PURIFY: a new approach to radio-interferometric imaging

R. E. Carrillo[1]⋆, J. D. McEwen[2,3] and Y. Wiaux[1,4,5,6]

[1]*Institute of Electrical Engineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland*
[2]*Department of Physics and Astronomy, University College London (UCL), London WC1E 6BT, UK*
[3]*Mullard Space Science Laboratory, University College London (UCL), Holmbury St Mary, Surrey RH5 6NT, UK*
[4]*Department of Medical Radiology, University Hospital Center (CHUV) and University of Lausanne (UNIL), CH-1011 Lausanne, Switzerland*
[5]*Department of Radiology and Medical Informatics, University of Geneva (UniGE), CH-1211 Geneva, Switzerland*
[6]*Institute of Sensors, Signals & Systems, Heriot-Watt University, Edinburgh EH14 4AS, UK*

ABSTRACT

In a recent article series, the authors have promoted convex optimization algorithms for radio-interferometric imaging in the framework of compressed sensing, which leverages sparsity regularization priors for the associated inverse problem and defines a minimization problem for image reconstruction. This approach was shown, in theory and through simulations in a simple discrete visibility setting, to have the potential to outperform significantly CLEAN and its evolutions. In this work, we leverage the versatility of convex optimization in solving minimization problems to both handle realistic continuous visibilities and offer a highly parallelizable structure paving the way to significant acceleration of the reconstruction and high-dimensional data scalability. The new algorithmic structure promoted relies on the simultaneous-direction method of multipliers (SDMM), and contrasts with the current major-minor cycle structure of CLEAN and its evolutions, which in particular cannot handle the state-of-the-art minimization problems under consideration where neither the regularization term nor the data term are differentiable functions. We release a beta version of an SDMM-based imaging software written in C and dubbed PURIFY (http://basp-group.github.io/purify/) that handles various sparsity priors, including our recent average sparsity approach SARA. We evaluate the performance of different priors through simulations in the continuous visibility setting, confirming the superiority of SARA.

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

1 INTRODUCTION

Radio interferometry is a powerful technique that allows observation of the radio emission from the sky with high angular resolution and sensitivity, providing valuable information for astrophysics, astrometry and cosmology (Ryle & Vonberg 1946; Blythe 1957; Ryle et al. 1959; Ryle & Hewish 1960; Thompson et al. 2001). The measurement equation for radio interferometry defines an ill-posed linear inverse problem in the perspective of signal reconstruction. Under restrictive assumptions of monochromatic non-polarized imaging on small fields of view (FOV), the measured visibilities relates to Fourier measurements of the observed signal. Next-generation radio telescopes, such as the new LOw Frequency ARray (LOFAR[1]), or the recently upgraded Karl G. Jansky Very Large Array (VLA[2]), or the future Square Kilometer Array (SKA[3]), will

achieve much higher dynamic range than current instruments, also at higher angular resolution. Also, these telescopes will acquire a massive amount of data, thus posing large-scale problems. Classical imaging techniques developed in the field, such as the CLEAN algorithm and its multi-scale variants (Högbom 1974; Bhatnagar & Cornwell 2004; Cornwell 2008), are known to be slow and to provide suboptimal imaging quality (Li et al. 2011; Carrillo et al. 2012). This state of things has triggered an intense research to reformulate imaging techniques for radio interferometry in the perspective of next-generation instruments.

The theory of compressed sensing (CS) introduces a signal acquisition and reconstruction framework that goes beyond the traditional Nyquist sampling paradigm (Donoho 2006; Candès 2006; Baraniuk 2007; Fornasier & Rauhut 2011). Recently, CS and convex optimization techniques have been applied to image deconvolution in radio interferometry (Wiaux et al. 2009a,b; Wenger et al. 2010; McEwen & Wiaux 2011; Li et al. 2011; Carrillo et al. 2012) showing promising results. These techniques promise improved image fidelity, flexibility and computation speed over traditional

---

⋆ E-mail: rafael.carrillo@epfl.ch
[1] http://www.lofar.org/
[2] https://science.nrao.edu/facilities/vla
[3] http://www.skatelescope.org/

approaches. This speed enhancement is crucial for the scalability of imaging techniques to very high dimensions in the perspective of next-generation telescopes. However, CS-based imaging techniques have only been studied for low dimensional discrete visibility coverages. The works in Wiaux et al. (2009a,b) and McEwen & Wiaux (2011) consider idealised random and discrete visibility coverages in order to remain as close to the CS theory as possible. First steps towards more realistic visibility coverages have been taken by Wenger et al. (2010) and Li et al. (2011), who consider coverages due to specific interferometer configurations but which remain discrete. Carrillo et al. (2012) consider variable density sampling patterns, which mimic common generic sampling patterns in radio-interferometric (RI) imaging but also remaining discrete. These preliminary works suggest that the performance of CS reconstructions is likely to hold for more realistic visibility coverages. Therefore, the extension of CS techniques to more realistic continuous interferometric measurements is of great importance.

In the present work, we extend the previously proposed imaging approaches in Wiaux et al. (2009a), Wiaux et al. (2009b) and Carrillo et al. (2012) to handle continuous visibilities and open the door to large-scale optimization problems. We propose a general algorithmic framework based on the simultaneous-direction method of multipliers (SDMM) to solve sparse imaging problems. The proposed framework offers a parallel implementation structure that decomposes the original problem into several small simple problems, hence allowing implementation in multicore architectures or in computer clusters, or on graphics processing units. These implementations provide both flexibility in memory requirements and a significant gain in terms of speed, thus enabling scalability to large-scale problems. SDMM stands in stark contrast with the current major-minor cycle structure of CLEAN and evolutions, which in particular cannot handle the state-of-the-art minimization problems under consideration (Carrillo et al. 2012), where neither the regularization term nor the data term are differentiable functions. We release a beta version of an SDMM-based imaging software written in C and dubbed PURIFY (http://basp-group.github.io/purify/) that handles various sparsity priors, including our recent average sparsity approach SARA (Carrillo et al. 2012), thus providing a new powerful framework for RI imaging. We evaluate the performance of different priors through simulations in the continuous visibility setting. Simulation results confirm the superiority of SARA for continuous Fourier measurements. Even though this beta version of PURIFY is not parallelized, we discuss in detail the extraordinary parallel and distributed optimization potential of SDMM, to be exploited in future versions.

The remainder of the paper is organized as follows. In Section 2, we review the theory of CS briefly. In Section 3, we recall the inverse problem for image reconstruction from RI data and describe the state-of-the-art image reconstruction techniques used in radio astronomy. Section 4 presents the SDMM-based algorithm for RI imaging, which enables the incorporation of any convex sparsity regularization prior. In Section 5 we describe the PURIFY package, including implementation details. Numerical results evaluating the different regularization priors included in PURIFY, in particular SARA, are presented in Section 6. Finally we conclude in Section 7.

2 COMPRESSED SENSING

CS introduces a signal acquisition framework that goes beyond the traditional Nyquist sampling paradigm (Donoho 2006; Candès 2006; Baraniuk 2007; Fornasier & Rauhut 2011), demonstrating that sparse signals may be recovered accurately from incomplete data. Consider a complex-valued signal $\boldsymbol{x} \in \mathbb{C}^N$, assumed to be sparse in some orthonormal basis $\boldsymbol{\Psi} \in \mathbb{C}^{N \times N}$ with $K \ll N$ nonzero coefficients, and also consider the measurement model $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x} + \boldsymbol{n}$, where $\boldsymbol{y} \in \mathbb{C}^M$ denotes the measurement vector, $\boldsymbol{\Phi} \in \mathbb{C}^{M \times N}$ is the sensing matrix and $\boldsymbol{n} \in \mathbb{C}^M$ represents the observation noise. The standard condition $M < N$ characterizes the incompleteness of the data. The most common approach to recover $\boldsymbol{x}$ from $\boldsymbol{y}$ is to solve the following convex problem (Fornasier & Rauhut 2011):

$$\min_{\bar{\boldsymbol{\alpha}} \in \mathbb{C}^N} \|\bar{\boldsymbol{\alpha}}\|_1 \text{ subject to } \|\boldsymbol{y} - \boldsymbol{\Phi}\boldsymbol{\Psi}\bar{\boldsymbol{\alpha}}\|_2 \leqslant \epsilon, \qquad (1)$$

where ϵ is an upper bound on the ℓ_2 norm of the noise and $\|\cdot\|_1$ denotes the ℓ_1 norm of a complex-valued vector. The signal is recovered as $\hat{\boldsymbol{x}} = \boldsymbol{\Psi}\hat{\boldsymbol{\alpha}}$, where $\hat{\boldsymbol{\alpha}}$ denotes the solution to the above problem. Such problems that solve for the representation of the signal in a sparsity basis are known as synthesis-based problems.

The standard theory of CS provides results for the recovery of $\boldsymbol{x}$ from $\boldsymbol{y}$ if $\boldsymbol{\Phi}$ obeys a Restricted Isometry Property (RIP) (Fornasier & Rauhut 2011). A sufficient condition is that M is larger than roughly the signal sparsity: $M > 2K \ll N$. Note that incomplete Fourier measurements, on discrete or continuous spatial frequencies, represent a good sampling approach in this context. In the continuous setting, the theory applies also for $M > N$. It is not strictly "compressed" sensing any more but the inverse problem remains ill-posed. The basic theory also requires $\boldsymbol{\Psi}$ to be orthonormal. However, signals often exhibit better sparsity in an overcomplete dictionary (Gribonval & Nielsen 2003; Bobin et al. 2007; Starck et al. 2010). Therefore recent works have begun to address the case of CS with redundant dictionaries. In this setting the signal $\boldsymbol{x}$ is expressed in terms of a dictionary $\boldsymbol{\Psi} \in \mathbb{C}^{N \times D}$, $N < D$, as $\boldsymbol{x} = \boldsymbol{\Psi}\boldsymbol{\alpha}$, $\boldsymbol{\alpha} \in \mathbb{C}^D$. Rauhut et al. (2008) find conditions on the dictionary $\boldsymbol{\Psi}$ such that the compound matrix $\boldsymbol{\Phi}\boldsymbol{\Psi}$ obeys the RIP to accurately recover $\boldsymbol{\alpha}$ by solving a synthesis-based problem. Note that the problem is now more severely undertermined since the dimensionality of the unknonw has increased from N to D.

As opposed to synthesis-based problems, analysis-based problems recover the signal itself solving:

$$\min_{\bar{\boldsymbol{x}} \in \mathbb{C}^N} \|\boldsymbol{\Psi}^\dagger \bar{\boldsymbol{x}}\|_1 \text{ subject to } \|\boldsymbol{y} - \boldsymbol{\Phi}\bar{\boldsymbol{x}}\|_2 \leqslant \epsilon, \qquad (2)$$

where $\boldsymbol{\Psi}^\dagger$ denotes the adjoint operator of $\boldsymbol{\Psi}$. In this paper the superscript $\dagger$ is used to denote both operator adjoint or conjugate transpose. Candès et al. (2010) provide a theoretical analysis of the ℓ_1 analysis-based problem, extending the standard CS theory to coherent and redundant dictionaries. They provide theoretical stability guarantees based on a general condition of the sensing matrix $\boldsymbol{\Phi}$, coined the Dictionary Restricted Isometry Property (D-RIP). Note that in the case when redundant dictionaries are used, the analysis problem does not increase the dimensionality of the problem as it solves for the signal itself. Empirical and theoretical studies have shown clear advantages of the analysis approach over the synthesis approach for imaging problems (Carrillo et al. 2013). See Nam et al. (2013) and references therein for further discussion of the analysis model.

3 RADIO-INTERFEROMETRIC IMAGING

3.1 Interferometric inverse problem

A radio interferometer is an array of spatially separated antennas that takes measurements of the radio emissions of the sky, the so-called visibilities. The visibility coordinates are given by the relative position between each pair of antennas. The baseline components (u, v, w) are measured in units of the wavelength λ of the incoming signal. The components $\boldsymbol{u} = (u, v)$ specify the planar baseline coordinates, while the third component w is associated with the basis vector of the coordinate pointing towards the center of the FOV of the telescope. The sky brightness distribution x can be described in the same coordinate system as the baseline, with components (l, m, n) where $\boldsymbol{l} = (l, m)$ denotes the coordinates on the image plane and $n(\boldsymbol{l}) = \sqrt{1 - l^2 - m^2}$. The general RI equation for monochromatic non-polarized imaging reads as:

$$y(\boldsymbol{u}) = \int A(\boldsymbol{l}, \boldsymbol{u}) \, x(\boldsymbol{l}) \, e^{-2\pi i \boldsymbol{u} \cdot \boldsymbol{l}} \, \mathrm{d}^2 \boldsymbol{l}, \qquad (3)$$

where $A(\boldsymbol{l}, \boldsymbol{u}) = A'(\boldsymbol{l}, \boldsymbol{u}) \, n^{-1}(\boldsymbol{l})$ and $A'(\boldsymbol{l}, \boldsymbol{u})$ stands for all contributions of direction dependent effects (DDE). Examples of DDEs are the primary beam, which limits the observed FOV, and the w-term $e^{-2\pi i w(n(\boldsymbol{l})-1)}$. This general equation defines a linear inverse problem in the perspective of recovering the intensity signal x from the measured visibilities (Rau et al. 2009). Under the assumptions of small FOV ($n \approx 1$) or when the array is coplanar ($w \approx 0$), each visibility corresponds to the measurement of the Fourier transform of a planar signal at the spatial frequency $\boldsymbol{u}$. This result is known as the van Cittert-Zernike theorem (Thompson et al. 2001). The total number of points $\boldsymbol{u}$ probed by all telescope pairs of the array during the observation provides some incomplete coverage in the Fourier plane, the so-called u-v coverage, characterizing the interferometer.

To recover the source image from incomplete visibility measurements, we pose the inverse problem (3) for a sampled version of the image. 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 in real space. The sampled intensity signal is denoted by the vector $\boldsymbol{x} \in \mathbb{R}^N$. We take M visibilities denoted by the vector $\boldsymbol{y} \in \mathbb{C}^M$, which are related to the discrete image by the following linear model:

$$\boldsymbol{y} = \Phi \boldsymbol{x} + \boldsymbol{n}, \qquad (4)$$

where $\Phi \in \mathbb{C}^{M \times N}$ represents the general linear map from the image space domain to the visibility domain, which defines an ill-posed inverse problem in the perspective of image reconstruction. In the particular case when the visibilities identify with Fourier samples the measurement essentially reduces to a Fourier matrix sampled on M spatial frequencies (see eq. (31) in Section 5). In a realistic continuous visibility setting, one usually has $M > N$ and sometimes $M \gg N$, which will be increasingly the case for next-generation telescopes.

3.2 State-of-the-art of classic imaging algorithms

The most standard image reconstruction algorithm from visibility measurements is called CLEAN, which is a non-linear deconvolution method based on local iterative beam removal (Högbom 1974; Schwarz 1978; Thompson et al. 2001). A sparsity prior on the original signal in real space is implicitly introduced thus already taking advantage of CS theory guarantees. Furthermore, as discussed in Cornwell (2008) and Wiaux et al. (2009a) the CLEAN algorithm

and its variants are examples of the Matching Pursuit algorithm (Mallat & Zhang 1993), which is well known in the CS community. CLEAN can be considered as a steepest descend algorithm to minimize the objective function $\chi^2 = \|\boldsymbol{y} - \Phi \boldsymbol{x}\|_2^2$ subject to an image model regularization (Rau et al. 2009). Most variants operate iteratively in two steps called the major and minor cycles. The major cycle computes the residual image $\boldsymbol{r}^{(t)} = \Phi^\dagger(\boldsymbol{y} - \Phi \boldsymbol{x}^{(t)})$, which is the gradient of the χ^2 objective function at iteration t. The minor cycle regularizes the image update by applying an operator T, which represents a deconvolution of the operator Φ, to the residual image yielding updates of the form

$$\boldsymbol{x}^{(t+1)} = \boldsymbol{x}^{(t)} + \mathsf{T}(\boldsymbol{r}^{(t)}). \qquad (5)$$

A multi-scale version of CLEAN, MS-CLEAN, has also been developed (Cornwell 2008), where the sparsity model is improved by multi-scale decomposition, hence enabling better recovery of the signal. The MS-CLEAN method was shown to perform better than the standard CLEAN, but still suffers from an empirical choice of basis profiles and scales. An adaptive scale pixel decomposition method called ASP-CLEAN was also introduced to improve on multi-scale CLEAN by relying on an adaptive choice of scales (Bhatnagar & Cornwell 2004). ASP-CLEAN models an image as a superposition of atoms in a redundant dictionary parametrized by amplitude, location and scale. Thus, this algorithm can be seen as a Matching Pursuit algorithm with an overcomplete dictionary. Note that these approaches are known to be slow, sometimes prohibitively so. Variants of CLEAN that addresses wide-band effects, or atmospheric effects have also been proposed in the literature (see Rau et al. (2009), Bhatnagar et al. (2008a), Bhatnagar et al. (2013) and references therein).

Another approach to the reconstruction of images from visibility measurements is the Maximum Entropy Method (MEM). In contrast to CLEAN, MEM solves a global optimization problem in which the inverse problem is regularized by the introduction of an entropic prior on the signal, but sparsity is not explicitly required (Cornwell & Evans 1985). In practice, CLEAN and variants have found more widespread application than MEM.

3.3 State-of-the-art of convex imaging algorithms

Reconstruction techniques based on CS and convex optimization have also been proposed. The relationship between CLEAN and ℓ_1 minimization coupled with a Dirac basis was first studied by Marsh & Richardson (1987). The first application of CS and convex optimization to radio interferometry was performed by Wiaux et al. (2009a), where the versatility of the approach and its superiority relative to standard interferometric imaging techniques was demonstrated. It was reported that an ℓ_1 minimization problem of the form of (1) coupled with a Dirac basis yields similar reconstruction quality to CLEAN, while including a positivity constraint in a convex formulation significantly enhances the reconstruction quality relative to CLEAN. The spread spectrum phenomenon associated with the w component on wide FOV observations was shown in Wiaux et al. (2009b) to underpin a significant enhancement of the imaging quality independently of the sparsity basis chosen. These considerations pave the way to potential optimization strategies at the acquisition level in terms of antenna distribution design. A CS approach was developed and evaluated by Wiaux et al. (2010) to recover the signal induced by cosmic strings in the cosmic microwave background. McEwen & Wiaux (2011) generalise the previous CS imaging techniques to a wide FOV, recovering interferometric images defined directly on the sphere, rather than a tangent plane. All

of these works consider uniformly random and discrete visibility coverage in order to remain as close to the CS theory as possible. Wiaux et al. (2010) and McEwen & Wiaux (2011) exploited the fact that many signals in nature are also sparse or compressible in the magnitude of their gradient space, in which case the total variation (TV) minimization problem,

$$\min_{\bar{\boldsymbol{x}} \in \mathbb{C}^N} \|\bar{\boldsymbol{x}}\|_{\mathrm{TV}} \text{ subject to } \|\boldsymbol{y} - \Phi\bar{\boldsymbol{x}}\|_2 \leqslant \epsilon, \qquad (6)$$

has been shown to yield superior reconstruction results. The TV norm is defined as $\|\bar{\boldsymbol{x}}\|_{\mathrm{TV}} = \|\nabla\bar{\boldsymbol{x}}\|_1$, where $\nabla\bar{\boldsymbol{x}}$ denotes the image gradient magnitude (Rudin et al. 1992).

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. The aforementioned works use the following unconstrained synthesis problem:

$$\min_{\bar{\boldsymbol{\alpha}} \in \mathbb{C}^N} \frac{1}{2}\|\boldsymbol{y} - \Phi\Psi\bar{\boldsymbol{\alpha}}\|_2^2 + \lambda\|\bar{\boldsymbol{\alpha}}\|_1, \qquad (7)$$

where λ is a regularization parameter that balances the weight between the fidelity term and the regularization term. Wenger et al. (2010) reports superior reconstruction quality relative to an automatic CLEAN reconstruction and similar results relative to a user-guided CLEAN reconstruction. Li et al. (2011) studied a CS imaging approach based on (7) and the isotropic undecimated wavelet transform, reporting results from discrete simulated coverages of ASKAP. The reconstruction quality of the isotropic undecimated wavelet transform method was reported to be superior to those of CLEAN and MS-CLEAN. Minimization of the problem (7) is done iteratively by a projected gradient algorithm with updates of the form:

$$\boldsymbol{\alpha}^{(t+1)} = \mathsf{S}_\lambda\left(\boldsymbol{\alpha}^{(t)} + \mu\Psi^\dagger\Phi^\dagger(\boldsymbol{y} - \Phi\Psi\boldsymbol{\alpha}^{(t)})\right), \qquad (8)$$

where $\mathsf{S}_\lambda(\cdot)$ is the soft-thresholding operator, which will be defined in Section 4.4. This algorithm can be seen as a major-minor cycle update where the major cycle computes the gradient of the χ^2 data fidelity term and the minor cycle regularizes the solution by applying the soft-thresholding operator.

Carrillo et al. (2012) proposed an imaging algorithm dubbed sparsity averaging reweighted analysis (SARA) based on average sparsity over multiple bases, showing superior reconstruction qualities relative to state-of-the-art imaging methods in the field. A sparsity dictionary composed of a concatenation of q bases, $\Psi = [\Psi_1, \Psi_2, \ldots, \Psi_q]$, with $\Psi \in \mathbb{C}^{N \times D}$, $N < D$, is used and average sparsity is promoted through the minimization of an analysis ℓ_0 prior, $\|\Psi^\dagger\bar{\boldsymbol{x}}\|_0$. The concatenation of the Dirac basis and the first eight orthonormal Daubechies wavelet bases (Db1-Db8) was proposed as an effective and simple candidate for a dictionary in the RI imaging context. See Carrillo et al. (2013) for further discussions on the average sparsity model, the dictionary selection and other applications to compressive imaging.

SARA adopts a reweighted ℓ_1 minimization scheme to promote average sparsity through the prior $\|\Psi^\dagger\bar{\boldsymbol{x}}\|_0$. The algorithm replaces the ℓ_0 norm by a weighted ℓ_1 norm and solves a sequence of weighted ℓ_1 problems where the weights are essentially the inverse of the values of the solution of the previous problem (Candès et al. 2008). The weighted ℓ_1 problem is defined as:

$$\min_{\bar{\boldsymbol{x}} \in \mathbb{R}_+^N} \|\mathsf{W}\Psi^\dagger\bar{\boldsymbol{x}}\|_1 \text{ subject to } \|\boldsymbol{y} - \Phi\bar{\boldsymbol{x}}\|_2 \leqslant \epsilon, \qquad (9)$$

where $\mathsf{W} \in \mathbb{R}^{D \times D}$ denotes the diagonal matrix with positive weights and $\mathbb{R}_+^N$ denotes the positive orthant in $\mathbb{R}^N$, which represents the positivity prior on $\boldsymbol{x}$. Note that problems of the form (6) and (9) involve the minimization of a constrained problem with non-differentiable functions, which rules out smooth optimization techniques and do not fit in the major-minor cycle structure of CLEAN and the projected gradient algorithm. Therefore one must resort to more sophisticated optimization techniques to solve these non-smooth problems.

4   A LARGE-SCALE OPTIMIZATION ALGORITHM

4.1   Proximal splitting methods

Convex optimization problems have many attractive properties, in particular the essential property that any local minimum must be a global minimum and thus there exist efficient methods to solve them. Among convex optimization methods, proximal splitting methods offer great flexibility and are shown to capture and extend several well-known algorithms in a unifying framework. Examples of proximal splitting algorithms include Douglas-Rachford, iterative thresholding, projected Landweber, projected gradient, forward-backward, alternating projections, alternating direction method of multipliers and alternating split Bregman (Combettes & Pesquet 2011). Proximal splitting methods solve optimization problems of the form

$$\min_{\boldsymbol{x} \in \mathbb{R}^N} f_1(\boldsymbol{x}) + \ldots + f_S(\boldsymbol{x}), \qquad (10)$$

where $f_1(\boldsymbol{x}), \ldots, f_S(\boldsymbol{x})$ are convex lower semicontinuous functions from $\mathbb{R}^N$ to $\mathbb{R}$, not necessarily differentiable. Note that any convex constrained problem can be formulated as an unconstrained problem by using the indicator function of the convex constraint set as one of the functions in (10), i.e. $f_k(\boldsymbol{x}) = i_C(\boldsymbol{x})$ where C represents the convex constraint set. The indicator function, defined as $i_C(\boldsymbol{x}) = 0$ if $\boldsymbol{x} \in C$ or $i_C(\boldsymbol{x}) = +\infty$ otherwise, belongs to the class of convex lower semicontinuous functions. Also, note that complex-valued vectors are treated as real-valued vectors with twice the dimension (accounting for real and imaginary parts).

Proximal splitting methods proceed by splitting the contribution of the functions $f_1(\boldsymbol{x}), \ldots, f_S(\boldsymbol{x})$ individually so as to yield an easily implementable algorithm. They are called proximal because each non-smooth function in (10) is incorporated in the minimization via its proximity operator. The proximity operator is an extension of the notion of the set projection operator to more general functions. Let f be a convex lower semicontinuous function from $\mathbb{R}^N$ to $\mathbb{R}$, then the proximity operator of f is defined as:

$$\mathrm{prox}_f(\boldsymbol{x}) \triangleq \arg\min_{\boldsymbol{z} \in \mathbb{R}^N} f(\boldsymbol{z}) + \frac{1}{2}\|\boldsymbol{x} - \boldsymbol{z}\|_2^2. \qquad (11)$$

Typically, the solution to (10) is reached iteratively by successive application of the proximity operator associated with each function. An important feature of proximal splitting methods is that they offer a powerful framework for solving convex problems in terms of speed and scalability of the techniques to very high dimensions. See Combettes & Pesquet (2011) for a review of proximal splitting methods and their applications in signal and image processing.

4.2   Shortcomings of previously used algorithms

The works in Wiaux et al. (2009b), McEwen & Wiaux (2011) and Carrillo et al. (2012) solved problems of the form in (9), whereas Wenger et al. (2010) and Li et al. (2011) solved the unconstrained

problem (7). Unconstrained problems are easier to handle and there exist fast algorithms to solve them, such as the FISTA algorithm (Beck & Teboulle 2009b). However, there is no optimal strategy to fix the regularization parameter even if the noise level is known, therefore constrained problems, such as (9), offer a stronger fidelity term when the noise power is known, or can be estimated *a priori*. Hence, we focus our attention on solving problem (9) efficiently. Wiaux et al. (2009b), McEwen & Wiaux (2011) and Carrillo et al. (2012) used a Douglas-Rachford splitting algorithm (Combettes & Pesquet 2007) to solve (9) in a simple discrete setting. However, in a realistic continuous setting this algorithm presents several shortcomings. In the following we discuss the main limitations of the Douglas-Rachford algorithm.

The Douglas-Rachford splitting algorithm solves the problem by iteratively minimizing the ℓ_1 norm and then projecting the result onto the constraint set $C' = \{\boldsymbol{x} \in \mathbb{C}^N : \|\boldsymbol{y} - \Phi\boldsymbol{x}\|_2 \leqslant \epsilon\} \cap \mathbb{R}_+^N$ until some stopping criteria is achieved. The projection onto the set C' is a hard optimization problem which in itself requires an iterative algorithm such as the generalized forward-backward algorithm. This iterative algorithm requires knowledge of the exact operator norm (maximum singular value) of Φ or at least a closed upper bound to guarantee convergence. In the discrete case the exact operator norm can be computed and the algorithm achieves a fast convergence rate. However, in the continuous case the operator norm is unknown and its estimation poses a new problem. If the estimate of the operator norm is not precise enough, the algorithm takes many sub-iterations to converge. Hence, it would be advantageous to have an algorithm that does not need prior knowledge of the operator norm to achieve a fast convergence rate. Another tenet of the Douglas-Rachford algorithm is that it does not offer a parallel structure, which is a desirable property when solving large scale-problems such as those envisaged for the upcoming telescopes. For these reasons, we propose to use the simultaneous-direction method of multipliers (SDMM) (Combettes & Pesquet 2011) which is also tailored to solve problems of the form of (10) and circumvents the shortcomings of a Douglas-Rachford approach.

4.3 Simultaneous Direction Method of Multipliers (SDMM)

SDMM has two important properties: (i) it does not require differentiability of any of the functions, and (ii) it offers a parallel implementation structure where all the proximity operators can be computed in parallel rather than sequentially (Combettes & Pesquet 2011). Such a parallel structure is useful when implementing the algorithms on multicore architectures or on graphics processing units, thus providing a significant gain in terms of speed and scalability to large-scale problems. SDMM is a generalization of the alternating-direction method of multipliers (Boyd et al. 2010) to a sum of more than two functions. As such, SDMM uses augmented Lagrangian techniques and duality arguments in its derivation. In the following we highlight the main steps in the derivation of SDMM tailored to solve (9).

First, observe that the problem in (9) can be reformulated as in (10) in the following way:

$$\min_{\boldsymbol{x}\in\mathbb{C}^N} f_1(\mathsf{L}_1\boldsymbol{x}) + f_2(\mathsf{L}_2\boldsymbol{x}) + f_3(\mathsf{L}_3\boldsymbol{x}), \qquad (12)$$

where $\mathsf{L}_1 = \Psi^\dagger \in \mathbb{C}^{D\times N}$, $\mathsf{L}_2 = \Phi \in \mathbb{C}^{M\times N}$ and $\mathsf{L}_3 = \mathsf{I} \in \mathbb{R}^{N\times N}$ is the identity matrix. In this formulation, $f_1(\boldsymbol{r}_1) = \|\mathsf{W}\boldsymbol{r}_1\|_1$ for $\boldsymbol{r}_1 \in \mathbb{C}^D$, $f_2(\boldsymbol{r}_2) = i_B(\boldsymbol{r}_2)$ with $B = \{\boldsymbol{r}_2 \in \mathbb{C}^M : \|\boldsymbol{y} - \boldsymbol{r}_2\|_2 \leqslant \epsilon\}$, and $f_3(\boldsymbol{r}_3) = i_C(\boldsymbol{r}_3)$ with $C = \mathbb{R}_+^N$. This

problem is also equivalent to solving

$$\min_{\substack{\boldsymbol{x}\in\mathbb{C}^N, \boldsymbol{r}_1\in\mathbb{C}^D, \\ \boldsymbol{r}_2\in\mathbb{C}^M, \boldsymbol{r}_3\in\mathbb{C}^N}} f_1(\boldsymbol{r}_1) + f_2(\boldsymbol{r}_2) + f_3(\boldsymbol{r}_3) \qquad (13)$$

subject to $\mathsf{L}_i\boldsymbol{x} = \boldsymbol{r}_i$, for $i = 1, 2, 3$.

The augmented Lagrangian associated with (13) is the saddle function

$$\mathcal{L}_\gamma(\boldsymbol{x}, \boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3, \boldsymbol{z}_1, \boldsymbol{z}_2, \boldsymbol{z}_3) = \qquad (14)$$

$$\sum_{i=1}^3 f_i(\boldsymbol{r}_i) + \frac{1}{\gamma}\boldsymbol{z}_i^\dagger(\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i) + \frac{1}{2\gamma}\|\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i\|_2^2,$$

where $\gamma > 0$ is a so-called penalty parameter and $\boldsymbol{z}_1 \in \mathbb{C}^D$, $\boldsymbol{z}_2 \in \mathbb{C}^M$ and $\boldsymbol{z}_3 \in \mathbb{C}^N$ are the dual variables or the Langrange multipliers. SDMM is a primal dual algorithm that proceeds iteratively by first minimizing $\mathcal{L}_\gamma$ with respect to the primal variables, $\boldsymbol{x}, \boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3$, and as second step, solving the dual problem

$$\max_{\boldsymbol{z}_1\in\mathbb{C}^D, \boldsymbol{z}_2\in\mathbb{C}^M, \boldsymbol{z}_3\in\mathbb{C}^N} \mathcal{J}(\boldsymbol{z}_1, \boldsymbol{z}_2, \boldsymbol{z}_3), \qquad (15)$$

where

$$\mathcal{J}(\boldsymbol{z}_1, \boldsymbol{z}_2, \boldsymbol{z}_3) = \min_{\substack{\boldsymbol{x}\in\mathbb{C}^N, \boldsymbol{r}_1\in\mathbb{C}^D, \\ \boldsymbol{r}_2\in\mathbb{C}^M, \boldsymbol{r}_3\in\mathbb{C}^N}} \mathcal{L}_\gamma(\boldsymbol{x}, \boldsymbol{r}_1, \boldsymbol{r}_2, \boldsymbol{r}_3, \boldsymbol{z}_1, \boldsymbol{z}_2, \boldsymbol{z}_3)$$

$$(16)$$

is the dual function. The main difference between SDMM and other primal-dual algorithms is that the optimization with respect to the primal variables is done in an alternating fashion by first minimizing $\mathcal{L}_\gamma$ with respect to $\boldsymbol{x}$ and then with respect to $\boldsymbol{r}_1$, $\boldsymbol{r}_2$, $\boldsymbol{r}_3$. The algorithm is shown to converge to a minimizer of (13). Convergence results of SDMM are based on convergence of the alternating-direction method of multipliers and can be found in Boyd et al. (2010).

The minimizer of $\mathcal{L}_\gamma$ with respect to $\boldsymbol{x}$ with fixed variables $\boldsymbol{r}_i$, $\boldsymbol{z}_i$ is given by

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}\in\mathbb{C}^N} \sum_{i=1}^3 \boldsymbol{z}_i^\dagger(\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i) + \frac{1}{2}\|\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i\|_2^2. \qquad (17)$$

Observe that the above problem is the minimization of a quadratic function, which is convex and differentiable. Therefore, necessary and sufficient optimality conditions are

$$\nabla_{\boldsymbol{x}}\mathcal{L}_\gamma(\boldsymbol{x}^*) = \sum_{i=1}^3 \left[\mathsf{L}_i^\dagger\boldsymbol{z}_i + \mathsf{L}_i^\dagger(\mathsf{L}_i\boldsymbol{x}^* - \boldsymbol{r}_i)\right] = 0 \qquad (18)$$

and the matrix $\mathsf{Q} = \sum_{i=1}^3 \mathsf{L}_i^\dagger\mathsf{L}_i \in \mathbb{C}^{N\times N}$ should be invertible. For our particular problem $\mathsf{Q} = \Phi^\dagger\Phi + \Psi\Psi^\dagger + I$, which is positive-definite and invertible. Solving (18) for $\boldsymbol{x}^*$ yields

$$\boldsymbol{x}^* = \mathsf{Q}^{-1}\sum_{i=1}^3 \mathsf{L}_i^\dagger(\boldsymbol{r}_i - \boldsymbol{z}_i). \qquad (19)$$

The minimization over $\boldsymbol{r}_i$ can be carried out for all i simultaneously since the problems are decoupled. Assume i is fixed and also assume that $\boldsymbol{x}$ and $\boldsymbol{z}_i$ are fixed. Then the minimizer of $\mathcal{L}_\gamma$ with respect to $\boldsymbol{r}_i$ is

$$\boldsymbol{r}_i^* = \arg\min_{\boldsymbol{r}_i\in\mathbb{C}^N} f_i(\boldsymbol{r}_i) + \frac{1}{\gamma}\boldsymbol{z}_i^\dagger(\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i) + \frac{1}{2\gamma}\|\mathsf{L}_i\boldsymbol{x} - \boldsymbol{r}_i\|_2^2. \quad (20)$$

After some algebraic manipulations and adding the term $\frac{1}{2}\boldsymbol{z}_i^H\boldsymbol{z}_i$ to (20) we get

$$\boldsymbol{r}_i^* = \arg\min_{\boldsymbol{r}_i\in\mathbb{C}^N} \gamma f_i(\boldsymbol{r}_i) + \frac{1}{2}\|\boldsymbol{r}_i - (\mathsf{L}_i\boldsymbol{x} + \boldsymbol{z}_i)\|_2^2, \qquad (21)$$

which is nothing but the proximity operator of γf_i applied to $\mathsf{L}_i \boldsymbol{x} + \boldsymbol{z}_i$. Thus, the minimizer with respect to $\boldsymbol{r}_i$ is computed as

$$\boldsymbol{r}_i^* = \text{prox}_{\gamma f_i}(\mathsf{L}_i \boldsymbol{x} + \boldsymbol{z}_i). \qquad (22)$$

The maximization over the dual variables is performed using a gradient ascend method. Again the optimization with respect to $\boldsymbol{z}_i$ can be carried out simultaneously for all i since the problems are decoupled. Thus, for a fixed i the problem becomes

$$\boldsymbol{z}_i^* = \arg\max_{\boldsymbol{z}_i} \mathcal{J} = \arg\max_{\boldsymbol{z}_i} \boldsymbol{z}_i^\dagger(\mathsf{L}_i \boldsymbol{x}^* - \boldsymbol{r}_i^*). \qquad (23)$$

The gradient of $\mathcal{J}$ with respect to $\boldsymbol{z}_i$ is given by $\mathsf{L}_i \boldsymbol{x}^* - \boldsymbol{r}_i^*$. Therefore, the dual ascend method yields updates of the form

$$\boldsymbol{z}_i^{(t)} = \boldsymbol{z}_i^{(t-1)} + \mathsf{L}_i \boldsymbol{x}^* - \boldsymbol{r}_i^*, \qquad (24)$$

for each iteration of the algorithm, where t denotes the iteration variable.

Note that the above described procedure can be easily extended for S functions, thus providing a flexible framework for incorporating additional prior information either in the form of convex constraints or as additional convex penalty functions. The expressions in (19), (22) and (24) constitutes the main iteration steps in our SDMM based solver, which is detailed in the next section.

4.4   Implementation details

The resulting algorithm is summarized in Algorithm 1 where $S = 3$. The algorithm is run for a fixed number of iterations, $T_{\max}$, or until a stopping criteria is met. The algorithm is stopped if the relative variation between the objective function evaluated at successive solutions, $\zeta = |f_1(\mathsf{L}_1 \hat{\boldsymbol{x}}^{(t)}) - f_1(\mathsf{L}_1 \hat{\boldsymbol{x}}^{(t-1)})|/|f_1(\mathsf{L}_1 \hat{\boldsymbol{x}}^{(t-1)})|$, is smaller than some bound $\xi \in (0, 1)$ and if the normalized residual $\nu = \|\boldsymbol{y} - \mathsf{L}_2 \hat{\boldsymbol{x}}^{(t)}\|_2/\epsilon$ is within the interval $[1 - \tau, 1 + \tau]$ for some tolerance $\tau \in (0, 1)$, $\tau \ll 1$. In our implementation we fix $\xi = 10^{-3}$ and $\tau = 10^{-1}$.

Algorithm 1 SDMM

1: Initialize $\gamma > 0$, $\hat{\boldsymbol{x}}^{(0)}$ and $\boldsymbol{z}_i^{(0)} = \boldsymbol{0}$, $i = 1, \ldots, S$.
2: $\boldsymbol{r}_i^{(0)} = \mathsf{L}_i \hat{\boldsymbol{x}}^{(0)}$, $i = 1, \ldots, S$.
3: $\boldsymbol{x}_i^{(0)} = \mathsf{L}_i^\dagger \boldsymbol{r}_i^{(0)}$, $i = 1, \ldots, S$.
4: **for** $t = 1, \ldots, T_{\max}$ **do**
5:     $\hat{\boldsymbol{x}}^{(t)} = \mathsf{Q}^{-1} \sum_{i=1}^{S} \boldsymbol{x}_i^{(t-1)}$.
6:     **for all** $i = 1, \ldots, S$ **do**
7:         $\boldsymbol{r}_i^{(t)} = \text{prox}_{\gamma f_i}(\mathsf{L}_i \hat{\boldsymbol{x}}^{(t)} + \boldsymbol{z}_i^{(t-1)})$.
8:         $\boldsymbol{z}_i^{(t)} = \boldsymbol{z}_i^{(t-1)} + \mathsf{L}_i \hat{\boldsymbol{x}}^{(t)} - \boldsymbol{r}_i^{(t)}$.
9:         $\boldsymbol{x}_i^{(t)} = \mathsf{L}_i^\dagger(\boldsymbol{r}_i^{(t)} - \boldsymbol{z}_i^{(t)})$.
10:    **end for**
11:    **if** $\hat{\boldsymbol{x}}^{(t)}$ meets halting criteria **then**
12:        Break.
13:    **end if**
14: **end for**
15: **return** $\hat{\boldsymbol{x}}^{(t)}$

In the following we detail the computation of the proximity operators used in Algorithm 1. To compute the proximity operator of f_1, let us first define it entrywise as follows: $f_1(\boldsymbol{r}_1) = \|\mathsf{W}\boldsymbol{r}_1\|_1 = \sum_{j=1}^{D} \omega_j |r_{1,j}|$, where $\omega_j = \mathsf{W}_{jj}$ (since W is a diagonal positive matrix) and $|\cdot|$ denotes the norm of a complex number. Since f_1 can be split as the sum of independent components of $\boldsymbol{r}_1$, the proximity operator of $\gamma f_1(\boldsymbol{r}_1)$ is given by

$$\text{prox}_{\gamma f_1}(\boldsymbol{r}_1) = \mathsf{S}_\gamma(\boldsymbol{r}_1) = \{\text{prox}_{\gamma\omega_j |\cdot|}(r_{1,j})\}_{1 \leqslant j \leqslant D}, \qquad (25)$$

where $\text{prox}_{\lambda|\cdot|}$ is the entrywise soft-thresholding operator defined as $\text{prox}_{\lambda|\cdot|}(a) = \frac{a}{|a|}(|a| - \lambda)^+$, with $(\cdot)^+ = \max(0, \cdot)$. The proximity operator of $f_2(\boldsymbol{r}_2) = i_B(\boldsymbol{r}_2)$ is the projector onto the convex set $B = \{\boldsymbol{r}_2 \in \mathbb{C}^M : \|\boldsymbol{y} - \boldsymbol{r}_2\|_2 \leqslant \epsilon\}$, and is computed as

$$\text{prox}_{\gamma f_2}(\boldsymbol{r}_2) = \min(1, \epsilon/\|\boldsymbol{r}_2\|_2)\boldsymbol{r}_2, \qquad (26)$$

which is independent of γ. The proximity operator of $f_3(\boldsymbol{r}_3)$ is the projector onto the positive orthant and is given by

$$\text{prox}_{\gamma f_3}(\boldsymbol{r}_3) = \{(r_{3,j})^+\}_{1 \leqslant j \leqslant N}, \qquad (27)$$

which is also independent of γ. See Combettes & Pesquet (2011) and references therein for derivation of these results.

The bottleneck of Algorithm 1, in terms of computational resources, is the inversion of the matrix Q. To invert this matrix we use the conjugate gradient algorithm (Saad 2003) to solve the system $\mathsf{Q}\hat{\boldsymbol{x}}^{(t)} = \sum_{i=1}^{3} \boldsymbol{x}_i^{(t-1)}$. The conjugate gradient algorithm is an iterative process that involves one matrix multiplication by Q at each iteration. Given that $\mathsf{Q} = \Phi^\dagger \Phi + \Psi\Psi^\dagger + I$, in general, each iteration requires one computation of the sensing operator Φ and its adjoint, and, one computation of the sparsity operator Ψ and its adjoint. If we restrict the algorithm to use Parseval frames, i.e. $\Psi\Psi^\dagger = I$, the computation time can be considerably reduced since now $\mathsf{Q} = \Phi^\dagger \Phi + 2I$. Examples of Parseval frames are orthogonal bases and the concatenation of orthogonal bases used in SARA.

Another important consideration in Algorithm 1 is the choice of the penalty parameter γ. In theory any $\gamma > 0$ guarantees convergence of the algorithm. However, in practice the convergence speed of the algorithm is severely affected by the value of this parameter. As it can be observed from the augmented Lagrangian function (14), small values of γ place a large penalty on violations of primal feasibility, thus enforcing fast convergence of the dual variables $\boldsymbol{z}_i$. Conversely, large values of γ place more weight on the original functions f_i, thus achieving a faster convergence rate on the objective function. Before discussing how to set the value of this parameter note that the proximity operators of f_2 and f_3, (26) and (27), are independent of the value of γ since f_2 and f_3 are indicator functions and the only effect of γ in Algorithm 1 is in the proximity operator of f_1. Therefore, γ should scale with $\Psi^\dagger \boldsymbol{x}^*$, where $\boldsymbol{x}^*$ denotes the true signal. Since $\boldsymbol{x}^*$ is unknown, we propose to set the penalty parameter as $\gamma = \beta\|\Psi^\dagger\Phi^\dagger\boldsymbol{y}\|_\infty$, i.e. a constant times the peak value of the dirty image in the sparsity domain. In our implementation we fix $\beta = 10^{-3}$.

4.5   Parallel and distributed optimization

The SDMM structure offers several degrees of parallelization that can be further exploited. Firstly, the proximity operators can be implemented in parallel providing an acceleration factor of three. Secondly, as can be seen from (25), (26) and (27), the computation of the proximity operators is very simple and could support a high level of parallelization since it mostly involves simple entrywise operations. Finally, in the case of large-scale data problems, i.e. large number of visibilities $M \gg N$, the visibilities can no longer be processed on a single computer but rather in a computer cluster thus requiring a distributed processing of the data for the image reconstruction task. In this distributed scenario the data vector $\boldsymbol{y}$ and the measurement operator can be partitioned into R blocks in the following manner:

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_R \end{bmatrix} \text{ and } \Phi = \begin{bmatrix} \Phi_1 \\ \vdots \\ \Phi_R \end{bmatrix}, \qquad (28)$$

where $\boldsymbol{y}_i \in \mathbb{C}^{M_i}$, $\Phi_i \in \mathbb{C}^{M_i \times N}$ and $M = \sum_{i=1}^{R} M_i$. Each $\boldsymbol{y}_i$ is modelled as $\boldsymbol{y}_i = \Phi_i \boldsymbol{x} + \boldsymbol{n}_i$, where $\boldsymbol{n}_i \in \mathbb{C}^{M_i}$ denotes the noise vector.

With this partition the optimization problem in (9) can be rewritten as

$$\min_{\bar{\boldsymbol{x}} \in \mathbb{R}_+^N} \|\mathsf{W}\Psi^\dagger \bar{\boldsymbol{x}}\|_1 \text{ subject to } \|\boldsymbol{y}_i - \Phi_i \bar{\boldsymbol{x}}\|_2 \leqslant \epsilon_i, i = 1, \dots, R, \quad (29)$$

where each ϵ_i is an appropriate bound for the ℓ_2 norm of the noise term $\boldsymbol{n}_i$. Observe that (29) can be solved by SDMM (Algorithm 1) if we reformulate the problem as

$$\min_{\boldsymbol{x} \in \mathbb{C}^N} f_1(\mathsf{L}_1 \boldsymbol{x}) + \dots + f_S(\mathsf{L}_S \boldsymbol{x}), \quad (30)$$

with $S = R + 2$. In this formulation f_1 and f_2 denote the ℓ_1 sparsity term and the positivity constraint respectively, and f_3 to f_S denote the R data fidelity constraints. Thus $\mathsf{L}_1 = \Psi^\dagger$, $\mathsf{L}_2 = \mathsf{I}$ and $\mathsf{L}_{i+2} = \Phi_i$ for $i = 1, \dots, S$. Note that steps 7 to 9 in Algorithm 1 can be computed in parallel for each i. The advantages of this distributed optimization approach are: (i) the visibilities $\boldsymbol{y}_i$ and the measurement operators Φ_i are local to each node in the cluster, therefore the memory requirements are distributed among R nodes, with a data dimensionality $M_i \ll M$; (ii) the measurement operators Φ_i, and their adjoint, are applied locally at each node thus distributing the processing load, for acceleration of the reconstruction process; (iii) the central processing node, where the global update $\hat{\boldsymbol{x}}^{(t)} = \mathsf{Q}^{-1} \sum_{i=1}^{S} \boldsymbol{x}_i^{(t-1)}$ is computed, and the parallel nodes, where the local updates $\boldsymbol{x}_i^{(t-1)}$ are computed, only need to exchange information of the size of the image vector at each iteration rather than of the size of the visibilities, thus alleviating the communication requirements to transfer information between nodes. Note that the composite operator $\Phi^\dagger \Phi$, needed in the conjugate gradient solver for the global update, can be applied in parallel by each node since $\Phi^\dagger \Phi = \sum_{i=1}^{R} \Phi_i^\dagger \Phi_i$. Although this approach would distribute the processing load of the conjugate gradient step into the parallel nodes, it would incur in a communication overhead since each parallel node needs to communicate its result at each iteration of the conjugate gradient algorithm. One approach that can be used to avoid this situation is to precompute and store the composite operator $\Phi^\dagger \Phi$ in the central processing node. The aforementioned distributed optimization approach could be very appealing for next-generation telescopes where massive amounts of data are acquired. These distributed optimization ideas are not implemented in the beta version of PURIFY, discussed in Section 5, and are the subject of ongoing work.

5   THE PURIFY PACKAGE

PURIFY[4] is a collection routines written in C that implements different tools for RI imaging including file handling (for both visibilities and fits images), implementation of the measurement operator and set-up of the different optimization problems used for image deconvolution. The code calls the generic Sparse OPTimization (SOPT[5]) package, which is also written in C, to solve the imaging optimization problems. In the following we describe the different

features included in PURIFY and SOPT. Note that the name PURIFY has no other meaning than that of a powerful alternative to CLEAN.

The optimization problems solved by SOPT within the SDMM structure are: (i) the weighted ℓ_1 minimization problem in (9) and (ii) the weighted TV minimization problem similar to (6) but with the TV norm replaced by a by a weighted TV norm defined as $\|\bar{\boldsymbol{x}}\|_{\mathrm{WTV}} = \|\mathsf{W}\nabla \bar{\boldsymbol{x}}\|_1$ where W is a matrix with positive weights applied to the image gradient. The non-reweighted problems can be solved just by setting the weight matrix to the identity matrix. In the case of the reweighted TV problem $f_1(\boldsymbol{x}) = \|\bar{\boldsymbol{x}}\|_{\mathrm{WTV}}$, with the proximity operator computed using the fast first order iterative method described in Beck & Teboulle (2009a). For the ℓ_1 problems a set of different dictionaries is supported, including: the Dirac basis, the Daubechies wavelets family and the concatenation of any of these bases.

For the measurement operator, PURIFY implements a non-uniform FFT that maps a discrete image into continuous visibilities (Greengard & Lee 2004). The operator is defined as

$$\Phi = \mathsf{GFDZB}. \quad (31)$$

The matrix $\mathsf{B} \in \mathbb{R}^{N \times N}$ is the diagonal matrix implementing the primary beam. The operator $\mathsf{Z} \in \mathbb{R}^{N' \times N}$ denotes the zero padding operator with $N' = kN$ and $k \geqslant 2$ needed to compute the discrete Fourier transform of $\boldsymbol{x}$ on an oversampled grid and achieve higher accuracy. The unitary matrix $\mathsf{F} \in \mathbb{C}^{N' \times N'}$ denotes the discrete Fourier transform. The matrix $\mathsf{G} \in \mathbb{R}^{M \times N'}$ represents a convolutional interpolation operator to model the map from a discrete frequency grid onto the continuous plane so that the FFT can be used to implement F. PURIFY supports a Gaussian kernel in the frequency domain with a compact support, but support for other convolutional interpolation kernels can easily be included. Due to the kernel's compact support, the matrix G is highly sparse therefore allowing fast matrix-vector multiplications. The operator $\mathsf{D} \in \mathbb{R}^{N' \times N'}$ is a diagonal matrix that in practice implements a discrete version of the reciprocal of the inverse Fourier transform of the interpolation kernel, i.e. $d = 1/\hat{g}$, where $\hat{g}$ denotes the inverse Fourier transform of the continuous interpolation kernel. The idea behind this procedure is to undo the effects of the convolution by the interpolation kernel in the frequency domain by dividing by the inverse Fourier transform of the interpolation kernel in the spatial domain. This operator and its adjoint are implemented in the package. Although the current version of PURIFY only supports the Gaussian kernel, other interpolation kernels, such as prolate spheroidal wave functions (Thompson et al. 2001), will be incorporated in future versions.

Also note that our framework can easily incorporate DDEs, in particular the w-component effect, as additional convolution kernels in the frequency plane entering the matrix G. Again, compact support of those kernels will ensure sparsity of G, in turn ensuring its necessary fast implementation. This represents an alternative to the w-projection and the A-projection algorithms (Bhatnagar et al. 2008b,a). See Wolz et al. (2013) for first steps in these directions.

Careful attention has been paid to the design of the interfaces of PURIFY. The solvers receive the measurement operators as pointers to functions implementing the forward and adjoint operators with a generic signature, thus other measurements operators can easily be used. Weighting matrices, such as complex antenna gains and natural or uniform weighting matrices, are not supported in the current implementation but their incorporation into the measurement operator is straightforward. The same philosophy is adopted for the sparsity operators allowing the incorporation of

---

[4] Package available at http://basp-group.github.io/purify/.
[5] Package available at http://basp-group.github.io/sopt/.

any sparsity dictionary. These interfaces will facilitate direct integration with standard packages for interferometric imaging such as CASA[6].

The current version of SOPT does not exploit the parallel structure of SDMM. Firstly, the proximity operators are implemented in a serial manner rather than in parallel. Secondly, the computation of each proximity operators is implemented serially rather than in parallel thus not exploiting its separable structure. The only parallel structure that is exploited is the implementation of the sparsity averaging operator in SARA, i.e. each decomposition on the basis in the operator are computed in parallel. Therefore the highly redundant dictionary in SARA has an implementation as fast as a single orthonormal basis, which already represents a significant advantage. As discussed in Section 4.4, the computation of the measurement operator Φ is a major bottleneck for very high dimensional problems. In this case the measurement operator Φ can be parallelized by implementing a parallel matrix-vector product for the sparse matrix G, e.g. partitioning G into several blocks G_i as done in (28) for Φ. Similar strategies might be adopted for the sparsity operator Ψ. As discussed in Section 4.4 the global update $\hat{x}^{(t)} = Q^{-1} \sum_{i=1}^{S} x_i^{(t-1)}$ is the main bottleneck of the algorithm. One approach that could be implemented here is to precompute and store the sparse matrix $G^\dagger G = \sum_{i=1}^{R} G_i^\dagger G_i$ to accelerate the conjugate gradient solver[7]. These optimizations are the subject of ongoing work.

6 SIMULATIONS AND RESULTS

In this section we illustrate the performance of the imaging algorithms implemented in PURIFY by recovering well known test images from simulated continuous frequency visibilities. The test images used in all simulations are M31, based on a HII region in the M31galaxy, and 30Dor, the 30 Doradus in the Large Magellanic Cloud. These images present different compact and extended structures thus being good candidates to evaluate different regularization priors. Figure 1 shows the 256×256 discrete models of M31 (left) and 30Dor (middle) used as ground truth images[8].

For our evaluation we compare constrained ℓ_1 and TV minimization problems, as well as their reweighted versions, in terms of reconstruction quality and computation time. For the ℓ_1 problems we study three different dictionaries Ψ in (9): the Dirac basis, the Daubechies 8 wavelet basis and the Dirac-Db1-Db8 concatenation highlited for the SARA algorithm in Section 3.3. The associated algorithms are respectively denoted BP, BPDb8 and BPSA for the non-reweighted case. The reweighted versions are respectively denoted RWBP, RWBPDb8 and SARA. We also study the TV minimization problem in (6) with the additional constraint that $\bar{x} \in \mathbb{R}_+^N$, denoted as TV, and its reweighted version, denoted as RWTV. Recall that ℓ_1 minimization with a Dirac basis yields reconstruction qualities similar to CLEAN, thus we use BP as a proxy for CLEAN. Also, we use BPDb8 as a proxy for MS-CLEAN reconstruction quality since Li et al. (2011) reported that the isotropic undecimated wavelet transform outperformed MS-CLEAN and Carrillo et al. (2012) reported that BPDb8 outperformed the isotropic undecimated wavelet transform in the discrete setting.

We use as reconstruction quality metric the signal to noise ratio (SNR), which is defined as:

$$\text{SNR} = 20 \log_{10} \left(\frac{\|x\|_2}{\|x - \hat{x}\|_2} \right) \qquad (32)$$

where x and $\hat{x}$ denote the the original image and the estimated image respectively. The visibilities are corrupted by complex Gaussian noise with a fixed input SNR set to 30 dB. The input SNR is defined as $\text{ISNR} = 20 \log_{10}(\|y_0\|_2/\|n\|_2)$, where y_0 identifies the clean measurement vector. Assuming visibilities corrupted by i.i.d. complex Gaussian noise with variance σ_n, the bound on the ℓ_2 norm term in (9), ϵ, is identical to a bound on a χ^2 distribution with $2M$ degrees of freedom. Therefore, we set this bound as $\epsilon^2 = (2M + 4\sqrt{M})\sigma_n^2/2$, where $\sigma_n^2/2$ is the variance of both the real and imaginary parts of the noise. This choice provides a likely bound for $\|n\|_2$ (Carrillo et al. 2012). We use the measurement operator described in (31) with $B = I$ and an oversampling factor $k = 2$.

The first experiment in this section considers incomplete visibility coverages generated by random variable density sampling profiles. Such profiles are characterized by denser sampling at low spatial frequencies than at high frequencies. This choice mimics common generic sampling patterns in radio interferometry. In order to make the simulated coverages more realistic we suppress the $(0, 0)$ component of the Fourier plane from the measured visibilities. This generic profile approach allows us to make a thorough study of the reconstruction quality of the imaging algorithms with a large numbers of simulations for arbitrary number of visibilities and without concern for various telescope configurations. We vary the number of visibilities from $M = 0.2N$ to $M = 2N$. Reconstruction results for M31 and 30Dor are reported in the top and bottom rows of Figure 2 respectively. Average values over 30 simulations and associated one standard deviation error bars are reported for all plots.

The left panel of Figure 2 shows SNR results for M31 (top) and 30Dor (bottom). The results show that SARA outperforms all other methods in reconstruction quality for both images. This confirms previous results reported by Carrillo et al. (2012) in the discrete case now for the more realistic continuous Fourier setting, including the case when $M > N$. Interestingly, BPSA shows the best reconstruction quality over all non-reweighted methods for both images. The results for M31, which exhibits a compact support with some extended structures, show that the second best method is RWBPDb8 having SNRs at most 4 dB below SARA. The results for 30Dor, which is a more complicated image with both extended structures and compact structures, show that TV and RWTV offer a good model for continuous extended structures achieving SNRs at most 2 dB below SARA. Note that BP and its reweighted version do not achieve good results for this image, as expected since the Dirac basis is not a good model for extended structures, achieving SNRs at least 4 dB below all other methods for coverages above $M = 0.2N$.

Computation times, on a 2.4 GHz Xeon quad core and using the current non-optimized software version, are reported in the right panel of Figure 2 for M31 (top) and 30Dor (bottom). As expected the reweighted methods are most costly having reconstruction times ranging from tens of minutes for $M = 0.2N$ to one hour for $M = 2N$. Even though the concatenation of bases in SARA makes the algorithm structure more costly in theory, the parallel implementation of the bases in SARA yields a competitive algorithm in terms of computation time. In fact, the results show that RWBPDb8, with a single wavelet basis, is the slowest
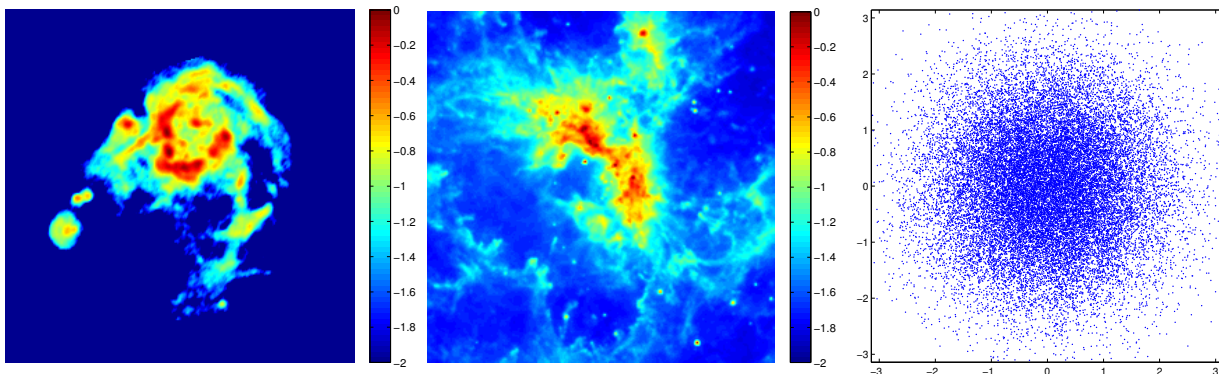
---

[6] http://casa.nrao.edu/.
[7] Note that Sullivan et al. (2012) also proposed to precompute $G^\dagger G$ to accelerate a CLEAN-based algorithm.
[8] Available at http://casaguides.nrao.edu/index.php.

Figure 1. Left and middle panels: original 256×256 test images, M31 (left) and 30Dor (middle), shown in a $\log_{10}$ scale with brightness values in the interval [0.01, 1]. Right panel: Example of a simulated variable density coverage in the Fourier plane ($M = 26374 \approx 0.4N$).
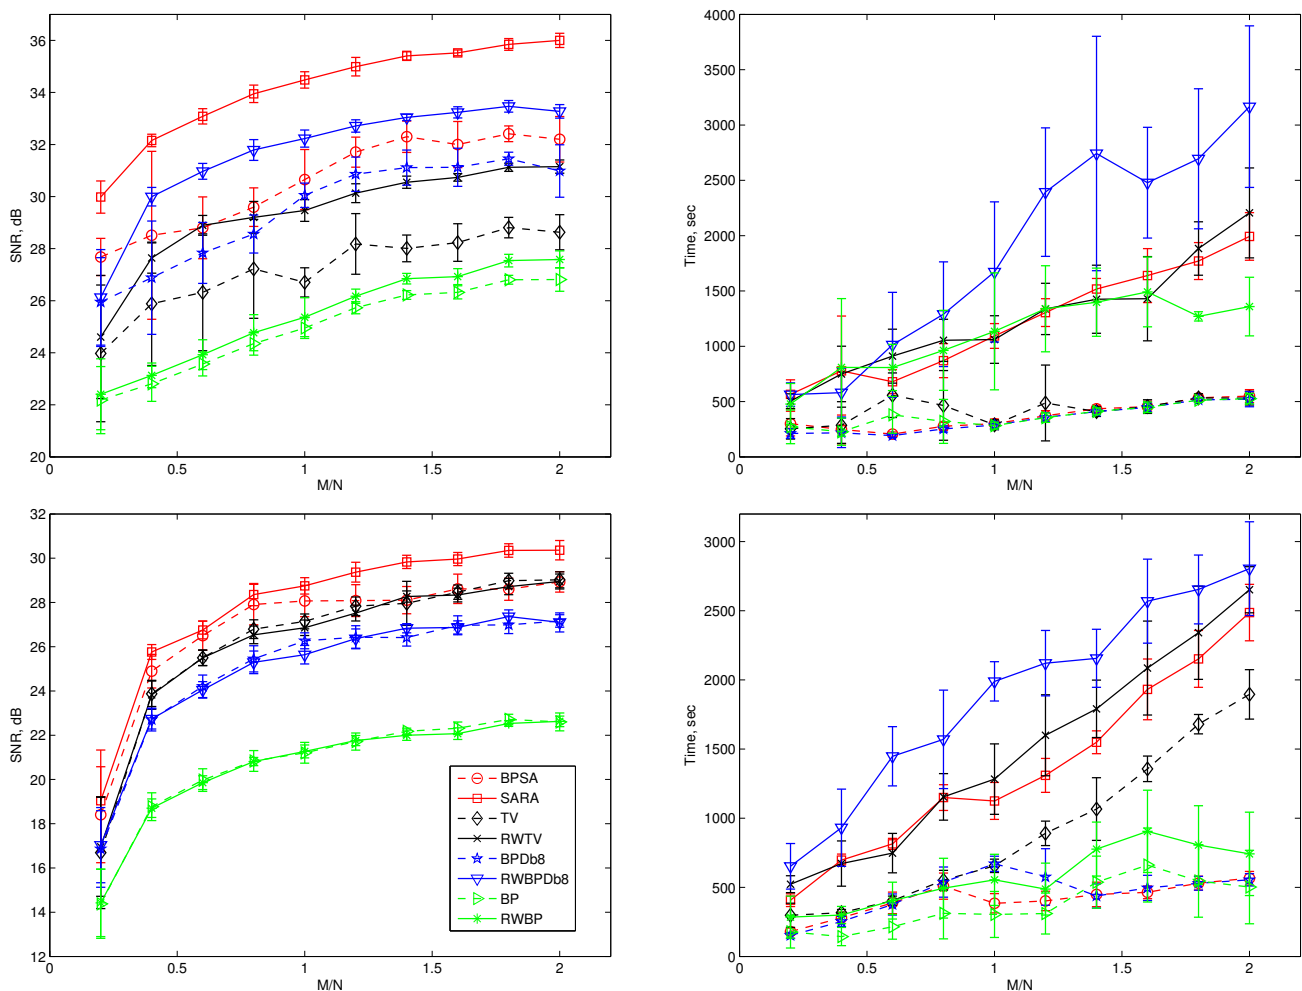

Figure 2. Reconstruction results for M31 (top row) and 30Dor (bottom row) 256×256 test images. Left column: average reconstruction SNR against normalized number of visibilities M/N. Right column: average computation time. Vertical bars identify one standard deviation errors around the mean over 30 simulations. The input SNR is set to 30 dB. The results show that SARA outperforms all other methods in terms of reconstruction quality for both images.

method and the most unstable with respect to convergence rate, as can be observed from the large error bars. This result indicates that RWBPDb8 might need more iterations to achieve convergence than other methods. RWTV reports similar reconstruction times to SARA. The results also show that the non-reweighted methods are

fast, achieving reconstruction times below 10 minutes for all coverages, except for TV in 30Dor which has a similar behaviour as the reweighted methods. An interesting observation is that the reconstruction times scale linearly with the number of visibilities for the reweighted methods. This is due to the fact that the complexity of

the SDMM algorithm is dominated by the cost of solving the linear system at step 5 of Algorithm 1, which needs to apply the sensing operator Φ and its adjoint at every iteration of the conjugate gradient algorithm. Therefore beyond having a fast implementation of Φ, alternative strategies to accelerate the solution of the linear system should be explored such as the use of preconditioned conjugate gradient solvers and faster implementations of the Gram matrix $\Phi^\dagger\Phi$.

Next we present a visual assessment of the reconstruction quality of the different algorithms. Figure 3 and Figure 4 show the results from M31 and 30Dor respectively for a u-v coverage of $M = 26374 \approx 0.4N$ visibilities. The results are shown from top to bottom for SARA, RWBPDb8, RWTV and RWBP respectively. The first column shows the reconstructed images in a $\log_{10}$ scale, the second column shows the error images, defined as $\boldsymbol{x} - \hat{\boldsymbol{x}}$, in linear scale, and, the third column shows the real part of the residual dirty images, defined as the difference between dirty images and dirty images constructed from recovered images, i.e. $\boldsymbol{r} = \Phi^\dagger\boldsymbol{y} - \Phi^\dagger\Phi\hat{\boldsymbol{x}}$, also in linear scale. These images confirm the previous results found by examining recovered SNR levels; SARA yields reconstructed images with fewer artifacts in the background regions and smaller errors in the structured inner regions than the other methods. Interestingly RWBPDb8 yields a nearly flat residual map for 30Dor. However, this does not necessarily translate into a better reconstruction quality as can be observed in the error image. This phenomenon can also be seen in the reconstructed image by RWTV of 30Dor, which shows a small error image compared to RWBPDb8 but showing a residual map with a lot of structures. This highlights the fact that the common criterion of flatness of residual image is not an optimal measure of reconstruction fidelity as emphasized in our previous work (Carrillo et al. 2012).

The last experiment presents an illustration with a realistic radio telescope coverage. We use a simulation of the Arcminute Microkelvin Imager (AMI) (Zwart et al. 2008) array to obtain a u-v coverage with $M = 9413$ points. For this experiment we use a low resolution 128×128 version of M31. The top row in Figure 5 shows the original test image in $\log_{10}$ scale, the u-v coverage and the corresponding dirty image in linear scale. The SNR of the recovered image for each algorithm is as follows: BP (10.7dB), RWBP (SNR=10.9 dB), BPDb8 (11.6 dB), RWBPDb8 (SNR=12.3 dB), TV (10.6 dB), RWTV (10.5 dB), BPSA (12.4 dB) SARA (14.3 dB). The second and third rows in Figure 5 show the reconstructed images along with the corresponding error and residual dirty images images for SARA, RWBPDb8 and RWBP. SARA provides not only a SNR increase but also a significant reduction of visual artifacts relative to all other methods.

7 CONCLUDING REMARKS

In this paper we have proposed an algorithmic framework based on the simultaneous-direction method of multipliers to solve sparse imaging problems in RI imaging. The new algorithm provides a parallel implementation structure, therefore offering an attractive framework to handle continuous visibilities and associated high dimensional problems. A variety of state-of-the-art sparsity regularization priors, including our recent average sparsity approach SARA, as well as discrete and continuous measurement operators are available in the new PURIFY software. Source code for PURIFY is publicly available. Experimental results confirm both the superiority of SARA for continuous Fourier measurements and the

fact that the new algorithmic structure offers a promising path to handle large-scale problems.

In future work we will extend the current PURIFY implementation to take full advantage of the parallel and distributed structure of SDMM as discussed in Section 4.5. We expect that parallel and hardware implementations of the measurement and sparsity operators as well as the proximity operators could achieve drastic accelerations of the algorithms. Also, different strategies will be explored to accelerate the convergence of the conjugate gradient solver, e.g. using preconditioners for the operator Q and precomputing the sparse matrix $\mathsf{G}^\dagger\mathsf{G}$ to avoid multiplications by G and $\mathsf{G}^\dagger$ separately, which involve an intermediate high dimensional vector of length $M > N$, at each iteration of the conjugate gradient solver. Finally, DDEs will be incorporated into PURIFY. Recall that DDEs can easily be included in the matrix G as additional convolution kernels in the frequency plane. Compact support kernels will ensure sparsity of G and a fast matrix-vector multiplication. Integration with standard packages for interferometric imaging, such as CASA, will allow to take advantage of their built-in real data handling and also to have a full comparison with standard algorithms such as MS-CLEAN and ASP-CLEAN.

References

Baraniuk R., 2007, IEEE Signal Process. Mag., 24, 4, 118

Beck A., Teboulle M., 2009a, IEEE Trans. Image Process., 18, 11, 2419

Beck A., Teboulle M., 2009b, SIAM Journal on Imaging Sciences, 2, 1, 183

Bhatnagar S., Cornwell T.J., 2004, A&A, 426, 747

Bhatnagar S., Cornwell T.J., Golap K., Uson J.M., 2008a, A&A, 487, 419

Bhatnagar S., Golap K., Cornwell T.J., 2008b, IEEE J. Sel. Top. Sig. Process., 2, 5, 647

Bhatnagar S., Rau U., Golap K., 2013, ApJ, 770, 91

Blythe J.H., 1957, MNRAS, 117, 644

Bobin J., Starck J.L., Fadili J., Moudden Y., Donoho D., 2007, IEEE Trans. Image Process., 16, 11, 2675

Boyd S., Parikh N., Chu E., Pelato B., Eckstein J., 2010, Foundations and Trends in Machine Learning, 3, 1, 1

Candès E.J., 2006, in *Proceedings, Int. Congress of Mathematics*, Madrid, Spain

Candès E.J., Eldar Y., Needell D., Randall P., 2010, Appl. Comp. Harmonic Anal., 31, 1, 59

Candès E.J., Wakin M., Boyd S., 2008, J. Fourier Anal. Appl., 14, 5, 877

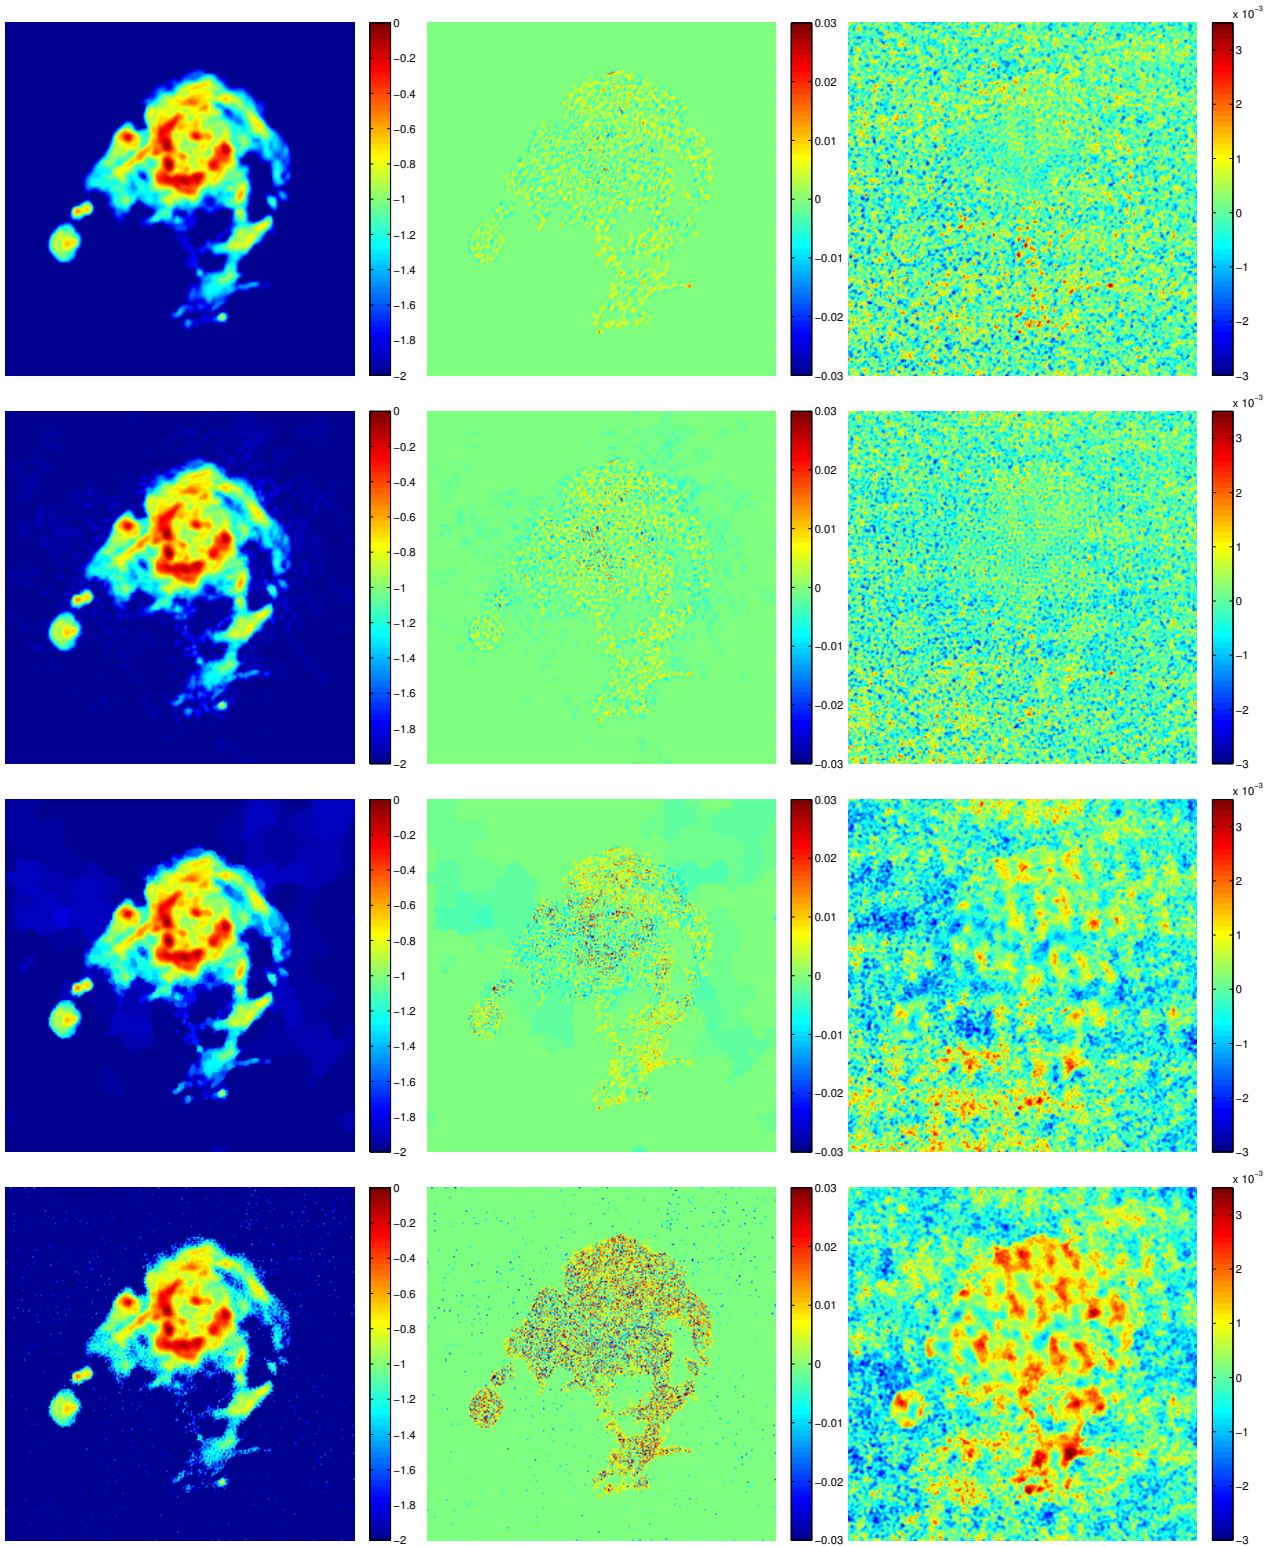**Figure 3.** Reconstruction example of M31 (256×256) for a u-v coverage with $M = 0.4N$ sampling frequencies. The results are shown from top to bottom for SARA (SNR=32.4 dB), RWBPDb8 (SNR=30.6 dB), RWTV (SNR=28.6 dB) and RWBP (SNR=23.4 dB) respectively. The first column shows the reconstructed images in a $\log_{10}$ scale, the second column shows the error images in linear scale, and the third column shows the residual dirty images also in linear scale.
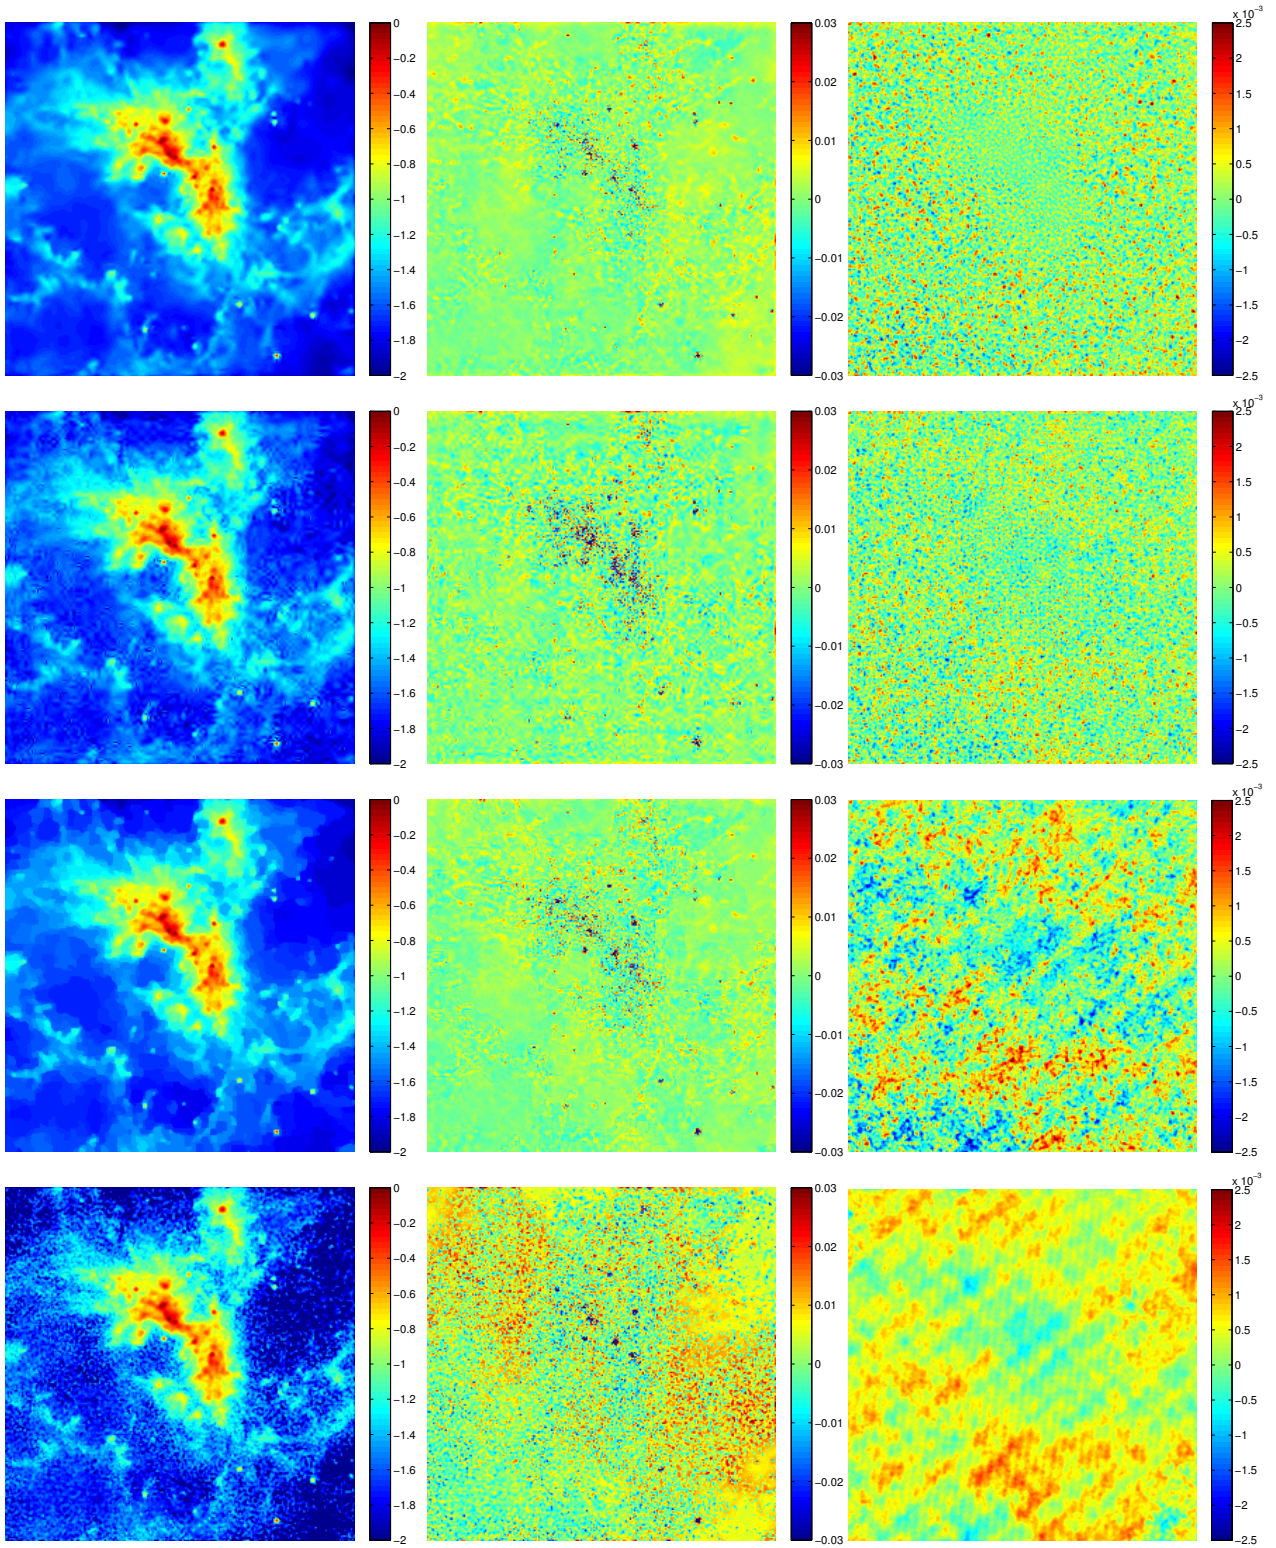
Figure 4. Reconstruction example of 30dor (256×256) for a u-v coverage with $M = 0.4N$ sampling frequencies. The results are shown from top to bottom for SARA (SNR=25.3 dB), RWBPDb8 (SNR=22.6 dB), RWTV (SNR=24.1 dB) and RWBP (SNR=18.8 dB) respectively. The first column shows the reconstructed images in a $\log_{10}$ scale, the second column shows the error images in linear scale, and the third column shows the residual dirty images also in linear scale.
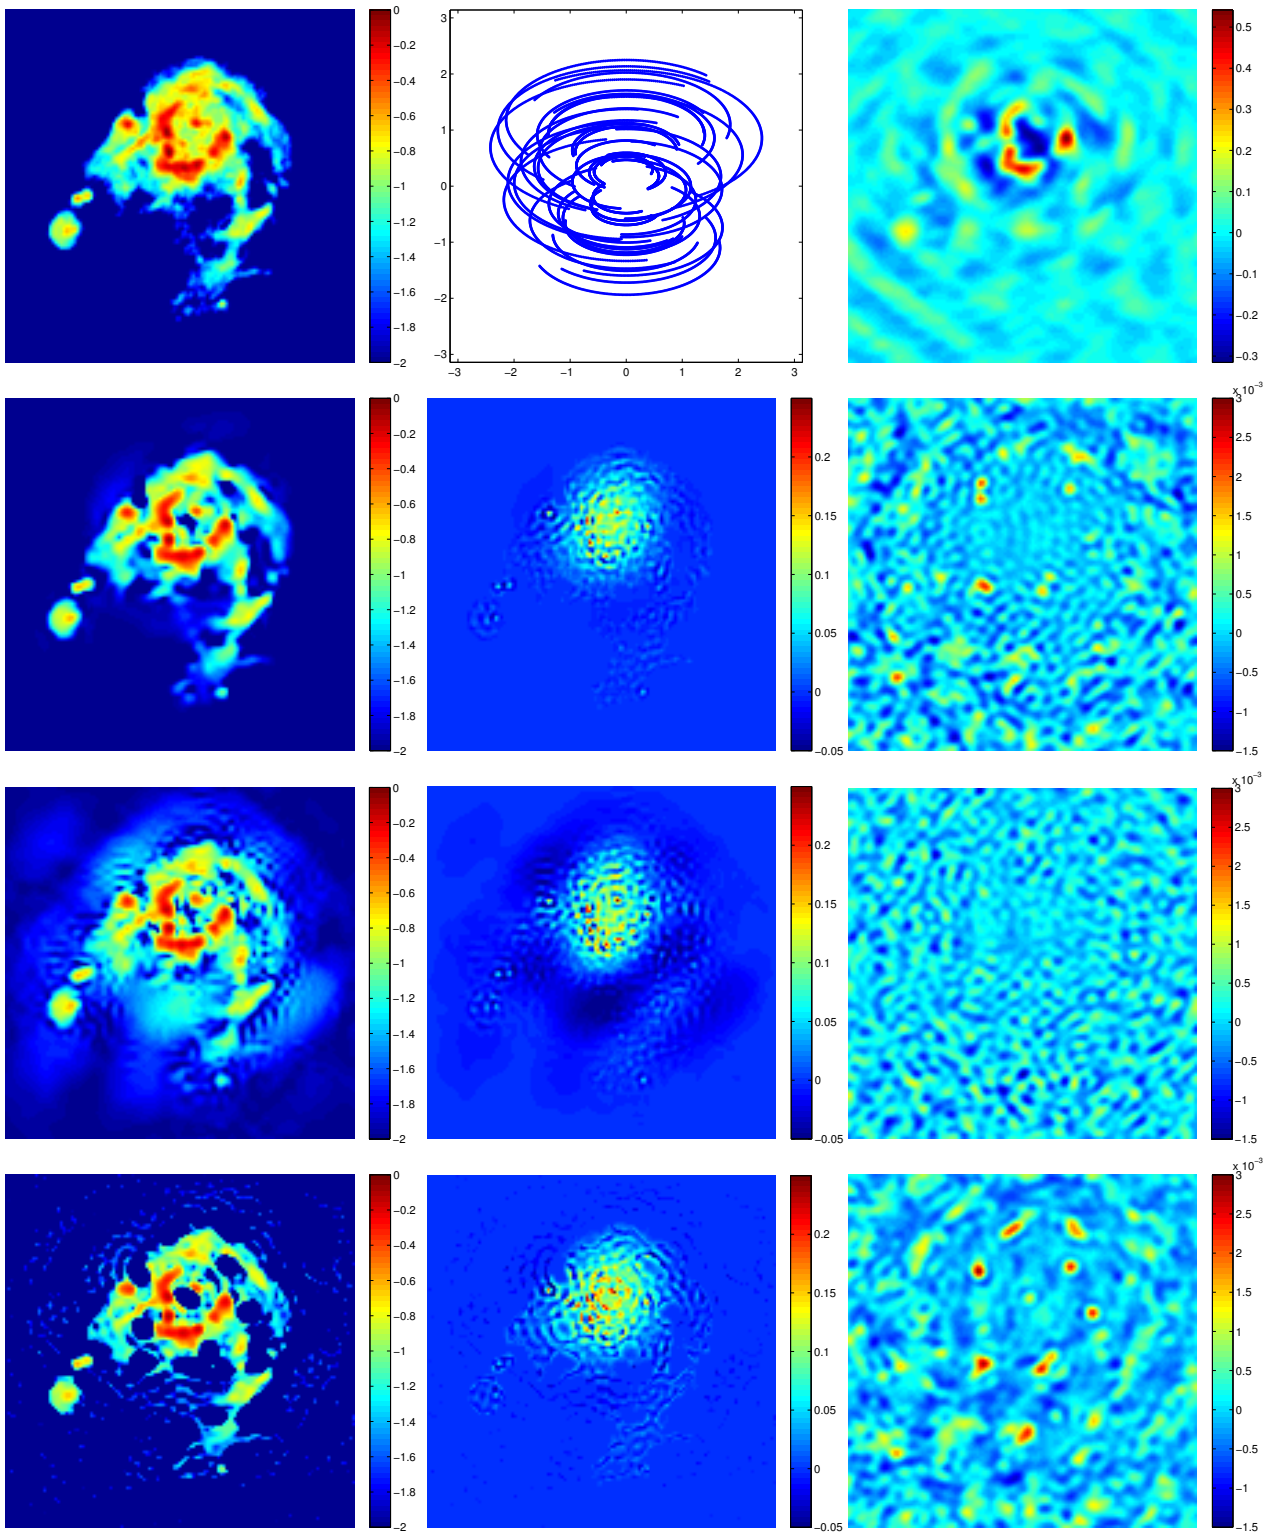
Figure 5. AMI coverage example. First row from left to right: original M31 128×128 test image in $\log_{10}$ scale, u-v coverage in normalized angular frequency units ($M = 9413$) and corresponding dirty image in linear scale. Second to last rows: reconstruction results for SARA (SNR=14.3 dB), RWBPDb8 (SNR=12.3 dB) and RWBP (SNR=10.9 dB). The first column shows the reconstructed images in a $\log_{10}$ scale, the second column shows the error images in linear scale, and the third column shows the residual dirty images also in linear scale.

Carrillo R.E., McEwen J.D., Ville D.V.D., Thiran J.P., Wiaux Y., 2013, IEEE Signal Process. Letters, 20, 6, 591

Carrillo R.E., McEwen J.D., Wiaux Y., 2012, MNRAS, 426, 2, 1223

Combettes P.L., Pesquet J.C., 2007, IEEE J. Sel. Top. Sig. Process., 1, 4, 564

Combettes P.L., Pesquet J.C., 2011, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, chapter Proximal splitting methods in signal processing, Springer, New York, 185–212

Cornwell T.J., 2008, IEEE J. Sel. Top. Sig. Process., 2, 5, 793

Cornwell T.J., Evans K.F., 1985, A&A, 143, 77

Donoho D.L., 2006, IEEE Trans. Inf. Theory, 52, 4, 1289

Fornasier M., Rauhut H., 2011, *Handbook of Mathematical Methods in Imaging*, chapter Compressed sensing, Springer

Greengard L., Lee J.Y., 2004, SIAM Review, 46, 3, 443

Gribonval R., Nielsen M., 2003, IEEE Trans. Inf. Theory, 49, 12, 3320

Högbom J.A., 1974, A&A, 15, 417

Li F., Cornwell T.J., de Hoog F., 2011, A&A, A31, 528

Mallat S., Zhang Z., 1993, IEEE Trans. Signal Process., 41, 12, 3397

Marsh K.A., Richardson J.M., 1987, A&A, 182, 174

McEwen J.D., Wiaux Y., 2011, MNRAS, 413, 2, 1318

Nam S., Davies M., R.Gribonval, Elad M., 2013, Applied and Computational Harmonic Analysis, 34, 1, 30

Rau U., Bhatnagar S., Voronkov M.A., Cornwell T.J., 2009, Proc. IEEE, 97, 1472

Rauhut H., Schnass K., Vandergheynst P., 2008, IEEE Trans. Inf. Theory, 54, 5, 2210

Rudin L.I., Osher S., Fatemi E., 1992, Physica D, 60, 259

Ryle M., Hewish A., 1960, MNRAS, 120, 220

Ryle M., Hewish A., Shakeshaft J., 1959, IRE Trans. Antenna Propag., 7, 120

Ryle M., Vonberg D.D., 1946, Nat., 158, 339

Saad Y., 2003, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Philadelphia

Schwarz U.J., 1978, A&A, 65, 345

Starck J., Murtagh F., Fadili J., 2010, *Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity*, Cambridge University Press, Cambridge, GB

Suksmono A., 2009, in *Proc. Int. Conf. on Electrical Eng. and Informatics*, volume 1, 110–116

Sullivan I.S., et al., 2012, ApJ, 759, 17

Thompson A.R., Moran J.M., Swenson G.W., 2001, *Interferometry and Synthesis in Radio Astronomy*, Wiley-Interscience, New York

Wenger S., Magnor M., Pihlströsm Y., Bhatnagar S., Rau U., 2010, Publ. Astron. Soc. Pac., 122, 897, 1367

Wiaux Y., Jacques L., Puy G., Scaife A.M.M., Vandergheynst P., 2009a, MNRAS, 395, 3, 1733

Wiaux Y., Puy G., Boursier Y., Vandergheynst P., 2009b, MNRAS, 400, 2, 1029

Wiaux Y., Puy G., Vandergheynst P., 2010, MNRAS, 402, 4, 2626

Wolz L., McEwen J.D., Abdalla F.B., Carrillo R.E., Wiaux Y., 2013, MNRAS, 463, 3, 1993

Zwart J.T.L., et al., 2008, MNRAS, 391, 4, 1545