

Jean-François Cardoso

# Precision Cosmology with the Cosmic Microwave Background

[A focus on issues related to spherical analysis]

© XXXXXX

**T**he cosmic microwave background (CMB) is a cornerstone of modern cosmology; extracting cosmological information from it requires high accuracy instruments as well as a great deal of signal processing. This article highlights a critical step of the CMB data processing pipeline: the construction of a likelihood function for the statistical exploitation of CMB data. Since this task is based on a sky map of CMB anisotropies, it requires some understanding of spherical analysis. The aim of this article is to introduce the reader to some of the issues related to spherical analysis for the CMB. It opens with a short overview of the CMB and explain how to build a spherical likelihood for it, starting with an idealized case of perfect (full-sky, noise-free) CMB observations and then moving to the actual computations with imperfect data using specific tools for signal processing on the sphere.

## INTRODUCTION

The CMB has become a “primary tool for determining the global properties, content, and history of the Universe” [11], with the most recent high precision measurements beautifully matching a

(relatively) simple “concordance model” (see Figure 1). <AU: Figures 1–7 were renumbered throughout the text so that they appear in numerical order.> These great achievements are obtained by combining extensive physical models of the early Universe, cutting edge instrumentation, and a lot of signal processing.

A data processing pipeline for a CMB instrument includes steps that consist of sky scans, sky maps, spectrum estimation, and statistical inference about the cosmological parameters describing the early Universe. <AU: please check that edited sentence retains original meaning> These steps involve challenging signal processing tasks. This article focuses on one of the last steps: writing down a tractable yet accurate likelihood function for CMB data once they have been processed into a map of temperature anisotropies, such as the one depicted in Figure 2. This figure, as well as all full-sky maps of this article, represents the celestial sphere using the Mollweide equal-area projection.

Digital Object Identifier 10.1109/MSP.2009.934715

From the signal processing point of view, a specific requirement is to perform spectral analysis on the sphere. This (maybe) unusual setting is the secondary focus of this article.

### PRECISION COSMOLOGY AND THE CMB

This section briefly describes the nature of the CMB and how the measure of its anisotropies has emerged as a pillar of modern cosmology. Statistical exploitation of these increasingly accurate measurements is outlined, which allows cosmologists to go from a CMB sky map to a statistical statement about, for instance, the flatness of the Universe.

### THE CMB

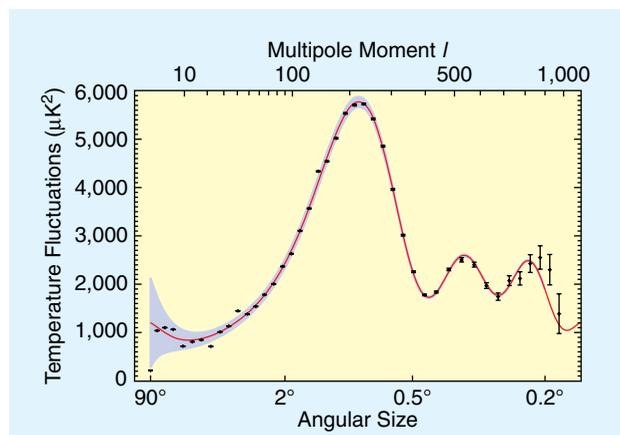
If you operate a radio-telescope in the centimeter domain, you will find that from any point of the sky comes a faint signal with a very uniform (direction independent) intensity. If you are thorough, you may convince yourself that this is not some instrumental noise but an actual sky signal, the CMB, but you will not earn a Nobel prize, because Penzias and Wilson already did in 1978 for their groundbreaking 1964 observations [30] at the Bell Labs.

The existence of the CMB signal was actually predicted by Gamow and his collaborators even before its unambiguous observation by Penzias and Wilson, based on the following episode of the Big Bang scenario [5]. At some time in the past, the Universe was hot enough to be a plasma of light nuclei (mostly protons), electrons and photons, closely interacting with each other and hence in thermal equilibrium. If such a plasma expands (due to the expansion of the Universe itself), it also cools down to the point where thermal agitation can no longer prevent nuclei and electrons from combining to form atoms. In that transition from an ionized state to a neutral state, the Universe turns from opaque to transparent because photons scatter much less off a neutral atom than off a free electron. Hence, matter and radiation decouple and the photons are set free to roam the Cosmos. Most of them have been doing so, traveling undisturbed since the time of last scattering. That was about 13 billion years ago, some 380,000 years after the Big Bang.

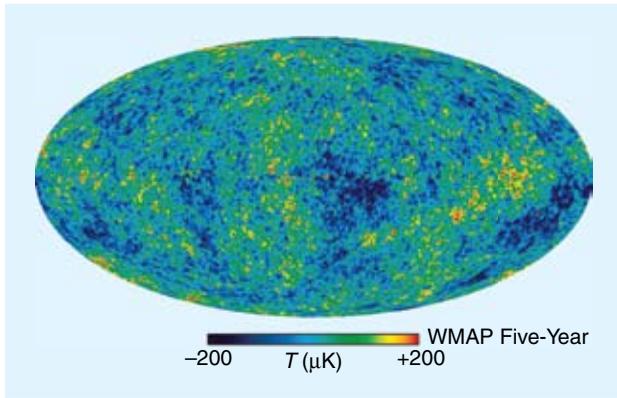
At decoupling time (also called recombination), photons were in thermal equilibrium with matter at temperature of a few thousand Kelvins with a mean wavelength of the order of the micron. But general relativity predicts that the expansion of the Universe also expands the wavelength of CMB photons by the same factor while preserving the shape of its spectral density. See Figure 3 and wonder at the exquisite fit between the spectral density of the CMB as measured by NASA's Cosmic Background Explorer (COBE)-far infrared absolute spectrophotometer in 1990 and its theoretical prediction by the Planck formula for the radiation/matter thermal equilibrium at temperature  $T = 2.725$  K. Very standard physics tells us that a plasma (as the one described above) undergoes a phase transition from ionized to neutral at a temperature around 3,000 K so that observing the CMB today at  $T \approx 3$  K teaches us very straightforwardly that the Universe has expanded by a factor of about 1,000 since recombination.

Penzias and Wilson had already measured the mean CMB temperature of about 3 K uniformly on the sky but one had to wait for NASA's COBE-Differential Microwave Radiometer (DMR) to discover long sought CMB anisotropies, that is, variations over the sky of CMB temperature. The DMR instrument <AU: ok to delete word "instrument?"> aboard COBE measured the CMB temperature anisotropies with an angular resolution of about  $7^\circ$  and a sensitivity of a few tens of  $\mu K$ . Figure 4(a) shows a featureless sky: the CMB temperature is apparently very uniform over the sky. Figure 4(b) displays the temperature with its mean value removed, showing essentially a dipolar pattern: it is the Doppler effect for the instrument flying through the CMB sea. The most interesting patterns emerge after this dipole is subtracted: a red stripe along the equator is due to the microwave emission from our Galaxy, the Milky Way, and, at the poles, tiny temperature fluctuations, of the order of a few tens of  $\mu K$ , which are mostly free from Galactic contamination. That is, blurred by the instrument and rescaled by cosmic expansion, the large-scale distribution of the temperature of the Universe when it was a few hundred thousand years old!

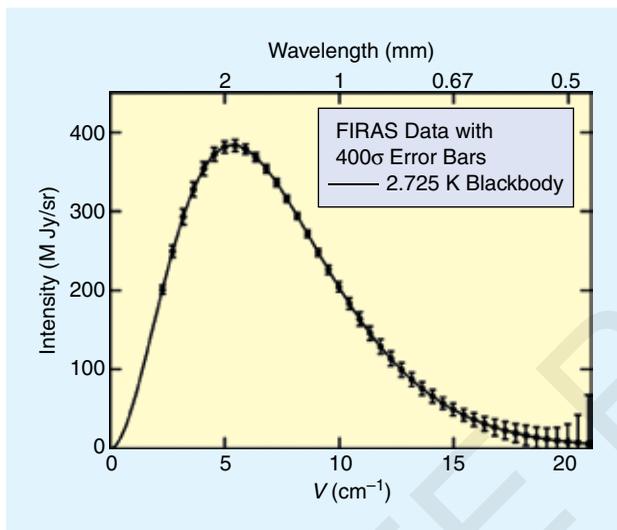
COBE images made a splash on the front page of the *New York Times* (and in 2006, G. Smoot and J. Mather earned the second CMB Nobel Prize) but even more exciting, physics had to come from ever more accurate CMB measurements. Today, the best full-sky observations come from NASA's Wilkinson microwave anisotropy probe (WMAP) mission with CMB observations at a resolution better than  $1^\circ$ , giving access to details about ten times finer than COBE's. Figure 2 shows a full-sky CMB map, with almost no visible Galactic contamination, thanks to a component separation procedure combining several frequency channels [13]. Measuring CMB anisotropies at the degree resolution (or better) is of utmost importance because one degree is small enough an angular separation to connect parts of the Universe close enough to be causally related. Hence, these fine features are signatures



**[FIG1]** Rescaled angular spectrum  $C_\ell \ell(\ell + 1)/2\pi$  in  $\mu K^2$  vs angular frequency (top horizontal axis, labelled 'multipole moment' here) or vs the angular size  $\theta = \pi/\ell$  (bottom horizontal axis). Black dots:  $\hat{C}_\ell$  measurements from WMAP (grey points). The grey area shows the cosmic variance. Red line is the best fit model  $C(\hat{\omega}_{ML})$ . Credit: NASA.



**[FIG2]** Map of the cosmic background in the Mollweide projection after foreground removal based on five years of observation by WMAP. Credit: NASA.



**[FIG3]** Solid line: Planck law of black body radiation (light in equilibrium with matter at temperature  $T$ ) has spectral density depending only on  $T$ . Data points with error bars: COBE measurements, fitted by Planck law for  $T = 2.725\text{ K}$  with wonderful accuracy: here, for readability, error bars are blown up by a factor of 400. Credit: NASA.

for the physics of the early Universe. Since they correspond to spatial fluctuations in the energy density, the CMB anisotropies can be understood as marking the seeds of large scale structures of the Universe.

The Planck space mission of the European Spatial Agency [10], launched in May 2009, will deliver even sharper images of the microwave sky in nine frequency channels from 30 to 857 GHz with a (frequency-dependent) resolution as good as 5 arcmin [9].

### PRECISION COSMOLOGY

How is science extracted from a map of CMB anisotropies? Quantitative statements are obtained statistically by considering the CMB map as a (single) realization of a random process. In this respect, the task of theoretical cosmologists is to under-

stand how the physics of the early Universe governs the probability distribution of CMB observations to build a likelihood function  $\Pr(\text{CMB}|\omega)$ , where  $\omega$  is a set of parameters for a given cosmological model. If the CMB is dubbed a “pillar of cosmology,” it is because its probability distribution is sensitive to key cosmological parameters. One may write informally

$$\omega = (\Omega_b, \Omega_c, \Omega_\Lambda, H, \dots),$$

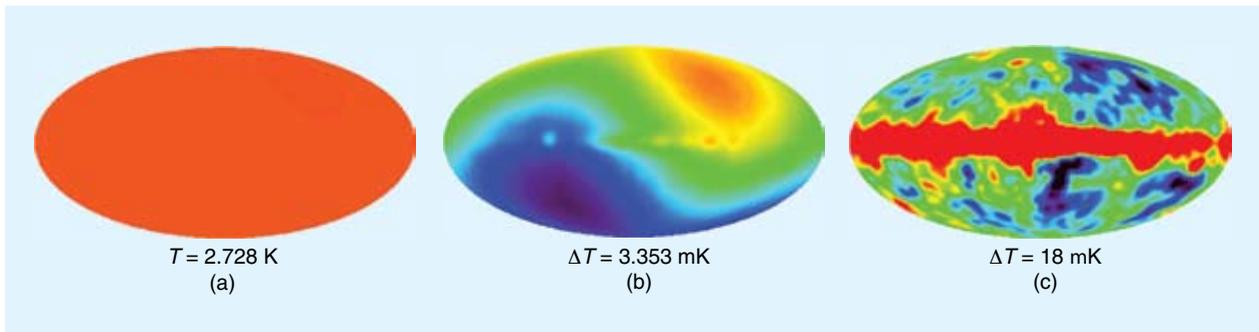
where  $\Omega_b$  ( $\Omega_c$ ,  $\Omega_\Lambda$ , respectively) <AU: edit ok here?> is the (properly normalized) energy density due to ordinary matter (mostly baryons) (to cold dark matter and dark energy, respectively), <AU: edit ok here?>  $H$  is the Hubble constant. As an example, Figure 5 shows contours of a CMB likelihood when two parameters are varied: the relative energy densities of matter  $\Omega_m = \Omega_b + \Omega_c$  and of dark energy  $\Omega_\Lambda$ . The contours are aligned along the  $\Omega_m + \Omega_\Lambda = 1$  line that correspond to a flat Universe. It is outside the scope of this article and beyond the skills of this author to provide a serious exposition of cosmological parameters. The interested reader may refer to some excellent cosmology primers from real cosmologists such as Wayne Hu [24], Simon White [37], Ned Wright [39], or Max Tegmark [34] to name a few, or from NASA [29].

Building a good probabilistic model  $\Pr(\text{CMB}|\omega)$  seems to be quite a feat, but it is not so difficult in the standard Big Bang model in which the CMB is predicted to be (a realization of) a Gaussian stationary field. In such a case, its probability distribution depends on the cosmological parameters only via its angular power spectrum  $\mathcal{C}$  (defined in the next section) so that the division of labor for the scientific exploitation of CMB data can be caricatured as combining the following:

- 1) a cosmological model (or competing models) predicting the dependence of the angular power spectrum on the cosmological parameters: software packages, such as the Code for Anisotropies in the Microwave Background (CAMB) [26] or CMBfast [31], are responsible for implementing  $\omega \rightarrow \mathcal{C}(\omega)$
- 2) an instrument and a lot of upstream signal processing to turn CMB data into the best possible, well-characterized CMB map
- 3) a likelihood function  $\mathcal{C} \rightarrow \Pr(\text{CMB map}|\mathcal{C})$  to express the probability of a given CMB map given an angular power spectrum  $\mathcal{C}$ .

There are many signal processing tasks required in the second point to go from the exquisitely delicate measures of a radiation temperature in many sky directions to producing a CMB map.

One critical task is going from the time-ordered data obtained by scanning the sky to spherical maps, one of for each detector or for each frequency channel. This so-called “map making process” must exploit redundancy (each pixel of the final map is seen several times during the instrumental survey), must take into account possibly asymmetric instrumental beams (point spread functions), and must also produce a good characterization of the residual noise. If the instrument response is well characterized, mapmaking can be addressed as



**[FIG4]** COBE-DMR temperature maps in Mollweide projection. (a) Plain temperature map: no visible deviation from uniform temperature of  $T_o \approx 2.725$  K. (b) After mean removal, a mostly dipolar pattern with amplitude in the  $mK$  range. (c) After dipole removal, visible anisotropies in the range of tens of  $\mu K$ . Credit: NASA.

a well-posed linear problem; it is mostly the sheer size of the problem (multimillion pixel maps) which makes it challenging. See [17] (and [15] dealing with asymmetric beams) and references therein for some approaches to mapmaking for the Planck mission.

Another critical task is component separation by which the diversity offered by multifrequency observations is used to sort out the cosmic background and the astrophysical foreground emissions which all contribute to the signal measured at the detectors. The foreground emissions, both from galactic (from our Milky Way) and extra-galactic origins are of astrophysical interest but, of course, the cosmological focus is on cleaning up the CMB. Figure 4(c) is a sky map without foreground cleaning (showing strong Galactic emission at the equator), while Figure 2 shows a CMB map after foreground removal. Component separation is a difficult task which requires combining a good statistical modeling of the data with as much prior information as possible. See [19] for the component separation performed by the WMAP team and [16] and references therein for a comparison of component separation methods for the Planck mission.

This article is not intended to address all aspects of CMB data processing. Rather, it aims at providing the reader with a glimpse of the third point mentioned earlier, that is, developing the likelihood of angular power spectra given a spherical map.

### SPHERICAL ANALYSIS

A full-sky CMB map may be seen as a real-valued function  $X(\xi)$  defined everywhere on the unit sphere  $S^2$ , where a point on the sphere (a direction in the sky) is represented by a  $3 \times 1$  unit vector  $\xi$ . This section introduces basic concepts of spherical analysis with emphasis on invariance: we look for invariant quantities, which are defined independently of a coordinate system.

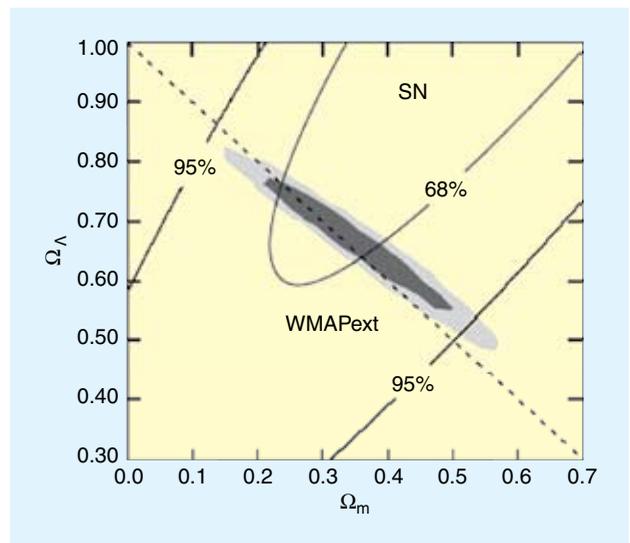
### MULTIPOLE DECOMPOSITION

A primary tool for exploring spherical fields is the multipole decomposition that reads

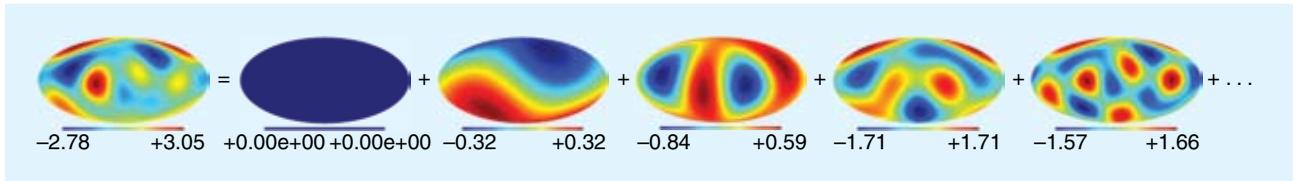
$$X(\xi) = \sum_{\ell \geq 0} X^{(\ell)}(\xi), \quad (1)$$

where index  $\ell$  is called the angular frequency and where the first components are called monopole, dipole, quadrupole, octopole, ... for  $\ell = 0, 1, 2, 3, \dots$  as illustrated by Figure 6. This is an orthogonal decomposition, that is,  $\int_{S^2} X^{(\ell)} X^{(\ell')} = 0$  for  $\ell \neq \ell'$ . Each multipolar component  $X^{(\ell)}$  in fact is the orthogonal projection of  $X$  onto a linear subspace  $\mathcal{H}_\ell$  of functions. These harmonic subspaces can be characterized as the smallest function spaces which are globally invariant under rotations. They have a very simple form at the first three frequencies:  $\mathcal{H}_0$  is the set of constant functions ( $\dim(\mathcal{H}_0) = 1$ );  $\mathcal{H}_1$  is the set of functions of the form  $X(\xi) = a^T \xi$  for all fixed  $3 \times 1$  vector  $a$  ( $\dim(\mathcal{H}_1) = 3$ );  $\mathcal{H}_2$  is the set of functions of the form  $X(\xi) = \xi^T A \xi$  for all fixed  $3 \times 3$  symmetric traceless matrices  $A$  ( $\dim(\mathcal{H}_2) = 5$ ). More generally, the harmonic subspace  $\mathcal{H}_\ell$  can be shown to have dimension  $2\ell + 1$ .

A beautiful general characterization of harmonic subspaces is via the spherical Laplacian  $\Delta_S$ . The spherical Laplacian



**[FIG5]** Likelihood of relative energy densities for matter  $\Omega_m$  and dark energy  $\Omega_\Lambda$  cosmological parameters, given WMAP first year data complemented with other fine-scale CMB measurements (hence the WMAPext label) as described in [12]. Also shown: confidence regions based on independent supernova (SN) observations). For updated results after five years of data collection by WMAP, see [14]. Credit: NASA.



**[FIG6]** Multipole decomposition of a spherical field. This is the pictorial version of (1), showing the first five multipoles.

$\Delta_S X(\xi)$  at any point  $\xi$  is the sum of any two second derivatives of  $X(\xi)$  taken in mutually orthogonal directions which are tangent to the sphere at that point. The spherical Laplacian, as a linear operator on spherical functions, can be shown to have eigenvalues equal to  $-\ell(\ell + 1)$  for all integers  $\ell \geq 0$  and the harmonic subspace  $\mathcal{H}_\ell$  is the associated eigen-space: the set of all spherical functions such that  $\Delta_S X(\xi) = -\ell(\ell + 1)X(\xi)$ . An explicit basis for the harmonic subspaces is described in the section “Signal Processing on the Sphere,” but is not needed for now, as long as we focus on invariant properties.

The multipole decomposition may also be understood as the result of a spherical convolution: the  $\ell$ th multipole term can be obtained as the convolution product

$$X^{(\ell)}(\xi) = \int_{\xi' \in S^2} X(\xi') \frac{2\ell + 1}{4\pi} P_\ell(\xi^\dagger \xi') d\sigma(\xi'), \quad (2)$$

where the scalar product  $\xi^\dagger \xi'$  is the cosine of the angle between directions  $\xi$  and  $\xi'$  and where  $P_\ell(\cdot)$  is the  $\ell$ th Legendre polynomial. Figure 7 shows the first six Legendre polynomials.

It may be enlightening to compare to the case of a circular function  $X(\theta)$ ,  $\theta \in [0, 2\pi)$  defined on  $S_1$ , the unit circle. The “circular Laplacian” reducing to the second derivative, its eigen-subspaces clearly are  $\mathcal{H}_n^c = \{\alpha \cos(n\theta) + \beta \sin(n\theta)\}$  for  $n \geq 0$ , of dimension  $w_n$  with  $w_0 = 1$  and  $w_n = 2$  for  $n > 0$ . If  $X \in \mathcal{H}_n^c$ , then  $X'' = -n^2 X$ , so the corresponding eigenvalue is  $-n^2$ , which is almost like  $-\ell(\ell + 1)$ . Hence, on  $S_1$ , the  $n$ th harmonic subspace  $\mathcal{H}_n^c$  is made of all periodic functions at frequency  $n$ , is globally invari-

ant under rotations, has dimension  $w_n$  and the  $n$ th harmonic component  $X^{(n)}(\theta)$  is extracted by a circular convolution

$$X^{(n)}(\theta) = \int_0^{2\pi} X(\theta') \frac{w_n}{2\pi} \cos(n(\theta - \theta')) d\theta' \quad (3)$$

in close analogy to the spherical convolution (2). We see here that the Legendre polynomials somehow are to the sphere what cosines are to the circle. See (5) and (6) for another instance of this parallelism.

### ISOTROPIC FIELDS, ANGULAR CORRELATION, AND SPECTRUM

A spherical random field is isotropic or stationary (these two terms are used indifferently here) if its probability distribution is invariant under all rotations. In particular, the correlation between two directions  $\xi$  and  $\xi'$  depends only on their angular separation

$$\text{Cov}(X(\xi), X(\xi')) = \rho(\xi^\dagger \xi'), \quad (4)$$

where, again, the scalar product  $\xi^\dagger \xi'$  is the cosine of the angle separating  $\xi$  and  $\xi'$ . Equation (4) defines, on the interval  $[-1, 1]$ , the angular correlation function  $\rho$  of the stationary field  $X$ .

Recall how, for a stationary process on the circle, the correlation function  $\rho^c(\cdot)$  defined by  $\text{Cov}(X(\theta), X(\theta')) = \rho^c(\theta - \theta')$  and the discrete power spectrum  $S_n$  are related via a cosine transform

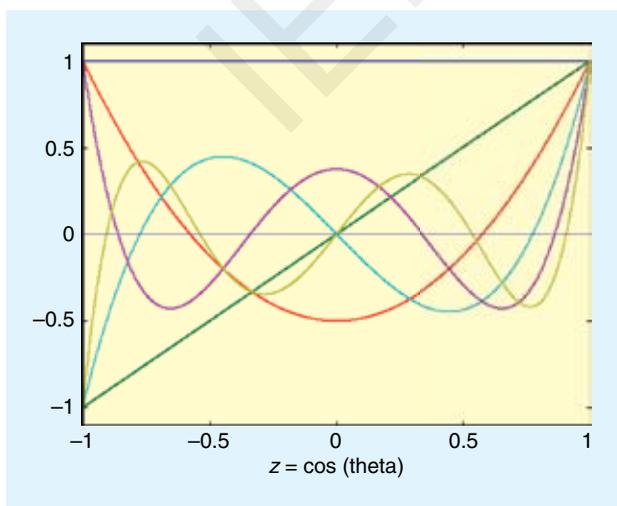
$$\rho^c(\tau) = \sum_{n \geq 0} S_n \frac{w_n}{2\pi} \cos(n\tau), \quad (5)$$

where, as in (3),  $w_0 = 1$  and  $w_n = 2$  for  $n > 0$ . Now, just as the spherical equivalent of (3) is (2), the spherical equivalent of (5) is to define the angular power spectrum  $\mathcal{C} = \{C_\ell\}_{\ell \geq 0}$  of an isotropic process on the sphere via a “Legendre transform” of the angular correlation, that is,

$$\rho(z) = \sum_{\ell \geq 0} C_\ell \frac{2\ell + 1}{4\pi} P_\ell(z). \quad (6)$$

How is that definition related to angular power? Since the multipole decomposition is orthogonal, the total energy  $\|X\|^2 = \int_{S^2} X^2$  of a given realization  $X$  decomposes as a sum over multipoles:  $\|X\|^2 = \sum_{\ell \geq 0} \|X^{(\ell)}\|^2$ . Define then the empirical angular power spectrum  $\hat{\mathcal{C}}$  of  $X$  as

$$\hat{\mathcal{C}} = \{\hat{C}_\ell\}_{\ell \geq 0} \text{ with } \hat{C}_\ell = \frac{1}{2\ell + 1} \|X^{(\ell)}\|^2,$$



**[FIG7]** Legendre polynomials  $P_\ell(z)$  for  $0 \leq \ell \leq 5$  on  $[-1, 1]$ .  $P_0(z) = 1$ ,  $P_1(z) = z$ ,  $P_2(z) = (3z^2 - 1)/2$ ,  $P_3(z) = (5z^3 - 3z)/2$ , ... The  $\ell$ th polynomial changes sign  $\ell$  times on  $[-1, 1]$ .

that is,  $\hat{C}_\ell$  is the average energy per dimension (since  $\dim(\mathcal{H}_\ell) = 2\ell + 1$ ) of the  $\ell$ th multipole of  $X$ . Now, if  $X$  is a realization of a stationary field with angular spectrum  $C_\ell$ , one can show that  $C_\ell = E(\hat{C}_\ell)$ , so that the angular spectrum indeed measures the mean (ensemble average) distribution of power across angular frequencies.

### SPHERICAL LIKELIHOOD

The empirical angular spectrum of a Gaussian isotropic field is an exhaustive statistic i.e., the probability  $p(X|\mathcal{C})$  of any realization  $X$  depends on  $X$  only through its empirical spectrum  $\hat{\mathcal{C}}$ . Indeed, the likelihood of a power spectrum  $\mathcal{C}$  takes the form

$$-2\log p(X|\mathcal{C}) = \sum_{\ell \geq 0} (2\ell + 1) \left( \frac{\hat{C}_\ell}{C_\ell} + \log C_\ell \right) + \text{cst.} \quad (7)$$

That expression is derived in the section “Statistics in Harmonic Space” using spherical harmonics but could also be established using only arguments of spherical invariance.

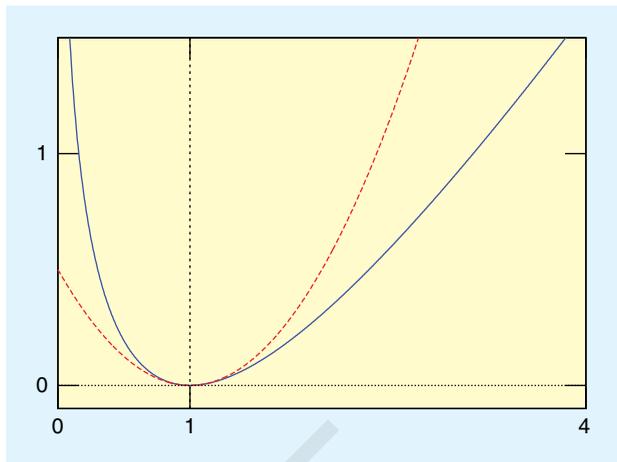
The cosmological principle postulates an isotropic and homogeneous Universe and therefore predicts that the CMB map is the realization of a stationary field on the sphere. Further, the standard Big Bang model also predicts normally distributed CMB anisotropies. The Big Bang model is also known as “ $\Lambda$ -CDM” model because it requires a cosmological constant  $\Lambda$  (or dark energy) and cold dark matter (CDM). Actually, that model predicts possible small deviations from Gaussianity, small enough to be ignored at this level of exposition. Hence, within the Gaussian isotropic framework, the likelihood of any model of the Universe with cosmological parameters  $\omega$  can be evaluated by confronting a theoretical spectrum  $\mathcal{C}(\omega)$  to the empirical spectrum  $\hat{\mathcal{C}}$  via (7). In other words, the empirical angular spectrum  $\hat{\mathcal{C}}$  of the CMB summarizes all cosmological information.

### WMAP MEASUREMENTS AND COSMIC VARIANCE

Figure 1 summarizes the spectral analysis of the CMB temperature field based on five years of WMAP observations [13]. It displays rescaled angular spectra, plotting  $C_\ell \ell(\ell + 1)/2\pi$  rather than plain  $C_\ell$  for the sake of readability. There is a very clean determination of two peaks around  $\ell \approx 220$  and  $\ell \approx 550$  as well as a pretty good capture of a third peak. The empirical spectrum is beautifully fitted (except maybe a puzzlingly low quadrupole  $\hat{C}_4$ ) by the likeliest “concordance model,” involving only a handful of cosmological parameters.

Here, we will be concerned with the error analysis. Figure 1 shows error bars (in black ink) that get larger with increasing  $\ell$ , reflecting the decreasing signal-to-noise ratio at the finest scales. At low  $\ell$ , the signal-to-noise ratio is much larger than one and error bars are very small. However, the figure also shows a grey area of uncertainty in this part of the spectrum. This is the so-called “cosmic variance,” which we now discuss.

To do so, it is enlightening to rewrite the likelihood (7) by noting that  $\hat{C}/C + \log C = \kappa(\hat{C}/C) + \text{cst}$  where cst does not



[FIG8] Function  $k(u) = u - 1 - \log u$  (solid blue line) and its quadratic approximation  $(u - 1)^2/2$  (dashed red line).

depend on  $C$  and where  $\kappa(u) = u - 1 - \log u$ . The  $\kappa$  function is depicted in Figure 8; it is nonnegative and equivalent to  $(u - 1)^2/2$  up to third order at  $u = 1$ . Hence, the likelihood function (7) can be rewritten as

$$-2\log p(X|\mathcal{C}) = \sum_{\ell} (2\ell + 1) \kappa\left(\frac{\hat{C}_\ell}{C_\ell}\right) + \text{cst} \quad (8)$$

and interpreted as a measure of the mismatch between the empirical spectrum  $\hat{\mathcal{C}}$  and the candidate spectrum  $\mathcal{C}$ . Equation (8) shows that the log-likelihood integrates the error (or mismatch) term  $\kappa(\hat{C}_\ell/C_\ell)$  over the whole multipole range with a weight equal to  $2\ell + 1$ . This reflects the increased accuracy of  $\hat{C}_\ell$  as an estimate of  $C_\ell$  at higher multipoles: it is easily found [from (14)] that

$$\text{Cov}(\hat{C}_\ell/C_\ell) = 2/(2\ell + 1) \quad (9)$$

as a simple consequence of the fact  $\hat{C}_\ell$  estimates  $C_\ell$  as an average over  $2\ell + 1$  independent terms. Note that (8) would read as an optimally weighted (inverse variance) mean-square error criterion if  $\kappa(u)$  was replaced by its quadratic approximation  $(u - 1)^2/2$ .

Cosmologists call cosmic variance the irreducible uncertainty expressed by (9). By nature, it is mostly significant for low values of  $\ell$  as seen in Figure 1. It is an ultimate limitation of spectral analysis based on a single realization of the random field. Indeed, there is no possibility to improve on it neither by using a better instrument (we are already assuming perfect, full-sky, noise-free observations here) nor by averaging over many realizations because we have only one realization available: our unique Universe.

### SIGNAL PROCESSING ON THE SPHERE

We saw that inferring cosmological parameters only requires knowing the empirical angular power spectrum  $\hat{\mathcal{C}}$  of the CMB, which is simply related to an ideal (full-sky, noise-free) CMB map. The section “Dealing with Imperfect Observations” briefly addresses the case of imperfect observations, but even with per-

fect observations, the procedure leading to the CMB likelihood (7) remains an abstraction until we have the practical tools to estimate  $\hat{C}_\ell$ .

So far, we have defined only rotationally invariant, coordinate-free quantities. That may be conceptually pleasing, but concrete implementations are also needed! This section introduces some of the signal processing tools for practical data processing on the sphere. We discuss a sampling scheme, the construction of a Fourier basis on the sphere, and semifast algorithms for computing the coefficients of a spherical map in this basis.

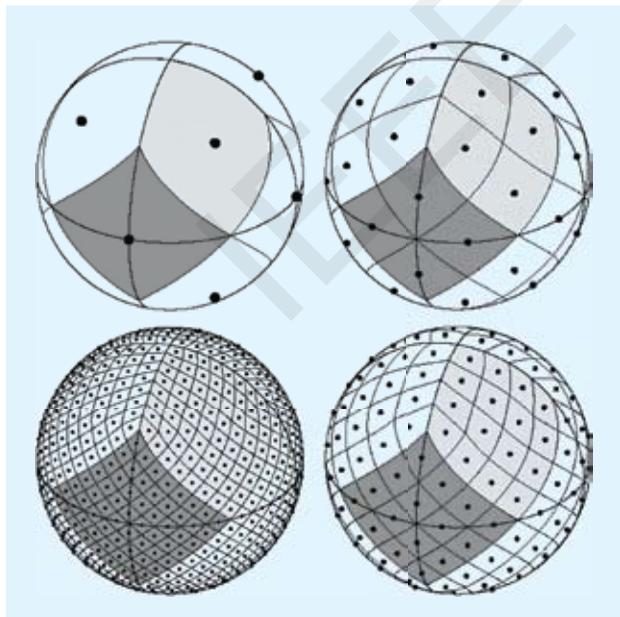
We use a coordinate system  $\xi = \xi(\theta, \phi)$  where a point  $\xi$  on the sphere is labeled by a pair  $(\theta, \phi)$  of angles: a longitude  $\phi \in [0, 2\pi]$  and a “colatitude” (angular distance from the North Pole)  $\theta \in [0, \pi]$ .

### SAMPLING

Sampling on the sphere is not a trivial problem. It is exacerbated in the case of CMB studies by the fact that the pixelization scheme must support the high resolution of today’s experiments that produce multimillion-pixel maps. The Hierarchical Equal Area isoLatitude Pixelization (HEALPix) scheme [20] has become a standard in CMB studies.

It is based on a first division of the sphere in 12 large pixels that can be further subdivided dyadically at the desired resolution, yielding maps of  $N_{\text{pix}} = 12 \times 2^{2r}$  pixels at resolution  $r$  (see Figure 9). WMAP produces maps at HEALPix resolution  $r = 10$  while Planck resolution will require  $r = 11$ .

Indeed, the pixels are carefully constructed to have the same area and to be evenly aligned on isolatitude rings. This last property is critical to implement fast Fourier transforms (FFTs)



**[FIG9] HEALPix sampling.** In (a), the 12 basic pixels are shown. Further subdivisions into  $N_{\text{pix}} = 12 \times 2^{2r}$  pixels at resolution levels  $r = 1, 2, 3$  are shown in (b)–(d), respectively. Figure courtesy of [20]. <AU: is edited caption ok? (a)–(d) will be added to figure later. Also, is credit line correct? Do you have permission from these authors to use this figure?>

(see below). <AU: do you mean the “A Fourier Basis: Spherical Harmonics” section?> Another nice property of HEALPix is its dyadic structure: the pixels can be arranged in memory in nested order, allowing for efficient local processing.

However, the HEALPix scheme is not perfect in every respect because it lacks an exact quadrature. Again, we can compare to the circle case: if  $X(\theta)$  is a properly band-limited function on the unit circle, its integral can be computed exactly from  $N$  equispaced points as

$$\int_0^{2\pi} X(\theta) d\theta = \sum_{i=1}^N X(\theta_i) \omega_i$$

using quadrature weights  $\omega_i = 2\pi/N$ . A similar form

$$\int_{S^2} X(\theta, \phi) d\sigma(\theta, \phi) = \sum_{i=1}^{N_{\text{pix}}} X(\theta_i, \phi_i) \Omega_i \quad (10)$$

exists for the sphere for some band limit on  $X$  (that is, if  $X^{(\ell)} = 0$  for  $\ell \geq \ell_{\text{max}}$ ) but the quadrature weights cannot be determined exactly in practice for large  $N_{\text{pix}}$  in the HEALPix scheme. A good approximation of course is  $\Omega_i = 4\pi/N_{\text{pix}}$ , since this is the common pixel area. The HEALPix package provides latitude-dependent weights that improve the quadrature significantly.

Many other pixelization schemes are possible but apparently no perfect pixelization has been found so far. Rectangular (equiangular) grids offer clean sampling theorems and exact quadrature (see [7], for instance) but have pixels of wildly varying sizes. Icosahedral grids offer pixels that are quite uniform in shape [35] but are not arranged on isolatitude rings, preventing fast spherical harmonic transform (see the next section). “Igloo” grids, e.g., [3] and [6], may allow for fast exact quadrature or have approximately equal-area pixels, but they are trickier to arrange hierarchically. Sampling and quadrature on the sphere is still an active area of research.

### A FOURIER BASIS: SPHERICAL HARMONICS

For harmonic computations, an explicit basis for each harmonic subspace  $\mathcal{H}_\ell$  is needed. Any orthonormal basis would do in theory, but a specific basis with nice computational properties can be obtained as follows. Recall that  $\mathcal{H}_\ell$  is characterized as the  $(2\ell + 1)$ -dimensional eigenspace of the spherical Laplacian  $\Delta_S$ , associated with eigenvalue  $-\ell(\ell + 1)$ . Define  $L = -i(\partial/\partial\phi)$  where  $i$  is the imaginary unit, an operator related to rotations around the pole axis. Its eigenfunctions clearly are of the form  $X(\theta, \phi) = g(\theta)e^{im\phi}$  for any function  $g(\theta)$ , associated with integer eigenvalues  $m$ . Now,  $\Delta_S$  and  $L$  are commuting Hermitian operators so that they share a common orthogonal basis of eigenfunctions. Hence, it must exist, for each  $\mathcal{H}_\ell$ , an orthonormal basis of functions for which the dependence on longitude factors as  $e^{im\phi}$ . These are the spherical harmonics, indexed by a pair of integers:  $\{Y_{\ell m}(\theta, \phi) | 0 \leq \ell, -\ell \leq m \leq \ell\}$  that form an orthonormal system on  $S^2$  and factor as

$$Y_{\ell m}(\theta, \phi) = P_{\ell m}(\cos\theta) e^{im\phi}, \quad (11)$$

where the  $P_{\ell m}$ s are proportional to the Legendre functions [1]. Space limitations prevent going into more detail, but the key point to keep in mind is the separation of the variables  $\theta$  and  $\phi$  in (11).

For our purpose, we would rather use real-valued spherical harmonics, denoted  $\mathcal{Y}_{\ell m}(\theta, \phi)$ , with  $\mathcal{Y}_{\ell 0} = P_{\ell 0}$  and

$$\begin{bmatrix} \mathcal{Y}_{\ell, +m} \\ \mathcal{Y}_{\ell, -m} \end{bmatrix} = G \begin{bmatrix} \mathcal{Y}_{\ell, +m} \\ \mathcal{Y}_{\ell, -m} \end{bmatrix} \quad \text{for } G = \sqrt{2} \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix}$$

for  $0 < m \leq \ell$ . Unitary matrix  $G$  preserves orthonormality while turning the complex pair  $[e^{im\phi} e^{-im\phi}]^T$  into the real pair  $[\sqrt{2}\cos(m\phi) \sqrt{2}\sin(m\phi)]^T$ . Figure 10 shows, for  $\ell = 3$ , the  $2\ell + 1 = 7$  real spherical harmonics making an orthobasis for  $\mathcal{H}_3$ .

Thanks to the orthonormality of spherical harmonics, we can now write a Fourier transform on the sphere

$$a_{\ell m} = \int_{S^2} \mathcal{Y}_{\ell m}(\theta, \phi) X(\theta, \phi) d\sigma(\theta, \phi), \quad (12)$$

$$X(\theta, \phi) = \sum_{0 \leq \ell} \sum_{-\ell \leq m \leq \ell} a_{\ell m} \mathcal{Y}_{\ell m}(\theta, \phi) \quad (13)$$

more often referred to as a spherical harmonic transform. The  $\ell$ th multipole component  $X^{(\ell)}(\theta, \phi)$  can be obtained by summing only over  $m$  with fixed  $\ell$  in the synthesis formula (13). However, this is not even needed to obtain the empirical angular spectrum: by orthonormality,  $\|X^{(\ell)}\|^2 = \sum_{-\ell \leq m \leq \ell} a_{\ell m}^2$ , so that

$$\hat{C}_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^2. \quad (14)$$

That expression, which can be seen as a spherical periodogram, is computed at virtually no cost once the  $a_{\ell m}$  coefficients are available.

### DISCRETE (AND FAST!) SPHERICAL TRANSFORMS

Imaging instruments have limited resolution, producing only smooth, band-limited versions of the underlying field. If a smooth map has a well-defined band limit  $L$ , that is,  $a_{\ell m} = 0$  for  $\ell \geq L$ , how should it be sampled to allow for the computation of its spherical harmonic coefficients  $a_{\ell m}$ ? Since the total number  $N_{\text{modes}}$  of nonzero coefficients is  $N_{\text{modes}} = \sum_{0 \leq \ell < L} (2\ell + 1) = L^2$ , such a band limited map has  $L^2$  degrees of freedom and thus must

be supported on a grid of at least  $L^2$  pixels. But is that number sufficient? This is a thorny issue because the practitioner needs a sampling grid, which is not only efficient (not many more points than the bare minimum  $L^2$ ), but also offers accurate quadrature for the computation of the  $a_{\ell m}$  coefficients from a discrete version of (12). Further, that computation must be fast because we deal with multimillion-pixel maps.

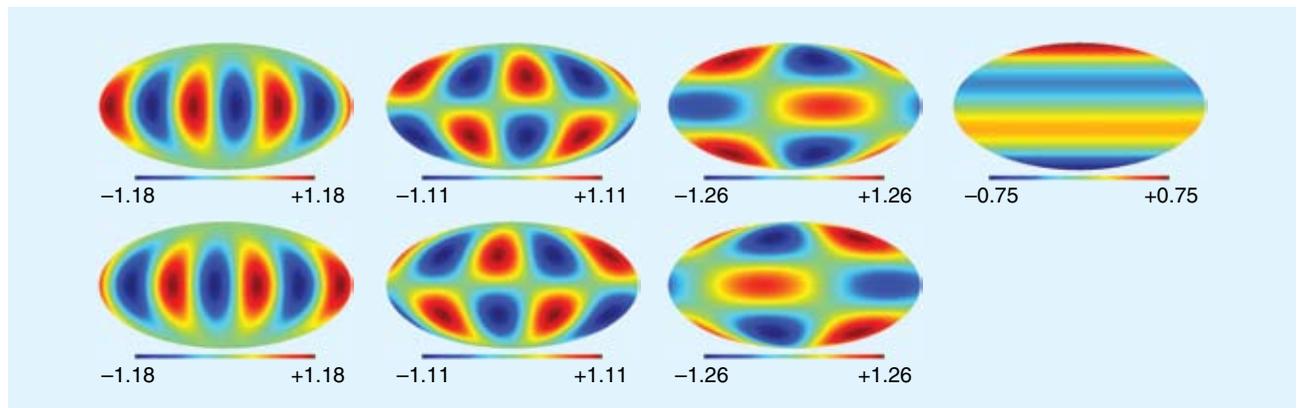
As of today, efficiency, exactness, and speed are slightly conflicting requirements.

For instance, an often-quoted sampling theorem of Driscoll and Healy guarantees exact harmonic coefficients  $a_{\ell m}$  from samples of a band-limited map on a  $(\theta, \phi)$ -rectangular grid of size  $2L \times 2L = 4L^2$  (even though that figure can be brought down to  $2L^2$ ) [7]. The HEALPix grid, and other well distributed grids, are probably more efficient but explicit weights for exact quadrature are missing. However, recent results show that weights could be obtained from a numerical procedure and indicate that  $4/3L^2$  samples are the approximate practical limit needed to access harmonic coefficients up to frequency  $L$  [21].

Regarding the speed of spherical harmonic transforms, a naive evaluation of (12), (13) scales as  $L^2 \times L^2 = L^4$ . This is prohibitive at large  $L$  (without even accounting for computing  $\mathcal{Y}_{\ell m}$  at the sampling points). Fortunately, any sampling scheme with pixels located on isolatitude rings can exploit the separation of variables  $\theta$  and  $\phi$  in the spherical harmonics (11) and organize the computations to scale as  $L^3$ . Further, if the pixels are evenly located one each ring, one-dimensional FFTs can be used to perform the summation over  $\phi$  or  $m$ . Even faster transforms can be designed by refining the integration in the  $\theta$  direction in the spirit of the discrete cosine transform. <AU: please check that edited sentence retains original meaning> One can achieve scaling of  $L^2 \log_2^2 L$  on rectangular grids [7], [38], but it is still unclear whether that is competitive with simpler approaches when all computing costs are accounted for.

### STATISTICS IN HARMONIC SPACE

We can now go back to the issue of producing a likelihood for a realization  $X$  of a Gaussian stationary field on the sphere. Gaussianity of the field entails the Gaussianity of its coefficients



**[FIG10]** Real spherical harmonics  $\mathcal{Y}_{\ell m}(\theta, \phi)$  for  $\ell = 3$  and, column-wise,  $\pm m = 3, 2, 1, 0$ .

and its stationarity entails the complete decorrelation of the harmonic coefficients

$$\text{Cov}(a_{\ell m}, a_{\ell' m'}) = C_\ell \delta(\ell, \ell') \delta(m, m') \quad (15)$$

and therefore their independence since they are Gaussian (see e.g., [2]). Thus, the probability density of a coefficient is  $p(a_{\ell m}) = 1/(\sigma\sqrt{2\pi}) \exp - (a_{\ell m}^2/2\sigma^2)$  with  $\sigma^2 = C_\ell$  or, equivalently,  $-2\log p(a_{\ell m}) = a_{\ell m}^2/C_\ell + \log C_\ell + \text{cst}$  and the log-density for the whole set, denoted  $\mathbf{a}$ , is just the sum of the log-densities. Therefore,

$$\begin{aligned} -2\log p(\mathbf{a}|\mathcal{C}) &= \sum_\ell \sum_m a_{\ell m}^2/C_\ell + \log C_\ell + \text{cst} \\ &= \sum_\ell (2\ell + 1) [\hat{C}_\ell/C_\ell + \log C_\ell] + \text{cst}, \end{aligned} \quad (16)$$

where last equality uses (14). With harmonic coefficients  $\mathbf{a}$  and CMB map  $X$  being related by the fixed-linear transforms (12), (13), their log-densities can differ only by a constant, establishing the likelihood form of (7).

### DEALING WITH IMPERFECT OBSERVATIONS

The availability of spherical harmonic coefficients makes it easy to compute the likelihood of any angular spectrum: as we just saw, the cosmological information contained in a sky map is losslessly compressed [33] in the empirical angular spectrum  $\hat{\mathcal{C}}$  in the sense that this is the only statistic needed to evaluate the likelihood function (7). Unfortunately, the form of the CMB likelihood obtained at (7) is only an idealization based on the assumption of spherical harmonic coefficients computed from an perfect CMB temperature map: full sky, full resolution, noise-free, and foreground-clean. With real CMB data, that is only approximately true and, since the cost of data collection and the greatness of the scientific objectives call for data processing of the highest quality, we must now turn to examining how to deal with imperfect observations.

### REAL CMB MAPS

Any real instrument has finite resolution and sensitivity and can only produce blurred and noisy CMB maps. The simplest possible description of these effects is to assume that one observes the true CMB map convolved by an axisymmetric beam and contaminated by an additive Gaussian stationary noise. The angular spectrum of the observed map then is  $W_\ell C_\ell + N_\ell$  where  $W_\ell$  is the (squared) beam transfer function and  $N_\ell$  is the angular spectrum of the noise. However, this is still a Gaussian stationary model so that the CMB likelihood (8) becomes

$$-2\log p(X|\mathcal{C}) = \sum_{\ell=0}^{\ell_{\max}} (2\ell + 1) \kappa \left( \frac{\hat{C}_\ell}{W_\ell C_\ell + N_\ell} \right) + \text{cst}.$$

The only difference, with respect to (8), is the theoretical angular spectrum  $W_\ell C_\ell + N_\ell$  now including beam and noise effects and an explicit truncation of the sum to some upper multipole above <AU: can you clarify "above"> that one expects  $W_\ell C_\ell \ll N_\ell$ , i.e.,

the signal is hopelessly buried in noise. For WMAP, the upper limit is around  $\ell_{\max} \approx 1,000$  and Planck is expected to go up to  $\ell_{\max} \approx 3,000$ .

However, changing  $C_\ell$  to  $W_\ell C_\ell + N_\ell$  is not accurate enough to represent the imperfections of the CMB map obtained with the best current instruments. The major complications all stem from some lack of isotropy: a) the beam is not exactly axisymmetric and furthermore, it varies over the sky so that  $W_\ell$  is not properly defined; b) the noise is nonstationary; and c) only part of the sky can be considered as free of foreground contamination.

Leaving point a) aside, regarding the beam, point b) is illustrated by the sky coverage achieved by WMAP during its first year of observation [11]. Figure 11 shows the number of times a given pixel is seen during the survey, which may vary from 500 to 7,000 as a result of the sky-scanning strategy of the instrument (itself governed by operational constraints of the satellite). The CMB temperature is estimated in a given pixel with a variance that is essentially inversely proportional to the number of times the pixel is visited by the beam. Hence, the noise variance is pixel-dependent and Figure 11 can be seen as a map of the noise inverse variance. Point c) is illustrated in Figure 12, showing the extent of various masks that may be applied to a CMB map before attempting spectral estimation. The reason for masking part of the sky is that the beautiful CMB map of Figure 2 appears more homogeneous than it probably is. Such a map, obtained after the component separation step (see the section "Precision Cosmology") cannot be perfectly free of foreground contamination. Hence, as a conservative measure, the angular spectrum is estimated only from the presumably cleanest fraction of the sky [13]. The excluded regions, as illustrated by Figure 12, are mostly determined from patterns of galactic emission (such a coarse pattern is visible in Figure 2) but other very bright, point-like sources are also masked away.

### A TWO-SCALE STRATEGY

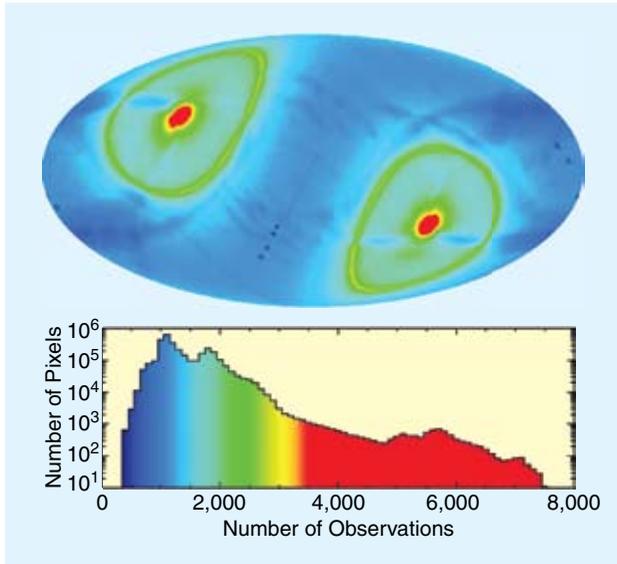
How does one perform spectral analysis with missing (masked away) data and nonuniform noise that both break isotropy and prevent Fourier magic from operating? One may write the likelihood in pixel space, as opposed to the harmonic-space version of (8). Indeed, if all available pixels are collected in a vector  $\mathbf{x}$ , the model  $X(\theta, \phi) = \sum_{\ell m} \mathcal{Y}_{\ell m}(\theta, \phi) a_{\ell m} + N(\theta, \phi)$  reads

$$\mathbf{x} = \mathbf{Y}\mathbf{a} + \mathbf{n},$$

where matrix  $\mathbf{Y}$  collects the values of the spherical harmonics at the sampling points  $(\theta_i, \phi_i)$  of the available pixels and vector  $\mathbf{a}$  collects the spherical harmonic coefficients  $a_{\ell m}$  up to some frequency  $\ell_{\max}$ . Then, the covariance matrix  $\mathbf{R} = \text{Cov}(\mathbf{x})$  of vector  $\mathbf{x}$  is

$$\mathbf{R} = \mathbf{R}(\mathcal{C}) = \mathbf{Y}\mathbf{C}\mathbf{Y}^t + \mathbf{N}, \quad (17)$$

where  $\mathbf{N} = \text{Cov}(\mathbf{n})$  is the noise covariance matrix and  $\mathbf{C} = \text{Cov}(\mathbf{a})$  is, by (15), a diagonal matrix with each value  $C_\ell$  of the angular spectrum repeated  $2\ell + 1$  times along the diagonal. With (17), specifying how the covariance matrix of the observations depends on the angular spectrum  $\mathcal{C}$ , the familiar



**[FIG11] Nonuniform sky coverage with the WMAP probe. (a) The hit map, that is, the number of independent observations per pixel and (b) the histogram of hit counts. Credit: NASA.**

likelihood for zero-mean Gaussian observations follows immediately as

$$-2\log p(x|\mathcal{C}) = \mathbf{x}^\dagger \mathbf{R}(\mathcal{C})^{-1} \mathbf{x} + \log \det \mathbf{R}(\mathcal{C}) + \text{cst.} \quad (18)$$

Unfortunately, this so-called “pixel-based likelihood” is computationally intractable for high-resolution maps because of the sheer size of the matrices involved.

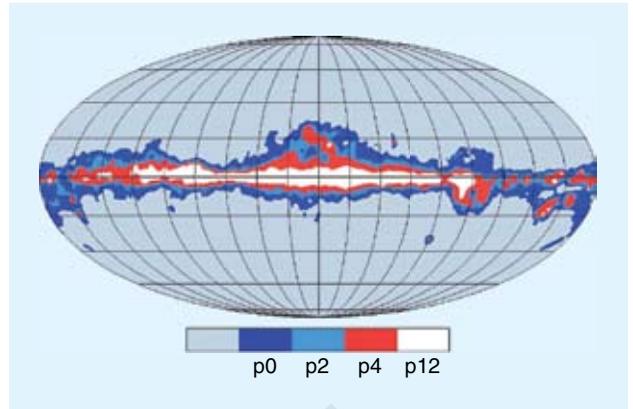
The current thinking to obtain a tractable likelihood from large incomplete CMB maps is to consider large scales and small scales independently (e.g., [8] and [22]). The observed map  $X$  is decomposed as  $X = X^L + X^H$  where  $X^L$  is a low-frequency (coarse scale) part and  $X^H$  is a high-frequency (fine scale) part. The total likelihood  $p(X|\mathcal{C})$  is the approximated as

$$\log p(X|\mathcal{C}) \approx \log p(X^L|\mathcal{C}^L) + \log p(X^H|\mathcal{C}^H),$$

where  $\mathcal{C}^L$  (respectively,  $\mathcal{C}^H$ ) is the low- $\ell$  (respectively high- $\ell$ ) part of the spectrum. Two approximations, one for each part of the likelihood, can be implemented, as now sketched.

1) *Coarse-scale likelihood:* The coarse-scale likelihood  $p(X^L|\mathcal{C}^L)$  can be computed in pixel space according to expression (18) because the low-resolution map  $X^L$  can be down-sampled to a reasonable size. For instance, if  $X^L$  contains multipoles up to  $\ell = 40$ , it can be sampled on an HEALPix grid at resolution  $r = 4$  supporting  $N_{\text{pix}} = 12 \cdot 16^2 = 3072$  pixels. The pixel-based form (18) of the likelihood does not constrain the noise covariance matrix  $N^L$  to have any specific form: the noise challenge then is not to include  $N^L$  in the likelihood but actually to obtain  $N^L$  itself by propagating the noise structure at the instrument level all the way to the low-resolution CMB map  $X^L$ .

2) *Fine-scale likelihood:* We describe a fine-scale approximation to the cut-sky likelihood [23] based on the so-called “pseudospectrum,” which is the empirical angular spectrum of a



**[FIG12] More or less conservative masks used by the WMAP team to guard against Galactic contamination. Credit: NASA.**

masked (or windowed) map. Let  $W(\theta, \phi)$  be a mask or some window function and consider  $\tilde{a}_{\ell m} = \int_{S^2} X(\theta, \phi) W(\theta, \phi) \mathcal{Y}_{\ell m}(\theta, \phi)$ , the spherical harmonic coefficients of the windowed map. One forms a “pseudospectrum”  $\tilde{\mathcal{C}}$

$$\tilde{\mathcal{C}}_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} \tilde{a}_{\ell m}^2 \quad (19)$$

and a vector  $\tilde{\mathcal{C}}$  made of such pseudospectra for high enough  $\ell$ . Since  $\tilde{\mathcal{C}}_\ell$  depends quadratically on the underlying field  $X$  (assumed Gaussian stationary for now), its mean value  $E(\tilde{\mathcal{C}}_\ell)$  is a linear function of the spectrum  $\mathcal{C}$  of  $X$ . Therefore, it must exist a matrix  $\mathcal{M}_W$ , depending only on the window function  $W$ , such that  $E(\tilde{\mathcal{C}}) = \mathcal{M}_W \mathcal{C}$ . Similarly, the covariance matrix  $\Sigma_W(\mathcal{C})$  of  $\tilde{\mathcal{C}}$  depends quadratically on  $\mathcal{C}$  (but this is not reflected by the notation).

Matrix  $\mathcal{M}_W$  is the identity matrix  $I$  on a full sky ( $W = 1$ ) and is roughly approximated by  $f_{\text{sky}} \times I$  on a cut sky if only a fraction  $f_{\text{sky}}$  is observed. In actuality, matrix  $\mathcal{M}_W$  is not diagonal and its off-diagonal coefficient  $[\mathcal{M}_W]_{\ell\ell'}$  measures (on average) the “leakage” of power from multipole  $\ell'$  into multipole  $\ell$ , induced by windowing. Note that since  $C_\ell$  scales typically as  $1/\ell^2$  for the CMB, the most significant spectral leakage by far is from low to high multipoles.

For high enough  $\ell$ , the distribution of  $\tilde{\mathcal{C}}_\ell$  is approximately Gaussian because  $\tilde{\mathcal{C}}_\ell$  is the sum of a sufficiently large number of independent terms. The distribution of vector  $\tilde{\mathcal{C}}$  is thus approximated as Gaussian with mean  $\mathcal{M}_W \mathcal{C}$  and covariance matrix  $\Sigma_W(\mathcal{C})$ . At this level of accuracy, one may even replace  $\Sigma_W(\mathcal{C})$  with a fixed covariance matrix  $\Sigma_\star$  based on a fiducial spectrum:  $\Sigma_\star = \Sigma_W(\mathcal{C}_{\text{fid}})$ . The distribution of the pseudospectrum is thus approximated as  $\mathcal{N}(\mathcal{M}_W \mathcal{C}, \Sigma_\star)$  resulting in an approximate fine-scale likelihood

$$-2\log p(X|\mathcal{C}) = (\tilde{\mathcal{C}} - \mathcal{M}_W \mathcal{C})^\dagger \Sigma_\star^{-1} (\tilde{\mathcal{C}} - \mathcal{M}_W \mathcal{C}) + \text{cst.}$$

This is tractable because both  $\mathcal{M}_W$  and  $\Sigma_\star$  can be computed for any given window function  $W$  and fiducial spectrum  $\mathcal{C}_{\text{fid}}$  [23]. For the sake of exposition, the additive noise has not been included but, needless to say, it should be done explicitly.

Interestingly, that approach also allows for a more sophisticated handling of the noise by using a window function tailored to the noise distribution. Indeed, if a map of pixel-dependent noise variance is available (see Figure 11 and comments), one may take  $W(\theta, \phi)$  to be proportional to the inverse noise variance in direction  $(\theta, \phi)$ , thus down-weighting the noisiest pixels. The pseudospectrum is perturbed by such a weird window function, but that effect is corrected by the transfer matrix  $\mathcal{M}_W$ . Even though this procedure seems ad hoc, some asymptotic results point towards its quasi optimality [8].

## WAVELETS

Wavelets seem to be a natural tool for dealing with nonuniform noise and sky coverage. Indeed, several types of spherical wavelets have been developed in recent years with a focus on CMB processing. We can only briefly touch upon this topic and provide a few references; see [38] and references therein for continuous wavelet transforms on the sphere and [32] for discrete systems based on the HEALPix grid. For papers specifically focused on spectral estimation using “needlets,” see, for instance, [27] and [18]. It is worth noting that spherical wavelets are being used on CMB data for other tasks than spectral estimation, e.g., [4], for foreground subtraction using needlets. For detecting non-Gaussianity in CMB maps, an early result [36] uses simple Mexican Hat wavelets with many follow-ups using more sophisticated systems, e.g., directional spherical wavelets [28].

## CONCLUSION

Processing data from cutting-edge experiments is often challenging because new science is always expected at the instrumental limits (otherwise, it would have been done already). It is the hope of this author that readers have been given a glimpse of some of the issues facing CMB data analysis. The reader must be aware that several complications have been left out and that signal processing research in this area is still in progress. For instance, the two-scale scheme described earlier is relatively well established, but it is not so clear yet how to combine fine-scale and coarse likelihood in a better way than just coadding the respective log-likelihoods. Another key topic is component separation (CMB cleaning or foreground removal). Many competing approaches are under study [16], but it will be difficult to single out the best one without knowing the ground truth or without very accurate models of all microwave emissions.

The article is mostly illustrated by the remarkable achievements of NASA’s WMAP mission. However, more data processing challenges are expected with the forthcoming Planck mission from the European Space Agency [9]. <AU: can “forthcoming” be updated?> Increased angular resolution and sensitivity will require increased accuracy in data processing. For instance, spectral leakage from the most energetic low- $\ell$  multipoles into the faintest high- $\ell$  multipoles cannot be controlled without a detailed knowledge of the instrument’s beams. More importantly, Planck’s instruments will measure the polarization of CMB [25] with unprecedented accuracy. Spectral analysis of polarized

fields on the sphere offers new challenges and may be the topic of another article in *IEEE Signal Processing Magazine*.

## ACKNOWLEDGMENTS

NASA is acknowledged for COBE and WMAP Figures 1–4, 7, 11, and 12 and the HEALPix authors for Figure 9 and their software package [20]. I thank the anonymous reviewers for very constructive comments.

## AUTHOR

*Jean-François Cardoso* (cardoso@tsi.enst.fr) is directeur de recherche with the French Centre National de la Recherche Scientifique (the National Center for Scientific Research) at Télécom ParisTech. Since 2001, he has worked on the Planck collaboration with a focus on the statistical and computational aspects of component separation, CMB likelihood, and cosmological parameter estimation.

## REFERENCES

- [1] G. B. Arfken and H. J. Weber, *Mathematical Methods for Physicists*, 5th ed. San Diego, CA: Harcourt/Academic Press, 2001.
- [2] J. R. Bond, A. H. Jaffe, and L. E. Knox, “Radical compression of cosmic microwave background data,” *Astrophys. J.*, vol. 533, pp. 19–37, 2000. <AU: please supply issue number or month>
- [3] R. G. Crittenden and N. G. Turok, “Exactly azimuthal pixelizations of the sky,” *ArXiv Astrophysics e-prints*, June 1998. <AU: Is this a journal? If yes, please provide the volume and complete page number.>
- [4] J. Delabrouille, J. F. Cardoso, M. Le Jeune, M. Betoule, G. Fay, and F. Guilloux. (2008). A full sky, low foreground, high resolution CMB map from WMAP. *Astron. Astrophys.* [Online]. Available: <http://arxiv.org/abs/0807.0773>
- [5] S. Dodelson, *Modern Cosmology*. New York: Academic, 2003. <AU: please supply page numbers>
- [6] 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. (2005). Gauss–Legendre sky pixelization (GLESP) for CMB maps. *Int. J. Modern Phys. D* [Online]. 14, 275. Available: <http://arxiv.org/abs/astro-ph/0305537>
- [7] J. R. Driscoll and D. M. Healy, “Computing Fourier transforms and convolutions on the 2-sphere,” *Adv. Appl. Math.*, vol. 15, no. 2, pp. 202–250, June 1994.
- [8] G. Efstathiou, “Myths and truths concerning estimation of power spectra: The case for a hybrid estimator,” *Monthly Notices Royal Astron. Soc.*, vol. 349, pp. 603–626, Apr. 2004.
- [9] G. Efstathiou, C. R. Lawrence, and J. Tauber. (2005 June). Planck scientific programme. *Eur. Space Agency* [Online]. ESA-SCI-1. Available: [http://www.rssd.esa.int/SA/PLANCK/docs/Bluebook-ESA-SCI\(2005\)1\\_V2.pdf](http://www.rssd.esa.int/SA/PLANCK/docs/Bluebook-ESA-SCI(2005)1_V2.pdf)
- [10] ESA. The Planck Web site at the European Spatial Agency [Online]. Available: <http://www.esa.int/Planck> <AU: Please provide the year.>
- [11] C. L. Bennett, et al., “First year Wilkinson microwave anisotropy probe (WMAP) observations: Preliminary maps and basic results,” *Astrophys. J.*, vol. 148, pp. 1, 2003. <AU: Please provide the complete list of author names and complete page range and year.>
- [12] D. Spergel, et al., “First-year Wilkinson microwave anisotropy probe (WMAP) observations: Determination of cosmological parameters,” *Astrophys. J.* (Suppl. Ser.), vol. 148, pp. 175–194, Sept. 2003. <AU: Please provide the complete list of author names.>
- [13] G. Hinshaw, et al., “Five-year Wilkinson microwave anisotropy probe (WMAP) observations: Data processing, sky maps, and basic results,” *eprint arXiv*, 0803.732, Mar. 2008. <AU: Please provide the complete list of author names, journal name, volume and complete page numbers.>
- [14] J. Dunkley, et al., “Five-year Wilkinson microwave anisotropy probe (WMAP) observations: Likelihoods and parameters from the WMAP data,” *Astrophys. J.* (Suppl. Ser.), vol. 180, pp. 306–376, Feb. 2009.
- [15] M. A. J. Ashdown, C. Baccigalupi, J. G. Bartlett, J. Borrill, C. Cantalupo, G. de Gasperis, G. de Troia, K. M. Gorski, E. Hivon, K. Huffenberger, E. Keihanen, R. Kesitalo, T. Kisner, H. Kurki-Suonio, C. R. Lawrence, P. Natoli, T. Poutanen, G. Prezeau, M. Reinecke, G. Rocha, M. Sandri, R. Stomp, F. Villa, and B. Wandelt (2009). Making maps from Planck LFI 30GHz data with asymmetric beams and cooler noise. *Astron. Astrophys.* [Online]. 493, 753. Available: <http://arxiv.org/abs/0806.3167> <AU: Please check that the list of author names is correct.>
- [16] S. M. Leach, J.-F. Cardoso, C. Baccigalupi, R. B. Barreiro, M. Betoule, J. Bobin, A. Bonaldi, J. Delabrouille, G. de Zotti, C. Dickinson, H. K. Eriksen, J. González-Nuevo, F. K. Hansen, D. Herranz, M. LeJeune, M. López-Cañiego, E.

Martinez-González, M. Massardi, J. -B. Melin, M.-A. Miville-Deschênes, G. Patanchon, S. Prunet, S. Ricciardi, E. Salerno, J. L. Sanz, J.-L. Starck, F. Stivoli, V. Stolyarov, R. Stompór, and P. Vielva. (2008 Nov.). Component separation methods for the Planck mission. *Astron. Astrophys.* 491, 597–615. Available: <http://arxiv.org/abs/0805.0269><AU: Please check that author names are correct.>

[17] T. Poutanen, G. de Gasperis, E. Hivon, H. Kurki-Suonio, A. Balbi, J. Borrill, C. Cantalupo, O. Doré, E. Keihanen, C. Lawrence, D. Maino, P. Natoli, S. Prunet, R. Stompór, and R. Teyssier (2006). Comparison of map-making algorithms for cmb experiments. *Astron. Astrophys.* [Online]. Available: <http://arxiv.org/abs/astro-ph/0501504> <AU: Please check that the complete list of author names is correct. Please clarify what 449, 1311–1322 are, page numbers? Volume?>

[18] G. Fay, F. Guilloux, M. Betoule, J. F. Cardoso, J. Delabrouille, and M. LeJeune. (2008). “CMB power spectrum estimation using wavelets.” *Phys. Rev. D* [Online]. 78, 083013. Available: <http://arxiv.org/abs/0807.1113>

[19] B. Gold, C. L. Bennett, R. S. Hill, G. Hinshaw, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, J. Dunkley, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, D. Larson, S. S. Meyer, M. R. Nolta, E. Wollack, and E. L. Wright (2009 Feb.). Five-year Wilkinson microwave anisotropy probe (WMAP) observations: Galactic foreground emission. *Astrophys. J. Suppl.* [Online]. 180, 265–282. Available: <http://arxiv.org/abs/0803.0715><AU: Please check that the complete list of author names is correct.>

[20] K. M. Górski, A. J. Banday, E. Hivon, and B. D. Wandelt, “HEALPix—A framework for high resolution, fast analysis on the sphere.” *Astron. Data Anal. Softw. Syst. XI*, vol. 281, p. 107, Jan. 2002.<AU: Please provide the complete page range.>

[21] M. Gräf, S. Kunis, and D. Potts, “On the computation of nonnegative quadrature weights on the sphere,” *Appl. Computat. Harmonic Anal.*, vol. 27, pp. 124–132, July 2009.

[22] S. Hamimeche and A. Lewis, “Properties and use of cmb power spectrum likelihoods,” *Phys. Rev. D*, vol. 79, p. 083012, 2009.<AU: Please provide the issue number.>

[23] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, “MASTER of the cosmic microwave background anisotropy power spectrum: A fast method for statistical analysis of large and complex cosmic microwave background data sets,” *Astrophys. J.*, vol. 567, pp. 2–17, Mar. 2002.

[24] W. Hu. CMB introduction [Online]. Available: <http://background.uchicago.edu/~whu/><AU: Please provide the year.>

[25] W. Hu and M. White, “A CMB polarization primer,” *New Astron.*, vol. 2, pp. 323, 1997.<AU: Please provide the issue number and complete page range.>

[26] A. Lewis and A. Challinor. CAMB: Code for anisotropies in the microwave background [Online]. Available: <http://camb.info><AU: Please provide the year.>

[27] D. Marinucci, D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkycharian, P. Natoli, D. Picard, and N. Vittorio, “Spherical needlets for CMB data analysis,” *Monthly Notices Royal Astron. Soc.*, vol. 383, no. 2, pp. 539–545, Jan. 2008.

[28] J. D. McEwen, M. P. Hobson, A. N. Lasenby, and D. J. Mortlock, “A high-significance detection of non-Gaussianity in the WMAP 3-yr data using directional spherical wavelets,” *MNRAS*, vol. 371, pp. L50–L54, Sept. 2006.

[29] NASA. Universe 101 [Online]. Available: <http://map.gsfc.nasa.gov/universe><AU: Please provide the year.>

[30] A. A. Penzias and R. W. Wilson, “A measurement of excess antenna temperature at 4080-Mc/s,” *Astrophys. J.*, vol. 142, pp. 419–421, 1965.<AU: Please provide the issue number.>

[31] U. Seljak and M. Zaldarriaga. CMBFast Web site [Online]. Available: <http://www.cfa.harvard.edu/~mzaldarr/CMBFAST/cmbfast.html><AU: Please provide the year.>

[32] J.-L. Starck, Y. Moudden, P. Abrial, and M. Nguyen, “Wavelets, ridgelets and curvelets on the sphere,” *Astron. Astrophys.*, vol. 446, pp. 1191–1204, Feb. 2006.

[33] M. Tegmark, “How to measure CMB power spectra without losing information,” *Phys. Rev. D*, vol. 55, pp. 5895–5907, May 1997.

[34] M. Tegmark. Max’s cosmic cinema [Online]. Available: <http://space.mit.edu/home/tegmark/cmb/movies.html><AU: Please provide the year.>

[35] M. Tegmark. (1996). An icosahedron-based method for pixelizing the celestial sphere. *APJ Lett.* [Online]. 470. Available: <http://arxiv.org/abs/astro-ph/9610094>

[36] P. Vielva, E. Martinez-Gonzalez, R. B. Barreiro, J. L. Sanz, and L. Cayon, “Detection of non-Gaussianity in the WMAP 1-year data using spherical wavelets,” *Astrophys. J.*, vol. 609, pp. 22–34, 2004.<AU: Please provide the issue number.>

[37] M. White. Online publications and movies [Online]. Available: <http://astro.berkeley.edu/~mwhite/htmlpapers.html><AU: Please provide the year.>

[38] Y. Wiaux, J. D. McEwen, P. Vanderghenst, and O. Blanc, “Exact reconstruction with directional wavelets on the sphere,” *Monthly Notices Royal Astron. Soc.*, vol. 388, pp. 770–788, Aug. 2008.

[39] N. Wright. Cosmology tutorial [Online]. Available: <http://www.astro.ucla.edu/~wright/cosmolog.htm><AU: Please provide the year.>

SP

## CALLOUTS

A critical step of the CMB data processing pipeline is the construction of a likelihood function for the statistical exploitation of CMB data.

The CMB anisotropies can be understood as marking the seeds of large scale structures of the Universe.

Another nice property of HEALPix is its dyadic structure: the pixels can be arranged in memory in nested order, allowing for efficient local processing.

Imaging instruments have limited resolution, producing only smooth, band-limited versions of the underlying field.

As of today, efficiency, exactness, and speed are slightly conflicting requirements.