Low- ℓ CMB analysis and inpainting

J.-L. Starck¹, M. J. Fadili², and A. Rassat^{3,1}

¹ Laboratoire AIM, UMR CEA-CNRS-Paris 7, Irfu, Service d'Astrophysique, CEA Saclay, 91191 Gif-sur-Yvette Cedex, France
e-mail: jstarck@cea.fr

² GREYC CNRS-ENSICAEN-Université de Caen, 6 Bd du Maréchal Juin, 14050 Caen Cedex, France

³ Laboratoire d'Astrophysique, École Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 1290 Versoix, Switzerland

Received 4 September 2012 / Accepted 26 October 2012

ABSTRACT

Reconstructing the cosmic microwave background (CMB) in the Galactic plane is extremely difficult due to the dominant foreground emissions such as dust, free-free or synchrotron. For cosmological studies, the standard approach consists in masking this area where the reconstruction is insufficient. This leads to difficulties for the statistical analysis of the CMB map, especially at very large scales (to study for instance the low quadrupole, integrated Sachs Wolfe effect, axis of evil, etc.). We investigate how well some inpainting techniques can recover the low- ℓ spherical harmonic coefficients. We introduce three new inpainting techniques based on three different kinds of priors: sparsity, energy, and isotropy, which we compare. We show that sparsity and energy priors can lead to extremely high-quality reconstruction, within 1% of the cosmic variance for a mask with a sky coverage larger than 80%.

Key words. methods: data analysis – techniques: image processing – cosmic background radiation

1. Introduction

Because component separation methods do not provide good estimates of the cosmic microwave background (CMB) in the Galactic plane and at locations of point sources, the standard approach for a CMB map analysis is to consider that the data are not reliable in these areas, and to mask them. This leads to an incomplete coverage of the sky that has to be handled properly for the subsequent analysis. This is especially true for analysis methods that operate in the spherical harmonic domain where localization is lost and full-sky coverage is assumed. For power spectrum estimations, methods such as MASTER (Hivon et al. 2002) solve a linear ill-posed inverse problem that allows to deconvolve the observed power spectrum of the masked map from the mask effect.

For a non-Gaussianity analysis, many approaches to deal with this missing data problem have been proposed. Methods to solve these problems are called gap-filling or inpainting methods.

The simplest approach is to set the pixel values in the masked area to zero. This is sometimes claimed to be a good approach by assuming it does not add any information. However, this is not correct, as setting the masked area to zero actually adds (or removes) information, which again differs from the true CMB values. A consequence of this gap-filling technique is the creation of many artificial high amplitude coefficients at the border of the mask in a wavelet analysis or a leakage between different multipoles in a spherical harmonic analysis. This effect can be reduced with an apodized mask. A slightly more sophisticated method consists of replacing each missing pixel by the average of its neighbors and iterating until the gaps are filled. This technique is called diffuse inpainting and has been used in *Planck*-LFI data pre-processing Zacchei et al. (2011). It is acceptable for treatment of point sources, but is far from being a reasonable solution for the Galactic plane inpainting in a CMB non-Gaussianity analysis.

In Abrial et al. (2007; 2008), the problem was considered as an ill-posed inverse problem, $z = \mathbf{M}\mathbf{x}$, where $\mathbf{x}$ is the unknown CMB, $\mathbf{M}$ is the masking operator, and z the masked CMB map. A sparsity prior in the spherical harmonic domain was used to regularize the problem. This sparsity-based inpainting approach has been successfully used for two different CMB studies, the CMB weak-lensing on *Planck* simulated data (Perotto et al. 2010; Plaszczynski et al. 2012), and the analysis of the integrated Sachs-Wolfe effect (ISW) on WMAP data (Dupé et al. 2011). In both cases, the authors showed from Monte-Carlo simulations that the statistics derived from the inpainted maps can be trusted at a high confidence level, and that sparsity-based inpainting can indeed provide an easy and effective solution to the large Galactic mask problem.

It was also shown that sparse inpainting is useful for weak-lensing data (Pires et al. 2009), Fermi data (Schmitt et al. 2010), and asteroseismic data (Sato et al. 2010). The sparse inpainting success has motivated the community to investigate more inpainting techniques, and other approaches have recently been proposed.

Bucher & Louis (2012) and Kim et al. (2012) sought a solution that complies with the CMB properties, i.e. to be a homogeneous Gaussian random field with a specific power spectrum. Good results were derived, but the approach presents the drawback that we need to assume a given cosmology, which may not be appropriate for non-Gaussianity studies. For large scale CMB anomaly studies, Feeney et al. (2011) and Ben-David et al. (2012), building upon the work of de Oliveira-Costa & Tegmark (2006), proposed to use generalized least-squares (GLS), which coincide with the maximum-likelihood estimator (MLE) under appropriate assumptions such as Gaussianity. This method also requires an input power spectrum. In addition, when the covariance in the GLS is singular, the authors proposed to regularize it by adding a weak perturbation term that must also to ensure

positive-definiteness. With this regularization, the estimator is no longer a GLS (nor a MLE).

Because the missing data problem is ill-posed, prior knowledge is needed to reduce the space of candidate solutions. Methods that handle this problem in the literature assume priors, either explicitly or implicitly. If we exclude the zero-inpainting and the diffuse inpainting methods, which are of little interest for the Galactic plane inpainting, the two other priors are:

- **Gaussianity:** the CMB is assumed to be a homogeneous and isotropic Gaussian random field, therefore it is adequate to use such a prior. In practice, methods using this prior require the theoretical power spectrum, which has either to be estimated using a method such as TOUSI (Paykari et al. 2012), or a cosmology has to be assumed, which is an even stronger assumption than the Gaussianity prior. We should also keep in mind that the goal of non-Gaussianity studies is to check that the observed CMB does not deviate from Gaussianity. Therefore, we should be careful with this prior.
- **Sparsity:** the sparsity prior on a signal consists of assuming that its coefficients in a given representation domain, when sorted in decreasing order of magnitude, exhibit a fast decay rate, typically a polynomial decay with an exponent that depends on the regularity of the signal. Spherical harmonic coefficients of the CMB show a decay in $O(\ell^{-2})$ up to ℓ around 900 and then $O(\ell^{-3})$ for larger multipoles. Thus, the sparsity prior is advocated, and this explains its success for CMB-related inverse problems such as inpainting or component separation (Bobin et al. 2012).

We here revisit the Gaussianity and sparsity priors, and introduce an additional one, the CMB isotropy, to recover the spherical harmonic coefficients at low ℓ (<10) from masked data. We describe novel and fast algorithms to solve the optimization problems corresponding to each prior. These algorithms originate from the field of non-smooth optimization theory and are efficiently applied to large-scale data. We then show that some of these inpainting algorithms are very efficient to recover the spherical harmonic coefficients for $\ell < 10$ when using the sparsity or energy priors. We also study the reconstruction quality as a function of the sky coverage, and we show that a very good reconstruction quality, within 1% of the cosmic variance, can be reached for a mask with a sky coverage better than 80%.

2. CMB inpainting

2.1. Problem statement

We observe $\mathbf{z} = \mathbf{M}\mathbf{x}$, where $\mathbf{x}$ is a real-valued centered and square-integrable field on the unit sphere $\mathbb{S}^2$, and $\mathbf{M}$ is a linear bounded masking operator. The goal is to recover $\mathbf{x}$ from $\mathbf{z}$.

The field $\mathbf{x}$ can be expanded as

$$\mathbf{x}(p) = \sum_{\ell=0}^{+\infty} \sum_{m=-\ell}^{\ell} \mathbf{a}_{\ell,m} Y_{\ell m}(p),$$

$$\text{where } \mathbf{a}_{\ell,m} = \int_{\mathbb{S}^2} \mathbf{x}(p) Y_{\ell m}^*(p) dp,$$

where the complex-valued functions $Y_{\ell m}$ are the co-called spherical harmonics, ℓ is the multipole moment, and m is the phase ranging from $-\ell$ to ℓ . The $\mathbf{a}_{\ell,m}$ are the spherical harmonic coefficients of $\mathbf{x}$. In the following we will denote $\mathbf{S}$ the spherical harmonic transform operator and $\mathbf{S}^*$ its adjoint (hence its inverse since spherical harmonics form an orthobasis of $L_2(\mathbb{S}^2)$).

If $\mathbf{x}$ is a wide-sense stationary (i.e., homogeneous) random field, the spherical harmonic coefficients are uncorrelated,

$$\mathbb{E}[\mathbf{a}_{\ell,m}^* \mathbf{a}_{\ell',m'}] = \delta_{\ell\ell'} \delta_{mm'} C_{\ell,m}.$$

Moreover, if the process is isotropic, then

$$\mathbb{E}[|\mathbf{a}_{\ell,m}|^2] = C_{\ell}, \quad -\ell \leq m \leq \ell,$$

where C_{ℓ} is the angular power spectrum, which depends solely on ℓ .

In the remainder of the paper, we consider a finite-dimensional setting, where the sphere is discretized. $\mathbf{x}$ can be rearranged in a column vector in $\mathbb{R}^n$, and similarly for $\mathbf{z} \in \mathbb{R}^m$, with $m < n$, and $\mathbf{a} \in \mathbb{C}^p$. Therefore, $\mathbf{M}$ can be seen as a matrix taking values in $\{0, 1\}$, i.e., $\mathbf{M} \in \mathcal{M}_{m \times n}(\{0, 1\})$. The goal of inpainting is to recover $\mathbf{x}$ from $\mathbf{z}$.

2.2. a general inpainting framework

The recovery of $\mathbf{x}$ from $\mathbf{z}$ when $m < n$ amounts to solving an underdetermined system of linear equations. Traditional mathematical reasoning, in fact the fundamental theorem of linear algebra, tells us not to attempt this: there are more unknowns than equations. However, if we have prior information about the underlying field, there are some rigorous results showing that such an inversion might be possible (Starck et al. 2010).

In general, we can write the inpainting problem as the following optimization program

$$\text{find } \widehat{\mathbf{x}} \in \underset{\mathbf{x} \in \mathbb{R}^n}{\text{Argmin}} R(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{z} - \mathbf{M}\mathbf{x} \in C, \quad (1)$$

where R is a proper lower-bounded penalty function reflecting some prior on $\mathbf{x}$, and C is closed constraint set expressing the fidelity term. Typically, in the noiseless case, $C = \{\mathbf{x} : \mathbf{z} = \mathbf{M}\mathbf{x}\}$. This is what we focus on in the remainder of the paper. We consider three types of priors, each corresponding to a specific choice of R .

3. Sparsity prior

Sparsity-based inpainting has been proposed for the CMB in (Abrial et al. 2007, 2008) and for weak-lensing mass map reconstruction in Pires et al. (2009, 2010). In Perotto et al. (2010), Dupé et al. (2011) and Rassat et al. (2012), it was shown that sparsity-based inpainting does not destroy CMB weak-lensing, ISW signals or some large-scale anomalies in the CMB, and is therefore an elegant way to handle the masking problem. The masking effect can be thought of as a loss of sparsity in the spherical harmonic domain because the information required to define the map has been spread across the spherical harmonic basis (leakage effect).

Optimization problem

If the spherical harmonic coefficients $\mathbf{a}$ of $\mathbf{x}$ (i.e. $\mathbf{a} = \mathbf{S}\mathbf{x}$) are assumed to be sparse, then a well-known penalty to promote sparsity is the l_q (pseudo- or quasi-)norm, with $q \in [0, 1]$. Therefore, Eq. (1) becomes

$$\text{find } \widehat{\mathbf{a}} \in \underset{\mathbf{a} \in \mathbb{C}^p}{\text{Argmin}} \|\mathbf{a}\|_q^q \quad \text{s.t.} \quad \mathbf{z} = \mathbf{M}\mathbf{S}^* \mathbf{a}, \quad (2)$$

where $\|\mathbf{a}\|_q^q = \sum_i |\mathbf{a}_i|^q$, and where $|z| = \sqrt{\Re(z)^2 + \Im(z)^2}$ for $z \in \mathbb{C}$. For $q = 0$, the l_0 pseudo-norm counts the number of non-zero entries of its argument. The inpainted map is finally reconstructed as $\widehat{\mathbf{x}} = \mathbf{S}^* \widehat{\mathbf{a}}$.

Fig. 1. Relative MSE per ℓ in percent for the two sparsity-based inpainting algorithms, IHT (dashed red line) and DR (continuous black line).

Solving Eq. (2) when $q = 0$ is known to be NP-hard. This is additionally complicated by the non-smooth constraint term. Iterative Hard thresholding (IHT), described in Appendix A, attempts to solve this problem. It is also well-known that the l_1 norm is the tightest convex relaxation (in the ℓ_2 ball) of the l_0 penalty (Starck et al. 2010). This suggests solving Eq. (2) with $q = 1$. In this case, the problem is well-posed: it has at least a minimizer, and all minimizers are global. Furthermore, although it is completely non-smooth, it can be solved efficiently with a provably converging algorithm that belongs to the family of proximal splitting schemes (Combettes & Pesquet 2011; Starck et al. 2010). This can be done with the Douglas-Rachford (DR) algorithm described in Appendix B.

Comparison between IHT and DR

We compared the IHT and DR sparsity-based inpainting algorithms in 100 Monte-Carlo simulations using a mask with sky coverage $F_{\text{sky}} = 77\%$. In all our experiments, we used 150 iterations for both iterative schemes, $\beta = 1$ and $\alpha_n \equiv 1$ ($\forall n$) in the DR scheme (see Appendix B). For each inpainted map $i \in \{1, \dots, 100\}$, we computed the relative mean squared-error (MSE)

$$e^{(i)}[\ell] = \left\langle \frac{|\mathbf{a}_{\ell,m}^{\text{True}} - \mathbf{a}_{\ell,m}^{(i)}|^2}{C_\ell} \right\rangle_m$$

and its version in percent per ℓ

$$E[\ell] = 100 \times \langle e^{(i)}[\ell] \rangle_i \quad (\%).$$

Figure 1 depicts the relative MSE in percent per ℓ for the two sparsity-based inpainting algorithms IHT (l_0) and DR (l_1). We see that at very low ℓ , l_1 -sparsity inpainting as provided by the DR algorithm yields better results. We performed the test with other masks and arrived at similar conclusions. Because we focus on low- ℓ , only the l_1 inpainting as solved by the DR algorithm is considered below.

Fig. 2. Set $\mathcal{B}_\ell$.

4. Energy prior

Optimization problem

If we know a priori the power-spectrum $(C_{\ell,m})_{\ell,m}$ (not the angular one, which implicitly assumes isotropy) of the Gaussian field $\mathbf{x}$, then using a maximum a posteriori (MAP) argument, the inpainting problem amounts to minimizing a weighted ℓ_2 -norm subject to a linear constraint

$$\text{find } \hat{\mathbf{x}} = \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \|\mathbf{S}\mathbf{x}\|_{C^{-1}} \quad \text{s.t.} \quad \mathbf{z} = \mathbf{M}\mathbf{x}, \quad (3)$$

where for a complex-valued vector $\mathbf{a}$, $\|\mathbf{a}\|_{C^{-1}}^2 = \sum_{\ell,m} \frac{|a_{\ell,m}|^2}{C_{\ell,m}}$, i.e. a weighted ℓ_2 norm. By strong convexity, Problem Eq. (3) is well-posed and has exactly one minimizer, therefore justifying equality instead of inclusion in the argument of the minimum in Eq. (3).

Since the objective is differentiable with a Lipschitz-continuous gradient, and the projector $\text{Proj}_{\{\mathbf{x}:\mathbf{z}=\mathbf{M}\mathbf{x}\}}$ on the linear equality set is known in closed-form, one can use a projected gradient-scheme to solve Eq. (3). However, it turns that the estimate of the descent step-size can be rather crude, which may hinder the convergence speed. One can think of using an inexact line-search, but this will complicate the algorithm unnecessarily. This is why we propose a new algorithm in Appendix C, based on the Douglas-Rachford splitting method, which is easy to implement and efficient.

In all our experiments in the remainder of the paper, we used 150 iterations for this algorithm with $\beta = 50$ and $\alpha_n \equiv 1$ ($\forall n$) (see Appendix C). This approach requires the power-spectrum $C_{\ell,m}$ as an input parameter. In practice, an estimate $\hat{C}_\ell$ of the power spectrum from the data using MASTER was used, and we set $C_{\ell,m} = \hat{C}_\ell$. The latter is reminiscent of an isotropy assumption, which is not necessarily true.

5. Isotropy prior

5.1. Structural constraints on the power spectrum

Strictly speaking, we would define the set of isotropic homogeneous random processes on the (discrete) sphere as the closed set

$$C_{\text{iso}} := \{\mathbf{x} : C_{\ell,m} = C_\ell, \forall -\ell \leq m \leq \ell\},$$

Fig. 3. Normalized critical thresholds at the significance level 0.05 as a function of ℓ .

where $C_{\ell,m} = \mathbb{E} \left[|(\mathbf{S}\mathbf{x})_{\ell,m}|^2 \right]$ and C_ℓ is the angular power spectrum, which depends solely on ℓ . Given a realization $\mathbf{x}$ of a random field, we can then naively state the orthogonal projection of $\mathbf{x}$ onto C_{iso} by solving

$$\min_{\mathbf{v} \in C_{\text{iso}}} \|\mathbf{v} - \mathbf{x}\|_2.$$

This formulation of the projection constraint set is not straightforward to deal with for at least the following reasons: (i) the isotropy constraint set involves the unknown true C_ℓ (through the expectation operator), which necessitates resorting to stochastic programming; (ii) the constraint set is not convex.

5.2. Projection with a deterministic constraint

An alternative to the constraint set C_{iso} would be to replace $C_{\ell,m}$ and C_ℓ by their empirical sample estimates, i.e., $C_{\ell,m}$ by $|\mathbf{a}_{\ell,m}|^2$ and C_ℓ by $\widehat{C}_\ell = \frac{1}{2\ell+1} \sum_m |\mathbf{a}_{\ell,m}|^2$. However, this hard constraint might be too strict in practice and we propose to make it softer by taking into account the variability inherent to the sample estimates, as we explain now.

In a nutshell, the goal is to formulate a constraint set, where for each ℓ , the entries $\mathbf{a}_{\ell,m}$ of the spherical harmonic coefficient vector $\mathbf{a}$ deviate the least possible in magnitude (up to a certain accuracy) from some pre-specified vector $\boldsymbol{\mu}$; typically we take

$\boldsymbol{\mu} = \sqrt{C_\ell}$ or its empirical estimate $\sqrt{\widehat{C}_\ell}$. Put formally, this reads

$$C_\varepsilon = \{\mathbf{x} : \mathbf{a} = \mathbf{S}\mathbf{x}, \|\mathbf{a}_{\ell,m} - \boldsymbol{\mu}_\ell\| \leq \varepsilon, \forall \ell, m\}.$$

C_ε is a compact set, although not convex.

We now turn to the projection on C_ε . We begin by noting that this set is separable, i.e., $C_\varepsilon = \times_{\ell,m} \mathcal{B}_\ell$,

where $\mathcal{B}_\ell$ is the band depicted Fig. 2. It turns out that for fixed $\boldsymbol{\mu}$, $\mathcal{B}_\ell$ is so-called prox-regular since its associated orthogonal projector is single-valued with a closed-form. Indeed, the projector onto C_ε is

$$\text{Proj}_{C_\varepsilon}(\mathbf{x}) = \mathbf{S}^* \left(\text{Proj}_{\mathcal{B}_\ell}(\mathbf{a}_{\ell,m}) \right)_{\ell,m}, \quad (4)$$

where

$$\text{Proj}_{\mathcal{B}_\ell}(\mathbf{a}_{\ell,m}) = \begin{cases} (\boldsymbol{\mu}_\ell + \varepsilon) \mathbf{a}_{\ell,m} / |\mathbf{a}_{\ell,m}| & \text{if } |\mathbf{a}_{\ell,m}| > \boldsymbol{\mu}_\ell + \varepsilon \\ (\boldsymbol{\mu}_\ell - \varepsilon) \mathbf{a}_{\ell,m} / |\mathbf{a}_{\ell,m}| & \text{if } |\mathbf{a}_{\ell,m}| < \boldsymbol{\mu}_\ell - \varepsilon \\ \mathbf{a}_{\ell,m} & \text{otherwise.} \end{cases} \quad (5)$$

Choice of the constraint radius: to devise a meaningful choice (from a statistical perspective) of the constraint radius ε , we first need to derive the null distribution¹ of $|\mathbf{a}_{\ell,m}| - \boldsymbol{\mu}_\ell, \forall(\ell, m)$, with the particular case when $\boldsymbol{\mu}_\ell = \sqrt{\widehat{C}_\ell}$. $\mathbf{a}_{\ell,m}$ being the spherical harmonic coefficients of a zero-mean stationary and isotropic Gaussian process whose angular power spectrum is C_ℓ , it is easy to show that

$$2|\mathbf{a}_{\ell,m}|^2 \sim C_\ell \chi^2(2) \quad \text{and}$$

$$2L\boldsymbol{\mu}_\ell^2 \sim C_\ell \chi^2(2L), \quad \forall \ell \geq 2,$$

where $L = 2\ell + 1$. By the Fisher asymptotic formula, we then obtain

$$\begin{aligned} |\mathbf{a}_{\ell,m}| &\xrightarrow{d} \sqrt{C_\ell} \mathcal{N}\left(\sqrt{\frac{3}{4}}, \frac{1}{4}\right) \quad \text{and} \\ \boldsymbol{\mu}_\ell &\xrightarrow{d} \sqrt{C_\ell} \mathcal{N}\left(\sqrt{1 - \frac{1}{4L}}, \frac{1}{4L}\right), \quad \forall \ell \geq 2, \end{aligned}$$

where $\xrightarrow{d}$ means convergence in distribution. Thus, ignoring the obvious dependency between $|\mathbf{a}_{\ell,m}|$ and $\boldsymbol{\mu}_\ell$, we obtain

$$|\mathbf{a}_{\ell,m}| - \boldsymbol{\mu}_\ell \xrightarrow{d} \sqrt{C_\ell} \mathcal{N}\left(\sqrt{\frac{3}{4}} - \sqrt{1 - \frac{1}{4L}}, \frac{L+1}{4L}\right).$$

The distribution of this difference can be derived properly using the delta method in law, but computing the covariance matrix of the augmented vector $(\mathbf{a}_{\ell,m})_m$, remains a problem. The quality of the above asymptotic approximation increases as ℓ increases.

The upper and lower critical thresholds at the double-sided significance level $0 \leq \alpha \leq 1$ are

$$\varepsilon_\pm = \sqrt{C_\ell} \left(\sqrt{\frac{3}{4}} - \sqrt{1 - \frac{1}{4L}} \pm \sqrt{\frac{L+1}{4L}} \Phi^{-1}\left(1 - \frac{\alpha}{2}\right) \right),$$

where Φ is the standard normal cumulative distribution function.

We depict in Fig. 3 the upper and lower critical thresholds normalized by $\sqrt{C_\ell}$ at the classical significance level 0.05 as a function of ℓ . One can observe that the thresholds are not symmetric. This entails that different values of ε should be used in Eq. (5). Furthermore, as expected, the two thresholds decrease in magnitude as ℓ increases. They attain a plateau for ℓ that is sufficiently high (typically ≥ 100). It is easy to see that

¹ That is the distribution under the isotropy hypothesis.

Fig. 4. Six masks with $F_{\text{sky}} \in \{0.98, 0.93, 0.87, 0.77, 0.67, 0.57\}$. The masked area is in blue.

the two limit values are $\sqrt{3/4} - 1 + 1.96/\sqrt{8} = 0.559$ and $\sqrt{3/4} - 1 - 1.96/\sqrt{8} = -0.8269$.

The isotropy-constrained inpainting algorithm is given in Appendix D.

6. Experiments

6.1. $a_{\ell,m}$ reconstruction

In this section, we compare the different inpainting methods described above: the DR-sparsity prior inpainting, the DR-energy prior inpainting, and the isotropy prior inpainting. We also test the case of no-inpainting by applying just an Fsky correction to the $a_{\ell,m}$ spherical harmonic coefficients of the map (i.e., data $a_{\ell,m}$ are corrected with a correction term equal to $1/\sqrt{F_{\text{sky}}}$). We ran 100 CMB Monte-Carlo simulations with a resolution corresponding to n_{side} equal to 32, which is precise enough for the large scales we are considering. We

considered six different masks with sky coverage F_{sky} of $\{0.98, 0.93, 0.87, 0.77, 0.67, 0.57\}$.

Figure 4 displays these masks. Point source masks were not considered here, since they should not affect $a_{\ell,m}$ estimation at very low multipole. Each of these simulated CMB maps was masked with the six masks. Figure 5 top right shows one CMB realization masked with the 77% Fsky mask. Figure 5 middle and bottom show the inpainting of the Fig. 5 top right image with the four methods (i.e., using sparsity, energy, isotropy priors, and Fsky correction). Figure 6 depicts the results given by the ℓ_1 sparsity-based inpainting method for the six different masks. We can clearly see that the quality of the reconstruction degrades with decreasing sky coverage.

The $a_{\ell,m}$ coefficients of the six hundred maps (100 realizations masked with the six different masks) were then estimated using the three inpainting methods and the Fsky correction methods. The approaches were compared in terms of the

Fig. 5. *Top:* input-smoothed simulated CMB map ($\ell_{\max} = 10$) and the same map, but masked ($F_{\text{sky}} = 77\%$) and not smoothed (i.e., input-simulated data). *Middle:* inpainting of the top right image up to $\ell_{\max} = 10$, using a sparsity prior (*left*) and an energy prior (*right*). *Bottom:* inpainting of the top right image up to $\ell_{\max} = 10$, using an isotropy prior (*left*) and a simple F_{sky} correction (*right*).

relative MSE and relative MSE per ℓ ; see end of Sect. 3 for their definition.

Figure 7 shows the relative MSE for $\ell = 2$ to 5 versus the sky coverage F_{sky} . The three horizontal lines correspond to 1%, 5%, and 10% of the cosmic variance (CV). Inpainting based on the sparsity and energy priors give relatively close results, which are better than the one assuming the isotropy prior or the F_{sky} correction. It is interesting to notice that a very high quality reconstruction (with 1% of the CV) can be obtained with both the sparsity- and energy-based inpainting methods up to $\ell = 4$ for a mask with an F_{sky} exceeding 80%. If WMAP data do not allow us to study non-Gaussianity with such a small mask, *Planck* component separation will certainly be able to achieve the required quality to enable the use of such a small Galactic mask. With a 77%-coverage Galactic mask, the error increases by a factor 5 at $\ell = 4$!

Figure 8 shows the same errors, but they are now plotted versus the multipoles for the six different masks. The ℓ_1 sparsity-based inpainting method seems to be slightly better than the one based on the energy prior, especially when the sky coverage decreases.

6.2. Anomalies in CMB maps

One of the main motivations of recovering full-sky maps is to be able to study statistical properties such as Gaussianity and statistical isotropy on large scales. Statistical isotropy is violated if there exists a preferred axis in the map. Mirror parity, i.e., parity with respect to reflections through a plane: $\hat{\mathbf{x}} = \hat{\mathbf{x}} - 2(\hat{\mathbf{x}} \cdot \hat{\mathbf{n}})\hat{\mathbf{n}}$, where $\hat{\mathbf{n}}$ is the normal vector to the plane, is an example of a

Fig. 6. ℓ_1 sparsity-based inpainting of a simulated CMB map with different masks. *Top:* Fsky = 98% and 93%, *middle:* 87% and 77%, and *bottom:* 67% and 57%.

statistic where a preferred axis can be sought (Land & Magueijo 2005; Ben-David et al. 2012).

With full-sky data, one can estimate the S -map for a given multipole in spherical harmonics by considering (Ben-David et al. 2012):

$$\tilde{S}_\ell(\hat{\mathbf{n}}) = \sum_{m=-\ell}^{\ell} (-1)^{\ell+m} \frac{|a_{\ell m}(\hat{\mathbf{n}})|^2}{\hat{C}_\ell}, \quad (6)$$

where $a_{\ell m}(\hat{\mathbf{n}})$ corresponds to the value of the $a_{\ell m}$ coefficients when the map is rotated to have $\hat{\mathbf{n}}$ as the z -axis. Positive (negative) values of $\tilde{S}_\ell(\hat{\mathbf{n}})$ correspond to even (odd) mirror parities in the $\hat{\mathbf{n}}$ direction. The same statistic can also be considered summed over all low multipoles one wishes to consider (e.g., focusing only on low multipoles as in Ben-David et al. 2012):

$$\tilde{S}_{\text{tot}}(\hat{\mathbf{n}}) = \sum_{\ell=2}^{\ell_{\text{max}}} \tilde{S}_\ell(\hat{\mathbf{n}}). \quad (7)$$

It is convenient to redefine the parity estimator as $S(\hat{\mathbf{n}}) = \tilde{S}_{\text{tot}}(\hat{\mathbf{n}}) - (\ell_{\text{max}} - 1)$, so that $\langle S \rangle = 0$.

The most even and odd mirror-parity directions for a given map can be considered by estimating (Ben-David et al. 2012)

$$S_+ = \frac{\max(S) - \mu(S)}{\sigma(S)}, \quad (8)$$

$$S_- = \frac{|\min(S) - \mu(S)|}{\sigma(S)}, \quad (9)$$

where $\mu(S)$ and $\sigma(S)$ are the mean and standard deviation of the S map. The quantities S_+ and S_- each correspond to different axes $\hat{\mathbf{n}}_+$ and $\hat{\mathbf{n}}_-$ and are the quantities which we consider when discussing mirror parity in CMB maps.

To use inpainted maps to study mirror parity, it is crucial that the inpainting method constitutes a bias-free reconstruction

Fig. 7. Relative MSE (in percent) per ℓ versus sky coverage for different multipoles. See text for details.

Fig. 8. Relative MSE in percent per sky coverage versus ℓ for the six masks.

Fig. 9. Percentage difference between anomalous (i.e., high) odd mirror parity scores (S_-) before (i.e., full-sky) and after inpainting for three different inpainting priors.

method, i.e., that it does not destroy existing mirror parities nor that it creates previously inexistent mirror parity anomalies.

In order to test this, we considered three sets of 1000 simulated CMB maps using a WMAP7 best-fit cosmology with $n_{\text{side}} = 32$, one set for each inpainting prior. For each simulation we calculated the mirror parity estimators $S_{\pm}$. We isolated the anomalous maps, defined as those whose $S_{\pm}$ value is higher than twice the standard deviation given by the simulations. This yielded 34 maps (3.4%) for the even mirror parity and 40 maps (4%) for the odd mirror parity in the full-sky simulations.

Since recent work showed a potential odd-mirror anomaly in WMAP data (Ben-David et al. 2012), we focused on testing for potential biases in odd-mirror anomalies. Figure 9 shows the percentage difference between the true S_- value and the one estimated from the inpainted map. For this statistical test, the result does not depend on the inpainting method because the three different priors give similar biases and error bars. For positive-mirror anomalies, we find similar results (not shown in figure). Similarly to the previous experiment, we can see that a very good estimation of the parity statistic can be achieved for masks with Fsky exceeding 0.8.

7. Software and reproducible research

To support reproducible research, the developed IDL code will be released with the next version of ISAP (Interactive Sparse astronomical data Analysis Packages) via the web site: <http://www.cosmostat.org>

All experiments were performed using the default options

– Sparsity:

```
Alm = cmb_lowl_alm_inpainting(CMBMap,
                             Mask, /sparsity)
```

– Energy:

```
Pd = mrs_powspec( CMBMap, Mask)
```

```
P = mrs_deconv_powspec( pd, Mask)
Alm = cmb_lowl_alm_inpainting(CMBMap,
                              /Energy, Prea=P)
```

– Isotropy:

```
Pd = mrs_powspec( CMBMap, Mask)
P = mrs_deconv_powspec( Pd, Mask)
Alm = cmb_lowl_alm_inpainting(CMBMap,
                              /Isotropy, Prea=P)
```

8. Conclusion

We have investigated three priors to regularize the CMB inpainting problem: Gaussianity, sparsity, and isotropy. To solve the corresponding minimization problems, we also proposed fast novel algorithms, based for instance on proximal splitting methods for convex non-smooth optimization. We found that both Gaussianity and ℓ_1 sparsity priors lead to very good results. The isotropy prior does not provide comparably good results. The sparsity prior seems to lead to slightly better results for $\ell > 2$, and more robust when the sky coverage decreases. Furthermore, unlike the energy-prior based inpainting, the sparsity-based one does not require a power spectrum as an input, which is a significant advantage.

Then we evaluated the reconstruction quality as a function of the sky coverage and saw that a high-quality reconstruction, within 1% of the cosmic variance, can be reached for a mask with an Fsky exceeding 80%. We also studied mirror-parity anomalies and found that mask with an Fsky exceeding 80% also permitted nearly bias-free reconstructions of the mirror parity scores. Such a mask seems unrealistic for a WMAP data analysis, but it could and should be a target for *Planck* component separation methods. To know if this 80% mask is realistic for *Planck*, we will need more realistic simulations including the ISW effect, and residual foregrounds at amplitudes similar to those achieved by the actual best *Planck* component separation methods. Another interesting study will consist of comparing the sparse inpainting to other methods such as the maximum likelihood or the Wiener filtering (Feeney et al. 2011; Ben-David et al. 2012).

Acknowledgements. The authors would like to thank Hiranya Peiris and Stephen Feeney for useful discussions. This work was supported by the French National Agency for Research (ANR -08-EMER-009-01), the European Research Council grant SparseAstro (ERC-228261), and the Swiss National Science Foundation (SNSF).

Appendix A: Algorithm for the l_0 problem

Solving Eq. (2) when $q = 0$ is known to be NP-hard. This is further complicated by the presence of the non-smooth constraint term. Iterative Hard thresholding (IHT) attempts to solve this problem through the following scheme

$$\mathbf{a}^{n+1} = \Delta_{\lambda_n}^H (\mathbf{a}^n + \mathbf{S}(\mathbf{z} - \mathbf{M}\mathbf{S}^* \mathbf{a}^n)), \quad (10)$$

where the nonlinear operator Δ_{λ}^H is a term-by-term hard thresholding, i.e. $\Delta_{\lambda}^H(\mathbf{a}) = (\rho_{\lambda}^H(\mathbf{a}_i))_i$, where $\rho_{\lambda}^H(\mathbf{a}_i) = \mathbf{a}_i$ if $|\mathbf{a}_i| > \lambda$ and 0 otherwise. The threshold parameter λ_n decreases with the iteration number and is supposed to play a role similar to the cooling parameter in simulated annealing, i.e. hopefully it allows to avoid stationary points and local minima.

Appendix B: Algorithm for the l_1 problem

It is now well-known that l_1 norm is the tightest convex relaxation (in the ℓ_2 ball) of the l_0 penalty (Starck et al. 2010). This suggests solving Eq. (2) with $q = 1$. In this case, the problem is well-posed: it has at least a minimizer, and all minimizers are global. Furthermore, although it is completely non-smooth, it can be solved efficiently with a provably convergent algorithm that belongs to the family of proximal splitting schemes (Combettes & Pesquet 2011; Starck et al. 2010).

In particular, we propose to use the Douglas-Rachford (DR) splitting scheme. Let $\beta > 0$, $(\alpha_n)_{n \in \mathbb{N}}$ be a sequence in $]0, 2[$ such that $\sum_{n \in \mathbb{N}} \alpha_n(2 - \alpha_n) = +\infty$. The DR recursion, applied to Eq. (2) with $q = 1$, reads

$$\begin{aligned} \mathbf{a}^{n+\frac{1}{2}} &= \text{Proj}_{\{a: z = \mathbf{M}^* a\}}(\mathbf{a}^n), \\ \mathbf{a}^{n+1} &= \mathbf{a}^n + \alpha_n \left(\text{prox}_{\beta \|\cdot\|_1} \left(2\mathbf{a}^{n+\frac{1}{2}} - \mathbf{a}^n \right) - \mathbf{a}^{n+\frac{1}{2}} \right). \end{aligned} \quad (11)$$

$\text{prox}_{\beta \|\cdot\|_1}$ is the proximity operator of the l_1 -norm, which can be easily shown to be soft-thresholding

$$\text{prox}_{\beta \|\cdot\|_1}(\mathbf{a}) = \Delta_\beta^S(\mathbf{a}), \quad (12)$$

where $\Delta_\beta^S(\mathbf{a}) = (\rho_\beta^S(\mathbf{a}_i))_i$ and $\rho_\beta^S(\mathbf{a}_i) = \text{sign}(\mathbf{a}_i) \max(0, |\mathbf{a}_i| - \beta)$. $\text{Proj}_{\{a: z = \mathbf{M}^* a\}}$ is the orthogonal projector on the corresponding set, and it consists of taking the inverse spherical harmonic transform of its argument, setting its pixel values to the observed ones at the corresponding locations, and taking the forward spherical harmonic transform. It can be shown that the sequence $(\mathbf{a}^{n+\frac{1}{2}})_{n \in \mathbb{N}}$ converges to a global minimizer of Eq. (2) for $q = 1$.

Appendix C: Algorithm for the l_2 problem,

This scheme is based again on Douglas-Rachford splitting applied to solve problem Eq. (3). Its steps are

$$\begin{aligned} \mathbf{x}^{n+\frac{1}{2}} &= \text{Proj}_{\{x: z = \mathbf{M}x\}}(\mathbf{x}^n), \\ \mathbf{x}^{n+1} &= \mathbf{x}^n + \alpha_n \left(\text{prox}_{\beta \|\mathbf{S}(\cdot)\|_{C^{-1}}^2} \left(2\mathbf{x}^{n+\frac{1}{2}} - \mathbf{x}^n \right) - \mathbf{x}^{n+\frac{1}{2}} \right), \end{aligned} \quad (13)$$

where β and α_n are defined as above, and the proximity operator of the squared weighted ℓ_2 -norm is

$$\text{prox}_{\beta \|\mathbf{S}(\cdot)\|_{C^{-1}}^2}(\mathbf{x}) = \mathbf{S}^* \left(\left(\mathbf{S}\mathbf{x} \right) \otimes \left(\frac{C}{\beta + C} \right) \right), \quad (14)$$

where $\otimes$ stands for the entry-wise multiplication between two vectors, and the division is also to be understood entry-wise. The projector $\text{Proj}_{\{x: z = \mathbf{M}x\}}$ has a simple closed-form and consists in setting pixel values to the observed ones at the corresponding

locations, and keeping the others intact. It can be shown that the sequence $(\mathbf{x}^{n+\frac{1}{2}})_{n \in \mathbb{N}}$ converges to the unique global minimizer of Eq. (3).

Appendix D: Algorithm for isotropy inpainting

The isotropy-constrained inpainting problem can be cast as the *non-convex* feasibility problem

$$\text{find } \widehat{\mathbf{x}} \in C_\varepsilon \cap \{\mathbf{x} : \mathbf{z} = \mathbf{M}\mathbf{x}\}. \quad (15)$$

The feasible set is nonempty since the constraint set is non-empty and closed; the CMB is in it under the isotropy hypothesis. To solve Eq. (15), we propose to use the von Neumann method of alternating projections onto the two constraint sets, whose recursion can be written

$$\mathbf{x}^{(n+1)} = \text{Proj}_{\{x: z = \mathbf{M}x\}} \left(\text{Proj}_{C_\varepsilon}(\mathbf{x}^{(n)}) \right). \quad (16)$$

Using closedness and prox-regularity of the constraints, and by arguments from Lewis & Mallick (2008), we can conclude that our non-convex alternating projections algorithm for inpainting converges locally to a point of the intersection $C_\varepsilon \cap \{\mathbf{x} : \mathbf{z} = \mathbf{M}\mathbf{x}\}$ which is non-empty.

References

- Abrial, P., Moudden, Y., Starck, J., et al. 2007, J. Fourier Anal. Appl., 13, 729
- Abrial, P., Moudden, Y., Starck, J., et al. 2008, Stat. Meth., 5, 289
- Ben-David, A., Kovetz, E. D., & Itzhaki, N. 2012, ApJ, 748, 39
- Bobin, J., Starck, J.-L., Sureau, F., & Basak, S. 2012, A&A, in press
DOI: [10.1051/0004-6361/201219781](https://doi.org/10.1051/0004-6361/201219781)
- Bucher, M., & Louis, T. 2012, MNRAS, 3271
- Combettes, P. L., & Pesquet, J.-C. 2011, Fixed-Point Algorithms for Inverse Problems in Science and Engineering (Springer-Verlag), 185
- de Oliveira-Costa, A., & Tegmark, M. 2006, Phys. Rev. D, 74, 023005
- Dupé, F.-X., Rassat, A., Starck, J.-L., & Fadili, M. J. 2011, A&A, 534, A51
- Feeney, S. M., Peiris, H. V., & Pontzen, A. 2011, Phys. Rev. D, 84, 103002
- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, ApJ, 567, 2
- Kim, J., Naselsky, P., & Mandolesi, N. 2012, ApJ, 750, L9
- Land, K., & Magueijo, J. 2005, Phys. Rev. Lett., 72, 101302
- Lewis, A., & Mallick, J. 2008, Math. Oper. Res., 33, 216
- Paykari, P., Starck, J.-L., & Fadili, M. J. 2012, A&A, 541, A74
- Perotto, L., Bobin, J., Plaszczyński, S., Starck, J., & Lavabre, A. 2010, A&A, 519, A4
- Pires, S., Starck, J., Amara, A., et al. 2009, MNRAS, 395, 1265
- Pires, S., Starck, J., & Refregier, A. 2010, IEEE Sign. Proc. Mag., 27, 76
- Plaszczyński, S., Lavabre, A., Perotto, L., & Starck, J.-L. 2012, A&A, 544, A27
- Rassat, A., Starck, J.-L., & Dupe, F. X. 2012, A&A, submitted
- Sato, K. H., Garcia, R. A., Pires, S., et al. 2010, Astron. Nachr., submitted
[arXiv:1003.5178]
- Schmitt, J., Starck, J. L., Casandjian, J. M., Fadili, J., & Grenier, I. 2010, A&A, 517, A26
- Starck, J.-L., Murtagh, F., & Fadili, M. 2010, Sparse Image and Signal Processing (Cambridge University Press)
- Zacchei, A., Maino, D., Baccigalupi, C., et al. 2011, A&A, 536, A5