

Spread spectrum for imaging techniques in radio interferometry

Y. Wiaux,^{1,2*} G. Puy,¹ Y. Boursier¹ and P. Vandergheynst¹

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

²*Centre for Particle Physics and Phenomenology, Université catholique de Louvain (UCL), B-1348 Louvain-la-Neuve, Belgium*

Accepted 2009 August 6. Received 2009 August 4; in original form 2009 July 6

ABSTRACT

We consider the probe of astrophysical signals through radio interferometers with a small field of view and baselines with a non-negligible and constant component in the pointing direction. In this context, the visibilities measured essentially identify with a noisy and incomplete Fourier coverage of the product of the planar signals with a linear chirp modulation. In light of the recent theory of compressed sensing and in the perspective of defining the best possible imaging techniques for sparse signals, we analyse the related spread spectrum phenomenon and suggest its universality relative to the sparsity dictionary. Our results rely both on theoretical considerations related to the mutual coherence between the sparsity and sensing dictionaries and on numerical simulations.

Key words: techniques: image processing – techniques: interferometric.

1 INTRODUCTION

Aperture synthesis in radio interferometry is a powerful technique in radio astronomy, allowing observations of the sky with otherwise inaccessible angular resolutions and sensitivities. This technique dates back to more than 60 years ago (Ryle & Vonberg 1946; Blythe 1957; Ryle, Hewish & Shakeshaft 1959; Ryle & Hewish 1960; Thompson, Moran & Swenson 2004). In this context, the portion of the celestial sphere around the pointing direction tracked by a radio telescope array during observation defines the original real signal or image probed x . The field of view observed is limited by a primary beam A . Standard interferometers are characterized by a small field of view, so that the signal and the primary beam are assumed to be planar. Considering non-polarized radiation, they, respectively, read as scalar functions $x(\mathbf{l})$ and $A(\mathbf{l})$ of the position $\mathbf{l} \in \mathbb{R}^2$, with components (l, m) . Each telescope pair at one instant of observation identifies a baseline defined as the relative position between the two telescopes. With each baseline $\mathbf{b}_\lambda \in \mathbb{R}^3$, with components (u, v, w) in units of the signal emission wavelength λ , is associated one measurement called visibility. In the simplest setting, one also considers baselines with a negligible component w in the pointing direction of the instrument. Under this additional assumption, if the signal is made up of incoherent sources, each visibility corresponds to the value of the Fourier transform of the signal multiplied by the primary beam at a spatial frequency $\mathbf{u} \in \mathbb{R}^2$, identified by the components (u, v) of the baseline projection on the plane of the signal. Radio-interferometric data are thus identified by incomplete and noisy measurements in the Fourier plane. In the perspective of the reconstruction of the original image, these data define an ill-posed inverse problem.

It is well known that a large variety of natural signals are sparse or compressible in multiscale dictionaries, such as wavelet frames. A band-limited signal may be expressed as the N -dimensional vector of its values sampled at the Nyquist–Shannon rate. By definition, a signal is sparse in some basis if its expansion contains only a small number $K \ll N$ of non-zero coefficients. More generally it is compressible if its expansion only contains a small number of significant coefficients, i.e. if a large number of its coefficients bear a negligible value. The theory of compressed sensing demonstrates that, for sparse or compressible signals, a small number $M \ll N$ of random measurements, in a sensing basis incoherent with the sparsity basis, will suffice for an accurate and stable reconstruction of the signals (Candès 2006; Candès, Romberg & Tao 2006a,b; Donoho 2006; Baraniuk 2007; Donoho & Tanner 2009). The mutual coherence between two bases may be defined as the maximum complex modulus of the scalar product between unit-norm vectors of the two bases. Random Fourier measurements of a signal sparse in real space are a particular example of a good sensing procedure in this context.

In a very recent work (Wiaux et al. 2009), we presented results showing that compressed sensing offers powerful image reconstruction techniques for radio-interferometric data, in the case of baselines with a negligible component w . These techniques are based on global Basis Pursuit (BP) minimization problems, which are solved by convex optimization algorithms. We particularly illustrated the versatility of the scheme relative to the inclusion of specific prior information on the signal in the minimization problems.

In the present work, we raise the important problem of the dependence of the image reconstruction quality as a function of the sparsity basis or more generally the sparsity dictionary. The larger the typical size of the waveforms constituting the dictionary in which the signal is sparse or compressible in real space, the smaller their extension in the Fourier plane and the smaller the incoherence

*E-mail: yves.wiaux@epfl.ch

between the sparsity and sensing dictionaries. In the context of compressed sensing, a loss of incoherence leads to a degradation of the reconstruction quality for a given sparsity K and a given number M of random measurements.

The detailed structure of radio-interferometric measurements might actually provide a natural response to this issue. The approximation of baselines with a negligible component w is a key assumption in order to identify visibilities with Fourier measurements of the original signal. This approximation actually sets a strong constraint on the field of view probed by the interferometer, requiring that it is small enough, not only for the planar approximation of the signal but also to neglect the complete effect of the component w in the visibilities. We relax this approximation and consider radio interferometers with a small field of view and baselines with a non-negligible component w . In this context, each visibility at spatial frequency \mathbf{u} identifies with the Fourier transform of a complex signal obtained as the product of the original planar signal and the primary beam with a linear chirp $e^{i\pi w|\mathbf{l}|^2}$ where the norm $|\mathbf{l}|$ identifies the distance to the centre of the image. Note that this chirp is a priori characterized by a different chirp rate for each baseline: $w = w(\mathbf{u})$. In the Fourier plane, the modulation amounts to the convolution of the Fourier transform of the chirp with that of the signal multiplied by the primary beam, which inevitably spreads the two-dimensional sample power spectrum of the constitutive waveforms. This spread spectrum phenomenon increases the incoherence between the sparsity and sensing dictionaries.

In this context, we define signals made up of Gaussian waveforms with equal size, identified by a standard deviation t , considered as a rough representation of any kind of astrophysical structures. We make the simplifying assumption that all baselines, identified by the spatial frequencies \mathbf{u} , have the same component w , so that all visibilities are affected by the same chirp. A theoretical computation of the mutual coherence between the sparsity and sensing dictionaries as a function of t and w naturally proves that the coherence decreases when w increases, through the spread spectrum phenomenon. Our theoretical relation also shows that the coherence may be decreased as close as required to the minimum coherence between the real and Fourier spaces through the use of a large enough chirp rate w . In light of the theory of compressed sensing, this suggests some universality of the spread spectrum phenomenon according to which, for a given sparsity K and a given number M of random measurements, the quality of the BP reconstruction would not only increase when w increases, but could also be rendered independent of the sparsity dictionary for a large enough chirp rate w . In other words, the BP reconstruction could be made as good for signals sparse in a dictionary of Gaussian waveforms of any size as for signals sparse in real space. We produce simulations of the signals considered and perform BP reconstructions from noisy visibility measurements. Our results confirm the universality of the spread spectrum phenomenon relative to the sparsity dictionary.

Let us acknowledge the fact that spread spectrum techniques are widely used in telecommunications. In this context, chirp modulations are used in particular to render transmitted signals more robust to noise. In the context of compressed sensing, spread spectrum techniques using pseudo-random modulations have very recently been highlighted as a means to enhance the signal reconstruction and render it stable relative to noise (Naini et al. 2009). The use of chirp modulations has also been proposed in compressed sensing radar to produce a convolution of a sparse signal in the sparsity basis itself before performing measurements in that same basis (Baraniuk & Steeghs 2007; Herman & Strohmer 2009). In this respect, these techniques are much more related to coded aperture techniques

(Gottesman & Fenimore 1989; Marcia & Willett 2008) as well as to compressed sensing by random convolution (Romberg 2008) than to spread spectrum techniques.

In Section 2, we formulate the interferometric inverse problem for image reconstruction on a small field of view in the presence of a non-negligible and constant component w of the baselines. In Section 3, we concisely recall results of the theory of compressed sensing regarding the definition of a sensing basis and the accurate reconstruction of sparse or compressible signals from BP. In Section 4, we establish the universality of the spread spectrum phenomenon for a sparsity dictionary made up of Gaussian waveforms, both from theoretical considerations and on the basis of simulations. We finally conclude in Section 5.

2 RADIO INTERFEROMETRY

In this section, we recall the general form of the visibility measurements and study the approximation of a small field of view and baselines with a non-negligible component w . We identify in particular the corresponding spread spectrum phenomenon in the presence of a constant component w . We also pose the corresponding interferometric inverse problem for image reconstruction.

2.1 General visibilities

In a tracking configuration, all radio telescopes of an interferometric array point in the same direction on the unit celestial sphere S^2 identified by a unit vector $\hat{s}_0 \in \mathbb{R}^3$. Arbitrary directions are denoted by unit vectors $\hat{s} \in \mathbb{R}^3$. The field of view observed is limited by a so-called primary beam $A(\boldsymbol{\tau})$, where the vector $\boldsymbol{\tau} \equiv \hat{s} - \hat{s}_0 \in \mathbb{R}^3$ identifies directions relative to the pointing direction (see Fig. 1). The size of its angular support is essentially inversely proportional to the size of the dishes of the telescopes. We consider a monochromatic signal $x(\boldsymbol{\tau})$ with an emission wavelength λ and made up of incoherent sources (see Fig. 1). Also considering non-polarized radiation, both the signal and the primary beam are scalar square-integrable functions on the sphere, $x(\boldsymbol{\tau}), A(\boldsymbol{\tau}) \in L^2(S^2, d\Omega(\boldsymbol{\tau}))$,




Figure 1. Illustration of notations. The signal probed $x(\boldsymbol{\tau})$, with $\boldsymbol{\tau} \equiv \hat{s} - \hat{s}_0$, extends in any direction on the celestial sphere S^2 identified by unit vectors $\hat{s} \in \mathbb{R}^3$, around the pointing direction $\hat{s}_0 \in \mathbb{R}^3$. The field of view of the radio interferometer is set by the primary beam $A(\boldsymbol{\tau})$. The baselines also characterizing the interferometer are vectors $\mathbf{b}_\lambda \in \mathbb{R}^3$ with norm $|\mathbf{b}_\lambda|$. The coordinate system $R : (O, \hat{e}_1, \hat{e}_2, \hat{e}_3)$ in \mathbb{R}^3 is illustrated, with \hat{e}_3 identified with \hat{s}_0 , \hat{e}_2 pointing towards north and \hat{e}_1 pointing towards east. The components of the direction vectors \hat{s} on the sphere are denoted by (l, m, n) , while those of a baseline are denoted by (u, v, w) .

where $d\Omega(\boldsymbol{\tau})$ stands for the invariant measure on the sphere. At each instant of observation, each telescope pair measures a complex visibility defined as the correlation between incoming electric fields at the positions of the two telescopes. This visibility only depends on the relative position between the two telescopes, defined as a baseline. The baseline in units of λ is denoted by a vector $\boldsymbol{b}_\lambda \in \mathbb{R}^3$ (see Fig. 1). Each visibility measured reads in a general form as (Thompson et al. 2004)

$$y(\boldsymbol{b}_\lambda) \equiv \int_{S^2} A x(\boldsymbol{\tau}) e^{-2i\pi \boldsymbol{b}_\lambda \cdot \boldsymbol{\tau}} d\Omega(\boldsymbol{\tau}). \quad (1)$$

We consider the conventional right-handed Cartesian coordinate system $\mathbf{R} : (O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2, \hat{\boldsymbol{e}}_3)$ in \mathbb{R}^3 centred on the Earth. The direction $\hat{\boldsymbol{e}}_3$ identifies with the pointing direction $\hat{\boldsymbol{s}}_0$. The directions $\hat{\boldsymbol{e}}_1$ and $\hat{\boldsymbol{e}}_2$ identify the orthogonal plane, parallel to the plane tangent to the celestial sphere at the pointing direction, with $\hat{\boldsymbol{e}}_2$ pointing towards north and $\hat{\boldsymbol{e}}_1$ pointing towards east. The unit vectors $\hat{\boldsymbol{s}}$ may be identified by their components (l, m, n) in \mathbf{R} (see Fig. 1). Each point is thus obviously identified by the two-dimensional vector $\boldsymbol{l} \in \mathbb{R}^2$ with components (l, m) in the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$, while the component in direction $\hat{\boldsymbol{e}}_3$ reads as $n(\boldsymbol{l}) = (1 - |\boldsymbol{l}|^2)^{1/2}$ with the norm $|\boldsymbol{l}| \equiv (l^2 + m^2)^{1/2}$ identifying the distance to the origin in the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$. In particular, $\hat{\boldsymbol{s}}_0$ bears by definition the components $(0, 0, 1)$ so that $\boldsymbol{\tau}$ reads as $(\boldsymbol{l}, n(\boldsymbol{l}) - 1)$. The signal and primary beam may also be seen as functions of this two-dimensional vector, respectively, $x(\boldsymbol{l})$ and $A(\boldsymbol{l})$. The invariant measure on the sphere reads as $d\Omega(\boldsymbol{\tau}) \equiv n^{-1}(\boldsymbol{l}) d^2\boldsymbol{l}$, where $d^2\boldsymbol{l} \equiv dldm$ stands for the canonical invariant measure in the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$. The components of \boldsymbol{b}_λ are in standard denoted by (u, v, w) , so that its norm reads as $|\boldsymbol{b}_\lambda| \equiv (u^2 + v^2 + w^2)^{1/2}$ (see Fig. 1). These may also be divided into the two-dimensional vector $\boldsymbol{u} \in \mathbb{R}^2$ with component (u, v) in the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$ and the component w in direction $\hat{\boldsymbol{e}}_3$, i.e. in the pointing direction of the radio interferometer.

The general form (1) of the visibilities reads in terms of these coordinates as

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

where the integration is performed inside the unit disc D^2 on the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$.

In the course of an observation, for each telescope pair, the baseline components u, v and w all follow a sinusoidal dependence in time thanks to the Earth rotation, with specific parameters linked to the parameters of observation. The total number of points (\boldsymbol{u}, w) probed by all telescope pairs of the array during the observation provides some coverage in \mathbb{R}^3 characterizing the interferometer.

2.2 Small field of view and non-negligible w

Note that for a given telescope array, the baseline component w may be seen as a function $w = w(\boldsymbol{u})$ of the two-dimensional vector \boldsymbol{u} with an implicit dependence on the distance between the two telescopes considered. Each visibility may thus be seen as a function of the two-dimensional vectors \boldsymbol{u} : $y(\boldsymbol{u}, w) = y(\boldsymbol{u})$. From this point of view, the points \boldsymbol{u} probed by all telescope pairs during the observation provide some coverage of the plane $(O, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_2)$ also characterizing the interferometer. The two-dimensional vector \boldsymbol{u} associated with each telescope pair actually runs over an arc of ellipse in this plane.

In this context, we make the two assumptions of a small field of view and baselines with a non-negligible component w . First, we consider a standard interferometer with a primary beam of small enough angular support, so that the field of view may be identified with a patch of the tangent plane at $\hat{\boldsymbol{s}}_0$. Technically, this amounts

to assuming that $|\boldsymbol{l}|^2 \ll 1$ on the field of view so that the zero-order expansion $n(\boldsymbol{l}) \simeq 1$ is valid in the invariant measure on the sphere, which therefore identifies with that on the plane: $d\Omega(\boldsymbol{\tau}) = d^2\boldsymbol{l}$. The signal and the primary beam hence become scalar square-integrable functions in the plane $x(\boldsymbol{l}), A(\boldsymbol{l}) \in L^2(\mathbb{R}^2, d^2\boldsymbol{l})$. Secondly, we consider interferometers with large enough values of the $w(\boldsymbol{u})$, so that the same approximation on $n(\boldsymbol{l})$ does not hold for the phase of the imaginary exponential (Thompson et al. 2004). In this case, any term discarded in $n(\boldsymbol{l})$ should be small relative to $w^{-1}(\boldsymbol{u})$. We however assume that the further constraint on the field of view $|\boldsymbol{l}|^4 \ll w^{-1}(\boldsymbol{u})$ holds for all \boldsymbol{u} , so that the second-order approximation $n(\boldsymbol{l}) \simeq 1 - |\boldsymbol{l}|^2/2$ is valid in the imaginary exponential. Under these conditions, each visibility takes the form

$$y(\boldsymbol{u}) = \int_{D^2} A x(\boldsymbol{l}) e^{i\pi w(\boldsymbol{u})|\boldsymbol{l}|^2} e^{-2i\pi \boldsymbol{u} \cdot \boldsymbol{l}} d^2\boldsymbol{l}. \quad (3)$$

In other words, the visibility in \boldsymbol{u} identifies with the value in \boldsymbol{u} of the two-dimensional Fourier transform of a complex signal obtained as the product of the original planar signal $Ax(\boldsymbol{l})$ with a linear chirp modulation $C^{(w(\boldsymbol{u}))}(|\boldsymbol{l}|) \equiv e^{i\pi w(\boldsymbol{u})|\boldsymbol{l}|^2}$. This chirp is characterized by a chirp rate $w(\boldsymbol{u})$ depending on \boldsymbol{u} , and $|\boldsymbol{l}|$ now identifies the distance to the origin in the plane of the signal, i.e. to the point identified by $\hat{\boldsymbol{s}}_0$.

Note that in the further approximation that $w(\boldsymbol{u}) = 0$ for all \boldsymbol{u} , the chirp indeed disappears and relation (4) reduces to the standard van Cittert–Zernike theorem. This theorem states that the visibilities measured identify with the two-dimensional Fourier transform of the signal multiplied by the primary beam: $y(\boldsymbol{u}) = \tilde{A}x(\boldsymbol{u})$ (Wiaux et al. 2009).

The fact that baselines may have a non-negligible component w is at the origin of important challenges from the computational point of view. This issue is related to the fact that the explicit computation of simulated visibilities in (3) from a signal and the corresponding inverse problem have a large complexity. In the case of baselines with a negligible component $w(\boldsymbol{u})$ at each \boldsymbol{u} , the simple two-dimensional Fourier transform relation between the signal and visibilities allows the use of the fast Fourier transform (FFT), which substantially lightens the computation. This is no longer possible in the case of baselines with a non-negligible component $w(\boldsymbol{u})$ as the chirp modulation is characterized by a chirp rate explicitly dependent on \boldsymbol{u} , and therefore needs to be imposed separately for each visibility. A number of algorithms have been studied in this case, among which are the faceting algorithms as well as the very recent w -projection algorithm (Cornwell, Golap & Bhatnagar 2008). In one word this w -projection algorithm simply computes a unique FFT, after which the visibility at each \boldsymbol{u} is obtained by the evaluation, at this \boldsymbol{u} , of the convolution of the Fourier transform of the signal with that of the chirp with chirp rate $w(\boldsymbol{u})$. Considerations on the extension of the corresponding convolution kernel in the Fourier plane allow us to drastically reduce the computational load in practice. These algorithmic issues are however not our concern here. We only aim at discussing how this natural and compulsory chirp modulation may drastically enhance the quality of signal reconstruction.

2.3 Constant w and spread spectrum phenomenon

We already acknowledged that, in the course of an observation, the two-dimensional vector \boldsymbol{u} associated with each telescope pair runs over an arc of ellipse and the corresponding component w follows a sinusoidal evolution. However, the whole baseline distribution in \mathbb{R}^3 is extremely dependent on the specific configuration of the radio telescope array under consideration. Visibilities from various

interferometers may also be combined, as well as visibilities from the same interferometer with different pointing directions through the mosaicking technique (Thompson et al. 2004). From this point of view, the baseline distributions are rather flexible.

Relying on this flexibility, we assume that all baselines have the same component in the pointing direction: $w(\mathbf{u}) = w$.¹ This assumption of constant component w allows us to discard considerations related to specific interferometers. Moreover, it also allows us to study the impact of a chirp modulation on the basis of a large range of simulations at light computational load. Under this assumption, all visibilities indeed identify with the values of the two-dimensional Fourier transform of a complex signal obtained as the product of the original planar signal $Ax(\mathbf{l})$ and the same linear chirp $C^{(w)}(|\mathbf{l}|) = e^{i\pi w|\mathbf{l}|^2}$ characterized by the chirp rate w independent of \mathbf{u} :

$$y(\mathbf{u}) = \widehat{C^{(w)}Ax}(\mathbf{u}). \quad (4)$$

Note that the multiplication by the chirp modulation amounts to the convolution of the Fourier transform of the chirp with that of the original signal multiplied by the primary beam: $\widehat{C^{(w)}Ax} = \widehat{C^{(w)}} \star \widehat{Ax}$. This convolution inevitably spreads the two-dimensional sample power spectrum of the signal multiplied by the primary beam, i.e. the square modulus of its Fourier transform, while preserving its norm.

In particular, the original signal $x(\mathbf{l})$ and the primary beam $A(\mathbf{l})$ can be approximated by band-limited functions on the finite field of view set by the primary beam itself. A linear chirp $C^{(w)}(|\mathbf{l}|)$ with chirp rate w is characterized by an instantaneous frequency $w\mathbf{l}$ at position \mathbf{l} . On the finite field of view set by the primary beam, it is therefore approximately a band-limited function. In each basis direction, the band limit of the signal after convolution $B^{(C^{(w)}Ax)}$ is the sum of the individual band limits of the original signal multiplied by the primary beam $B^{(Ax)}$ and of the chirp modulation $B^{(C^{(w)})}$: $B^{(C^{(w)}Ax)} = B^{(C^{(w)})} + B^{(Ax)}$. This simple consideration already quantifies the spread spectrum phenomenon associated with the chirp modulation. Note for completeness that the primary beam is by definition limited at much lower frequencies than the signal so that it does not significantly affect its band limit: $B^{(Ax)} = B^{(A)} + B^{(x)} \simeq B^{(x)}$.

2.4 Interferometric inverse problem

The band-limited functions considered are completely identified by their Nyquist–Shannon sampling on a discrete uniform grid of $N = N^{1/2} \times N^{1/2}$ points $\mathbf{l}_i \in \mathbb{R}^2$ in real space with $1 \leq i \leq N$. The sampled signal is denoted by a vector $\mathbf{x} \in \mathbb{R}^N \equiv \{x_i \equiv x(\mathbf{l}_i)\}_{1 \leq i \leq N}$ and the primary beam is denoted by the vector $\mathbf{A} \in \mathbb{R}^N \equiv \{A_i \equiv A(\mathbf{l}_i)\}_{1 \leq i \leq N}$. The chirp is complex and reads as the vector $\mathbf{C}^{(w)} \in \mathbb{C}^N \equiv \{C_i^{(w)} \equiv C^{(w)}(|\mathbf{l}_i|)\}_{1 \leq i \leq N}$. Because of the assumed finite field of view, the functions may equivalently be described by their complex Fourier coefficients on a discrete uniform grid of $N = N^{1/2} \times N^{1/2}$ spatial frequencies \mathbf{u}_i with $1 \leq i \leq N$. This grid is limited at some maximum frequency defining the band limit. Due to the fact that the chirp is complex, the Fourier transform of the

product $C^{(w)}Ax$ does not bear any specific symmetry property in the Fourier plane.

As in Wiaux et al. (2009), we assume that the spatial frequencies \mathbf{u} probed by all telescope pairs during the observation belong to the discrete grid of points \mathbf{u}_i . The Fourier coverage provided by the $M/2$ spatial frequencies probed \mathbf{u}_b , with $1 \leq b \leq M/2$, can simply be identified by a binary mask in the Fourier plane equal to 1 for each spatial frequency probed and 0 otherwise. The visibilities measured may be denoted by a vector of $M/2$ complex Fourier coefficients $\mathbf{y} \in \mathbb{C}^{M/2} \equiv \{y_b \equiv y(\mathbf{u}_b)\}_{1 \leq b \leq M/2}$, possibly affected by complex noise of astrophysical or instrumental origin, identified by the vector $\mathbf{n} \in \mathbb{C}^{M/2} \equiv \{n_b \equiv n(\mathbf{u}_b)\}_{1 \leq b \leq M/2}$. Formally, the measured visibilities may equivalently be denoted by a vector of M real measures $\mathbf{y} \in \mathbb{R}^M \equiv \{y_r\}_{1 \leq r \leq M}$ consisting of the real and imaginary parts of the complex measures, affected by the corresponding real noise values $\mathbf{n} \in \mathbb{R}^M \equiv \{n_r\}_{1 \leq r \leq M}$.

In this discrete setting, the Fourier coverage is in general incomplete in the sense that the number of real constraints M is smaller than the number of unknowns N : $M < N$. An ill-posed inverse problem is thus defined for the reconstruction of the signal \mathbf{x} from the measured visibilities \mathbf{y} :

$$\mathbf{y} \equiv \Phi^{(w)}\mathbf{x} + \mathbf{n} \text{ with } \Phi^{(w)} \equiv \mathbf{MFC}^{(w)}\mathbf{A}, \quad (5)$$

where the matrix $\Phi^{(w)} \in \mathbb{C}^{(M/2) \times N}$ identifies the complete linear relation between the signal and the visibilities. The matrix $\mathbf{A} \in \mathbb{R}^{N \times N} \equiv \{A_{ij} \equiv A_i \delta_{ij}\}_{1 \leq i, j \leq N}$ is the diagonal matrix implementing the primary beam. The matrix $\mathbf{C}^{(w)} \in \mathbb{C}^{N \times N} \equiv \{C_{ij}^{(w)} \equiv C_i^{(w)} \delta_{ij}\}_{1 \leq i, j \leq N}$ is the diagonal matrix implementing the chirp modulation. The unitary matrix $\mathbf{F} \in \mathbb{C}^{N \times N} \equiv \{F_{ij} \equiv e^{-2i\pi \mathbf{u}_i \cdot \mathbf{x}_j / N^{1/2}}\}_{1 \leq i, j \leq N}$ implements the discrete Fourier transform providing the Fourier coefficients. The matrix $\mathbf{M} \in \mathbb{R}^{(M/2) \times N} \equiv \{M_{bj}\}_{1 \leq b \leq M/2, 1 \leq j \leq N}$ is the rectangular binary matrix implementing the mask characterizing the interferometer. It contains only one non-zero value on each line, at the index of the Fourier coefficient corresponding to each of the spatial frequencies probed \mathbf{u} .

In the perspective of the reconstruction of the signal \mathbf{x} , relation (5) represents the measurement constraint. We take a statistical point of view and consider independent Gaussian noise on each real measure y_r . Considering a candidate reconstruction $\bar{\mathbf{x}}$, the residual noise reads as $\bar{\mathbf{n}}^{(w)} \equiv \mathbf{y} - \Phi^{(w)}\bar{\mathbf{x}}$. The residual noise level estimator, defined as twice the negative logarithm of the likelihood associated with $\bar{\mathbf{x}}$, reads as

$$\chi^2(\bar{\mathbf{x}}; \Phi^{(w)}, \mathbf{y}) \equiv \sum_{r=1}^M \left(\frac{\bar{n}_r^{(w)}}{\sigma^{(n_r)}} \right)^2, \quad (6)$$

with $\bar{\mathbf{n}}^{(w)} \in \mathbb{R}^M \equiv \{\bar{n}_r^{(w)}\}_{1 \leq r \leq M}$, and where $\sigma^{(n_r)}$ stands for the standard deviation of the noise component n_r . This noise level estimator follows a chi-square distribution with M degrees of freedom. Typically, this estimator should be minimized by the good candidate reconstruction. The measurement constraint on the reconstruction may be defined as a bound $\chi^2(\bar{\mathbf{x}}; \Phi^{(w)}, \mathbf{y}) \leq \epsilon^2$, with ϵ^2 corresponding to some percentile of the chi-square distribution. Let us note that the expectation value of the χ^2 is equal to its number of degrees of freedom M , while its standard deviation is $(2M)^{1/2}$. In other words, for a large number of degrees of freedom the distribution is extremely peaked around its expectation value. This fact is related to the well-known phenomenon of the concentration of measure (Candès et al. 2006b). As a consequence, the value ϵ^2 is anyway extremely close to M .

In this context, many signals may formally satisfy the measurement constraint. A regularization scheme that encompasses enough prior information on the original signal is needed in order to find

¹ Note for illustration that, for a unique interferometer pointing to the north celestial pole, each telescope pair would exhibit constant w during observation. For particular configurations of telescopes in two east–west linear arrays, all telescope pairs with each telescope belonging to a different array would produce baselines with an identical value of w .

a unique solution. All image reconstruction algorithms will differ through the kind of regularization considered.

3 COMPRESSED SENSING

In this section, we concisely recall the inverse problem for sparse signals considered in the compressed sensing framework. We describe the BP minimization problem set for reconstruction and recall the randomness and incoherence properties for a suitable sensing basis.

3.1 Inverse problem for sparse signals

In the framework of compressed sensing (Candès 2006; Candès et al. 2006a,b; Donoho 2006; Baraniuk 2007; Donoho & Tanner 2009), the signals probed are firstly assumed to be sparse or compressible in some basis. Technically, one considers a real signal identified by its Nyquist–Shannon sampling, denoted by the vector $\mathbf{x} \in \mathbb{R}^N \equiv \{x_i\}_{1 \leq i \leq N}$. A real orthonormal basis $\Psi \in \mathbb{R}^{N \times N} \equiv \{\Psi_{ij}\}_{1 \leq i, j \leq N}$ is also considered, in which the decomposition $\boldsymbol{\alpha} \in \mathbb{R}^N \equiv \{\alpha_i\}_{1 \leq i \leq N}$ of the signal is defined by

$$\mathbf{x} \equiv \Psi \boldsymbol{\alpha}. \quad (7)$$

The signal is said to be sparse or compressible in this basis if it only contains a small number $K \ll N$ of non-zero or significant coefficients α_i , respectively.

Secondly, the signal is assumed to be probed by M real linear measurements denoted by a vector $\mathbf{y} \in \mathbb{R}^M \equiv \{y_r\}_{1 \leq r \leq M}$ in some real sensing basis $\Phi \in \mathbb{R}^{M \times N} \equiv \{\Phi_{rj}\}_{1 \leq r \leq M, 1 \leq j \leq N}$ and possibly affected by independent and identically distributed noise $\mathbf{n} \in \mathbb{R}^M \equiv \{n_r\}_{1 \leq r \leq M}$. This defines an inverse problem

$$\mathbf{y} \equiv \Theta \boldsymbol{\alpha} + \mathbf{n} \text{ with } \Theta \equiv \Phi \Psi \in \mathbb{R}^{M \times N}, \quad (8)$$

where the matrix Θ identifies the sensing basis as seen from the sparsity itself. The number M of constraints is typically assumed to be smaller than the dimension N of the vector defining the signal, so that the inverse problem is ill-posed.

3.2 BP reconstruction

The framework proposes the global BP minimization problem for the signal recovery. This problem regularizes the originally ill-posed inverse problem by an explicit sparsity or compressibility prior on the signal. In the presence of noise, the so-called BP denoise (BP ϵ) problem is the minimization of the ℓ_1 norm $\|\tilde{\boldsymbol{\alpha}}\|_1$ of the coefficients of the signal in the sparsity basis under a constraint on the ℓ_2 norm $\|\tilde{\mathbf{n}}\|_2$ of the residual noise, with $\tilde{\mathbf{n}} \equiv \mathbf{y} - \Theta \tilde{\boldsymbol{\alpha}}$:

$$\min_{\tilde{\boldsymbol{\alpha}} \in \mathbb{R}^N} \|\tilde{\boldsymbol{\alpha}}\|_1 \text{ subject to } \|\mathbf{y} - \Theta \tilde{\boldsymbol{\alpha}}\|_2 \leq \epsilon. \quad (9)$$

Let us recall that the ℓ_1 norm of $\tilde{\boldsymbol{\alpha}}$ is defined as $\|\tilde{\boldsymbol{\alpha}}\|_1 \equiv \sum_{i=1}^N |\tilde{\alpha}_i|$. The ℓ_2 norm of the residual noise is the standard norm of the corresponding vector: $\|\tilde{\mathbf{n}}\|_2 \equiv (\sum_{i=1}^M |\tilde{n}_i|^2)^{1/2}$. In these relations, the notation $|a|$ for a scalar a stands for the complex modulus when applied to a complex number and the absolute value when applied to a real number. The BP ϵ problem is solved by application of non-linear and iterative convex optimization algorithms (Combettes & Pesquet 2007; van den Berg & Friedlander 2008).

Note that the ℓ_2 norm term in the BP ϵ problem is identical to a bound on the χ^2 distribution with M degrees of freedom governing the noise level estimator.

3.3 Randomness and incoherence

Among other approaches (Donoho & Tanner 2009), the theory of compressed sensing defines the explicit restricted isometry property (RIP) that the matrix $\Theta = \Phi \Psi$ should satisfy in order to allow an accurate recovery of sparse or compressible signals (Candès 2006; Candès et al. 2006a,b). In that regard, the issue of the design of the sensing matrix Φ is of course fundamental. The theory offers various ways to design suitable sensing matrices, showing in particular that a small number of measurements is required relative to a naive Nyquist–Shannon sampling: $M \ll N$.

One can actually show that incoherence of Φ with the sparsity or compressibility basis Ψ and randomness of the measurements will ensure that the RIP is satisfied with overwhelming probability, provided that the number of measurements is large enough relative to the sparsity K considered (Candès et al. 2006b; Candès 2006, 2008). In particular, measurements may be performed through a uniform random selection of Fourier frequencies. In this case, the precise condition for the RIP depends on the degree of incoherence between the Fourier basis and the sparsity or compressibility basis. The mutual coherence μ between the orthonormal Fourier and sparsity bases may be defined as the maximum complex modulus of the scalar product between basis vectors of the two bases. If the unit-normed basis vectors corresponding to the lines of \mathbf{F} and the columns of Ψ are denoted by $\{\mathbf{f}_i\}_{1 \leq i \leq N}$ and $\{\boldsymbol{\psi}_j\}_{1 \leq j \leq N}$, respectively, then the mutual coherence between the bases reads as

$$\mu(\mathbf{F}, \Psi) \equiv \max_{1 \leq i, j \leq N} |\mathbf{f}_i \cdot \boldsymbol{\psi}_j|. \quad (10)$$

In other words, this mutual coherence identifies with the maximum complex modulus of the Fourier coefficient values of the sparsity basis vectors. Note that the inverse of the mutual coherence, $\mu^{-1}(\mathbf{F}, \Psi)$, can also be called the mutual incoherence between the Fourier and sparsity bases.

In this context, the RIP is already satisfied for a small number of measurements satisfying the constraint

$$K \leq \frac{cM}{N\mu^2(\mathbf{F}, \Psi) \ln^4 N}, \quad (11)$$

for some constant c . Under this condition,² the BP ϵ reconstruction of signals well approximated by a K sparse signal is shown to be accurate and stable relative to noise as well as relative to compressibility, in the sense of a departure from exact sparsity. Let us emphasize the fact that the mutual coherence plays an essential role in relation (11) as for fixed M the sparsity recovered increases with the mutual incoherence, as $K \propto \mu^{-2}(\mathbf{F}, \Psi)$.

The incoherence is notably maximum between the Fourier basis and the real space basis identified by a sparsity matrix $\Psi \equiv \Delta$ made up of unit spikes: $\mu(\mathbf{F}, \Delta) = N^{-1/2}$. In the continuous limit $N \rightarrow \infty$ the real space basis identifies with a Dirac basis and the maximum incoherence is infinite, corresponding to zero coherence:

$$\lim_{N \rightarrow \infty} \mu(\mathbf{F}, \Delta) = 0. \quad (12)$$

Hence, for a fixed number M of pure Fourier measurements, the sparsity recovered K is maximum for signals sparse in real space.

² Let us also acknowledge the fact that this bound is not tight. Empirical results (Candès & Romberg 2005; Lustig, Donoho & Pauly 2007) suggest that ratios M/K between 3 and 5 already ensure the expected reconstruction quality.

4 SPREAD SPECTRUM UNIVERSALITY

In this section, we define simple astrophysical signals sparse in a dictionary of Gaussian waveforms and explicitly identify our sensing dictionary for radio-interferometric measurements. We compute the theoretical coherence between the sparsity and sensing dictionaries as a function of the chirp rate and of the size of the Gaussian waveforms, which suggests the universality of the spread spectrum phenomenon relative to the sparsity dictionary. We define an observational set-up and describe our simulations and specific reconstruction procedures. We finally expose the results of the analysis and comment on the need for future work along these lines.

4.1 Sparsity and sensing dictionaries

We consider simple astrophysical signals sparse in a dictionary of Gaussian waveforms, all with equal and fixed size given by a standard deviation $t \in \mathbb{R}_+$. These signals are assumed to be probed by radio interferometers with a Gaussian primary beam $A^{(t_0)}$ with a size identified by a standard deviation $t_0 \in \mathbb{R}_+$. The corresponding matrix thus reads as $\mathbf{A}^{(t_0)}$. The value t_0 sets the size of the field of view of interest, which must naturally be larger than the size of the Gaussian waveforms: $t_0 > t$. We consider a small field of view and baselines with a non-negligible constant component w , so that the visibilities measured $y_b = y(\mathbf{u}_b)$ take the form (4). Relying on the previous discussion relative to the flexibility of realistic baseline distributions, we also assume that the spatial frequencies \mathbf{u}_b probed arise from a uniform random selection of Fourier frequencies. As for the assumption of constant w , this allows us to discard considerations related to specific interferometers. It also allows us to place our discussion in a setting which complies directly with the requirement of the theory of compressed sensing for random measurements. Let us however recall that specific deterministic distributions of a low number of linear measurements might in fact also allow accurate signal reconstruction in the context of compressed sensing (Matei & Meyer 2008).

In this context, the interferometric inverse problem (5) simply arises from a uniform random selection of spatial frequencies with a sensing dictionary $\Phi \equiv \Phi^{(w,t_0)} \equiv \mathbf{MFC}^{(w)} \mathbf{A}^{(t_0)}$ also parametrized by the size t_0 of the primary beam. The sparsity dictionary identifying Gaussian waveforms of size t may be denoted by $\Psi \equiv \Psi^{(t)} \equiv \text{mathsf{Gamma}}^{(t)}$. The sensing dictionary as seen from the sparsity dictionary itself therefore reads as $\Theta^{(w,t_0,t)} \equiv \Phi^{(w,t_0)} \Gamma^{(t)} \equiv \mathbf{MFC}^{(w)} \mathbf{A}^{(t_0)} \Gamma^{(t)}$.

Let us emphasize the fact that the sparsity dictionary is obviously not an orthogonal basis. Moreover, the sensing dictionary does no longer correspond exactly to a random selection of vectors in an orthogonal basis. No precise compressed sensing result similar to the bound (11) relating sparsity to mutual coherence between such dictionaries was yet obtained. However, in the line of first results already proved with redundant dictionaries (Rauhut, Schnass & Vandergheynst 2008), one can intuitively conjecture that a bound similar to (11) holds in our context, still exhibiting the same trade-off between sparsity and mutual coherence. In the perspective of assessing the BP ϵ reconstruction quality, it is therefore essential to understand how the mutual coherence between the sensing and sparsity dictionaries depends on the size t of the Gaussian waveforms and on the chirp rate w .

4.2 Theoretical coherence

After normalization of the vectors of the sparsity and sensing dictionaries in the ℓ_2 norm, a simple analytical computation gives the

mutual coherence between the sensing and sparsity dictionaries as

$$\mu \left(\mathbf{FC}^{(w)} \mathbf{A}^{(t_0)}, \Gamma^{(t)} \right) = \frac{2tt_0}{t^2 + t_0^2} \left[1 + \left(\frac{2\pi w t^2 t_0^2}{t^2 + t_0^2} \right)^2 \right]^{-\frac{1}{2}}. \quad (13)$$

Note that one can formally re-organize the decomposition of the matrix $\Theta^{(w,t_0,t)}$ into modified sensing and sparsity dictionaries, respectively, $\tilde{\Phi} \equiv \mathbf{MF}$ and $\tilde{\Psi}^{(w,t_0,t)} \equiv \mathbf{C}^{(w)} \mathbf{A}^{(t_0)} \Gamma^{(t)}$. In this perspective, the mutual coherence $\mu(\mathbf{FC}^{(w)} \mathbf{A}^{(t_0)}, \Gamma^{(t)})$ identifies with the maximum complex modulus of the Fourier coefficient values of the modified sparsity dictionary vectors, which depend on the chirp modulation and on the primary beam. Consequently, the requirement for a lower coherence can be fulfilled by an operation that spreads the two-dimensional sample power spectrum of the signal in the Fourier plane, while preserving its norm.

The relation (13) is valid both for finite N and in the continuous limit $N \rightarrow \infty$. In this continuous limit, analysing the evolution of the coherence as a function of the parameters t , t_0 , and w is very enlightening.

First, we consider finite non-zero values of the size t_0 of the primary beam and of the chirp rate w . When the size t of the Gaussian waveforms tends to zero, the mutual coherence tends to zero:

$$\lim_{t \rightarrow 0} \mu \left(\mathbf{FC}^{(w)} \mathbf{A}^{(t_0)}, \Gamma^{(t)} \right) = 0 \text{ for all } w, t_0 \in \mathbb{R}_+. \quad (14)$$

In the absence of chirp modulation and primary beam, this null value is expected from relation (12) as the modified sparsity dictionary identifies with the Dirac basis. But we also show that the mutual coherence gets to this minimum value independently of the chirp rate w in this limit and less importantly independently of the size of the primary beam t_0 .

Secondly, we consider finite non-zero values of the size t_0 of the primary beam and of the size t of the Gaussian waveforms, with $t < t_0$. In the absence of chirp, i.e. for $w = 0$, the coherence is strictly positive: $\mu(\mathbf{FC}^{(0)} \mathbf{A}^{(t_0)}, \Gamma^{(t)}) = 2tt_0/(t^2 + t_0^2)$. The finite size t of a Gaussian waveform implies a Gaussian profile of its Fourier transform with a standard deviation $(2\pi t)^{-1}$. The larger the t , the smaller the extent of the two-dimensional sample power spectrum in the Fourier plane and the larger the coherence. As discussed in Section 2.3, in the presence of a chirp, i.e. for $w \neq 0$, the spectrum is spread by means of the corresponding convolution in the Fourier plane, while the norm of the signal multiplied by the primary beam is preserved. This spread spectrum phenomenon reduces the coherence which, in light of relation (11), should enhance the quality of the BP ϵ reconstruction in the context of compressed sensing. In the limit of an infinite chirp rate, the coherence also tends to its minimum null value independently of the size t of the Gaussian waveforms:

$$\lim_{w \rightarrow \infty} \mu \left(\mathbf{FC}^{(w)} \mathbf{A}^{(t_0)}, \Gamma^{(t)} \right) = 0 \text{ for all } t, t_0 \in \mathbb{R}_+. \quad (15)$$

This limit provides a strong result in the sense that the incoherence lost by considering a sparsity dictionary of Gaussian waveforms of an arbitrary non-zero size t may be completely recovered, thanks to a chirp modulation with a high enough chirp rate w . Still from relation (11), this result suggests the universality of the spread spectrum phenomenon according to which the quality of the BP ϵ reconstruction can be rendered independent of the sparsity dictionary for a large enough component w of the baselines.

4.3 Observational set-up

The signals x considered are built as the superposition of $K = 10$ Gaussian waveforms, with random positions and a random central




Figure 2. From left to right, maps and graphs are associated with increasing values of the size $t_d \in \{2, 4, 8, 16\}$ of the Gaussian waveforms constituting the sparsity dictionaries $\Gamma^{(t)}$ in which the astrophysical signals considered are defined. The top panels represent original signal samples multiplied by the primary beam. Dark and light regions, respectively, correspond to large and small intensity values. The middle panels represent the one-dimensional sample power spectra in the square of the intensity units (iu^2) and as a function of the radial frequency $L^{1/2}|\mathbf{u}|$ with $0 \leq L^{1/2}|\mathbf{u}| \leq 32\sqrt{2}$ for the original signals multiplied by the primary beam in the above panel (continuous black curve), for the chirp modulation (continuous blue curve) and for the signals spread by the chirp modulation (continuous red curve). The bottom panels represent the SNR of the reconstructions multiplied by the primary beam as a function of the coverage identified by the number of complex visibilities $M/2 \in \{50, 100, 200, 300, 400, 500, 1000\}$. The curves correspond to the $\Delta\text{BP}_\epsilon 0$ reconstructions (dot-dashed black curve), the $\Delta\text{BP}_\epsilon 1$ reconstructions (dot-dashed red curve), the $\Gamma\text{BP}_\epsilon 0$ reconstructions (continuous black curve) and the $\Gamma\text{BP}_\epsilon 1$ reconstructions (continuous red curve). The latter illustrate the spread spectrum universality relative to the sparsity dictionary. Each curve precisely represents the mean SNR over the 30 simulations and the vertical lines identify the error at 1 standard deviation.

value in the interval $[0, 1]$ in some arbitrary intensity units (iu). The signals are sampled as images \mathbf{x} on a grid of $N = 64 \times 64 = 4096$ pixels. Fields of view $L = L^{1/2} \times L^{1/2}$ with $L^{1/2}$ not larger than several degrees of angular opening may be considered for the approximation of a planar signal to be valid. A primary beam $A^{(t_0)}$ with a full width at half-maximum (FWHM) $2t_0(2 \ln 2)^{1/2} = L^{1/2}$ is also considered, centred on the signal and with a unit value at the centre. In the discrete setting defined in Section 2.4, the inverse problem is transparent to the precise value of the field of view so that we do not need to fix it. Observations are considered for $M/2$ complex visibilities with $M/2 \in \{50, 100, 200, 300, 400, 500, 1000\}$. The corresponding coverages of the spatial frequencies \mathbf{u}_i in the Fourier plane up to the accessible band limit on the grids $B^{(N,L)} = N^{1/2}/2L^{1/2}$, are roughly between 1 per cent and 25 per cent. Instrumental noise is also added as independent identically distributed Gaussian noise. The corresponding standard deviation $\sigma^{(n_r)}$, identical for all r with $1 \leq r \leq M$, is set to $(\sigma^{(n_r)})^2 = 10^{-3}(\Sigma(\mathbf{FA}^{(t_0)\mathbf{x}}))^2$, where $\Sigma(\mathbf{FA}^{(t_0)\mathbf{x}})$ stands for the sample standard deviation of the real and imaginary parts of the Fourier transform of the original signal multiplied by the primary beam $\mathbf{FA}^{(t_0)\mathbf{x}}$. Note that this noise measure is independent of the chirp modulation, which preserves the norm of the signal.

The size t of the Gaussian waveforms may be written in terms of a discrete size t_d as $t = t_d L^{1/2}/\pi N^{1/2}$. We consider the values $t_d \in \{2, 4, 8, 16\}$ such that the Gaussian waveforms are well sampled on the grid. For $t_d = 2$ the approximate band limit $B^{(A^{(t_0)\mathbf{x}})}$ of the signal is around the same value as $B^{(N,L)}$, while for $t_d = 16$ it is much smaller. All values also satisfy the constraint that the size of the Gaussian waveforms remains smaller than that of the field of view: $t < t_0$. Original signal samples multiplied by the primary beam are reported in the top panels of Fig. 2, one for each value of t_d considered. The middle panels notably represent, for each value of t_d , the one-dimensional sample power spectrum of the original signal multiplied by the primary beam, i.e. the average of the two-dimensional sample power spectrum on all spatial frequencies \mathbf{u} on the same annulus at fixed radial frequency $|\mathbf{u}| \equiv (u^2 + v^2)^{1/2}$.

The component w may also be written in terms of a discrete component w_d as $w = w_d N^{1/2}/L$. We consider two extreme values. The case of baselines with a negligible component w is identified by $w_d = 0$. The case of baselines with a non-negligible and constant component w is identified by $w_d = 1$, corresponding to a linear chirp modulation with a maximum instantaneous frequency $wL^{1/2}/2 = B^{(N,L)}$, i.e. with an approximate band limit $B^{(C^{(w)})}$ not much larger than $B^{(N,L)}$ (see Fig. 3). In other words, the component w is a factor




Figure 3. Real part (left-hand panel) and imaginary part (right-hand panel) of the chirp modulation with $w_d = 1$ assumed in the simulations. The chirp is sampled on a grid of $4N$ pixels in order to avoid any aliasing artefact. Dark and light regions, respectively, correspond to positive and negative values.

Table 1. Values with three significant figures of the mutual coherence $\mu(\mathbf{FC}^{(w)}\mathbf{A}^{(t_0)}, \Gamma^{(t)})$ between the sparsity and sensing dictionaries, for sizes of the Gaussian waveforms identified by $t_d \in \{2, 4, 8, 16\}$, both in the absence of chirp modulation, i.e. for $w_d = 0$, and for a chirp modulation with $w_d = 1$.

Gaussian size	$t_d = 2$	$t_d = 4$	$t_d = 8$	$t_d = 16$
Coherece for $w_d = 0$	0.0518	0.103	0.205	0.400
Coherece for $w_d = 1$	0.0517	0.102	0.174	0.151

of $2/L^{1/2}$ larger than the maximum value of the components u and v in the plane of the signal, so that for the small fields of view assumed the baselines are strongly aligned with the pointing direction. One can check that this value of w is consistent with the constraint for relations (3) and (4).

Table 1 contains the values of the mutual coherence $\mu(\mathbf{FC}^{(w)}\mathbf{A}^{(t_0)}, \Gamma^{(t)})$ for the values of t_d and w_d considered. These values are computed numerically from the definition (10). A computation from the theoretical relation (13) provides the same results, up to a relative discrepancy of 10 per cent related to the fact that in our setting the primary beam is not completely contained in the field view. These numerical values are thus completely in the line of our discussion of Section 4.2. Additionally, the one-dimensional sample power spectra of the original signals multiplied by the primary beam, as well as those of the chirp modulation and of the signals spread by the chirp modulation are represented in the middle panels of Fig. 2, directly below the corresponding original signal for each value of t_d . These graphs illustrate the spread spectrum phenomenon due to the chirp modulation and the reduction of the mutual coherence intimately related to it. Up to some normalization, the mutual coherence is indeed essentially the maximum value of the square root of the one-dimensional sample power spectrum.

4.4 Simulations and BP_ϵ reconstruction procedures

A number of 30 simulations is generated for each value of t_d and w_d considered. The visibilities are simulated,³ and the signals are

³ In practice, the grids considered allow the sampling at the Nyquist–Shannon rate for signals with band limits up to $B^{(N,L)}$ in each direction. This is not the case for the original signals after the chirp modulation. In order to avoid artificial aliasing effects in the discrete Fourier transform at the level of the frequencies probed, it is essential to introduce an operator increasing the resolution of the grids, by zero-padding in the Fourier plane, between the sparsity and sensing matrices.

reconstructed through the BP_ϵ problem, which is solved by convex optimization.⁴ The quality of reconstruction is analysed in terms of the signal-to-noise ratio (SNR) of the reconstructions multiplied by the primary beam $\text{SNR} \equiv -20 \log_{10}(\Sigma(\mathbf{A}^{(t_0)}(\mathbf{x} - \bar{\mathbf{x}})) / \Sigma(\mathbf{A}^{(t_0)}\mathbf{x}))$, where $\Sigma(\mathbf{A}^{(t_0)}\mathbf{x})$ stands for the sample standard deviation of the original signal $\mathbf{A}^{(t_0)}\mathbf{x}$ and $\Sigma(\mathbf{A}^{(t_0)}(\mathbf{x} - \bar{\mathbf{x}}))$ for that of the discrepancy signal $\mathbf{A}^{(t_0)}(\mathbf{x} - \bar{\mathbf{x}})$.

We actually compare two different settings for the reconstructions. In the first setting, called ΓBP_ϵ , we assume that the Gaussian waveform dictionary with appropriate t is known, and we use it explicitly as sparsity dictionary: $\Psi \equiv \Gamma^{(t)}$. As a consequence, the ΓBP_ϵ problem deals with the best possible sparsity value $K = 10$. The reconstructions in the absence ($w_d = 0$) and in the presence ($w_d = 1$) of the chirp modulation are, respectively, denoted by $\Gamma\text{BP}_\epsilon 0$ and $\Gamma\text{BP}_\epsilon 1$. It is the precise setting in which we just brought up the spread spectrum phenomenon and its universality on the basis of considerations relative to the mutual coherence between the sensing and sparsity dictionaries. In the second setting, called ΔBP_ϵ , we assume that the sparsity dictionary is the real space basis: $\Psi \equiv \Delta$. As a consequence, the ΔBP_ϵ problem deals with the best possible coherence value $\mu(\mathbf{F}, \Delta) = N^{-1/2}$. However, the sparsity computed in real space increases drastically with the size of the Gaussian waveforms, suggesting that the reconstruction quality should clearly decrease when the Gaussian size increases. The reconstructions in the absence ($w_d = 0$) and in the presence ($w_d = 1$) of the chirp modulation are, respectively, denoted by $\Delta\text{BP}_\epsilon 0$ and $\Delta\text{BP}_\epsilon 1$. But the mutual coherence here remains unaffected by the chirp modulation, so that this modulation should fail to enhance the reconstruction quality in this case.

This analysis structure is strongly suggested by our interest in understanding the behaviour of the standard reconstruction algorithms used in radio interferometry. On the one hand, the standard CLEAN algorithm is essentially a Matching Pursuit (MP) algorithm (Mallat & Zhang 1993; Mallat 1998) which works by iterative removal of the so-called dirty beam, i.e. the inverse Fourier transform of the mask, in real space (Högbon 1974; Schwarz 1978; Thompson et al. 2004). It is also known that CLEAN provides reconstruction qualities very similar to what we just called the ΔBP_ϵ approach (Marsh & Richardson 1987; Wiaux et al. 2009). On the other hand, multiscale versions of CLEAN can in principle account for the fact that the signal may be sparser in some multiscale dictionary (Cornwell 2008). Assuming that the signals considered have a very sparse expansion in such dictionaries, the corresponding performances would probably be very close to that of our ΓBP_ϵ approach, provided that the equivalence between MP and BP still holds in a multiscale setting.

4.5 Results

The results of the analysis are reported in the bottom panels of Fig. 2. The curves represent the mean SNR over the 30 simulations and the vertical lines identify the error at 1 standard deviation.

First, for each size t_d of the Gaussian waveforms and each reconstruction procedure, the SNR obviously benefits from an increase

⁴ The algorithm is based on the Douglas–Rachford splitting method (Combettes & Pesquet 2007) in the framework of proximal operator theory (Moreau 1962). Under our constant w assumption, the computation complexity is driven at each iteration by the complexity of the FFT, i.e. $O(N \log N)$. For the small value of N considered, the reconstruction typically takes less than 1 min on a single standard processor. The algorithm therefore easily scales to much larger values of N .

in the number $M/2$ of visibilities measured, corresponding to an increase in the information directly available on the signal.

Secondly, in the absence of chirp it clearly appears that, for each size t_d of the Gaussian waveforms and each number of visibilities measured $M/2$, the $\Gamma\text{BP}_\epsilon 0$ reconstruction exhibits a significantly better SNR than the $\Delta\text{BP}_\epsilon 0$ reconstruction. This suggests that for the signals considered, it is better to optimize the sparsity by accounting for the proper dictionary than to optimize the mutual coherence by formally postulating that the signal lives in real space. It also appears that, for each number of visibilities measured $M/2$, the SNR of both kinds of reconstructions significantly degrades when the size t_d of the Gaussian waveforms increases. This behaviour is due to the fact that either the mutual coherence, in the case of $\Gamma\text{BP}_\epsilon 0$, or the sparsity, in the case of $\Delta\text{BP}_\epsilon 0$, gets farther from their optimal values.

Thirdly, the SNR of the $\Delta\text{BP}_\epsilon 0$ and $\Delta\text{BP}_\epsilon 1$ reconstructions remain undistinguishable from one another at 1 standard deviation, independently of the size t_d of the Gaussian waveforms and of the number of visibilities measured $M/2$. This illustrates the fact that the mutual coherence is already optimal in the absence of chirp modulation, so that any chirp modulation will fail to enhance the reconstruction quality. Let us also emphasize that we have implemented the standard CLEAN algorithm for comparison with this ΔBP_ϵ setting. For each value of t_d and M and for both $w_d = 0$ and $w_d = 1$, the corresponding SNR of reconstruction (not shown in Fig. 2) remains, as expected, undistinguishable from both the $\Delta\text{BP}_\epsilon 0$ and $\Delta\text{BP}_\epsilon 1$ reconstructions at 1 standard deviation.

Fourthly, for any number of visibilities measured $M/2$ and for large enough size t_d of the Gaussian waveforms, the SNR of the $\Gamma\text{BP}_\epsilon 1$ is significantly larger than that of the $\Gamma\text{BP}_\epsilon 0$ reconstruction. This is the spread spectrum phenomenon related to the reduction of the mutual coherence in the presence of the chirp modulation. Moreover, the SNR of the $\Gamma\text{BP}_\epsilon 1$ reconstruction is essentially independent of the sparsity dictionary identified by t_d . This supports very strongly the principle of universality of the spread spectrum phenomenon relative to the sparsity dictionary, in perfect agreement with our theoretical considerations. The reconstruction quality also appears to be more stable around the mean SNR values in the presence of the chirp modulation.

Finally, let us note that our pure considerations on sparsity and mutual coherence led to relation (15), which suggests the spread spectrum universality and associated optimal reconstructions in the ΓBP_ϵ setting in the limit of an infinite chirp rate. In practice though, the SNR of a ΓBP_ϵ reconstruction should saturate at some finite value, and progressively degrade above this value, due to a leakage phenomenon. The larger the w_d , the larger the spreading of the spectrum of the signal. But the extent of the Fourier coverages considered is limited by the band limit $B^{(N,L)}$. Consequently, when the band limit of the signal after chirp modulation significantly exceeds the band limit where the visibility distributions are defined, a significant part of the energy of the signal remains unprobed.

Simulations and reconstructions have actually been performed for a range of chirp modulations with chirp rates between $w_d = 0$ and $w_d = 1$, as well as for a higher chirp rate $w_d = 1.5$, in both the ΓBP_ϵ setting and the ΔBP_ϵ setting. In the ΓBP_ϵ setting, the SNR of reconstruction (not shown in Fig. 2) undergoes a natural continuous increase for values $w_d \leq 1$ from the $\Gamma\text{BP}_\epsilon 0$ curve to the $\Gamma\text{BP}_\epsilon 1$ curve, confirming the spread spectrum phenomenon for all values of t_d and M considered. The saturation of the SNR occurs around $w_d = 1$, and its degradation is for example already significant for a chirp rate $w_d = 1.5$, for $t_d = 2$ and for all values of M considered. In the ΔBP_ϵ setting, the SNR of reconstruction (not shown

in Fig. 2) exhibits no evolution for values $w_d \leq 1$, independently of the values of t_d and M considered, confirming the fact that the chirp modulation has no impact on the mutual coherence. However, the leakage phenomenon obviously also affects the reconstructions. The corresponding degradation is observed in the same conditions as for the ΓBP_ϵ setting.

4.6 Comments

A huge amount of work still needs to be envisaged for analysing the effect of baselines with a non-negligible component w in radio interferometry. We comment here on three important points.

First, all our results should be confirmed for various levels of instrumental noise before stronger conclusions are drawn. In particular, possible implications of the spreading of the signal energy on all spatial frequencies probed due to the chirp modulation should be studied as a function of the noise level on each visibility.

Secondly, we assumed perfect knowledge of a sparsity dictionary made up of simple Gaussian waveforms. A large range of multiscale dictionaries, notably wavelet frames, may be used in which a large variety of natural signals are known to be sparse or compressible. Let us however acknowledge the fact that a non-optimal choice of a sparsity dictionary will of course have an effect on the sparsity and possibly degrade the reconstruction. In this perspective, further analyses should be performed in order to assess the suitability of specific dictionaries.

For signals with sparse or compressible gradients, a total variation (TV) norm may be substituted for the ℓ_1 norm of the coefficients in the sparsity dictionary, in the very definition of the minimization problem. The TV norm of a signal is simply defined as the ℓ_1 norm of the magnitude of its gradient (Rudin, Osher & Fatemi 1992). A theoretical result of exact reconstruction holds for such TV norm minimization problems in the case of Fourier measurements of signals with exactly sparse gradients in the absence of noise (Candès et al. 2006a). But no proof of stability relative to noise and compressibility exists. Such minimization is also accessible through an iterative scheme from convex optimization algorithms (Candès & Romberg 2005). Even though our signals would not primarily be thought to have very sparse gradients, the Gaussian waveforms have in practice well-defined contours. Preliminary reconstruction results using this scheme actually show promising reconstruction qualities.

Thirdly, further analyses should also be performed in order to understand the effect of the chirp modulation for realistic distributions of the baseline components (u , v , w) and, in particular, for a non-constant component w . However, one can already note from the bottom panels of Fig. 2 that, for each size t_d of the Gaussian waveforms, the SNR of the $\Gamma\text{BP}_\epsilon 1$ reconstruction rises very steeply with the number of visibilities measured $M/2$, before saturating around some multiple of the sparsity $K = 10$. This suggests a selection procedure in cases where the total number of visibilities measured is very large relative to the sparsity considered. If a suitable fraction of visibilities is associated with the suitable coverage of frequencies \mathbf{u}_b in the Fourier plane and with a suitable identical value of w , only these visibilities might be retained in the problem. In such a case, our present considerations could apply directly.

Let us also recall that the value $w_d = 1$ requires a strong alignment of the baselines with the pointing direction for small fields of view, possibly corresponding to an unrealistically large component w . In this respect, the $\Gamma\text{BP}_\epsilon 1$ reconstruction results are asymptotic values. However, a given baseline component w will actually correspond to a larger w_d on a larger field of view L and for a given band limit

$B^{(N,L)}$, so that $w_d = 1$ will become more realistic. This is actually the exact opposite argument to the one used for the definition of faceting algorithms, which decompose a given field of view on sub-fields where the effect of w is negligible. In that regard, the extension of our results on wide fields of view on the celestial sphere will be essential (Cornwell et al. 2008; McEwen & Scaife 2008), notably with regard to forthcoming radio interferometers such as the Square Kilometer Array (SKA; Carilli & Rawlings 2004).⁵

5 CONCLUSION

We have focused our attention on radio interferometers with a small field of view and baselines with a non-negligible and constant component in the pointing direction, for which a linear chirp modulation affects the astrophysical signals probed. Considering simple sparse signals made up of Gaussian waveforms, we have discussed the sensitivity of imaging techniques relative to the sparsity dictionary in the context of the theory of compressed sensing. A theoretical computation of the mutual coherence between the sparsity and sensing dictionaries, as well as the results of our numerical simulations, suggests the universality of the spread spectrum phenomenon relative to the sparsity dictionary, in terms of the achievable quality of reconstruction through the BP_ϵ problem.

ACKNOWLEDGMENTS

The authors wish to thank L. Jacques for productive discussions as well as M. J. Fadili for private communication of results on optimization by proximal methods. The authors also thank the reviewer for his valuable comments. YW is Postdoctoral Researcher of the Belgian National Science Foundation (F.R.S.-FNRS). YB a Postdoctoral Researcher funded by the APIDIS European Project.

REFERENCES

Baraniuk R., 2007, *IEEE Signal Process. Mag.*, 24, 118
 Baraniuk R., Steeghs P., 2007, *Proc. IEEE Radar Conf. IEEE Signal Proc. Soc.*, p. 128
 Blythe J. H., 1957, *MNRAS*, 117, 644

Candès E. J., 2006, in Sanz-Solé M., Soria J., Varona J. L., Verdera J., eds, *Proc. Int. Congress Math. Vol. 3. Euro. Math. Soc., Zürich*, p. 1433
 Candès E. J., 2008, *C. R. Acad. Sci.*, 346, 589
 Candès E. J., Romberg J., 2005, preprint (arXiv:math/0411273v1)
 Candès E. J., Romberg J., Tao T., 2006a, *IEEE Trans. Inf. Theory*, 52, 489
 Candès E. J., Romberg J., Tao T., 2006b, *Comm. Pure Appl. Math.*, 59, 1207
 Carilli C., Rawlings S., 2004, *New Astron. Rev.*, 48, 979
 Combettes P. L., Pesquet J.-C., 2007, *IEEE Sel. Top. Signal Process.*, 1, 564
 Cornwell T. J., 2008, *IEEE Sel. Top. Signal Process.*, 2, 793
 Cornwell T. J., Golap K., Bhatnagar S., 2008, *IEEE Sel. Top. Signal Process.*, 2, 647
 Donoho D. L., 2006, *IEEE Trans. Inf. Theory*, 52, 1289
 Donoho D. L., Tanner J., 2009, *J. Am. Math. Soc.*, 22, 1
 Gottesman S. R., Fenimore E. E., 1989, *Appl. Opt.*, 28, 4344
 Herman M. A., Strohmer T., 2009, *IEEE Trans. Signal Process.*, 57, 2275
 Högbom J. A., 1974, *A&AS*, 15, 417
 Lustig M., Donoho D., Pauly J. M., 2007, *Magn. Reson. Med.*, 58, 1182
 McEwen J. D., Scaife A. M. M., 2008, *MNRAS*, 389, 1163
 Mallat S. G., 1998, *A Wavelet Tour of Signal Processing*. Academic Press, San Diego
 Mallat S. G., Zhang Z., 1993, *IEEE Trans. Signal Process.*, 41, 3397
 Marcia R. F., Willett R. M., 2008, *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Process. IEEE Signal Process. Soc.*, p. 833
 Marsh K. A., Richardson J. M., 1987, *A&A*, 182, 174
 Matei B., Meyer Y., 2008, preprint (<http://dsp.rice.edu/cs>)
 Moreau J. J., 1962, *A&A*, 255, 2897
 Naini F. M., Gribonval R., Jacques L., Vandergheynst P., 2009, *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Process. IEEE Signal Process. Soc.*, p. 2877
 Rauhut H., Schnass K., Vandergheynst P., 2008, *IEEE Trans. Inf. Theory*, 54, 2210
 Romberg J., 2008, preprint (<http://dsp.rice.edu/cs>)
 Rudin L. I., Osher S., Fatemi E., 1992, *Physica D*, 60, 259
 Ryle M., Hewish A., 1960, *MNRAS*, 120, 220
 Ryle M., Vonberg D. D., 1946, *Nat*, 158, 339
 Ryle M., Hewish A., Shakeshaft J. R., 1959, *IRE Trans. Antennas Propag.*, 7, 120
 Schwarz U. J., 1978, *A&A*, 65, 345
 Thompson A. R., Moran J. M., Swenson G. W. Jr., 2004, *Interferometry and Synthesis in Radio Astronomy*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim
 van den Berg E., Friedlander M. P., 2008, *SIAM J. Sci. Comput.*, 31, 890
 Wiaux Y., Jacques L., Puy G., Scaife A. M. M., Vandergheynst P., 2009, *MNRAS*, 395, 1733

⁵ <http://www.skatelescope.org/>

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.