8$, $j = 0$ as they have no support over the range of angular scales considered.

other ten features, the needlet type and frequency which yielded the maximum significance is recorded in Table 3.4.

Analysis of bubble collision simulations

To assess how robustly the needlet transform can pick out a collision in the temperature map, we have performed an analysis of the simulated collisions described in Section 3.5. If the later steps of the analysis are to have a chance at detecting a simulated collision, it must be contained within the set of pixels defined by the blob. To determine if the needlet analysis has detected a bubble collision, we therefore require that a blob exists which fully contains the region affected by the collision, and that the true center of the collision lies within the set of pixels passing the significance threshold.

The results of this analysis for the 5° and 10° collisions are shown in Fig. 3.11. We define the "exclusion region" of these plots as the part of parameter space for which all six realiza-

feature	blob	B	j	S
1	1	2.5	3	3.83
1	2	1.8	5	3.55
2	1	1.8	4	3.99
3	1	1.8	5	3.28
4	1	1.8	5	3.33
4	2	1.4	10	3.77
5	1	1.8	6	3.96
6	1	1.8	7	4.13
7	1	1.8	7	3.97
8	1	1.8	7	4.34
9	1	1.4	11	3.71
10	1	1.4	12	4.14

Table 3.4: Blobs found by the needlet transform in the WMAP end-to-end simulations with the 7-year KQ75 mask.

tions/locations yield a detection. If there were a bubble collision in the WMAP 7-year data with these parameters, it would be detected with high significance regardless of its location on the sky (as long as the collision was not significantly masked). The "sensitivity region" is defined as the part of parameter space for which *any* of the six realizations/locations yields a detection. A bubble collision in this range of parameter space would be detected only for a favorable location or realization of the background fluctuations. The exclusion and sensitivity regions for the 25° collisions are identical to those for the 10° collisions.

Looking at the simulations in detail, there are a few general trends. First, for the needlets to pick out a collision, it is sufficient to have either a relatively large central amplitude z_0 or a relatively large temperature discontinuity z_{crit} at the causal edge. This is clear from the shape of the exclusion region in Fig. 3.11. From the size of the sensitivity region in this plot, one can also see that the amount of instrumental noise and particular realization of the background fluctuations can greatly affect the significance of the needlet coefficients in the vicinity of a collision. Many more collisions were detected in the low-noise region than the high noise region of the sky. For collisions in the exclusion region, there is a significant needlet response for all three needlet types over a range of frequencies, with an average significance exceeding $S > 5$. Collisions in the sensitivity region are typically detected by only one needlet type and frequency, with an average significance near $S \simeq 4$. As our ability to detect 5°, 10°, and 25° collisions is nearly the same, we conclude that these results are fairly representative of our detection limits over all angular scales $\gtrsim 5°$.

These general trends can be contrasted with the response to features found in the end-to-end simulations. Here, blobs are typically detected with a single needlet type and are near the
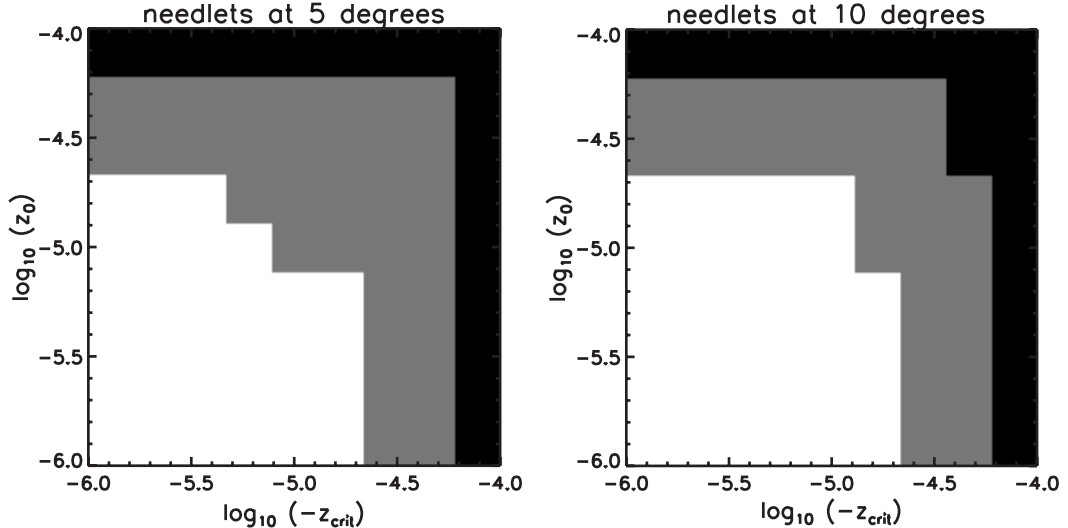
Figure 3.11: Exclusion (black) and sensitivity (grey) regions for the needlet step of the analysis pipeline applied to a set of 5° (left) and 10° (right) simulated bubble collisions. If all realizations at the high and low noise locations yield a detection, we include the collision in the exclusion region; such collisions would certainly be detected as long as they were not significantly masked. If a detection is made in any realization/location, we include the collision in the sensitivity region; such collisions could be found if they were in a favorable location of the sky (i.e., low noise, or a region with a specific realization of background CMB fluctuations).

significance threshold (not surprisingly, as the threshold was chosen to have this property). A feature detected in the data by multiple needlet types and/or frequencies at a significance of $S \geq 4$ would be a good bubble collision candidate. However, we stress that many different underlying features could give rise to such a signal. The following steps in the analysis pipeline, which we now describe, are designed to verify if these candidates are in fact bubble collisions.

3.6.2 Edge detection

The first of the two parallel verification steps of the pipeline tests whether features highlighted by the blob detection stage have circular temperature discontinuities. The unambiguous detection of a circular temperature discontinuity would strongly suggest that a particular feature highlighted by the needlets is in fact a bubble collision. We employ the Canny edge detection algorithm (Canny, 1986), whereby the gradient of an image is generated and thinned into single-pixel proto-edges, the best of which are stitched together into "true" edges. We also use an adaptation of the Circular Hough Transform (CHT) algorithm (Kimme et al., 1975), which assigns a "score" dependent on how many edge pixels lie on circles of varying centers and radii. In this section, we describe our edge detection algorithm, and study its performance on an end-to-end simulation of the experiment, as well as on simulated bubble collisions.

The Canny edge detection algorithm

The Canny edge detector is *the* standard edge detection algorithm in image-processing software, and has recently been used to search for cosmic strings (Danos and Brandenberger, 2010; Amsel et al., 2008). Designed as the optimal algorithm for localized, duplicate-free detection of edges within a noisy image, it uses three steps – smoothed gradient generation, non-maximal suppression and hysteresis thresholding – to extract contiguous edge sections. In Fig. 3.12, we depict each of these three steps as applied to a temperature map containing a circular discontinuity; each of these steps are in turn described below.

1. **Smoothed gradient generation:** The gradient of a Cartesian image is traditionally generated by convolving the image with two small symmetric filters, each determining one orthogonal component of the gradient. A number of simple filters – typically 3×3 pixels – perform the job adequately, but the optimally adaptable filters are the first partial derivatives of the two-dimensional Gaussian (Canny, 1986). Using these filters is equivalent to first convolving the image with a small Gaussian filter (and thus smoothing out the effects of small-scale noise on the gradient calculation, an important step given the small number of pixels involved in the calculation) and then finding its gradient components.

   Unfortunately it is impossible to construct symmetric pixel-based gradient filters that cover the whole sphere. We therefore carry out both the Gaussian smoothing and gradient generation steps in harmonic space, making use of HEALPix's in-built `alm2map_der` subroutine to calculate the magnitude and direction of the gradient at each pixel. We smooth with a Gaussian filter of FWHM $0.22°$ – approximately two pixels' width at the resolution of our input maps – to minimize features due to pixel noise while retaining longer edges.

   The gradient maps are generated before masking to reduce "ringing" from the sharp sky cut back into the map. The smoothing step ensures that any leakage from masked features is restricted to areas a few pixels within the sky cut, and affects only areas a few pixels outside of the cut. Nevertheless, any features found close to the mask should be carefully examined to check if they are generated from within the mask.

2. **Non-maximal suppression:** The second step of the Canny algorithm reduces the smoothed gradients produced by the first step into local maxima. At this stage, all pixels are assumed to belong to a local edge, whose direction is defined to be perpendicular to the local gradient direction. Taking each pixel at a time, the two direct neighbors which lie closest to the local edge are found. The gradient magnitudes of the three pixels are compared, and the central pixel's gradient magnitude is set to zero unless it is the largest of the

three. Processing each pixel in turn reduces the gradient magnitude map to only the local maxima (see the central panels of Fig. 3.12).

As an example, consider the simplest case of crossing a sharp discontinuity along a perpendicular path. The gradient direction is constant at each pixel, whereas the smoothed gradient magnitudes increase until the edge is crossed, when they start to decrease. A nonmaximal suppression algorithm tracks along the path setting all of the gradient magnitudes to zero apart from the pixel on (or closest to) the edge.

3. **Hysteresis thresholding:** At this stage of the Canny algorithm, we have gradient magnitudes and directions for a set of local maxima of varying amplitude, some corresponding to true edges (which may have been disrupted by noise) and others to runs of noisy pixels or to more slowly-varying boundaries of CMB patches. To filter out true edges from the noise, and stitch together any edges that have been split, the final step of the algorithm takes advantage of the fact that, unlike randomly-oriented noise, true edges conserve their gradient magnitude and direction (to an extent affected by the shape of the edge, the pixelization scheme, and the noise level) over their path.

Hysteresis thresholding involves first setting an upper threshold for the gradient magnitude: any pixels surpassing this threshold are considered to be part of true edges, and a new "true edge" map is created with these pixels' positions marked. A second, lower threshold is then set. Hysteresis thresholding then proceeds as follows:

(a) The map is scanned until a true edge pixel is found.

(b) The next potential edge pixel is defined to be the direct neighbor closest to 90° clockwise from the local gradient direction. The gradient direction of this pixel is compared to the current pixel's. To do so, the local phase angles to the current pixel's nearest neighbors are found, and used to define the neighbor closest to the current gradient direction. The neighbors adjacent to this pixel are then determined. The gradient direction at the next potential edge pixel is required to lie between the phase angles of these neighbors. This rather loose requirement allows the algorithm to step along pixelated curved edges.

   i. If the two pixels' gradient directions match within the tolerance,

      A. If the neighbor's gradient magnitude passes the low threshold but not the high, it is considered to be part of a potential true edge. Its position is marked in a history array; then the algorithm "steps onto" this pixel and the process is repeated from step (b) until one of the other conditions is met.

B. If the neighbor's gradient magnitude passes the high threshold, all pixels found on the way from the source pixel are confirmed as a true edge. Their positions are marked on the true edge map, and the algorithm returns to step (a).

C. If the neighbor's gradient magnitude fails both thresholds, the edge is considered to be false: the history of potential edge pixels found on the way from the source pixel is erased, and the algorithm returns to step (a).

ii. If the two pixels' gradient directions do not match within the required tolerance,

A. If the neighbor's gradient magnitude passes the high threshold or the neighbor is already marked in the history array, all pixels found on the way from the source pixel are confirmed as a true edge. Their positions are marked on the true edge map, and the algorithm returns to step (a). This ensures simple branched and looped edges can be reconstructed.

B. If the neighbor's magnitude does not pass the high threshold and the pixel is not already marked in the history array, the edge is considered to be false, the history of potential edge pixels found on the way from the source pixel is erased, and the algorithm returns to step (a).

The entire process is then repeated, choosing the neighbor closest to $90°$ *counter-clockwise* to the local gradient direction in step (b).

The end product of the Canny algorithm is a Boolean map of stitched true edges (see right-hand panel of Fig. 3.12). To reduce computation time, the pipeline from hysteresis thresholding onwards is restricted to the blobs produced by extending the regions passing the needlet significance test to ensure any discontinuity is fully contained.

Care must be taken when setting the thresholds used in the hysteresis thresholding step. If either threshold is set too high, very few edges are confirmed. If the low threshold is set too low, a huge amount of potential edges are considered, and the algorithm runs extremely slowly. As the edges we could potentially find must be comparable in amplitude to the CMB signal and detector noise (as they have not yet been discovered by eye) and have been smoothed by the WMAP beam, we set low thresholds to err on the side of caution. Low and high thresholds of 30% and 40% of the maximum gradient magnitude found in each search region are found empirically to confirm edges generated in simulations in feasible computation timescales. This means that the gradients associated with the strongest CMB features – typically $\sim 1°$ in scale – are classified as edges, as is shown in Fig. 3.12.

Figure 3.12: An illustration of the Canny algorithm for edge detection. The input temperature map contains a circular discontinuity, which can be seen in a map of the gradient magnitude as a local maximum. Non-maximal suppression selects for local maxima in the gradient map. The hysteresis thresholding step finds stitched edges by comparing the local direction of the gradient at adjacent pixels.

Circular Hough Transform

The maps of stitched candidate edges found using the Canny algorithm are processed using the Circular Hough Transform to search for the presence of circular edges. The basic idea, as shown in Fig. 3.13, is to count the number of intersections between circular arcs of varying radii centered on each of the candidate edge pixels and oriented along the local gradient direction. If there is a circular edge in the map, the number of intersections will be maximized at the center of the circular edge when the radius of the circular arc matches that of the edge.

Assuming an edge pixel forms part of a circular edge of angular radius θ_{crit}, one can define a prescription for the set of pixels that are the potential centers of the edge. The two most likely candidates in this set are the pixel θ_{crit} away in the direction of the local gradient, and the pixel the same distance in the opposite direction; the edge could be a step up or a step down. Building in flexibility to cope with pixelation and noise effects, this set is expanded to two annular arcs of radius θ_{crit}, oriented in the direction of the local gradient and centred on the edge pixel.

The CHT works by assuming that *all* edge pixels are part of circular edges. The most likely circle center at a given radius is defined to be the pixel that is included in the greatest number of these arcs, counted using an accumulator array. If the search radius matches the radius of a circular edge within the map, we expect all of the circular edge pixels' arcs to include the true center, and the CHT accumulator will show a single clear peak. If the search and true radii are discrepant, fewer of the circular edge pixels' arcs will intersect, and this peak will appear as a ring with decreased amplitude (see Fig. 3.13). When the search and true radii are very discrepant, any rings will disappear beneath the background due to randomly-oriented noise and other non-circular edges. Note that "non-circular edges" will include the $\sim 1°$ CMB fluctuations that are

Figure 3.13: A depiction of the Circular Hough Transform (CHT). On the left is a Boolean map of edge pixels, as output by the Canny algorithm. Centering a pair of arcs oriented in the direction of the local gradient about each edge pixel, the CHT counts the number of times each pixel is intersected. The presence of a circular edge is indicated by a maximum in the CHT score – the hit count divided by the arc radius – as the arc radii are varied. On the left, we show a set of arcs, centered on four pixels on the circular edge we wish to detect; there will be no clear peak in the CHT score for this radius. Increasing the arc radius to match that of the circular edge (center), there will be a large number of hits at the true center of the edge. On the right, we show the actual map of the CHT score at each pixel for this example. As this has been scanned at the correct angular scale, there is a large peak at the center of the circular edge.

qualified as edges by the hysteresis thresholding step.

To compare the CHT results at different search radii, one must divide out the approximately linear growth with angular radius of the number of pixels in each annular arc. We call this normalized quantity the "CHT score". The most likely center and radius of a circular edge within a map can therefore be found by scanning the map with the CHT at a range of radii and determining the maximum CHT score.

The blob detection step provides the range of scales $\theta_{\mathrm{crit,min}} \leq \theta_{\mathrm{crit},i} \leq \theta_{\mathrm{crit,max}}$ of potential circular edges present in each blob. To determine whether a blob contains a circular edge, we compare the CHT scores obtained by scanning at every $0.25°$ increment within this range, using annular arcs that are $0.25°$ thick and which cover $45°$ of phase about each edge pixel. The annular arcs are therefore approximately two pixels thick, and are fairly wide to account for the effects of pixelation on the gradient direction. The thickness of the CHT annular arcs leads to an uncertainty in the CHT radius of $0.25°$ and position of $0.50°$. If a circular edge is detected, we expect a clear peak in the CHT results for a particular blob.

In Fig. 3.14 we show the output of the edge detection algorithm on our illustrative example bubble collision simulation (see Fig. 3.10). On the left is the portion of the temperature map containing the collision. On the right we plot the CHT score in the pixels that passed the needlet significance threshold for $\theta_{\mathrm{crit},i} = 10°$ (the input value). There is a clear peak at the location of the true center of the simulated bubble collision, which is ~ 3 times the average response at other pixels. Since this feature was flagged in the blob detection step for standard needlets with

−539 ▬▬▬▬▬ 539 μK   0.3 ▬▬▬▬▬ 17.1

Figure 3.14: The temperature map (left) and CHT score (right) for our illustrative example of a 10° bubble collision simulation. The CHT score is recorded at each pixel passing the needlet significance threshold. For a search radius of 10° there is a clear peak in the CHT score at the center of the simulated collision.

$B = 2.5$ at $j = 3$, the range of radii scanned during the CHT step is determined from Table 3.1 to be $5° \leq \theta_{crit} \leq 14°$. This range contains the true radius $\theta_{crit} = 10°$. In Fig. 3.15 we plot the maximum CHT score found in the map for each circular radius, which contains a clear peak at the true radius of the causal boundary. This signal is a clear and unambiguous signature of a bubble collision. From visual inspection of the temperature map, it can be seen that we are able to clearly detect the edge even though the background fluctuations, noise and beam drastically reduce the sharpness of the observed temperature discontinuity.

Analysis of the WMAP end-to-end simulation

We expect strong circular edges to be extremely rare in a purely Gaussian CMB temperature map. However, it is possible that foregrounds, instrumental noise, the mask, and other experimental artifacts could lead to a spurious detection of a circular edge. To evaluate this, we have performed the edge detection step of our analysis pipeline on the features that passed the significance threshold in the WMAP end-to-end simulation (see Table 3.4) with the KQ75 mask applied.

Comparing each feature in the end-to-end simulation with the bubble collision example studied above, the peak structure of the CHT score as a function of angle and morphology in pixel space are both drastically different. Examining the maximum CHT score as a function of circular radius, although there are several peaks, the clearest of which is shown in Fig. 3.16, their amplitude relative to the average score is nowhere near that of the collision example shown in
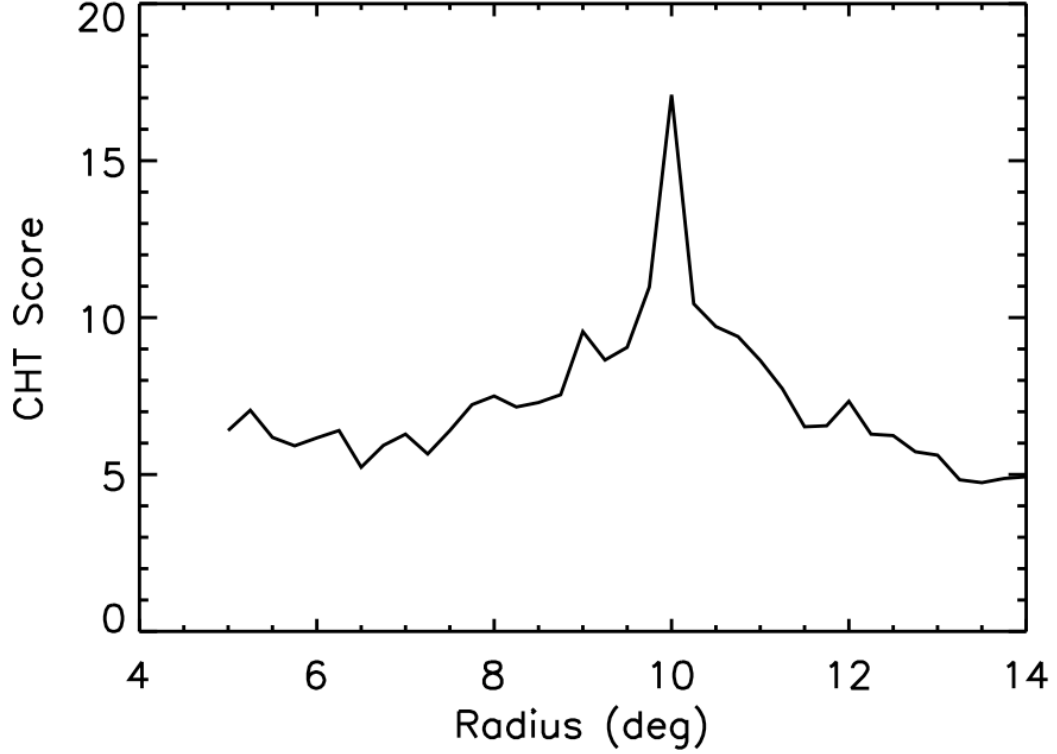
Figure 3.15: The global maximum of the CHT score at each circle radius for the collision simulation shown in Fig. 3.14. The collision has a maximum response for standard needlets with $B = 2.5$ at $j = 3$, which from Table 3.1 sets the search range to be $5° \leq \theta_{\text{crit}} \leq 14°$. The peak of the CHT score at 10° clearly identifies the correct angular scale of the simulated collision.

Fig. 3.15. In addition, from the plots of the CHT score at each pixel, there are typically a number of fairly broad local maxima at different locations with approximately the same score. This is in contrast to the collision example of Fig. 3.14, which yields a highly peaked score around a small number of pixels.

Analysis of bubble collision simulations

To better understand the response of our edge detection algorithm to the signal from a bubble collision in WMAP-quality data, we have analyzed the simulations described in Section 3.5. We use as inputs the blobs found using the first step of the pipeline, and search for circular edges over the range of angular scales appropriate to the needlet type and frequency for each blob (see Table 3.1). We conclude that a true causal edge has been detected if there is a global maximum for the CHT score at the radius of the true edge and the pixel with the highest score is within a typical CHT error (0.5°) of the actual center. We again present our results in the form of a contour plot denoting exclusion and sensitivity regions in the parameter space of z_0 and z_{crit}. This is shown in Fig. 3.17 for simulated bubbles with $\theta_{\text{crit}} = 5°$ and 10°. The plot for the 25° collisions is identical to the plot for 10° collisions.
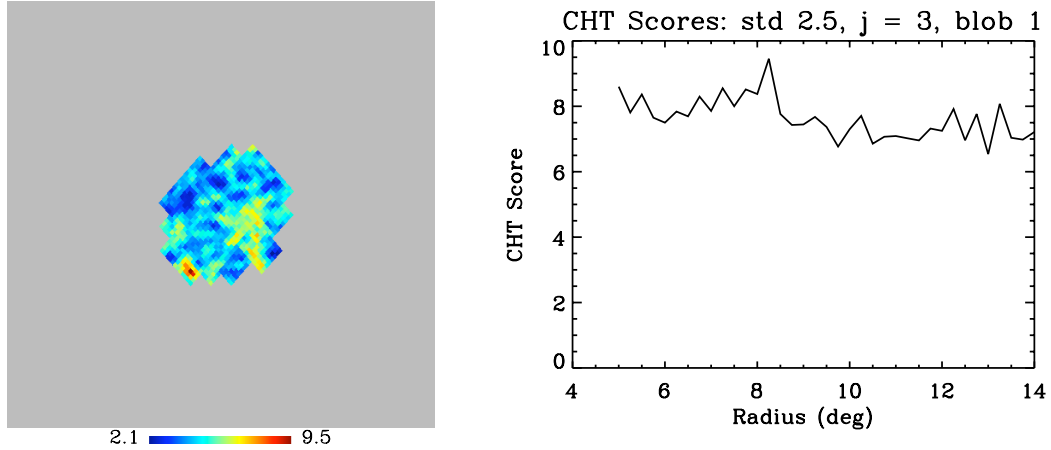
Figure 3.16: The most edge-like feature in the WMAP end-to-end simulation. The contrast in scores as a function of position (left) and radius (right) is greatly reduced compared to the collision example (Figs. 3.14 and 3.15).

Again, from the size of the sensitivity region, we conclude that our ability to make a detection is dependent on the location of the collision and the particular realization of the background CMB fluctuations. The exclusion region for the 10° (and 25°) collisions is far larger than for the 5° collisions. We attribute this to the proliferation of $\sim 1°$-sized features in the background fluctuations, which can both disrupt a significant fraction of the edge pixels in a small collision and swamp the collision signal with their own strong gradients. We therefore expect our sensitivity to edges at small angular scales $\theta_{\rm crit} \lesssim 5°$ to be quite poor at WMAP resolution. As the performance of the edge detection algorithm for 10° and 25° collisions is identical, we conclude that the 10° results are fairly representative of our sensitivities over a wide range of angular scales $\theta_{\rm crit} \gtrsim 10°$. Most of the collisions we mark as a detection have a clear peak in the CHT score of the type seen in Fig. 3.15. If a collision has parameters in the exclusion region, it would be reliably detected. Based on these results, the first two steps of our pipeline can detect bubble collisions with central modulations $|z_0| \gtrsim 3 \times 10^{-5}$ and causal edges $|z_{\rm crit}| \gtrsim 3 \times 10^{-5}$ at $\theta_{\rm crit} \gtrsim 5°$ in WMAP-quality data.

3.6.3 Parameter estimation and model selection

In many CMB anomaly analyses (but not all – see, e.g. Groeneboom et al. (2010); Hoftuft et al. (2009); Cruz et al. (2008)), the significance of a signal is quantified by calculating the frequentist *P-value* of some relevant statistic. This typically involves doing a large number of Monte Carlo realizations of the standard cosmological model (i.e., the "null hypothesis"), calculating the above statistic for each, and finding the fraction for which the statistic has a "more extreme" value than was actually observed. There are several problems with this approach. First, the

Figure 3.17: Exclusion (black) and sensitivity (grey) regions (see Fig. 3.11) for the edge detection step of the analysis pipeline applied to a set of 5° (left) and 10° (right) simulated bubble collisions.

calculated P-value is only a measure of how (un)likely the measured data were given the null hypothesis of the standard model; no comparison is made to an alternative model (which is what we are primarily interested in here). Second, the notion of "more extreme" is fundamentally ambiguous – both "more discrepant" (i.e., values of the statistic which are further from some fiducial expected value than that which was measured) or "less likely" are common choices[9]. The heart of the problem is that all such P-values are integrals over the likelihood, whereas it is only the likelihood of the actual data that is relevant. The fact that the likelihood and its integral generally have a similar qualitative dependence in the tail(s) of the distribution (i.e., both tend to zero for extreme values of the statistic) can mask this problem. In particular, if the tails of the likelihood are Gaussian then the integral that gives the P-value falls off more rapidly than the likelihood itself, and so the resultant P-values are unreasonably harsh on the null hypothesis. A related problem is that that many attempts to identify CMB anomalies using frequentist P-values are overly sensitive to *a posteriori* selection effects (see, e.g. Bennett et al. (2011) and Pontzen and Peiris (2010) for a discussion of this effect). Here the issue – that the statistics being applied to the data are often chosen on the basis of interesting features initially identified in the same data – is not intrinsic to frequentist methods (which, correctly, do not permit any data to be used more than once); but the need to invent a statistic from which to calculate a P-value can make it hard to avoid this trap. For these reasons we do not use P-values in our analysis.

Instead, we adopt a Bayesian approach. Bayes' theorem provides a prescription for parameter

---

[9]For simple, single-peaked, likelihoods these two definitions are at least equivalent, but in some cases (e.g., a likelihood that is constant over some finite range) neither definition is satisfactory.

estimation. In addition, given that we have two well-defined hypotheses, we can utilize Bayesian model comparison to make probabilistic statements about the degree to which the available data (and theoretical prior information) imply that bubble collisions have been observed. As shown by Cox (1946), Bayesian methods are the only self-consistent framework for such calculations. The optimal Bayesian calculation would be to evaluate the likelihood of the entire WMAP dataset under the two models; however this not computationally feasible at present. In Appendix A.1, we outline a set of simplifications that allow us to approximate the optimal Bayesian result. As outlined in Sec. 3.4, we utilize the information on the location and scales of the most probable bubble collision sites obtained in the blob detection step of the pipeline to implement this procedure. Even this reduced problem is computationally demanding: analysis of the blobs detected in the WMAP 7-year data during the first steps of the pipeline requires three days' processing on 28 cores. Working at full resolution is necessary to ensure that any possible circular temperature discontinuities are examined.

These computational limitations also mean we are only able to process a limited number of simulated temperature maps with and without bubble collisions. The WMAP end-to-end simulation provides a great asset at this stage, giving the best possible measure of what false detections are to be expected from experimental effects and any systematic errors that are not included in our likelihood. We also analyze a small number of representative bubble collision simulations to obtain an estimate of the strength of signal we are looking for.

We now describe our methods and the results from simulations in greater detail.

Bayesian formalism

A model of eternal inflation predicts the average number of collisions $\bar{N}_s$ that are, in principle, detectable by our pipeline on the full sky [10]. The ultimate goal of our Bayesian analysis is to evaluate the full posterior probability distribution for $\bar{N}_s$, given a CMB data set $\mathbf{d}$ covering a sky fraction f_{sky}. Using Bayes' theorem, this can be written as

$$\Pr(\bar{N}_s|\mathbf{d}, f_{sky}) = \frac{\Pr(\bar{N}_s)\Pr(\mathbf{d}|\bar{N}_s, f_{sky})}{\Pr(\mathbf{d}|f_{sky})}. \tag{3.26}$$

The form of the posterior depends on the model prior $\Pr(\bar{N}_s)$ and the *evidence* (also known as the model likelihood) $\Pr(\mathbf{d}|\bar{N}_s, f_{sky})$. The evidence is defined by marginalizing the likelihood, $\Pr(\mathbf{d}|\mathbf{m}, \bar{N}_s, f_{sky})$, over the n parameters describing a collision, as specified by the model $\mathbf{m}$. Once the shape of the posterior has been determined, it is normalized using $\Pr(\mathbf{d}|f_{sky})$. The posterior leads directly to constraints on the values of $\bar{N}_s$ consistent with a CMB data set.

---

[10] The number of detectable sources $\bar{N}_s$ is a subset the total number of sources on the sky $\bar{N}$ (Eq. 3.1).

In a landscape scenario, $\bar{N}_{\rm s}$ can be considered as a continuous parameter and the prior $\Pr(\bar{N}_{\rm s})$ would be determined from a measure over the possible values of $\bar{N}_{\rm s}$. We can also view $\bar{N}_{\rm s}$ as a proxy for different models of eternal inflation (i.e., selecting a single value of $\bar{N}_{\rm s}$), as described further in Sec. 3.3. The standard cosmological model without bubble collisions is specified by the case $\bar{N}_{\rm s} = 0$. Using Eq. 3.26, the probability of a model which predicts $\bar{N}_{\rm s}$ collisions (on average) relative to that of the standard cosmological model is

$$\frac{\Pr(\bar{N}_{\rm s}|\mathbf{d}, f_{\rm sky})}{\Pr(0|\mathbf{d}, f_{\rm sky})} = \frac{\Pr(\bar{N}_{\rm s})\Pr(\mathbf{d}|\bar{N}_{\rm s}, f_{\rm sky})}{\Pr(0)\Pr(\mathbf{d}|0, f_{\rm sky})}. \tag{3.27}$$

The model priors and the evidence values play an equal role in this relationship, but in the absence of a detailed understanding of the former, it is often useful to proceed under the assumption that the two models are equally probable *a priori*. A theory predicting an expected $\bar{N}_{\rm s}$ collisions is favoured over the standard model when the relative probability on the LHS of Eq. 3.27 is greater than unity.

It is also useful to provide heuristic conversions between the Bayesian evidence ratio and other commonly used model comparison quantities. The number of "sigma" of an anomaly statistic, $N\sigma$, is often used to characterize the deviation from a null model, but it is unambiguously defined only in the case that the null distribution of the chosen statistic is Gaussian with zero mean. In such a case the probability of measuring an $N\sigma$ deviation is $P(N) \propto \exp(-N^2/2)$, which can be identified approximately with the inverse of the ratio in Eq. 3.27, so that, e.g., a 3σ detection is comparable to a ratio of approximately one hundred. However we emphasize that both the number of sigma and related statistics such as $\Delta\chi^2$ are of limited utility in the context of all but the most trivial model comparison problems.

Computing $\Pr(\mathbf{d}|\bar{N}_{\rm s}, f_{\rm sky})$ by marginalizing over the likelihood for the full prior volume is an immense computational task, requiring the inversion of the full sky WMAP covariance matrix at full resolution and marginalizing over all possible numbers, locations, and sizes of collisions. However, taking advantage of the fact that bubble collisions produce discrete localized effects on the CMB sky, it is possible to approximate the full-sky Bayesian analysis by a patch-wise analysis if the most promising candidate signatures can be identified in advance. We describe in detail in Appendix A.1 an algorithm to perform such a patch-wise approximation to this full multidimensional integral.

The key ingredient is determining the regions of parameter space where the likelihood is significantly peaked, and hence gives the most significant contributions to the evidence. If these regions can be identified, the integral need only be performed over the restricted ranges to obtain an estimate of the evidence at greatly-reduced computational cost. We use the results of the

blob detection step of the analysis pipeline to identify the patches which are likely make the most significant contributions to the integral. Assuming that the bubble collision model likelihood is peaked in the $N_{\rm b}$ detected blobs, we show in Appendix A.1 that the unnormalized posterior can be approximated as

$$\Pr(\bar{N}_{\rm s}|\mathbf{d}, f_{\rm sky}) \propto \Pr(\bar{N}_{\rm s})\, e^{-f_{\rm sky}\bar{N}_{\rm s}} \sum_{N_{\rm s}=0}^{N_{\rm b}} \frac{(f_{\rm sky}\bar{N}_{\rm s})^{N_{\rm s}}}{N_{\rm s}!} \sum_{b_1,b_2,\ldots,b_{N_{\rm s}}=1}^{N_{\rm b}} \left[\prod_{s=1}^{N_{\rm s}} \rho_{b_s} \prod_{i,j=1}^{N_{\rm s}} (1 - \delta_{s_i,s_j}) \right],$$
(3.28)

where the pre-factors reflect the fact that the number of collisions present on the observable sky, $N_{\rm s}$, is the realization of a Poisson-like process (of mean $f_{\rm sky}\bar{N}_{\rm s}$), and ρ_b is the evidence ratio evaluated within a candidate collision region (with data sub-set $\mathbf{d}_b$) using a single bubble collision template

$$\rho_b = \frac{\Pr(\mathbf{d}_b|1)}{\Pr(\mathbf{d}_b|0)}.$$
(3.29)

The posterior can therefore be built from local measures of how well the data are described by the standard model with and without a collision template. Once Eq. 3.28 is obtained in a particular case, it can be normalized, although this is not strictly necessary to perform the parameter estimation and model selection analyses.

To illustrate some possibilities, in Fig. 3.18 we plot the normalized posterior assuming $f_{\rm sky} = 0.7$ (from the KQ75 mask) and a uniform prior on $\bar{N}_{\rm s}$, for the case where there is a single detected blob (left panel), and four detected blobs (right panel). A theory predicting a particular value of $\bar{N}_{\rm s}$ will be preferred to a theory without bubble collisions if the ratio in Eq. 3.27 is larger than one. This amounts to comparing the posterior at some value of $\bar{N}_{\rm s}$ to the posterior at $\bar{N}_{\rm s} = 0$ (dashed line). To prefer *any* theory with bubble collisions, in the one-blob case it is necessary for the blob to yield a local evidence ratio larger than one (here, we plot the posterior assuming $\rho_b = 4$). This is not true when there are multiple blobs, as can be seen in the right panel of Fig. 3.18, where we plot the posterior assuming each blob has a local evidence ratio $\rho_b = 0.5$. The bubble collision hypothesis (for some values of $\bar{N}_{\rm s}$) is preferred even when the local evidence ratios are less than one: a number of marginal detections can be significant when considered together. We can also obtain any desired confidence intervals on $\bar{N}_{\rm s}$ by examining the shape of the posterior (although it is always the whole distribution that is the full answer to any parameter estimation problem).

When the local evidence ratios are large, the posterior can be approximated by Eq. A.16, appropriately normalized. In Fig. 3.19, we plot the posterior in the limit of large evidence ratios (again assuming $f_{\rm sky} = 0.7$) for no blobs, two blobs, and four blobs. Even in the presence of large local evidence ratios, it can be seen that the posterior has a significant spread due to cosmic

Figure 3.18: The normalized posterior $\Pr(\bar{N}_\mathrm{s}|N_\mathrm{b}, f_\mathrm{sky})$ (see Eq. A.16) assuming $f_\mathrm{sky} = 0.7$. In the left panel, we show the posterior obtained for one blob $N_b = 1$ for a local evidence ratio $\rho_b = 4$. Comparing with the posterior at $\bar{N}_\mathrm{s} = 0$ (dashed line), we see that any theory predicting $\bar{N}_\mathrm{s} \lesssim 4$ will be preferred over the theory without bubble collisions. In the right panel, we show the posterior obtained for four blobs with identical local evidence ratios $\rho_b = 1/2$. Again, comparing with the posterior at $\bar{N}_\mathrm{s} = 0$, any theory with $\bar{N}_\mathrm{s} \lesssim 7$ will be preferred over the theory without bubble collisions. When there are multiple blobs, the bubble collision hypothesis can be supported even when the evidence ratio for each blob is less than one.

Figure 3.19: The full posterior $\Pr(\bar{N}_\mathrm{s}|N_\mathrm{b}, f_\mathrm{sky})$ (Eq. A.16) that would be obtained from a conclusive detection (i.e., $\rho_{b_s} \gg 1$) of $N_b = 0, 2, 4$ (solid, dashed, and dot-dashed curves) blobs containing bubble collisions assuming $f_\mathrm{sky} = 0.7$. The presence of a sky cut skews the distribution towards $\bar{N}_\mathrm{s} > N_b$. Note that even when features are conclusively detected, there is an intrinsic uncertainty in $\bar{N}_\mathrm{s}$; this is a form of cosmic variance.

variance: we only have access to one realization of bubble collisions on the CMB sky. Note that this is true even when there are no detected blobs. When there are multiple decisively detected blobs, the posterior correctly assigns a very small probability to $\bar{N}_\mathrm{s} = 0$.

Our analysis also provides constraints on the parameter values of each candidate collision. The constraints on the n template parameters $\mathbf{m}$ are encoded in their joint posterior distribution

$$\Pr(\mathbf{m}|\mathbf{d}_b, 1) = \frac{\Pr(\mathbf{m})\Pr(\mathbf{d}_b|\mathbf{m}, 1)}{\Pr(\mathbf{d}_b, 1)}. \tag{3.30}$$

The marginal distribution of any subset of the parameters is given by integrating $\Pr(\mathbf{m}|\mathbf{d}_b, 1)$ over the remaining parameters which are not of interest. For the bubble collision model the parameters should include both those describing the collision and the global cosmological parameters; marginalizing over the latter would give constraints on the properties of a (putative)

detected collision. We now discuss the analysis of the likelihood and evidence ratios for a patch in greater detail.

Analysis of candidate bubble collision patches

At the heart of the above formalism for assessing the full posterior for $\bar{N}_\mathrm{s}$ is the evaluation of the patch likelihood for a single collision, $\mathrm{Pr}(\mathbf{d}_b|\mathbf{m}, 1)$. Here the data, $\mathbf{d}_b$, are the measured temperature values of the pixels in the vicinity of the detected blob that are not in the sky cut. The bubble collision model parameters, $\mathbf{m}$, should include both those that describe the collision, $\{z_0, z_\mathrm{crit}, \theta_\mathrm{crit}, \theta_0, \phi_0\}$, as well as the cosmological parameters which determine the CMB power spectrum. However any plausible bubble collision would be sufficiently localized that the cosmological parameters are essentially uncorrelated with them; moreover they are sufficiently tightly constrained from CMB measurements that their uncertainties are minimal in the context of a template-matching problem like this. Hence we fix the cosmological parameters to their best-fit WMAP values (Larson et al., 2011) and only the bubble collision parameters are varied. Hence $\mathbf{m} = \{z_0, z_\mathrm{crit}, \theta_\mathrm{crit}, \theta_0, \phi_0\}$ for the bubble collision model, and there are no free parameters in the null model. Indeed, the no-collision model can be treated as a special case of the collision model in which the collision has zero amplitude.

As both the CMB signal and the WMAP noise are Gaussian, the likelihood has the form

$$\mathrm{Pr}(\mathbf{d}_b|\mathbf{m}, 1) \propto \exp\left(-\frac{1}{2}\chi^2\right) = \exp\left\{-\frac{1}{2}[\mathbf{d}_b - \mathbf{t}(\mathbf{m})]^\mathrm{T}\mathsf{C}_b^{-1}[\mathbf{d}_b - \mathbf{t}(\mathbf{m})]\right\}, \qquad (3.31)$$

where $\mathbf{t}(\mathbf{m})$ is the temperature modulation caused by the collision and C_b is the pixel-pixel covariance matrix. The temperature modulation of the p^th pixel is given from Eq. 3.3 as $t_p = 1 + f(\hat{\mathbf{n}}_p)$, where $\hat{\mathbf{n}}$ is the position on the sky. The covariance matrix includes CMB cosmic variance, Gaussian smoothing approximating the WMAP W-band beam, and the pixel-dependent WMAP noise. The covariance between two pixels p and q with angular positions $\hat{\mathbf{n}}_p$ and $\hat{\mathbf{n}}_q$ is hence given by

$$C_{p,q} = N_{p,q} + \sum_\ell \frac{2\ell + 1}{4\pi}\bar{C}_\ell P_\ell(\hat{\mathbf{n}}_p \cdot \hat{\mathbf{n}}_q), \qquad (3.32)$$

where $\bar{C}_\ell$ is the best-fit WMAP CMB power spectrum convolved with a Gaussian beam of FWHM $0.22°$, $P_\ell(x)$ is the Legendre polynomial of degree ℓ, and $N_{p,q}$ is the noise covariance between pixels. This is taken to be

$$N_{p,q} = \delta_{p,q}\frac{\sigma_\mathrm{W}^2}{N_{\mathrm{obs},p}}, \qquad (3.33)$$

where $\delta_{p,q}$ is the Kronecker delta function, $\sigma_{\mathrm{W}} = 6.549\,\mathrm{mK}$ is the RMS noise of the W-band detectors, and $N_{\mathrm{obs},p}$ is the number of times WMAP has observed the p^{th} pixel. To preserve any edges, we must invert C_b at full resolution. Given available computational resources, the maximum area of the sky we can study at any one time is limited to patches of radius $\sim 11°$ surrounding the center of each detected blob.

The evaluation of the evidence integral Eq. A.12 and the full characterization of the posterior distribution of the parameters are both computationally challenging – even when restricted to small patches – as they require a large number of likelihood evaluations. In all but the simplest of cases it is fatally inefficient to evaluate the likelihood over a multi-dimensional grid and so a variety of sampling algorithms have been developed in which the likelihood is only evaluated in the high posterior regions that are of most interest. For both parameter estimation and evidence calculations we use the nested sampling algorithm (Skilling, 2004) as implemented in the publicly available MultiNest package (Feroz et al., 2009). MultiNest performs numerical integration in order to estimate the evidence values; the required convergence of the integration can be adjusted to balance computation speed with accuracy. At the settings we use, the evidence values returned by MultiNest are accurate to $\sim 10\%$. We use the `getdist` routine included in CosmoMC (Lewis and Bridle, 2002) to extract parameter estimates and uncertainties.

The parameter prior $\Pr(\mathbf{m})$ in Eq. 3.30 is derived from theory, previous experimental results, and the limitations of the data-set and pipeline: it encompasses the full prior understanding of what defines a detectable collision. Because we lack a detailed theoretical prediction for the amplitude parameters in each template (as discussed in Sec. 3.3), we assume a uniform prior on z_0 and z_{crit} over the ranges $-10^{-4} \leq z_0 \leq 10^{-4}$ and $-10^{-4} \leq z_{\mathrm{crit}} \leq 10^{-4}$, set by the observed temperature fluctuations in the CMB. Bubbles with larger values of these parameters would have been visible to the naked eye in any existing CMB data-set. Bubble collisions are expected to be distributed isotropically on the CMB sky, and so we assume uniform priors on the full ranges of $\{\cos\theta_0, \phi_0\}$ to ensure that the probability of finding a bubble per unit area is constant across the sphere. Theory predicts that bubble collision radii should range from $0°$ to half-sky, but our pipeline's sensitivity is restricted by CMB power at small scales and computational requirements at large scales. The non-zero prior range for *detectable* bubble collisions is accordingly restricted, and we assume uniform priors on θ_{crit} values in the range $2° \leq \theta_{\mathrm{crit}} \leq 11°$. [11]

To minimize computation time, the evidence integrals are only calculated over the parameter ranges within which the priors are non-zero and the likelihood is peaked. For each feature, the

---

[11] Eq. 3.2 predicts that the angular scale distribution for all bubbles falling within our past light cone varies with $\sin\theta_{\mathrm{crit}}$. However, this is derived under the assumption that collisions do not affect our bubble interior, and a more careful treatment might lead to a correlation between the values of z_0, z_{crit}, and θ_{crit}. To retain consistency with our uniform priors on z_0 and z_{crit}, we assume a uniform prior on θ_{crit}. Regardless, both choices for the prior lead to identical conclusions for the WMAP 7-year data.

angular scale lookup tables (Table 3.1) indicate the range of interest for θ_{crit}. Merging all of the sets of significant pixels found for each feature yields the ranges for $\{\theta_0, \phi_0\}$. As the needlets are equally sensitive to cold and hot features with varying profiles, little information about the ranges of interest for $\{z_0, z_{\text{crit}}\}$ can be extracted from the blob detection results, so the full prior volume must be considered. The accuracy of this procedure has been verified by performing the integral on the same patch of sky using different parameter ranges for $\{\theta_{\text{crit}}, \theta_0, \phi_0\}$. As long as the likelihood peak is encompassed by the parameter ranges given to Multinest, the returned evidence values agree to within numerical accuracy.

For the fiducial collision example shown in Fig. 3.10, our analysis yields an evidence ratio of $\ln \rho_b = 119.8 \pm 0.3$: the collision model is a very good fit to the data. The full-sky posterior would favour any theory predicting bubble collisions over a large range of $\bar{N}_{\text{s}}$. The marginalized bounds on the parameters are compared to the input parameters in the first row of Table 3.6: the agreement is excellent. However, to make a final judgment about a detection, we must ask what types of evidence ratios we get for false detections in the WMAP end-to-end simulation.

Analysis of the WMAP end-to-end simulation

We have performed the full Bayesian parameter estimation and model selection analysis on the blobs found in the WMAP end-to-end simulation (see Table 3.4). The total processing time for the full pipeline to run on this single map is on the order of 12 hours on 28 cores. Our results for the evidence ratios and marginalized parameter constraints for $\{z_0, z_{\text{crit}}, \theta_{\text{crit}}, \theta_0, \phi_0\}$ for each feature are recorded in Table 3.5.

The evidence ratios for the features identified in the blob-detection step of the pipeline are all significantly less than one. We can therefore approximate the full posterior for $\bar{N}_{\text{s}}$ by Eq. A.20, and rule out $\bar{N}_{\text{s}} \gtrsim 1.6$ at the 68% confidence level. The posterior is maximized at $\bar{N}_{\text{s}} = 0$, and we therefore correctly conclude that the data from the end-to-end simulation does *not* warrant augmenting ΛCDM with bubble collisions.

These results from the end-to-end simulation yield quantitative information on the degree to which systematics and foregrounds could mimic the signal from a bubble collision. Reassuringly, no features yield evidence ratios greater than one. To be distinguishable from systematics and foregrounds, we require the evidence ratios that we find for any feature to at least exceed the evidence ratios found in the end-to-end simulation at similar needlet frequencies.

feature	$\theta_{\rm crit}$ range	$\ln \rho_b$	z_0	$z_{\rm crit}$	$\theta_{\rm crit}$	θ_0	ϕ_0
1	$5-14$	-7.9 ± 0.1	$-3.3^{+1.1}_{-1.0} \times 10^{-5}$	$0.0^{+0.4}_{-0.4} \times 10^{-5}$	$8.6^{+1.5}_{-1.5}$	$120.4^{+1.5}_{-1.3}$	$77.8^{+1.8}_{-1.4}$
2	$10-21$	-9.9 ± 0.1	$-2.4^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.4}_{-0.4} \times 10^{-5}$	$10.6^{+0.5}_{-0.6}$	$46.9^{+3.2}_{-2.4}$	$152.0^{+3.0}_{-2.8}$
3	$6-12$	-11.9 ± 0.1	$-2.9^{+1.0}_{-1.0} \times 10^{-5}$	$0.0^{+0.4}_{-0.5} \times 10^{-5}$	$8.5^{+1.9}_{-1.9}$	$13.9^{+0.3}_{-0.3}$	$72.1^{+1.6}_{-1.5}$
4	$4-12$	-6.9 ± 0.2	$4.6^{+1.2}_{-1.2} \times 10^{-5}$	$-0.1^{+0.7}_{-0.6} \times 10^{-5}$	$5.2^{+0.7}_{-1.2}$	$50.3^{+0.6}_{-0.6}$	$221.7^{+0.8}_{-0.9}$
5	$3-7$	-10.7 ± 0.1	$4.1^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.7}_{-0.7} \times 10^{-5}$	$4.4^{+0.5}_{-0.6}$	$80.5^{+0.3}_{-0.3}$	$218.3^{+0.3}_{-0.4}$
6	$2-4$	-11.3 ± 0.1	$-4.2^{+1.5}_{-1.5} \times 10^{-5}$	$0.1^{+1.0}_{-1.0} \times 10^{-5}$	$2.7^{+0.4}_{-0.5}$	$62.6^{+0.4}_{-0.3}$	$146.1^{+0.3}_{-0.4}$
7	$2-4$	-6.6 ± 0.1	$6.5^{+1.5}_{-1.4} \times 10^{-5}$	$-0.1^{+0.9}_{-0.8} \times 10^{-5}$	$3.0^{+0.3}_{-0.3}$	$62.6^{+0.2}_{-0.2}$	$142.0^{+0.3}_{-0.3}$
8	$2-4$	-8.2 ± 0.1	$6.1^{+1.7}_{-1.7} \times 10^{-5}$	$-0.2^{+1.1}_{-1.1} \times 10^{-5}$	$2.5^{+0.2}_{-0.3}$	$111.5^{+0.2}_{-0.2}$	$69.7^{+0.2}_{-0.2}$
9	$3-6$	-9.7 ± 0.1	$5.0^{+1.3}_{-1.3} \times 10^{-5}$	$0.0^{+0.6}_{-0.6} \times 10^{-5}$	$3.7^{+0.4}_{-0.4}$	$36.4^{+0.3}_{-0.3}$	$131.0^{+0.5}_{-0.5}$
10	$2-4$	-9.9 ± 0.2	$5.6^{+1.6}_{-1.5} \times 10^{-5}$	$-0.8^{+0.9}_{-0.9} \times 10^{-5}$	$2.6^{+0.2}_{-0.4}$	$160.2^{+0.2}_{-0.2}$	$235.9^{+0.6}_{-0.7}$

Table 3.5: Results of the Bayesian parameter estimation and model selection analysis for the WMAP end-to-end simulation. The ranges of $\theta_{\rm crit}$ are determined from the needlet response (see Table 3.1). By computational necessity, the evidence integral is truncated at $11°$. Reported error bars are at 68% CL. Angular quantities are quoted in degrees.

example	z_0	$z_{\rm crit}$	$\ln \rho_b$	$\hat{z}_0$	$\hat{z}_{\rm crit}$	$\hat{\theta}_{\rm crit}$	$\hat{\theta}_0$	$\hat{\phi}_0$
1	5.0×10^{-5}	-5.0×10^{-5}	119.8 ± 0.3	$5.1^{+1.0}_{-1.0} \times 10^{-5}$	$-5.0^{+0.3}_{-0.3} \times 10^{-5}$	10.0	57.7	99.2
2	1.0×10^{-5}	-5.6×10^{-5}	136.0 ± 0.2	$2.3^{+1.1}_{-1.1} \times 10^{-5}$	$-5.4^{+0.3}_{-0.3} \times 10^{-5}$	10.0	57.7	99.2
3	1.0×10^{-4}	-1.0×10^{-6}	28.9 ± 0.3	$9.4^{+0.6}_{-0.4} \times 10^{-5}$	$-0.2^{+0.8}_{-0.8} \times 10^{-5}$	$10.0^{+0.9}_{-0.9}$	$57.7^{+0.5}_{-0.7}$	$99.8^{+0.5}_{-0.6}$
4	3.2×10^{-5}	-3.2×10^{-5}	9.0 ± 0.3	$6.4^{+1.0}_{-1.1} \times 10^{-5}$	$-2.0^{+0.3}_{-0.3} \times 10^{-5}$	$8.92^{+0.09}_{-0.02}$	$57.01^{+0.03}_{-0.03}$	$100.27^{+0.04}_{-0.04}$
5	1.0×10^{-5}	-3.2×10^{-5}	-1.0 ± 0.3	$4.2^{+0.6}_{-0.6} \times 10^{-5}$	$-1.2^{+0.3}_{-0.4} \times 10^{-5}$	$8.6^{+0.4}_{-0.9}$	$57.3^{+0.1}_{-0.3}$	$100.1^{+0.3}_{-0.1}$
6	1.0×10^{-5}	-1.8×10^{-5}	-7.2 ± 0.2	$4.8^{+1.2}_{-1.1} \times 10^{-5}$	$-0.1^{+0.7}_{-0.7} \times 10^{-5}$	$6.5^{+0.4}_{-0.7}$	$58.2^{+0.3}_{-0.3}$	$99.9^{+0.3}_{-0.3}$

Table 3.6: The input and marginalized 68% CL parameter bounds for the representative sample of simulated $10°$ collisions. All simulated collisions are located at $\theta_0 = 57.7°$, $\phi_0 = 99.2°$. Hatted quantities are estimates, and un-hatted quantities are inputs. No errors are quoted for the estimated central positions and radii for the cases where there was an extremely strong detection. This is due to the pixelization of the map: variations in collision-centre coordinates or radius of much less than a pixel's width will not affect the pixelated template, and hence will not affect the likelihood. Angular quantities are quoted in degrees.

Analysis of bubble collision simulations

The long processing time, even for a single map, prohibits us from running the Bayesian parameter estimation and model selection analysis on the full set of bubble collision simulations. We therefore choose a small number of representative examples from the set of simulated collisions passing the needlet significance threshold (drawn from the exclusion and sensitivity regions of Fig. 3.11). Six 10° collision simulations were chosen to sample distinct areas of our parameter space, specifically collisions with:

1. a large central amplitude and edge;

2. a small central amplitude but large edge;

3. a large central amplitude but small edge;

4. a medium central amplitude and medium edge;

5. a small central amplitude and medium edge, and;

6. a small central amplitude and small edge.

The first two collisions lie in the CHT exclusion zone, the third in the needlets exclusion zone, and the others in the sensitivity region. All collisions were placed at the low-noise location to maximize the chance of a detection.

The results of the Bayesian analysis of the collision simulations are displayed in Table 3.6. The first example corresponds to the collision in Fig. 3.10, and is clearly a highly significant detection with an evidence ratio of $\ln \rho_b \simeq 120$. The second example is, again, an extremely clear detection, with $\ln \rho_b \simeq 136$. While the evidence for the third example is numerically lower than for the strongly discontinuous cases, at $\ln \rho_b \simeq 29$, it is again a conclusive detection. In each of these examples, the full-sky posterior assuming $N_b = 1$ (which is well approximated by Eq. A.18), would prefer models with bubble collisions over a wide range of $\bar{N}_s$.

For the collisions sampled in the sensitivity region, the maximum needlet significance recorded in each case was around $S \simeq 4$, which is on the upper end of the significances found in the end-to-end simulation: similar features in the data would be passed to the Bayesian analysis section of the pipeline. The evidence ratios were found to be $\ln \rho_b \simeq 9$ for the collision with a medium central amplitude and a medium edge, $\ln \rho_b \simeq -1$ for the collision with the medium edge but a smaller central amplitude, and $\ln \rho_b \simeq -7$ for the collision with a small edge. Since the latter two templates differ only by the value of z_{crit}, this is further proof that the presence of a detectable causal boundary increases our ability to distinguish a collision. In addition, comparing examples 3 and 4, it can be seen that changing the central amplitude by a bit more than a factor of

two yields an evidence ratio that is orders of magnitude larger. Apparently, there are rather sharply defined limits of detection. For these marginal cases, the parameter uncertainties are significantly underestimated due to the relative strength of the CMB and noise. Only in the case of the collision with a medium amplitude and medium edge could we conclude that models with bubble collisions are preferred over those without over a modest range in $\bar{N}_{\rm s}$.

In conclusion, for the simulated collisions in the needlet and CHT exclusion regions of parameter space, our pipeline can clearly determine that the bubble collision hypothesis is favoured for a variety of $\bar{N}_{\rm s}$. In the other cases we have studied, where the collision lies in the needlet sensitivity region, the conclusion is less clear. The evidence ratios are higher than most of those in the end-to-end simulation, but not much greater. They are also small in magnitude, and therefore do not yield full-sky posteriors that favor the bubble collision hypothesis. Thus, while we might rule these features out as being systematics or foregrounds, better data would be needed to definitively establish the bubble collision hypothesis. Furthermore, the bounds on parameter values in detections associated with the sensitivity regions of parameter space should be regarded as rough estimates only. Note also that since the data sets we consider for each blob are restricted to patches of the sky smaller than $11°$, the gain in sensitivity that arises from the existence of a circular temperature discontinuity will not be present for modulations with $\theta_{\rm crit} \gtrsim 11°$. For large features with an edge, the evidence ratios we obtain would therefore be an *underestimate*.

3.6.4 Summary of the analysis pipeline and conditions for claiming a detection

We now summarize the analysis pipeline and the interpretation of its outputs. First, the analysis pipeline segments the sky into "blobs," each of which corresponds to a region which, for some needlet type and frequency, passes our needlet significance threshold. A specific region of the temperature map can be covered by multiple blobs if there is a response for multiple needlet types/frequencies at the same location. The output of this first step in our pipeline is the location, size, and maximum significance associated with each blob. The edge detection step of our pipeline finds the CHT score as a function of assumed circle size and pixel. If there is a clearly peaked global maximum for the CHT score, this can be processed into the most likely circle center and angular scale. In parallel, we calculate the marginalized constraints on the parameters $\{z_0, z_{\rm crit}, \theta_{\rm crit}, \theta_0, \phi_0\}$ and Bayesian evidence ratio ρ_b for each feature. These evidence ratios are then used to construct the full-sky posterior $\Pr(\bar{N}_{\rm s}|N_{\rm b}, f_{\rm sky})$ (Eq. 3.28) which is a function of $\bar{N}_{\rm s}$.

The posterior allows us to put constraints on the possible values of $\bar{N}_\mathrm{s}$ that are consistent with the data. Comparing the value of the posterior at $\bar{N}_\mathrm{s} = 0$ and some particular value of $\bar{N}_\mathrm{s}$ specifies whether or not the ΛCDM model should be superseded by a model that also predicts on average $\bar{N}_\mathrm{s}$ bubble collisions. If a large ratio of the posteriors is obtained, a conclusive detection of the bubble collision hypothesis can be claimed (provided a model that predicts an appropriate value of $\bar{N}_\mathrm{s}$ exists). A clear peak in the CHT score would indicate the presence of a circular temperature discontinuity in the CMB. This is a clear signature of bubble collisions, and would be nearly conclusive evidence for the eternal inflation scenario. We have found using simulations that a clear edge also yields large evidence ratios, indicating that these two tests are complementary. However, an edge is not necessary to verify the bubble collision hypothesis. There is a clear expectation obtained from the end-to-end simulation for the contribution from false detections due to systematics and foregrounds: the absence of a clear peak in the CHT score, and evidence ratios for each blob not exceeding $\ln \rho_b \sim -6.6$ at detectable scales.

3.7 Analysis of the WMAP 7-year data

Our analysis of the W-band WMAP 7-year foreground-reduced temperature map with the KQ75 mask produces a total of 38 blobs passing our needlet sensitivity thresholds. These blobs can be grouped into 15 distinct features, four of which either intersect or are within a few pixels of the main Galactic cut; these features are assumed to be responses to the mask, and we do not consider them further. The properties of the blobs belonging to the 11 remaining features are given in Table 3.7.

A number of these features have been noted previously. Feature 2 is at the same position as the famous Cold Spot (Cruz et al., 2005). In addition, features 1 and 3 are coincident with the most significant hot spots identified in the needlet analysis of Pietrobon et al. (2008). The number of features we have found is consistent with the results from the WMAP end-to-end simulation, although the simulation does not contain as many high-significance features at low j. In addition, the most significant features in the WMAP 7-year data generate responses from multiple needlet types at multiple frequencies (e.g., the Cold Spot is picked out by seven needlet frequencies), whereas features in the end-to-end simulation tend to be highlighted only by a single needlet. Interestingly, 9 of the 11 features identified as significant are in the Southern Galactic hemisphere.

The CHT scores do not have a clear peak at any angular scale or location for any of the detected features. Indeed, the detailed outputs for the data are completely consistent with those obtained for the end-to-end simulation. The largest CHT peak found in the data is

feature	blob	B	j	θ_0	ϕ_0	blob radius	S
1	1	2.5	2	140.1	173.7	4.5	3.76
1	2	1.8	3	140.9	174.4	5.3	3.48
2	1	2.5	3	147.8	209.5	4.1	4.49
2	2	1.8	4	148.2	207.7	5.5	4.55
2	3	1.8	5	148.5	210.2	1.4	3.37
2	4	1.4	7	147.8	209.5	2.5	3.81
2	5	1.4	8	147.8	209.5	4.1	4.58
2	6	1.4	9	147.4	208.1	2.8	4.30
2	7	1.4	10	146.6	207.5	1.3	3.74
3	1	2.5	3	123.2	321.3	2.5	4.09
3	2	1.8	4	122.8	322.4	4.9	3.82
3	3	1.4	7	122.8	321.0	1.9	3.59
3	4	1.4	8	122.8	321.0	3.2	4.01
3	5	1.4	9	122.8	321.0	2.7	4.30
3	6	1.4	10	122.4	320.6	1.5	3.78
4	1	2.5	4	145.1	33.0	0.9	4.20
4	2	1.8	6	145.5	32.4	0.7	3.72
4	3	1.4	11	145.1	33.0	0.9	3.95
5	1	1.8	5	32.2	74.0	1.2	3.41
6	1	1.8	5	128.7	91.1	1.2	3.37
7	1	1.8	5	169.8	181.6	2.3	3.82
7	2	1.8	6	169.0	187.5	0.8	3.76
7	3	1.4	10	169.4	184.7	1.5	4.12
7	4	1.4	11	168.7	187.3	1.1	4.07
8	1	1.8	6	57.9	115.7	0.7	3.78
9	1	1.8	7	152.3	241.8	0.6	4.12
10	1	1.4	10	167.2	268.7	1.0	3.99
10	2	1.4	11	166.8	271.3	1.0	4.09
11	1	1.4	11	115.0	22.5	0.5	3.80
11	2	1.4	12	114.6	22.1	0.5	4.32

Table 3.7: Features found by the needlet transform in the WMAP 7-year data. Features 1 and 3 correspond to the hot spots found in Pietrobon et al. (2008); feature 2 is the Cold Spot (Cruz et al., 2005). Angular quantities are reported in degrees.
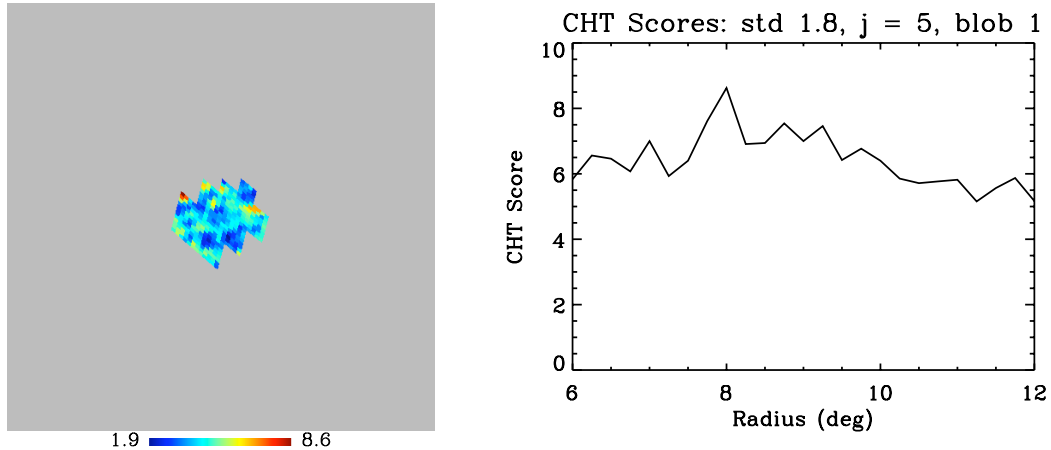
Figure 3.20: The clearest peak found during the edge-detection analysis of the WMAP data. The contrasts in scores as a function of position (left) and radius (right) are comparable to those obtained in the analysis of the end-to-end simulation (Fig. 3.16), and greatly reduced compared to the collision example (Figs. 3.14 and 3.15).

shown in Fig. 3.20 (which should be compared to the most peak-like feature found in the end-to-end simulation, shown in Fig. 3.16). We therefore find no evidence for circular temperature discontinuities in the WMAP 7-year data, and can rule out bubble collisions in the CHT exclusion region defined by simulated collisions shown in Fig. 3.17.

The marginalized parameter constraints and local evidence ratio for each of the features is recorded in Table 3.8. [12] Features 2, 3, 7, and 10 have evidence ratios significantly larger than those found in the collision-free end-to-end simulation ($\ln \rho_b \sim -6.6$), specifically -4.6, -4.1, -5.4 and -3.8 respectively. Assuming $N_b = 4$, and using these values for the local evidence ratios in Eq. 3.28, we find that the posterior is maximized at $\bar{N}_s = 0$, and we can constrain $\bar{N}_s < 1.6$ at 68% CL. One would need roughly $\ln \rho_b \sim -1$ for each of the four features to prefer the bubble collision hypothesis for any value of $\bar{N}_s$. Therefore, the WMAP 7-year data does not warrant adding bubble collisions to ΛCDM.

Although the local evidence ratios found for the WMAP 7-year data were not large enough to yield support for the bubble collision hypothesis, they are about an order of magnitude larger than what was expected from systematics based on the end-to-end simulation. The analysis of future data sets may increase the significance of these blobs if they are indications of bubble collisions, or else they will decrease in significance if they are not; in any case they are the most significant features on our sky, and thus take priority in being further investigated with better data. Thus, we now examine these four most significant features in more detail. The location of each of the four features on the sky is shown in Fig. 3.21. A closer view of each feature is shown

[12]Since we are limited to patches of the sky 11° in radius, the evidence ratios we have obtained for features whose θ_{crit} priors extend beyond $\sim 11°$ will be underestimated if a weak edge exists outside the patch of sky considered.

in Fig. 3.22, along with plots of the needlet significances S triggering the Bayesian analysis step, collision templates for the marginalized parameter constraints found in each case, and the CMB sky as it would appear with these template contributions removed.

To confirm that these features are not due to residual foregrounds, we have also applied our suite of needlet transforms to the WMAP 7-year Q (41 GHz) and V (61 GHz) band foreground-reduced maps. Taking all of the needlets which generate a significant response for the four most significant features, we calculate the average of the needlet coefficients within the regions described by the estimated bubble templates. The results are plotted in Fig. 3.23. We show, for each blob forming part of a feature, the W-band-normalized needlet coefficient averages given by

$$\Delta\beta_{jk,\mathrm{Q/V}} = \frac{\bar{\beta}_{jk,\mathrm{Q/V}} - \bar{\beta}_{jk,\mathrm{W}}}{\bar{\beta}_{jk,\mathrm{W}}}, \tag{3.34}$$

where $\bar{\beta}_{jk,\mathrm{Q/V/W}}$ is the pixel-averaged needlet coefficient value in a given WMAP frequency band. The plots are consistent with no change in the strength of the signal with frequency, suggesting that the features are not due to foreground contamination.

feature	$\theta_{\rm crit}$ range	$\ln \rho_b$	z_0	$z_{\rm crit}$	$\theta_{\rm crit}$	θ_0	ϕ_0
1	$12-38$	N/A	N/A	N/A	N/A	N/A	N/A
2	$4-21$	-4.6 ± 0.2	$-4.8^{+1.4}_{-1.4} \times 10^{-5}$	$-0.1^{+0.7}_{-0.7} \times 10^{-5}$	$6.4^{+1.7}_{-1.1}$	$147.3^{+0.7}_{-0.6}$	$208.0^{+1.5}_{-1.4}$
3	$4-21$	-4.1 ± 0.2	$5.2^{+1.2}_{-1.2} \times 10^{-5}$	$0.1^{+0.5}_{-0.5} \times 10^{-5}$	$6.5^{+0.7}_{-0.7}$	$123.0^{+0.7}_{-0.7}$	$320.8^{+1.0}_{-1.0}$
4	$2-7$	-7.3 ± 0.1	$-5.2^{+1.3}_{-1.3} \times 10^{-5}$	$0.0^{+0.9}_{-0.9} \times 10^{-5}$	$3.2^{+0.8}_{-0.8}$	$145.3^{+0.5}_{-0.4}$	$32.8^{+0.9}_{-1.1}$
5	$6-12$	-9.2 ± 0.2	$3.7^{+1.2}_{-1.1} \times 10^{-5}$	$-0.1^{+0.4}_{-0.4} \times 10^{-5}$	$6.9^{+0.5}_{-0.9}$	$32.6^{+0.6}_{-0.6}$	$74.3^{+1.3}_{-1.1}$
6	$6-12$	-9.7 ± 0.1	$3.0^{+1.1}_{-1.2} \times 10^{-5}$	$-0.1^{+0.4}_{-0.4} \times 10^{-5}$	$7.9^{+2.2}_{-1.9}$	$128.6^{+0.7}_{-0.7}$	$91.8^{+1.2}_{-1.3}$
7	$3-12$	-5.4 ± 0.2	$-5.0^{+1.3}_{-1.4} \times 10^{-5}$	$0.0^{+0.6}_{-0.6} \times 10^{-5}$	$4.6^{+0.9}_{-1.1}$	$169.0^{+0.6}_{-0.6}$	$185.7^{+3.9}_{-3.9}$
8	$3-7$	-8.6 ± 0.1	$4.4^{+1.3}_{-1.2} \times 10^{-5}$	$0.0^{+0.7}_{-0.7} \times 10^{-5}$	$4.4^{+0.6}_{-0.6}$	$58.0^{+0.4}_{-0.4}$	$115.7^{+0.5}_{-0.6}$
9	$2-4$	-9.0 ± 0.2	$-6.0^{+1.6}_{-1.6} \times 10^{-5}$	$0.0^{+1.0}_{-1.0} \times 10^{-5}$	$2.3^{+0.2}_{-0.3}$	$152.1^{+0.2}_{-0.2}$	$241.9^{+0.4}_{-0.4}$
10	$3-8$	-3.8 ± 0.2	$-6.1^{+1.3}_{-1.3} \times 10^{-5}$	$0.1^{+0.9}_{-0.8} \times 10^{-5}$	$4.2^{+0.4}_{-0.4}$	$167.2^{+0.3}_{-0.3}$	$269.1^{+1.6}_{-1.4}$
11	$2-6$	-8.1 ± 0.1	$-5.9^{+1.5}_{-1.6} \times 10^{-5}$	$0.2^{+0.8}_{-0.8} \times 10^{-5}$	$2.5^{+0.4}_{-0.5}$	$114.9^{+0.4}_{-0.4}$	$22.4^{+0.4}_{-0.4}$

Table 3.8: Results of the Bayesian parameter estimation and model selection analysis for the WMAP 7-year data. Reported error bars are at 68% CL. Angular quantities are reported in degrees. The ranges of $\theta_{\rm crit}$ are determined from the needlet response (see Table 3.1). By computational necessity, the evidence integral is truncated at $11°$. Hence, an evidence ratio for feature 1 could not be calculated as its $\theta_{\rm crit}$ range lies entirely beyond this upper bound. The angular positions $\{\theta_0, \phi_0\}$ can be related to Galactic coordinates through longitude $l_0 = \phi_0$ and latitude $b_0 = 90° - \theta_0$.
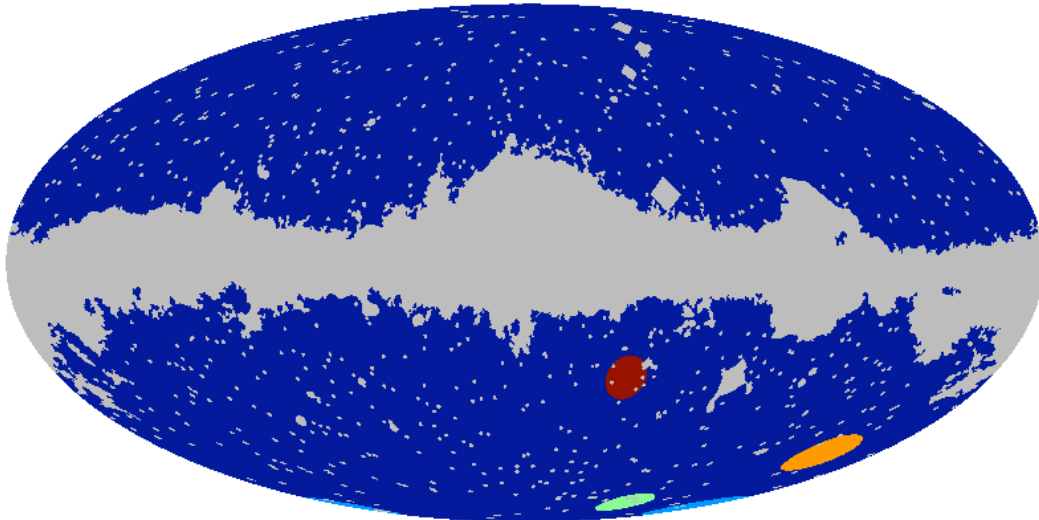
Figure 3.21: Full-sky map showing the positions and sizes of the four features with largest evidence ratios, alongside the 7-year KQ75 sky cut. Feature 2 is plotted in orange, feature 3 in red, feature 7 in light blue and feature 10 in light green.

3.8 Conclusions and outlook

An exciting opportunity to confront the eternal inflation scenario with experiment lies in the observation of collisions between other bubble universes and our own. In this paper, we have described an algorithm to search for the imprint of bubble collisions on the cosmic microwave background, and applied it to the WMAP 7-year data. Our search algorithm targets the generic signatures expected from bubble collisions: azimuthal symmetry, long-wavelength modulation of the temperature confined to discs on the sky, and circular temperature discontinuities. For this reason, we expect our analysis to be fairly robust under changing assumptions about the underlying theory, which is presently rather poorly understood.

The analysis pipeline we have developed takes a two-pronged approach, applied in parallel. The first uses heuristic techniques to test for the presence of features specific to bubble collisions. The second is a fully Bayesian algorithm for the general problem of non-Gaussian source detection, implemented as a patch-wise approximation to the full-sky model selection and parameter estimation problem. The data set is segmented in a completely automated way, allowing us to avoid *a posteriori* selection effects associated with choosing the most "interesting" features on the CMB sky by hand. The algorithm is tested and thresholds at each step are calibrated using extensive simulations, and then frozen before ever looking at the data, to follow as much as possible the philosophy of a blind analysis. Candidate collisions are identified from an input temperature map based on the response to a suite of needlet transforms (calibrated using simulations with and without bubble collisions), and grouped into "blobs." These blobs are scrutinized for circular

113

Figure 3.22: Maps of the four features with largest evidence ratios. The top row shows the W-band temperature map in the locality of the four features, masked with the KQ75 mask. Overlaid are circles indicating the estimated position and angular scale found in each case. The second row contains plots of the masked needlet significances for the needlets whose θ_{crit} priors produced the largest evidence ratios. These plots appear pixelated as the blob detection step is carried out at reduced resolution. The third row shows the bubble collision templates corresponding to the estimated model parameters; these templates are subtracted from the W-band data in the fourth row. The width of each plot is $\sim 16.7°$.

temperature discontinuities using an edge detection algorithm. The quantitative significance of an edge is characterized using the Circular Hough Transform (CHT). The blobs are also used to construct an approximation to the full-sky Bayesian parameter estimation and model selection problem for bubble collisions. The posterior probability distribution over the expectation value for the number of detectable collisions, $\bar{N}_{\mathrm{s}}$, is then obtained. This allows us to quantify which of the two models – a theory which predicts on average $\bar{N}_{\mathrm{s}}$ bubble collision signatures described by temperature modulations of the form given in Eq. 3.4, or else the standard model (specified by $\bar{N}_{\mathrm{s}} = 0$) with CMB plus realistic noise and beam effects – better explains the data.

Applying our analysis pipeline to simulations, we have found that a circular temperature discontinuity at the causal boundary is a clear signature of bubble collisions.[13] Although our

---

[13]The observational detection of a circular temperature discontinuity is so unlikely to arise spuriously that it

Figure 3.23: WMAP channel frequency dependence of the features highlighted by our pipeline. The W-band-normalized difference in pixel-averaged needlet coefficients between the Q and W bands (triangles) and V and W bands (diamonds) are plotted for each blob making up a feature, as highlighted by the full suite of needlet transforms.

analysis can identify collisions without temperature discontinuities, their presence greatly increases our ability to make a conclusive detection. Both the edge-detection and Bayesian model selection steps have the ability to identify a causal boundary in the patches of the sky that are highlighted as candidate collisions by the blob detection step of our analysis pipeline. We have found no evidence for circular temperature discontinuities in the WMAP 7-year data using either method. Based on our analysis of simulations, this allows us to rule out the presence of collisions in the exclusion region of Fig. 3.17. For collisions larger than $\theta_{\mathrm{crit}} \gtrsim 10°$, this corresponds to $10^5 |z_{\mathrm{crit}}| \lesssim 3$–$6$ for the amplitude of the circular temperature discontinuity defined in Eq. 3.4. For collisions on smaller scales, the CHT step loses sensitivity due to the proliferation of degree-scale blobs in the background CMB fluctuations.

The posterior evaluated using the WMAP 7-year data is maximized at $\bar{N}_{\mathrm{s}} = 0$, and constrains $\bar{N}_{\mathrm{s}} < 1.6$ at 68% confidence. We therefore conclude that this data set does not favor the bubble collision hypothesis for any value of $\bar{N}_{\mathrm{s}}$. In light of this null detection, comparing with

provides conclusive evidence of a detection.

the simulated bubble collisions, we can constrain the central amplitude of the temperature modulation caused by the collision (defined in Eq. 3.4) to be $z_0 \lesssim 1 \times 10^{-4}$ over the range of scales $\theta_{\text{crit}} \gtrsim 5°$ we have simulated. If the collision is described by a single super-Hubble wavelength mode confined to a disc on the sky, from Eq. 3.12 we can use these bounds (with the largest collision size we have simulated: $\theta_{\text{crit}} = 25°$) to constrain $\Omega_k^{1/2}\Phi(0) \lesssim 7 \times 10^{-4}$ (where Ω_k is the present component in curvature and $\Phi(0)$ is the initial magnitude of the Newtonian potential caused by the collision). More generally, Eq. 3.13 bounds the nucleation rate of bubbles in our parent vacuum, provided gravitational waves and negative curvature are observed with future experiments.

The prior on the average number of collisions is chosen to be uniform, parametrizing our theoretical ignorance of this parameter and allowing easy reinterpretation of the results should a better-motivated prior be uncovered by ongoing research. Nevertheless, the fact that the posterior is so strongly peaked at $\bar{N}_s = 0$ means that ΛCDM will be favored unless there is a strong reason *a priori* to prefer models producing multiple collisions.

Although we have obtained a null result, our analysis pipeline has identified four features in the WMAP 7-year data that have Bayesian evidence ratios that are significantly larger than expected for false detections from an end-to-end simulation of the WMAP experiment. Two of these features (features 2 and 3) have been noted in previous literature. Feature 2 corresponds to the WMAP Cold Spot (Cruz et al., 2005) (see Cruz et al. (2010) for a review of its properties), and feature 3 was identified using standard needlets in Pietrobon et al. (2008). All four features are far from the Galactic cut of the KQ75 7-year mask (see Fig. 3.21), and none appear to be responses to the point source components of the mask (see Fig. 3.22). We have confirmed that the signal in each case is not strongly dependent on the frequency band used (see Fig. 3.23), providing evidence that these features are not due to astrophysical foregrounds. A number of analyses, most recently the redshift analysis of Bremer et al. (2010), suggest that the Cold Spot is primordial and not associated with the integrated Sachs-Wolfe effect of a large void. Further studies of the other three features would be needed to confirm that they are truly primordial.

Our ability to detect bubble collisions will improve greatly with data from the *Planck* satellite. Firstly, *Planck*'s high-precision measurement of the CMB power spectrum will provide a near-ideal characterization of the dominant "noise" in the analysis, i.e. the CMB itself, increasing the algorithm's accuracy. Secondly, *Planck*'s decreased instrumental noise will enlarge the exclusion and sensitivity regions in parameter space for the needlet step of the analysis, as evidenced by our ability to detect more simulated collisions in low-noise regions of the WMAP data. While measurements of the temperature power spectrum are cosmic-variance-limited at large

scales, measurements of individual patches of data are not, and so decreasing the instrumental noise increases sensitivity. Furthermore, the threefold increase in resolution offered by *Planck* will greatly improve our ability to detect circular edges. In addition, the polarization data from *Planck* will be of sufficient resolution to look for complementary signatures of bubble collisions (Czech et al., 2010; Dvorkin et al., 2008). Such an analysis should be able to confirm if the features we have identified are in fact bubble collisions.

It is also important to determine if other theories predicting azimuthally symmetric features in the CMB (Cornish et al., 1998; Cruz et al., 2008; Afshordi et al., 2010; Kovetz et al., 2010) are better fits to the data. The blob and edge detection steps in our analysis pipeline are sensitive to a variety of possible signatures, and given a model, the Bayesian model comparison step could be easily tailored to accommodate different forms of the temperature modulation. Because our pipeline is automated, we can compare the evidence ratios obtained for different models to decide which is a better fit, without recourse to *a posteriori* choices of which features to analyze.

In conclusion, we have presented a powerful algorithm for analyzing CMB data for signatures of bubble collisions. Applying this pipeline to the WMAP 7-year data, we have constrained the possible parameter space of bubble collisions, as well as identifying interesting candidate signatures in the data for further investigation. Future data from the *Planck* experiment will allow us to greatly improve on these results. If confirmed, the presence of bubble collisions in the CMB would be an extraordinary insight into the origins of our universe.

3.9 Acknowledgements

Chapter 4

Robust constraint on cosmic textures from the cosmic microwave background

4.1 Abstract

Fluctuations in the cosmic microwave background (CMB) contain information which has been pivotal in establishing the current cosmological model. These data can also be used to test well-motivated additions to this model, such as cosmic textures. Textures are a type of topological defect that can be produced during a cosmological phase transition in the early universe, and which leave characteristic hot and cold spots in the CMB. We apply Bayesian methods to carry out a rigorous test of the texture hypothesis, using full-sky data from the Wilkinson Microwave Anisotropy Probe. We conclude that current data do not warrant augmenting the standard cosmological model with textures. We rule out at 95% confidence models that predict more than 6 detectable cosmic textures on the full sky.

4.2 Introduction

Precision measurements of anisotropies in the cosmic microwave background (CMB) radiation have been instrumental in establishing the standard "ΛCDM" model of cosmology: that the universe is composed mostly of dark energy and dark matter, with structures seeded by nearly scale-invariant Gaussian density fluctuations. In addition to establishing ΛCDM, the CMB is also an ideal observable for determining if there are departures from this baseline model.

In this paper, we present a novel algorithm for the analysis of CMB data from the Wilkinson Microwave Anisotropy Probe (WMAP) (Bennett et al., 2003a) to search for the presence of a class of topological defects known as cosmic textures (Turok, 1989). Although textures (and other topological defects, such as cosmic strings) have been ruled out as the dominant source for the primordial perturbations (Albrecht et al., 1999, 1997), their production is inevitable in theories in which a non-Abelian global symmetry is broken (Kibble, 1976). Previous work (Cruz et al., 2007; Cruz et al., 2008; Cruz et al., 2010; Vielva et al., 2011) presented evidence, based on the properties of a single feature in the CMB, that ΛCDM should be augmented by adding cosmic textures. Implementing Bayesian model selection using data on the full sky, we are able to put the texture hypothesis to a much more stringent test. Incorporating this extra information, we conclude that the WMAP 7-year data do not warrant augmenting ΛCDM with cosmic textures, and place constraints on theories giving rise to textures. Our algorithm is easily extendable to incorporate better data, multiple datasets, and a more complete theoretical understanding of the properties and evolution of cosmic textures.

4.3   Cosmic texture theory

The theory of cosmic textures posits a phase transition in the early universe in which a non-Abelian global symmetry is broken. In an expanding universe, different regions of the universe can be out of causal contact, obstructing the symmetry-breaking phase transition from occurring in the same manner everywhere in space (Kibble, 1976). Therefore, a scale-invariant set of knots in the symmetry-breaking order parameter inevitably form: these are cosmic textures. Subsequent to the phase transition, knots from the distribution come into causal contact with their surroundings and undergo collapse (Turok, 1989; Turok and Spergel, 1991; Spergel et al., 1991; Pen et al., 1994; Turok and Spergel, 1990). Upon collapse, textures unwind when the gradient energy of the field configuration exceeds the energy required to restore the global symmetry. As the field re-orders, the energy of the texture configuration is released as an outgoing shell of scalar field radiation.

The gravitational potential associated with a cosmic texture varies in time as it collapses and subsequently explodes. CMB photons passing through an evolving texture will be redshifted if they pass through a collapsing texture, and blueshifted if they pass through an exploding texture (Turok and Spergel, 1990). Each texture unwinding event therefore produces an additive hot or cold spot on the sky, which can be approximated as a disc whose angular size, θ_c, depends on the distance to the texture unwinding event and whose amplitude, $\epsilon \equiv 8\pi^2 G\eta^2$, depends on the scale of symmetry breaking η. The temperature profile in the central region of an unwinding

event situated at the Galactic North Pole can be approximated as (Turok and Spergel, 1990)

$$t(\theta, \phi) = \frac{(-1)^p \, \epsilon}{\sqrt{1 + 4 \left(\frac{\theta}{\theta_c} \right)^2}}, \tag{4.1}$$

where θ and ϕ are Galactic co-latitude and longitude, respectively, and $p = \{0, 1\}$. Here $p = 0$ corresponds to a hot spot and $p = 1$ to a cold spot. The form of the modulation for large θ is presently unknown; following Cruz et al. (2007) and Cruz et al. (2008), we match onto a Gaussian profile at the half-maximum radius, $\theta_* = \sqrt{3}\theta_c/2$.

The angular scale distribution is determined by the evolution of the cosmological horizon during the matter-dominated era (Turok and Spergel, 1990), which is fixed by the late-time cosmological parameters of the ΛCDM model. In addition, each feature is equally likely to be: (i) hot or cold, and (ii) located at any point on the sky, allowing us to define the prior over the "local" template parameters as

$$\Pr(p, \theta_c, \theta_0, \phi_0 | \epsilon) = \frac{\sin \theta_0}{4\pi\theta_c^3} \left(\frac{1}{(2°)^2} - \frac{1}{(50°)^2} \right)^{-1}, \tag{4.2}$$

where $0 \leq \theta_0 \leq \pi$, $0 \leq \phi_0 < 2\pi$, and we take $2° \leq \theta_c \leq 50°$. The lower limit on θ_c results from the large power on degree scales in the CMB; the upper limit stems from the fact that templates with $\theta_c > 50°$ are large enough to cover the whole sky and overlap themselves, rendering Eq. 4.1 invalid.

Different theories giving rise to textures yield different predictions for the symmetry breaking scale and frequency of texture unwinding events; however, all mechanisms produce CMB modulations of the form described in Eq. 4.1. Observationally, theories giving rise to textures are therefore differentiated only by the expected number of *detectable* texture unwinding events on the CMB sky, $\bar{N}_s$, and their amplitude, ϵ. In our analysis, the background CMB fluctuations dominate the definition of detectability. The prior probability $\Pr(\bar{N}_s, \epsilon)$ is set by using simulations to determine the parameter space to which our algorithm is sensitive, as we will discuss shortly.

The ΛCDM+texture model can therefore be fully described by: the standard ΛCDM parameters; a set of "global" texture parameters, $\mathbf{m}_0 = \{\bar{N}_s, \epsilon\}$, labelling theories; a set of "local" parameters, $\mathbf{m}_i = \{p, \theta_c, \theta_0, \phi_0\}_i$, describing each texture; and theoretical priors on these parameters, $\Pr(\mathbf{m}_0)$ and $\Pr(\mathbf{m}_i | \mathbf{m}_0)$. To test the ΛCDM+textures model against vanilla ΛCDM, we need only vary those parameters that are unique to the more complex model (Dickey, 1971). We therefore fix the ΛCDM parameters to their best-fit values from the analysis of WMAP

7-year data (Komatsu et al., 2011) (hereafter referred to as WMAP7). We will now describe the specifics of our search algorithm.

4.4 Searching for textures

The fundamental question posed by this analysis is: are the WMAP7 data better described by the standard ΛCDM cosmological model or ΛCDM plus cosmic textures? The goal is to calculate the joint posterior distribution of $\bar{N}_s$ and ϵ, given the available data. Pure ΛCDM corresponds to $\bar{N}_s = 0$. We avoid the *a posteriori* selection effects associated with postdicting an explanation for anomalous portions of the data (see Bennett et al. (2011) for an in-depth discussion) by performing an analysis of the full dataset. This is important, given that previous evidence (Cruz et al., 2008) for cosmic textures in the CMB was based on the analysis of a single anomalous feature, the so-called CMB Cold Spot (Cruz et al., 2010).

Given an expected number of detectable textures over the whole sky, $\bar{N}_s$, the actual number of detectable textures, N_s, is drawn from a Poisson distribution with mean $f_{\text{sky}}\bar{N}_s$, where f_{sky} is the fraction of the sky covered by the observations. The full posterior probability distribution of the global parameters describing the texture model, ϵ and $\bar{N}_s$, is given by marginalizing the likelihood, $\Pr(\mathbf{d}|\mathbf{m}_1,\ldots \mathbf{m}_{N_s}, \epsilon, N_s, f_{\text{sky}})$, weighted by the prior, over the (unknown) actual number of textures and their individual properties. This is an extremely challenging integral to evaluate directly, but a good approximation to it can be found by identifying the regions of this parameter space in which the likelihood is appreciable and only including these contributions (Feeney et al., 2011a,b). Extending this formalism to also incorporate the global parameter ϵ allows us to self-consistently combine the evidence that each candidate is a texture into a global constraint on the texture theory. The resultant expression (cf. Feeney et al. (2011a) and Feeney et al. (2011b)) is

$$
\begin{aligned}
\Pr(\epsilon, \bar{N}_s | \mathbf{d}, f_{\text{sky}}) \quad \simeq \quad & \frac{\Pr(\epsilon, \bar{N}_s)\Pr(\mathbf{d}|N_s = 0, f_{\text{sky}})}{\Pr(\mathbf{d}|f_{\text{sky}})} e^{-f_{\text{sky}}\bar{N}_s} \times \\
& \sum_{N_s=0}^{N_b} \frac{(f_{\text{sky}}\bar{N}_s)^{N_s}}{N_s!} \sum_{b_1,b_2,\ldots,b_{N_s}=1}^{N_b} \Delta^{b_1 b_2 \ldots b_{N_s}} \prod_{s=1}^{N_s} \rho_{b_s}(\epsilon) ,
\end{aligned}
\tag{4.3}
$$

where $\Pr(\epsilon, \bar{N}_s)$ is the prior (the properties of which are discussed below), $\Pr(\mathbf{d}|N_s = 0, f_{\text{sky}})$ is the likelihood for ΛCDM (i.e. the likelihood assuming no textures), and N_b denotes the number of regions on the sky, or "blobs," containing candidate signatures, each labeled by b_i. The actual number of detectable textures N_s lies between 0 and N_b. The quantity $\Delta^{b_1 b_2 \ldots b_{N_s}}$ is one when all indices take distinct values and zero otherwise: it generates all permutations of N_s textures

located in N_b blobs, assuming no more than one texture per blob. Finally, the quantity $\rho_{b_i}(\epsilon)$, defined as

$$\rho_{b_i}(\epsilon) \equiv \sum_{p=0,1} \frac{\int_{b_i} \mathrm{d}\theta_0 \mathrm{d}\phi_0 \int \mathrm{d}\theta_c \mathrm{Pr}(p, \theta_c, \theta_0, \phi_0 | \epsilon) \mathrm{Pr}(\mathbf{d}_{b_i} | p, \theta_c, \theta_0, \phi_0, \epsilon, N_s = 1, f_{\mathrm{sky}})}{f_{\mathrm{sky}} \mathrm{Pr}(\mathbf{d}_{b_i} | N_s = 0, f_{\mathrm{sky}})}, \qquad (4.4)$$

is a patch-based evidence ratio evaluated in each blob: this is a measure of how much better ΛCDM plus a single texture fits the data than pure ΛCDM, considering only the data in blob b_i. The factor of f_{sky} appearing in the denominator accounts for the fact that we are restricted to detecting textures outside the sky cut. Unless the data provide strong support for the presence of a texture, the evidence ratio penalizes this more complicated model through the larger volume of parameter space that must be considered in constructing the priors, thus self-consistently implementing Occam's razor.

The likelihood for blob b_i is

$$\mathrm{Pr}(\mathbf{d}_{b_i} | p, \theta_c, \theta_0, \phi_0, \epsilon, N_s = 1, f_{\mathrm{sky}}) = \frac{1}{(2\pi)^{N_{\mathrm{pix},b_i}/2} |\mathbf{C}_{b_i}|} e^{-[\mathbf{d}_{b_i} - \mathbf{t}(\epsilon, \mathbf{m}_1)] \mathbf{C}_{b_i}^{-1} [\mathbf{d}_{b_i} - \mathbf{t}(\epsilon, \mathbf{m}_1)]^{\mathrm{T}}/2}, \quad (4.5)$$

where N_{pix,b_i} is the total number of pixels in the blob, $\mathbf{d}_{b_i}$ are the data points in the blob, and $\mathbf{C}_{b_i}$ is the pixel-pixel covariance matrix using only pixels contained in the blob, which includes the fluctuations due to ΛCDM as well as instrumental noise and the effects of the beam.

4.5 Locating texture candidates

To evaluate Eq. 4.3, we must first identify the most promising candidates in the map. We do so by employing the suite of spherical needlet transforms (Marinucci et al., 2008; Pietrobon et al., 2008; Scodeller et al., 2011) defined in Feeney et al. (2011a). Filtering CMB temperature maps with spherical needlets yields information about both the position and angular size of interesting features. The statistics of the filtered field (established using 3000 simulated Gaussian CMB realizations) can then be used to assess the significance of a candidate. Applying the needlet transform to texture templates (Eq. 4.1) of various sizes yields a lookup table specifying the needlet whose response is maximal at each texture size. This table can then be used to identify peaks in a filtered input map with a texture candidate of a certain size. To minimize the number of false detections, while not discarding potentially interesting signals, we determine a set of size-dependent thresholds (identical to those in Feeney et al. (2011a)) using an end-to-end simulation of the WMAP experiment (see Jarosik et al. (2011) and Gold et al. (2011)) containing a ΛCDM CMB as well as realistic foregrounds and systematics that we cannot include in our likelihood

function. The thresholds chosen restrict the number of candidate textures – by definition false detections – to be of order ten. All thresholds and parameters in the needlet transform are fixed at this point.

4.6 Sensitivity testing

To determine our ability to detect textures given our thresholds, we generate a set of CMB maps from the WMAP7 best-fit power spectrum, and place textures of varying ϵ and θ_c in both a region with low and high instrumental noise (as the noise properties of the WMAP experiment vary according to position on the sky). We find that, for $2° \leq \theta_c \leq 50°$, the significance threshold is certainly exceeded (and therefore a candidate identified) for $\epsilon > 10^{-4}$. For a favorable realization of the background CMB and instrumental noise, candidates are detected for $\epsilon > 2.5 \times 10^{-5}$ at scales $\theta_c \gtrsim 5°$ and $\epsilon > 5 \times 10^{-5}$ at somewhat smaller scales. We use $\epsilon = 2.5 \times 10^{-5}$ as a lower limit for detectable textures, and neglect the effect of θ_c on our candidate detection efficiency as it is far less important than the factor of θ_c^{-3} in Eq. 4.2.

4.7 Calculating the texture posterior probability

Once the candidate textures have been identified, the posterior probability distribution Eq. 4.3 can be calculated by first evaluating the patch-based evidence ratio Eq. 4.4 for each blob using the MultiNest (Feroz et al., 2009; Feroz and Hobson, 2008) nested sampling software. This requires calculating the inverse covariance matrix $\mathbf{C}_{b_i}^{-1}$, which is extremely memory-intensive at full WMAP resolution: the necessary storage capacity scales with size as θ_c^4. We therefore employ an adaptive-resolution analysis, processing each blob at the highest resolution possible given its size and the available computational resources. This removes the limitation on blob size of Feeney et al. (2011a) and Feeney et al. (2011b).

The only remaining quantity to evaluate in Eq. 4.3 is the prior $\Pr(\epsilon, \bar{N}_s)$. We choose a uniform prior for ϵ between $2.5 \times 10^{-5} \leq \epsilon \leq 1.0 \times 10^{-4}$. The lower bound is an estimate of what is detectable with our pipeline, determined by the simulations described above. The upper bound comes from requiring that the symmetry-breaking scale for textures, η, is below the scale of cosmological inflation. To be consistent with the lack of observed B-mode polarization in the CMB (Larson et al., 2011; Komatsu et al., 2011), the scale of inflation must be less than approximately 10^{16} GeV, constraining ϵ to be less than roughly 10^{-4} (this agrees with the prior of Cruz et al. (2007)).

We adopt a uniform prior on $\bar{N}_s$ between $0 \leq \bar{N}_s \leq 10$. The comoving density of textures

produced in a phase transition depends on the particular texture model in question, and can be determined from simulations. The total number of unwinding events is obtained by integrating this density over the four-volume swept out by the CMB photons. For example, simulations (Cruz et al., 2007) of SU(2) textures indicate that we can expect to have causal access to roughly 7 textures with $\theta_c > 2°$ in the CMB. The number of these unwinding events which are then *detectable* is mainly a function of our particular realization of the background CMB. Our choice of a uniform prior accounts for our ignorance of both the precise theory giving rise to textures and the precise number of detectable textures in the context of a specific theory. This allows us to compare ΛCDM and *all* models that give rise to textures. Under the assumption of a uniform prior the posterior is simply proportional to the likelihood; results for a different prior on $\bar{N}_s$ could be obtained easily by reweighting the current posterior.

The significance required to favor the ΛCDM+textures model can be understood by evaluating Eq. 4.3 using a set of hypothetical evidence ratios $\rho_{b_i}(\epsilon)$, assuming that two candidate textures have been located in the data. The evidence ratios are chosen to be either low-amplitude and flat in ϵ (the case where each blob yields no support for the texture model), or Gaussian with varying amplitude (indicating varying degrees of support for the texture model). In all cases, the Gaussians are chosen to peak at the same value, $\epsilon = 5 \times 10^{-5}$, and have the same standard deviation, $\sigma = 5 \times 10^{-6}$. The amplitudes of the Gaussian peaks are selected so that $\int \rho_{b_i}(\epsilon) \Pr(\epsilon) \, d\epsilon$ is 1/20, 1, or 20. These values are indicative of weak, intermediate, and strong texture signals, respectively. No maps or textures are simulated for this step: only the evidence ratios themselves.

The posteriors for all combinations of the hypothetical evidence ratios are shown in Fig. 4.1. When none of the candidate features support the texture hypothesis (top row), the posterior is exponentially decreasing in $\bar{N}_s$. In this case, we would correctly conclude that pure ΛCDM is strongly favored, and no constraints on ϵ could be extracted. When one or two blobs produce a peaked evidence ratio (central and bottom rows), it becomes possible to make a detection. As the amplitudes of the evidence ratios are increased (left to right), the posterior begins to bulge, before ultimately becoming peaked. We would correctly conclude that the data favor ΛCDM+textures over pure ΛCDM if a peak in the posterior at $\bar{N}_s \neq 0$ was sufficiently higher than the value of the posterior at $\bar{N}_s = 0$. Comparing the central and bottom rows of Fig. 4.1, a detection can be made either in the case where there is a single strong candidate, or the case where there is a number of moderately strong candidates (provided each $\rho_{b_i}(\epsilon)$ is peaked in the same range of ϵ). For the one- and two-strong-candidate cases (the centre- and bottom-right plots in Fig. 4.1), the peaks of the posterior (after marginalizing over ϵ) are 10 and 1500 times

Figure 4.1: Regions containing 68% (dark blue) and 95% (light blue) of the posterior probability distribution, Eq. 4.3, for the hypothetical situations described in the text. Each case contains two hypothetical texture candidates. Top row: both candidates are ΛCDM only, and have low-amplitude, flat evidence ratios, $\rho_{b_i}(\epsilon)$. Middle row: one candidate is "texture-like", and has a Gaussian evidence ratio whose amplitude increases from left to right. Bottom row: both candidates are texture-like; again, their evidence ratio amplitudes increase from left to right.

that of the value at $\bar{N}_s = 0$, respectively.

Calculating the posterior for the end-to-end simulation of the WMAP experiment yields the constraints shown as solid and dashed lines in Fig. 4.2. This posterior resembles the top row of Fig. 4.1, is peaked at $\bar{N}_s = 0$, and is not significantly different from the input priors on the global texture parameters. We therefore correctly conclude that the end-to-end simulation does not contain textures.

4.8 Results and conclusions from WMAP

We perform our analysis on the foreground-subtracted 94 GHz W-band temperature map from the 7-year release of the WMAP experiment (Larson et al., 2011) (prepared by subtracting a model of known astrophysical foregrounds, as described in Hinshaw et al. (2007)). The W band has the highest resolution of the five measured by WMAP, with a full-width at half maximum of $0.22°$. To minimize the effects of residual foregrounds, we apply the KQ75 mask, which yields a sky coverage of $f_{\text{sky}} = 0.706$. The candidate textures are the same as those identified in Feeney et al. (2011a), minus one which lies outside our prior on θ_c. The features range in size from $2°$ to

17.25°, and we are able to process seven at full WMAP resolution, two at half WMAP resolution and the largest at a quarter WMAP resolution. Although the lower-resolution computation of the likelihood for these three largest features does result in reduced accuracy, the impact on the overall posterior is minimal as the prior for textures of such large size is very low (cf. Eq. 4.2).

Evaluating Eq. 4.3 yields the posterior for cosmic textures in the WMAP7 data shown in Fig. 4.2 as dark- and light-blue regions. The posterior is clearly peaked at $\bar{N}_s = 0$, and we find the marginalized constraint on the expected number of detectable textures to be $\bar{N}_s < 5.9$ (at 95% confidence). We therefore conclude that the WMAP7 data do not warrant augmenting ΛCDM with textures. The marginalized constraint on the scale of symmetry breaking is found to be $2.6 \times 10^{-5} \leq \epsilon \leq 1.0 \times 10^{-4}$ (at 95% confidence).

While the posterior is peaked at $\bar{N}_s = 0$, there is also a clear difference between the WMAP7 posterior and that of the end-to-end simulation (over-plotted in Fig. 4.2). Comparing the WMAP7 posterior to the example plots in Fig. 4.1, our result is also consistent with a signal that is present, but too weak to provide a detection. The different shape of the posterior is determined almost entirely by two features, located at $(l = 185°, b = -79°)$ and $(l = 209°, b = -57°)$ in Galactic coordinates, the second of which is the Cold Spot (Cruz et al., 2005, 2010). As in Feeney et al. (2011a), we use information from the multiple frequency bands of the WMAP instrument to confirm that there is no detectable residual foreground contamination in these features. This strongly motivates an analysis with better data, as will soon be provided by the Planck satellite (Planck Collaboration et al., 2011), or a better candidate-location technique, such as one utilizing optimal filters (McEwen et al., 2012). There is also the possibility of including CMB polarization data, as textures would not induce a polarization signal, unlike the primary CMB perturbations (Vielva et al., 2011)). All of these efforts are currently in progress. These and other tests will lead to better constraints on – or, if a signal is present, a confirmation of – the texture hypothesis.

4.9 Acknowledgements

Figure 4.2: Regions containing 68% (dark blue) and 95% (light blue) of the posterior probability distribution, Eq. 4.3, for the WMAP7 data, along with the corresponding contours of the 68% (solid line) and 95% (dashed line) of the posterior probability for the end-to-end simulation of the WMAP experiment, based on a pure ΛCDM model.

Chapter 5

Hierarchical Bayesian detection algorithm for early-universe relics in the cosmic microwave background

5.1 Abstract

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 are constrained to be fewer than 4.0 and 5.2 at 95% confidence, respectively.

5.2 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 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 (Turok, 1989). Textures are not stable, but instead undergo collapse as they come within the expanding cosmological horizon, eventually unwinding into scalar radiation (Turok, 1989; Turok and Spergel, 1991; Spergel et al., 1991; Pen et al., 1994; Turok and Spergel, 1990). 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 (Turok and Spergel, 1990). 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., Aguirre (2008); Guth (2007)). 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 (Aguirre et al., 2007).

In previous work, we presented the first constraints on theories giving rise to cosmic bubble

collisions (Feeney et al., 2011a,b) and the first full-sky constraints on cosmic textures (Feeney et al., 2012). The present paper focuses on:

- Generalizing the statistical formalism and approximation scheme used in Feeney et al. (2011a,b) and Feeney et al. (2012);

- Implementing an adaptive-resolution analysis, allowing us to overcome the limitations in Feeney et al. (2011a,b) on the size of candidate bubble collisions;

- Including and refining the candidate detection scheme using optimal filters presented in McEwen et al. (2012);

- 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. 5.3, we outline the formalism and approximations we use. In Sec. 5.4, we describe the theoretical predictions for cosmic textures and bubble collisions. The algorithm used to calculate the approximated posterior is described in Secs. 5.5-5.6 and tested in Secs. 5.7 and 5.8. A null test of the pipeline is carried out in Sec. 5.9 before the pipeline is applied to CMB data from the Wilkinson Microwave Anisotropy Probe (WMAP) (Bennett et al., 2003a) in Sec. 5.10. The results of this analysis are compared with previous analyses in Sec. 5.11, and our conclusions are summarised in Sec. 5.12.

5.3 Hierarchical Bayesian source detection formalism

5.3.1 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 (Sunyaev and Zeldovich, 1972)), 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

$$\frac{\Delta T}{T}(\theta, \phi) = \delta(\theta, \phi) + n(\theta, \phi) + \sum_{i=1}^{N_{\rm s}} t_i(\theta, \phi). \tag{5.1}$$

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. 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. Loredo (2012)) 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

---

[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.

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

$$C_{ij} \equiv C(\theta_{ij}) = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell(\mathbf{m}_{\Lambda\text{CDM}}) P_\ell(\cos(\theta_{ij})), \tag{5.2}$$

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.

5.3.2 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 $\Omega_{\text{obs}} = 4\pi f_{\text{sky}}$ of the sky (and, optionally, any statistics derived from them). Bayes' theorem gives the posterior as

$$\Pr(\mathbf{m}_0|\mathbf{d}, f_{\text{sky}}) = \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|\mathbf{m}_0, f_{\text{sky}})}{\Pr(\mathbf{d}|f_{\text{sky}})}, \tag{5.3}$$

where $\Pr(\mathbf{m}_0)$ is the prior distribution on the global parameters, $\mathbf{m}_0$, $\Pr(\mathbf{d}|\mathbf{m}_0, f_{\text{sky}})$ is the likelihood of getting the observed data, and $\Pr(\mathbf{d}|f_{\text{sky}})$ ensures that the posterior is normalized.

---

[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 Freivogel (2011) and Salem (2012) for recent reviews).

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 $\mathrm{Pr}(\mathbf{d}|\mathbf{m}_0, f_\mathrm{sky})$ 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:

$$
\begin{aligned}
\mathrm{Pr}(\mathbf{d}|\mathbf{m}_0, f_\mathrm{sky}) \;=\; & \sum_{N_\mathrm{s}=0}^{\infty} \frac{(f_\mathrm{sky}\bar{N}_\mathrm{s})^{N_\mathrm{s}} e^{-f_\mathrm{sky}\bar{N}_\mathrm{s}}}{N_\mathrm{s}!} \\
& \times \int \mathrm{d}\mathbf{m}_{\Lambda\mathrm{CDM}} \mathrm{Pr}(\mathbf{m}_{\Lambda\mathrm{CDM}}) \int \mathrm{d}\mathbf{m}_1 \ldots \mathrm{d}\mathbf{m}_{N_\mathrm{s}} \mathrm{Pr}(\mathbf{m}_1, \ldots \mathbf{m}_{N_\mathrm{s}}) \\
& \times \mathrm{Pr}(\mathbf{d}|N_\mathrm{s}, f_\mathrm{sky}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots \mathbf{m}_{N_\mathrm{s}}),
\end{aligned}
\tag{5.4}
$$

where $\mathrm{Pr}(\mathbf{m}_{\Lambda\mathrm{CDM}})$ is the prior over the parameters of ΛCDM, $\mathrm{Pr}(\mathbf{m}_1, \ldots \mathbf{m}_{N_\mathrm{s}})$ is the prior over the local model parameters for N_s sources, and $\mathrm{Pr}(\mathbf{d}|N_\mathrm{s}, f_\mathrm{sky}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots \mathbf{m}_{N_\mathrm{s}})$ 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

$$
\mathrm{Pr}(\mathbf{d}|N_\mathrm{s}, f_\mathrm{sky}, \mathbf{m}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots \mathbf{m}_{N_\mathrm{s}}) = \frac{e^{-(\mathbf{d}-\sum_{i=1}^{N_\mathrm{s}}\mathbf{t}(\mathbf{m}_i))\mathsf{C}^{-1}(\mathbf{d}-\sum_{i=1}^{N_\mathrm{s}}\mathbf{t}(\mathbf{m}_i))^\mathrm{T}/2}}{(2\pi)^{n_\mathrm{pix}/2}|\mathsf{C}|},
\tag{5.5}
$$

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

One must evaluate Eq. 5.5 in order to construct the full posterior, Eq. 5.3. Evaluating this expression is impossible, as it requires marginalizing over a formally infinite dimensional parameter space. Even if the parameter space were finite, the enormous size of modern CMB datasets, such as those produced by the WMAP or Planck (Planck Collaboration et al., 2011) 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.

5.3.3 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

$$\Pr(\mathbf{d}|\mathbf{m}_0, f_{\mathrm{sky}}) \simeq \sum_{N_{\mathrm{s}}=0}^{\infty} \frac{(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^{N_{\mathrm{s}}} e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}}{N_{\mathrm{s}}!} \Pr(\mathbf{d}|N_{\mathrm{s}}), \tag{5.6}$$

where we define the helpful short-hand

$$\Pr(\mathbf{d}|N_{\mathrm{s}}) \equiv \int d\mathbf{m}_1 \ldots d\mathbf{m}_{N_{\mathrm{s}}} \Pr(\mathbf{m}_1, \ldots \mathbf{m}_{N_{\mathrm{s}}}) \Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}, \bar{\mathbf{m}}_{\Lambda\mathrm{CDM}}, \mathbf{m}_0, \mathbf{m}_1, \ldots \mathbf{m}_{N_{\mathrm{s}}}). \tag{5.7}$$

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. 5.1. 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. 5.6 term by term. The likelihood in the first term, for $N_{\mathrm{s}} = 0$, is simply given by

$$\Pr(\mathbf{d}|N_{\mathrm{s}} = 0) = \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2}, \tag{5.8}$$

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 $\bar{b}$:

$$\begin{aligned}
\Pr(\mathbf{d}|N_{\mathrm{s}} = 1) = &\int_b d\mathbf{m}_1 \Pr(\mathbf{m}_1) \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2} \\
&+ \int_{\bar{b}} d\mathbf{m}_1 \Pr(\mathbf{m}_1) \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}.
\end{aligned} \tag{5.9}$$

Figure 5.1: A schematic depicting the approximation scheme we employ. By locating a set of candidate sources, it is possible to determine in which regions of parameter space the likelihood function is appreciably different from zero. Eq. 5.7 can be approximated by integrating over only those regions where the likelihood is large. This in addition collapses the sum in Eq. 5.6 to a finite number of terms. Finally, we neglect correlations of the random Gaussian background CMB between pixels inside and outside each blob.

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

$$\Pr(\mathbf{d}|N_\mathrm{s} = 1) \quad \simeq \quad \int_b d\mathbf{m}_1 \Pr(\mathbf{m}_1) \frac{1}{(2\pi)^{n_\mathrm{pix}/2}|\mathsf{C}|} e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^\mathrm{T}/2}. \tag{5.10}$$

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

$$\frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^\mathrm{T}/2}}{(2\pi)^{n_\mathrm{pix}/2}|\mathsf{C}|} = \frac{1}{(2\pi)^{n_\mathrm{pix}/2}|\mathsf{C}|} e^{-(d_i-t_i)C_{ij}^{-1}(d_j-t_j)/2} \times e^{-d_\alpha C_{\alpha\beta}^{-1}d_\beta/2} \times e^{-(d_i-t_i)C_{i\alpha}^{-1}d_\alpha} \tag{5.11}$$

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

$$
\begin{aligned}
\frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} \;=\;& \frac{e^{-(d_i-t_i)C_{ij}^{-1}(d_j-t_j)/2}}{e^{-d_i C_{ij}^{-1} d_j/2}} \times \left(\frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2} \right) \times \left(e^{t_i C_{i\alpha}^{-1} d_\alpha} \right) \\
\simeq\;& \frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)^{-1}}\mathbf{d}^{\mathrm{T}}/2}} \times \left(\frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2} \right) \times \left(e^{t_i C_{i\alpha}^{-1} d_\alpha} \right) \\
\simeq\;& \frac{e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)^{-1}}\mathbf{d}^{\mathrm{T}}/2}} \times \Pr(\mathbf{d}|N_{\mathrm{s}}=0).
\end{aligned}
\tag{5.12}
$$

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

$$
e^{2t_i C_{i\alpha}^{-1} d_\alpha^{\mathrm{T}}} \simeq 1.
\tag{5.13}
$$

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

$$
(d_i - t_i)C_{ij}^{-1}(d_j - t_j) \rightarrow (\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2,
\tag{5.14}
$$

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 5.8.

Finally, performing the integral over $\mathbf{m}_1$ in Eq. 5.10, we obtain

$$
\Pr(\mathbf{d}|N_{\mathrm{s}}=1) \simeq \Pr(\mathbf{d}|N_{\mathrm{s}}=0)\,\rho_b(\mathbf{m}_0),
\tag{5.15}
$$

where the patch-based evidence ratio ρ_b is given by

$$
\rho_b(\mathbf{m}_0) \equiv \frac{\int_b \mathrm{d}\mathbf{m}_1 \Pr(\mathbf{m}_1) e^{-(\mathbf{d}-\mathbf{t})\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)^{-1}}\mathbf{d}^{\mathrm{T}}/2}}.
\tag{5.16}
$$

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 two-source term. Again, we approximate the full integral as the integral

over region b alone. This yields

$$\Pr(\mathbf{d}|N_{\mathrm{s}}=2) = \int_b \int_b \mathrm{d}\mathbf{m}_1 \mathrm{d}\mathbf{m}_2 \Pr(\mathbf{m}_1)\Pr(\mathbf{m}_2)\frac{e^{-(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))\mathsf{C}^{-1}(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))^{\mathrm{T}}/2}}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|}. \quad (5.17)$$

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

$$\begin{aligned}
\Pr(\mathbf{d}|N_{\mathrm{s}}=2) \quad\simeq\quad & \Pr(\mathbf{d}|N_{\mathrm{s}}=0) \quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad (5.18)\\
& \times\frac{\int_b \int_b \mathrm{d}\mathbf{m}_1 \mathrm{d}\mathbf{m}_2 \Pr(\mathbf{m}_1)\Pr(\mathbf{m}_2)e^{-(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))\mathsf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t}(\mathbf{m}_1)-\mathbf{t}(\mathbf{m}_2))^{\mathrm{T}}/2}}{e^{-\mathbf{d}\mathsf{C}^{(bb)^{-1}}\mathbf{d}^{\mathrm{T}}/2}}.
\end{aligned}$$

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 $\mathbf{m}_2$). This will introduce a natural "Occam factor" that favors the simpler model (i.e., the model with one template). 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

$$\Pr(\mathbf{m}_i) = \frac{1}{4\pi}. \quad (5.19)$$

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 $N_{\mathrm{s}}=1$ and $N_{\mathrm{s}}=2$ terms can be estimated as

$$\frac{\Pr(\mathbf{d}|N_{\mathrm{s}}=1)}{\Pr(\mathbf{d}|N_{\mathrm{s}}=2)} \simeq \frac{4\pi}{\Omega_b}. \quad (5.20)$$

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

For a single blob, we can therefore approximate the full-sky posterior Eq. 5.3 as

$$\Pr(\mathbf{m}_0|\mathbf{d}, f_{\mathrm{sky}}) \simeq \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|N_{\mathrm{s}}=0)\,e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}\left(1+(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})\,\rho_b(\mathbf{m}_0)\right)}{\Pr(\mathbf{d}|f_{\mathrm{sky}})}, \quad (5.21)$$

where

$$\Pr(\mathbf{d}|f_{\mathrm{sky}}) \equiv \int \mathrm{d}\mathbf{m}_0 \Pr(\mathbf{m}_0)e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}\left(1+(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})\,\rho_b(\mathbf{m}_0)\right) \quad (5.22)$$

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

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

$$
\Pr(\mathbf{d}|N_{\rm s}) = \begin{cases} 0, & \text{if} \quad N_{\rm s} > N_{\rm b}, \\[2em] \Pr(\mathbf{d}|N_{\rm s}=0) \sum_{b_1,b_2,\ldots,b_{N_{\rm s}}=1}^{N_{\rm b}} \Delta^{b_1 b_2 \ldots b_{N_{\rm s}}} \prod_{s=1}^{N_{\rm s}} \rho_{b_s}(\mathbf{m}_0), & \text{if} \quad N_{\rm s} \leq N_{\rm b}. \end{cases} \tag{5.23}
$$

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.

Substituting Eq. 5.23 into Eq. 5.6, the expression for the approximation to the full posterior is given by

$$
\begin{aligned} \Pr(\mathbf{m}_0|\mathbf{d}, f_{\rm sky}) \quad &\simeq \quad \frac{\Pr(\mathbf{m}_0)\Pr(\mathbf{d}|N_{\rm s}=0)}{\Pr(\mathbf{d}|f_{\rm sky})} e^{-f_{\rm sky}\bar{N}_{\rm s}} \\ &\times \sum_{N_{\rm s}=0}^{N_{\rm b}} \frac{(f_{\rm sky}\bar{N}_{\rm s})^{N_{\rm s}}}{N_{\rm s}!} \sum_{b_1,b_2,\ldots,b_{N_{\rm s}}=1}^{N_{\rm b}} \Delta^{b_1 b_2 \ldots b_{N_{\rm s}}} \prod_{s=1}^{N_{\rm s}} \rho_{b_s}(\mathbf{m}_0). \end{aligned} \tag{5.24}
$$

Eq. 5.24 is the main result of this calculation, from which all following results are derived. In the limit of a single isolated observation, Eq. 5.24 reproduces the Bayesian source detection formalism developed in Hobson and McLachlan (2003) and Hobson et al. (2010).

5.4   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 $\bar{N}_{\rm s}$, 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 Feeney et al. (2011a,b), where, in addition to the location, three local model parameters (size, edge discontinuity, and

amplitude) were included. Cosmic textures have two global parameters: $\bar{N}_{\mathrm{s}}$ 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 Feeney et al. (2012). 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 Cruz et al. (2005), Cruz et al. (2007), Cruz et al. (2008), Cruz et al. (2010) and Vielva et al. (2011). 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 Feeney et al. (2011a,b) in light of improved theoretical understanding, and remove the edge discontinuity parameter from our analysis.

5.4.1 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 Aguirre (2008) and Aguirre and Johnson (2009). For a detailed discussion of the expected signature of bubble collisions in the CMB, we refer the reader to Chang et al. (2009), Feeney et al. (2011a), Czech et al. (2010), Gobbetti and Kleban (2012) and Kleban et al. (2011); 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 (Chang et al., 2009; Feeney et al., 2011a,b)

$$t(\theta, \phi) = \left(\frac{z_{\mathrm{crit}} - z_0 \cos \theta_{\mathrm{crit}}}{1 - \cos \theta_{\mathrm{crit}}} + \frac{z_0 - z_{\mathrm{crit}}}{1 - \cos \theta_{\mathrm{crit}}} \cos \theta \right) \Theta(\theta_{\mathrm{crit}} - \theta), \qquad (5.25)$$

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. 5.1. Theoretical work (Gobbetti and Kleban, 2012; Kleban et al., 2011) 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 $z_{\mathrm{crit}} = 0$.

The bubble collision model contains only one global parameter, the expected number of detectable sources $\bar{N}_{\mathrm{s}}$. 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, $\bar{N}_\mathrm{s}$ 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 $\bar{N}_\mathrm{s} \gg 1$, in which case the approximation of looking for widely separated sources is not valid; see Kozaczuk and Aguirre (2012) for a discussion of the observational signatures in this case. In the absence of a detailed theory, we assume that $\bar{N}_\mathrm{s}$ can take values of $\mathcal{O}(1)$, and set $\bar{N}_\mathrm{s}$ to be uniform in the range $0 \leq \bar{N}_\mathrm{s} \leq 10$; 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 $\bar{N}_\mathrm{s} = 0$.

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 $\sin \theta_\mathrm{crit}$ (Aguirre et al., 2007; Aguirre and Johnson, 2008; Freivogel et al., 2009; Aguirre and Johnson, 2009). Further assuming no correlation between the various local parameters, the final *normalized* prior on the local parameters is

$$ \Pr(\mathbf{m}_1) = \Pr(z_0) \Pr(\theta_0, \phi_0) \Pr(\theta_\mathrm{crit}) = \frac{1}{2 \times 10^{-4}} \left(\frac{\sin \theta_0}{4\pi} \right) \left(\frac{\sin \theta_\mathrm{crit}}{\cos \theta_\mathrm{crit}^\mathrm{min} - \cos \theta_\mathrm{crit}^\mathrm{max}} \right), \qquad (5.26) $$

where $0 \leq \theta_0 \leq \pi$, $0 \leq \phi_0 < 2\pi$, and the extrema of the size distribution, $\{\theta_\mathrm{crit}^\mathrm{min}, \theta_\mathrm{crit}^\mathrm{max}\}$ are chosen such that the bubble collisions are detectable. The lower limit, $\theta_\mathrm{crit}^\mathrm{min} = 2°$, 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, $\theta_\mathrm{crit}^\mathrm{max} = 90°$, would be indistinguishable from the signature of a collision of size $180° - \theta_\mathrm{crit}$.

5.4.2 Cosmic textures

For a detailed discussion of the production, evolution, and observational signature of cosmic textures we refer the reader to the original literature (Cruz et al., 2007; Pen et al., 1994; Spergel et al., 1991; Turok, 1989; Turok and Spergel, 1990, 1991). 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

$$t(\theta,\phi) = \frac{(-1)^p \, \epsilon}{\sqrt{1 + 4\left(\frac{\theta}{\theta_{\mathrm{crit}}}\right)^2}}, \quad \theta \leq \theta_*; \quad t(\theta,\phi) = \frac{(-1)^p \epsilon}{2} \exp\left(-\frac{(\theta^2 - \theta_*^2)}{2\theta_{\mathrm{crit}}^2}\right), \quad \theta > \theta_*; \quad (5.27)$$

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, $\theta_* = \sqrt{3}\theta_{\mathrm{crit}}/2$, $p = \{0,1\}$ 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 $\bar{N}_{\mathrm{s}}$. 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 (Cruz et al., 2007) of SU(2) textures indicate that we can expect to have causal access to ~ 7 textures with $\theta_{\mathrm{crit}} > 2°$ 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 $\bar{N}_{\mathrm{s}}$ 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. 5.5). 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 $\theta_{\mathrm{crit}}^{-3}$ (see Cruz et al. (2007) for a derivation), and so the normalized prior on the local texture parameters is

$$\Pr(\mathbf{m}_1) = \frac{\sin\theta_0}{4\pi\theta_{\mathrm{crit}}^3}\left(\frac{1}{(2°)^2} - \frac{1}{(50°)^2}\right)^{-1}, \quad (5.28)$$

where $0 \leq \theta_0 \leq \pi$, $0 \leq \phi_0 < 2\pi$, and we take $2° \leq \theta_{\mathrm{crit}} \leq 50°$. 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 $\theta_{\mathrm{crit}} > 50°$ are large enough to cover the whole sky and overlap themselves,

rendering Eq. 5.27 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. 5.3. 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. 5.5 below. In the second step, the patch-based evidence ratios, Eq. 5.16, are calculated for each candidate using the nested sampler `MultiNest` (Feroz and Hobson, 2008; Feroz et al., 2009), before combining to form the posterior, Eq. 5.24, on the global model parameters. We describe this part of the calculation in Sec. 5.6.

5.5 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 McEwen et al. (2012), 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.

5.5.1 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 (Larson et al., 2011) (hereafter referred to as the lensed WMAP7+BAO+H0 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

(a) $\theta_{\text{crit}} = 5°$ (b) $\theta_{\text{crit}} = 20°$ (c) $\theta_{\text{crit}} = 60°$

Figure 5.2: Matched filters optimized to bubble collision signatures of varying size embedded in a ΛCDM CMB background.

are made in the presence of a Gaussian beam of full-width-half-maximum FWHM = 13.2 arcminutes and isotropic white noise of $N_\ell = 0.02$ μK^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 McEwen et al. (2008).

In Fig. 5.2 and Fig. 5.3 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 McEwen et al. (2012). In Fig. 5.4 we plot the SNR computed for bubble collision and texture profiles for the unfiltered field, for the optimal filters constructed here, and for needlets (Marinucci et al., 2008; Scodeller et al., 2011), which have been used previously to detect candidate sources (Feeney et al., 2011a,b; Feeney et al., 2012). Note that the lack of a sharp transition in the texture template means that the SNR for textures are lower than those of bubble colli-

(a) $\theta_{\mathrm{crit}} = 5°$ (b) $\theta_{\mathrm{crit}} = 10°$ (c) $\theta_{\mathrm{crit}} = 20°$

Figure 5.3: Matched filters optimized to cosmic texture signatures of varying size embedded in a ΛCDM CMB background.

sions. Nevertheless, it is clear that matched filters yield the highest SNR, in accordance with expectations.

5.5.2   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 McEwen et al. (2012) 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. McEwen et al. (2012)). 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 $\{R_{\mathrm{adj}} \in \{\theta_{\mathrm{crit}}^k\}_{k=1}^{N_{\theta_{\mathrm{crit}}}} : |R_{\mathrm{adj}} - R| \le \theta_{\mathrm{adj}}\}$, 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 (McEwen et al., 2008) (which in turn relies on the codes S2 (McEwen et al., 2007) and HEALPix (Górski et al., 2005)), while the COMB code (McEwen et al., 2008) has been used to simulate bubble collision signatures

embedded in a CMB background.[3] 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 (McEwen et al., 2012). For texture profiles, however, we have found this phenomenon to hold below scales $\theta_{\rm crit} \sim 10°$ 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. 5.4), 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 $\theta_{\rm crit} = 10°$ we look across adjacent scales as usual to find the most significant potential candidate source, whereas for sizes above $\theta_{\rm crit} = 10°$ we do not (by a judicious choice of the adjacency parameter $\theta_{\rm adj}$ 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. 5.4 (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 $\theta_{\rm crit}$ grid is sampled sufficiently densely, the matched filters for both bubble collision and texture profiles remain effective and are superior to needlets.

5.5.3 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 (Gold et al., 2011), 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 (Feeney et al., 2011c). This choice has the additional advantage of revealing more of the sky and hence

---

[3]S2FIL, S2 and COMB are available from http://www.jasonmcewen.org/, while HEALPix is available from http://healpix.jpl.nasa.gov/.

(a) Bubble collision signatures of known source size

(b) Bubble collision signatures of unknown source size

(c) Cosmic texture signatures of known source size

(d) Cosmic textures of unknown source size

Figure 5.4: 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 (Marinucci et al., 2008) 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.

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 McEwen et al. (2012), as defined in Table 5.1 (left). The θ_{crit} prior range is smaller for textures than for bubble collisions – the texture profile extends well beyond θ_{crit}, covering the full sky for $\theta_{\mathrm{crit}} \gtrsim 50°$ – and hence for textures we consider a smaller grid of scales, also defined in Table 5.1 (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. 5.4 (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 $\theta_{\mathrm{adj}} = 5°$ for both bubble collisions and textures. For textures this ensures that we look across scales for sizes below $\theta_{\mathrm{crit}} = 10°$ 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 (Feeney et al., 2011c), where we adopt a Gaussian prior for the harmonic coefficients specified by the lensed WMAP7+BAO+H0 power spectrum. Note that this differs from the maximum likelihood reconstruction (de Oliveira-Costa and Tegmark, 2006) of harmonic coefficients performed by Feeney et al. (2011a,b) and Feeney et al. (2012) and the cut-sky estimation performed by McEwen et al. (2012). 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 (Feeney et al., 2011c).

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 (Gold et al., 2011). 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

$\theta_{\mathrm{crit}}(°)$	N_{σ_R}		$\theta_{\mathrm{crit}}(°)$	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 5.1: The θ_{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.

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 Feeney et al. (2011a,b), Feeney et al. (2012) and McEwen et al. (2012), 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 5.1. 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 (McEwen et al., 2012). 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 (Feeney et al., 2011a,b; McEwen et al., 2012). Here we instead take a probabilistic approach based on the analytic SNRs of the filters computed previously (see Fig. 5.4). 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. 5.5 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 5.1. This level defines the 50% completeness curves shown in Fig. 5.5 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 (Feeney et al., 2011a,b; McEwen et al., 2012), 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 (Feeney et al., 2011a,b; McEwen et al., 2012).

5.6 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 $\sim 5'$ (Planck Collaboration et al., 2011), resulting in maps pixelated on the arcminute scale. While this is necessary for pinning down the secondary CMB anisotropies at small scales, it means that calculating pixel-space covariance

(a) Bubble collisions	(b) Cosmic textures

Figure 5.5: 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 5.1. 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.

matrices becomes extremely memory-intensive. We illustrate this point in Fig. 5.6, which shows the memory needed to calculate covariance matrices from $1°$ to $180°$ in radius at `HEALPix` resolutions ranging from $N_{\mathrm{side}} = 8$ to $N_{\mathrm{side}} = 2048$ (i.e., Planck resolution).[4] 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 (Feeney et al., 2011a,b), 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 $N_{\mathrm{side}} = 64$: 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.

---

[4]The quantity plotted corresponds to a total number of matrix elements equal to $\sim 1.5 \times N_{\mathrm{pix}}^2$. Our algorithm calculates the full N_{pix}^2 covariance matrix in order to make use of the `LAPACK` inversion routines (Anderson et al., 1999), then compresses the inverted matrix to a 1-D array containing its upper triangle to reduce memory costs while sampling.

Figure 5.6: The memory needed to calculate the covariance matrix for patches of a given size at `HEALPix` resolutions ranging from $N_{\text{side}} = 8$ (lightest-blue, lowest curve) to $N_{\text{side}} = 2048$ (highest, darkest curve). The effects of pixel size are visible at small patch radii and low resolutions.

5.6.1 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

$$n_{\text{deg}} = \text{ceiling}\left(\frac{\log\left(m_{\text{est}}/m_{\text{max}}\right)}{\log 16}\right), \tag{5.29}$$

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

$$N_{\text{side,deg}} = \frac{N_{\text{side,full}}}{2^{n_{\text{deg}}}}. \tag{5.30}$$

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

$$f_{\mathrm{deg}} = \sqrt{\frac{8 \log 2 \, \log 100}{\ell_{\mathrm{max}}(\ell_{\mathrm{max}} + 1)}}. \tag{5.31}$$

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 5.2. 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.

5.6.2 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,

$N_{\mathrm{side,deg}}$	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 5.2: 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.

C, into its CMB, S, and noise, N, constituents, as they are morphologically different.

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. 5.7, 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 $N_{\mathrm{side,deg}} = 64$.

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

$$S_{ij} = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell P_\ell(\cos\theta_{ij}) B_{\ell,\mathrm{WMAP}}^2 \qquad (5.32)$$

if the patch is processed at full resolution, and

$$S_{ij} = \sum_\ell \frac{2\ell + 1}{4\pi} C_\ell P_\ell(\cos\theta_{ij}) B_{\ell,\mathrm{WMAP}}^2 B_{\ell,\mathrm{deg}}^2 W_{\ell,\mathrm{deg}}^2 \qquad (5.33)$$

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

$$B_\ell = \exp\left(\frac{-\ell(\ell + 1)f^2}{16\log 2}\right), \qquad (5.34)$$

where f is the relevant FWHM.

Figure 5.7: 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 $N_{\mathrm{side,deg}} = 128$ (dashed), and the pixel window function at this resolution (dotted). Note that the pixel window function is only defined up to $\ell = 4N_{\mathrm{side,deg}}$, the maximum multipole allowed by the `HEALPix` software.

Instrumental noise covariance

At full resolution, the WMAP noise is uncorrelated in pixel space (see, e.g., Hinshaw et al. (2003)). The noise covariance matrix is therefore diagonal, and can be written as

$$N_{ij} = \frac{\sigma_{\mathrm{WMAP}}^2}{N_{\mathrm{obs},i}} \delta_{ij}, \tag{5.35}$$

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

$$N_\ell = \frac{\sigma_{\mathrm{WMAP}}^2 \Omega_{\mathrm{pix}}}{\bar{N}_{\mathrm{obs}}}, \tag{5.36}$$

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

$$N_{ij} = \frac{\bar{N}_{\mathrm{obs}}}{\sqrt{N_{\mathrm{obs},i} N_{\mathrm{obs},j}}} \sum_\ell \frac{2\ell + 1}{4\pi} N_\ell P_\ell(\cos\theta_{ij}) B_{\ell,\mathrm{deg}}^2 W_{\ell,\mathrm{deg}}^2, \tag{5.37}$$

where the separation into components residing in pixel and harmonic space is clear. Note that

Figure 5.8: 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 $\theta_{\mathrm{crit}} = 7°$. The power spectra are plotted at full WMAP resolution, $N_{\mathrm{side}} = 512$ (left), and after smoothing and degradation to $N_{\mathrm{side,deg}} = 128$ (right).

this expression does not include the instrumental beam, as this is *detector* noise.

Full covariance

The sum over ℓ in the covariance matrix calculations is strictly a sum to infinity, but is approximated with a sum from $0 \leq \ell \leq \ell_{\mathrm{max,cov}}$, truncated at a point where all significant contributions have been included. At full resolution, the data are noise-dominated at high-ℓ (see Fig. 5.8), so the covariance is not particularly sensitive to the precise value of $\ell_{\mathrm{max,cov}}$, provided it is larger than ~ 1000. We use $\ell_{\mathrm{max,cov}} = 1024$.[5] 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 $\ell_{\mathrm{max}} = 2N_{\mathrm{side,deg}}$, 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 $2N_{\mathrm{side,deg}}$ yields over-estimates of ~ 5 in the log-evidence at $N_{\mathrm{side,deg}} = 64$. Tests reveal that evidence ratios are stable provided multipoles $\ell \gtrsim 3N_{\mathrm{side,deg}}$ are included in the covariance matrix. We choose to be conservative and push this to $\ell_{\mathrm{max,cov}} = 4N_{\mathrm{side,deg}}$.

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

$$\mathsf{C} = \mathsf{S} + \mathsf{N}. \tag{5.38}$$

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

---

[5]The WMAP7 observations (Larson et al., 2011) 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 $\ell_{\mathrm{max,cov}} = 1024$ ensures the CMB contribution is characterized well into the noise-dominated regime.

covariance matrix has a prominent diagonal due to the uncorrelated pixel noise, and hence is readily invertible. However, at lower resolutions the additional smoothing can make the covariance matrix nearly singular. Diagonal regularizing noise is therefore added at the level of $2\,\mu\mathrm{K}^2$ to all degraded-resolution covariance matrices to aid inversion: we have checked that this does not affect the results of the likelihood calculation.

5.6.3   Calculation of patch-based evidence ratios with MultiNest

To form the full posterior (Eq. 5.24), we must calculate the patch-based evidence ratios (Eq. 5.16) for each blob. Our pipeline uses the `MultiNest` sampler (Feroz and Hobson, 2008; Feroz et al., 2009), 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. 5.16, 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

$$\mathrm{Pr}_b(\mathbf{m}_0, \mathbf{m}_1|\mathbf{d}) = \frac{\mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}_b(\mathbf{m}_1)\mathrm{Pr}_b(\mathbf{d}|\mathbf{m}_0, \mathbf{m}_1)}{\mathrm{Pr}_b(\mathbf{d})}. \tag{5.39}$$

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

$$\mathrm{Pr}_b(\mathbf{d}|\mathbf{m}_0, \mathbf{m}_1) \propto e^{-(\mathbf{d}-\mathbf{t})\mathbf{C}^{(bb)^{-1}}(\mathbf{d}-\mathbf{t})^{\mathrm{T}}/2} \tag{5.40}$$

as the quantity necessary for the patch-based evidence ratio. The normalization assumed in Eq. 5.39 is

$$\int_b \mathrm{d}\mathbf{m}_0\mathrm{d}\mathbf{m}_1\mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}_b(\mathbf{m}_1) = 1, \tag{5.41}$$

which implies that $\mathrm{Pr}_b(\mathbf{m}_1)$ is related to the prior $\mathrm{Pr}(\mathbf{m}_1)$ in Eq. 5.16 by an overall normalization:

$$F_b = \frac{\mathrm{Pr}(\mathbf{m}_1)}{\mathrm{Pr}_b(\mathbf{m}_1)} = 1 - \int_{\bar{b}} \mathrm{d}\mathbf{m}_0\mathrm{d}\mathbf{m}_1\mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}(\mathbf{m}_1). \tag{5.42}$$

We also have

$$\mathrm{Pr}_b(\mathbf{d}) = \int_b \mathrm{d}\mathbf{m}_0\mathrm{d}\mathbf{m}_1\mathrm{Pr}(\mathbf{m}_0)\mathrm{Pr}_b(\mathbf{m}_1)\mathrm{Pr}_b(\mathbf{d}|\mathbf{m}_0, \mathbf{m}_1). \tag{5.43}$$

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

$$\rho_b(\mathbf{m}_0) = F_b \times \frac{\mathrm{Pr}_b(\mathbf{d})}{\mathrm{Pr}_b(\mathbf{d}|0,0)} \times \frac{1}{\mathrm{Pr}(\mathbf{m}_0)}\int_b \mathrm{d}\mathbf{m}_1\mathrm{Pr}_b(\mathbf{m}_0, \mathbf{m}_1|\mathbf{d}), \tag{5.44}$$

where $\text{Pr}_b(\mathbf{d}|0,0)$ 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.

5.7 Adaptive-resolution tests

5.7.1 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, $N_{\text{side}} = 64$. 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 $N_{\text{side}} = 512$, 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. 5.9. 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

7° Bubble Collision	15° Bubble Collision	30° Bubble Collision

-504 ⬛ 655 μK -504 ⬛ 655 μK -504 ⬛ 689 μK

Figure 5.9: 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 located 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.

– 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.[6]

The results of the tests are shown in Table 5.3. 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 Skilling (2004) and Feroz and Hobson (2008)). We find these estimates (typically ~ 0.1 in log-evidence) to be subdominant to the variations due to sampling.

5.7.2 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

---

[6]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.

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 5.3: 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.

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.[7] While it is possible to mitigate this effect by extending the mask (and potentially using a smoothing kernel that is localized in pixel space) (Feeney et al., 2011c), 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 (Gold et al., 2011) 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 Bennett et al. (2003b)), 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 (Gold et al., 2011), 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., $N_{\mathrm{side}} = \{256, 128, 64\}$), and masked using the degraded point-source mask at each resolution. The resulting maps, plotted in Fig. 5.10, can then be scaled to mimic a

---

[7]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.

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 $100 - 1000\,\mu\mathrm{K}$ (Wright et al., 2009), these results suggest that our degradation technique induces contaminants of at most $1 - 2\,\mu\mathrm{K}$ 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 Pontzen and Peiris (2010), 1% of the difference between the V-band signal and the WMAP7 Internal Linear Combination (ILC) map (Gold et al., 2011) 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_{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. 5.11, 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 $0.3 - 0.4\,\mu\mathrm{K}$. Scaling this up by a factor of 50 yields contaminants of $\sim 20\,\mu\mathrm{K}$. 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.

5.8 Testing the formalism's approximations

The formalism set out in Section 5.3 includes two main approximations: that the likelihood need only be integrated over ranges of $\{\theta_0, \phi_0\}$ (and, indeed, θ_{crit}) corresponding to patches containing candidate sources, and that correlations between data inside and outside these patches can be neglected. The error made in each of these approximations will depend on the nature of the source templates. In particular, the accuracy of the first 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

Figure 5.10: 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 $N_{\mathrm{side,deg}} = 256$, 128 and 64. All low-resolution plots are shown with a degraded mask applied, which limits any residuals to $\mathcal{O}(\mu\mathrm{K})$. All plots are 12.5° on a side.

Figure 5.11: 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.

particular regions.

To explore these issues, consider the single-source contribution to the posterior Eq. 5.10. Neglecting correlations between regions inside and outside the patch (see Eq. 5.13) corresponds to the assumption that

$$\mathbf{t}\mathbf{C}^{-1}\mathbf{d}_{\bar{b}}^{\mathrm{T}} \ll 1\,, \tag{5.45}$$

where the template has support in a blob b, and the data consist of pixels in region $\bar{b}$ outside the blob. In Fig. 5.12, 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. 5.13, 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

Figure 5.12: 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.

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. 5.14: 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.

5.8.1 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 $N_\mathrm{side} = 64$. At 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 $z_0 = 8 \times 10^{-5}$

Figure 5.13: 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.

(as plotted in Fig. 5.9). 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. 5.9), 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. 5.12 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 5.4 and 5.5. In each case, increasing

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

Table 5.4: 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.

$\theta_{\mathrm{patch}}(°)$	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 5.5: The evidence ratios obtained when the size of the patch covariance matrix is incrementally increased until all correlations are included.

the patch size does not 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.

5.8.2 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. 5.16. 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 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 5.6. 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. 5.14, and indeed is very strongly peaked about

Center Range $/\theta_{\text{patch}}$	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 5.6: The evidence ratios obtained when the range of collision centers sampled is incrementally increased until the collision can be placed anywhere within the patch.

Figure 5.14: The localization-test likelihood as a function of co-latitude θ_0 (left) and longitude ϕ_0 (right), plotted on a logarithmic scale. 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 very strongly peaked in both angular coordinates.

the position highlighted by the candidate detection stage.

5.9 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 5.7.

ID	θ_{crit} range (°)	θ_0(°)	ϕ_0(°)	$\log \rho$
1	2-4	63.1	142.0	-9.68 ± 0.3
2	3-5	50.2	221.5	-10.87 ± 0.3
3	6-8	126.4	214.8	-8.75 ± 0.3
4	8-12	43.8	152.7	-8.57 ± 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

ID	θ_{crit} range (°)	θ_0(°)	ϕ_0(°)	$\log \rho$
1	2-5	50.2	221.5	-4.41 ± 0.3
2	4-9	45.8	153.1	-6.01 ± 0.3
3	4-15	143.2	3.2	-3.99 ± 0.3
4	7-50	38.7	94.8	-6.10 ± 0.3

Table 5.7: 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 $b = 90° - \theta$.

Figure 5.15: 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.

The patch evidence ratios obtained by applying the adaptive-resolution Bayesian evidence calculation to these candidates are also presented in Table 5.7. 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. 5.15. The posteriors for both models are clearly peaked at $\bar{N}_s = 0$, 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.

5.10 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 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 5.8 and plotted in Fig. 5.16.

ID	θ_{crit} range (°)	θ_0(°)	ϕ_0(°)	$\log \rho$
1	2-6	169.0	187.5	-8.82 ± 0.3
2	3-5	167.2	268.7	-7.79 ± 0.3
3	4-14	147.4	207.1	-5.80 ± 0.3
4	4-8	123.2	321.3	-7.17 ± 0.3
5	6-8	62.7	220.4	-9.33 ± 0.3
6	14-22	136.6	176.5	-5.96 ± 0.3
7	12-20	118.0	212.0	-7.39 ± 0.3
8	20-24	75.8	168.8	-9.46 ± 0.3
9	28-35	85.8	166.3	-7.04 ± 0.3
10	22-40	126.8	220.1	-5.67 ± 0.3
11	80-90	69.6	61.9	-6.66 ± 0.3

ID	θ_{crit} range (°)	θ_0(°)	ϕ_0(°)	$\log \rho$
1	1-3	114.6	22.1	-6.08 ± 0.3
2	1.5-5	168.7	184.4	-3.12 ± 0.3
3	1.5-4	166.8	268.8	-4.58 ± 0.3
4	1.5-3	72.4	150.1	-5.57 ± 0.3
5	2-15	123.5	321.0	-3.31 ± 0.3
6	2-7	123.2	79.5	-4.73 ± 0.3
7	2-4	128.3	93.5	-5.56 ± 0.3
8	2-15	147.4	210.2	-2.06 ± 0.3
9	5-10	69.3	202.1	-8.04 ± 0.3
10	8-15	119.3	155.7	-7.71 ± 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 5.8: 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.

Applying the adaptive-resolution evidence calculation to the candidates produces the patch evidence ratios also reported in Table 5.8. No single candidate is strong enough to claim a detection on its own. However, as demonstrated in Feeney et al. (2011a) and Feeney et al. (2012), 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 and texture models, derived by combining the results from the candidates, are plotted in Fig. 5.17: 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.

5.11 Discussion

In Feeney et al. (2011a,b) and Feeney et al. (2012), 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 (Gobbetti and Kleban, 2012; Kleban et al., 2011), 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 (Dickey, 1971): the change in evidence will be the ratio of the posterior and

Figure 5.16: 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).


Figure 5.17: 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, $\bar{N}_\mathrm{s}$, for bubble collisions, and two parameters, $\bar{N}_\mathrm{s}$ 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 $\bar{N}_\mathrm{s} = 0$.

prior probabilities of the edge amplitude, evaluated at $z_{\rm crit} = 0$ using the results of the previous analysis, i.e.,

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

As there was little evidence to support the edge parameter in the earlier analysis, the ratio of posterior to prior at $z_{\rm crit} = 0$ 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 5.9 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 (Feeney et al., 2011a,b) required that the $\theta_{\rm crit}$ range sampled be truncated (compare Tables VII in Feeney et al. (2011a) and 5.8 in the current analysis). The new adaptive-resolution algorithm allows the full range of $\theta_{\rm crit}$ 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.

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 5.9: 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.

The only differences between the current texture analysis and that of Feeney et al. (2012) 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 Feeney et al. (2012) and the right-hand plot of Fig. 5.17) are almost indistinguishable by eye.

5.12 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 (Feeney et al., 2011a,b), as

well as optimising the candidate detection stages of previous bubble collision and texture analyses (Feeney et al., 2011a,b; Feeney et al., 2012). 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 per 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 (Tauber et al., 2010) 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 (Niemack et al., 2010), SPTPol (McMahon et al., 2009) and Spider (Crill et al., 2008). Textures do not induce a polarization signal (Vielva et al., 2011), but bubble collisions are expected to create characteristic imprints (Czech et al., 2010), 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.

5.13   Acknowledgements

Appendix A

Appendix to Chapter 3

A.1 Statistical formalism

A.1.1 Posterior

In this appendix, we discuss how Bayesian parameter estimation and model selection for theories which predict localized sources can be approximated by a patch-wise analysis. Consider astronomical observations covering solid angle $\Omega_{\text{obs}} = 4\pi f_{\text{sky}}$ that are of sufficient depth/resolution to identify sources with a particular range of properties (which can then be deemed "detectable"). Given a theory that predicts an expectation value of $\bar{N}_{\text{s}}$ sources over the whole sky, we want to know both: what constraints the available data place on $\bar{N}_{\text{s}}$; and whether the data favour a model which predicts one value of $\bar{N}_{\text{s}}$ over another. all the relevant information is encoded in the posterior distribution $\Pr(\bar{N}_{\text{s}}|\mathbf{d}, f_{\text{sky}})$, where $\mathbf{d}$ are the pixelized flux or temperature measurements (and, optionally, any statistics derived from them). Bayes' theorem allows the posterior to be written as

$$\Pr(\bar{N}_{\text{s}}|\mathbf{d}, f_{\text{sky}}) = \frac{\Pr(\bar{N}_{\text{s}}) \Pr(\mathbf{d}|\bar{N}_{\text{s}}, f_{\text{sky}})}{\Pr(\mathbf{d}|f_{\text{sky}})}, \tag{A.1}$$

where $\Pr(\bar{N}_{\text{s}})$ is the prior distribution on $\bar{N}_{\text{s}}$, $\Pr(\mathbf{d}|\bar{N}_{\text{s}}, f_{\text{sky}})$ is the likelihood of getting the observed data given the area of observation and the expected number of sources, and $\Pr(\mathbf{d}|f_{\text{sky}})$ ensures that the posterior is normalized over $\bar{N}_{\text{s}}$. Constraints on $\bar{N}_{\text{s}}$ can be drawn directly from this normalized posterior; the relative probability of models predicting different values of $\bar{N}_{\text{s}}$ can be found by picking out the posterior at two values of $\bar{N}_{\text{s}}$:

$$\frac{\Pr(\bar{N}_{\text{s},1}|\mathbf{d}, f_{\text{sky}})}{\Pr(\bar{N}_{\text{s},2}|\mathbf{d}, f_{\text{sky}})} = \frac{\Pr(\bar{N}_{\text{s},1})\Pr(\mathbf{d}|\bar{N}_{\text{s},1}, f_{\text{sky}})}{\Pr(\bar{N}_{\text{s},2})\Pr(\mathbf{d}|\bar{N}_{\text{s},2}, f_{\text{sky}})}. \tag{A.2}$$

In the absence of a prescriptive theory, it is useful to emphasize the role of the data, which

can be done by adopting a flat prior on $\bar{N}_{\rm s}$; further assuming that the data will give an upper limit on $\bar{N}_{\rm s}$, it is possible to adopt an improper uniform prior $\Pr(\bar{N}_{\rm s}) = \Theta(\bar{N}_{\rm s})$ without any high-$\bar{N}_{\rm s}$ cut-off. The resultant posterior has the form

$$\Pr(\bar{N}_{\rm s}|\mathbf{d}, f_{\rm sky}) \propto \Theta(\bar{N}_{\rm s})\Pr(\mathbf{d}|\bar{N}_{\rm s}, f_{\rm sky}), \tag{A.3}$$

up to a normalization constant that depends on the data and $f_{\rm sky}$ but not on $\bar{N}_{\rm s}$.

In general $\bar{N}_{\rm s}$ is not directly measurable, even for perfect data, because the number of sources present in the observable sky, $N_{\rm s}$, is the realization of a Poisson-like process (of mean $f_{\rm sky}\bar{N}_{\rm s}$). The possibility that $N_{\rm s}$ is itself subject to some uncertainty (e.g., due to noisy data or confusion problems) can be incorporated by marginalizing over $N_{\rm s}$ to give

$$
\begin{aligned}
\Pr(\mathbf{d}|\bar{N}_{\rm s}, f_{\rm sky}) &= \sum_{N_{\rm s}=0}^{\infty} \Pr(N_{\rm s}|\bar{N}_{\rm s}, f_{\rm sky})\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky}) \\
&= \sum_{N_{\rm s}=0}^{\infty} \frac{(f_{\rm sky}\bar{N}_{\rm s})^{N_{\rm s}} e^{-f_{\rm sky}\bar{N}_{\rm s}}}{N_{\rm s}!}\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky}),
\end{aligned} \tag{A.4}
$$

where the second formula explicitly assumes that the number of observable sources is drawn from a Poisson process. Inserting this second expression into Eq. A.3 then gives

$$\Pr(\bar{N}_{\rm s}|\mathbf{d}, f_{\rm sky}) \propto \Theta(\bar{N}_{\rm s})\, e^{-f_{\rm sky}\bar{N}_{\rm s}} \sum_{N_{\rm s}=0}^{\infty} \frac{(f_{\rm sky}\bar{N}_{\rm s})^{N_{\rm s}}}{N_{\rm s}!}\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky}). \tag{A.5}$$

The form of the likelihood $\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky})$ is treated largely in abstract here, with the specific details of the likelihood calculation for the bubble collision hypothesis given WMAP 7-year data described in Sec. 3.6.3. Assuming the measurements take the form of flux/counts at different positions on the sky (as for a CMB experiment) and that they are subject to (possibly correlated) Gaussian noise, the likelihood would have the form

$$\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky}) = \int \mathrm{d}\mathbf{m}_1 \ldots \mathrm{d}\mathbf{m}_{N_{\rm s}} \prod_{s=1}^{N_{\rm s}} \Pr(\mathbf{m}_s) \frac{e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m}_1)\ldots-\mathbf{t}(\mathbf{m}_{N_{\rm s}})]\mathsf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m}_1)\ldots-\mathbf{t}(\mathbf{m}_{N_{\rm s}})]^{\mathrm{T}}/2}}{(2\pi)^{n_{\rm pix}/2}|\mathsf{C}|}, \tag{A.6}$$

where $\mathbf{t}(\mathbf{m})$ is the data template that would result from a source whose position and profile/scale are defined by the model parameters $\mathbf{m}$, $\Pr(\mathbf{m}_s)$ is the prior distribution of source parameters for "detectable" sources, and C is the pixel-pixel covariance matrix of the non-source noise (which could include contributions that are considered signal in other contexts, such as the CMB).

Evaluating the full sum in Eq. A.5 is not always practical or even possible. In addition, the evaluation of individual terms in this sum will be computationally limited by the size of

the covariance matrix $\mathbf{C}$ and the cost of performing the integral over model parameters for each template. However, it is possible to circumvent these problems, and estimate the posterior Eq. A.5 if one knows in advance some of the properties of the integrand in Eq. A.6.

To see how this works, assume that one has located a set of N_{b} "blobs" on the sky that are candidate sources. Segment the sky into $N_{\mathrm{r}} = N_{\mathrm{b}} + 1$ regions consisting of those containing blobs, and the rest of the sky. Given N_{b}, we can now evaluate Eq. A.5 term-by-term. The likelihood in the first term, for $N_{\mathrm{s}} = 0$, is simply given by:

$$\Pr(\mathbf{d}|0, f_{\mathrm{sky}}) = \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-\mathbf{d}\mathsf{C}^{-1}\mathbf{d}^{\mathrm{T}}/2}, \tag{A.7}$$

which is the likelihood for the null-hypothesis with no sources. Moving on to the $N_{\mathrm{s}} = 1$ term, we first expand the integral over source positions to cover each of the N_{r} regions:

$$\Pr(\mathbf{d}|1, f_{\mathrm{sky}}) = \sum_{r=1}^{N_{\mathrm{r}}} \int_{\mathrm{region}\ r} d\mathbf{m}\, \Pr(\mathbf{m}) \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m})]\mathsf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m})]^{\mathrm{T}}/2}, \tag{A.8}$$

We now assume that the blobs containing candidate sources include all of the significant contributions to the integral, and replace N_{r} in the sum by N_{b}. This will give us a *lower bound* on the likelihood, even if a number of actual sources are not contained within the blobs defined by the candidate sources. We further assume that sources do not overlap. If the covariance matrix is small enough to invert, we could stop here. However, in cases where the covariance matrix is too large to feasibly invert (as is the case for the WMAP 7 year data), we can make one further approximation:

$$\int_{\mathrm{region}\ \mathrm{b}} d\mathbf{m}\, \Pr(\mathbf{m}) \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}|} e^{-[\mathbf{d}-\mathbf{t}(\mathbf{m})]\mathsf{C}^{-1}[\mathbf{d}-\mathbf{t}(\mathbf{m})]^{\mathrm{T}}/2} \simeq \int_{\mathrm{region}\ \mathrm{b}} d\mathbf{m}\, \Pr(\mathbf{m}) \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{m}), \tag{A.9}$$

where the product is over N_{r} disjoint regions on the sky and

$$L_r(\mathbf{m}) = \frac{1}{(2\pi)^{n_{\mathrm{pix}}/2}|\mathsf{C}_r|} e^{-[\mathbf{d}_r-\mathbf{t}_r(\mathbf{m})]\mathsf{C}_r^{-1}[\mathbf{d}_r-\mathbf{t}_r(\mathbf{m})]^{\mathrm{T}}/2}. \tag{A.10}$$

is the contribution to the likelihood of data in region r, defined in terms of the covariance of the pixels in this region, C_r, the data in this region, $\mathbf{d}_r$, and the source template in this region, $\mathbf{t}_r(\mathbf{m})$. This is exact for a diagonal covariance, but is only approximate in the case where there are off-diagonal elements. Using the assumption that the integral has a significant contribution only inside the blobs, in the rest of the sky we can replace $\mathbf{t} = \mathbf{0}$. The one-source model likelihood

then becomes

$$\Pr(\mathbf{d}|1, f_{\mathrm{sky}}) \simeq \sum_{b=1}^{N_{\mathrm{b}}} \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0})\rho_b, \tag{A.11}$$

where

$$\rho_b = \frac{\int_{\mathrm{region}\ b} \mathrm{d}\mathbf{m}\, \Pr(\mathbf{m}) L_b(\mathbf{m})}{L_b(\mathbf{0})}. \tag{A.12}$$

This is the evidence ratio for a single source template centered in region b.

For a general number of blobs and sources, the model likelihood is

$$\Pr(\mathbf{d}|N_{\mathrm{s}}, f_{\mathrm{sky}}) = \begin{cases} 0, & \text{if} \quad N_{\mathrm{s}} > N_{\mathrm{b}}, \\[2ex] \sum_{b_1, b_2, \ldots, b_{N_{\mathrm{s}}}=1}^{N_{\mathrm{b}}} \left[\prod_{s=1}^{N_{\mathrm{s}}} \rho_{b_s} \prod_{i,j=1}^{N_{\mathrm{s}}} (1 - \delta_{s_i, s_j}) \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0}) \right], & \text{if} \quad N_{\mathrm{s}} \leq N_{\mathrm{b}}, \end{cases} \tag{A.13}$$

where the combinatorics require some explanation. If there are fewer blobs on the sky than proposed sources then the likelihood is very small: by assumption, the likelihood evaluated outside of a blob is small. 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_{s} sources with the N_{b} blobs, provided that no two sources are matched to the same blob. Hence the multiple sum generates all possible combinations of source-blob associations and the product over evidence ratios gives the relevant weightings; the product over delta functions removes the terms in which any two sources are attached to the same blob.

Inserting the likelihood given in Eq. A.13 into Eq. A.5 yields the unnormalized posterior on $\bar{N}_{\mathrm{s}}$:

$$\Pr(\bar{N}_{\mathrm{s}}|\mathbf{d}, f_{\mathrm{sky}}) \propto \Theta(\bar{N}_{\mathrm{s}})\, e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}} \sum_{N_{\mathrm{s}}=0}^{N_{\mathrm{b}}} \frac{(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^{N_{\mathrm{s}}}}{N_{\mathrm{s}}!} \sum_{b_1, b_2, \ldots, b_{N_{\mathrm{s}}}=1}^{N_{\mathrm{b}}} \left[\prod_{s=1}^{N_{\mathrm{s}}} \rho_{b_s} \prod_{i,j=1}^{N_{\mathrm{s}}} (1 - \delta_{s_i, s_j}) \right], \tag{A.14}$$

under the assumption that

$$\Pr(\mathbf{d}|0, f_{\mathrm{sky}}) = \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0}), \tag{A.15}$$

in which case regions that do not contain a blob are irrelevant for determining the posterior. Eq. A.14 is the main result of this calculation, from which all following results can be derived. In the limit of a single isolated observation Eq. A.14 reproduces the Bayesian source detection formalism developed in Hobson and McLachlan (2003) and Hobson et al. (2010).

A.1.2 Special cases

Perfect data

With infinite, perfect data the number of sources on the sky would be directly determined by counting the $N_{\rm b}$ "blobs" in the data and so $\Pr(\mathbf{d}|N_{\rm s}, f_{\rm sky}) = \delta_{N_{\rm s}, N_{\rm b}}$. The posterior in Eq. A.14 would become

$$\Pr(\bar{N}_{\rm s}|N_{\rm b}, f_{\rm sky}) \propto \Theta(\bar{N}_{\rm s})\bar{N}_{\rm s}^{N_{\rm b}}e^{-f_{\rm sky}\bar{N}_{\rm s}}, \tag{A.16}$$

the standard result for constraining a rate variable from a single measurement, modified slightly to account for the fact that the constraint on $\bar{N}_{\rm s}$ is weakened if $f_{\rm sky} \ll 1$. In the even more particular case that no blobs were detected in perfect data, the posterior would be

$$\Pr(\bar{N}_{\rm s}|0, f_{\rm sky}) = \Theta(\bar{N}_{\rm s})\, f_{\rm sky}e^{-f_{\rm sky}\bar{N}_{\rm s}}. \tag{A.17}$$

If a single blob was detected unequivocally then the posterior would be

$$\Pr(\bar{N}_{\rm s}|1, f_{\rm sky}) = \Theta(\bar{N}_{\rm s})\, f_{\rm sky}(f_{\rm sky}\bar{N}_{\rm s})e^{-f_{\rm sky}\bar{N}_{\rm s}}. \tag{A.18}$$

If two blobs were detected unequivocally then the posterior would be

$$\Pr(\bar{N}_{\rm s}|2, f_{\rm sky}) = \Theta(\bar{N}_{\rm s})\frac{f_{\rm sky}}{2}(f_{\rm sky}\bar{N}_{\rm s})^2 e^{-f_{\rm sky}\bar{N}_{\rm s}}. \tag{A.19}$$

No blobs

If there are no identified blobs then $N_{\rm b} = 0$ and there is no evidence for any sources at all. This is really a weaker constraint than the above situation if the data are perfect, but in the approximation used here the final result is the same. In Eq. A.14 the first sum is truncated at the first term and so

$$\Pr(\bar{N}_{\rm s}|0, \mathbf{d}, f_{\rm sky}) = \Theta(\bar{N}_{\rm s})\, f_{\rm sky}e^{-f_{\rm sky}\bar{N}_{\rm s}}, \tag{A.20}$$

matching Eq. A.17.

Adopting a flat prior on $\bar{N}_{\rm s}$, the posterior probability ratio for a model predicting a generic $\bar{N}_{\rm s} > 0$ versus one predicting no collisions in this case is given by

$$\frac{\Pr(\bar{N}_{\rm s}\,|\,0, \mathbf{d}, f_{\rm sky})}{\Pr(0\,|\,0, \mathbf{d}, f_{\rm sky})} = e^{-f_{\rm sky}\bar{N}_{\rm s}}. \tag{A.21}$$

This is always less than one, and so as expected, a theory which predicts $\bar{N}_{\rm s}$ sources on the sky is always disfavoured when compared to a theory that predicts no sources on the sky.

One blob

Probably the most important simple case is where there is a single identified blob, which might represent a first detection of this class of source. Inserting $N_{\mathrm{b}} = 1$ into Eq. A.14, the sum includes the possibilities of either one source on the (observed) sky or no sources; the posterior evaluates to

$$\Pr(\bar{N}_{\mathrm{s}}|1,\mathbf{d},f_{\mathrm{sky}}) = \Theta(\bar{N}_{\mathrm{s}})\, f_{\mathrm{sky}} e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}} \frac{1 + f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}\rho_b}{1 + \rho_b}, \tag{A.22}$$

In the limit that the data in this region are much better fit by a source then $\rho_b \gg 1$, and the posterior becomes

$$\Pr(\bar{N}_{\mathrm{s}}|1,\mathbf{d},f_{\mathrm{sky}}) = \Theta(\bar{N}_{\mathrm{s}})\, f_{\mathrm{sky}} (f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}) e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}, \tag{A.23}$$

which matches A.18 above. Conversely, in the limit that the source is a worse fit to the data (possible given that the source has been forced to be detectable), then $\rho_b \ll 1$ and

$$\Pr(\bar{N}_{\mathrm{s}}|1,\mathbf{d},f_{\mathrm{sky}}) = \Theta(\bar{N}_{\mathrm{s}})\, f_{\mathrm{sky}} e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}, \tag{A.24}$$

matching Eq. A.17 which was obtained under the assumption that there was no blob in the first place.

Adopting again a flat prior on $\bar{N}_{\mathrm{s}}$, the posterior probability ratio for a model predicting a generic $\bar{N}_{\mathrm{s}} > 0$ versus the no-bubble case is given by

$$\frac{\Pr(\bar{N}_{\mathrm{s}}|\,1,\mathbf{d},f_{\mathrm{sky}})}{\Pr(0|1,\mathbf{d},f_{\mathrm{sky}})} = e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}} \left(1 + f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}\rho_b\right). \tag{A.25}$$

Here, it can be seen that two things are necessary to favour the theory with sources given one detection: $\bar{N}_{\mathrm{s}} \sim \mathcal{O}(1)$ *and* $\rho_b \gg 1$.

Two blobs

If two blobs are identified then the sum in A.13 has three terms, for which the likelihoods are:

$$\Pr(2,\mathbf{d}|0,f_{\mathrm{sky}}) = \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0}), \tag{A.26}$$

$$\Pr(2,\mathbf{d}|1,f_{\mathrm{sky}}) = [\rho_{b_1} + \rho_{b_2}] \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0}) \tag{A.27}$$

and

$$\Pr(2,\mathbf{d}|2,f_{\mathrm{sky}}) = \rho_{b_1}\rho_{b_2} \prod_{r=1}^{N_{\mathrm{r}}} L_r(\mathbf{0}). \tag{A.28}$$

Hence the (unnormalized) posterior is

$$\mathrm{Pr}(\bar{N}_{\mathrm{s}}|2, \mathbf{d}, f_{\mathrm{sky}}) \propto \Theta(\bar{N}_{\mathrm{s}})e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}\left\{1 + f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}\left[\rho_{b_1} + \rho_{b_2}\right] + (f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^2 \rho_{b_1}\rho_{b_2}\right\}. \tag{A.29}$$

In the limit that the evidence for both sources is strong (i.e., $\rho_{b_1} \gg 1$ and $\rho_{b_2} \gg 1$) then the third term in the curly braces dominates and

$$\mathrm{Pr}(\bar{N}_{\mathrm{s}}|2, \mathbf{d}, f_{\mathrm{sky}}) = \Theta(\bar{N}_{\mathrm{s}})\frac{f_{\mathrm{sky}}}{2}(f_{\mathrm{sky}}\bar{N}_{\mathrm{s}})^2 e^{-f_{\mathrm{sky}}\bar{N}_{\mathrm{s}}}, \tag{A.30}$$

which matches the perfect data case with $N_{\mathrm{b}} = 2$, as expected. In the limit where one blob is a false candidate, but the other yields a strong evidence (e.g., $\rho_{b_1} \gg 1$ and $\rho_{b_2} \ll 1$), then we recover the perfect data case with $N_{\mathrm{b}} = 1$.

Appendix B

Optimal filters for detecting cosmic bubble collisions

B.1   Abstract

A number of well-motivated extensions of the ΛCDM concordance cosmological model postulate the existence of a population of sources embedded in the cosmic microwave background (CMB). One such example is the signature of cosmic bubble collisions which arise in models of eternal inflation. The most unambiguous way to test these scenarios is to evaluate the full posterior probability distribution of the global parameters defining the theory; however, a direct evaluation is computationally impractical on large datasets, such as those obtained by the Wilkinson Microwave Anisotropy Probe (WMAP) and Planck. A method to approximate the full posterior has been developed recently, which requires as an input a set of candidate sources which are most likely to give the largest contribution to the likelihood. In this article, we present an improved algorithm for detecting candidate sources using optimal filters, and apply it to detect candidate bubble collision signatures in WMAP 7-year observations. We show both theoretically and through simulations that this algorithm provides an enhancement in sensitivity over previous methods by a factor of approximately two. Moreover, no other filter-based approach can provide a superior enhancement of these signatures. Applying our algorithm to WMAP 7-year observations, we detect eight new candidate bubble collision signatures for follow-up analysis.

B.2 Introduction

Precision observations of the cosmic microwave background (CMB) provide the most accurate picture of the early universe that is available currently. The standard ΛCDM concordance cosmological model – which states that we live in a universe composed mostly of dark energy and dark matter, whose structure was seeded by adiabatic and very nearly Gaussian and scale-invariant density perturbations – describes the statistics of temperature fluctuations in the CMB extremely well (Komatsu et al., 2011; Larson et al., 2011). However, there are many theoretically well-motivated extensions of ΛCDM that predict detectable secondary signals in the CMB.

One example, which has been the subject of a number of recent studies (Garriga et al., 2007; Aguirre et al., 2007; Aguirre and Johnson, 2008; Aguirre et al., 2009; Chang et al., 2008, 2009; Czech et al., 2010; Dahlen, 2010; Freivogel et al., 2009; Larjo and Levi, 2010; Kleban et al., 2011; Gobbetti and Kleban, 2012), is the signature of cosmic bubble collisions which arise in models of eternal inflation (see Aguirre and Johnson (2009) for a review). In the model of eternal inflation, our observable universe is contained inside one member of an ensemble of bubbles. Collisions between bubbles disturb the homogeneity and isotropy of the very early universe, leaving possibly detectable imprints on the CMB. In the limit where the number of detectable collisions on the CMB sky is relatively small, the signature is a set of azimuthally-symmetric modulations of the temperature (Garriga et al., 2007; Aguirre et al., 2007), varying as the cosine of the angular distance from the collision centre (Chang et al., 2009), with a size-distribution peaking at half-sky scales (Freivogel et al., 2009). Other examples of secondary signals arise in theories with topological defects such as cosmic strings (see e.g. Vilenkin and Shellard (1986) for a review) or textures (Turok and Spergel, 1990); a less exotic example is the signature of clusters of galaxies produced by the Sunyaev-Zel'dovich (SZ) effect (Sunyaev and Zeldovich, 1980).

In each of these examples, a population of sources is hypothesized to exist on top of the background CMB, the members of which have properties drawn from a calculable probability distribution. The most unambiguous way to test these scenarios is to utilize the most general predictions for the population of sources on the full-sky, and determine the posterior probability distribution over the global parameters defining the theory (such as the total number of features expected, their intrinsic amplitude, *etc.*). The enormous size of modern CMB datasets, such as those obtained by the Wilkinson Microwave Anisotropy Probe (Bennett et al., 2003a) (WMAP) and those currently being obtained by the Planck satellite (Tauber et al., 2010), provide a unique challenge for such an analysis. Indeed, a direct pixel-based evaluation of the posterior at full resolution is computationally intractable.

Recently, however, Feeney et al. (2011a,b) outlined a method for approximating the full

posterior describing source populations in the context of the bubble collision hypothesis. The method is generalized to the detection of other sources easily. This approach requires pre-processing of the data to recover a set of candidate sources which are most likely to give the largest contribution to the likelihood. The preprocessing stage of this method is thus crucial to its overall effectiveness. Candidate source detection aims to minimize the number of false detections while remaining sensitive to a weak signal; a manageable number of false detections is thus tolerated, as the subsequent Bayesian processing step will discriminate these from true signals. To detect candidate bubble collision signatures, Feeney et al. (2011a,b) employ a suite of needlet transforms (Marinucci et al., 2008; Scodeller et al., 2011). Needlets are a form of azimuthally-symmetric wavelet[1] defined on the sphere, that render the location and scale of candidate features simultaneously accessible[2]. While the effectiveness of needlets for detecting candidate features has been demonstrated already (Feeney et al., 2011a,b), needlets are generic and are not adapted to the signal of interest; consequently, they are not optimal. A better approach is to enhance the effectiveness of candidate detection by exploiting knowledge of the source signature.

Optimal filters have found widespread application in many branches of physics and signal processing for the detection of compact objects embedded in a stochastic background. In the context of astrophysics, the matched filter has been applied to detect point sources and SZ emission in CMB observations (Tegmark and de Oliveira-Costa, 1998; Haehnelt and Tegmark, 1996). Alternative optimal filters, such as the scale-adaptive filter, have also been derived (Sanz et al., 2001; Herranz et al., 2002) and applied to CMB observations (Barreiro et al., 2003). In all of these cases, optimal filters are applied to small patches of the sky, where a flat tangent plane approximation of the celestial sphere in the region of interest is made. To analyze full-sky CMB observations these techniques must be extended from Euclidean space to a spherical manifold. Optimal filter theory has been extended to the sphere by Schaefer et al. (2006) (and applied to detect SZ emission (Malte Schafer and Bartelmann, 2007)) for the case of azimuthally-symmetric source signatures and by McEwen et al. (2008) for the general directional setting.

In this article we develop an alternative candidate source detection algorithm using optimal filters. We focus on the problem of detecting the signatures of bubble collisions in observations of the CMB, but our approach generalizes to other sources and backgrounds trivially. Since the angular scale of a typical bubble collision is expected to be large (Freivogel et al., 2009; Aguirre et al., 2007; Aguirre and Johnson, 2009), tangent plane approximations are not valid, and we

---

[1] Note that Mexican needlets (Scodeller et al., 2011) are not formally wavelets since exact synthesis is not possible, even in theory.

[2] Needlets are in fact the azimuthally-symmetric restriction of exact steerable wavelets defined on the sphere (Wiaux et al., 2008), which render the orientation of directional features also accessible.

instead consider optimal filters defined on the sphere (Schaefer et al., 2006; McEwen et al., 2008). We describe and evaluate our new candidate source detection algorithm in Sec. B.3 and show it to be superior to the needlet approach considered previously (Feeney et al., 2011a,b). Finally, we apply our algorithm to WMAP observations in Sec. B.4, resulting in the detection of a number of new candidate bubble collision signatures in the WMAP 7-year data. Concluding remarks are made in Sec. B.5.

B.3   Optimal detection of candidate bubble collisions

Filter based approaches to enhance a signal in a background process are common due to their effectiveness and efficiency. Indeed, a wavelet transform, such as needlets, is merely a filtering operation with a carefully constructed set of filter kernels (to allow the exact reconstruction of the original signal). In this section we consider filters that provide the maximal enhancement of the source signature in a given stochastic background. The filters are optimal in the sense that no other filter can yield a greater enhancement in the signal-to-noise ratio (SNR) of the filtered field. Our optimal-filter-based method is general: in this work, we focus on its application to the problem of detecting signatures of bubble collisions. Firstly, we define the signatures of the bubble collision remnants that we search for. We then construct and evaluate optimal filters for detecting candidate bubble collision signatures when the size of the signature is known, before describing an algorithm for detecting multiple candidate bubble collision signatures of unknown and differing sizes.

B.3.1   Bubble collision signatures

Bubble collisions induce a modulative and additive contribution to the temperature fluctuations of the CMB (Chang et al., 2009), however the modulative component is second order and may be safely ignored. The additive contribution induced in the CMB by a bubble collision is given by the azimuthally-symmetric profile

$$\Delta T_{\mathrm{b}}(\theta, \phi) = [c_0 + c_1 \cos(\theta)] \, s(\theta; \theta_{\mathrm{crit}}) \,,$$

when centered on the North pole, where $(\theta, \phi) \in \mathrm{S}^2$ denote the spherical coordinates of the unit sphere S^2, with colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi)$, and c_0 and c_1 are free parameters (not to be confused with the power spectrum monopole and dipole). A typical bubble collision signature is illustrated in Fig. B.1. Following the parameterization of Feeney et al. (2011a,b), we describe the bubble collision signature by its amplitude at its centre and at its causal boundary,

(a) Bubble collision signature radial profile

(b) Bubble collision signature on the sphere

(c) Power spectrum of bubble collisionsignature and CMB

Figure B.1: Panels (a) and (b) show the radial profile and 3D surface plot (lit from top-left, with specular highlight), respectively, of a bubble collision signature with parameters $\{z_0, \theta_{\mathrm{crit}}, \theta_0, \phi_0\} = \{100\ \mu\mathrm{K}, 10°, 0°, 0°\}$ (throughout we consider $z_{\mathrm{crit}} \sim 0\ \mu\mathrm{K}$). In panel (c) the power spectrum of the bubble collision signature (solid blue curve) is compared with the best-fit WMAP7+BAO+H0 CMB power spectrum (red dashed curve). Matched filters for azimuthally symmetric templates promote harmonic modes where the source template power spectrum is large and suppress modes where the CMB power spectrum is large.

given by $z_0 = c_0 + c_1$ and $z_{\mathrm{crit}} = c_0 + c_1 \cos(\theta_{\mathrm{crit}})$ respectively, and by its size θ_{crit}. We replace the discontinuous Heaviside step function of the bubble collision profile with a "Schwartz" step function $s(\theta; \theta_{\mathrm{crit}})$ that is infinitely differentiable but nevertheless exhibits a smooth but rapid transition to zero about θ_{crit}. As theoretical work suggests that the temperature discontinuity parameter should be negligible (Gobbetti and Kleban, 2012; Kleban et al., 2011) (an observation that is supported by the candidate bubble collision signatures detected previously (Feeney et al., 2011a,b)), we restrict our attention to $z_{\mathrm{crit}} \sim 0\ \mu K$. Bubble collision signatures may occur at any position on the sky (θ_0, ϕ_0) and at a range of sizes θ_{crit} and amplitudes z_0. We denote by ΔT_i the temperature contribution induced by a candidate bubble collision i with parameters $\{z_0, \theta_{\mathrm{crit}}, \theta_0, \phi_0\}$.

B.3.2 Optimal bubble collision filters

We define optimal filters to enhance the contributions of compact sources embedded in a stochastic background, focusing on the case of locating candidate bubble collision signatures in the CMB. Firstly, we discuss filtering on the sphere in general, before defining the optimal matched filter. We compute the matched filter for detecting bubble collision signatures and compare the SNR for the matched filter to alternatives, such as needlets and the unfiltered field itself.

Filtering

Filtering on the sphere is the natural analogue of the filtering operation in Euclidean space and is defined by the projection of a function, such as the CMB temperature fluctuations ΔT, onto rotated filter kernels. Consequently, filtering on the sphere is defined through the spherical convolution

$$
\begin{aligned}
F_R(\rho) &= \langle \Delta T, \, \mathcal{R}(\rho)\Psi_R \rangle \\
&= \int_{\mathrm{S}^2} \mathrm{d}\Omega(\theta', \phi') \, \Delta T(\theta', \phi') \, [\mathcal{R}(\rho)\Psi_R]^*(\theta', \phi') \,,
\end{aligned}
\tag{B.1}
$$

where Ψ_R is the filter kernel at scale R, $\mathcal{R}$ is the rotation operator describing a rotation by the Euler angles $\rho \in \mathrm{SO}(3)$, $\langle \cdot, \, \cdot \rangle$ denotes the inner product on the sphere, * denotes complex conjugation and $\mathrm{d}\Omega(\theta, \phi) = \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$ is the usual rotation-invariant measure on the sphere. The filtering operation given by Eq. (B.1) is general in the sense that directional filter kernels are considered. Since we are concerned with bubble collision signatures, which are azimuthally-symmetric, we henceforth restrict our attention to azimuthally-symmetric filter kernels such that $\Psi_R(\theta, \phi) = \Psi_R(\theta)$. In this case, the filter kernel is invariant under rotations about its own axis of symmetry and the set of distinct rotations is restricted from the rotation group $\mathrm{SO}(3)$ to the sphere S^2, *i.e.* $\rho = (\theta, \phi) \in \mathrm{S}^2$.

Just like in the Euclidean setting, filtering on the sphere can be computed more efficiently in harmonic space than through an evaluation of Eq. (B.1) by direct quadrature. The CMB temperature fluctuations may be represented by their expansion in the basis of spherical harmonics $Y_{\ell m}$, given by

$$
\Delta T(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi) \,,
\tag{B.2}
$$

where the harmonic coefficients are given by the usual projection onto the basis functions: $a_{\ell m} = \langle \Delta T, Y_{\ell m} \rangle$. In practice, we consider a maximum band-limit $\ell_{\max}$, such that the summation over ℓ in Eq. (B.2) may be truncated to $\ell_{\max}$. Similarly, the filter kernel may be decomposed into its spherical harmonic expansion, with coefficients given by $(\Psi_R)_{\ell m} = \langle \Psi_R, Y_{\ell m} \rangle$. For an

azimuthally-symmetric kernel the filter coefficients are non-zero for harmonic indices $m = 0$ only, i.e. $(\Psi_R)_{\ell m} = \delta_{m0}(\Psi_R)_{\ell 0}$, where δ_{ij} is the Kronecker delta symbol. In this setting, the harmonic coefficients of the filtered field are given by

$$(F_R)_{\ell m} = \sqrt{\frac{4\pi}{2\ell + 1}}\, a_{\ell m}\, (\Psi_R)^*_{\ell 0}\,. \tag{B.3}$$

Fast spherical harmonic transforms (e.g. Driscoll and Healy (1994); Górski et al. (2005); Doroshkevich et al. (2005); McEwen and Wiaux (2011)) may then be employed to reduce the complexity of filtering with an azimuthally-symmetric kernel from $\mathcal{O}(\ell_{\max}^4)$ to $\mathcal{O}(\ell_{\max}^3)$.[3]

The purpose of filtering the observed signal on the sphere is to enhance source signatures relative to the stochastic background; we thus require a quantitive measure of the effectiveness of filtering. We define the SNR of the filtered field for scale R by the ratio of its mean to its dispersion in the presence of a source located at (θ_0, ϕ_0):

$$\Gamma_R = \frac{\mu_R(\theta_0, \phi_0)}{\sigma_R(\theta_0, \phi_0)}\,, \tag{B.4}$$

where the mean and variance of the filtered field are defined, respectively, by

$$\mu_R(\theta, \phi) = \mathbb{E}[F_R(\theta, \phi)]$$

and

$$\sigma_R^2(\theta, \phi) = \mathbb{E}[|F_R(\theta, \phi)|^2] - \mu_R^2(\theta, \phi)\,.$$

Optimal filters

The observed CMB temperature fluctuations ΔT are assumed to be comprised of a number of compact sources ΔT_i, such as bubble collision signatures, embedded in a stochastic background noise process n:

$$\Delta T(\theta, \phi) = \sum_i \Delta T_i(\theta, \phi) + n(\theta, \phi)\,.$$

We decompose the sources into their amplitude A_i and normalized template profile τ_i by

$$\Delta T_i(\theta, \phi) = A_i\, \tau_i(\theta, \phi);$$

for the case of bubble collision signatures we make the association $A = z_0$. The stochastic noise process n is assumed to be zero-mean, isotropic and homogeneous and is defined by its power

---

[3]Filtering with directional filter kernels can also be performed more efficiently in harmonic space than in real space (Risbo, 1996; Wandelt and Górski, 2001; McEwen et al., 2007).

spectrum:

$$\mathbb{E}[n_{\ell m} n_{\ell' m'}^{*}] = C_{\ell}\, \delta_{\ell \ell'}\, \delta_{mm'}\,,$$

where $n_{\ell m} = \langle n,\, Y_{\ell m}\rangle$. The source population is the signal of interest, hence the noise is comprised of primary and secondary CMB anisotropies.

We filter the observed CMB temperature fluctuations ΔT with the aim of enhancing the source contributions ΔT_i relative to the background noise n. The matched filter Ψ_R^{MF} is defined to maximize the SNR of the filtered field given by Eq. (B.4), while ensuring that the amplitude of the filtered field at the source position gives an unbiased estimator of the source amplitude. Thus, the matched filter defined on the sphere is recovered by solving the constrained optimization problem:

$$\min_{\mathrm{w.r.t.}\ \Psi_R}\ \sigma_R^2(\theta_0, \phi_0)\ \text{such that}\ \mu_R(\theta_0, \phi_0) = A\,.$$

The resulting matched filter is given by (McEwen et al., 2008)

$$(\Psi_R^{\mathrm{MF}})_{\ell m} = \frac{\tau_{\ell m}}{\alpha\, C_{\ell}}\,, \tag{B.5}$$

where

$$\alpha = \sum_{\ell m} C_{\ell}^{-1} |\tau_{\ell m}|^2$$

and $(\Psi_R^{\mathrm{MF}})_{\ell m} = \langle \Psi_R^{\mathrm{MF}},\, Y_{\ell m}\rangle$. Here and subsequently we use the shorthand notation $\sum_{\ell m} = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell}$. On inspection of the filtering operation in harmonic space given by Eq. (B.3), the matched filter given by Eq. (B.5) is justified intuitively since the filter promotes harmonic modes where the source template $\tau_{\ell m}$ is large and suppresses modes where the noise power C_{ℓ} is large.

In Fig. B.2 we plot the matched filters that are optimized to bubble collision signatures of varying size embedded in a CMB background defined by the ΛCDM power spectrum that best fits WMAP 7-year, baryon acoustic oscillations and supernovae observations (hereafter we refer to this spectrum as the best-fit WMAP7+BAO+H0 power spectrum) (Larson et al., 2011). Notice that on smaller scales the matched filter contains a central broad hot region to enhance the main bubble collision contribution, surrounded by hot and cold rings to enhance the collision edge. However, on larger scales notice that the matched filter contains only the hot and cold rings that enhance the collision edge. Since the CMB has more power on large scales, the matched filters on large scales do not look for the large-scale features of the bubble collision signature but rather the transition region near the location where the template goes to zero. Note that the transition region is the best place to look even though the matched filter is constructed for

templates with $z_{\text{crit}} \sim 0\,\mu\text{K}$.

Alternative optimal filters have also been proposed, such as the scale-adaptive filter, defined in Euclidean space by Sanz et al. (2001) and Herranz et al. (2002) and extended to the sphere by Schaefer et al. (2006) and McEwen et al. (2008). Like the matched filter, the scale-adaptive filter minimizes the variance of the filtered field while still providing an unbiased estimate of the source amplitude, but it also imposes a local peak in the filtered field over scale R. Since an additional constraint is imposed when solving the optimization problem that defines the scale-adaptive filter, the SNR for the scale-adaptive filter is inevitably lower than that for the matched filter. However, in the case of (i) a scale-invariant background and (ii) a template profile that changes size simply through a scaling of θ, the peak in the scale-adaptive filter field can help to find sources of unknown size. When criteria (i) and (ii) hold, the scale-adaptive filter for a given source size can be constructed by scaling the scale-adaptive filter for a source of a different size. A filter of incorrect size (since the underlying size of the source is unknown), and scaled variants of it, may then be applied; the peak imposed in scale when constructing the filter can then be used to estimate the unknown source size. However, neither criterion holds for the case of bubble collision signatures embedded in the CMB. Furthermore, although the scale-adaptive filter has been derived on the sphere by Schaefer et al. (2006) and (McEwen et al., 2008), small-angle approximations are made in these derivations; hence the scale-adaptive filter constraints may break down for sources of very large size, such as the bubble collision signatures of interest. Indeed, we have performed numerical experiments that have shown this to be the case. Consequently, we do not consider the scale-adaptive filter further. The problem of detecting sources of unknown size is considered further in Sec. B.3.3.

Signal-to-noise ratio comparison

We compare the SNR for the matched filter, which by definition is optimal, with the SNR for needlets and the unfiltered field. For a arbitrary filter Ψ_R, such as needlets, the SNR defined by Eq. (B.4) becomes

$$\Gamma_R^{\Psi} = \frac{A \sum_{\ell m} \tau_{\ell m} (\Psi_R)^*_{\ell m}}{\sqrt{\sum_{\ell m} C_\ell \left| (\Psi_R)_{\ell m} \right|^2}} \, ,$$

where $(\Psi_R)_{\ell m} = \langle \Psi_R, Y_{\ell m} \rangle$. For the case of the matched filter this expression reduces to (McEwen et al., 2008)

$$\Gamma_R^{\text{MF}} = \alpha^{1/2} \, A \, .$$

(a) $\theta_{\mathrm{crit}} = 5°$

(b) $\theta_{\mathrm{crit}} = 10°$

(c) $\theta_{\mathrm{crit}} = 20°$

(d) $\theta_{\mathrm{crit}} = 30°$

(e) $\theta_{\mathrm{crit}} = 60°$

(f) $\theta_{\mathrm{crit}} = 90°$

Figure B.2: Matched filters optimized to bubble collision signatures of varying size embedded in a ΛCDM CMB background defined by the best-fit WMAP7+BAO+H0 power spectrum.

Finally, we also consider the SNR of the unfiltered field, defined by the ratio of its mean and dispersion at the location of a source, given by

$$\Gamma^{\text{orig}} = \frac{A \sum_{\ell m} \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} \, \tau_{\ell m}}{\sqrt{\sum_\ell \frac{2\ell+1}{4\pi} C_\ell}} \, .$$

The SNRs computed for bubble collision signatures of varying size embedded in a CMB background defined by the ΛCDM best-fit WMAP7+BAO+H0 power spectrum are plotted in Fig. B.3 (a). Notice the superiority of the matched filter to both needlets and the original unfiltered field.

B.3.3 Candidate bubble collision detection

We have selected the optimal matched filter as the filter of choice, since the matched filter optimizes the SNR of the filtered field at the position of a source, but thus far we have only considered source profiles of known size. Here we describe an algorithm using the matched filter to detect multiple sources of unknown and differing size. The algorithm proceeds as follows.

1. Construct matched filters optimized to the source signatures for a grid of scales, *i.e.* $R \in \{\theta_{\text{crit}}^k\}_{k=1}^{N_{\theta_{\text{crit}}}}$.

2. Filter the sky with the matched filter for each scale R.

3. Compute significance maps

$$S_R(\theta, \phi) = \frac{|F_R(\theta, \phi) - \mu_R(\theta, \phi)|}{\sigma_R(\theta, \phi)} \, , \tag{B.6}$$

   for each filter scale R. The mean and dispersion of the filtered field is computed over realisations of the noise process in the absence of sources.

4. Threshold the significance maps for each filter scale R, setting all values of $S_R(\theta, \phi) < N_{\sigma_R}$ to zero.

5. Find localized peaks in the thresholded significance maps for each filter scale R and associate each with a potential detection of a source.

6. For each potential detection at a given scale R, look across adjacent scales $R_{\text{adj}} \in \{R_{\text{adj}} \in \{\theta_{\text{crit}}^k\}_{k=1}^{N_{\theta_{\text{crit}}}} : |R_{\text{adj}} - R| \leq \theta_{\text{adj}}\}$ and eliminate the potential detection if a stronger potential detection is made on an adjacent scale. Potential detections are eliminated as follows. If adjacent scales contain an overlapping non-zero thresholded region, and if the pixel with the maximum absolute value of the filtered field in the thresholded region is the same sign

(a) Known source size

(b) Unknown source size

Figure B.3: SNRs of bubble collision signatures of varying size with amplitude $A = z_0 = 100\ \mu$K embedded in a ΛCDM CMB background defined by the best-fit WMAP7 + BAO + H0 power spectrum. SNR curves are plotted for matched filters (solid blue curve), needlets 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). Notice the clear superiority of the matched filter. In panel (b) SNR curves for the matched filters constructed at a given scale and applied at all other scales are also shown (light 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 θ_{crit} grid is sampled sufficiently densely, the matched filters remain superior to needlets.

as the corresponding value at the current scale, but greater in magnitude, then discard the potential detection at the current scale. Otherwise retain the potential detection and classify it as a detected source.

7. For all detected sources, estimate the parameters of the source size, location and amplitude, using 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 (McEwen et al., 2008) (which in turn relies on the codes S2 (McEwen et al., 2007) and HEALPix (Górski et al., 2005)), while the COMB code (McEwen et al., 2008) has been used to simulate bubble collisions 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 theoretical guarantee that the peak in the filtered field across scales will coincide with the scale of the unknown source. Nevertheless, for bubble collision signatures embedded in the CMB we have found, through numerical simulations, that there is indeed such a peak at the scale of an underlying source, as illustrated in Fig. B.4. Thus, the algorithm outlined above is an effective approach to detecting multiple bubble collision signatures of unknown and differing size. In situations where a peak does not occur at the scale of an underlying source, numerical simulations may be performed to fit the curve of the filtered field across scales to an underlying source size. The algorithm outlined above would therefore remain applicable, with only minor alterations.

Although this algorithm considers a grid of candidate scales $R \in \{\theta_{\mathrm{crit}}^k\}_{k=1}^{N_{\theta_{\mathrm{crit}}}}$, it is likely that a source may exist at scales between the samples of the grid. It is thus important to examine how sensitive the matched filter is to small errors in the source size. In Fig. B.3 (b) we plot SNR curves for matched filters constructed on the grid of candidate scales for bubble collision signatures embedded in the CMB. A degradation in the SNR away from the scale used to construct each filter is clearly apparent; however, provided that the θ_{crit} grid is sampled sufficiently densely, the matched filters remain effective and are superior to needlets.

The algorithm described above has just two parameters. The first is the distance θ_{adj} for which scales are considered to be adjacent, which can be set relative to the grid of candidate sizes. The second parameter is the threshold level N_{σ_R}, which may be allowed to vary for each filter scale R. The threshold levels may be calibrated from simulations in order to allow a manageable number of false detections, while remaining sensitive to weak source signals.

---

[4]S2FIL, S2 and COMB are available from http://www.jasonmcewen.org/, while HEALPix is available from http://healpix.jpl.nasa.gov/.

Figure B.4: Amplitude of the filtered field at the position of a bubble collision signature versus the scale used to construct the corresponding matched filter. The underlying bubble collision signature has parameters $\{z_0, \theta_{\mathrm{crit}}, \theta_0, \phi_0\} = \{100\ \mu\mathrm{K}, 20°, 0°, 0°\}$ and is embedded in a ΛCDM CMB background defined by the best-fit WMAP7 + BAO + H0 power spectrum. The solid curve shows the mean value obtained over 100 CMB realizations, while the error bars show the corresponding standard deviation. Notice that a peak is clearly visible at the scale of the underlying bubble collision signature. Furthermore, the amplitude of the filtered field at the source scale gives an unbiased estimate of the collision amplitude, as imposed through the construction of the matched filter.

B.4 Bubble collision candidates in WMAP 7-year observations

In this section we apply the optimal-filter-based source detection algorithm described in Sec. B.3 to WMAP 7-year observations of the CMB to search for signatures of bubble collisions. Firstly, we construct optimal filters matched to WMAP observations and then calibrate the detection algorithm on a realistic WMAP end-to-end simulation. We then study the sensitivity of the optimal-filter-based detection algorithm. Finally, we apply the algorithm to WMAP observations, resulting in the detection of a number of new candidate bubble collision signatures.

B.4.1 Optimal bubble collision filters for WMAP

We analyze foreground-cleaned WMAP 7-year W-band observations since this band has the highest resolution beam (with full-width-half-maximum FWHM = 13.2 arcmin) and suffers from the least foreground contamination (Jarosik et al., 2011). We restrict our analysis to the band-limit $\ell_{\max} = 256$ since this is sufficient to represent the bubble collision signatures of interest, which are relatively large scale. The stochastic background in which the bubble collision signatures live, and that is used to derive matched filters, is defined by the CMB power spectrum, where we assume the best-fit WMAP7+BAO+H0 best-fit ΛCDM power spectrum. The noise considered in the derivation of the matched filter is assumed to be homogenous and isotropic, whereas WMAP observations exhibit anisotropic noise that varies over the sky. We therefore neglect WMAP noise when constructing optimal filters. This approximation is valid since the W-band instrumental noise is subdominant relative to the CMB contribution in the harmonic region of interest ($\ell_{\max} \leq 256$).[5]

The optimal filters matched to WMAP W-band observations are then computed by Eq. (B.5), where the noise power spectrum C_ℓ is given by the CMB spectrum, and the harmonic coefficients of the normalized template profile $\tau_{\ell m}$ are modulated by the Legendre coefficients of an azimuthally-symmetric Gaussian beam with FWHM = 13.2 arcmin. The matched filters computed in this setting are very similar to those displayed in Fig. B.2, that were computed in the absence of a beam.[6]

For the algorithm to detect candidate bubble collision signatures of unknown and varying size described in Sec. B.3.3, we must construct matched filters for a grid of scales. We consider

---

[5]We have tested the validity of this assumption by successfully detecting synthetic bubble collision signatures embedded in simulated WMAP observations that do include anisotropic noise.

[6]The Gaussian beam employed in this work is an approximation to the true W-band beam (Jarosik et al., 2011). As the matched filters computed in the absence of a beam are very similar to those computed with a Gaussian beam, any effects due to the approximated beam are negligible.

the scales $R \in \{$ 1°, 1.5°, 2°, 3°, 4°, 5°, 6°, 7°, 8°, 9°, 10°, 12°, 14°, 16°, 18°, 20°, 22°, 24°, 26°, 28°, 30°, 35°, 40°, 45°, 50°, 55°, 60°, 65°, 70°, 75°, 80°, 85°, 90°$\}$. The SNR curves for the matched filters constructed for these scales are shown in Fig. B.3 (b) (albeit in the absence of a beam, although the SNR curves do not change markedly when these effects are included). This grid of scales is thus sufficiently sampled to ensure that the matched filters remain effective for scales between the samples of the grid.

B.4.2 Calibration

It is necessary to calibrate the optimal-filter-based bubble collision detection algorithm to realistic WMAP observations. Throughout the calibration we apply the WMAP KQ75 mask (Gold et al., 2011) since we will adopt this conservative mask when analyzing the WMAP data. Firstly, for each scale R, we use 3,000 Gaussian CMB WMAP simulations with W-band beam and anisotropic instrumental noise to compute the mean and dispersion of the filtered field in the absence of sources, as required to compute significance maps of each filtered field through Eq. (B.6). Based on the sampling of the grid of scales we set the adjacency parameter to $\theta_{\mathrm{adj}} = 5°$. We then calibrate the threshold levels N_{σ_R} for each scale R from a realistic WMAP simulation that does not contain bubble collision signatures. The thresholds are chosen to allow a manageable number of false detections while remaining sensitive to weak bubble collision signatures. For this calibration we use a complete end-to-end simulation of the WMAP experiment provided by the WMAP Science Team (Gold et al., 2011). 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 is obtained separately from simulated sources including diffuse Galactic foregrounds, CMB fluctuations, 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 two 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). When running the fully-calibrated candidate bubble collision detection algorithm on the WMAP W-band end-to-end simulation, 13 false detections are made (note that this is an identical number of false detections to that obtained using needlets (Feeney et al., 2011a,b)). Processing a single map through the algorithm, including filtering at all 33 scales, requires on the order of seconds on a standard desktop computer.

B.4.3 Sensitivity

Before applying the calibrated candidate bubble collision detection algorithm to WMAP observations, we first assess its sensitivity by applying it to simulated observations where bubble collision signatures are present. We repeat the sensitivity analysis performed by Feeney et al. (2011a,b), where we lay down known bubble collision signatures on low-noise and high-noise regions of the sky, given by locations $(\theta_0, \phi_0) = (57.7°, 99.2°)$ and $(\theta_0, \phi_0) = (56.6°, 193.0°)$ respectively, where throughout we use Galactic coordinates. For each collision scale and amplitude that we consider, in each of the low-noise and high-noise regions, we simulate three Gaussian CMB WMAP W-band observations. We then run the calibrated bubble collision detection algorithm on these six simulations. If the underlying bubble collision signature is detected in all simulations, we classify the amplitude and scale parameter pair as living in an exclusion region. If the underlying bubble collision is detected in some but not all simulations, we classify the parameter pair as living in a sensitivity region. If the underlying bubble collision is not detected in any simulation, we classify the parameter pair as living in an unprobed region. These regions describe the sensitivity of the bubble collision detection algorithm and are plotted in Fig. B.5 for a range of scale and amplitude parameter pairs.

Bubble collision signatures that lie in exclusion regions would certainly be detected by the optimal-filter-based bubble collision detection algorithm provided they were not significantly masked, while collision signatures that lie in sensitivity regions would be detected if they were in a favorable location on the sky. When compared to the exclusion and sensitivity regions recovered using needlets (Feeney et al., 2011a,b), the regions recovered using optimal filters are extended to lower temperatures by a factor of ~ 1.7 in ΔT for scales $\theta_{\rm crit} \sim 10°$ and most likely further for larger scales (note that the regions plotted in (Feeney et al., 2011a,b) are for $\Delta T/T_0$, where T_0 is the average temperature of the CMB, while here they are plotted for ΔT). Optimal filters thus provide an enhancement in sensitivity by a factor of approximately two when compared with needlets, in line with expectations from the SNR curves plotted in Fig. B.3. This improvement in sensitivity will be important for uncovering the necessarily weak bubble collision signatures that may be embedded in CMB observations.

B.4.4 Candidate bubble collisions

The calibrated bubble collision detection algorithm is applied to foreground-cleaned WMAP 7-year W-band observations (Jarosik et al., 2011), with the conservative KQ75 mask applied (Gold et al., 2011). Sixteen candidate bubble collision signatures are detected. The WMAP W-band data that are analyzed and the detected bubble collision candidates are plotted on the

Figure B.5: Exclusion (black) and sensitivity (grey) regions for the optimal-filter-based bubble collision detection algorithm. Bubble collision signatures that lie in exclusions regions would certainly be detected by the algorithm provided they were not significantly masked, while collision signatures that lie in sensitivity regions would be detected if they were in a favorable location on the sky.

full-sky in Fig. B.6. A list of the parameters recovered for each detected candidate is given in Table B.1, where the bubble collision candidate labels match those of Fig. B.6 (c). In Table B.1 we also give the significance level of each detection and state whether a feature with similar parameters was detected using needlets (Feeney et al., 2011a,b). We detect eight new candidate bubble collisions that have not been reported previously.

As a very preliminary analysis to check that residual foreground contributions are not responsible for the detected candidate bubble collision signatures, we also apply the bubble collision detection algorithm to the foreground-cleaned V-band and Q-band WMAP 7-year observations. Since foreground contributions are frequency-dependent, one would expect a large difference between the regions detected on different bands if they were due to foreground contributions. Whether each candidate bubble collision signature is detected in the other WMAP bands is listed in the final two columns of Table B.1. All of the new regions detected in the W-band are detected in at least one of the other bands, suggesting residual foregrounds are unlikely to be responsible for the new bubble collision candidates that we detect.

Let us remark that the combination of bubble collision candidates with labels 14 and 15 look somewhat like a dipole contribution. However, this resemblance is likely to be a coincidence: we know that the matched filters on these large scales enhance ring-like features (see Fig. B.2). Indeed, since the prior on the expected angular size of bubble collision signatures in the CMB is peaked at 90° (Freivogel et al., 2009), very large candidate bubble collisions are of particular interest. A subsequent Bayesian analysis, following the method of Feeney et al. (2011a,b), will be able to discriminate whether these features are spurious ΛCDM fluctuations, or else provide evidence for the bubble collision hypothesis.

(a) WMAP 7-year W-band observations

(b) Candidate bubble collision signatures

(c) Labelled candidate bubble collision signatures

Figure B.6: WMAP data analyzed by the bubble collision detection algorithm and the resulting candidate bubble collision signatures detected (in units of mK). In panels (a) and (b) the conservative KQ75 mask is applied. Full-sky maps are plotted using the Mollweide projection.

Table B.1: Candidate bubble collisions detected in WMAP 7-year W-band observations.

Label	Bubble collision parameters					Detected previously	Detected in other bands	
	z_0 (mK)	θ_0 (°)	ϕ_0 (°)	θ_{crit} (°)	$S_{\theta_{crit}}(\theta_0, \phi_0)$		V-band	Q-band
0	0.24	119.0	304.5	1.5	4.25	N	Y	N
1	0.20	78.3	176.5	2	4.15	N	N	Y
2	0.20	112.3	264.4	2	4.08	Y	Y	Y
3	-0.19	145.1	33.0	2	4.05	Y	N	N
4	-0.17	169.0	187.5	3	4.26	Y	Y	Y
5	0.17	72.4	150.8	3	4.02	Y	Y	Y
6	-0.16	167.2	268.7	4	4.56	Y	Y	Y
7	-0.16	147.4	207.1	5	4.67	Y	Y	Y
8	0.15	123.2	321.3	5	4.43	Y	Y	Y
9	0.14	62.7	220.4	7	4.39	N	Y	Y
10	0.11	136.6	172.6	20	3.94	Y	Y	Y
11	-0.09	127.2	216.9	26	3.07	N	N	Y
12	0.09	116.3	31.6	35	3.33	N	Y	N
13	0.10	136.6	282.0	40	3.07	N	N	Y
14	0.15	69.6	62.6	85	3.03	N	Y	N
15	-0.16	88.5	277.7	90	3.11	N	Y	Y

B.5 Conclusions

The problem of detecting the existence of a population of sources embedded in the CMB is of widespread interest. The most unambiguous method of doing so is through a direct evaluation of the full posterior probability distribution of the global parameters of the theory giving rise to the source population. However, such an approach is computationally impractical for large datasets, such as WMAP and Planck. A method to approximate the full posterior has been developed recently by Feeney et al. (2011a,b). This approach requires preprocessing of the data to recover a set of candidate sources which are most likely to give the largest contribution to the likelihood. The preprocessing stage of this method is thus crucial to its overall effectiveness. Previously needlets were used for candidate source detection (Feeney et al., 2011a,b). In this article we have developed a new algorithm, based on optimal filtering, to detect candidate sources of unknown and differing angular sizes embedded in full-sky observations of the CMB.

This method is optimal in the sense that no other filter-based approach can provide a superior enhancement of the source contribution. However, as we have emphasized, the parameters of our algorithm are set to allow some false detections: there is no guarantee that the candidates picked out are the signatures of bubble collisions. The filters will also respond to similar temperature patterns resulting from rare ΛCDM fluctuations. A further Bayesian model selection step (implementing Occam's razor via a self-consistent penalty for extra model parameters) is required to determine the most likely explanation for the data – be it a bubble collision, a rare statistical fluctuation of ΛCDM or something else entirely.

Although our source detection algorithm has general applicability, in this case we have applied it to the problem of detecting candidate bubble collision signatures in WMAP 7-year observations, where we have demonstrated its superiority. After calibrating our algorithm on a realistic WMAP end-to-end simulation, we have shown both theoretically and through simulations that it provides an enhancement in sensitivity over the previous needlet approach by a factor of approximately two, for an identical number of false detections on the WMAP end-to-end simulation. Applying our algorithm to WMAP 7-year observations, we detect eight candidate bubble collision signatures that have not been reported previously.

In a follow-up analysis, we intend to compute the full posterior probability distribution of the number of bubble collision signatures in WMAP data using the method developed by Feeney et al. (2011a,b), in light of these new candidate bubble collision signatures. However, this method was previously restricted to candidate collisions of size $\theta_{\mathrm{crit}} \leq 11°$ due to computational memory requirements, while we have detected a number of candidate bubble collision signatures at larger scales. To handle these large candidate bubble collision signatures, an adaptive-resolution

refinement of the method has been developed which processes each candidate at the highest resolution possible given its size and the available computational resources. It was previously shown that the WMAP 7-year data do not warrant augmenting ΛCDM with bubble collisions (Feeney et al., 2011a,b). However, the enhanced sensitivity of our optimal-filter-based candidate collision detection algorithm will improve the accuracy of the approximated posterior distribution, and has the potential to uncover evidence for bubble collisions in WMAP observations of the CMB, as well as in next-generation datasets.

B.6 Acknowledgements

Bibliography

C. E. Aalseth, P. S. Barbeau, N. S. Bowden, B. Cabrera-Palmer, J. Colaresi, J. I. Collar, S. Dazeley, P. de Lurgio, J. E. Fast, N. Fields, C. H. Greenberg, T. W. Hossbach, M. E. Keillor, J. D. Kephart, M. G. Marino, H. S. Miley, M. L. Miller, J. L. Orrell, D. C. Radford, D. Reyna, O. Tench, T. D. Van Wechel, J. F. Wilkerson, and K. M. Yocum. Results from a search for light-mass dark matter with a p-type point contact germanium detector. *Phys. Rev. Lett.*, 106:131301, Mar 2011. doi: 10.1103/PhysRevLett.106.131301. URL `http://link.aps.org/doi/10.1103/PhysRevLett.106.131301`.

Niayesh Afshordi, Anze Slosar, and Yi Wang. A Theory of a Spot. 2010.

Anthony Aguirre. Eternal Inflation, past and future. In *Beyond the Big Bang*. Springer, 2008.

Anthony Aguirre and Matthew C Johnson. Towards observable signatures of other bubble universes II: Exact solutions for thin-wall bubble collisions. *Phys. Rev.*, D77:123536, 2008. doi: 10.1103/PhysRevD.77.123536.

Anthony Aguirre and Matthew C. Johnson. A status report on the observability of cosmic bubble collisions. 2009.

Anthony Aguirre, Matthew C Johnson, and Assaf Shomer. Towards observable signatures of other bubble universes. *Phys. Rev.*, D76:063509, 2007.

Anthony Aguirre, Matthew C. Johnson, and Martin Tysanner. Surviving the crash: assessing the aftermath of cosmic bubble collisions. *Phys. Rev.*, D79:123514, 2009. doi: 10.1103/PhysRevD.79.123514.

Andreas Albrecht and Paul J. Steinhardt. Cosmology for grand unified theories with radiatively induced symmetry breaking. *Phys. Rev. Lett.*, 48:1220–1223, Apr 1982. doi: 10.1103/PhysRevLett.48.1220. URL `http://link.aps.org/doi/10.1103/PhysRevLett.48.1220`.

Andreas Albrecht, Richard A. Battye, and James Robinson. The Case against scaling defect

models of cosmic structure formation. *Phys. Rev. Lett.*, 79:4736–4739, 1997. doi: 10.1103/ PhysRevLett.79.4736.

Andreas Albrecht, Richard A. Battye, and James Robinson. Detailed study of defect models for cosmic structure formation. *Phys. Rev.*, D59:023508, 1999. doi: 10.1103/PhysRevD.59.023508.

R. Allahverdi, R. Brandenberger, F.-Y. Cyr-Racine, and A. Mazumdar. Reheating in Inflationary Cosmology: Theory and Applications. *Annual Review of Nuclear and Particle Science*, 60: 27–51, November 2010. doi: 10.1146/annurev.nucl.012809.104511.

Stephen Amsel, Joshua Berger, and Robert H. Brandenberger. Detecting Cosmic Strings in the CMB with the Canny Algorithm. *JCAP*, 0804:015, 2008. doi: 10.1088/1475-7516/2008/04/ 015.

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.* Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999. ISBN 0-89871-447-8 (paperback).

J. Angle, E. Aprile, F. Arneodo, L. Baudis, A. Bernstein, A. Bolozdynya, P. Brusov, L. C. C. Coelho, C. E. Dahl, L. DeViveiros, A. D. Ferella, L. M. P. Fernandes, S. Fiorucci, R. J. Gaitskell, K. L. Giboni, R. Gomez, R. Hasty, L. Kastens, J. Kwong, J. A. M. Lopes, N. Madden, A. Manalaysay, A. Manzur, D. N. McKinsey, M. E. Monzani, K. Ni, U. Oberlack, J. Orboeck, G. Plante, R. Santorelli, J. M. F. dos Santos, P. Shagin, T. Shutt, P. Sorensen, S. Schulte, C. Winant, and M. Yamashita. First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory. *Phys. Rev. Lett.*, 100:021303, Jan 2008. doi: 10.1103/PhysRevLett.100.021303. URL http://link.aps.org/doi/10.1103/PhysRevLett.100.021303.

R. Aurich and S. Lustig. Can one reconstruct masked CMB sky? *Mon. Not. Roy. Astron. Soc.*, 411:124–136, 2011. doi: 10.1111/j.1365-2966.2010.17667.x.

P. Baldi, G. Kerkyacharian, D. Marinucci, and D. Picard. Asymptotics for spherical needlets. *ArXiv Mathematics e-prints*, June 2006.

R. B. Barreiro, J. L. Sanz, D. Herranz, and E. Martínez-González. Comparing filters for the detection of point sources. 342:119–133, 2003. doi: 10.1046/j.1365-8711.2003.06520.x.

Thomas Bayes and Richard Price. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions*, 53:370–418, 1763. doi: 10.1098/rstl.1763.0053. URL http://rstl.royalsocietypublishing.org/content/53/370.short.

C. L. Bennett, M. Bay, M. Halpern, G. Hinshaw, C. Jackson, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, D. T. Wilkinson, E. Wollack, and E. L. Wright. The Microwave Anisotropy Probe Mission. *Astrophys. J.*, 583:1–23, January 2003a. doi: 10.1086/345346.

C. L. Bennett, R. S. Hill, G. Hinshaw, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, E. L. Wright, M. Halpern, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. Wollack. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission. *Astrophys. J. Suppl. Ser.*, 148:97–117, September 2003b. doi: 10.1086/377252.

C. L. Bennett, R. S. Hill, G. Hinshaw, D. Larson, K. M. Smith, J. Dunkley, B. Gold, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Are There Cosmic Microwave Background Anomalies? *Astrophys. J. Suppl. Ser.*, 192:17, February 2011. doi: 10.1088/0067-0049/192/2/17.

R. Bernabei, P. Belli, F. Cappella, R. Cerulli, C. Dai, A. d'Angelo, H. He, A. Incicchitti, H. Kuang, J. Ma, F. Montecchia, F. Nozzoli, D. Prosperi, X. Sheng, and Z. Ye. First results from DAMA/LIBRA and the combined results with DAMA/NaI. *The European Physical Journal C - Particles and Fields*, 56:333–355, 2008. ISSN 1434-6044. URL `http://dx.doi.org/10.1140/epjc/s10052-008-0662-y`. 10.1140/epjc/s10052-008-0662-y.

Pawel Bielewicz, K. M. Gorski, and A. J. Banday. Low order multipole maps of CMB anisotropy derived from WMAP. *Mon. Not. Roy. Astron. Soc.*, 355:1283, 2004. doi: 10.1111/j.1365-2966.2004.08405.x.

M. N. Bremer, J. Silk, L. J. M. Davies, and M. D. Lehnert. A redshift survey towards the CMB Cold Spot. *ArXiv e-prints*, April 2010.

Martin Bucher, Alfred S. Goldhaber, and Neil Turok. An open universe from inflation. *Phys. Rev.*, D52:3314–3337, 1995.

E. F. Bunn, K. B. Fisher, Y. Hoffman, O. Lahav, J. Silk, and S. Zaroubi. Wiener filtering of the COBE Differential Microwave Radiometer data. *Astrophys. J.*, 432:L75–L78, September 1994. doi: 10.1086/187515.

J Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, 1986. ISSN 0162-8828. doi: http://dx.doi.org/10.1109/TPAMI.1986.4767851.

Spencer Chang, Matthew Kleban, and Thomas S. Levi. When Worlds Collide. *JCAP*, 0804:034, 2008. doi: 10.1088/1475-7516/2008/04/034.

Spencer Chang, Matthew Kleban, and Thomas S. Levi. Watching Worlds Collide: Effects on the CMB from Cosmological Bubble Collisions. *JCAP*, 0904:025, 2009. doi: 10.1088/1475-7516/2009/04/025.

W. Z. Chao. Gravitational effects in bubble collisions. *Phys. Rev. D*, 28:1898–1906, October 1983. doi: 10.1103/PhysRevD.28.1898.

Alain Coc, Elisabeth Vangioni-Flam, Pierre Descouvemont, Abderrahim Adahchour, and Carmen Angulo. Updated Big Bang nucleosynthesis confronted to WMAP observations and to the abundance of light elements. *Astrophys. J.*, 600:544–552, 2004. doi: 10.1086/380121.

Sidney R. Coleman. The Fate of the False Vacuum. 1. Semiclassical Theory. *Phys. Rev.*, D15: 2929–2936, 1977. doi: 10.1103/PhysRevD.15.2929.

Sidney R. Coleman and Frank De Luccia. Gravitational Effects on and of Vacuum Decay. *Phys. Rev.*, D21:3305, 1980. doi: 10.1103/PhysRevD.21.3305.

Craig J. Copi, Dragan Huterer, Dominik J. Schwarz, and Glenn D. Starkman. Bias in low-multipole CMB reconstructions. 2011.

Neil J. Cornish, David N. Spergel, and Glenn D. Starkman. Circles in the Sky: Finding Topology with the Microwave Background Radiation. *Class. Quant. Grav.*, 15:2657–2670, 1998. doi: 10.1088/0264-9381/15/9/013.

R. T. Cox. Probability, frequency and reasonable expectation. *Am. J. Th. Phys*, 14:1, 1946. doi: 10.1119/1.1990764.

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. SPIDER: a balloon-borne large-scale CMB polarimeter. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7010, August 2008. doi: 10.1117/12.787446.

M. Cruz, E. Martinez-Gonzalez, P. Vielva, and L. Cayon. Detection of a non-Gaussian Spot in WMAP. *Mon. Not. R. Astron. Soc.*, 356:29–40, 2005. doi: 10.1111/j.1365-2966.2004.08419. x/abs/.

M. Cruz, N. Turok, P. Vielva, E. Martinez-Gonzalez, and M. Hobson. A Cosmic Microwave Background feature consistent with a cosmic texture. *Science*, 318:1612–1614, 2007. doi: 10.1126/science.1148694.

M. Cruz, E. Martínez-Gonzalez, P. Vielva, J. M. Diego, M. Hobson, and N. Turok. The CMB cold spot: texture, cluster or void? *Mon. Not. R. Astron. Soc.*, 390:913–919, November 2008. doi: 10.1111/j.1365-2966.2008.13812.x.

M. Cruz, E. Martinez-Gonzalez, and P. Vielva. The WMAP cold spot. In *Highlights of Spanish Astrophysics V*, 2010.

B. Czech, M. Kleban, K. Larjo, T. S. Levi, and K. Sigurdson. Polarizing bubble collisions. *JCAP*, 12:023, December 2010. doi: 10.1088/1475-7516/2010/12/023.

Alex Dahlen. Odds of observing the multiverse. *Phys.Rev.*, D81:063501, 2010. doi: 10.1103/PhysRevD.81.063501.

Rebecca J. Danos and Robert H. Brandenberger. Canny Algorithm, Cosmic Strings and the Cosmic Microwave Background. *Int. J. Mod. Phys.*, D19:183–217, 2010. doi: 10.1142/S0218271810016324.

S. Das, T. A. Marriage, P. A. R. Ade, P. Aguirre, M. Amiri, J. W. Appel, L. F. Barrientos, E. S. Battistelli, J. R. Bond, B. Brown, B. Burger, J. Chervenak, M. J. Devlin, S. R. Dicker, W. Bertrand Doriese, J. Dunkley, R. Dünner, T. Essinger-Hileman, R. P. Fisher, J. W. Fowler, A. Hajian, M. Halpern, M. Hasselfield, C. Hernández-Monteagudo, G. C. Hilton, M. Hilton, A. D. Hincks, R. Hlozek, K. M. Huffenberger, D. H. Hughes, J. P. Hughes, L. Infante, K. D. Irwin, J. Baptiste Juin, M. Kaul, J. Klein, A. Kosowsky, J. M. Lau, M. Limon, Y.-T. Lin, R. H. Lupton, D. Marsden, K. Martocci, P. Mauskopf, F. Menanteau, K. Moodley, H. Moseley, C. B. Netterfield, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, B. Partridge, B. Reid, N. Sehgal, B. D. Sherwin, J. Sievers, D. N. Spergel, S. T. Staggs, D. S. Swetz, E. R. Switzer, R. Thornton, H. Trac, C. Tucker, R. Warne, E. Wollack, and Y. Zhao. The Atacama Cosmology Telescope: A Measurement of the Cosmic Microwave Background Power Spectrum at 148 and 218 GHz from the 2008 Southern Survey. *Astrophys. J.*, 729:62, March 2011. doi: 10.1088/0004-637X/729/1/62.

M. Davis, J. Huchra, D. W. Latham, and J. Tonry. A survey of galaxy redshifts. II - The large scale space distribution. *Astrophys. J.*, 253:423–445, February 1982. doi: 10.1086/159646.

T. M. Davis and C. H. Lineweaver. Expanding Confusion: Common Misconceptions of Cosmological Horizons and the Superluminal Expansion of the Universe. *Publications of the Astronomical Society of Australia*, 21:97–109, 2004. doi: 10.1071/AS03040.

P. de Bernardis, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, B. P. Crill, G. De Gasperis, P. C. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, A. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, S. Prunet, S. Rao, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforna, and N. Vittorio. A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404:955–959, April 2000.

Angelica de Oliveira-Costa and Max Tegmark. CMB multipole measurements in the presence of foregrounds. *Phys. Rev.*, D74:023005, 2006. doi: 10.1103/PhysRevD.74.023005.

Andrea De Simone and Michael P. Salem. The distribution of Ω_k from the scale-factor cutoff measure. *Phys.Rev.*, D81:083527, 2010. doi: 10.1103/PhysRevD.81.083527.

R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. *Astrophys. J.*, 142:414–419, July 1965. doi: 10.1086/148306.

J. Dickey. The weighted likelihood ratio, linear hypotheses on normal location parameters. *Ann. Mathemat. Statist.*, 42:204, 1971.

A. G. Doroshkevich, P. D. Naselsky, O. V. Verkhodanov, D. I. Novikov, V. I. Turchaninov, I. D. Novikov, P. R. Christensen, and L. Y. Chiang. Gauss–Legendre Sky Pixelization (GLESP) for CMB maps. *Int. J. Mod. Phys. D.*, 14(2):275–290, 2005.

J. R. Driscoll and D. M. Jr. Healy. Computing Fourier transforms and convolutions on the sphere. *Advances in Applied Mathematics*, 15:202–250, 1994.

Cora Dvorkin, Hiranya V. Peiris, and Wayne Hu. Testable polarization predictions for models of CMB isotropy anomalies. *Phys. Rev.*, D77:063008, 2008. doi: 10.1103/PhysRevD.77.063008.

Richard Easther and Hiranya V. Peiris. Bayesian Analysis of Inflation II: Model Selection and Constraints on Reheating. *Phys. Rev.*, D85:103533, 2012. doi: 10.1103/PhysRevD.85.103533.

Richard Easther, John T. Giblin, Jr, Lam Hui, and Eugene A. Lim. A New Mechanism for Bubble Nucleation: Classical Transitions. *Phys. Rev.*, D80:123519, 2009. doi: 10.1103/PhysRevD.80.123519.

G. Efstathiou, Y.-Z. Ma, and D. Hanson. Large-angle correlations in the cosmic microwave background. *Mon. Not. R. Astron. Soc.*, 407:2530–2542, October 2010. doi: 10.1111/j.1365-2966.2010.17081.x.

Adrienne L. Erickcek, Sean M. Carroll, and Marc Kamionkowski. Superhorizon Perturbations and the Cosmic Microwave Background. *Phys. Rev.*, D78:083012, 2008. doi: 10.1103/PhysRevD.78.083012.

S. M. Feeney, M. C. Johnson, D. J. Mortlock, and H. V. Peiris. Robust Constraint on Cosmic Textures from the Cosmic Microwave Background. *Physical Review Letters*, 108(24):241301, June 2012. doi: 10.1103/PhysRevLett.108.241301.

Stephen M. Feeney, Matthew C. Johnson, Daniel J. Mortlock, and Hiranya V. Peiris. First Observational Tests of Eternal Inflation: Analysis Methods and WMAP 7-Year Results. *Phys. Rev.*, D84:043507, 2011a. doi: 10.1103/PhysRevD.84.043507.

Stephen M. Feeney, Matthew C. Johnson, Daniel J. Mortlock, and Hiranya V. Peiris. First Observational Tests of Eternal Inflation. *Phys. Rev. Lett.*, 107:071301, 2011b. doi: 10.1103/PhysRevLett.107.071301.

Stephen M. Feeney, Hiranya V. Peiris, and Andrew Pontzen. Avoiding bias in reconstructing the largest observable scales from partial-sky data. *Phys. Rev.*, D84:103002, 2011c. doi: 10.1103/PhysRevD.84.103002.

F. Feroz, M. P. Hobson, and M. Bridges. Multinest: an efficient and robust bayesian inference tool for cosmology and particle physics. *Mon. Not. R. Astron. Soc.*, 398:1601–1614, 2009.

Farhan Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis. *Mon. Not. R. Astron. Soc.*, 384:449, 2008. doi: 10.1111/j.1365-2966.2007.12353.x.

Ben Freivogel. Making predictions in the multiverse. *Class.Quant.Grav.*, 28:204007, 2011. doi: 10.1088/0264-9381/28/20/204007.

Ben Freivogel, Matthew Kleban, Maria Rodriguez Martinez, and Leonard Susskind. Observational consequences of a landscape. *JHEP*, 03:039, 2006.

Ben Freivogel, Matthew Kleban, Alberto Nicolis, and Kris Sigurdson. Eternal Inflation, Bubble Collisions, and the Disintegration of the Persistence of Memory. *JCAP*, 0908:036, 2009. doi: 10.1088/1475-7516/2009/08/036.

A. Friedman. Über die Krümmung des Raumes. *Zeitschrift fur Physik*, 10:377–386, December 1922. doi: 10.1007/BF01332580.

Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, K. Kobayashi, Y. Kobayashi, Y. Koshio, K. Martens, M. Miura, M. Nakahata, S. Nakayama, A. Okada, M. Oketa, K. Okumura, M. Ota, N. Sakurai, M. Shiozawa, Y. Suzuki, Y. Takeuchi, Y. Totsuka, S. Yamada, M. Earl, A. Habig, J. T. Hong, E. Kearns, S. B. Kim, M. Masuzawa, M. D. Messier, K. Scholberg, J. L. Stone, L. R. Sulak, C. W. Walter, M. Goldhaber, T. Barszczak, W. Gajewski, P. G. Halverson, J. Hsu, W. R. Kropp, L. R. Price, F. Reines, H. W. Sobel, M. R. Vagins, K. S. Ganezer, W. E. Keig, R. W. Ellsworth, S. Tasaka, J. W. Flanagan, A. Kibayashi, J. G. Learned, S. Matsuno, V. Stenger, D. Takemori, T. Ishii, J. Kanzaki, T. Kobayashi, K. Nakamura, K. Nishikawa, Y. Oyama, A. Sakai, M. Sakuda, O. Sasaki, S. Echigo, M. Kohama, A. T. Suzuki, T. J. Haines, E. Blaufuss, R. Sanford, R. Svoboda, M. L. Chen, Z. Conner, J. A. Goodman, G. W. Sullivan, M. Mori, J. Hill, C. K. Jung, C. Mauger, C. McGrew, E. Sharkey, B. Viren, C. Yanagisawa, W. Doki, T. Ishizuka, Y. Kitaguchi, H. Koga, K. Miyano, H. Okazawa, C. Saji, M. Takahata, A. Kusano, Y. Nagashima, M. Takita, T. Yamaguchi, M. Yoshida, M. Etoh, K. Fujita, A. Hasegawa, T. Hasegawa, S. Hatakeyama, T. Iwamoto, T. Kinebuchi, M. Koga, T. Maruyama, H. Ogawa, A. Suzuki, F. Tsushima, M. Koshiba, M. Nemoto, K. Nishijima, T. Futagami, Y. Hayato, Y. Kanaya, K. Kaneyuki, Y. Watanabe, D. Kielczewska, R. Doyle, J. George, A. Stachyra, L. Wai, J. Wilkes, and K. Young. Measurements of the solar neutrino flux from super-kamiokande's first 300 days. *Phys. Rev. Lett.*, 81:1158–1162, Aug 1998. doi: 10.1103/PhysRevLett.81.1158. URL `http://link.aps.org/doi/10.1103/PhysRevLett.81.1158`.

Juan Garcia-Bellido. Single-bubble open inflation: An overview. 1997.

Juan Garcia-Bellido, Andrew R. Liddle, David H. Lyth, and David Wands. The Open Universe Grishchuk-Zel'dovich Effect. *Phys. Rev.*, D52:6750–6759, 1995. doi: 10.1103/PhysRevD.52. 6750.

Jaume Garriga, Alan H. Guth, and Alexander Vilenkin. Eternal inflation, bubble collisions, and the persistence of memory. *Phys. Rev.*, D76:123512, 2007. doi: 10.1103/PhysRevD.76.123512.

John T. Giblin, Jr, Lam Hui, Eugene A. Lim, and I-Sheng Yang. How to Run Through Walls: Dynamics of Bubble and Soliton Collisions. *Phys. Rev.*, D82:045019, 2010. doi: 10.1103/PhysRevD.82.045019.

Roberto Gobbetti and Matthew Kleban. Analyzing Cosmic Bubble Collisions. 2012.

B. Gold, N. Odegard, J. L. Weiland, R. S. Hill, A. Kogut, C. L. Bennett, G. Hinshaw, X. Chen, J. Dunkley, M. Halpern, N. Jarosik, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Galactic Foreground Emission. *Astrophys. J. Suppl. Ser.*, 192:15, February 2011. doi: 10.1088/0067-0049/192/2/15.

Christopher Gordon, Wayne Hu, Dragan Huterer, and Tom Crawford. Spontaneous isotropy breaking: A mechanism for CMB multipole alignments. *Phys. Rev. D*, 72(10):103002, Nov 2005. doi: 10.1103/PhysRevD.72.103002.

K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophys. J.*, 622:759–771, April 2005. doi: 10.1086/427976.

J. R. Gott and T. S. Statler. Constraints on the formation of bubble universes. *Phys. Lett.*, B136:157–161, 1984.

J. R. Gott, III. Creation of open universes from de Sitter space. *Nature*, 295:304–306, January 1982. doi: 10.1038/295304a0.

Nicolaas E. Groeneboom, Lotty Ackerman, Ingunn Kathrine Wehus, and Hans Kristian Eriksen. Bayesian analysis of an anisotropic universe model: systematics and polarization. *Astrophys. J.*, 722:452–459, 2010. doi: 10.1088/0004-637X/722/1/452.

Frederic Guilloux, Gillesn Fay, and Jean-Francois Cardoso. Practical wavelet design on the sphere. 2007.

A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, January 1981. doi: 10.1103/PhysRevD.23.347.

A. H. Guth. Eternal inflation and its implications. *Journal of Physics A Mathematical General*, 40:6811–6826, June 2007. doi: 10.1088/1751-8113/40/25/S25.

A. H. Guth and E. J. Weinberg. Cosmological consequences of a first-order phase transition in the SU$_5$ grand unified model. *Phys. Rev. D*, 23:876–885, February 1981. doi: 10.1103/PhysRevD.23.876.

Alan H. Guth and Erick J. Weinberg. Could the Universe Have Recovered from a Slow First Order Phase Transition? *Nucl. Phys.*, B212:321, 1983. doi: 10.1016/0550-3213(83)90307-3.

Martin G. Haehnelt and Max Tegmark. Using the kinematic Sunyaev-Zel'dovich effect to determine the peculiar velocities of clusters of galaxies. 279:545–556, 1996.

S. Hanany, P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, V. V. Hristov, A. H. Jaffe, A. E. Lange, A. T. Lee, P. D. Mauskopf, C. B. Netterfield, S. Oh, E. Pascale, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of $10' - 5°$. *Astrophys. J.*, 545:L5–L9, December 2000. doi: 10.1086/317322.

S. W. Hawking, I. G. Moss, and J. M. Stewart. Bubble Collisions in the Very Early Universe. *Phys. Rev.*, D26:2681, 1982. doi: 10.1103/PhysRevD.26.2681.

D. Herranz, J. L. Sanz, R. B. Barreiro, and E. Martínez-González. Scale-adaptive filters for the detection/separation of compact sources. 580:610–625, 2002.

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. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: The Angular Power Spectrum. *Astrophys. J. Suppl. Ser.*, 148:135–159, September 2003. doi: 10.1086/377225.

G. Hinshaw, M. R. Nolta, C. L. Bennett, R. Bean, O. Doré, M. R. Greason, M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, N. Odegard, S. S. Meyer, L. Page, H. V. Peiris, D. N. Spergel, G. S. Tucker, L. Verde, J. L. Weiland, E. Wollack, and E. L. Wright. Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Temperature Analysis. *Astrophys. J. Suppl. Ser.*, 170:288–334, June 2007. doi: 10.1086/513698.

M. P. Hobson and C. McLachlan. A Bayesian approach to discrete object detection in astronomical data sets. *Mon. Not. R. Astron. Soc.*, 338:765–784, January 2003. doi: 10.1046/j.1365-8711.2003.06094.x.

M. P. Hobson, G. Rocha, and R. S. Savage. *Bayesian source extraction*, pages 167–+. 2010.

J. Hoftuft, H. K. Eriksen, A. J. Banday, K. M. Górski, F. K. Hansen, and P. B. Lilje. Increasing Evidence for Hemispherical Power Asymmetry in the Five-Year WMAP Data. *Astrophys. J.*, 699:985–989, July 2009. doi: 10.1088/0004-637X/699/2/985.

Edwin Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *Proc.Nat.Acad.Sci.*, 15:168–173, 1929.

N. Jarosik, C. L. Bennett, J. Dunkley, B. Gold, M. R. Greason, M. Halpern, R. S. Hill, G. Hinshaw, A. Kogut, E. Komatsu, D. Larson, M. Limon, S. S. Meyer, M. R. Nolta, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. *Astrophys. J. Suppl. Ser.*, 192:14, February 2011. doi: 10.1088/0067-0049/192/2/14.

Matthew C. Johnson and I-Sheng Yang. Escaping the crunch: gravitational effects in classical transitions. *Phys. Rev.*, D82:065023, 2010. doi: 10.1103/PhysRevD.82.065023.

Nick Kaiser and A. Stebbins. Microwave Anisotropy Due to Cosmic Strings. *Nature*, 310:391–393, 1984. doi: 10.1038/310391a0.

R. Keisler, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, R. Williamson, and O. Zahn. A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope. *Astrophys. J.*, 743:28, December 2011. doi: 10.1088/0004-637X/743/1/28.

T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J. Phys. A*, A9:1387–1398, 1976. doi: 10.1088/0305-4470/9/8/029.

Carolyn Kimme, Dana Ballard, and Jack Sklansky. Finding circles by an array of accumulators. *Commun. ACM*, 18(2):120–122, 1975. ISSN 0001-0782. doi: http://doi.acm.org/10.1145/360666.360677.

Matthew Kleban, Thomas S. Levi, and Kris Sigurdson. Observing the Multiverse with Cosmic Wakes. 2011.

E. Komatsu, A. Kogut, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, L. Verde, E. Wollack, and E. L.

Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Tests of Gaussianity. *Astrophys. J. Suppl. Ser.*, 148:119–134, September 2003. doi: 10.1086/377220.

E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *Astrophys. J. Suppl. Ser.*, 192:18, February 2011. doi: 10.1088/0067-0049/192/2/18.

Ely D. Kovetz, Assaf Ben-David, and Nissan Itzhaki. Giant Rings in the CMB Sky. *Astrophys.J.*, 724:374–378, 2010. doi: 10.1088/0004-637X/724/1/374.

Jonathan Kozaczuk and Anthony Aguirre. Multiple Cosmic Collisions and the Microwave Background Power Spectrum. 2012.

Klaus Larjo and Thomas S. Levi. Bubble, Bubble, Flow and Hubble: Large Scale Galaxy Flow from Cosmological Bubble Collisions. *JCAP*, 1008:034, 2010. doi: 10.1088/1475-7516/2010/08/034.

D. Larson, J. Dunkley, G. Hinshaw, E. Komatsu, M. R. Nolta, C. L. Bennett, B. Gold, M. Halpern, R. S. Hill, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, L. Page, K. M. Smith, D. N. Spergel, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Power Spectra and WMAP-derived Parameters. *Astrophys. J. Suppl. Ser.*, 192:16, February 2011. doi: 10.1088/0067-0049/192/2/16.

Antony Lewis and Sarah Bridle. Cosmological parameters from CMB and other data: a Monte-Carlo approach. *Phys. Rev.*, D66:103511, 2002. doi: 10.1103/PhysRevD.66.103511.

A.D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108 (6):389 – 393, 1982. ISSN 0370-2693. doi: 10.1016/0370-2693(82)91219-9. URL `http://www.sciencedirect.com/science/article/pii/0370269382912199`.

Amy S. Lo and Edward L. Wright. Signatures of cosmic strings in the cosmic microwave background. 2005.

T. J. Loredo. Bayesian astrostatistics: a backward look to the future. *ArXiv e-prints*, August 2012.

Bjorn Malte Schafer and Matthias Bartelmann. Detecting sunyaev-zel'dovich clusters with planck. 3. properties of the expected sz-cluster sample. *Mon.Not.Roy.Astron.Soc.*, 377:253–268, 2007. doi: 10.1111/j.1365-2966.2007.11596.x.

D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkyacharian, P. Natoli, D. Picard, and N. Vittorio. Spherical needlets for cosmic microwave background data analysis. *Mon. Not. R. Astron. Soc.*, 383:539–545, January 2008. doi: 10.1111/j.1365-2966.2007.12550.x.

J. C. Mather, E. S. Cheng, D. A. Cottingham, R. E. Eplee, Jr., D. J. Fixsen, T. Hewagama, R. B. Isaacman, K. A. Jensen, S. S. Meyer, P. D. Noerdlinger, S. M. Read, L. P. Rosen, R. A. Shafer, E. L. Wright, C. L. Bennett, N. W. Boggess, M. G. Hauser, T. Kelsall, S. H. Moseley, Jr., R. F. Silverberg, G. F. Smoot, R. Weiss, and D. T. Wilkinson. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *Astrophys. J.*, 420:439–444, January 1994. doi: 10.1086/173574.

J. D. McEwen and Y. Wiaux. A Novel Sampling Theorem on the Sphere. *IEEE Transactions on Signal Processing*, 59:5876–5887, December 2011. doi: 10.1109/TSP.2011.2166394.

J. D. McEwen, M. P. Hobson, D. J. Mortlock, and A. N. Lasenby. Fast directional continuous spherical wavelet transform algorithms. 55(2):520–529, 2007. doi: 10.1109/TSP.2006.887148.

J. D. McEwen, M. P. Hobson, and A. N. Lasenby. Optimal Filters on the Sphere. *IEEE Transactions on Signal Processing*, 56:3813–3823, August 2008. doi: 10.1109/TSP.2008.923198.

J. D. McEwen, S. M. Feeney, M. C. Johnson, and H. V. Peiris. Optimal filters for detecting cosmic bubble collisions. *Phys. Rev. D*, 85(10):103502, May 2012. doi: 10.1103/PhysRevD.85.103502.

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. SPTpol: an instrument for CMB polarization. In B. Young, B. Cabrera, and A. Miller, editors, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 511–514, December 2009. doi: 10.1063/1.3292391.

Michael J. Mortonson, Hiranya V. Peiris, and Richard Easther. Bayesian Analysis of Inflation: Parameter Estimation for Single Field Models. *Phys. Rev.*, D83:043505, 2011. doi: 10.1103/PhysRevD.83.043505.

Viatcheslav F. Mukhanov and G.V. Chibisov. Quantum Fluctuation and Nonsingular Universe. (In Russian). *JETP Lett.*, 33:532–535, 1981.

C. B. Netterfield, M. J. Devlin, N. Jarolik, L. Page, and E. J. Wollack. A Measurement of the Angular Power Spectrum of the Anisotropy in the Cosmic Microwave Background. *Astrophys. J.*, 474:47, January 1997. doi: 10.1086/303438.

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. ACTPol: a polarization-sensitive receiver for the Atacama Cosmology Telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741, July 2010. doi: 10.1117/12.857464.

M. R. Nolta, J. Dunkley, R. S. Hill, G. Hinshaw, E. Komatsu, D. Larson, L. Page, D. N. Spergel, C. L. Bennett, B. Gold, N. Jarosik, N. Odegard, J. L. Weiland, E. Wollack, M. Halpern, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Angular Power Spectra. *Astrophys. J. Suppl. Ser.*, 180:296–305, February 2009. doi: 10.1088/0067-0049/180/2/296.

Jorge Norena, Christian Wagner, Licia Verde, Hiranya V. Peiris, and Richard Easther. Bayesian Analysis of Inflation III: Slow Roll Reconstruction Using Model Selection. *Phys. Rev.*, D86: 023505, 2012. doi: 10.1103/PhysRevD.86.023505.

J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *Bull. Astron. Inst. Neth.*, 6:249, August 1932.

H. V. Peiris, E. Komatsu, L. Verde, D. N. Spergel, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications For Inflation. *Astrophys. J. Suppl. Ser.*, 148:213–231, September 2003. doi: 10.1086/377228.

Ue-Li Pen, David N. Spergel, and Neil Turok. Cosmic structure formation and microwave anisotropies from global field ordering. *Phys. Rev.*, D49:692–729, 1994. doi: 10.1103/PhysRevD.49.692.

A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophys. J.*, 142:419–421, July 1965. doi: 10.1086/148307.

S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and Supernova Cosmology Project. Measurements of Omega and Lambda from 42 High-Redshift Supernovae. *Astrophys. J.*, 517: 565–586, June 1999. doi: 10.1086/307221.

D. Pietrobon, A. Amblard, A. Balbi, P. Cabella, A. Cooray, and D. Marinucci. Needlet detection of features in the WMAP CMB sky and the impact on anisotropies and hemispherical asymmetries. *Phys. Rev. D*, 78(10):103504, November 2008. doi: 10.1103/PhysRevD.78.103504.

Davide Pietrobon, Amedeo Balbi, and Domenico Marinucci. Integrated Sachs-Wolfe effect from the cross-correlation of WMAP 3 year and NVSS: new results and constraints on dark energy. *Phys. Rev.*, D74:043524, 2006. doi: 10.1103/PhysRevD.74.043524.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, M. Baker, A. Balbi, A. J. Banday, and et al. Planck early results. I. The Planck mission. *Astron. Astrophys.*, 536:A1, December 2011. doi: 10.1051/0004-6361/201116464.

Andrew Pontzen and Hiranya V. Peiris. The cut-sky cosmic microwave background is not anomalous. *Phys. Rev.*, D81:103008, 2010. doi: 10.1103/PhysRevD.81.103008.

J. P. Preskill. Cosmological production of superheavy magnetic monopoles. *Physical Review Letters*, 43:1365–1368, November 1979. doi: 10.1103/PhysRevLett.43.1365.

M. J. Rees and D. W. Sciama. Large-scale Density Inhomogeneities in the Universe. *Nature*, 217:511–516, February 1968. doi: 10.1038/217511a0.

B. A. Reid, W. J. Percival, D. J. Eisenstein, L. Verde, D. N. Spergel, R. A. Skibba, N. A. Bahcall, T. Budavari, J. A. Frieman, M. Fukugita, J. R. Gott, J. E. Gunn, Ž. Ivezić, G. R. Knapp, R. G. Kron, R. H. Lupton, T. A. McKay, A. Meiksin, R. C. Nichol, A. C. Pope, D. J. Schlegel, D. P. Schneider, C. Stoughton, M. A. Strauss, A. S. Szalay, M. Tegmark, M. S. Vogeley, D. H.

Weinberg, D. G. York, and I. Zehavi. Cosmological constraints from the clustering of the Sloan Digital Sky Survey DR7 luminous red galaxies. *Mon. Not. R. Astron. Soc.*, 404:60–85, May 2010. doi: 10.1111/j.1365-2966.2010.16276.x.

A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astron. J.*, 116:1009–1038, September 1998. doi: 10.1086/300499.

T. Risbo. Fourier transform summation of Legendre series and D-functions. *Journal of Geodesy*, 70:383–396, July 1996. doi: 10.1007/BF01090814.

Vera C. Rubin and Jr. Ford, W. Kent. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *Astrophys. J.*, 159:379–403, 1970. doi: 10.1086/150317.

R. K. Sachs and A. M. Wolfe. Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. *Astrophys. J.*, 147:73–+, January 1967. doi: 10.1086/148982.

Michael P. Salem. Bubble collisions and measures of the multiverse. *JCAP*, 1201:021, 2012. doi: 10.1088/1475-7516/2012/01/021.

J. L. Sanz, D. Herranz, and E. Martínez-Gónzalez. Optimal Detection of Sources on a Homogeneous and Isotropic Background. *Astrophys. J.*, 552:484–492, May 2001. doi: 10.1086/320550.

Bjoern Malte Schaefer, Christoph Pfrommer, Reinhard M. Hell, and Matthias Bartelmann. Detecting Sunyaev-Zel'dovich clusters with Planck: II. Foreground components and optimised filtering schemes. 370:1713–1736, August 2006. doi: 10.1111/j.1365-2966.2006.10622.x.

B. M. Schäfer, C. Pfrommer, R. M. Hell, and M. Bartelmann. Detecting Sunyaev-Zel'dovich clusters with Planck - II. Foreground components and optimized filtering schemes. *Mon. Not. R. Astron. Soc.*, 370:1713–1736, August 2006. doi: 10.1111/j.1365-2966.2006.10622.x.

S. Scodeller, Ø. Rudjord, F. K. Hansen, D. Marinucci, D. Geller, and A. Mayeli. Introducing Mexican Needlets for CMB Analysis: Issues for Practical Applications and Comparison with Standard Needlets. *Astrophys. J.*, 733:121, June 2011. doi: 10.1088/0004-637X/733/2/121.

Yasuhiro Sekino, Stephen Shenker, and Leonard Susskind. On the Topological Phases of Eternal Inflation. *Phys. Rev.*, D81:123515, 2010. doi: 10.1103/PhysRevD.81.123515.

J. Silk. Cosmic Black-Body Radiation and Galaxy Formation. *Astrophys. J.*, 151:459, February 1968. doi: 10.1086/149449.

J. Skilling. Nested sampling. In *AIP Conference Proceedings of the 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 735, pages 395–405, 2004.

Anze Slosar, Christopher Hirata, Uros Seljak, Shirley Ho, and Nikhil Padmanabhan. Constraints on local primordial non-Gaussianity from large scale structure. *JCAP*, 0808:031, 2008. doi: 10.1088/1475-7516/2008/08/031.

G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *Astrophys. J.*, 396: L1–L5, September 1992. doi: 10.1086/186504.

D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. *Astrophys. J. Suppl. Ser.*, 148:175–194, September 2003. doi: 10.1086/377226.

David N. Spergel, Neil Turok, William H. Press, and Barbara S. Ryden. Global texture as the origin of large scale structure: numerical simulations of evolution. *Phys. Rev.*, D43:1038–1046, 1991. doi: 10.1103/PhysRevD.43.1038.

R. A. Sunyaev and I. B. Zeldovich. Microwave background radiation as a probe of the contemporary structure and history of the universe. 18:537–560, 1980. doi: 10.1146/annurev.aa.18. 090180.002541.

R. A. Sunyaev and Y. B. Zeldovich. The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173, November 1972.

Leonard Susskind. The anthropic landscape of string theory. In *Universe or Multiverse*. Cambridge University Press, 2003.

N. Suzuki, D. Rubin, C. Lidman, G. Aldering, R. Amanullah, K. Barbary, L. F. Barrientos, J. Botyanszki, M. Brodwin, N. Connolly, K. S. Dawson, A. Dey, M. Doi, M. Donahue, S. Deustua, P. Eisenhardt, E. Ellingson, L. Faccioli, V. Fadeyev, H. K. Fakhouri, A. S.

Fruchter, D. G. Gilbank, M. D. Gladders, G. Goldhaber, A. H. Gonzalez, A. Goobar, A. Gude, T. Hattori, H. Hoekstra, E. Hsiao, X. Huang, Y. Ihara, M. J. Jee, D. Johnston, N. Kashikawa, B. Koester, K. Konishi, M. Kowalski, E. V. Linder, L. Lubin, J. Melbourne, J. Meyers, T. Morokuma, F. Munshi, C. Mullis, T. Oda, N. Panagia, S. Perlmutter, M. Postman, T. Pritchard, J. Rhodes, P. Ripoche, P. Rosati, D. J. Schlegel, A. Spadafora, S. A. Stanford, V. Stanishev, D. Stern, M. Strovink, N. Takanashi, K. Tokita, M. Wagner, L. Wang, N. Yasuda, H. K. C. Yee, and T. Supernova Cosmology Project. The Hubble Space Telescope Cluster Supernova Survey. V. Improving the Dark-energy Constraints above $z > 1$ and Building an Early-type-hosted Supernova Sample. *Astrophys. J.*, 746:85, February 2012. doi: 10.1088/0004-637X/746/1/85.

J. A. Tauber, N. Mandolesi, J.-L. Puget, T. Banos, M. Bersanelli, F. R. Bouchet, R. C. Butler, J. Charra, G. Crone, J. Dodsworth, and et al. Planck pre-launch status: The Planck mission. *Astron. Astrophys.*, 520:A1, September 2010. doi: 10.1051/0004-6361/200912983.

M. Tegmark, M. R. Blanton, M. A. Strauss, F. Hoyle, D. Schlegel, R. Scoccimarro, M. S. Vogeley, D. H. Weinberg, I. Zehavi, A. Berlind, T. Budavari, A. Connolly, D. J. Eisenstein, D. Finkbeiner, J. A. Frieman, J. E. Gunn, A. J. S. Hamilton, L. Hui, B. Jain, D. Johnston, S. Kent, H. Lin, R. Nakajima, R. C. Nichol, J. P. Ostriker, A. Pope, R. Scranton, U. Seljak, R. K. Sheth, A. Stebbins, A. S. Szalay, I. Szapudi, L. Verde, Y. Xu, J. Annis, N. A. Bahcall, J. Brinkmann, S. Burles, F. J. Castander, I. Csabai, J. Loveday, M. Doi, M. Fukugita, J. R. Gott, III, G. Hennessy, D. W. Hogg, Ž. Ivezić, G. R. Knapp, D. Q. Lamb, B. C. Lee, R. H. Lupton, T. A. McKay, P. Kunszt, J. A. Munn, L. O'Connell, J. Peoples, J. R. Pier, M. Richmond, C. Rockosi, D. P. Schneider, C. Stoughton, D. L. Tucker, D. E. Vanden Berk, B. Yanny, D. G. York, and SDSS Collaboration. The Three-Dimensional Power Spectrum of Galaxies from the Sloan Digital Sky Survey. *Astrophys. J.*, 606:702–740, May 2004. doi: 10.1086/382125.

Max Tegmark. How to make maps from CMB data without losing information. *Astrophys. J.*, 480:L87–L90, 1997a. doi: 10.1086/310631.

Max Tegmark. How to measure CMB power spectra without losing information. *Phys. Rev.*, D55:5895–5907, 1997b. doi: 10.1103/PhysRevD.55.5895.

Max Tegmark and Angelica de Oliveira-Costa. Removing point sources from CMB maps. 500: L83–L86, 1998.

Roberto Trotta. Applications of Bayesian model selection to cosmological parameters. *Mon. Not. R. Astron. Soc.*, 378:72–82, 2007. doi: 10.1111/j.1365-2966.2007.11738.x.

Michael S. Turner. A Tilted universe (and other remnants of the preinflationary universe). *Phys. Rev.*, D44:3737–3748, 1991. doi: 10.1103/PhysRevD.44.3737.

Neil Turok. Global Texture as the Origin of Cosmic Structure. *Phys. Rev. Lett.*, 63:2625, 1989. doi: 10.1103/PhysRevLett.63.2625.

Neil Turok and David Spergel. Global texture and the microwave background. *Phys. Rev. Lett.*, 64:2736, 1990. doi: 10.1103/PhysRevLett.64.2736.

Neil Turok and David N. Spergel. Scaling solution for cosmological sigma models at large N. *Phys. Rev. Lett.*, 66:3093–3096, 1991. doi: 10.1103/PhysRevLett.66.3093.

P. Vielva, E. Martinez-Gonzalez, M. Cruz, R. B. Barreiro, and M. Tucci. Cosmic microwave background polarization as a probe of the anomalous nature of the cold spot. *Mon. Not. R. Astron. Soc.*, 410:33–38, January 2011. doi: 10.1111/j.1365-2966.2010.17418.x.

A. Vilenkin and E. P. S. Shellard. *Cosmic strings and other topological defects*. Cambridge University Press, 1986.

B. D. Wandelt and K. M. Górski. Fast convolution on the sphere. *Phys. Rev. D*, 63(12):123002, June 2001. doi: 10.1103/PhysRevD.63.123002.

Steven Weinberg. The cosmological constant problem. *Rev. Mod. Phys.*, 61:1–23, Jan 1989. doi: 10.1103/RevModPhys.61.1. URL `http://link.aps.org/doi/10.1103/RevModPhys.61.1`.

Y. Wiaux, J. D. McEwen, P. Vandergheynst, and O. Blanc. Exact reconstruction with directional wavelets on the sphere. 388(2):770–788, 2008. doi: 10.1111/j.1365-2966.2008.13448.x.

Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT Press, 1964.

E. L. Wright, X. Chen, N. Odegard, C. L. Bennett, R. S. Hill, G. Hinshaw, N. Jarosik, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, J. L. Weiland, E. Wollack, J. Dunkley, B. Gold, M. Halpern, A. Kogut, D. Larson, M. Limon, S. S. Meyer, and G. S. Tucker. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Source Catalog. *Astrophys. J. Suppl. Ser.*, 180:283–295, February 2009. doi: 10.1088/0067-0049/180/2/283.

S. Zaroubi, Y. Hoffman, K. B. Fisher, and O. Lahav. Wiener Reconstruction of The Large Scale Structure. *Astrophys. J.*, 449:446–459, 1995. doi: 10.1086/176070.

Jun Zhang and Yun-Song Piao. Preheating in Bubble Collision. *Phys.Rev.*, D82:043507, 2010. doi: 10.1103/PhysRevD.82.043507.

J. P. Zibin and Douglas Scott. Gauging the cosmic microwave background. *Phys. Rev.*, D78: 123529, 2008. doi: 10.1103/PhysRevD.78.123529.

F. Zwicky. Spectral displacement of extra galactic nebulae. *Helv.Phys.Acta*, 6:110–127, 1933.