

Compressed sensing for wide-field radio interferometric imaging

J. D. McEwen^{1*} and Y. Wiaux^{1,2,3}

¹*Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

²*Institute of Bioengineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*

³*Department of Radiology and Medical Informatics, University of Geneva (UniGE), CH-1211 Geneva, Switzerland*

Accepted 2010 December 16. Received 2010 November 25; in original form 2010 October 18

ABSTRACT

For the next generation of radio interferometric telescopes it is of paramount importance to incorporate wide field-of-view (WFOV) considerations in interferometric imaging, otherwise the fidelity of reconstructed images will suffer greatly. We extend compressed sensing techniques for interferometric imaging to a WFOV and recover images in the spherical coordinate space in which they naturally live, eliminating any distorting projection. The effectiveness of the spread spectrum phenomenon, highlighted recently by one of the authors, is enhanced when going to a WFOV, while sparsity is promoted by recovering images directly on the sphere. Both of these properties act to improve the quality of reconstructed interferometric images. We quantify the performance of compressed sensing reconstruction techniques through simulations, highlighting the superior reconstruction quality achieved by recovering interferometric images directly on the sphere rather than the plane.

Key words: techniques: interferometric – methods: numerical – cosmology: observations.

1 INTRODUCTION

Incorporating wide field-of-view (WFOV) contributions in the images reconstructed from radio interferometric observations is becoming increasingly important. Next-generation radio interferometers, such as the Square Kilometre Array¹ (SKA) (Carilli & Rawlings 2004), will inherently observe very large fields of view about the pointing direction of the telescope. Wide fields introduce two important distinctions to standard interferometric imaging: firstly, interferometric images are inherently spherical and planar projections necessarily introduce distortions; and, secondly, non-zero baseline components in the pointing direction of the telescope must be taken into account. If these contributions are ignored, the fidelity of reconstructed images will suffer greatly.

WFOV contributions have been considered by McEwen & Scaife (2008) in simulating the visibilities observed by an interferometer. Full-sky interferometric formalisms were derived using a number of different signal representations, including representations in real, spherical harmonic and wavelet spaces. Real and spherical harmonic space representations were shown to be numerically infeasible for simulating realistic observations, while a fast wavelet space method was developed, reducing the computational cost considerably and rendering realistic simulations feasible. However, the forward problem only was considered in this work. Very recently Carozzi & Woan (2009) developed a generalised radio interferometric measurement equation that is valid

for partially polarised sources over a WFOV, however this framework has not yet been applied in practice. The inverse WFOV imaging problem has traditionally been tackled by faceting the sky into a number of regions which are sufficiently small that standard Fourier imaging is possible (Cornwell & Perley 1992). More recently, the w -projection algorithm has been developed by Cornwell et al. (2008) to incorporate WFOV contributions by applying the modulating so-called w -term that appears in the interferometric integral as a convolution in the Fourier plane. This provides an order of magnitude speed enhancement over traditional facet-based approaches. Nevertheless, both of these methods recover *planar* images in the space of directional cosines; this necessarily distorts the image. None of the current methods recover images in the spherical coordinate space in which they live. In this article we recover *spherical* interferometric images parameterised by colatitude and longitude on the celestial sphere. By recovering images defined directly on the sphere, and thereby eliminating the distortion due to a projection to the plane, the performance of reconstruction is enhanced. We develop reconstruction algorithms in the context of the theory of compressed sensing.

Compressed sensing (Candès et al. 2006a,b; Candès 2006; Donoho 2006; Baraniuk 2007) is a recent development in the field of information theory, which goes beyond the usual Nyquist-Shannon sampling theorem. It relies on the fact that many signals in Nature are *sparse* (or approximately so) and may be represented in a basis requiring many fewer non-zero coefficients than the dimensionality of the signal itself. Compressed sensing theory shows that a sparse signal may be recovered from many fewer measurements than Nyquist-Shannon sampling would suggest and

* E-mail: jason.mcewen@epfl.ch

¹ <http://www.skatelescope.org/>

thus aims to merge data acquisition and compression. These results also hold for signals that are approximately sparse only (*i.e.* signals which contain many coefficients of small but non-zero value), so-called *compressible* signals.

The first application of compressed sensing to radio interferometry was performed by Wiaux et al. (2009a), where the problem of image reconstruction from incomplete visibility measurements was considered. Wiaux et al. (2009a) demonstrated the versatility of the approach, through the ability to incorporate additional signal priors easily, and its superiority relative to standard interferometric imaging techniques, such as CLEAN (Högbom 1974). The spread spectrum phenomenon, which arises by partially relaxing the small field-of-view (FOV) assumption and including a first order w -term, was introduced by Wiaux et al. (2009b), enhancing the performance of image reconstruction for sparsity bases that are not maximally incoherent with the measurement basis (see Sec. 2.2 for a review of sparsity and measurement bases and their coherence). Furthermore, a compressed sensing based approach was developed and evaluated by Wiaux et al. (2010) to recover the signal induced by cosmic strings in the cosmic microwave background, exploiting the sparse nature of line-like discontinuities due to the string signal. All of these works consider uniformly random and discrete visibility coverage in order to remain as close to the theory of compressed sensing as possible. First steps towards more realistic visibility coverages have been taken by Suksmono (2009) and Wenger et al. (2010), who consider coverages due to specific interferometer configurations but which remain discrete. These preliminary works suggest that the performance of compressed sensing reconstructions is unlikely to deteriorate significantly for more realistic visibility coverages.

In this article we generalise the compressed sensing imaging techniques developed by Wiaux et al. (2009a) and Wiaux et al. (2009b) to a WFOV. To do so we recover interferometric images defined directly on the sphere, rather than a tangent plane. Recovering images in the space in which the signal naturally lives eliminates any distortion due to projection. Furthermore, projection effects also act to hamper the sparsity of the signal (*i.e.* act to make it less sparse), impeding the performance of compressed sensing based reconstruction. To remain close to the theoretical compressed sensing framework we follow Wiaux et al. (2009a) and Wiaux et al. (2009b) by considering uniformly random and discrete planar visibility coverage. We also assume non-zero but constant baseline components in the pointing direction of the telescope (*i.e.* non-zero but constant w , where w is defined explicitly in Sec. 2.1), in order to study the enhancement of reconstruction quality due to the spread spectrum phenomenon. This assumption allows us to discard considerations related to specific interferometer configurations and to study the impact of the spread spectrum phenomenon at light computational load. As we shall see, the effectiveness of the spread spectrum phenomenon is improved in the WFOV setting considered. Due to all of these considerations, the quality of reconstruction on the sphere is enhanced considerably relative to planar reconstructions.

The remainder of this article is structured as follows. In Sec. 2 we review radio interferometric imaging in the context of compressed sensing. In Sec. 3 we generalise these techniques to a WFOV. Simulations are presented in Sec. 4 to evaluate the performance of the compressed sensing based reconstruction of spherical images compared to planar images. Concluding remarks are made in Sec. 5.

2 PLANAR INTERFEROMETRIC IMAGING

In this section we review radio interferometric imaging in the context of compressed sensing. Radio interferometry is reviewed in the WFOV setting, before small FOV assumptions are discussed explicitly. A brief review of compressed sensing is given, highlighting various reconstruction techniques and the importance of sparsity and coherence on reconstruction performance. Finally, these frameworks are merged and compressed sensing techniques for radio interferometric imaging are discussed.

2.1 Radio interferometry

In order to image a region on the sky, all radio telescopes of an interferometric array are oriented towards the same pointing direction \hat{s}_0 on the unit celestial sphere S^2 . The FOV observed by the interferometer is limited by the primary beam of the telescope $A(\tau)$, where the beam is defined relative to the pointing direction, *i.e.* $\tau \equiv \hat{s} - \hat{s}_0$ for arbitrary directions \hat{s} . The sky intensity to be imaged $x(\tau)$ is defined in the same coordinates. The complex visibility measured by each telescope pair of the interferometer is given by (Thompson et al. 2001)

$$y(\mathbf{b}) = \int_{S^2} A(\tau) x(\tau) e^{-i2\pi\mathbf{b}\cdot\tau} d\Omega(\tau), \quad (1)$$

where $d\Omega(\hat{s}) = \sin\theta d\theta d\varphi$ is the usual rotation invariant measure on the sphere and (θ, φ) denote the spherical coordinates of \hat{s} , with colatitude $\theta \in [0, \pi]$ and longitude $\varphi \in [0, 2\pi)$. The measured visibility depends only on the relative positions between telescope pairs, denoted in units of wavelengths by the baseline vector $\mathbf{b} = (u, v, w)$. As the Earth rotates relative to the celestial sphere the interferometer tracks the source position as it traverses the sky, with each position corresponding to a rotation of the baseline. In this manner, visibility measurements are accumulated for various baselines, with each antenna pair generating an elliptical track of baseline values over the course of the observation. Visibility coverage is therefore incomplete.

The beam and sky intensity are typically expressed in the same Cartesian coordinate system as the baseline, which is centred on the Earth and aligned with \hat{s}_0 . Consider the coordinates (l, m, n) of \hat{s} in this system, noting $n \equiv n(\mathbf{l}) = (1 - \|\mathbf{l}\|_2^2)^{1/2}$ with $\mathbf{l} = (l, m)$ and $\|\mathbf{l}\|_2^2 = l^2 + m^2$. The primary beam and sky intensity may then be seen as functions of \mathbf{l} . Following this change of coordinates the visibility integral of (1) reads

$$y(\mathbf{u}, w) = \int_{D^2} A(\mathbf{l}) x(\mathbf{l}) e^{-i2\pi[\mathbf{u}\cdot\mathbf{l} + w(n(\mathbf{l})-1)]} \frac{d^2\mathbf{l}}{n(\mathbf{l})}, \quad (2)$$

where $\mathbf{u} = (u, w)$ and we have noted that $\tau = (\mathbf{l}, n(\mathbf{l}) - 1)$. Through the change of coordinates the invariant measure on the sphere becomes $d\Omega(\hat{s}) = n^{-1}(\mathbf{l}) d^2\mathbf{l}$, where $d^2\mathbf{l} = dl dm$ is the canonical invariant measure on the plane, and the integration is now performed over the unit disk D^2 . Note that (2) remains general and does not rely on any small FOV assumption. However, the directional cosines $\mathbf{l} = (l, m)$ express a projection onto the equatorial plane of the signal, which is inherently defined on the celestial sphere. This projection is trivial in the continuous setting but will become important once we reach the discrete setting required for interferometric imaging.

Often small FOV assumptions are made, in which case it is convenient to represent the so-called w -term

$$C(\|\mathbf{l}\|_2) \equiv e^{i2\pi w(1 - \sqrt{1 - \|\mathbf{l}\|_2^2})} \quad (3)$$

explicitly, from which it follows

$$y(\mathbf{u}, w) = \int_{D^2} A(\mathbf{l}) x(\mathbf{l}) C(\|\mathbf{l}\|_2) e^{-i2\pi\mathbf{u}\cdot\mathbf{l}} \frac{d^2\mathbf{l}}{n(\mathbf{l})}. \quad (4)$$

Two types of approximation regarding small FOVs are used to relate (4) to the Fourier transform of the beam-modulated intensity $A(\mathbf{l})x(\mathbf{l})$. Firstly, the assumption $\|\mathbf{l}\|_2^2 \ll 1$ implies $n(\mathbf{l}) \simeq 1$, so that the Jacobian term $n^{-1}(\mathbf{l})$ is reduced to unity. Alternatively, $n^{-1}(\mathbf{l})$ may be absorbed into the beam. Secondly, various assumptions are made to approximate the w -term. A zeroth order approximation of the w -term is made by assuming $\|\mathbf{l}\|_2^2 w \ll 1$, so that $C(\|\mathbf{l}\|_2) \simeq 1$. When incorporating both approximations, (4) reduces to the Fourier transform and standard Fourier imaging may be used to recover images from measured visibilities. Alternatively, a first order approximation $\|\mathbf{l}\|_2^4 w \ll 1$ reduces the w -term to $C(\|\mathbf{l}\|_2) \simeq e^{i\pi w \|\mathbf{l}\|_2^2} \equiv C_1(\|\mathbf{l}\|_2)$. This assumption gives rise to the linear chirp modulation responsible for the spread spectrum phenomenon (Wiaux et al. 2009b). Although not considered here, direction dependent beam effects may also introduce a phase modulation, which would provide an alternative source of the spread spectrum phenomenon. Since we intend to consider WFOVs, no small FOV assumptions will be made herein; we include the Jacobian term $n^{-1}(\mathbf{l})$ and consider the full w -term given by (3).

2.2 Compressed sensing

Compressed sensing (Candès et al. 2006a,b; Candès 2006; Donoho 2006; Baraniuk 2007) is concerned with the recovery of sparse or compressible signals from a small number of measurements. The signal to be recovered is defined by its Nyquist-Shannon sampling, denoted by the vector $\mathbf{x} \in \mathbb{C}^N$, and is assumed to be sparse in some orthogonal basis $\{\boldsymbol{\psi}_i\}_{1 \leq i \leq N}$, where $\boldsymbol{\psi}_i \in \mathbb{C}^N, \forall i$. The signal may then be represented by its coefficients $\boldsymbol{\alpha} \in \mathbb{C}^N$ in this basis: $\mathbf{x} = \Psi\boldsymbol{\alpha}$, where $\Psi \in \mathbb{C}^{N \times N}$ is the $N \times N$ matrix with columns $\boldsymbol{\psi}_i$. Formally, \mathbf{x} is said to be sparse or compressible if $\boldsymbol{\alpha}$ contains $K \ll N$ non-zero or significant elements respectively. Hereafter, we refer to enhancing (hampering) sparsity as synonymous with decreasing (increasing) the sparsity value K . Measurement of \mathbf{x} is assumed to be made by the projection onto measurement basis vectors $\{\boldsymbol{\phi}_r\}_{1 \leq r \leq M}$, for M measurements, belonging to an orthogonal sensing basis where $\boldsymbol{\phi}_r \in \mathbb{C}^N, \forall r$. This is a very flexible framework which allows a wide range of acquisition procedures to be modelled. The vector of measurements $\mathbf{y} \in \mathbb{C}^M$ may then be expressed by

$$\mathbf{y} = \Phi\mathbf{x} + \mathbf{n} = \Phi\Psi\boldsymbol{\alpha} + \mathbf{n}, \quad (5)$$

where $\Phi \in \mathbb{C}^{M \times N}$ is the $M \times N$ matrix with rows $\boldsymbol{\phi}_r$ and $\mathbf{n} \in \mathbb{C}^M$ represents noise. Compressed sensing suggests that \mathbf{x} can be recovered with a number of measurements $M \sim K \ll N$. However, recovering \mathbf{x} in this setting involves solving the inverse problem (5), which becomes ill-posed for $M < N$.

Compressed sensing techniques are generally based on global minimisation problems, which are solved by greedy algorithms or convex non-linear optimisation algorithms. The ill-posed inverse problem described above can be defined by a constrained optimisation problem explicitly regularised by a sparsity or compressibility prior. This results in the Basis Pursuit denoising (BP) problem², which consists of minimising the ℓ_1 -norm of the coefficients of the

signal in the sparsity basis $\|\boldsymbol{\alpha}\|_1$ under a constraint on the ℓ_2 -norm of the residual noise $\|\mathbf{y} - \Phi\Psi\boldsymbol{\alpha}\|_2$:

$$\min_{\boldsymbol{\alpha}} \|\boldsymbol{\alpha}\|_1 \quad \text{such that } \|\mathbf{y} - \Phi\Psi\boldsymbol{\alpha}\|_2 \leq \epsilon. \quad (6)$$

Recall that the ℓ_1 -norm is simply given by the sum of the absolute values of the elements of a vector $\|\boldsymbol{\alpha}\|_1 = \sum_{i=1}^N |\alpha_i|$, whereas the ℓ_2 -norm is the standard norm defined previously. The constraint ϵ may be related to a residual noise level estimator. Assuming independent identically distributed Gaussian noise, a residual noise level estimator is given by twice the negative logarithm of the likelihood associated with the candidate reconstruction, which follows a χ^2 -distribution with $2M$ degrees of freedom. The measurement constraint ϵ may then be chosen to correspond to some (100p)th percentile of the χ^2 -distribution, for $0 \leq p \leq 1$ (Wiaux et al. 2009a). The BP problem may be solved by the application of non-linear, iterative convex optimisation algorithms (e.g. Combettes & Pesquet 2007). If the solution of the optimisation problem is denoted $\boldsymbol{\alpha}^*$, then the signal is reconstructed through the synthesis $\mathbf{x}^{*\text{BP}} = \Psi\boldsymbol{\alpha}^*$.

The BP denoising problem is appropriate for signals sparse or compressible in a given basis. However, many signals in Nature are also sparse or compressible in the magnitude of their gradient, in which case the Total Variation (TV) minimisation problem applies (Candès et al. 2006a). The TV problem involves replacing the ℓ_1 -norm sparsity prior in the BP problem with a prior on the TV norm of the signal itself:

$$\min_{\mathbf{x}} \|\mathbf{x}\|_{\text{TV}} \quad \text{such that } \|\mathbf{y} - \Phi\mathbf{x}\|_2 \leq \epsilon, \quad (7)$$

where the TV norm is defined by the ℓ_1 -norm of the gradient of the signal $\|\mathbf{x}\|_{\text{TV}} = \|\nabla\mathbf{x}\|_1$ (Rudin et al. 1992; Chambolle 2004). The TV problem may also be solved by the application of non-linear, iterative convex optimisation algorithms (e.g. Chambolle 2004; Durand et al. 2010). The signal is directly recovered from the solution to the optimisation problem, denoted $\mathbf{x}^{*\text{TV}}$.

Finally, we review the two fundamental criteria that drive the performance of compressed sensing reconstructions. The first has already been central to our discussion of compressed sensing: sparsity. The more sparse a signal the fewer measurements required to recover it, or similarly, the better the reconstruction quality for a given number of measurements. In addition, the matrix $\Theta = \Phi\Psi$ must satisfy the restricted isometry property (RIP) (Candès et al. 2006a,b; Candès 2006) to ensure accurate recovery. It has been shown that this property can be satisfied by acquiring random measurements (i.e. random visibility coverage in the terminology of interferometry), if the measurement and sparsity bases are *incoherent*. Coherence is therefore the second criterion driving reconstruction performance; as the coherence between the two bases increases, the reconstruction performance degrades. Incoherence ensures that the rows of Φ cannot sparsely represent the columns of Ψ , ensuring that signal content is sufficiently probed by random measurements. Note that coherence is only defined strictly in the presence of a sparsity basis, hence it cannot be studied for the TV problem (nevertheless, approximate coherence arguments may still be made in this setting to gain intuition). The *mutual coherence* between the measurement and sparsity bases is defined by the maximum absolute inner product of all combinations of their normalised basis vectors:

$$\mu(\Phi, \Psi) = \max_{r,i} |\langle \boldsymbol{\phi}_r | \boldsymbol{\psi}_i \rangle|,$$

with the incoherence given by the reciprocal $\mu^{-1}(\Phi, \Psi)$. Note that the Dirac and Fourier bases are maximally incoherent. The number of measurements required to satisfy the RIP and ensure accurate

² In the absence of noise the problem is simply called Basis Pursuit.

recovery, satisfies the bound

$$M \geq c K \mu^2(\Phi, \Psi) N \log^4 N,$$

for some constant c (note that the bound is not tight). Although useful for theoretical and intuitive considerations, the mutual coherence is a blunt numerical instrument as it captures only the most extreme inner product between the measurement and sparsity bases (Tropp 2004). An alternative measure of coherence is given by the *cumulative coherence* (Romberg 2009)

$$\nu(\Phi, \Psi, \Gamma) = \max_r \left(\sum_{i \in \Gamma} |\langle \phi_r | \psi_i \rangle|^2 \right)^{1/2},$$

where Γ represents the set of all sparsity basis vectors that contribute to the signal of interest. The cumulative coherence is therefore signal dependent and incorporates sparsity information, which renders it more robust than mutual coherence to departures from the pure compressed sensing framework. The number of measurements required for signal recovery may also be expressed in terms of cumulative coherence and can be shown to evolve as (Romberg 2009; Rauhut 2010)

$$M \sim \nu^2(\Phi, \Psi, \Gamma) N. \quad (8)$$

2.3 Interferometric imaging

We express the interferometric framework discussed in Sec. 2.1 in a discrete setting which is amenable to the compressed sensing reconstruction techniques discussed in Sec. 2.2, following the approach taken by Wiaux et al. (2009a) and Wiaux et al. (2009b) previously. Although we consider WFOV considerations we restrict ourselves to baselines with constant w , thus restricting visibilities to a Fourier plane.³ In reality, w will vary and the impact of the spread spectrum phenomenon will lie somewhere between the extreme cases that we consider here. Nevertheless, this assumption allows us to probe the expected impact of the spread spectrum phenomenon at light computational load and to discard considerations related to specific interferometer configurations. Considering the signal projected onto the equatorial plane, we consider the $\mathbf{l} = (l, m)$ coordinate system discretised on a uniform grid of $N = N^{1/2} \times N^{1/2}$ points $\mathbf{l}_i \in \mathbb{R}^2$ in real space, with integer $1 \leq i \leq N$, and the corresponding grid of discrete spatial frequencies \mathbf{u}_i defined by Nyquist-Shannon sampling. The band-limited intensity and primary beam functions are defined on the spatial grid and denoted by \mathbf{x} , $\mathbf{A} \in \mathbb{C}^N$, respectively. The complex w -term is also defined on this grid $\mathbf{C} \in \mathbb{C}^N$. Since the w -term is complex, the Fourier transform of $C(\|\mathbf{l}\|_2)A(\mathbf{l})x(\mathbf{l})$ does not bear any specific symmetry in the Fourier plane, even for a real beam and real source intensity. Following Wiaux et al. (2009a) and Wiaux et al. (2009b), we assume that the spatial frequencies \mathbf{u} probed by the interferometer fall on the discrete grid points \mathbf{u}_i . The Fourier coverage provided by the M spatial frequencies probed \mathbf{u}_r , with integer $1 \leq r \leq M$, can be identified by a binary mask in the Fourier plane equal to unity for each spatial frequency probed and zero otherwise.⁴ The visibilities

measured are denoted by the vector $\mathbf{y} \in \mathbb{C}^M$, which may be affected by complex noise $\mathbf{n} \in \mathbb{C}^M$.

The visibility integral (4) may be represented in this discrete setting by the linear system

$$\mathbf{y} = \Phi_p \mathbf{x}_p + \mathbf{n}, \quad (9)$$

with

$$\Phi_p = \mathbf{W} \mathbf{M} \mathbf{F} \mathbf{C} \mathbf{A},$$

where Gaussian noise and a whitening operator have also been included. The measurement matrix $\Phi_p \in \mathbb{C}^{M \times N}$ identifies the complete linear relation between the signal and the visibilities. The matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$ is the diagonal matrix implementing the primary beam, with the Jacobian term $n^{-1}(\mathbf{l})$ due to the change to Cartesian coordinates absorbed. The matrix $\mathbf{C} \in \mathbb{C}^{N \times N}$ is the diagonal matrix implementing the modulation by the w -term. The unitary matrix $\mathbf{F} \in \mathbb{C}^{N \times N}$ implements the discrete Fourier transform. The matrix $\mathbf{M} \in \mathbb{R}^{M \times N}$ is the rectangular binary matrix implementing the mask that characterises visibility coverage, containing one non-zero value on each row only, at the index of the Fourier coefficient corresponding to the spatial frequency probed by each interferometric measurement. We also augment the measurement matrix by a whitening operator $\mathbf{W} \in \mathbb{R}^{M \times M}$ (Wiaux et al. 2009a). Whitening corresponds to dividing each measured visibility by the standard deviation of the corresponding noise component so that the final observed visibilities are then affected by independent identically distributed Gaussian noise with unit variance. The subscript p is used to denote the planar setting since we subsequently generalise this framework to functions defined on the sphere (throughout subscripts p and s are used to denote the plane and sphere respectively). For incomplete visibility coverage $M < N$, (9) defines an ill-posed inverse problem, which we solve using the BP and TV reconstruction approaches described in Sec. 2.2. Finally, let us note that the compressed sensing framework is defined strictly for orthogonal sparsity and sensing bases. The application of the w -term modulation necessitates an upsampling operation in practice to ensure that the modulated signal is Nyquist-Shannon sampled, which breaks the orthogonality of the sensing basis. Compressed sensing techniques for interferometric imaging therefore depart from the theoretical compressed sensing framework but nevertheless have been demonstrated to perform very well.

We conclude this section with a review of the spread spectrum phenomenon. As discussed in Sec. 2.2, the coherence between the measurement and sparsity bases is a key criterion effecting the performance of compressed sensing based reconstructions. In the interferometric setting described above the measurement basis can essentially be identified with the Fourier basis, which will aid our intuitive understanding of the spread spectrum phenomenon. In this case, the mutual coherence is given by the maximum modulus of the Fourier coefficient of the sparsity basis vectors. Consequently, an operation that acts to reduce the maximum Fourier coefficient, reduces the coherence and thus improves the quality of compressed sensing reconstructions. Modulation by the w -term corresponds to a norm-preserving convolution in the Fourier plane, spreading the spectrum of the sparsity basis vectors, and achieves exactly this. Hence, the greater the frequency content of the w -term, the more effective the spread spectrum phenomenon.

³ The w -projection algorithm (Cornwell et al. 2008) could be applied in future to remove this restriction, although the extension of the spread spectrum phenomenon to a varying w would need to be studied extensively.

⁴ Conventional gridding (Thompson et al. 2001) of the visibilities measured at continuous spatial frequencies would result in a mask with non-binary weights, which could be incorporated easily in the framework described here.

3 SPHERICAL INTERFEROMETRIC IMAGING

In this section we generalise the compressed sensing imaging techniques developed by Wiaux et al. (2009a) and Wiaux et al. (2009b) to a WFOV. We recover interferometric images defined directly on the sphere, rather than the equatorial plane, which enhances the effectiveness of compressed sensing reconstructions.

The measurement operator transforming the sky intensity defined on the sphere to visibilities, consists of augmenting the usual interferometric measurement operator with an initial projection from the sphere to the plane. This initial projection corresponds to a change from spherical to Cartesian coordinates in much the same way (4) is defined from (1). Similarly, the framework remains general and does not rely on any small FOV assumptions. Practically, however, the projection which implements the change of variable is complicated by the discrete setting and the desire to recover a regular grid on the plane to allow the use of fast Fourier transforms (FFTs).⁵ In order to ensure information is not lost, we define the resolution of the planar grid so as to support the maximum projected frequency content of a band-limited signal on the sphere. Although our WFOV framework still involves a projection to the plane, this is included in the measurement operator, and so will be regularised when solving the interferometric inverse problem. Moreover, we recover the sky intensity directly on the sphere, where it is most sparse, and it is the signal space that is important for the impact of sparsity on the performance of compressed sensing reconstructions.

In the remainder of this section we first discuss considerations relating to band-limited signals on the sphere and the plane, operators used to project the sphere to the plane in a discrete setting and discrete gradient operators defined on the sphere (required for TV reconstructions on the sphere). Finally, we define interferometric imaging in the WFOV setting, while commenting on the impact of sparsity, coherence and the spread spectrum phenomenon.

3.1 Band-limited signals

We require a regular grid on the plane to enable the application of FFTs to reduce considerably the computational load of reconstructing images from visibility measurements. To ensure information is not lost when projecting a signal from the sphere to the plane in this discrete setting, we define a planar grid that supports a band-limit corresponding to the maximum projected frequency content of a band-limited signal defined on the sphere. The maximum projected frequency content arises at the extents of the FOV, where signal content defined on the surface of the sphere is projected onto much higher frequency content in the $\mathbf{l} = (l, m)$ plane. In the typical small FOV setting the relationship between the band-limit of a signal on the sphere ℓ_{\max} and its tangent plane u_{\max} is given by $\ell_{\max} = 2\pi u_{\max}$. In the WFOV setting simple geometric considerations at the extent of the FOV lead to the relationship

$$\ell_{\max} = 2\pi \cos(\theta_{\text{FOV}}/2) u_{\max}, \quad (10)$$

where θ_{FOV} denotes the angular opening of the FOV, corresponding to a planar FOV of $L = 2 \sin(\theta_{\text{FOV}}/2)$. The band-limit relation is now dependent on the size of the FOV. Note that for a WFOV a

higher band-limit on the plane is required to support a given band-limit on the sphere, as expected intuitively, and as $\theta_{\text{FOV}} \rightarrow 0$ the usual relationship for a small FOV is recovered.

Once the band-limit of a signal is defined, on both the sphere and the plane, sampling considerations then dictate the resolution of the discrete grid required to accurately represent the signal. Consequently, once the band-limit of the signal on the sphere and the FOV is set, the required band-limit on the plane is determined through (10), and the sampling resolutions of both the sphere and the plane follow. The Nyquist-Shannon sampling theorem on the plane states that $N_p^{1/2} = 2u_{\max}L$, where N_p is the number of samples on the planar grid. The relationship between the number of samples within the FOV on the sphere N_s and the harmonic band-limit ℓ_{\max} , depends on the pixelisation of the sphere adopted. We choose the HEALPix⁶ pixelisation of the sphere (Górski et al. 2005) due to the equal area of each pixel element and its ubiquitous use in the astrophysical community. The resolution of the HEALPix pixelisation is controlled by the parameter N_{side} , where the number of pixels at a given resolution is specified by $N_{\text{pix}} = 12N_{\text{side}}^2$. Functions that are band-limited on the sphere at ℓ_{\max} can be resolved on a HEALPix pixelisation at a resolution corresponding to $\ell_{\max} = 3N_{\text{side}}$, albeit with integration errors⁷ (Hivon et al. 2010) (since no exact quadrature rule exists for HEALPix). For the correspondence $\ell_{\max} = 2N_{\text{side}}$ integration errors are minimal (Hivon et al. 2010). Consequently, harmonic transforms are typically performed for $\ell_{\max} = kN_{\text{side}}$, with $k \sim 2-3$. In order to ensure that we do not favour either the plane or sphere in our analysis, we choose k such that in the limit of a small FOV, the number of samples on the plane and within the FOV on the sphere are the same. It can be shown trivially that

$$\frac{N_p}{N_s} = \frac{2k^2 \tan^2(\theta_{\text{FOV}}/2)}{3\pi^2(1 - \cos(\theta_{\text{FOV}}/2))},$$

where ℓ_{\max} drops out of the expression. Consequently, we choose the value $k = \sqrt{3}\pi/2$ to ensure

$$\lim_{\theta_{\text{FOV}} \rightarrow 0} N_p/N_s = 1.$$

The ratio N_p/N_s is plotted in Fig. 1 for various values of k . Notice how the ratio increases rapidly with the FOV. Naively, one might expect the ratio of sparsities between the plane and sphere K_p/K_s to evolve in a similar manner, which would highlight the superiority of spherical reconstructions compared to planar ones (this is necessarily a very approximate prediction and the exact ratio of sparsities will depend highly on the signal examined).

3.2 Projection operators

Now that discrete grids are defined on the sphere and the plane, it is necessary to determine how to project a signal between these grids. In the continuous setting this is trivial and simply corresponds to a change of variable, however the matter is complicated in practice due to the discrete nature of the sampled signals.

In order to project onto a regular grid on the plane, it is necessary to re-grid the pixelisation on the sphere to recover sample values at spherical positions that project directly onto the planar grid, as illustrated in Fig. 2. We perform a convolution on the sphere to achieve this re-gridding. This convolutional re-gridding on the sphere is similar to the re-gridding often performed when mapping the visibilities observed at continuous coordinates to a regular

⁵ Fast algorithms have been developed to compute a discrete Fourier transform on non-equispaced spatial frequencies (Potts et al. 2001), which could in principle be used to avoid explicit gridding.

⁶ <http://healpix.jpl.nasa.gov/>

⁷ Integration errors may be reduced by running the analysis iteratively.




Figure 1. Ratio of number of samples on the plane to the sphere (N_p/N_s) to sufficiently sample a band-limited signal defined on the sphere when projected onto the plane (θ_{FOV} is specified in degrees). Curves are plotted for a HEALPix pixelisation of the sphere, with $\ell_{\text{max}} = kN_{\text{side}}$, for values $k = 3$ (blue/dashed line); $k = \sqrt{3}\pi/2$ (black/solid line); $k = 2$ (red/dot-dashed line). To ensure a ratio of unity is obtained as θ_{FOV} tends to zero, we select $k = \sqrt{3}\pi/2$ throughout this work.

grid, also to afford the use of FFTs (Thompson et al. 2001). The interferometric measurement operator on the sphere is therefore augmented by prepending a projection operator \mathbf{P} that includes a spherical convolution followed by a project from the sphere to the tangent plane $\mathbf{l} = (l, m)$.

We consider three convolution kernels on the sphere: a box kernel, corresponding to a nearest-neighbour interpolation; a sinc-like kernel; and a Gaussian kernel. The box kernel is well localised on the sphere but has infinite support in harmonic space. The sinc-like kernel corresponds to evaluating the spherical harmonic transform of the sampled signal on the sphere, followed by evaluating the signal value at the new sample position from its spherical harmonic coefficients (we refer to this as a sinc-like kernel due to analogy with the plane). This procedure results in a kernel that is well localised in harmonic space but has support in real space that extends over the entire sphere. These two kernels represent the extremes between spatial and harmonic space localisation and each produces ringing in the projected signal in the domain in which the kernel is not well localised. We therefore seek a kernel that provides a compromise between this trade-off and is reasonably well localised in both space and frequency. In this article we settle on a simple Gaussian kernel, however alternative kernels such as the spherical equivalent of the Gaussian-sinc or spheroidal functions could also be considered (Thompson et al. 2001).

3.3 Gradient operators on the sphere

A discrete gradient operator must be defined on the sphere in order to compute the magnitude of the gradient of a function to solve the TV minimisation problem of (7) on the sphere. The discrete gradient operator on the plane is defined simply through finite differences (Rudin et al. 1992; Chambolle 2004). However, it is not possible to define a discrete spherical gradient operator on the HEALPix pixelisation in this manner since sample positions do not lie on both rings of constant latitude and rings of constant longitude (only equiangular pixelisations of the sphere satisfy this property). One alternative is to consider the continuous gradient operator on the sphere, for which the magnitude of the gradient of a function




Figure 2. Projection of a sampled signal from the sphere to the plane. In order to project onto a regular grid on the plane (to reduce significantly the computational load of subsequent analyses through the use of FFTs), a re-gridding operation is required. The value of the point on the plane (blue/dark dot) is determined by convolving the points on the sphere (red/grey dots) with a suitable kernel.

on the sphere f reads

$$|\nabla_s f| = \sqrt{\left(\frac{\partial f}{\partial \theta}\right)^2 + \frac{1}{\sin^2 \theta} \left(\frac{\partial f}{\partial \varphi}\right)^2},$$

which may be computed in harmonic space (Wandelt et al. 1998) to eliminate pixelisation concerns. However, such an approach requires a spherical harmonic transform, which necessarily requires global support on the sphere and hence is non-local in nature, and also is of high computational cost.

We define a discrete gradient operator on the sphere by analogue with the continuous gradient but computed through finite differences. Convolutional re-gridding is performed to obtain samples on the sphere at the positions required to compute finite difference values. The same Gaussian convolution kernel that is used to project the sphere to the plane (as discussed in Sec. 3.2) is used. Such an approach actually corresponds to computing the gradient of a smoothed version of the original signal. However, the smoothing is minimal in practice and numerical experiments have shown that the discrete gradient defined in this manner is a good approximation to the continuous gradient computed in harmonic space but does not suffer from the global nature and high computational cost of the spherical harmonic transform.

3.4 Interferometric imaging

We are now in a position to define explicitly the WFOV interferometric imaging framework developed in this work. We attempt to recover the sky intensity directly on the sphere $\mathbf{x}_s \in \mathbb{C}^{N_s}$ by solving the inverse problem

$$\mathbf{y} = \Phi_s \mathbf{x}_s + \mathbf{n}, \quad (11)$$

with

$$\Phi_s = \mathbf{W} \mathbf{M} \mathbf{F} \mathbf{C} \mathbf{A} \mathbf{P}.$$

The measurement operator on the sphere $\Phi_s \in \mathbb{C}^{M \times N_s}$ simply consists of augmenting the operator on the plane Φ_p by prepending a

projection from the sphere to the plane \mathbf{P} , which incorporates a convolutional re-gridding as discussed in Sec. 3.2. Careful consideration is given to samplings on the sphere and plane to ensure that the planar grid is sampled sufficiently to accurately represent the projection of a band-limited signal defined on the sphere (as discussed in Sec. 3.1). For incomplete visibility coverage the inverse problem defined by (11) becomes ill-posed. We solve this ill-posed inverse problem in a compressed sensing framework by solving the BP and TV minimisation problems defined by (6) and (7) respectively. In this work we consider the Dirac sparsity basis on the sphere for BP reconstructions.

We expect the performance of compressed sensing reconstructions to be enhanced by recovering the sky intensity in the space where it naturally lives (*i.e.* on the sphere), since we expect sparsity to be reduced in this space. Furthermore, the effectiveness of the spread spectrum phenomenon is enhanced by going to a WFOV. As discussed in Sec. 2.3, the greater the frequency content of the w -term, the more effective the spread spectrum phenomenon. By examining the w -term plotted in Fig. 3, we see that its frequency content increases with distance from the origin. Consequently, the larger the FOV, the higher the frequency content achieved by the w -term and the more effective the spread spectrum phenomenon. Moreover, a secondary enhancement arises by eliminating the first order assumption $\|I\|_2^4 w \ll 1$ made previously by Wiaux et al. (2009b), since this assumption reduces the maximum frequency content of the w -term achieved on a given FOV (the band-limits of the various w -terms are defined explicitly in Sec. 4.1). Finally, we comment on the impact of mutual coherence on reconstruction performance. Since we consider the Dirac sparsity basis, coherence is optimal on the plane (recall that the Dirac and Fourier bases are maximally incoherent). However, the projection of Dirac basis functions defined on the sphere to the plane does not result in Dirac functions on the plane (due to the convolutional re-gridding). Consequently, coherence in the spherical setting is suboptimal. The spread spectrum phenomenon acts to increase incoherence only in the case where it is not already maximally incoherence, thus we expect the spread spectrum phenomenon to be ineffective in the planar setting but to improve reconstruction performance in the spherical setting. For TV reconstructions, although a sparsity basis does not exist, one may gain some intuition regarding the impact of coherence by the following argument. For a piecewise constant signal that is sparse in the magnitude of its gradient (*i.e.* has a gradient defined by Dirac functions), the spectrum of the magnitude of its gradient must be flat. The gradient operator in space essentially corresponds to a multiplication by frequency in Fourier space, hence the original spectrum of the piecewise constant signal must evolve as the inverse of frequency. Since this differs to the optimally incoherent spectrum which is flat, the coherence must be suboptimal for TV reconstructions. We therefore expect the spread spectrum phenomenon to provide improvements both on the sphere and plane for TV reconstructions, with a greater enhancement expected on the sphere due to the spatial spreading of the projection operator. In any case, sparsity typically has a greater impact on the performance of compressed sensing reconstructions than coherence. These considerations pertaining to sparsity, mutual coherence and the spread spectrum phenomenon lead us to expect an improvement in the performance of compressed sensing reconstructions when recovering interferometric images in the WFOV framework developed here.

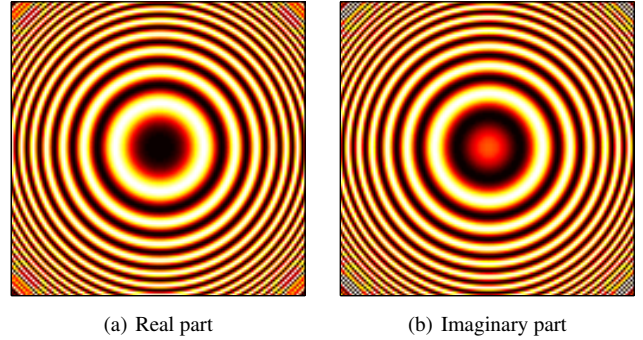


Figure 3. Real and imaginary parts of the w -term modulation $C(\|I\|_2)$ (for $w_d = 1/\sqrt{2}$; see text Sec. 4.1). Notice that the frequency content of the w -term modulation increases with distance from the origin (image centre). Consequently, for a WFOV the w -term modulation spreads the spectrum more effectively, enhancing the performance of the spread spectrum phenomenon. Dark and light regions correspond to positive and negative values respectively.

4 SIMULATED RECONSTRUCTIONS

We evaluate the performance of the WFOV interferometric imaging framework defined on the sphere, as outlined in Sec. 3, making a direct comparison with planar reconstructions. After describing the observational set-up, performance is quantified thoroughly in a low-resolution setting on sets of simulations of sources with a Gaussian profile. A more realistic setting at a higher resolution is then considered, where reconstruction performance is evaluated on a single simulated observation of Galactic dust emission.

4.1 Observational set-up

Simulated observations of real signals are made on the FOV defined by the angular opening $\theta_{\text{FOV}} = 90^\circ$, corresponding to a planar FOV of $L = \sqrt{2}$. A real Gaussian primary beam is assumed, with full-width-half-maximum (FWHM) of one half of the field of view. Random visibility coverage is considered, with visibility measurements falling on the discrete planar grid of spatial frequencies u_i . We consider incomplete visibility coverage, with only 2–25% of the discrete visibilities measured (since we consider real signals, measuring 50% of the complex visibilities corresponds to a number of measurements identical to the number of unknowns in the real signal that we attempt to recover).

As discussed previously, although we consider a WFOV we restrict ourselves to the case of constant w . We parameterise the continuous w in terms of the discrete component w_d , following the parameterisation used by Wiaux et al. (2009a) of $w = w_d N_p^{1/2}/L^2$. The band-limit of the modulating w -term $C(\|I\|_2)$ is given approximately by

$$u_{\text{max}}^{C(w)} \approx \frac{w_d N_p^{1/2}}{2L \sqrt{1 - (L/2)^2}},$$

corresponding to its maximum instantaneous frequency. Under the first order small FOV assumption $\|I\|_2^4 w \ll 1$, the w -term reduces to the linear chirp $C_1(\|I\|_2)$ with approximate band-limit

$$u_{\text{max}}^{C_1(w)} \approx \frac{w_d N_p^{1/2}}{2L}.$$

Notice that the band-limit for given non-zero values of w_d , N_p and

L is greater in the absence of the small FOV assumption, justifying rigorously the secondary enhancement discussed in Sec. 3.4 of the spread spectrum phenomenon due to the WFOV. The band-limit of the spread signal is given by the sum of the original band-limit $u_{\max} = N_p^{1/2}/2L$ and the band-limit of the w -term. Previously Wiaux et al. (2009b) considered the linear chirp $C_1(\|I\|_2)$ and the value $w_d = 1$, corresponding to spreading by a factor of two. We also consider spreading by a factor of two, but since we consider the exact w -term $C(\|I\|_2)$, this corresponds to a value of $w_d = 1/\sqrt{2}$. The corresponding continuous w is of the same order as the maximum visibility measurements in u and v , *i.e.* $w = u_{\max}$, hence it is an appropriate value to consider when studying the spread spectrum phenomenon. Note that in the absence of a WFOV, the value of w considered by Wiaux et al. (2009b) to achieve the same spreading was a factor of $2/L$ greater than u_{\max} . Since the band-limit of the original signal is doubled due to modulation by the w -term, to avoid aliasing we apply an upsampling operator to increase the resolution of the planar grid by a factor of two prior to application of the modulation (upsampling is performing by zero-padding in the Fourier domain). In the subsequent analysis we consider the exact w -term $C(\|I\|_2)$, with values of $w_d = \{0, 1/\sqrt{2}\}$ to highlight the effectiveness of the spread spectrum phenomenon.

Instrumental noise is also added to the simulated visibilities. Independent identically distributed Gaussian noise is assumed with variance $\sigma_n^2 = 10^{-3}\sigma_y^2$, where σ_y^2 is the variance of the visibilities in the absence of noise and w -term modulation. The added instrumental noise results in observed visibilities at a signal-to-noise ratio of $\text{SNR}_n = 10 \log_{10}(\sigma_y^2/\sigma_n^2) = 30\text{dB}$.

4.2 Gaussian sources

The WFOV interferometric imaging framework is evaluated thoroughly in this section on simulated observations of Gaussian sources. These simulations are first described, before we analyse their sparsity properties. Reconstruction performance is then evaluated both in the spherical and planar settings. Finally, reconstruction performance is compared to calculations of cumulative coherence.

4.2.1 Simulations

In order to perform a thorough evaluation, we analyse sets of simulations of Gaussian sources of various size, with each set containing 30 simulations, for all variations of reconstruction procedures (BP and TV reconstructions on the plane and the sphere, both with and without application of the spread spectrum phenomenon) and for various visibility coverages. Due to the large number of simulated reconstructions, we restrict these simulations to a low resolution. We consider a HEALPix resolution parameter of $N_{\text{side}} = 32$, corresponding to a harmonic band-limit of $\ell_{\max} = 88$. The remaining parameters defining the resolution of the simulations then follow from the considerations discussed in Sec. 3.1; we find $N_s = 1740$, $u_{\max} = 19.6$ and $N_p = 58 \times 58 = 3364$. The size of the Gaussian kernel in the convolutional re-gridding of the projection operator, defined by its standard deviation σ_p , is chosen to ensure that the kernel is well sampled. Since a kernel with small support is required, a local tangent plane approximation is made to relate its Fourier size to its spatial size through $\sigma_F = (2\pi\sigma_p)^{-1}$. We ensure $2\sigma_F$ is within the band-limit supported by the planar grid, resulting in the kernel size $\sigma_p = 0.02$ radians.

For each simulation, 10 Gaussian sources of the same size

are laid down at random positions within the FOV and with random amplitudes $A_s \in [0.5, 1.0]$. For each set of 30 simulations, source objects of a different size are considered, with sizes defined by the standard deviation of the Gaussian source $\sigma_s \in \{0.01, 0.02, 0.04, 0.10\}$ radians. For some cases the standard deviation of the source is smaller than the pixel size, hence the Gaussian sources are not necessarily well sampled on the spherical grid. However, this is of no concern since the simulated signals are defined by their discrete version and no contact is subsequently made with the continuous representation from which they originate. To simulate compact Gaussian objects on the sphere we use the COMB (McEwen et al. 2008) and S2 (McEwen et al. 2007) packages⁸. Random visibility coverages are considered with only $\kappa \in \{2, 4, 6, 10, 15, 20, 25\}\%$ of the discrete visibilities measured.

4.2.2 Sparsity

We perform tests to determine whether our hypothesis holds that sparsity is reduced by going to the space in which the signal inherently lives. Although sparsity levels will depend highly on the signal considered, since each set of Gaussian simulations has similar properties we expect sparsity levels to be reasonably consistent over the set of simulations. We therefore average the sparsities measured over the set of 30 simulations for each source size and also provide an indication of their spread.

Since the signals analysed are compressible rather than exactly sparse, we require a measure of compressibility. We construct such a measure as follows. For the BP problem, we first order the coefficients of the signal in the sparsity basis, which in the case of the Dirac sparsity basis that we adopt corresponds to simply ordering the sampled signal values. We then set the smallest coefficient to zero and measure both the sparsity of the resultant signal and also the SNR of the original signal relative to the error between the original and resultant signal. We set the next smallest coefficient to zero and repeat these two measurements, repeating the procedure until all signal values are set to zero. Following this approach we build a curve of sparsity against SNR. For the TV problem, we repeat the same procedure but in the space of the magnitude of the gradient of the signal, rather than a sparsity basis.

Sparsity measurements are computed for the signal on the sphere and its projected version on the plane in order to make comparisons. Curves are plotted in Fig. 4 for the two extreme source sizes (curves for intermediate source sizes are essentially interpolations between these extremes). Sparsity is clearly enhanced on the sphere. Moreover, the ratio of sparsities between the plane and sphere (K_p/K_s) does indeed approximately follow the ratio of the number of samples between the plane and sphere ($N_p/N_s \sim 1.9$ for $\theta_{\text{FOV}} = 90^\circ$), as predicted naively in Sec. 3.1. Although sparsity will always be highly dependent on the signal under investigation, we have at least demonstrated for Gaussian sources that the sparsity of the signal is indeed enhanced on the sphere.

4.2.3 Reconstruction performance

We evaluate reconstruction performance in the WFOV setting by recovering interferometric images directly on both the sphere and the plane, in order to make a direct comparison. Furthermore, we consider BP and TV reconstructions, both with and without application of the spread spectrum phenomenon. It is of particular

⁸ <http://www.jasonmcewen.org/>

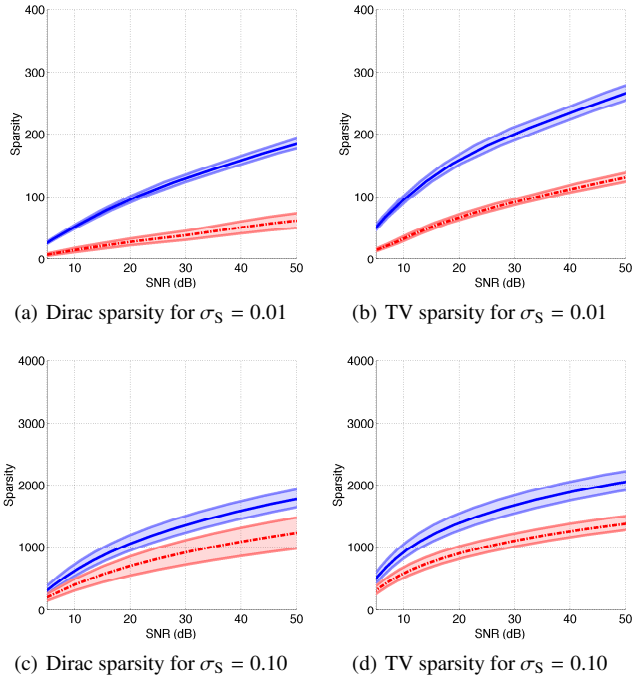


Figure 4. Imposed sparsity of simulations of Gaussian sources as a function of the SNR of the original compressible signal compared to the imposed sparse signal (as explained in the text). Curves correspond to the mean of 30 simulations, while shaded regions show one standard deviation confidence intervals. Curves are plotted for the signal defined on the sphere (red/dot-dashed line) and for its projected version on the plane (blue/solid line). Sparsity is clearly enhanced on the sphere.

interest to see whether our predictions are demonstrated through reconstruction quality.

The measurement constraint ϵ is defined by the level $p = 0.99$ (as described in Sec. 2.2) and we solve the BP and TV minimisation problems using the algorithms derived by Combettes & Pesquet (2007) and Durand et al. (2010) respectively. For the low-resolution simulations performed here, the computation time required to solve the optimisation problems are typically of the order of one minute on a standard laptop with a 2.66GHz Intel Core 2 Duo processor with 4GB of memory.

The quality of reconstruction is measured by the SNR of the original signal relative to the difference between the original and reconstructed signal, defined explicitly by $\text{SNR}_s = 10 \log_{10}(\sigma_{x_s}^2 / \sigma_{x_s - x_s^*}^2)$, where $\sigma_{x_s}^2$ is the variance of the original signal x_s and $\sigma_{x_s - x_s^*}^2$ is the variance of the discrepancy signal $x_s - x_s^*$. Since the sky intensity to be recovered lives inherently on the sphere, we consider the SNR defined on the sphere. It is therefore necessary to lift the image reconstructed on the plane to the sphere in order to make a comparison. We do this through a projection operator that is the direct analogue of the operator \mathbf{P} that projects the sphere to the plane, using an identical kernel in the convolutional re-gridding. We then compare SNR_s , measured on the sphere, for both the spherical and planar based reconstructions. Fig. 5 shows reconstruction quality measured by SNR_s for various visibility coverages, averaged over the 30 simulations for each source size. Reconstruction quality in the spherical setting is clearly superior to the quality of planar reconstructions. However, for sources of small size, lifting the planar reconstruction to the sphere introduces error and limits the effectiveness of the recon-

struction. Before discussing reconstruction performance in more detail, we consider the SNR defined on the plane.

We also examine the SNR defined on the plane by $\text{SNR}_p = 10 \log_{10}(\sigma_{x_p}^2 / \sigma_{x_p - x_p^*}^2)$, where $\sigma_{x_p}^2$ and $\sigma_{x_p - x_p^*}^2$ are the variances of the original and discrepancy signals on the plane respectively. To compute SNR_p for interferometric images recovered on the sphere, the spherical reconstructions are projected onto the plane by the projection operator \mathbf{P} . Fig. 6 shows reconstruction quality measured by SNR_p for various visibility coverages, averaged over the 30 simulations for each source size. The superiority of reconstructions on the sphere is again clear. Even if one were interested in planar reconstructions, superior reconstruction quality is achieved by first recovering interferometric images on the sphere, before projecting the recovered spherical image to the plane. In any case, we advocate the direct use of spherical reconstructions since signal content is not distorted by any projection.

The reconstruction performance observed in Fig. 5 and Fig. 6 is now discussed in more detail and related to the predictions that we made through intuitive reasoning. For BP reconstructions, the improvement in spherical reconstruction quality due to the spread spectrum phenomenon is apparent, whereas the spread spectrum phenomenon is clearly ineffective on the plane. For TV reconstructions, the spread spectrum phenomenon is effective both on the sphere and the plane, although the enhancement on the sphere is slightly larger than that on the plane. Also note that the variance of reconstruction quality (as indicated by the size of the error bars) is reduced by the spread spectrum phenomenon in the cases where it is effective. This effect was also observed by Wiaux et al. (2009b). Notice that the performance of BP reconstructions drops as the size of the Gaussian sources increase due to the corresponding increase in sparsity value. However, this is not the case for TV reconstructions. Although TV reconstructions do not perform as well for signals that are extremely sparse in the Dirac basis, they are more stable as sparsity varies and are certainly superior for reconstructing diffuse signal content. In summary, all of our intuitive predictions are indeed manifest in the reconstruction quality observed and the superiority of reconstructing the sky intensity directly on the sphere in the WFOV interferometric imaging framework that we develop is clear.

4.2.4 Comparison of performance with sparsity and coherence

Although mutual coherence is a useful measure for theoretical and intuitive considerations, it is not well suited to numerical computation (as discussed in Sec. 2.2). Cumulative coherence is a more robust measure of coherence and thus is more suitable for numerical evaluation. Furthermore, cumulative coherence is signal dependent and incorporates sparsity information. We therefore use cumulative coherence to study the combined sparsity and coherence of the interferometric imaging framework in the context of Gaussian sources and relate this to the reconstruction performance presented in the previous section. However, we may only study cumulative coherence in the presence of a sparsity basis, thus the analysis and discussion here is restricted to BP reconstructions only.

In a strict compressed sensing framework with orthogonal measurement and sparsity bases and in the absence of noise, the number of measurements required to reconstruct a signal accurately from incomplete random measurements evolves as (8). The square of the normalised cumulative coherence $\nu \sqrt{N}$ provides an approximate relative measure of the number of measurements required to recover a signal, or similarly, the quality of the reconstruction

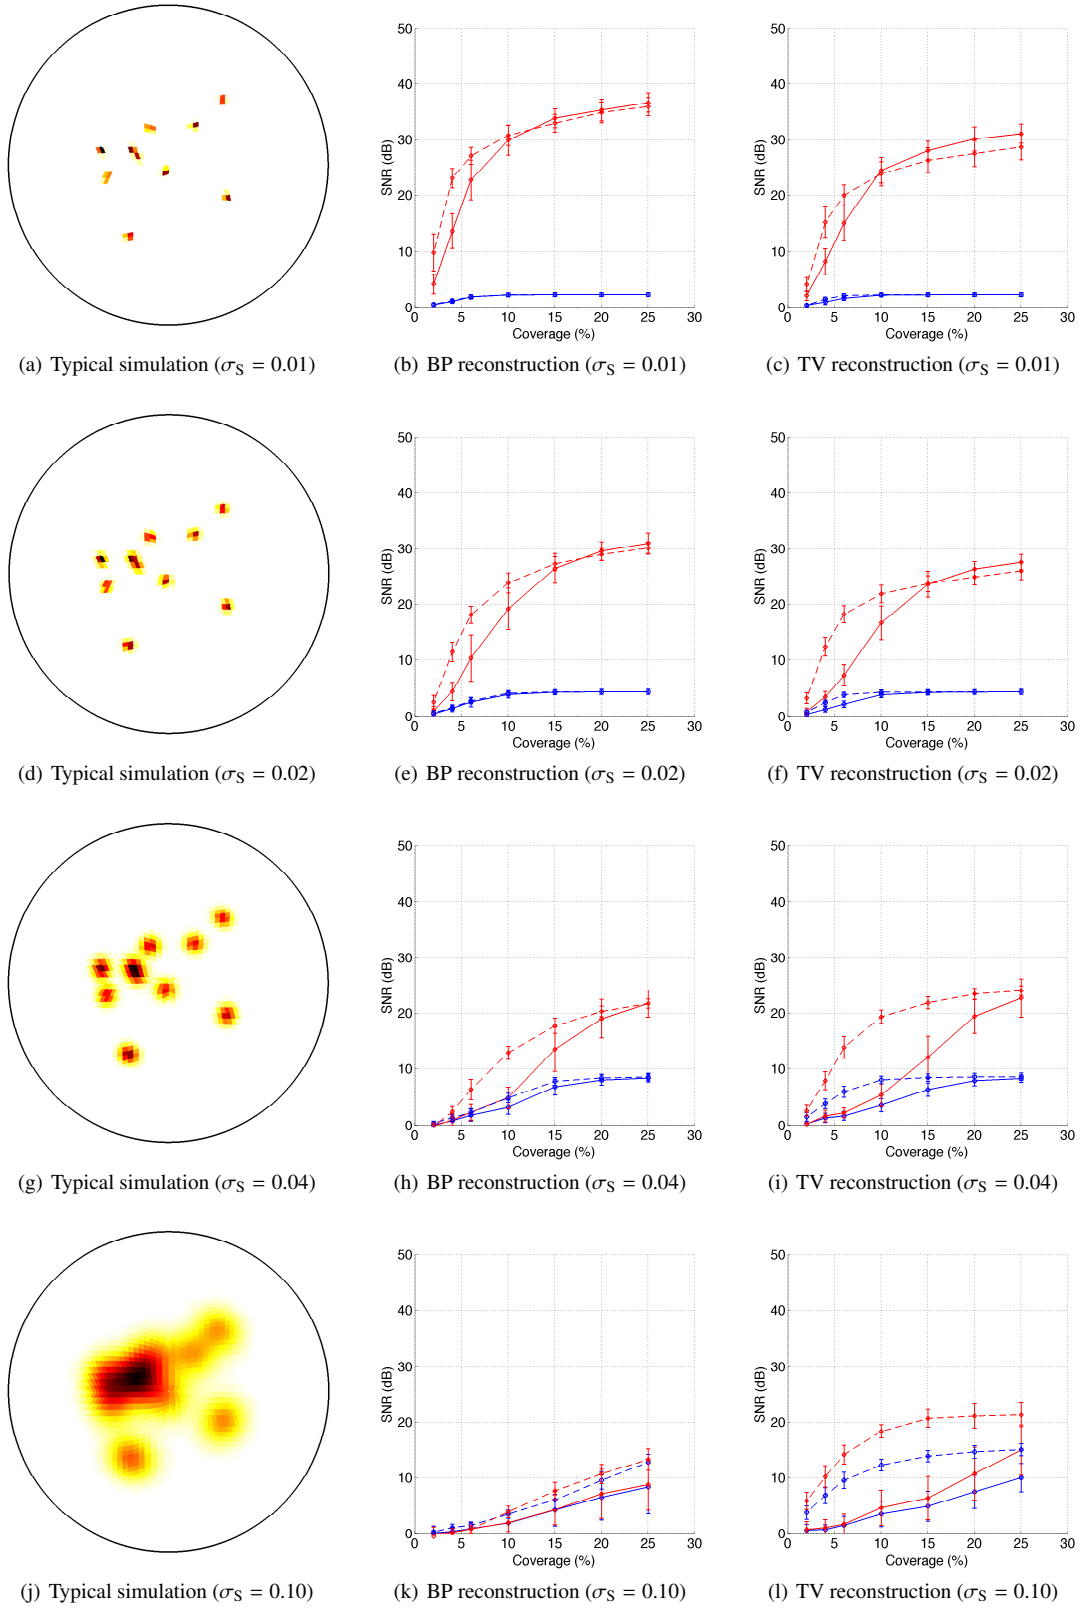


Figure 5. Reconstruction quality for simulated Gaussian sources measured by the SNR defined on the sphere (SNR_S). Each row of panels corresponds to Gaussian sources of a given size. The first column of panels shows a typical simulation of Gaussian sources on the sphere. The remaining columns of panels show reconstruction quality, with the second column illustrating the performance of BP reconstructions, while the third column illustrates the performance of TV reconstructions. Curves are plotted for reconstructions performed on the sphere (red/diamonds) and on the plane (blue/circles), both in the absence (solid lines) and presence (dashed lines) of the spread spectrum phenomenon. Reconstruction quality is averaged over 30 simulations for each source size and error bars corresponding to one standard deviation are shown.

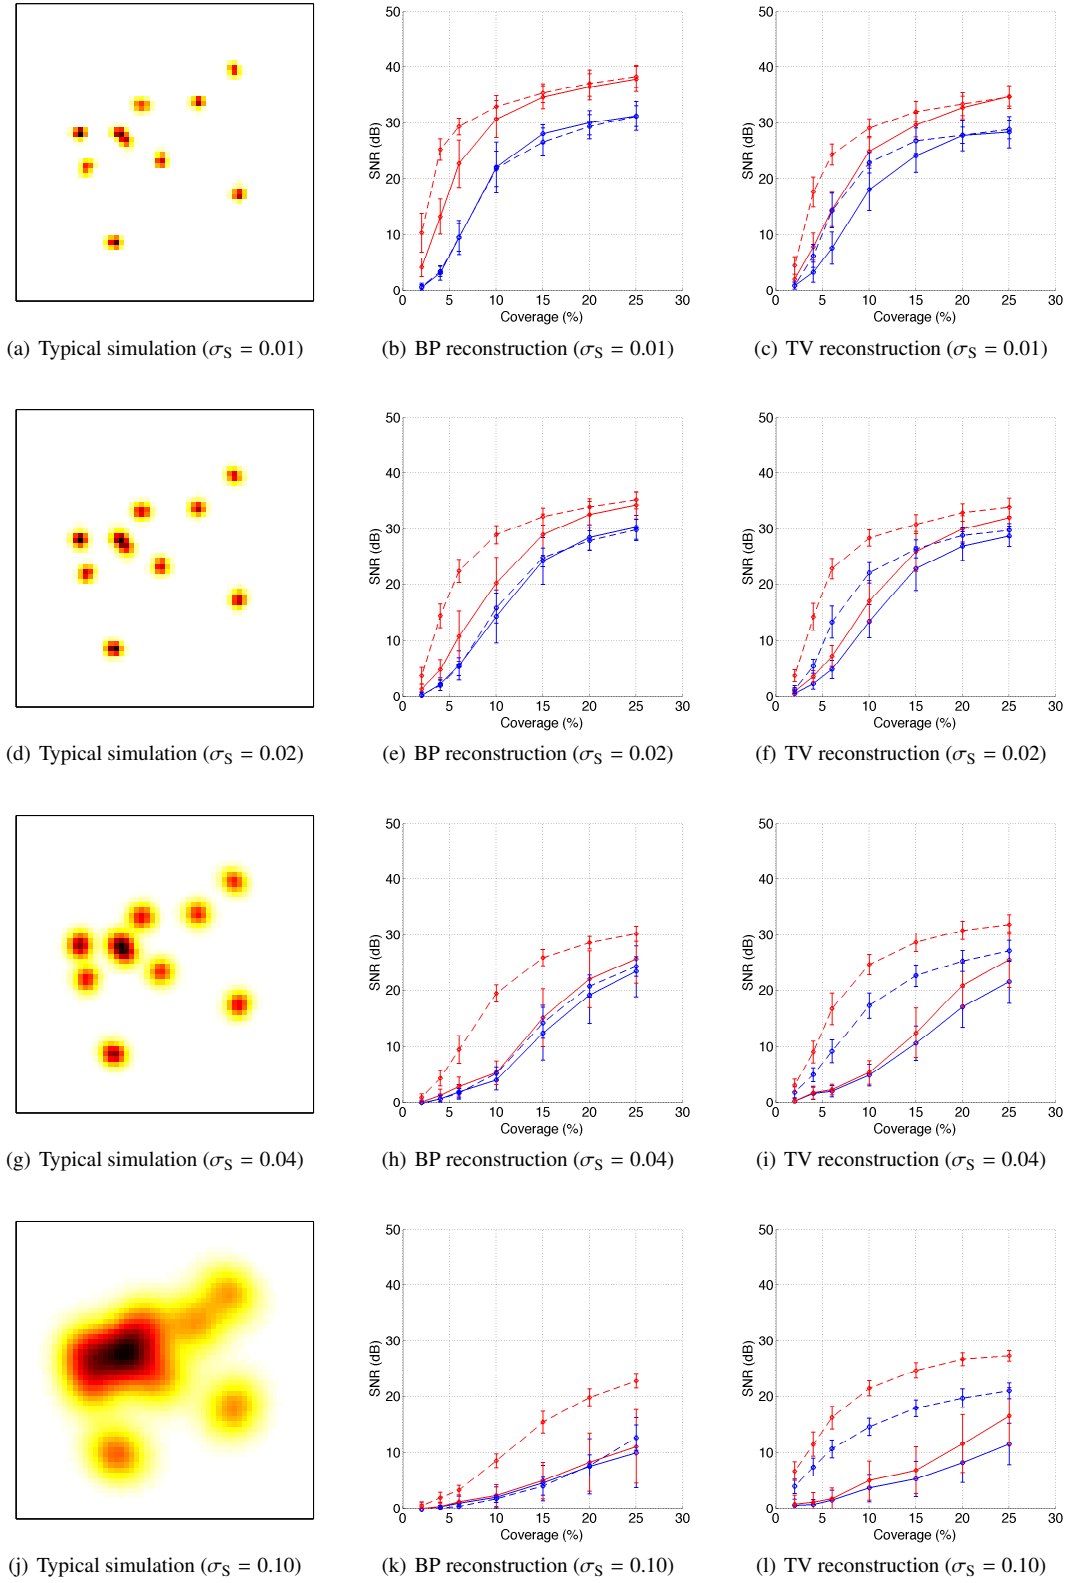


Figure 6. Reconstruction quality for simulated Gaussian sources measured by the SNR defined on the plane (SNR_p). Each row of panels corresponds to Gaussian sources of a given size. The first column of panels shows a typical simulation of Gaussian sources projected onto the plane. The remaining columns of panels show reconstruction quality, with the second column illustrating the performance of BP reconstructions, while the third column illustrates the performance of TV reconstructions. Curves are plotted for reconstructions performed on the sphere (red/diamonds) and on the plane (blue/circles), both in the absence (solid lines) and presence (dashed lines) of the spread spectrum phenomenon. Reconstruction quality is averaged over 30 simulations for each source size and error bars corresponding to one standard deviation are shown.

performance for a given number of measurements. Although the interferometric imaging framework we consider differs from the theoretical compressed sensing framework, the normalised cumulative coherence nevertheless provides a measure of the impact of sparsity and coherence on reconstruction performance.

We compute cumulative coherences for all of the sets of Gaussian simulations, averaging over the 30 simulations for each source size. Normalised cumulative coherences are computed on the plane and sphere, denoted $\nu_p(w_d) \sqrt{N_p}$ and $\nu_s(w_d) \sqrt{N_s}$ respectively, both in the presence ($w_d = 1/\sqrt{2}$) and absence ($w_d = 0$) of the spread spectrum phenomenon. Results are displayed in Table 1. Four insights can be made by studying the coherence values reported in Table 1. Firstly, the coherences on the plane are not significantly altered in the presence of the spread spectrum phenomenon, indicating that the spread spectrum phenomenon is likely to be ineffective in improving reconstruction performance in this setting. Secondly, the coherences on the sphere are reduced by application of the spread spectrum phenomenon, indicating that the spread spectrum phenomenon is likely to be effective in this setting. Thirdly, we note that coherences on the sphere are always lower than the corresponding values on the plane, highlighting the superiority of spherical reconstruction. Fourthly, we note that cumulative coherences increase with source size, indicating that reconstruction performance should reduce. All of these findings verify our intuitive expectations and are consistent with the reconstruction performance presented in the previous section.

4.3 Galactic dust

Now that the WFOV interferometric imaging framework developed in this article has been evaluated thoroughly, in this section we simply illustrate the reconstruction, at a higher resolution, of a more realistic simulation of Galactic dust emission. The simulation is first described, before reconstruction performance is evaluated.

4.3.1 Simulation

For this higher resolution simulation we consider the 94GHz map of predicted submillimeter and microwave emission of diffuse interstellar Galactic dust (Finkbeiner et al. 1999), hereafter referred to as the FDS map. This predicted map is based on the merged Infrared Astronomy Satellite (IRAS) and Cosmic Background Explorer Diffuse Infrared Background Experiment (COBE-DIRBE) observations produced by Schlegel et al. (1998). An undersampled version of the FDS map is available from the Legacy Archive for Microwave Background Data Analysis⁹ (LAMBDA) in the HEALPix pixelisation at resolution $N_{\text{side}} = 512$.

We consider the same observational set-up discussed in Sec. 4.1, focusing on the FOV centred on Galactic coordinates $(l, b) = (210^\circ, -20^\circ)$. The full-sky FDS map and the FOV considered are illustrated in Fig. 7. We downgrade the original FDS map to a resolution of $N_{\text{side}} = 128$ for our simulated observations, corresponding to a harmonic band-limit of $\ell_{\text{max}} = 348$. All other resolution parameters follow once the harmonic band-limit on the sphere and the size of the FOV are chosen; we find $N_s = 28560$, $u_{\text{max}} = 78.4$ and $N_p = 214 \times 214 = 45796$. The size of the Gaussian kernel in the convolutional re-gridding of the projection operation is chosen by the same condition as before, resulting in the value

⁹ <http://lambda.gsfc.nasa.gov/>

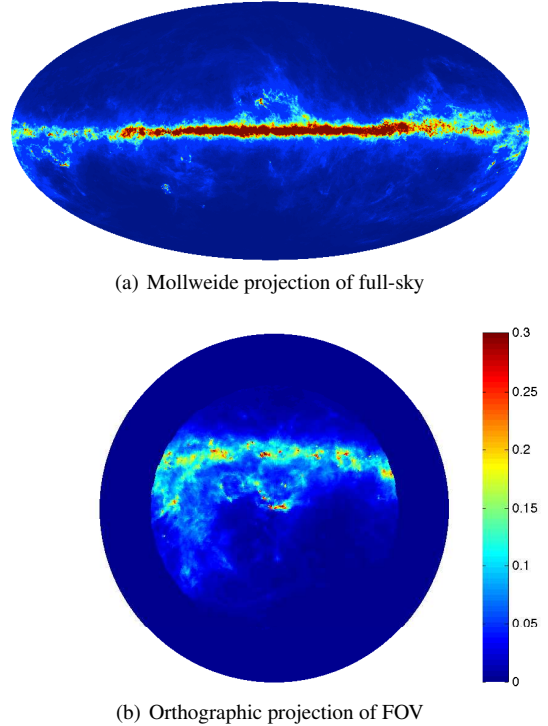


Figure 7. FDS map of predicted submillimeter and microwave emission of diffuse interstellar Galactic dust. The FOV of angular opening $\theta_{\text{FOV}} = 90^\circ$ and centred on Galactic coordinates $(l, b) = (210^\circ, -20^\circ)$ is considered for simulating interferometric observations. In panel (a) the FOV is located near the Galactic plane towards the left edge of the image.

$\sigma_p = 0.004$ radians. Random visibility coverage is again considered, with only 25% of the discrete visibilities measured. Note that the incomplete visibility coverage creates a very challenging setting for the recovery of a realistic, diffuse image.

4.3.2 Reconstruction performance

All of the reconstruction techniques discussed previously are applied to reconstruct the FDS map on the sphere (*i.e.* BP and TV reconstructions in the spherical and planar settings, both with and without application of the spread spectrum phenomenon). The same optimisation algorithms and measurement constraint level p considered previously are applied. For these high-resolution simulations, the computation time required to solve the optimisation problems are typically of the order of 15 minutes on the same machine described previously. The results presented in Sec. 4.2 indicate that for diffuse signals the SNR measured on the sphere (SNR_s) is not adversely affected significantly by lifting planar reconstructions to the sphere. Due to the diffuse nature of the underlying signal (and also since it is inherently defined on the sphere), we therefore measure reconstruction quality by SNR_s only.¹⁰ Note that for such signals the SNR metric is not a perfect measure of reconstruction quality and a visual inspection remains important. The diffuse nature of the FDS map and the results presented previously suggest

¹⁰ Although we only quote SNR measured on the sphere, the SNR measured on the plane (SNR_p) were also examined. Corresponding SNR values do change marginally, nevertheless the conclusions drawn remain the same regardless of whether SNR_s or SNR_p is examined.

Source σ_s	Plane			Sphere			Difference $\Delta(w_d)$	
	$w_d = 0$	$v_p(w_d) \sqrt{N_p}$ $w_d = \frac{1}{\sqrt{2}}$	Δ_p	$w_d = 0$	$v_s(w_d) \sqrt{N_s}$ $w_d = \frac{1}{\sqrt{2}}$	Δ_s	$w_d = 0$	$w_d = \frac{1}{\sqrt{2}}$
0.01	8.14 ± 0.72	8.31 ± 0.78	0.17 ± 0.10	5.07 ± 0.52	4.69 ± 0.55	-0.38 ± 0.06	-3.08 ± 0.23	-3.62 ± 0.25
0.02	11.05 ± 0.95	11.17 ± 1.00	0.12 ± 0.08	8.23 ± 0.83	7.64 ± 0.88	-0.59 ± 0.08	-2.82 ± 0.15	-3.53 ± 0.16
0.04	16.32 ± 1.19	16.36 ± 1.23	0.05 ± 0.05	13.85 ± 1.17	12.84 ± 1.25	-1.01 ± 0.11	-2.46 ± 0.07	-3.52 ± 0.13
0.10	22.32 ± 0.25	22.33 ± 0.25	0.01 ± 0.01	19.79 ± 0.25	18.23 ± 0.25	-1.73 ± 0.03	-2.53 ± 0.02	-4.12 ± 0.05

Table 1. Normalised cumulative coherences, on the plane and sphere, of simulations of Gaussian sources for various size σ_s . Coherences values computed both in the presence ($w_d = 1/\sqrt{2}$) and absence ($w_d = 0$) of the spread spectrum phenomenon are displayed. Normalised coherence differences between the presence and absence of the spread spectrum phenomenon are computed on the plane and sphere (with Δ_p and Δ_s denoting, respectively, the coherence for $w_d = 0$ subtracted from the coherence for $w_d = 1/\sqrt{2}$). In addition, normalised coherence differences between the plane and sphere are also computed, with $\Delta(w_d) = v_p(w_d) \sqrt{N_p} - v_s(w_d) \sqrt{N_s}$. Quoted values correspond to the mean of 30 simulations, with errors corresponding to one standard deviation. Notice that coherence differences are more stable than coherence values over the simulations. These measures incorporate both the impact of sparsity and coherence on reconstruction performance (with lower values indicating superior performance).

that the TV reconstruction on the sphere in the presence of spread spectrum phenomenon will be most effective.

Reconstructed spherical interferometric images of the FDS map are displayed in Fig. 8, with SNR_s of each recovery also specified. The quality of BP reconstructions is relatively poor (as expected since the FDS map is not particularly sparse in the Dirac basis), and a significant difference between planar and spherical reconstructions is not readily apparent. Spherical reconstructions show greater detail than the planar reconstructions, particularly within the centre of the FOV, but suffer near the extremes of the FOV since the signal is poorly constrained here due to the low magnitude of the primary beam. Planar reconstructions also suffer from this effect, however it is not as apparent in reconstructed spherical images since a small degree of smoothing is performed in the projection operator when lifting the reconstructed planar signal to the sphere. The quality of reconstruction is improved considerably for the TV case, again as expected since the diffuse FDS map is much sparser in the magnitude of its gradient than it is in the Dirac basis. Spherical reconstructions again show greater detail than the planar reconstructions, however in the absence of the spread spectrum phenomenon SNR_s is in fact lower for the spherical reconstruction. On visual inspection the spherical reconstruction is clearly superior as structure of much finer detail is reconstructed, highlighting the weakness of the SNR metric. In any case, the superiority of the spherical reconstruction is clearly apparent for the TV reconstruction when the spread spectrum phenomenon is present, both through visual inspection and comparison of SNR_s . This is the most effective reconstruction technique for the FDS map for both planar and spherical reconstructions, as expected. Comparing the most effective reconstruction on the plane and sphere, recovering the sky intensity directly on the sphere improves reconstruction quality from 13.7dB to 19.3dB.

5 CONCLUSIONS

Incorporating WFOV contributions when reconstructing interferometric images is becoming increasingly important, particularly for next-generation interferometers, such as the SKA and upcoming pathfinder telescopes. If these contributions are ignored, the fidelity of reconstructed images may suffer. In this article we have extended the compressed sensing interferometric imaging techniques developed by Wiaux et al. (2009a) and Wiaux et al. (2009b) to a WFOV. In doing so, we recover interferometric images defined directly on

the celestial sphere. We augment the usual measurement operator with a projection from the sphere to the plane, which essentially corresponds to a change from Cartesian to spherical coordinates. Practically, however, the projection is complicated by the discrete setting and careful consideration is given to the sampling of signals on the sphere and plane to ensure that the planar grid is sufficiently sampled to support the maximum projected frequency content of a band-limited signal on the sphere. Although a projection is incorporated in this framework, it is included in the measurement operator and thus is regularised when solving the interferometric inverse problem. Moreover, it is the space where recovery is performed that drives the performance of compressed sensing reconstructions, through sparsity and coherence considerations. The framework remains general and does not rely on any small FOV assumptions.

The effectiveness of the spread spectrum phenomenon is enhanced when going to a WFOV, while sparsity is promoted by recovering images directly on the sphere. These predictions have been verified by numerical tests and are also manifest in reconstruction performance. Low-resolution simulations of Gaussian sources were considered to quantify reconstruction performance thoroughly. Interferometric images were recovered directly on the sphere and the plane in order to make comparisons. For simulated images extremely sparse in the Dirac basis, BP reconstructions were shown to perform very well. However, as Dirac sparsity was reduced the quality of BP reconstructions fall, while the quality of TV reconstructions remained relatively stable. For diffuse images, TV reconstructions were shown to be superior since such signals are much more sparse in the magnitude of their gradient than in the Dirac basis. In all cases, the superior quality of recovering interferometric images directly on the sphere was clear. A simulation of diffuse interstellar Galactic dust was then performed to demonstrate WFOV reconstruction techniques in a more realistic, higher resolution setting. For the diffuse FDS map considered, TV reconstruction on the sphere in the presence of the spread spectrum phenomenon was most effective, as expected. For this case, reconstruction quality was improved from 13.7dB for the planar reconstruction to 19.3dB when recovering the interferometric image directly on the sphere.

The compressed sensing techniques developed for interferometric imaging by Wiaux et al. (2009a) and Wiaux et al. (2009b), and extended here to a WFOV, remain somewhat idealised. Random visibility coverage is assumed, with the spatial frequencies probed by the interferometer also assumed to fall on discrete grid points. Furthermore, to study the spread spectrum phenomenon a

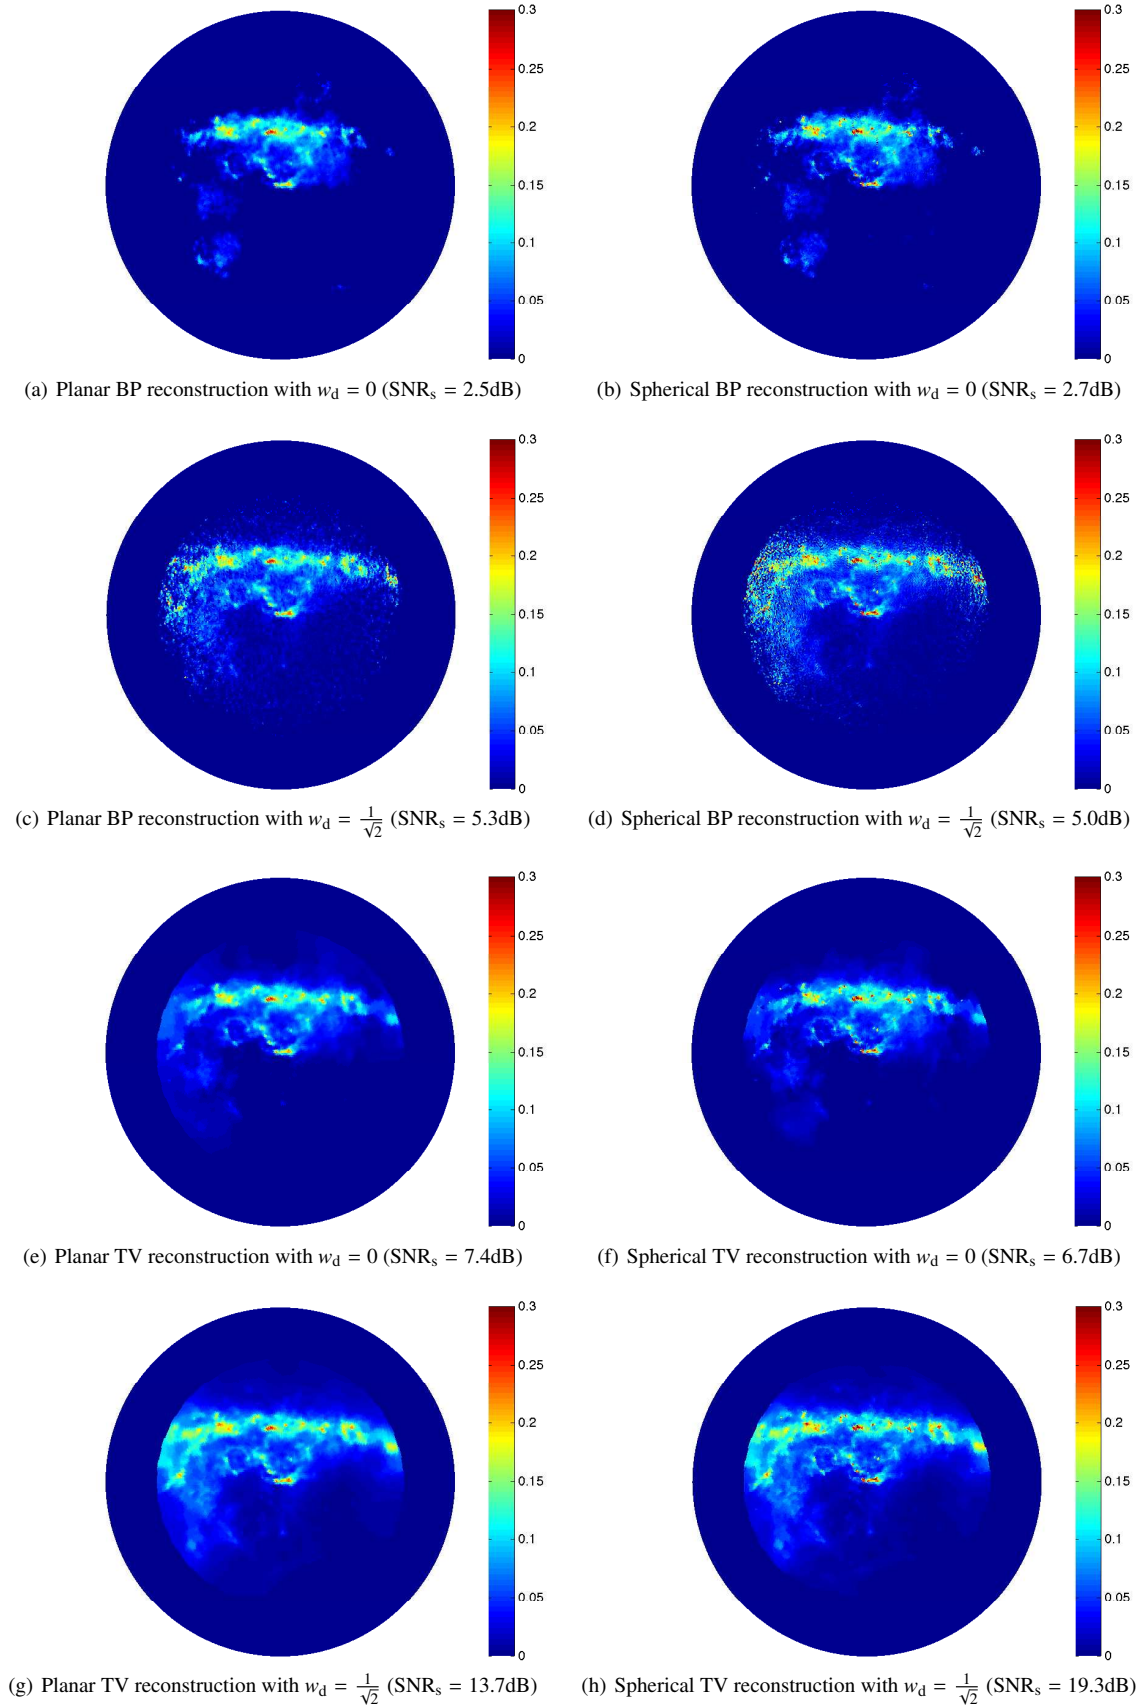


Figure 8. Reconstructed spherical interferometric images of the FDS map. The first column of panels shows images of planar recoveries lifted to the sphere, while the second column shows images recovered on the sphere directly. Reconstructions both in the presence ($w_d = 1/\sqrt{2}$) and absence ($w_d = 0$) of the spread spectrum phenomenon are displayed.

constant w is assumed. These restrictions have been necessary to remain as close to the theory of compressed sensing as possible during the development and evaluation of interferometric imaging techniques. In reality, however, w will not be constant and the performance of the spread spectrum phenomenon will lie between the extreme cases that we have considered of $w = 0$ and $w = u_{\max}$. Extensions to realistic and continuous visibility coverage and their impact on compressed sensing based interferometric imaging are now of considerable importance. In general, compressed sensing addresses imaging by optimising both reconstruction and acquisition, while we have essentially focused on reconstruction only. The possibility of optimising the configuration of interferometers to enhance the spread spectrum phenomenon for compressed sensing reconstruction is an exciting avenue of research at the level of acquisition. In addition, direction dependent beam effects may also provide an alternative source of the spread spectrum phenomenon. All of these issues should be studied in future works. Furthermore, the performance of other sparsity bases on the sphere, such as scale discretised wavelets (Wiaux et al. 2008), should also be studied. Next-generation radio interferometers will inherently observe very large fields of view, thus WFOV interferometric imaging techniques, such as the compressed sensing techniques developed in this article and the future research outlined here, are of increasing importance to ensure that the fidelity of reconstructed images keeps pace with the capabilities of new instruments.

ACKNOWLEDGEMENTS

We thank Gilles Puy for providing some code and for useful discussions. We also thank Pierre Vanderghenst, Jean-Philippe Thiran and Dimitri Van De Ville for providing the infrastructure to support our research. JDM is supported by the Swiss National Science Foundation (SNSF) under grant 200021-130359. YW is supported in part by the Center for Biomedical Imaging (CIBM) of the Geneva and Lausanne Universities, EPFL, and the Leenaards and Louis-Jeantet foundations, and in part by the SNSF under grant PP00P2-123438. We acknowledge the use of the COMB (McEwen et al. 2008), S2 (McEwen et al. 2007) and HEALPix (Górski et al. 2005) packages. We also acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

REFERENCES

- Baraniuk R., 2007, *IEEE Sig. Proc. Mag.*, 24, 4, 118
- Candès E., 2006, in *Proc. Int. Congress Math.*, volume III of *Euro. Math. Soc.*, 1433
- Candès E., Romberg J., Tao T., 2006a, *IEEE Trans. Inform. Theory*, 52, 2, 489
- Candès E., Romberg J., Tao T., 2006b, *Comm. Pure and Appl. Math.*, 59, 8, 1207, [arXiv:math/0503066](#)
- Carilli C., Rawlings S., 2004, *Science with the Square Kilometre Array*, Elsevier, Oxford
- Carozzi T.D., Woan G., 2009, *Mon. Not. Roy. Astron. Soc.*, 395, 1558, [arXiv:0812.0141](#)
- Chambolle A., 2004, *J. Math. Imaging Vis.*, 20, 1–2, 89
- Combettes P., Pesquet J.C., 2007, *IEEE J. Selected Top. in Sig. Proc.*, 1, 4, 564
- Cornwell T.J., Golap K., Bhatnagar S., 2008, *IEEE J. Selected Top. in Sig. Proc.*, 2, 5, 647, [arXiv:0807.4161](#)
- Cornwell T.J., Perley R.A., 1992, *Astron. & Astrophys.*, 261, 353
- Donoho D., 2006, *IEEE Trans. Inform. Theory*, 52, 4, 1289
- Durand S., Fadili J., Nikolova M., 2010, *J. Math. Imaging Vis.*, 36, 3, 201, [arXiv:0812.1697](#)
- Finkbeiner D.P., Davis M., Schlegel D.J., 1999, *Astrophys. J.*, 524, 867, [astro-ph/9905128](#)
- Górski K.M., Hivon E., Banday A.J., Wandelt B.D., Hansen F.K., Reinecke M., Bartelmann M., 2005, *Astrophys. J.*, 622, 759, [astro-ph/0409513](#)
- Hivon E., Hansen F.K., Wandelt B.D., Gorski K.M., Banday A.J., Reinecke M., 2010, *HEALPix Fortran Facility User Guidelines*, V2.15a
- Högbom J.A., 1974, *Astron. & Astrophys. Supp.*, 15, 417
- McEwen J.D., Hobson M.P., Lasenby A.N., 2008, *IEEE Trans. Sig. Proc.*, 56, 8, 3813, [astro-ph/0612688](#)
- McEwen J.D., Hobson M.P., Mortlock D.J., Lasenby A.N., 2007, *IEEE Trans. Sig. Proc.*, 55, 2, 520, [astro-ph/0506308](#)
- McEwen J.D., Scaife A.M.M., 2008, *Mon. Not. Roy. Astron. Soc.*, 389, 3, 1163, [arXiv:0803.2165](#)
- Potts D., Steidl G., Tasche M., 2001, in J.J. Benedetto, P.J.S.G. Ferreira, editors, *Modern Sampling Theory: Mathematics and Applications*, Birkhäuser, Boston, 247–270
- Rauhut H., 2010, in M. Fornasier, editor, *Theoretical Foundations and Numerical Methods for Sparse Recovery*, volume 9 of *Radon Series Comp. Appl. Math.*, de Gruyter, 1–92
- Romberg J., 2009, *SIAM J. Imaging Sciences*, 2, 4, 1098
- Rudin L.I., Osher S., Fatemi E., 1992, *Physica D: Nonlinear Phenomena*, 60, 1–4, 259
- Schlegel D.J., Finkbeiner D.P., Davis M., 1998, *Astrophys. J.*, 500, 525, [astro-ph/9710327](#)
- Suksmono A., 2009, in *ICEEI*, volume 1, 110–116
- Thompson A.R., Moran J.M., Jr. G.W.S., 2001, *Interferometry and synthesis in radio astronomy*, Wiley-VCH, Weinheim, 2nd edition
- Tropp J., 2004, *IEEE Trans. Inform. Theory*, 50, 10, 2231
- Wandelt B.D., Hivon E., Gorski K.M., 1998, in *XXXIIIrd Rencontres de Moriond*, [astro-ph/9803317](#)
- Wenger S., Darabi S., Sen P., Glassmeier K.H., Magnor M., 2010, in *ICIP*, 1381–1384
- Wiaux Y., Jacques L., Puy G., Scaife A.M.M., Vanderghenst P., 2009a, *Mon. Not. Roy. Astron. Soc.*, 395, 3, 1733, [arXiv:0812.4933](#)
- Wiaux Y., McEwen J.D., Vanderghenst P., Blanc O., 2008, *Mon. Not. Roy. Astron. Soc.*, 388, 2, 770, [arXiv:0712.3519](#)
- Wiaux Y., Puy G., Boursier Y., Vanderghenst P., 2009b, *Mon. Not. Roy. Astron. Soc.*, 400, 2, 1029, [arXiv:0907.0944](#)
- Wiaux Y., Puy G., Vanderghenst P., 2010, *Mon. Not. Roy. Astron. Soc.*, 402, 4, 2626, [arXiv:0908.4179](#)