LOFAR sparse image reconstruction

Garsden, H.; Girard, J.; Starck, J.; Corbel, S.; Tasse, C.; Woiselle, A.; ... ; Vocks,

Astronomy & Astrophysics

LOFAR sparse image reconstruction

H. Garsden[1], J. N. Girard[1,*], J. L. Starck[1], S. Corbel[1], C. Tasse[2,3], A. Woiselle[4,1], J. P. McKean[5],
A. S. van Amesfoort[5], J. Anderson[6,7], I. M. Avruch[8,9], R. Beck[10], M. J. Bentum[5,11], P. Best[12], F. Breitling[7],
J. Broderick[13], M. Brüggen[14], H. R. Butcher[15], B. Ciardi[16], F. de Gasperin[14], E. de Geus[5,17], M. de Vos[5], S. Duscha[5],
J. Eislöffel[18], D. Engels[19], H. Falcke[20,5], R. A. Fallows[5], R. Fender[21], C. Ferrari[22], W. Frieswijk[5], M. A. Garrett[5,23],
J. Grießmeier[24,25], A. W. Gunst[5], T. E. Hassall[13,26], G. Heald[5], M. Hoeft[18], J. Hörandel[20], A. van der Horst[27],
E. Juette[28], A. Karastergiou[21], V. I. Kondratiev[5,29], M. Kramer[10,26], M. Kuniyoshi[10], G. Kuper[5], G. Mann[7],
S. Markoff[27], R. McFadden[5], D. McKay-Bukowski[30,31], D. D. Mulcahy[10], H. Munk[5], M. J. Norden[5], E. Orru[5],
H. Paas[32], M. Pandey-Pommier[33], V. N. Pandey[5], G. Pietka[21], R. Pizzo[5], A. G. Polatidis[5], A. Renting[5],
H. Röttgering[23], A. Rowlinson[27], D. Schwarz[34], J. Sluman[5], O. Smirnov[2,3], B. W. Stappers[26], M. Steinmetz[7],
A. Stewart[21], J. Swinbank[27], M. Tagger[24], Y. Tang[5], C. Tasse[35], S. Thoudam[20], C. Toribio[5], R. Vermeulen[5], C. Vocks[7],
R. J. van Weeren[36], S. J. Wijnholds[5], M. W. Wise[5,27], O. Wucknitz[10], S. Yatawatta[5], P. Zarka[35], and A. Zensus[10]

(Affiliations can be found after the references)

ABSTRACT

Context. The LOw Frequency ARray (LOFAR) radio telescope is a giant digital phased array interferometer with multiple antennas distributed in Europe. It provides discrete sets of Fourier components of the sky brightness. Recovering the original brightness distribution with aperture synthesis forms an inverse problem that can be solved by various deconvolution and minimization methods.
Aims. Recent papers have established a clear link between the discrete nature of radio interferometry measurement and the "compressed sensing" (CS) theory, which supports sparse reconstruction methods to form an image from the measured visibilities. Empowered by proximal theory, CS offers a sound framework for efficient global minimization and sparse data representation using fast algorithms. Combined with instrumental direction-dependent effects (DDE) in the scope of a real instrument, we developed and validated a new method based on this framework.
Methods. We implemented a sparse reconstruction method in the standard LOFAR imaging tool and compared the photometric and resolution performance of this new imager with that of CLEAN-based methods (CLEAN and MS-CLEAN) with simulated and real LOFAR data.
Results. We show that i) sparse reconstruction performs as well as CLEAN in recovering the flux of point sources; ii) performs much better on extended objects (the root mean square error is reduced by a factor of up to 10); and iii) provides a solution with an effective angular resolution 2−3 times better than the CLEAN images.
Conclusions. Sparse recovery gives a correct photometry on high dynamic and wide-field images and improved realistic structures of extended sources (of simulated and real LOFAR datasets). This sparse reconstruction method is compatible with modern interferometric imagers that handle DDE corrections (A- and W-projections) required for current and future instruments such as LOFAR and SKA.

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

1. Introduction

Recent years have seen the development and planning of very large radio interferometers such as the LOw Frequency ARray (LOFAR; van Haarlem et al. 2013) in Europe, and the Square Kilometre Array (SKA; Dewdney et al. 2009) in Australia and South Africa (through its various precursors and pathfinders).

These new digital instruments have a very high sensitivity and a large field of view (FoV), as well as very large angular, temporal, and spectral resolutions in the radio spectrum observable from Earth. In particular, the very low frequency window (in the VHF – the very high frequency band between ∼10 and 250 MHz) is being explored (or revisited) with LOFAR, within the scope of various key science projects, spanning the search for fast transients (Stappers et al. 2011) to the study of early cosmology (de Bruyn & LOFAR EoR Key Science Project Team 2012).

1.1. LOFAR

LOFAR is a digital radio interferometer composed of 48 stations: 40 stations constitute the core and remote stations (two of which are future stations) that are distributed in the Netherlands (NL), and eight international stations that are located in Germany, Great Britain, France, and Sweden. Poland has planned the construction of three international stations that will enrich the array in the east-west direction. LOFAR has two working bands: the LBA-band (low-band antenna) from ≥30 to 80 MHz (that can be switched to 10−80 MHz), and the HBA-band (high-band antenna) from 110 to 250 MHz, lying on either-side of the FM band. One station consists of two arrays composed of fixed crossed dipole antennas, each offering a large FoV (therefore low directivity) and broadband properties[1]. The HBA field of the core stations is split into two to increase the number of baselines. The stations measure induced electric signals that undergo pre-processing operations in the station backend, consisting of digitization, filtering, phasing, and summing.

---

* Corresponding author: J. N. Girard,
e-mail: julien.girard@cea.fr

---

[1] See details at www.lofar.org

All these steps constitute the beamforming step of the phased antenna array. The output signal of one station is thus similar to that of a synthetic antenna whose beam is electronically pointed (rather than mechanically) in the direction of interest. Since most steps are performed in the digital side, multi-beam observations are possible and are only limited by the electronic hardware (i.e. in field-programmable gate arrays – FPGAs, by making trade-offs between the observed bandwidth and the number of numerical beams pointing at the sky).

At the interferometer level, the signal from every station is combined in a central correlator in the NL that performs a phased sum or inter-station cross-correlation. The latter operation enables LOFAR to build up a digital interferometer that samples a wide range of baselines (from ~70 m up to ~1300 km). The high flexibility of the instrument comes with the necessity of using advanced calibration strategies depending on the type of observations and its expected final sensitivity.

Since late 2012, LOFAR has been opened to the astronomical community by a regular public call for proposals[2]. Early results with LOFAR demonstrated that it can reach a very high dynamic range (DR) on a wide FoV (i.e. several tens of degrees, see for example Yatawatta et al. 2013) and a very high angular resolution (Wucknitz 2010; Shulevski 2010) at low frequencies.

LOFAR capabilities are geared to key science projects on deep extragalactic surveys, transient radio phenomena and pulsars, the epoch of reionization, high-energy cosmic rays, cosmic magnetism, and solar physics and space weather.

1.2. Increased complexity of low-frequency imaging

Since the beginning of radio interferometry imaging, various imaging methods have been designed to fit the requirements of different types of (extended) radio objects. The availability of high-performance computing and the need for efficient, fast, and accurate imaging for new wide-field interferometers has motivated the implementation of new imaging algorithms. Given the recording time and frequency resolutions, the integration time, and the diversity of baselines of wide-field interferometers, large amounts of data storage are required to save the telescope data. These data must then be transformed into a scientifically exploitable form (typically into images cubes). Substantial computational power is required for this as well (Begeman et al. 2011).

Because of the nature and the dimensions of the LOFAR array, direction-dependent effects (DDE; Tasse et al. 2012) occur during the span of a LOFAR observation and add up to the usual other effects intervening in classical radio interferometers. These effects require a direction-dependent calibration before imaging. In particular, the classical compact planar array and small FoV assumptions are no longer valid, especially for a wide-field instrument such as LOFAR. Modern calibration and imaging at radio wavelengths require a similar approach to that used in modern optical telescope with adaptive optics.

The problem can be generalized and expressed in the Measurement Equation framework (Sect. 2.5). The calibration problem therefore manifests as an inverse problem that should be solved to independently determine all the parameters and coefficients that describe each observation dataset.

Among the recent developments of data processing and reconstruction algorithms, the discovery of compressed sensing (CS; Candès et al. 2006) has led to new approaches to solving these problems. It has been proposed for radio interferometry (e.g. Wiaux et al. 2009a; Wenger et al. 2010; Li et al. 2011a,b;

Carrillo et al. 2012) because this constitutes a relevant practical case as a result of the sparse nature of the interferometric sampling (Sect. 2.4).

The implementation of sparse radio image reconstruction methods is expected to produce better results on large extended objects with high angular resolution than other classical deconvolution methods. In Sect. 2 we first apply the theory of sparse reconstruction within the scope of radio aperture synthesis imaging and then implement it in the LOFAR imaging software. We also relate it to previous CLEAN-based deconvolution methods. We present in Sect. 3 the results of benchmark tests using simulated and real LOFAR datasets by focusing on the quality of the image reconstruction compared to that of usual CLEAN-based algorithms. We then discuss the practical advantages and limitations of the current implementation and possible future developments.

2. Image reconstruction: from CLEAN to compressed sensing

2.1. Introduction

An ideal radio interferometer, composed of co-planar and identical antennas, samples the sky in the Fourier domain (Wilson et al. 2009). In other words, each pair of antennas that forms one baseline gives access to the measure of the sky brightness as seen through a set of fringes with characteristics that depend on the frequency, the baseline length, and orientation with respect to the source. As for optical interferometers, the measured quantity (after correlating the pair of signals) is the fringe contrast, also called (fringe) visibility. In the radio domain where the electric field varies slowly, we can sample both phase and amplitude with time and frequency. In the scope of radio interferometry, the visibility function is a complex function.

An N-antenna ideal interferometer provides $N(N-1)/2$ independent instantaneous visibility measurements. We define the spatial frequency coordinates plane, the (u, v) plane, which is orthogonal to the direction of observation containing all projected baselines (the third coordinate, w, is omitted in the small-field approximation). This direction defines the origin of the Fourier conjugated coordinate system on the sky, the direction cosines (l, m).

At any given time and frequency, one projected baseline samples one spatial frequency represented by one point in the (u, v) plane. With time, each (u, v) point sweeps the (u, v) plane, which enriches it with new samples (i.e. *Earth rotation synthesis*). Figure 1 presents the (u, v) coverage of one typical LOFAR observation integrated over six hours. Each track is associated with one antenna pair, and its shape depends on the antenna configuration, the instrument latitude, and the direction of observation.

The number of visibilities tends to increase with the building of million-element interferometers such as SKA, but stays small in case of snapshot (short integration time) observations. This collection of Fourier samples is the starting point for the imaging process, which consists of using the measurements to recover the true visibility function in the Fourier domain. After instrumental calibration, a dirty image can be generated by gridding the calibrated visibilities to a 2D plane by inverting this approximate image of the true Fourier transform to the image plane. The missing samples contribute to corrupt the image because it is a rough approximation of the sky brightness. This image is the convolution of the true sky brightness (represented as a 2D image) with the interferometer point spread function (PSF, or dirty beam).
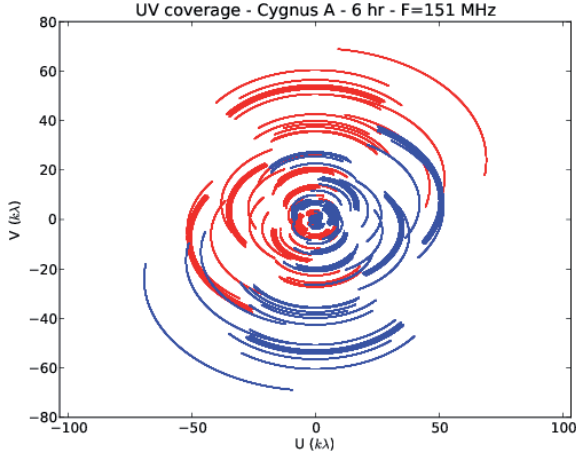

Fig. 1. Visibility coverage from a six-hour observation of Cygnus A (McKean et al. 2010). Visibilities are plotted in the Fourier space with the U (x-axis) and V (y-axis) being the spatial frequencies in wavelength unit, λ (here $f = 151$ MHz, $\lambda \approx 2$ m) determined by the baseline projection on the sky. Each (u, v) point (red) has its symmetric $(-u, -v)$ point (blue) corresponding to the same baseline. The lines indicate (u, v) points where a visibility was recorded. The arcs are built from varying the baseline projections with the rotation of Earth during the observation.

We now discuss the different approaches that can be used to approach the true sky brightness from the measured visibilities, starting with the CLEAN method, before describing the method based on the sparsity of the measured signal.

2.2. CLEAN

For many years, deconvolution has been achieved through the CLEAN algorithm (Högbom 1974) and its variants (Clark 1980; Schwab 1984). CLEAN considers the dirty image to be constructed from point sources convolved with the PSF; extended objects are decomposed as point sources as closely as possible. CLEAN operates directly on the dirty image (original versions like Högbom do) by locating the maximum of the image and iteratively subtracting a fraction of the dirty beam centred and scaled to the located maximum. The detected sources are indexed as CLEAN components to form a model image enriched by the successive subtraction steps. The source detection and subtraction continues until a threshold is reached on the residual image (typically representing the background level), which is assumed to contain no remaining sources. The model image contains a pixel-wise description of the levels and locations of the detected sources. To remove the unphysically high spatial frequency components (associated with the pixel size) that are introduced by the CLEAN algorithm, the model image is usually convolved by an elliptical Gaussian 2D fit of the centre of the dirty beam (which is the CLEAN beam) to provide a final angular resolution corresponding to that accessible by the interferometer. The residual image, containing no other source, is added to the convolved CLEAN image to form the restored image and represents the noisy sky background.

This image can be described in the following manner:

$$\mathcal{I} = \mathcal{M} * C + \mathcal{R}, \tag{1}$$

where $\mathcal{I}$ is the restored image, $\mathcal{M}$ is the model composed of CLEAN components, C is the CLEAN beam, $\mathcal{R}$ is the residual, and $*$ is the convolution operator.

The CLEAN method has several variants, one of which is Cotton-Schwab CLEAN (CoSch-CLEAN) from Schwab (1984),

which was used for all experiments in this article. CLEAN can be described as a least-squares minimization that solves normal equations and its variants that are usually split into *major* and *minor* cycles by combining a steepest-descent method (major cycle) and a greedy algorithm (minor cycle), which quickly calculates an approximate solution of the normal equations.

CoSch-CLEAN uses the steepest-descent method (during the major cycle) combined with a greedy algorithm (during the minor cycle) that calculates an approximate solution of the normal equations.

With LOFAR, the PSF varies over the FoV and is also difficult to determine accurately because of instrumental DDEs (Sect. 2.5). Moreover, the sensitivity and variety of baseline **s** brought by the instrument often limit the quality of the restored images using these methods.

2.3. Toward multi-resolution

The CLEAN method is well known to produce poor solutions when the image contains large-scale structures. Wakker & Schwarz (1988) introduced the concept of multi-resolution CLEAN (M-CLEAN) to alleviate difficulties occurring in CLEAN for extended sources. The M-CLEAN approach consists of building two intermediate images, the first one (called the smooth map) by smoothing the data to a lower resolution with a Gaussian function, and the second one (called the difference map) by subtracting the smoothed image from the original data. Both images are then processed separately. By using a standard CLEAN algorithm on them, the smoothed clean map and difference clean map are obtained. The recombination of these two maps gives the clean map at the full resolution.

This algorithm may be viewed as an artificial recipe, but it has been shown that it is linked to wavelet analysis (Starck & Bijaoui 1991), leading to a wavelet CLEAN (W-CLEAN) method (Starck et al. 1994). Furthermore, in the W-CLEAN algorithm, the final solution was derived using a least-squares iterative reconstruction algorithm applied to the set of wavelet coefficients detected by applying CLEAN on different wavelet scales. This can be interpreted as a debiased post-processing of the peak amplitudes found at the different scales. A positivity constraint on the wavelet coefficients was imposed during this iterative scheme by nullifying the negative coefficients.

Other multi-scale approaches exist, such as the adaptive scale pixel (Asp) deconvolution (Bhatnagar & Cornwell 2004) (which fits for parameters of Gaussian sets), the multi-scale CLEAN (MS-CLEAN; Cornwell 2008), and the multi-scale multi-frequency (MS-MF) deconvolution (Rau & Cornwell 2011). On the one hand, Asp is not implemented in the LOFAR imager, and on the other hand, we are imaging single-frequency channel datasets that do not justify the use of MS-MF even though they are in the LOFAR imager. We therefore only use MS-CLEAN in this paper as a scale-sensitive algorithm. It consists of fitting a collection of extended patches (or blobs). Instead of subtracting the dirty beam at a given location of the residuals at each iteration, as in the CLEAN algorithm, MS-CLEAN subtracts a blob, estimating first the most adequate blob size. MS-CLEAN presents several problems that do not exist in M-CLEAN or W-CLEAN, however. Indeed, it is unclear which function the algorithm minimizes, or even if it minimizes anything at all. The varying background is also problematic in MS-CLEAN (as in CLEAN), while it is not in W-CLEAN. MS-CLEAN also relies on an arbitrary manual setting of the characteristic scales of the image.

However, these different CLEAN-based algorithms (hereafter denoted x-CLEAN) have all shown to significantly improve the results over CLEAN when the image contains extended sources.

Other methods, based on a statistical Bayesian approach, have been developed to image extended emission (Junklewitz et al. 2013; Sutter et al. 2014) but were not included in the present work.

2.4. Compressed sensing and sparse recovery

Compressed sensing (Candès et al. 2006; Donoho 2006) is a sampling and compression theory based on the sparsity of an observed signal, which shows that under certain conditions, one can exactly recover a k-sparse signal (a signal for which only k coefficients have values different from zero out of n total coefficients, where $k < n$) from $m < n$ measurements. CS requires the data to be acquired through a random acquisition system, which is not the case in general. However, even if the CS theorem does not apply from a rigorous mathematical point of view, some links can still be considered between real life applications and CS. In astronomy, CS has already been studied in many applications such as data transfer from satellites to Earth (Bobin et al. 2008; Barbey et al. 2011), 3D weak lensing (Leonard et al. 2012), next-generation spectroscopic instrument design (Ramos et al. 2011), and aperture synthesis. Indeed, there is a close relationship between CS principles and the aperture-synthesis image reconstruction problem, which was first addressed in Wiaux et al. (2009a), Wenger et al. (2010), Li et al. (2011b) and Carrillo et al. (2012). Wide-field observations were subsequently studied in McEwen & Wiaux (2011), and different antenna configurations (Fannjiang 2013) and non-coplanar effects (Wiaux et al. 2009a,b; Wolz et al. 2013) were analysed in a compressed-sensing framework. Aperture synthesis presents the three main ingredients that are fundamental in CS:

– *Underdetermined problem:* we have fewer measurements (i.e. visibilities) than unknowns (i.e. pixel values of the reconstructed image).
– *Sparsity* of the signal: the signal to reconstruct can be represented with a small number of non-zero coefficients. For point source observations, the solution is even strictly sparse (in the Dirac domain) since it is only composed of a list of spikes. For extended objects, sparsity can be obtained in another representation space such as wavelets.
– *Incoherence* between the acquisition domain (i.e. Fourier space) and the sparsity domain (e.g. wavelet space). Point sources, for instance, are localized in the pixel domain, but spread over a large domain of the visibility plane. Conversely, each visibility contains information about all sources in the FoV.

From the CS perspective, the best way to reconstruct an image X from its visibilities is to use sparse recovery by solving the following optimization problem:

$$\min \|X\|_p \quad \text{subject to} \quad \|V - \mathbf{A}X\|_2^2 < \epsilon, \tag{2}$$

where V is the measured visibility vector, $\mathbf{A}$ the operator that embodies the realistic acquisition of the sky brightness components (instrumental effects, DDE, etc.), and $\|z\|_p^p = \sum_i |z_i|^p$. To reinforce the sparsity of the solution, p is positive but must be lower than or equal to 1. In particular, for $p = 0$, we derive the ℓ_0 pseudo-norm that counts the number of non-zero entries in z. The first part of Eq. (2) enforces sparsity, the second part indicates that the reconstruction matches the visibilities within some

error ϵ. Most natural signals, however, are not exactly sparse but rather concentrated within a small set. Such signals are termed *compressible* or *weakly sparse*, in the sense that the sorted coefficient values decay quickly according to a power law. The faster the amplitudes of coefficients decay, the more compressible the signal is.

It is interesting to note that the CLEAN algorithm can be interpreted as a *matching-pursuit algorithm* (Lannes et al. 1997) that minimizes the ℓ_0 pseudo-norm of the sparse recovery problem of Eq. (2), but recent progress in the field of numerical optimization (see application with the SARA algorithm in Carrillo et al. 2014 and references therein) allows us to have much faster algorithms. The sparsity model is extremely accurate for a field containing point sources, since the true sky can in this case be represented just by a list of spikes. This explains the very good performance of Högbom's CLEAN for point sources recovery, and why astronomers still use it 40 years after it has been published. But CS provides a context in which we can understand the limitation of CLEAN for extended object reconstructions. The notion of sparsity in radio astronomy is not new and may be recognized in algorithms that apply a transform to the data. For example, extended objects are not sparse in the pixel basis, therefore sparse recovery algorithms cannot provide good solutions on this basis. Indeed, as depicted in Wakker & Schwarz (1988) and Starck et al. (1994) with wavelets, representing the data in another domain where the solution is sparse was shown to be a good approach. More generally, we can assume that the solution X can be represented as the linear expansion

$$X = \mathbf{\Phi}\alpha = \sum_{i=1}^t \phi_i \alpha_i, \tag{3}$$

where α are the synthesis coefficients of X, $\mathbf{\Phi} = (\phi_1, \ldots, \phi_t)$ is the dictionary whose columns are t elementary waveforms ϕ_i also called *atoms*. The dictionary $\mathbf{\Phi}$ is a $b \times t$ matrix whose columns are the normalized atoms (of size b), assumed here to be normalized to a unit ℓ_2-norm, that is, $\forall i \in [1, t], \|\phi_i\|_2^2 = \sum_{n=1}^N |\phi_i[n]|^2 = 1$. The minimization problem of Eq. (2) can now be reformulated in two ways, the synthesis framework

$$\min_\alpha \|\alpha\|_p \quad \text{subject to} \quad \|V - \mathbf{A}\mathbf{\Phi}\alpha\|_2^2 < \epsilon, \tag{4}$$

and the analysis framework

$$\min_\alpha \|\mathbf{\Phi}^t X\|_p \quad \text{subject to} \quad \|V - \mathbf{A}X\|_2^2 < \epsilon. \tag{5}$$

When the matrix $\mathbf{\Phi}$ is orthogonal, both analysis and synthesis frameworks lead to the same solution, but the choice of the synthesis framework is generally the best if X has a clear and unambiguous representation in $\mathbf{\Phi}$.

In that scope, the recent proximal theory can provide proof of convergence of algorithms constructed in this framework (in contrast to some CLEAN derivates). It also led to the development of fast optimization algorithms combined with fast transforms enabling a fast (usually $N \log N$) evaluation of atom coefficients. The best dictionary provides the sparsest (i.e. the most economical) representation of the signal. In practice, implicit dictionaries with fast transforms (such as the wavelet or curvelet transforms, etc.) allow us to directly obtain the coefficients and reconstruct the signal using fast algorithms running in linear or almost linear time (unlike matrix-vector multiplications).

2.5. Imaging with LOFAR

For a large multi-element digital interferometer such as LOFAR observing in a wide FoV, the small-field approximation is no longer valid, and the sampled data (represented by the operator **A**) are no longer a discrete set of 2D Fourier components of the sky. The instrumental (direction-independent) and natural direction-dependent effects have to be taken into account for proper image reconstruction. These effects include

– the instrumental effects such as inter-station clock shifts,
– the non-coplanar nature of the baselines (Cornwell & Perley 1992) and the effect of their projections (the sample visibility function has a non-zero third coordinate, i.e. $w \neq 0$),
– the anisotropic directivity of the phased-array beam (Bhatnagar et al. 2008), and the non-trivial dipole projection effects with time and frequency,
– the sparsity in the sampling of the visibility function (i.e. the limited number of baselines and the time/freq integration), and
– the effect of the interstellar- and interplanetary media, and Earth's ionosphere, on the incoming plane waves.

These effects can be modelled in the framework of the radio interferometry measurement equation (RIME; see Hamaker et al. 1996; Sault et al. 1996; Smirnov 2011 and following papers), which describes the relation between the sky and the four-polarization visibilities associated with each pair of antennas in a time-frequency bin. It can model all the DDE mentioned above, cumulated in 2×2 *Jones* matrices that influence the electric field and voltage measurements. One of the most basic ways to express a four-polarization visibility of one baseline using the RIME formalism is as follows (in a given time-frequency bin):

$$\mathbf{V}_{pq}^{\text{real}} = \mathbf{J}_p^{-1} \mathbf{V}_{pq}^{\text{mes}} (\mathbf{J}_q^{\text{H}})^{-1}, \tag{6}$$

with

$\mathbf{V}^{\text{mes}}$ the four-polarization (XX, XY, YX, YY) measurement matrix corrupted by the DDE,
$\mathbf{V}^{\text{true}}$ the true visibility matrix uncorrupted by the DDE,
$\mathbf{J}_p$ the cumulated *Jones* matrix of antenna p, and
$\mathbf{J}_q^{\text{H}}$ the Hermitian transpose (conjugated transpose) of the cumulated *Jones* matrix of antenna q.

The corrected visibility $\mathbf{V}_{pq}^{\text{corr}}$ matrix can be expressed as a four-dimensional vector (which also depends on the time t and frequency ν) as in Eq. (1) of Tasse et al. (2013),

$$\text{Vec}(\mathbf{V}_{pq}^{\text{corr}}) = \int_S (\mathbf{D}_{q,s}^{t\nu,*} \otimes \mathbf{D}_{p,s}^{t\nu}) \text{Vec}(I_s) \times e^{-2i\pi\phi(u,v,w,s)} ds, \tag{7}$$

where I_s is the four-polarization sky, s the pointing direction, $\otimes$ the Kronecker product producing a 4×4 matrix (referred to as the Mueller matrix), Vec() the operator that transforms a 2×2 matrix into a four-dimensional vector, $\mathbf{D}_{q,s}^{t\nu,*} \otimes \mathbf{D}_{p,s}^{t\nu}$, the 4×4 Mueller matrix, containing the accumulation of the direction-dependent terms and the array geometry (see details in Appendix A in Tasse et al. 2013), and $\phi(u,v,w,s) = ul + vm + w(\sqrt{1 - l^2 - m^2} - 1)$, the baseline and direction-dependent phase factor. Because the sky is described in terms of Stokes intensity (I, Q, U, V), and not in terms of the electric field value, the RIME is linear in its sky-term. By considering the total set of visibilities over baselines, time, and frequency bins, and by including the measurement noise ϵ, we can therefore reduce Eq. (7) to

$$V = \mathbf{A}I_s + \epsilon, \tag{8}$$

with

I_s the four-polarization sky;

A the transformation matrix from the sky to the visibilities, including all DDE; and
ϵ the measurement noise.

The structure of matrix **A** and its connection to the RIME and Jones formalism shown above is described in detail in Tasse et al. (2012). The linear operator **A** contains (i) the Fourier kernels and the information on, (ii) the time-frequency-baseline dependence of the effective 2×2 Jones matrices, and (iii) the array configuration that is represented by the (u, v, w)-sampling over time and frequency. It is important to note that both (ii) and (iii) cause **A** to be non-unitary, which is a very important fact for the work presented in this paper. In practice, it is virtually impossible to make **A** explicit, essentially because of its $N_{\text{vis}} \times N_{\text{pixels}}$ dimension. Instead, **A** and $\mathbf{A}^{\text{T}}$ can be applied to any sky or visibility vector respectively using *A-projection* (Bhatnagar et al. 2008; Tasse et al. 2013). To cope with the non-coplanar effect, the *W-projection* algorithm (Cornwell et al. 2008) is used to turn the 3D recorded visibilities into a 2D Fourier transform. In the scope of the LOFAR project and its data reduction pipeline, the AWimager (Tasse et al. 2013) program was developed for imaging the LOFAR data by taking into account both A- and W-projections. It is therefore of high interest to implement CS in this type of new-generation imagers.

2.6. Algorithm

To perform the sparse image reconstruction, we need to

1. select a minimization method to solve Eq. (4) or Eq. (5),
2. select a dictionary **Φ**, and
3. select the parameter related to the minimization method.

Several minimization methods have been used for aperture synthesis, the FISTA method (fast iterative shrinkage-thresholding algorithm in Beck & Teboulle 2009) in Li et al. (2011b), Wenger et al. (2010, 2013), and Hardy (2013), the OMP (orthogonal matching pursuit; Davis et al. 1997) in Fannjiang (2013), Douglas-Rachford splitting in Wiaux et al. (2009b), McEwen & Wiaux (2011) and Carrillo et al. (2012) or the SDMM simultaneous-direction method of multipliers in Combettes & Pesquet (2011) and Carrillo et al. (2014). In our experiments, we investigated mainly two algorithms, FISTA, and a recent algorithm proposed by Vũ (2013), which works in the analysis framework. As they were providing very similar results, we chose to report here results derived with the FISTA algorithm. Full details can be found in Beck & Teboulle (2009), Wenger et al. (2010), and Starck et al. (2010).

The choice of dictionary is critical. Optimal dictionaries should contain atoms that represent the content of the data well. In this sense, the starlet transform, also called isotropic undecimated wavelet transform (Starck et al. 2010), which was used in Li et al. (2011b), is a very good choice, since this decomposition has shown to be extremely useful for astronomical image restoration (Starck & Murtagh 2006). In contrast, orthogonal or bi-orthogonal wavelets, even using several decompositions, are well known to produce artefacts due to critical sampling. Other dictionaries such as curvelets (Starck et al. 2010) are a good alternative if the data contain directional features, such as jets, which will be poorly represented with wavelets. The block discrete cosine transform (BDCT) used in Fannjiang (2013) is relatively hard to justify for astronomical images, since its atoms present an oscillatory pattern. In our study, we focus on the starlet dictionary.

The convex optimization problem formulated in Eq. (4) can be recast in an augmented Lagrangian form with

$$\min_{\alpha} \|V - \mathbf{A\Phi}\alpha\|_2^2 + \sum_j \lambda_j |\alpha_j|. \tag{9}$$

The first part of Eq. (9) indicates that the reconstruction should match the data and the second term enforces sparsity (through controlling the regularization parameter λ_j at each scale j of $\mathbf{\Phi}$). The information about the error ϵ in Eq. (2) is included in λ_j. The problem can then be solved using the FISTA algorithm (Beck & Teboulle 2009), following the implementation of algorithm 1.

The final problem remains: choosing parameters that are needed to control the algorithm. Most minimization methods, using l_0 and l_1, have a single thresholding step, where coefficients in the dictionary have to be soft or hard-thresholded using a threshold value, which is a value λ, common to projection in Φ (unlike Eq. (9)) (see Starck et al. 2010 for more details on hard and soft thresholding). This parameter controls the trade-off between the fidelity to the observed visibility and the sparsity of the reconstructed solution. In Li et al. (2011b) and Wenger et al. (2013), λ was fixed to an arbitrary value, different for each experiment, and certainly after several tests. This approach may be problematic for real data where the true solution is not known.

Carrillo et al. (2014) addressed the constrained problem of Eq. (5) directly and the ϵ parameter is a bound on the ℓ_2-norm data fidelity term, which is related explicitly to the noise level under the assumption that the noise is *i.i.d.* Gaussian. As the noise is generally not white, a whitening operator has to be applied first.

We propose in this paper another strategy, where the threshold is fixed only from the noise distribution at each wavelet scale j. Indeed, estimating a threshold from the residuals in a pixel basis is far from being optimal. On the one hand, the residual is not necessarily of zero mean and the background level is not known, and, on the other hand, the signal can be at the noise level in the pixel basis, but much stronger in the wavelet basis.

This has two main advantages: the default threshold value should always give reasonable results, and it optimizes the probability detection of faint objects. Indeed, an arbitrary threshold value could lead to many false detections in the case where the value is too low, and in contrast, many objects may be missed when the value is too high.

At the nth iteration of the FISTA algorithm, the residual image R^n is calculated by

$$\mathbf{R}^{(n)} = \mu \mathbf{A}^{\mathrm{T}} \left(V - \mathbf{A\Phi}\alpha^{(n)} \right), \tag{10}$$

where V are the measured visibilities, $\alpha^{(n)}$ are the coefficients in the dictionary $\mathbf{\Phi}$ of the solution at the nth iteration, and μ is the FISTA relaxation parameter that depends on the matrix $\mathbf{F} = \mathbf{A\Phi}$ (μ must verify $0 < \mu < \frac{1}{\|F\|}$). If the chosen dictionary Φ is the starlet or the curvelet transform, then the noise level can be estimated for each scale j of the residual $R^{(n)}$ using a reliable estimator such as the MAD (median of the absolute deviation):

$$\sigma_j^{(n)} = \frac{\mathrm{MAD}\left(\alpha_j^{(R^{(n)})}\right)}{0.6745}, \tag{11}$$

denoting $\alpha_j^{(R^{(n)})}$ the coefficients of $R^{(n)}$ of the band j in the dictionary $\mathbf{\Phi}$. The thresholds λ_j (see Eq. (9)) are then different for each band, with $\lambda_j = k\sigma_j$, where k fixes the probabilities of false detections. This provides the advantage of properly handling correlated noise. The constant factor in Eq. (11) is derived from the quantile function of a normal distribution taken at probability 3/4 (Hoaglin et al. 1983).

Algorithm 1 FISTA implementation

Require:
   A dictionary $\mathbf{\Phi}$.
   The implicit operator $\mathbf{A}$.
   The original visibility data V.
   A detection level k.
   The soft-thresholding operator:

   $\mathrm{SoftThresh}_\lambda(\alpha) = \mathrm{prox}_\lambda(\alpha) = \left(\left(1 - \frac{\lambda_j}{|\alpha_j|} \right)_+ \alpha_j \right)_{1 \le j \le N}$

1: Initialize $\alpha^{(0)}=0$, $t^{(0)}=0$
2: **for** $n = 0$ to $N_{\max} - 1$ **do**
3:      $\boldsymbol{\beta}^{(n+1)} = \mathbf{x}^{(n)} + \mu\mathbf{\Phi}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}(V - \mathbf{A\Phi}\alpha^{(n)})$
4:      $\mathbf{x}^{(n+1)} = \mathrm{SoftThresh}_{\mu\lambda_j}\boldsymbol{\beta}^{(n+1)}$
5:      $t^{(n+1)} = (1 + \sqrt{1 + 4(t^{(n)})^2})/2$
6:      $\gamma^{(n)} = (t^{(n)} - 1)/t^{(n+1)}$
7:      $\alpha^{(n+1)} = \mathbf{x}^{(n)} + \gamma^{(n)}(\mathbf{x}^{(n)} - \mathbf{x}^{(n-1)})$
8: **end for**
9: **Return:** image $\mathbf{\Phi}\alpha^{(N_{\max})}$.

3. Benchmarking of the sparse recovery

3.1. Benchmark data preparation

The performance of CS can be determined and compared to that of other imaging methods such as CoSch-CLEAN and MS-CLEAN. We used here three different LOFAR datasets (Table 1) to compare CS and CLEAN-based algorithms in simulated and real situations.

We implemented CS by creating a new method in AWimager based on the current CoSch-CLEAN implementation. This new method serves as an entry point to the LOFAR imager to connect an external library (developed by the authors, see Sect. 5.3) containing the dictionaries (wavelets and curvelets), transforms, and minimization methods described above. The CS cycle involves a continual passing of data back and forth between AWimager and the library (before and after *gridding*, *degridding* of the data, and applying the *A*- and *W-projections*). At step i in the CS cycle, AWimager performs the following operations:

- it takes as input an image reconstruction obtained at iteration $i - 1$ (represented in the selected dictionary),
- it simulates the visibilities using the same $\mathbf{A}$ operator as for CoSch-CLEAN (and perform *degridding*),
- it compares the simulated visibilities to the original ones, and
- it passes the difference back to the dictionary using $\mathbf{A}^{\mathrm{T}}$, and performs a step of the CS minimization, thresholding, and produces a reconstructed image of iteration i.

The process stops if the maximum number of iterations is reached or if a convergence criterion (based on the noise of the residual) is met. The current implementation fits the current framework used by recent algorithms where the major cycle performs the CS operations (especially line 3 of Algo 1, which enables computing Eq. (10)). In the future, multiple steps of CS will be possible (in a minor cycle) to improve performance. AWimager outputs a similar set of files for both CoSch-CLEAN and CS (the restored, residual, model, and a PSF image). CLEAN output is a CLEAN-beam-convolved image; in the case of CS, the reconstructed image is the solution $\mathbf{\Phi}\alpha^n$ (Eq. (10)). The program user may choose to convolve the compressed sensing output with the CLEAN beam and/or add the residual image (see discussion in Sect. 4.1). The CS residual is $R^{(n)}$ (Eq. (10)) at the last iteration of CS and is analogous to the residual at the output of x-CLEAN. By selecting the CS method,

Table 1. Summary information of the LOFAR datasets used in the present study.

Name	**A**	**B**	**C**
Sources	Two point sources	Source grid & W50	Cygnus A
Nature of the data	simulated	simulated	real
Pointing RA (hms)	02h45m00.0″	14h11m20.9s	19h59m28.0s
Pointing Dec (dms)	+52°54m54.4s	+52°13m55.2s	+40°44m02.0s
Station configuration	Core	*Core+Remote*	*Core+Remote*
$N_{stations}$	48	55	36
$N_{correlations}$	1176	1540	666
UV_{max} ($k\lambda$)	3.7 $k\lambda$	~65 $k\lambda$	~39 $k\lambda$
Maximum resolution (′)	~1′	~3″	~5″
Polarization		(linear) XX, YY, XY, YX	
t_{start} (UT)	6-Nov-2013/00:00:00.5	13-Apr-2012/18:25:05.0	2-Mar-2013/05:50:50.0
t_{end} (UT)	6-Nov-2013/00:59:59.5	14-Apr-2012/04:44:56.6	2-Mar-2013/11:50:40.0
δt (s)	1	10.01	10.01
N_{times}	3600 (=1 Hr)	37 192 (~10 Hrs)	21 590 (~6 Hrs)
ν_0 (MHz)	150.1	116.2	151.2
$\Delta\nu$ (MHz)	0.195	0.183	0.195
Section	Sect. 3.2.1	Sects. 3.2.2 and 3.3.1	Sect. 3.3.2

Notes. $N_{stations}$ and $N_{correlations}$ are the number of stations and the corresponding number of independent correlation measurements (computed as $N_{correlations} = N_{stations}(N_{stations} - 1)/2 + N_{stations}$). Autocorrelations were systematically flagged before imaging. UV_{max} is the largest radial distance in the (u, v) plane that limits the maximum theoretical resolution of the instrument. t_{start} and t_{end} observation times are in Universal Time (UT), providing a total number of time samples N_{times} at a time resolution δt. We considered only one frequency channel covering a bandwidth $\Delta\nu$ centred on the frequency ν_0. All three datasets come in *measurements sets* that follow the NRAO standards. Dataset **A** has been generated for a one-hour synthetic zenith observation using *mkant* and *makeMS* routines that are part of the LOFAR software package. Dataset **B** comes from a LOFAR commissioning observation and is filled with simulated data. Dataset **C** contains actual calibrated data from a real LOFAR observation of Cygnus A.

the meaning of some classical imager parameters are changed in the scope of the previous definitions: the gain in x-CLEAN controls the fraction of the PSF subtracted at each iteration. In CS, the gain is the relaxation parameter μ of Eq. (10), which controls the convergence rate of the algorithm. The threshold value in x-CLEAN is a flux density value usually associated with n times the level of noise measured (or expected) in the residuals (or in an source-free patch of the dirty image). Setting this threshold to zero would basically lead x-CLEAN to false detections from the background features. As discussed above in Sect. 2.6, the CS thresholding parameter is defined for each or all bands in the chosen dictionary. The number of iterations for CoSch-CLEAN algorithm is set to a high value ($>10^4$ to set an unrestrained convergence to the threshold). The number of iterations for CS can be set to an arbitrary value or be controlled by a convergence criterion.

The following figures of merit will measure the quality of the reconstruction using CS and CoSch-CLEAN methods:

- the total flux density computed in a region **S** around the source,
- the residual image standard deviation (Std-dev) and root mean square (rms) computed over **S** in the residual image or over the full image **I**,
- the error image and the mean square error (MSE) computed over **I** when a model image is available, and
- the dynamic range (DR) defined as the ratio of the maximum of the image to the residual Std-dev.

We here present different benchmark reconstructions arranged in increasing source complexity. In Sect. 3.2, we reconstruct point sources in two typical situations implied by new-generation radio interferometers: at small scale, by reconstructing the image of two point sources at different angular separation (Sect. 3.2.1), and at large scale, by reconstructing high-dynamic and wide-field images using a grid of point sources (Sect. 3.2.2).

In Sect. 3.3 we monitor the reconstruction quality of extended emissions: from a simulated W50 observation (Sect. 3.3.1) and a real LOFAR observation of Cygnus A (Sect. 3.3.2).

3.2. Unresolved sources

3.2.1. Angular separation between two sources

When considering point sources, the x-CLEAN algorithms are usually the most frequently adapted reconstruction method. The first test is therefore to compare the performance of CS reconstruction to that of CoSch-CLEAN for that simple type of source. We generated empty datasets (dataset **A** in Table 1) describing a simulated one-hour LOFAR observation. We only used core LOFAR stations, since they are included in most LOFAR observation. This justifies the evaluation of the ground performance of CS on this array subset. The layout of the core stations is slightly elongated (~10%, van Haarlem et al. 2013) in the north-south direction, providing a symmetric beam pattern to a direction close to zenith. In this dataset, we pointed the array to the local zenith, which will induce an a priori elongated beam pattern. The core stations in dataset **A** give a highest resolution of 1′ at $\nu_0 = 150$ MHz (in the HBA band). In addition, we restricted the (u, v) coverage to the radial distances $[r_{(u,v)}^{min}, r_{(u,v)}^{max}] = [0.1 \text{ k}\lambda, 1.6 \text{ k}\lambda]$ to artificially impose an angular resolution of ~2′ (which is effectively close to ~2.8′ as deduced from the PSF).

We simulated two one-Jansky point sources at frequency ν_0. With this simple sky model, we used the blackboard self-cal program (BBS; Pandey et al. 2009) to fill the LOFAR dataset with simulated visibilities. One source is located at the phase centre, the second is located at varying angular distances $\delta\theta$ ranging from 10″ up to 5′. As a consequence, we span various angular separation across the instrumental limit given by the PSF. We can define three different regimes where the performances of CLEAN and CS are compared: i) sources are unresolved;
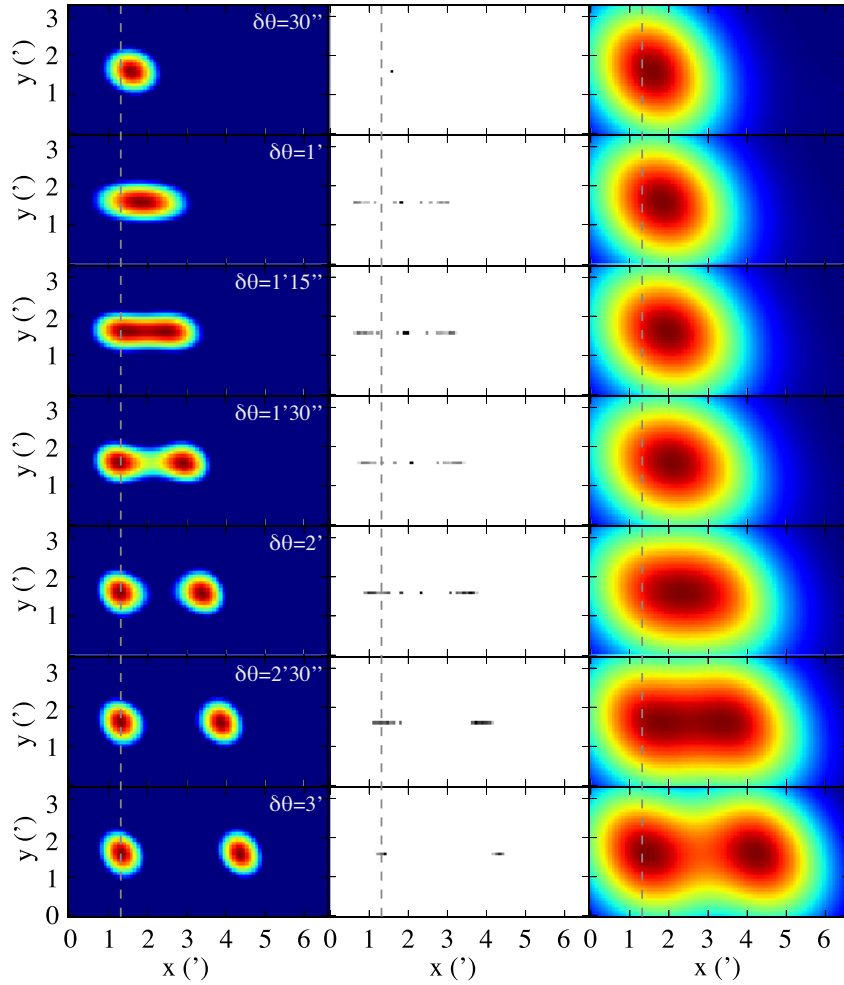
Fig. 2. Resulting reconstructed CS image (*left column*), CoSch-CLEAN model (*middle column*), and convolved (*right column*) images obtained from a simulation of two point sources with different angular separation $\delta\theta$ from $30''$ to $3'$ by steps of $30''$ (*from top to bottom*) plus the $\delta\theta = 1'15''$ case. The *third column* was obtained by convolving the CLEAN components (*middle column*) with the CLEAN beam of *FWHM* $3.18' \times 2.55'$. The effective separation of the two sources, after imaging by the different methods, occurs at smaller angular separation with CS (between $\delta\theta = 1'$ and $\delta\theta = 1'15''$) than with CoSch-CLEAN (between $\delta\theta = 2'$ and $\delta\theta = 2'30''$). The unambiguous separation of the two sources was obtained within few pixel errors, starting from $\delta\theta = 1'30''$ for CS and $\delta\theta = 3'$ for CoSch-CLEAN. Each cropped image was originally of size 512×512 pixels of $5''$. The (u, v) coverage has been restricted to the $[0.1\,\mathrm{k}\lambda, 1.6\,\mathrm{k}\lambda]$ domain to enforce an artificial resolution of $\sim 2'$. The colour map scales are normalized to the maximum of each image, but colour bars were not represented for clarity. The contrast of the CLEAN components (*middle column*) was enhanced to ease their visibility. The grey vertical dash line marks the true position of the source located at the phase centre of the dataset.

ii) partially resolved; or iii) unambiguously resolved. As we are close to the phase centre, the effects mentioned in Sect. 2.5 are negligible (which will no longer be the case in Sect. 3.2.2). For different angular separations, we injected Gaussian noise on the visibilities to obtain effective signal-to-noise ratio (S/N) levels of ~3, ~9, ~16, and ~2000 (noise-free case) in the image.

We used the original CoSch-CLEAN and our CS method implemented in AWimager, which serves here as a common testing environment for this study. We performed 10^4 iterations for the CoSch-CLEAN and 200 iterations for CS. As different levels of noise were simulated, we set the k parameter (Eq. (11)) to three times the noise level. We used for both algorithms the (u, v)-truncated dataset as described above and produced 512×512 pixel images at $5''$ pixel resolution. We imposed a *natural* weighting scheme to improve the S/N of the resulting images.

Figure 2 illustrates the results obtained with the noise-free dataset: the CS reconstruction (left column), the CoSch-CLEAN model image (middle column), and the corresponding *convolved* image (right column). The rows correspond to seven values of

angular separation $\delta\theta$ from $30''$ to $3'$ by steps of $30''$ and the $\delta\theta = 1'15''$ case. For the CS reconstruction, there is no model image in the CLEAN sense because the output of CS is directly the best representation of the true sky at a finite resolution (similar to the comparison of models in Li et al. 2011b). Different angular separation criteria can be chosen. One can apply the Rayleigh criterion based on the separation of two (pixel) CLEAN components on the model image or use source-finders on reconstructed images (see below). Here (and hereafter), the CLEAN beam, setting the highest angular resolution of CLEAN deconvolved images, has a size of $3.18' \times 2.55'$ (major $\times$ minor axes). In the CoSch-CLEAN model image, when $\delta\theta$ is small, only one group of CLEAN components is detected at the mean position of the two sources. Starting from $\delta\theta = 1'30''$, additional CLEAN components are being detected on the *wings* of the main component, but are located correctly on the source position. With $\delta\theta$ increasing, the amplitudes of these side CLEAN components increase and contribute in the resulting elongated shape on the convolved CLEAN image. Starting from $\delta\theta = 2'30''$, the two sources are unambiguously separated into two distinct groups

of components around the correct position. The position uncertainty of these two sources is still high ($\sim 30''$). For separations larger than $\delta\theta = 3'$, the astrometric and flux density errors start to decrease. In the regime where CoSch-CLEAN cannot unambiguously resolve the two sources, the features present in the model images (the central components as well as the wing components) cannot be associated with real sources. In the scope of this method, an exploitable image with limited resolution can only be obtained after convolving with the CLEAN beam.

With CS, we directly obtain the best estimate of the sky, as shown in the reconstructed image in the left column of Fig. 2. In the partially resolved regime, we also note that the elongation of the source size occurs at $\delta\theta = 1'$. We note that the two point sources are resolved at lower angular separation ($\delta\theta \sim 1'15''$) than for CoSch-CLEAN. An elliptical fit of the FWHM of the sources gives the source size, which can be seen as the *effective* CS convolution beam. In this case, its dimensions are $1.55' \times 1.09'$, representing a beam of cross-section approximately $1.39'$, which is smaller than that of the CLEAN beam. In addition to the improved angular separability, the CS reconstructed sources are also correctly located in the image (e.g. the central source lies close to the vertical mark in Fig. 2 for $\delta\theta \leq 1'30''$) as compared to the CoSch-CLEAN model image, where the CLEAN components are only correctly positioned starting from $\delta\theta \leq 2'30''$. In the noise-free case, by using exactly the same dataset and imaging parameters, these results suggest that CS is able to recover information on the true sky beyond the theoretical resolution limit (which is constant in this dataset). Because there is always a non-zero level of noise in the calibrated interferometric data, we tested the reliability of this characteristic against the image S/N by using the different noisy datasets mentioned above.

To compute the effective angular separability of the sources using CoSch-CLEAN and CS, we used the LOFAR pyBDSM[3] package (*python blob detection and source measurement*), which consists of island detection, fitting, and characterization of all structures in the image. We defined a detection threshold so that no other artefact could be detected as a source in the images because we only focused on the two simulated point sources. We used the same set of detection parameters for CoSch-CLEAN and CS images.

At a specific S/N level, the limit value of angular separation at which two distinct sources can be distinguished by the source finders gives an estimate of the separability angle of the sources and therefore of the effective resolution that limits the typical size of genuine physical features. The CoSch-CLEAN beam and the effective CS beam are not circular, therefore the absolute effective angular resolution may depend on the orientation of these beams with respect to the direction joining the two sources. To have a measure that is independent of this direction, it is probably better to compute the ratio between the beam elliptical parameters obtained with CoSch-CLEAN and that obtained with CS.

Using an appropriate sampling of $\delta\theta$ and noise levels, we can build a graph of this effective resolution deduced from the separability of the two sources for different S/N levels. We present in Fig. 3 the resulting curves obtained with CoSch-CLEAN (black) and CS (red). For each S/N we derived the effective improvement brought by CS by noting the angle $\delta\theta_{\min}$ at which the two sources are separated. On the one hand, the CoSch-CLEAN separability has a limited dependence on the S/N, which makes it a
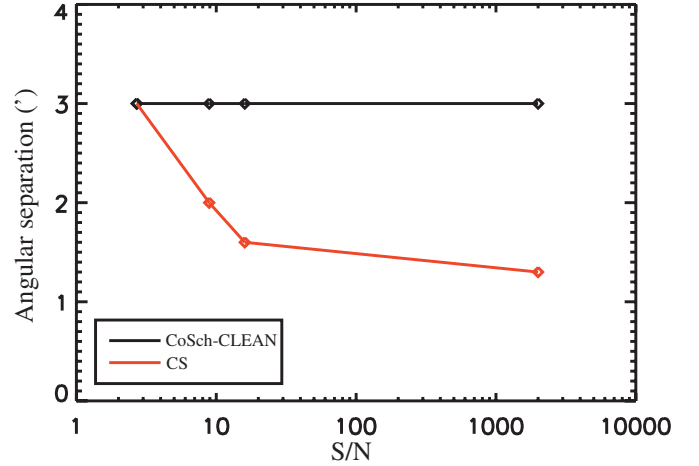

Fig. 3. Values of $\delta\theta_{\min}$ of source separability as a function of the S/N for CS (red curve) and CoSch-CLEAN (black curve). In the moderate ($S/N \geq 5$) and high S/N regimes, the resulting source separability is improved by a factor 1.3 to $\lesssim 2$.

stable and reliable algorithm for detecting point sources in various S/N regimes. On other hand, CS behaves differently. At a high S/N regime, the separability values outperforms that of CoSch-CLEAN by a factor of 2–3 in the [10, 2000] S/N range. When the S/N decreases, the separability of CS tends to that of CoSch-CLEAN.

In terms of astrometric error, the detected point source locations can be compared to the (α, δ) coordinates of the input sky model (where the sources were placed at pixel centres). The relative position errors do not exceed $1'$ for CS and $3'$ for CLEAN for all the different noise levels.

In addition to the angular separability and astrometric errors, we inspected the flux density of the source.

Naturally, the flux density error of the reconstructed sources is affected by the level of the background noise and scales with the S/N. It is found to be 3% in the low-noise regime and up to 25% in the high-noise regime for CS and 3% to 23% for CoSch-CLEAN.

From this perspective, it appears that CS provides results that are almost as good as the results by CLEAN by default, and provides an improved angular separability with high and moderate S/N data. The low astrometric and flux density error confirm the superior resolution capability of CS, which suggests the possibility of dramatically improved angular resolution of extended emission from poorly sampled interferometric data.

3.2.2. Wide-field imaging of a grid of sources

Interferometry imaging at low frequencies implies a larger FoV because of the size of the LOFAR station analogue beam. By using AWimager, which deals with wide FoV (Tasse et al. 2013), we checked the ability of CoSch-CLEAN and CS to recover the correct flux density of point sources that are away from the phase centre. We simulated a 10×10 grid of point sources over a region of $8° \times 8°$ around the phase centre. Their flux densities range from 1 to 10^4 Jy. We used BBS to fill dataset **B** by enabling the simulation of the beam (for the *A-projection*). The *W-projection* only depends on the layout of the interferometer, which is included in the dataset. Noise was injected into the dataset to provide an rms value of 10^{-4} (so ~ 1 Jy) relative to the peak of the dirty image. The points are located at equidistant vertices in the grid, which enabled us to examine the distribution of the flux
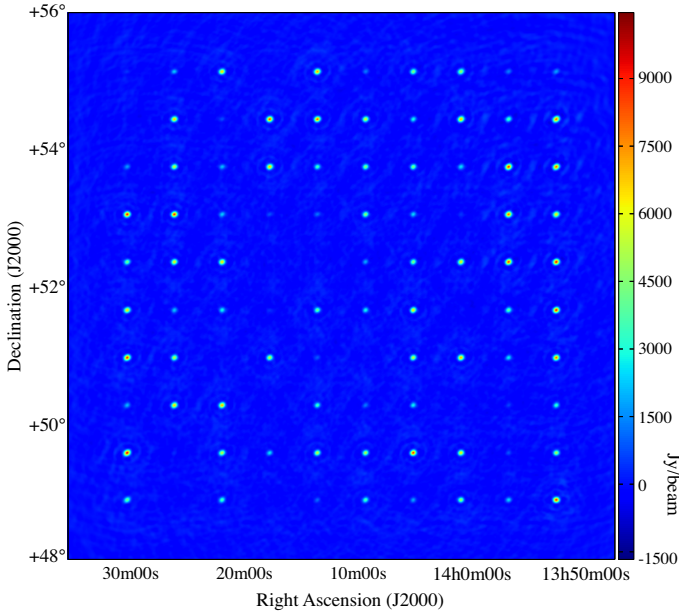
Fig. 4. Dirty image derived from a set of 100 simulated point sources with flux density ranging from 1 to 10^4 Jy. Simulated visibilities are generated with dataset **B**. Some sources are not yet visible because of the dynamic of the source and the noise level.

Fig. 5. *Top*: total flux density reconstruction for a set of point sources (Fig. 4) with original flux density spanning over 4 orders of magnitude with the original flux density (*x*-axis) vs. the recovered flux density (*y*-axis). *Bottom*: scatter plot of the absolute the error for each source. The recovered flux densities for CoSch-CLEAN (red) and CS (blue) are represented on a linear scale, whereas the absolute error is on a logarithmic scale for clarity. Perfect reconstruction lies along the black line. The output flux density values and errors are similar with both CoSch-CLEAN and CS.

density and the potential residual distortions over the field without overlapping of the source. While this arrangement is unrealistic, we can still monitor the astrometric and flux density accuracy versus the radial distance from the phase centre. We applied CoSch-CLEAN and CS to the simulated dataset containing the grid. The dirty image generated from visibilities is shown in Fig. 4.

We used 1024×1024 pixel images with a pixel size of $28''$ and the full (u, v) coverage of dataset **B** was used for imaging, giving an effective angular resolution of $\sim 3''$. For CoSch-CLEAN, we used 10^8 iterations and for CS, we used 200 iterations. With CS, the sources were reconstructed with the starlet dictionary. We report in Fig. 5 the output integrated flux density of all detected sources in the reconstructed images against the input flux densities of the sky model point sources. A sound reconstruction should put every source on the first bisector. We again used pyBDSM to perform the source detection and characterization (including the 2D elliptical Gaussian fit of the source, position of the Gaussian barycentre, and photometry along with providing respective errors). The error bars of each point were obtained directly from the photometry and are negligible given the low level of the background noise. However, these error bars do not include the bias of poorly reconstructed sources.

For the CS reconstructed image, the effective source size is reduced to smaller patches with a few pixels each around the sources (comparable to the CLEAN model, only constituting a collection of pixels). The flux is more efficiently gathered around the source position, resulting in an increase of all pixel size values in the CS images (in Jy/beam). The resulting source size was $3.2' \times 3.7'$ (i.e. the dimension of the CLEAN beam) for CoSch-CLEAN and between $\sim 30'' - 1'$ (1–2 pixels) for CS. This result corroborates that of the previous study. The superior resolution here yields an improvement of a factor of 3 to 6 of the source size compared to the size of the CLEAN beam, using exactly the same dataset (and weighting) for the two methods.

The flux density rms error was derived from the residual images and accounts for 3.6 Jy/beam Std-dev of CS and
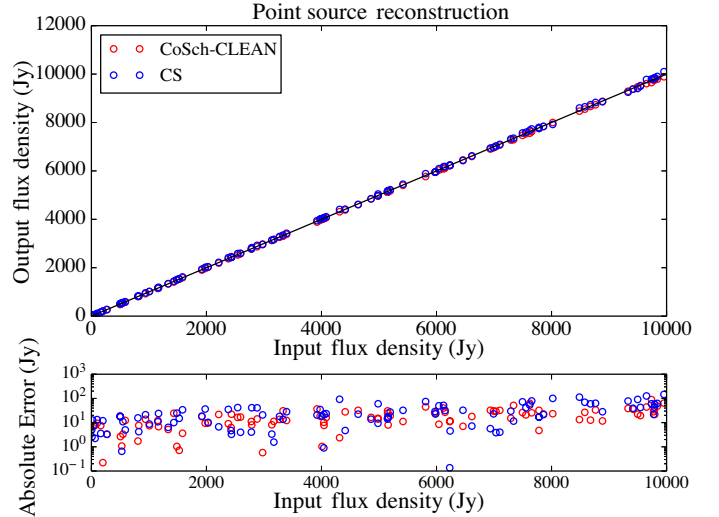
1.7 Jy/beam Std-dev for CoSch-CLEAN. The error bars are not reported in the plot for clarity. The relative error (compared to the input flux density of each source) is not larger than 10% in most cases for both CoSch-CLEAN and CS. The flux density error slightly increases with the source flux density, as depicted by the scatter plot represented in a log-scale in Fig. 5 (bottom).

As a result of the many strong sources, CoSch-CLEAN and CS were unable to reconstruct some faintest sources barely above the background level. CoSch-CLEAN presents slightly better performances for reconstructing correct flux densities with a lower error (the mean absolute error is 19 Jy for CLEAN and 29 Jy for CS). Nevertheless, CS led to the detection of more faint sources that were missed by CoSch-CLEAN, but with a larger error on their flux density values. In spite of its improved angular resolution and its detection capability, CS provides a larger Std-dev (3.6 Jy/beam) on the residual image than CoSch-CLEAN (1.7 Jy/beam). This could be an effect of the thresholding taking place in CS or the choice of the dictionary, which is not perfectly fitted to represent point sources.

We did not note any particular dependence on the radial distance from the phase centre for either CoSch-CLEAN or CS. This suggest that the *A*- and *W-projections* in AWimager correctly prevented a radial dependence of the flux, the astrometric error, and the source distortion (the last two being at the scale of one pixel). This tell us that the CS method is effectively compatible with the RIME framework. Moreover, x-CLEAN algorithms remain competitive with CS because the spatial extension of wavelet atoms is not as efficient as Dirac atoms (pixel) in representing single point sources.

3.3. Extended sources

We have shown that CS and CLEAN are clearly competing when imaging point sources. We now address the reconstruction of extended radio emission, which are not sparsely represented in a pixel dictionary. We therefore continue to use the starlet
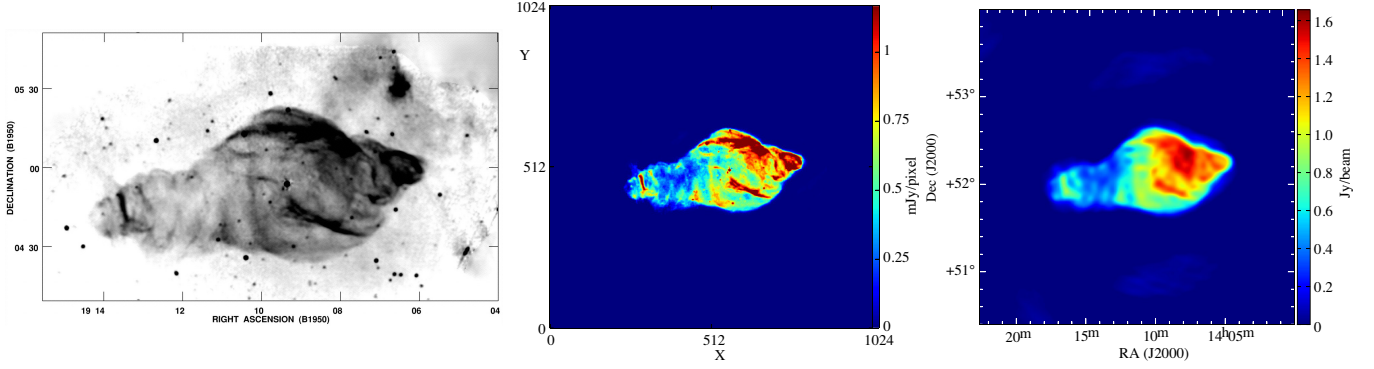
Fig. 6. *Left*: original 1.4 GHz radio image from Dubner et al. (1998) (pixel size = 6.75″, ang. resolution = 55″), *middle*: prepared input model image scaled to the approximate total flux density of 230 Jy (1024 × 1024 with pixel size of 13.5″), and *right* the dirty image obtained with dataset **B**. All emission falling outside the extended emission of W50 was set to zero, but point sources inside W50 were kept in the simulation.

Table 2. Relevant imaging parameters used by AWImager for CoSch-CLEAN, MS-CLEAN, and CS, for the two extended objects.

Algorithm	W50			Cygnus A		
	CoSch-CLEAN	MS-CLEAN	CS	CoSch-CLEAN	MS-CLEAN	CS
Specific parameter	–	5 scales	–	–	5 scales	–
Weight	Briggs (robust = 1)			Uniform		
Planes for W-projection	11					
Image size	1024×1024			512×512		
Padding	1.6			1.5		
Pixel size	13.5″			1″		
Threshold	10^{-3} Jy	10^{-3} Jy	10^{-3}	0.5 Jy	0.5 Jy	0.5
Gain	0.01	0.4	0.4	0.05	0.05	0.05
N_{iter}	10^8	10^3	200	3.10^4	10^4	200

Notes. For MS-CLEAN, we have used scales at [0, 10, 30, 70, 100] pixels for W50, and [0, 2, 5, 10, 15] pixels for Cygnus A.

dictionary for the CS reconstruction to ensure the best sparsity of the signal. We first study a simulation of W50 (Sect. 3.3.1) and then discuss the results of CS on a real LOFAR observation of Cygnus A (Sect. 3.3.2).

3.3.1. Simulated observation: the W50 nebula

The W50 nebula (hosting the SS 433 microquasar) is an extended supernova remnant of large dimension ($\sim 2° \times \sim 1°$) with internal filamentous structuring that makes it proper for benchmarking CS and CLEAN-based algorithms. First, we used the W50 brightness image from Dubner et al. (1998) (Fig. 6 left). This image at 1.4 GHz is rich in information at various angular scales down to an angular resolution of $\sim 55″$ (the image at 327.5 MHz, taken with the VLA in the D configuration, was also available, but no structures smaller than 70′ are resolved). We assumed that the higher spatial frequency features also emit in the LOFAR band (see W50 observed with LOFAR HBA in Broderick et al. in prep. and previous work). We focused on the central extended emission by removing part of the extraneous foreground and background features. We set the flux scale of the total model image to match a total flux density of ~250 Jy at 116 MHz as interpolated from Dubner et al. (1998). Second, we used the *predict* task of AWImager to convert the input model to visibilities in dataset **B**. No artificial DDEs were inserted in the simulation but *A-projection* and *W-projection* were enabled for all runs. Dataset **B** provides a theoretical angular resolution of $\sim 3″$, which is higher than the 55″ of the original image. We therefore expect to have good results for all three methods. The *predict* step samples the image at the resolution of dataset **B**.

Given the size of W50, the simulated field is $3.8° \times 3.8°$. Third, artificial Gaussian noise was added to the simulated visibilities and has an effective rms of 3×10^{-4} of the peak of the dirty *noise-free* image (corresponding DR ~ 4000). We reconstructed images using three methods: CoSch-CLEAN, MS-CLEAN, and CS in AWImager. MS-CLEAN was imported from the LWimager, the standard LOFAR imager, superseded by AWimager.

With the input model image, we computed the error image and inspected the residual images to track the effective angular resolution and the convergence of the methods. Table 2 gathers all imaging parameters (as well as for Cygnus A). We compare in Fig. 7 the outputs of CoSch-CLEAN (left column), MS-CLEAN (central column), and CS (right column). From top to bottom we present the reconstructed image, the model, the residual image, and the error image. The figure of merit from these images is gathered in Table 3. The total reconstructed flux densities over the source are 229.2 Jy, 229.4 Jy, and 229.7 Jy for CoSch-CLEAN, MS-CLEAN, and CS. These values are extremely close to the 230 Jy of the model, which verifies that all three methods conserve the total flux density. This fact is a basic requirement that any new imaging method should verify.

The extended emissions were rendered with different accuracy, and a particularly good image was produced with CoSch-CLEAN because of the low noise level in the data, the high number of iterations, and the size of the PSF. In comparison, the MS-CLEAN image presents a lower angular resolution despite taking into account the various scales of the image. Because it is a user parameter, we tried different sets of scales, thresholds, and number of iterations, but we did not improve the result, and sometimes it did not converge. This potential divergence may be
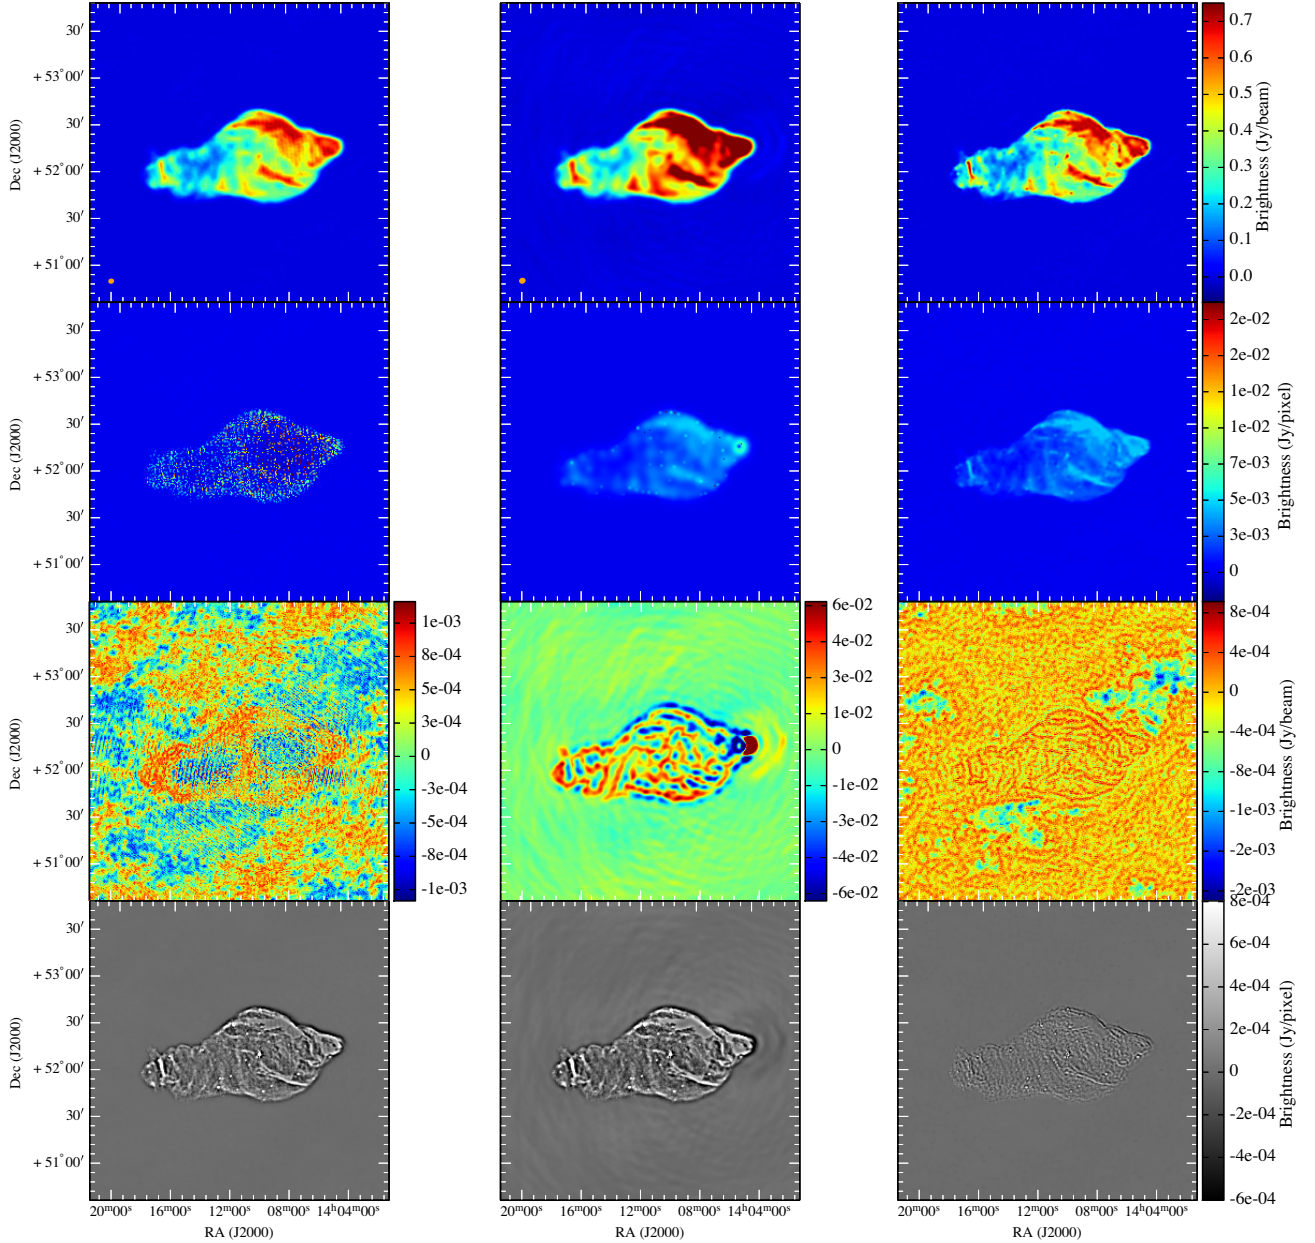
Fig. 7. Reconstructed images of W50 from the simulated LOFAR observation (dataset **B**) using CoSch-CLEAN (*left column*), MS-CLEAN (*middle column*), and CS (*right column*). *From top to bottom*: the restored (in Jy/beam), the model (in Jy/pixel), the residual, and the error images. The CS restored and model images only differ by their scaling (the former is in Jy/beam using the beam area of the CoSch-CLEAN beam, the latter is in Jy/pixel). The colour scales of the images are different for each row, as indicated by the colour bar on the right. The residual images are displayed with their respective colour bar. The CS reconstruction contains high spatial frequency information restored from the dataset. The effective angular resolution of the CS image (1′) is close to that of the original input image (55″). In addition, the error image shows a closer proximity between CS and the original input image.

caused by an inappropriate thresholding when the background residual level is reached, the algorithm then start to get signal from the background noise. In comparison, we also used the MS-CLEAN method of LWimager and obtained equivalently poor results. MS-CLEAN results might be improvable with additional controls and masks, but we chose a straightforward imaging (as our algorithm) to highlight the robustness of our algorithm with a limited number of user controls. Moreover, the implementation of MS-CLEAN in AWimager is still experimental and needs to be more extensively validated. This method is known to usually improve the representation of extended sources.

The restoring beam was 2.8′ × 2.4′ for CoSch-CLEAN and 2.8′ × 3.3′ for MS-CLEAN. The CS image is visually sharper

than the other two and contains high-frequency features that were recovered during the imaging. From the typical point source size in the CS image, the effective angular resolution was ∼1′ × ∼1′, nearly equal to the 55″ original resolution, and represents an improvement of almost a factor 2.5–3 over x-CLEAN images.

With MS-CLEAN, the extended features are represented by large blobs in the model, along with some pixel-sized CLEAN components. As described earlier, the output of CS comes directly as the best estimate of the sky and is in units of Jy/pixel in the model image. We multiplied the CS model image by the beam area to obtain the same brightness units (Jy/beam) as for the other images. Therefore, the model image and the restored

Table 3. Statistical results obtained from applying CoSch-CLEAN, MS-CLEAN, and CS to all the datasets **A**, **B**, and **C**.

Point source (Sect. 3.2.1)	CoSch-CLEAN	CS
Relative flux density error high S/N, low S/N	3%–23%	3%–25%
Effective angular resolution (′)	$3.18' \times 2.55'$	$1.55' \times 1.09'$
Astrometric error (′)	<2′	<1′
Grid of point sources (Sect. 3.2.2)	CoSch-CLEAN	CS
Relative flux density error (Jy)	Fig. 4	
Residuals: Std-dev in **I** (Jy/beam)	1.7	3.6

W50 (Sect. 3.3.1)	CoSch-CLEAN	MS-CLEAN	CS
Effective angular resolution (′)	$2.8' \times 2.4'$	$2.8' \times 3.3'$	$\sim 1' \times \sim 1'$
Reconstructed: total flux density in **S** (Jy)	229.19	229.41	229.72
Reconstructed: peak flux density in **I** (Jy/beam)	0.72	0.96	0.79
Residuals: rms in **S** (Jy/beam)	4.6×10^{-4}	2.71×10^{-2}	3.1×10^{-4}
Residuals: Std-dev in **I** (Jy/beam)	4.1×10^{-4}	1.3×10^{-2}	3.2×10^{-4}
Error: mean square error in **I** (Jy/pixel)	8.9×10^{-5}	1.0×10^{-4}	4.8×10^{-5}
DR	1729:1	75:1	2446:1
Cygnus A (Sect. 3.3.2)	CoSch-CLEAN	MS-CLEAN	CS
Effective angular resolution (′)	$6.1'' \times 5.0''$	$6.1'' \times 5.0''$	$2.8'' \times 2.8''$
Reconstructed: total flux density in **S** (Jy)	10 576	10 560	10 506
Reconstructed: Peak flux density in **I** (Jy/beam)	315.5	309.3	465.7
Residuals: rms in **S** (Jy/beam)	0.37	0.35	0.09
Residuals: Std-dev in **I** (Jy/beam)	0.17	0.19	0.05
DR	1799:1	1619:1	8392:1

Notes. The statistics are defined in the text and appear when applicable. **I** means that the quantity has been evaluated on the entire image, **S** in a defined region of the image (typically around the source).

image only differ by the residual image that was added to the latter. CS tends to concentrate even more the flux of point sources around the source (see Sect. 3.2.1), and the starlet dictionary is able to correctly represent the extended emission.

Any vestigial structures in the residual image show a lower rate of convergence, which might be due to an insufficient number of iterations (which is currently not the case for the x-CLEAN algorithms), or a limitation due to the imaging method itself. In our case, the CoSch-CLEAN image, supposed to perform better with point sources than with extended emission, presents a lower residual Std-dev level (4.1×10^{-4} Jy/beam) than that of MS-CLEAN (1.3×10^{-2} Jy/beam) in the present situation. However, the CS residual Std-dev level is slightly better (3.2×10^{-4} Jy/beam) across the image. The error images are shown in the last row of Fig. 7 and are the difference between the restored image (converted to Jy/pixel) with the initial full resolution input image (also in Jy/pixel). The input image was not degraded by convolution to match a particular resolution. From the different error images, CS visually provides the lowest error. The high-resolution features were partly restored by the CS reconstruction. The CS algorithm has the lowest mean square error value over the image (4.8×10^{-5} Jy/pixel), which represents an improvement of a factor ~2 compared to that of present CoSch-CLEAN and MS-CLEAN reconstructed images.

3.3.2. Real LOFAR observation: the Cygnus A radio galaxy

Cygnus A (3C 405) is one of the most powerful galaxies observable in the radio domain. It serves as a good test case for benchmarking our method: it contains extended emission, originating from two main radio lobes (representing a projected size of ~2′ × 1′) as well as compact emissions at the extremities of each lobe (the hotspots A and B in the western lobe and hotspot D in the eastern lobe, as labelled in Hargrave & Ryle 1974). Cygnus A has recently been observed by LOFAR during its commissioning phase (McKean et al. 2010), but has also been long observed in the past at low frequencies for instrumental calibration or science (e.g. Lazio et al. 2006 and references therein, who observed Cygnus at 74 MHz and 327 MHz with the VLA). We used a real calibrated LOFAR dataset (**C**) at 151 MHz containing real noisy measurements. The data were previously calibrated on the Perley-Butler (2010) absolute flux scale (McKean et al. 2010).

Similar imaging parameters were used (Table 2). With the FoV being relatively small and the source being the dominant source in the dataset, the DDEs were expected to be low. In Fig. 8, we show the same output images as in Sect. 3.3.1. However, we do not have an input model image to compute the error image. From left to the right, columns are the reconstructed image with CoSch-CLEAN, MS-CLEAN, and CS. From top to bottom, the rows present the reconstructed images in Jy/beam, the model images in Jy/pixel, and the residual image.

All the three methods rendered a proper image of the emission, given the extremely good quality of the data and the strong source brightness. The CoSch-CLEAN image shows residual high-frequency structures that lead to distortions of the source. These artefacts mainly come from CLEAN components collected from the noise background residual image. With MS-CLEAN, we obtained a similar brightness distribution to that in McKean et al. (2010). The restoring beam size was $6.1'' \times 5.0''$ with both x-CLEAN methods. The CS effective angular resolution was ~$2.8'' \times 2.8''$, representing again an improvement of
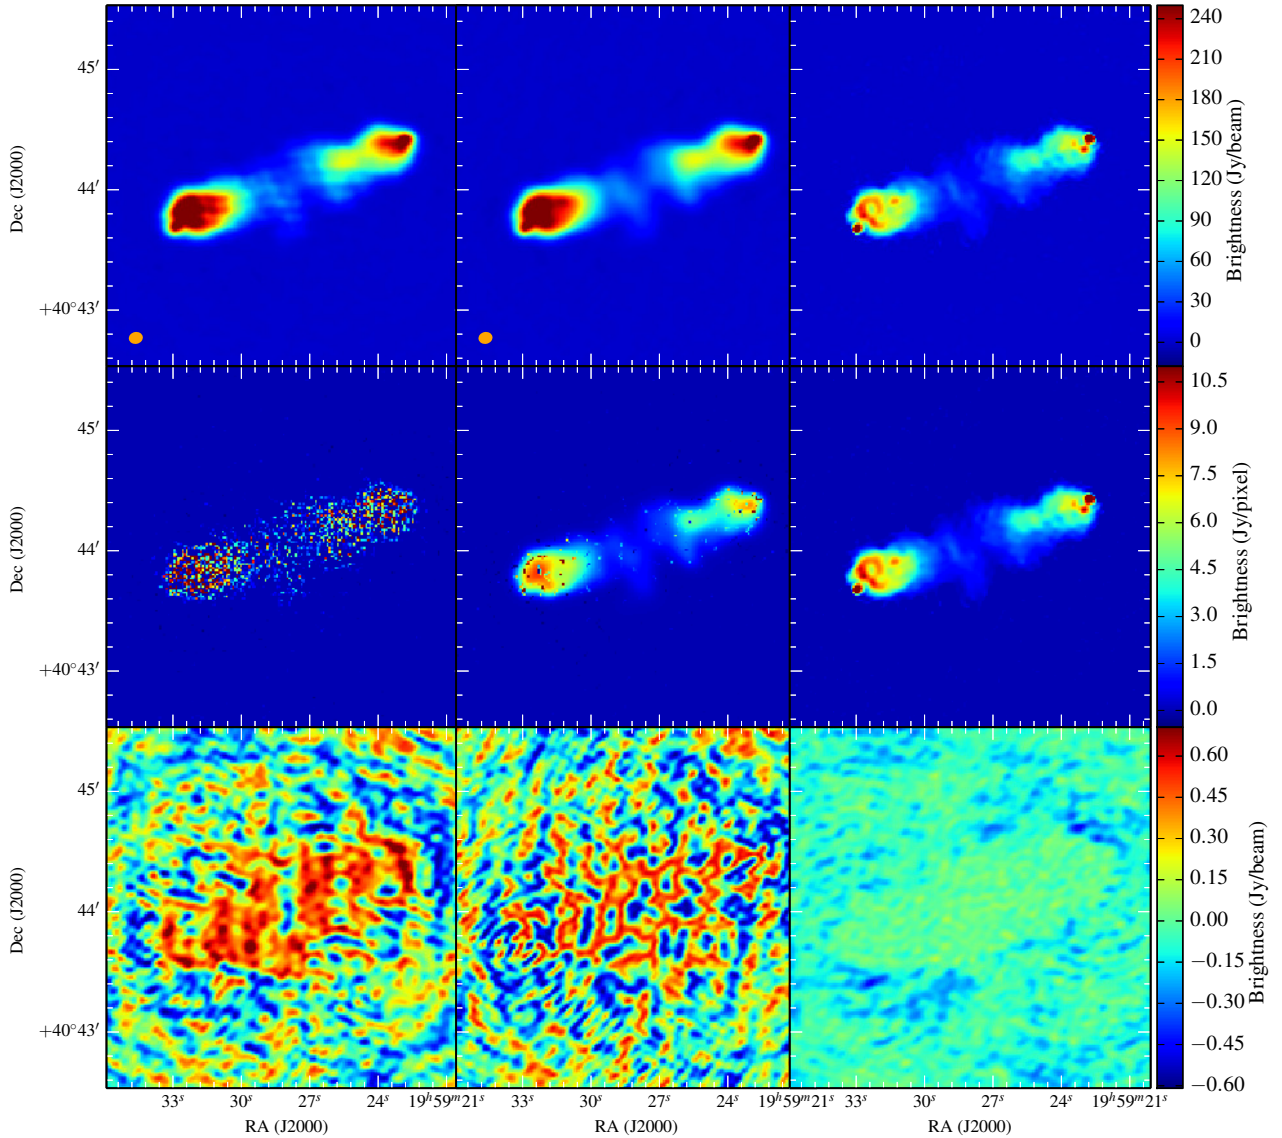
Fig. 8. Reconstructed images of Cygnus A from the real LOFAR observation (dataset **C**) using CoSch-CLEAN (*left column*), MS-CLEAN (*middle column*), and CS (*right column*). *From top to bottom*: the restored images, the model images, and the residual images. There is no error image, as its calculation would require a highly resolved image of the true brightness distribution of Cygnus A to compare with the recovered image. The colour scales of the images are different in each row, as indicated by the colour bar on the right. The CS reconstruction presents a higher angular resolution and a lower residual level (see text) than that obtained with the two other methods.

a factor of 2. The resulting CS image shows a much sharper reconstruction of the structures inside the lobes. In particular, hotspots A, B, and D were correctly rendered without any ambiguity compared to the result obtained with both CoSch-CLEAN and MS-CLEAN. To validate the reality of the features in the CS image, we overlaid contours from VLA data at 327.5 MHz (resolution of 2.5″) in Fig. 9. The location of hotspots and filaments and holes inside each radio lobe match that of the VLA image. This confirms the improvement in angular resolution brought by CS.

As for W50, the CoSch-CLEAN model image is only composed of pixel CLEAN component sources, the MS-CLEAN image displays a smoother distribution of the source, based on the multi-scale decomposition and the CS (model) image is mainly dominated by the signal of the hotspots situated at each extremity of the lobes. We checked the scientific relevance of the reconstructed images by inspecting the figures of merit gathered in Table 3. The first quantity is the total integrated flux density of

the source. This value can be measured with short spacings (ideally with a 0-length baseline or with the autocorrelations of the antennas) or by fitting the expected visibility function in amplitude vs. the (u, v) radial distance and measuring the Y-intercept of this curve. In the present case, the observing frequency was 151 MHz and the total expected flux density is $\gtrsim 10\,500$ Jy from the visibility data. A correct image reconstruction is therefore obtained when the total flux density present in the reconstructed image is close to that value. We found 10 576 Jy, 10 560 Jy, and 10 507 Jy for CoSch-CLEAN, MS-CLEAN, and CS. This range of total flux densities is compatible with previous MERLIN and LOFAR observations taken at the same frequency (Steenbrugge et al. 2010; McKean et al. 2010), within a 2% accuracy. We cannot conclude on the superiority of one method over another, but these values suggest again that CS conserves the total flux (as seen for W50) even with (real) noise, similarly to the two other methods. The residual images are represented with the same colour scale in Fig. 8. We measured the rms level over the same
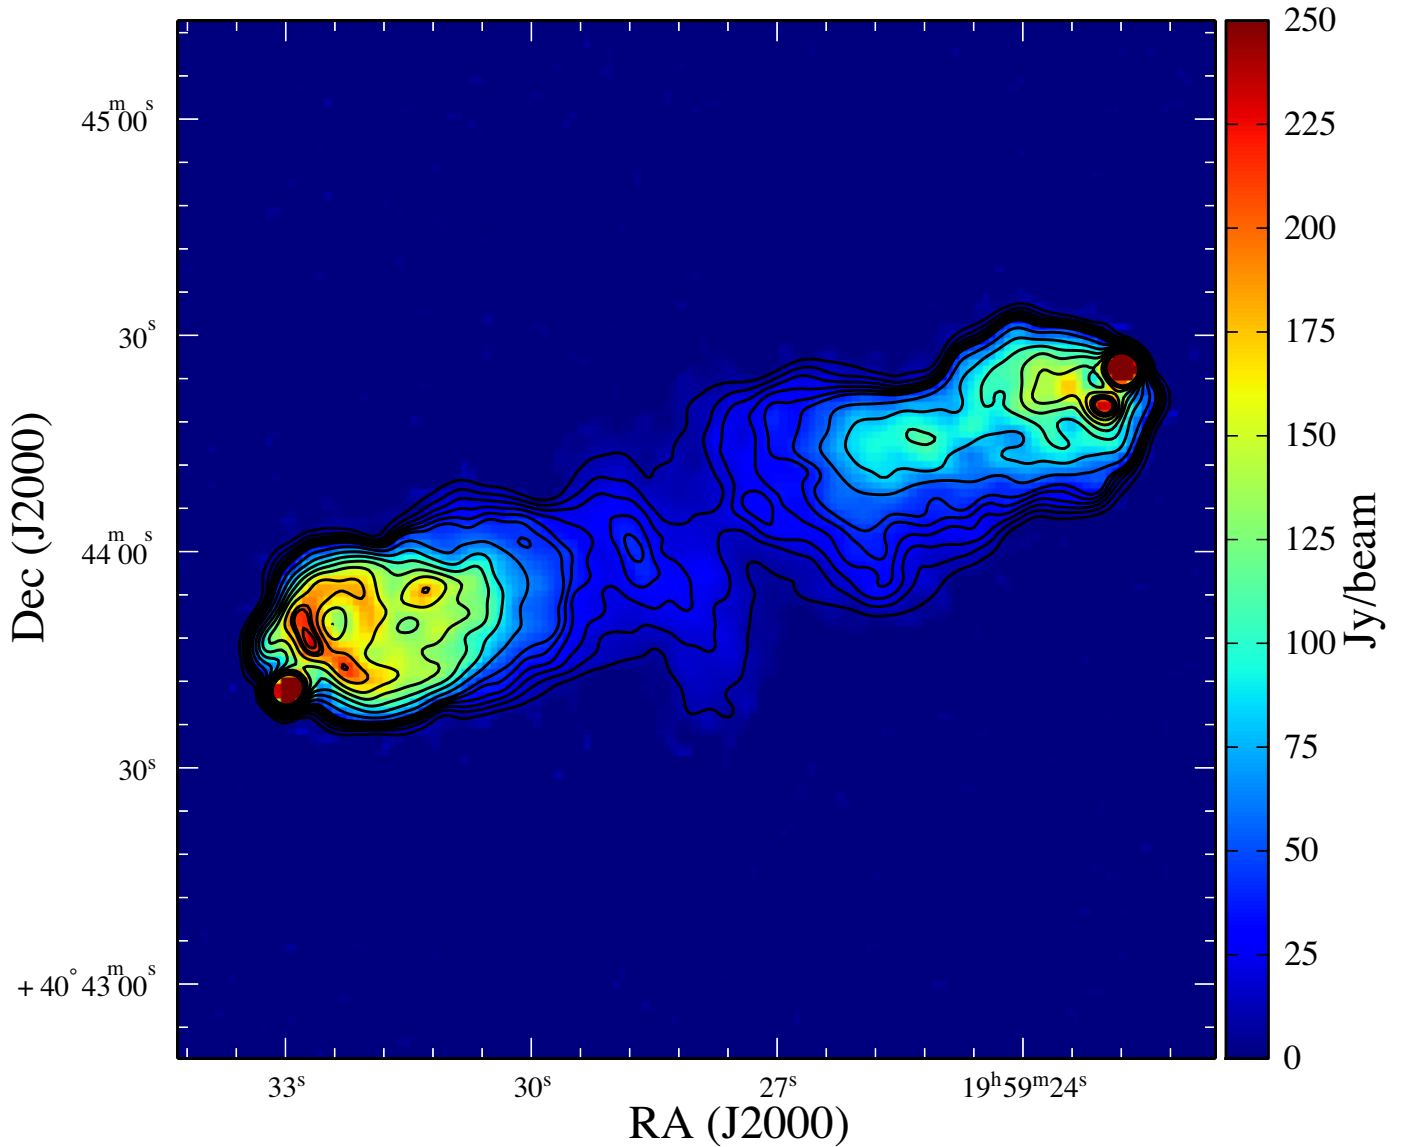
Fig. 9. Colour scale: reconstructed 512×512 image of Cygnus A at 151 MHz (with resolution $2.8''$ and a pixel size of $1''$). Contour levels (Toggle contours) are $[1, 2, 3, 4, 5, 6, 9, 13, 17, 21, 25, 30, 35, 37, 40]$ Jy/Beam from a 327.5 MHz Cyg A VLA image (Project AK570) at $2.5''$ angular resolution and a pixel size of $0.5''$. Most of the recovered features in the CS image correspond to real structures observed at higher frequencies.

region **S** of the flux density integration and the Std-dev of the entire 512×512 residual image. The CS residuals show an improvement of about one order of magnitude over CoSch-CLEAN and MS-CLEAN. This result is characteristic of a reliable reconstruction. The dynamic range (DR), computed as the ratio of the peak flux density to the residuals Std-dev, was computed for each image. The DRs are 1799:1, 1619:1, and 8392:1, suggesting that CS enhances the DR of the image. This is achieved by combining two effects: first, CS tends to concentrate the flux at the correct astrometric position, resulting in a higher peak flux density of the image; second, the low standard deviation of the residuals demonstrates a better convergence of the image reconstruction. The DR of the MS-CLEAN image is slightly lower than that of CoSch-CLEAN. This can be explained by the fact that the data are dominated by the source, making it easier to decompose in unambiguous CLEAN components, or by an unoptimized choice of the scale decomposition. From the present results, it appeared that the CS method is reliable enough to reconstruct

extended radio emissions coming from simulated data as well as real LOFAR data. The high flux density of Cygnus A and the quality of its calibration contributed to the quality of the CS reconstruction.

4. Discussion

4.1. Instrumental limitations and final convolution

The final step of CLEAN includes convolving the model image with the CLEAN beam in the pixel domain. Each detected point source finally has the shape of the CLEAN beam, resulting in a voluntary degradation of the information included in the model image. This is a way to express the "natural" angular resolution of the instrument, imposing a boundary on the size of the smallest "relevant" feature in the image. The choice of the final convolution beam (or any other beam) results also from an aesthetic rendering at the expense of realness. The CS output image

is assumed to be (or at least close to) the true sky under certain conditions (see Sect. 2.4). We demonstrated that it was possible to push the limit of angular resolution. The resulting CS model image does not have an infinite resolution (it would require an infinite Fourier support) and the equivalent CLEAN beam size is limited by the quality of the Fourier reconstruction at high spatial frequencies. Preliminary studies suggested that this size may improve with number of iterations. In addition, we chose not to convolve the final CS result with the CLEAN beam to preserve the recovered super-resolved features. As a consequence, the flux density of point sources is concentrated in smaller patches around the sources, resulting in higher pixel flux density values on Jy/beam images obtained with CS (and scaled with the CLEAN beam area).

4.2. Performance and convergence compared to CLEAN-based methods

We conducted several preliminary tests, which led to an educated guess of each CS parameter value used in this study. CoSch-CLEAN and CS only differ by the core of their minimization algorithm (iterative subtraction and iterative soft-thresholding). The run time of the imager is dominated by the time spent with operators $\mathbf{A}$ and $\mathbf{A}^{\mathrm{T}}$, used to convert *gridded* image data to *ungridded* visibilities (and vice versa) in the context the of RIME. As an example, with the W50 dataset, a single iteration of CoSch-CLEAN in AWimager) takes 1.88 s to run where CS takes 1.76 s. The number of iterations can vary significantly between CoSch-CLEAN, MS-CLEAN and CS and leads to different convergence rates that are not easily comparable. It tends to be more stable for CS because the latter always operates on the entire image during a single cycle. It appeared that the MAD estimation (Eq. (11)) is particularly reliable in the presence of a high level of noise, but produces similar results to fixed thresholding when the signal of the source is dominant in the data.

5. Conclusions and perspectives

5.1. Conclusion on the benchmark

The overall performance of CS, highlighted in this study, is very promising in the scope of LOFAR, and more generally in classical radio interferometry imaging. Throughout our studies, we can draw conclusions on the behaviour of CS:

– CS performs relatively well compared to CLEAN in almost all situations. It was able to correctly recover the positions and flux densities of single point sources while bringing an improved angular resolution at high and moderate S/N regimes, breaking the limit imposed by the chosen CLEAN beam (or more widely by the instrumental PSF). At low S/N regimes ($S/N \lesssim 5$), the angular separability of the CS applied to point sources falls back to that of CoSch-CLEAN.
– CS is compatible with the RIME framework (thus enabling the correction for ionospheric effects, complex and varying antenna beams, pointing errors, etc.) and it provides satisfying photometry when imaging wide FoV. Both CS and CoSch-CLEAN had the same difficulty in recovering the lowest flux density sources. The dynamic range might be improved by a correct fine-tuning of the imaging parameters with the different methods. A more sophisticated way of studying point flux density recovery would be to simulate many observations, each with a single point in the FoV, and at different random locations.

– CS is able to reconstruct extended emission with improved angular resolution. From the simulated dataset (W50) and the real LOFAR dataset (Cygnus A), CS has demonstrated its ability to recover unsampled regions of the visibility function, leading to the rendering of realistic features at higher spatial frequencies than those of the two other methods. The latter are more adapted to sparse distributions of point sources than to extended continuous sources, as depicted by the vestiges of sources that remain in both error and residual images. But even if CoSch-CLEAN is limited when applied to extended emission, MS-CLEAN is still a very good alternative to reduce the image distortion and improve the sparsity as compared to CoSch-CLEAN.
– We very carefully integrated our code in the existing LOFAR tools, that is, made it compatible with an immediate and a standard use with LOFAR. In particular, our current implementation of CS exists within AWimager and LWimager, which is the standard LOFAR imager that can handle DDE (by performing both *A*- and *W-projection*) as well as real data.

5.2. Future developments

In the CS algorithm, the – iterative soft thresholding – we used is not sophisticated, and other more sophisticated algorithms exist, such as SARA (Carrillo et al. 2012, 2014). There is much scope for modifying and experimenting with different minimization algorithms and their parameters, and this will be continued in the future and reported in future publications. In particular, we will integrate a reweighted-L1 scheme, which was demonstrated to reduce the bias of the signal reconstruction (Candes et al. 2008).

Performance improvements will be made to the computation of $\mathbf{A}$, for example, by using multi-threading and GPGPUs (Hardy 2013). This will benefit both CoSch-CLEAN and CS. There is a possibility of splitting CS into a major and minor cycle, similarly to CoSch-CLEAN, which will improve its speed. For our current CS implementation, a rigorous convergence criterion should be defined. The cycle stops either when reaching N_{max} or if the improvement between two consecutive iterations falls below a predefined threshold ϵ (Eq. (2)). We set this number to 200, as derived from educated experimentation on the presented datasets. Future implementation will include the tracking of the convergence based on the residual rms, and the determination of a robust stopping criterion.

The previous results suggest that CS methods are now mature enough for wider applications involving real datasets from LOFAR or other radio interferometers. To our knowledge, considering other completed developments or those that are in preparation as well as those discussed in the literature, this is the first time that CS has been integrated into an existing imaging system that handles DDEs through *A* and *W-projection* and operates on real data produced by a current giant radio telescope. To generalize it to other radio interferometers, we are now continuing this work by integrating our code into the standard imager of the common astronomy software applications (CASA) NRAO software package (*casapy*).

We plan to study the reliability of CS when applied to different visibility data sampling (i.e. sparse measurements brought by snapshot observations vs. non-sparse measurements of long time-integrated observations), different DDEs occurring with radio interferometers such as LOFAR, and various types of (resolved) objects with a high dynamic range, and requiring different dictionaries (e.g. imaging of other complex and extended objects at low frequencies such as Virgo A, which is

composed of a mix of extended emission and point sources). Reconstruction with a curvelet dictionary (Starck et al. 2003) is also possible in our implementation and leads, for the moment, to similar results. This specific dictionary further improves the sparsity of the solution when the source contains filamentous structures. Its extensive use will be a subject of future CS investigations in radio interferometry.

Future applications of CS will address multi-frequency imaging and imaging of transient sources that change during the course of an observation, and hence are imprinted only partially on the visibilities. CS is expected to provide the capability to recover snapshot images from sparsely sampled uv-coverage data slices. Multi-frequency polarization imaging can also be an application of this extension at low frequencies, see Li et al. (2011a).

Another aspect would be to investigate the use of CS for very long baseline interferometry (VLBI) imaging in LOFAR. By construction, an interferometer such as LOFAR presents a statistically lower number of very long baselines than the high density of shorter baselines. The highest angular resolution of the image depends on the quality of the measurement performed at these baselines. The signal from the longest baselines is sparse by nature, noisy, and especially sullied with DDEs (e.g. the radio antennas are situated under different ionospheres). It is therefore of interest to find methods such as CS that can reconstruct the true signal from these baselines.

Currently, the next generation of radiotelescopes are developed, namely SKA. The computational effort required to perform the observation, data acquisition, distribution, and processing calls for developing hierarchical and distributed real-time processing. These developments rely on the success of the current SKA pathfinders and precursors[4]. In that context, CS can be used for instrumental calibration. As an example, the Murchison Wide-field Array (MWA) has routinely deployed a fast-imaging mode to perform real-time image-plane calibration (Mitchell et al. 2008). Such a calibration mode, combined with CS methods that can approximate the content of the true sky from very sparse measurements, can represent a major advance in instrumental calibration. The NenuFAR project (Zarka et al. 2012), the independent phased-array interferometer formed from the extension of the French LOFAR station *FR606*, will have similar requirements. Moreover, this local interferometer offers a densely populated uv-space below $B \sim 400$ m with rare E-W and N-S longer baselines, which may benefit from CS reconstruction methods.

A larger community of astronomers (and especially imagers) are using and confirming the relevance of CS reconstruction methods. In the context of radio interferometry, we may use the sparse reconstruction method as a new standard way to image or calibrate data, complementary to x-CLEAN algorithms. The coming generation of radio instruments will be resource-demanding and will provide many potential applications for CS.

5.3. Software availability

An independent version of the CS software is available[5] for testing the CS code with different toy models. This package is called sparse aperture synthesis interferometry reconstruction (SASIR) and comes with its source code. It demonstrates the feasibility of the sparse reconstruction using either an original model

---

[4] https://www.skatelescope.org/technology/precursors-pathfinders-design-studies
[5] http://www.cosmostat.org/software.html

image and a Fourier mask in FITS format (representing the (u, v) sampling of the instrument), or a complex FITS image coming from gridding the data from any given measurement set. The operator **A**, central to the experiments of this paper, is computed directly inside the AWimager and cannot be extracted explicitly from the LOFAR framework, as it changes for every measurement set and refers to specific LOFAR dependencies for computing the DDE (antenna beam, polarization, projection effects, etc.). The source code of AWimager can be obtained by contacting Cyril Tasse and colleagues (Tasse et al. 2013).

References

Barbey, N., Sauvage, M., Starck, J.-L., Ottensamer, R., & Chanial, P. 2011, A&A, 527, A102
Beck, A., & Teboulle, M. 2009, SIAM J. Imag. Sci., 2, 183
Begeman, K., Belikov, A. N., Boxhoorn, D. R., et al. 2011, Future Gener. Comput. Syst., 27, 319
Bhatnagar, S., & Cornwell, T. J. 2004, A&A, 426, 747
Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, A&A, 487, 419
Bobin, J., Starck, J.-L., & Ottensamer, R. 2008, IEEE J. Select. Topics Signal Process., 2, 718
Candès, E., Romberg, J., & Tao, T. 2006, IEEE Trans. Information Theory, 52, 489
Candes, E. J., Wakin, M. B., & Boyd, S. P. 2008, J. Fourier Anal. Appl., 14, 877
Carrillo, R. E., McEwen, J. D., & Wiaux, Y. 2012, MNRAS, 426, 1223
Carrillo, R. E., McEwen, J. D., & Wiaux, Y. 2014, MNRAS, 439, 3591
Clark, B. G. 1980, A&A, 89, 377
Combettes, P. L., & Pesquet, J.-C. 2011, in Fixed-Point Algorithms for Inverse Problems in Science and Engineering (Springer-Verlag), 185
Cornwell, T. J. 2008, IEEE J. Select. Topics Signal Process., 2, 793
Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, IEEE J. Select. Topics Signal Process., 2, 647
Cornwell, T. J., & Perley, R. A. 1992, A&A, 261, 353
Davis, G. Mallat, S. & Avellaneda. M. 1997, J. Constructive Approximation, 13, 57
de Bruyn, A. G., & LOFAR EoR Key Science Project Team 2012, in AAS Meet. Abstracts 219, 214.05
Dewdney, P., Hall, P., Schilizzi, R., & Lazio, T. 2009, Proc. IEEE, 97, 1482
Donoho, D. 2006, IEEE Trans. Information Theory, 52, 1289
Dubner, G. M., Holdaway, M., Goss, W. M., & Mirabel, I. F. 1998, AJ, 116, 1842
Fannjiang, C. 2013, A&A, 559, A73
Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, A&AS, 117, 137
Hardy, S. J. 2013, A&A, 557, A134
Hargrave, P. J., & Ryle, M. 1974, MNRAS, 166, 305
Hoaglin, D., Mosteller, F., & Tukey, J. 1983, Understanding robust and exploratory data analysis, Wiley series in probability and mathematical statistics: Applied probability and statistics (Wiley)
Högbom, J. A. 1974, A&AS, 15, 417
Junklewitz, H., Bell, M. R., Selig, M., & Enßlin, T. A. 2013, A&A, submitted [arXiv:1311.5282]
Lannes, A., Anterrieu, E., & Marechal, P. 1997, A&AS, 123, 183

Lazio, T. J. W., Cohen, A. S., Kassim, N. E., et al. 2006, ApJ, 642, L33

Leonard, A., Dupé, F.-X., & Starck, J.-L. 2012, A&A, 539, A85

Li, F., Brown, S., Cornwell, T. J., & de Hoog, F. 2011a, A&A, 531, A126

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

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

McKean, J., Ker, L., van Weeren, R. J., et al. 2010, PoS(10th EVN Symposium)033

Mitchell, D. A., Greenhill, L. J., Wayth, R. B., et al. 2008, IEEE J. Select. Topics Signal Process., 2, 707

Pandey, V. N., van Zwieten, J. E., de Bruyn, A. G., & Nijboer, R. 2009, in The Low-Frequency Radio Universe, eds. D. J. Saikia, D. A. Green, Y. Gupta, & T. Venturi, ASP Conf. Ser., 407, 384

Ramos, E. P. R. G., Vio, R., & Andreani, P. 2011, A&A, 528, A75

Rau, U., & Cornwell, T. J. 2011, A&A, 532, A71

Sault, R. J., Hamaker, J. P., & Bregman, J. D. 1996, A&AS, 117, 149

Schwab, F. R. 1984, AJ, 89, 1076

Shulevski, A. 2010, Ph.D. Thesis, Leiden University, The Netherlands

Smirnov, O. M. 2011, A&A, 527, A106

Stappers, B. W., Hessels, J. W. T., Alexov, A., et al. 2011, A&A, 530, A80

Starck, J., & Bijaoui, A. 1991, in High Resolution Imaging by Interferometry (European Southern Observatory), 853

Starck, J.-L., & Murtagh, F. 2006, Astronomical Image and Data Analysis, 2nd edn. (Springer)

Starck, J.-L., Bijaoui, A., Lopez, B., & Perrier, C. 1994, A&A, 283, 349

Starck, J. L., Donoho, D. L., & Candès, E. J. 2003, A&A, 398, 785

Starck, J.-L., Murtagh, F., & Fadili, M. 2010, Sparse Image and Signal Processing (Cambridge University Press)

Steenbrugge, K. C., Heywood, I., & Blundell, K. M. 2010, MNRAS, 401, 67

Sutter, P. M., Wandelt, B. D., McEwen, J. D., et al. 2014, MNRAS, 438, 768

Tasse, C., van Diepen, G., van der Tol, S., et al. 2012, Comptes Rendus Physique, 13, 28

Tasse, C., van der Tol, S., van Zwieten, J., van Diepen, G., & Bhatnagar, S. 2013, A&A, 553, A105

van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2

Vũ, B. C. 2013, Adv. Comput. Math., 38, 667

Wakker, B., & Schwarz, U. 1988, ARA&A, 200, 312

Wenger, S., Magnor, M., Pihlström, Y., Bhatnagar, S., & Rau, U. 2010, PASP, 122, 1367

Wenger, S., Rau, U., & Magnor, M. 2013, Astron. Comput., 1, 40

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

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

Wilson, T. L., Rohlfs, K., & Hüttemeister, S. 2009, Tools of Radio Astronomy (Springer-Verlag)

Wolz, L., McEwen, J. D., Abdalla, F. B., Carrillo, R. E., & Wiaux, Y. 2013, MNRAS, 436, 1993

Wucknitz, O. 2010, PoS(ISKAF2010)058

Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, A&A, 550, A136

Zarka, P., Girard, J. N., Tagger, M., & Denis, L. 2012, in SF2A-2012: Proc. of the Annual meet. of the French Society of Astronomy and Astrophysics, eds. S. Boissier, P. de Laverny, N. Nardetto, et al., 687

[1] Laboratoire AIM, Université Paris Diderot Paris 7/CNRS/CEA-Saclay, DSM/IRFU/SAp, 91191 Gif-sur-Yvette, France

[2] Department of Physics and Electronics, Rhodes University, PO Box 94, 6140 Grahamstown, South Africa

[3] SKA South Africa, 3rd Floor, The Park, Park Road, 7405 Pinelands, South Africa

[4] Sagem (Safran), 27 rue Leblanc, 75512 Paris Cedex 15, France

[5] Netherlands Institute for Radio Astronomy (ASTRON), Postbus 2, 7990 AA Dwingeloo, The Netherlands

[6] Helmholtz-Zentrum Potsdam, Deutsches GeoForschungsZentrum GFZ, Department 1: Geodesy and Remote Sensing, Telegrafenberg, A17, 14473 Potsdam, Germany

[7] Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

[8] SRON Netherlands Institute for Space Research, PO Box 800, 9700 AV Groningen, The Netherlands

[9] Kapteyn Astronomical Institute, PO Box 800, 9700 AV Groningen, The Netherlands

[10] Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

[11] University of Twente, Drienerlolaan 5, 7522 Enschede, The Netherlands

[12] Institute for Astronomy, University of Edinburgh, Royal Observatory of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK

[13] School of Physics and Astronomy, University of Southampton, Southampton, SO17 1BJ, UK

[14] University of Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany

[15] Research School of Astronomy and Astrophysics, Australian National University, Mt Stromlo Obs., via Cotter Road, Weston, A.C.T. 2611, Australia

[16] Max Planck Institute for Astrophysics, Karl Schwarzschild Str. 1, 85741 Garching, Germany

[17] SmarterVision BV, Oostersingel 5, 9401 JX Assen, The Netherlands

[18] Thüringer Landessternwarte, Sternwarte 5, 07778 Tautenburg, Germany

[19] Hamburger Sternwarte, Gojenbergsweg 112, 21029 Hamburg, Germany

[20] Department of Astrophysics/IMAPP, Radboud University Nijmegen, PO Box 9010, 6500 GL Nijmegen, The Netherlands

[21] Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

[22] Laboratoire Lagrange, UMR7293, Université de Nice Sophia-Antipolis, CNRS, Observatoire de la Côte d'Azur, 06300 Nice, France

[23] Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, The Netherlands

[24] LPC2E – Universite d'Orléans/CNRS, 45071 Orléans Cedex 2, France

[25] Station de Radioastronomie de Nançay, Observatoire de Paris – CNRS/INSU, USR 704 – Univ. Orléans, OSUC, route de Souesmes, 18330 Nançay, France

[26] Jodrell Bank Center for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

[27] Astronomical Institute "Anton Pannekoek", University of Amsterdam, Postbus 94249, 1090 GE Amsterdam, The Netherlands

[28] Astronomisches Institut der Ruhr-Universität Bochum, Universitaetsstrasse 150, 44780 Bochum, Germany

[29] Astro Space Center of the Lebedev Physical Institute, Profsoyuznaya str. 84/32, 117997 Moscow, Russia

[30] Sodankylä Geophysical Observatory, University of Oulu, Tähteläntie 62, 99600 Sodankylä, Finland

[31] STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot OX11 0QX, UK

[32] Center for Information Technology (CIT), University of Groningen, PO Box 11044, 9700 CA Groningen, The Netherlands

[33] Centre de Recherche Astrophysique de Lyon, Observatoire de Lyon, 9 av Charles André, 69561 Saint Genis Laval Cedex, France

[34] Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany

[35] LESIA, UMR CNRS 8109, Observatoire de Paris, 92195 Meudon, France

[36] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA