A multiple-beam CLEAN for imaging intra-day variable radio sources

I. M. Stewart^{1,2}, D. M. Fenech³, and T. W. B. Muxlow¹

¹ Jodrell Bank Centre for Astrophysics, University of Manchester, Oxford Road, Manchester M13 9PL, UK
e-mail: Ian.Stewart-2@manchester.ac.uk

² Astrophysics, Cosmology and Gravity Centre, Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

³ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

Received 27 October 2010 / Accepted 12 July 2011

ABSTRACT

The CLEAN algorithm, widely used in radio interferometry for the deconvolution of radio images, performs well only if the raw radio image (dirty image) is, to good approximation, a simple convolution between the instrumental point-spread function (dirty beam) and the true distribution of emission across the sky. An important case in which this approximation breaks down is during frequency synthesis if the observing bandwidth is wide enough for variations in the spectrum of the sky to become significant. The convolution assumption also breaks down, in any situation but snapshot observations, if sources in the field vary significantly in flux density over the duration of the observation. Such time-variation can even be instrumental in nature, for example due to jitter or rotation of the primary beam pattern on the sky during an observation. An algorithm already exists for dealing with the spectral variation encountered in wide-band frequency synthesis interferometry. This algorithm is an extension of CLEAN in which, at each iteration, a set of N “dirty beams” are fitted and subtracted in parallel, instead of just a single dirty beam as in standard CLEAN. In the wide-band algorithm the beams are obtained by expanding a nominal source spectrum in a Taylor series, each term of the series generating one of the beams. In the present paper this algorithm is extended to images which contain sources which vary over both frequency and time. Different expansion schemes (or bases) on the time and frequency axes are compared, and issues such as Gibbs ringing and non-orthogonality are discussed. It is shown that practical considerations make it often desirable to orthogonalize the set of beams before commencing the cleaning. This is easily accomplished via a Gram-Schmidt technique.

Key words. methods: data analysis – techniques: image processing – radio continuum: general – techniques: interferometric

1. Introduction

It has been known for some decades that some radio sources vary in intensity over time (Dent 1965). The range of time scales detected so far, roughly from minutes to decades, is set by the practicalities of measurement rather than by any intrinsic properties of the source populations. The present paper is concerned with sources which exhibit significant variation in flux density on time scales of less than 24 h (so-called Intra-Day Variables or IDVs) and which are also associated with extended structure of a size resolvable by present-day radio interferometers. The structure makes it of interest to image such sources, but the rapid variability can hinder attempts to do so. It is these difficulties in the interferometric imaging of IDV sources which the present paper is designed to address.

How many such sources are there, and of what types? Any estimation of the incidence of variability among the radio source population is complicated by the extra degrees of freedom implicit in a non-flat light curve. It is also inevitable that there will be selection effects due to the inability of the instrument to detect modulation depths below a certain cutoff, or the insensitivity of the observing regime to time scales outside a certain range. The difference between the population of objects observable in our galaxy and those which are extra-galactic introduces a dependence of incidence on galactic latitude (further complicated by the greater propensity of compact, extragalactic sources to

scintillate at low galactic latitudes). It is even more difficult to estimate the fraction of variable sources which are IDV, since surveys of the depth and cadence necessary to determine this quantity demand extravagant amounts of observing time.

Of presently-available survey results, De Vries et al. (2004), in a two-epoch comparison of a high latitude field, found only $\sim 0.2 \text{ deg}^{-2}$ at 1.4 GHz flux densities of $\geq 2 \text{ mJy}$. Bower et al. (2007) have processed a large amount of VLA 5 and 8.4 GHz archival data and find a “snapshot” incidence of 1.5 deg^{-2} radio transients at flux densities greater than $350 \mu\text{Jy}$. Becker et al. (2010) compared 3 epochs of VLA observations of the galactic plane at 6 cm and found $\sim 1.6 \text{ deg}^{-2}$ sources between 1 and 100 mJy which had a modulation depth greater than 50%. Regarding the IDV fraction, Kedziora-Chudczer et al. (2001) found, in their multi-frequency study (between 1.4 and 8.6 GHz), that about half of their 13 BL Lac objects, and a smaller fraction of quasars, exhibited IDV up to about 10% modulation; more recently Lovell et al. (2008), in a more comprehensive survey at 5 GHz, report that 37% of their 443 flat-spectrum sources showed significant variability on a 2-day time scale.

Sources which exhibit a combination of IDV plus some resolved structure fall into several classes, which are briefly reviewed in the following paragraphs.

Novae are expected to exhibit significant changes in radio flux density at intraday time scales. In the standard model, the flux density from an expanding isothermal shell is expected

initially to increase proportional to the square of time since the outburst (Hjellming 1996), although some recent VLBA observations appear to be inconsistent with that picture (Krauss et al. 2011). However, most novae at least in the past decade have not been resolved by radio interferometry until of order 100 days after the outburst, by which time the fractional change in flux density has fallen to about a percent or two per day (see e.g. Eyres et al. 2000; Eyres et al. 2005; Heywood & O’Brien 2007). A recent exception is the 2006 outburst of the recurrent nova RS Ophiuchi (O’Brien et al. 2006a). This was observed with the VLBA first on day 14 after the outburst, and the European VLBI Network (EVN) from day 20 (O’Brien et al. 2006b). The size of the source at 6 cm wavelength grew from approximately 20 to 35 mas between these observations. The flux density was no longer varying rapidly by that time, but the earliest observations with MERLIN on day 4 after the outburst show it then varying by a factor of 2 over less than a day. Model calculations indicate that RS Oph would have been resolvable by the EVN at that time.

X-ray binaries form an interesting class of galactic IDV sources. Several of these are associated with resolved structure and have emitted jansky-strength radio flares which show pronounced evolution in intensity on intraday time scales. Examples include Circinus X-1 (Haynes et al. 1978; Tudose et al. 2008), Cygnus X-3 (Tudose et al. 2009) and GRS 1915+105 (Fender et al. 1999; Rushton et al. 2010).

Extragalactic IDV sources fall into fewer categories and are in fact dominated by flat-spectrum AGNs, or objects nowadays collectively described as blazars (Wagner & Witzel 1995; Lovell et al. 2008). One could almost say that being a blazar is both a necessary and a sufficient condition for an extragalactic source to exhibit IDV, since these objects have the combination of brightness and small size which allows them on the one hand to be detectable at great distances and on the other to be compact enough to exhibit IDV. In many of these objects the compact core appears to be embedded in radio structure which is resolvable by current interferometers (see for example Kovalev et al. 2005; Ojha et al. 2006).

IDV among blazars is an area of current interest. IDV mechanisms for these sources fall into two categories: variation which is intrinsic to the object, and variation caused by the propagation medium (scintillation in the local interstellar medium, or ISS). Both appear to play some role but the relative proportions are not clear. Spectral dependence of the variability at cm wavelengths, annual modulation of the time scales and time shifts between measurements made at different locations all offer strong evidence for the dominance of ISS for some sources (Bignall et al. 2009, and references within). Correlation between variation at radio and other wavelengths (Quirrenbach et al. 1991) or a bad fit to ISS spectral models (Fuhrmann et al. 2008) tend to support the intrinsic origin of IDV in others. Intrinsic origin presents theoretical challenges because it is difficult to explain it without invoking either an incidence of extreme Doppler factors which is hard to accept on statistical grounds, or brightness temperatures in excess of the inverse-Compton limit of about 10^{12} K (Jauncey et al. 2001). Some questions remain also in the ISS picture, such as a relatively weak correlation with galactic latitude (Kedziora-Chudczer et al. 2001; Lovell et al. 2008); the connection between scintillation and scatter-broadening is also still unclear (Ojha et al. 2006; Lazio et al. 2008).

The only non-blazar category of extragalactic IDV source which is relevant to the present paper is the maser. McCallum et al. (2009) for example show that the H₂O megamasers in Circinus are both resolvable and exhibit IDV. The techniques

described in the present paper can make it easier to image IDV masers which have superimposed lines, meaning they must both be present in the same channel map.

It can thus be seen that there are many objects which it is desirable to image at radio wavelengths with as high resolution as possible, but whose variation on intraday timescales interferes with imaging via earth-rotation aperture synthesis. Broadly speaking, the problem occurs because the IDV causes the pattern of lobes around the time-variable component of the source to be poorly matched to the dirty beam. Such structure cannot be completely removed by the standard CLEAN algorithm, and will therefore tend to limit the dynamic range of the image. As further described in Sect. 2.2, the situation is analogous to the problem in multi-frequency synthesis (MFS), caused by differences in spectral index among objects in the field. The argument for seeking an improved way to image IDV sources is the same as for this MFS case. In fact, as is shown in the present paper, the same technique, originally developed for MFS by Sault and Wieringa (1994), can be applied to both situations.

The plan of the paper is as follows. Section 2 contains a brief outline of aperture synthesis theory. In Sects. 3.1 and 3.2, the generalized N -beam Sault-Wieringa theory is described in detail. The desirability of orthogonalizing the beams is discussed in Sect. 3.3. Different time and frequency basis functions are compared in Sect. 4.1, with emphasis on the avoidance of Gibbs-like phenomena. In Sect. 4.2.2, a dual expansion in both frequency and time axes is shown to be both necessary and effective in cleaning a simulated wide-band data set which includes sources which vary significantly both in frequency and time. Finally, in Sect. 4.2.3, a simple 2-beam expansion is derived which is shown to be a powerful means of increasing dynamic range in many IDV cases.

Brief descriptions of the methods described here have already been presented elsewhere in earlier stages of development (Stewart 2008, and 2010).

2. Interferometry

2.1. Aperture synthesis

A radio interferometer measures cross-correlations between the voltage signals detected by pairs of antennas. Each correlation yields a complex number which encodes information about the amplitude and phase of the signals from all sources of radio waves in the field of view of the antennas. Each complex correlation can be considered to be a point sample, plus added noise, of a continuous function $V(\mathbf{u})$, known as the visibility function. The vector $\mathbf{u}$, known as a baseline, is the separation vector between any given pair of antennas, expressed as a number of wavelengths. For an array of antennas which is physically coplanar, V becomes a 2-dimensional function, and can be shown to be the Fourier transform of a function $I(\mathbf{r})$ which is related to the sky brightness distribution at the detector wavelength. Here $\mathbf{r} = (l, m)$ are the sines of the zenith angles of a given sky location. The relation between I and the true sky brightness distribution is given by

$$I(\mathbf{r}) = A(\mathbf{r}) \frac{I_{\text{true}}(\mathbf{r})}{\sqrt{1 - \mathbf{r} \cdot \mathbf{r}}}, \quad (1)$$

where A , known as the primary beam, is the receptivity of the individual antenna elements as a function of $\mathbf{r}$.

If delays are introduced into the signal chains such that signals from a single point in the sky (known as the phase centre)

arrive at the correlator exactly in phase, then the Fourier relationship to the sky brightness distribution becomes approximately true even for non-coplanar arrays, provided the portion of sky to be imaged is restricted to a sufficiently small region around the phase centre. In this general, non-coplanar case, the baseline vectors $\mathbf{u}$ are considered to be the separation vectors between the pairs of antennas projected onto a plane normal to the phase centre. The direction cosines l and m are also in this case taken in a basis in which $(l, m) = (0, 0)$ lies in the direction of the phase centre, rather than the zenith.

If the sampling function (which is, to first approximation, a sum of delta functions) is denoted by S then the output of the interferometer is VS . The Fourier inversion $\mathcal{F}^{-1}(VS)$ gives $D = I * B$ where D , known as the dirty image, is a convolution of the true sky image I with the “dirty beam” $B = \mathcal{F}^{-1}(S)$. B plays the same role for an interferometer as the point-spread function of a traditional telescope.

In general, the larger the number of samples of V , the more compact the distribution of flux density in B , and therefore the closer the correspondence between D and the true sky I ; the sensitivity of the observation will also be increased. An interferometer containing N antennas will generate $N(N - 1)/2$ independent samples of V at each observation. The number of samples can be further expanded via the technique of Earth-rotation synthesis. In this, the rotation of the Earth over the course of a day is used to generate a sequence of different baselines $\mathbf{u}$ for each antenna pair. Another way to increase the number of samples is to observe at several frequencies, a technique known as frequency synthesis. Because baselines are expressed in wavelengths, the fixed spatial separation between a pair of antennas generates baseline vectors $\mathbf{u}$ of different length at different frequencies. Implicit in these two synthesis techniques is the assumption that the sky brightness will be constant over time in the first case and over frequency in the second.

A more detailed description of the fundamentals of interferometry can be found in many sources, for example Thompson et al. (2001) or Taylor et al. (1999).

2.2. CLEAN

The CLEAN algorithm was invented by Högbom (1974) and has been further elaborated by Clark (1980) and Cotton and Schwab (Schwab 1984) among others. The essence of the algorithm is to perform many iterations of a process in which a small amount of the dirty beam is subtracted, centred at the highest remaining point in the dirty image, ideally until nothing remains in this image but noise. The positions and amounts subtracted are recorded as “clean components” and used afterwards to reconstruct an approximation to the true sky image I . The gain factor by which the dirty beam is multiplied, and the number of iterations to perform, are parameters which are chosen ahead of time by the user.

CLEAN was originally presented as an empirical algorithm which appeared to produce results, although it required some experience to judge how best to apply it. It is known not to perform well when applied to extended objects (Cornwell 1983), although a modified algorithm has been shown to yield improvements here (Steer et al. 1984). CLEAN is still in wide use as a practical method of removing sampling artifacts from interferometry images.

The reasons for CLEAN’s success were unclear for some time (Cornwell et al. 1999), and its theoretical basis has come under occasional criticism (Tan 1986; Lannes et al. 1997). Only relatively recently has it been shown to be related to compressive

sampling (Candès & Wakin 2008) and been given a sound theoretical underpinning.

In order for Högbom CLEAN to work, the convolution relation $D = I * B$ must be valid. There are cases where this assumption breaks down however. One of the most severe departures occurs in frequency synthesis with large fractional bandwidths. It is common for a single observational field to contain objects whose spectral indices differ by several tens of percent or more. Differences in spectra of this order are unimportant if the observation fractional bandwidth is small but may significantly degrade the convolution assumption where the fractional bandwidth approaches 1.

Conway et al. (1990) showed that, in the wide-band case in which the dirty image D no longer well approximates a convolution of the sky brightness distribution I , it was nevertheless possible to express D as a sum over a relatively small number N of component images D_i , each of which individually obeys a convolution relation $D_i = I_i * B_i$. Conway et al. arrived at this by expanding the nominal spectrum at each sky location in a Taylor series, each i th term of the series generating a respective “spectral dirty beam” B_i . In this case each image I_i is simply a sky map of the value of the i th Taylor coefficient. Conway et al. suggested a coordinate transform for better application of this technique to commonly-found power-law radio spectra, and presented an approximate method to solve for the coefficient images I_i when N is restricted to 2. This method consists of a 2-step sequential CLEAN and relies on the beams B_0 and B_1 being approximately orthogonal.

Sault & Wieringa (1994) retained the Taylor expansion but devised a new solution algorithm which can be thought of as a generalization of the Högbom CLEAN algorithm from its original “scalar” context, in which a single image D is iteratively deconvolved, to a new “vector” context in which N images D_i are deconvolved in parallel. Orthogonality of the beams is no longer required (although the technique fails if 2 or more beams are identical). Sault and Wieringa elaborated their theory as it applied to the $N = 2$ case but, as the authors themselves suggest, the extension to $N > 2$ is not difficult.

The CLEAN algorithm continues to be a subject of active development. Recent work includes an extension of CLEAN to produce clean components with a range of sizes (Cornwell 2008), a modification to clean tomographic LISA images (Hayama et al. 2006), and an adaption of the algorithm to reconstruct RHESSI images of solar flares (Schwartz 2009). The compressive-sampling formalism has recently also generated some promising new approaches (Wiaux et al. 2009; Li et al. 2011).

3. A generic multiple-beam CLEAN

3.1. Generalized Conway decomposition

In its most general form, the aim of the treatment described by Conway et al. (1990) is to find a way to decompose the dirty image into a sum of convolutions

$$D \sim \sum_{i=0}^{N-1} D_i = \sum_{i=0}^{N-1} I_i * B_i. \quad (2)$$

The problem is to calculate a set of beams B_i for which this relation is true. There are no doubt many ways to do this, but in the present section we are going to consider only the route which comes from expanding the frequency and time dependence of

the sky brightness in a sum of basis functions. Such an expansion is written

$$I(\mathbf{r}, \nu, t) \sim \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} I_{p,q}(\mathbf{r}) F_p(\nu) T_q(t), \quad (3)$$

where $I(\mathbf{r}, \nu, t)$ is the brightness function; $\mathbf{r} = (l, m)$ are direction cosines of the angular displacement from the phase centre; ν is frequency and t time; $I_{p,q}(\mathbf{r})$ is a sky map of the (p, q) th component; and F_p and T_q are p th and q th members respectively of sets of basis functions. Earth-rotation and frequency synthesis will return a set of $j = 1$ to M measurements V_j of the cross-correlations between antennas. Ignoring for the sake of simplicity non-coplanarity, weights, the shape of the primary beam, and the spherical projection correction $(1 - \mathbf{r} \cdot \mathbf{r})^{-0.5}$ (see Eq. (1)), each V_j is related to I by a transform:

$$V_j = \int d\mathbf{r} \exp(-2\pi i \mathbf{u} \cdot \mathbf{r}) I(\mathbf{r}, \nu_j, t_j).$$

Although each V_j represents an average over a small time $\times$ bandwidth window, if this window is small enough, V_j may be represented as a point sample of the visibility function at the spatial frequency $\mathbf{u}_j$. The total set $\mathbf{V}$ of visibility measurements can thus be written

$$\mathbf{V} = \sum_{j=1}^M \delta(\mathbf{u} - \mathbf{u}_j, \nu - \nu_j, t - t_j) \int d\mathbf{r} \exp(-2\pi i \mathbf{u} \cdot \mathbf{r}) I(\mathbf{r}, \nu, t).$$

Expanding I in its basis functions gives, with some rearrangement:

$$\mathbf{V} = \sum_{p,q} \sum_{j=1}^M \delta(\mathbf{u} - \mathbf{u}_j) F_p(\nu_j) T_q(t_j) \int d\mathbf{r} \exp(-2\pi i \mathbf{u} \cdot \mathbf{r}) I_{p,q}(\mathbf{r}).$$

Fourier inversion of $\mathbf{V}$ yields the dirty image D . This is indeed found to be a sum of convolutions:

$$D = \mathfrak{F}^{-1}(\mathbf{V}) = \sum_{p,q} I_{p,q} * B_{pq}, \quad (4)$$

where

$$B_{pq} = \mathfrak{F}^{-1} \left(\sum_{j=1}^M \delta(\mathbf{u} - \mathbf{u}_j) F_p(\nu_j) T_q(t_j) \right).$$

Here the $\mathfrak{F}$ symbol represents the Fourier transform. The two sums over p and q are easily collapsed to one, in which case Eq. (4) maps directly to Eq. (2). Clearly this requires that $N = P \times Q$.

3.2. The N -beam Sault-Wieringa algorithm

Sault & Wieringa (1994) developed an algorithm which, although it was set down only for $N = 2$, is easily generalized to the following:

- generate initial R_i from $D \star B_i$ (the $\star$ here represents correlation);
- generate, for $i \in [0 : N - 1]$ and $j \in [0 : N - 1]$, cross-correlation images $Z_{i,j} = B_i \star B_j$;
- calculate a matrix $\mathbf{M}$ such that its (i, j) th element is given by $Z_{i,j}(\mathbf{r} = 0)$;

- perform a number of iterations of the following procedure:
  1. find the location $\mathbf{r}_{\max}$ of the maximum value of the image formed from $R = \sum_{i,j} R_i R_j M_{i,j}^{-1}$;
  2. perform, for each i , the subtractions

$$R_i(\mathbf{r}) = R_i(\mathbf{r}) - \lambda \sum_{j=0}^{N-1} R_j(\mathbf{r}_{\max}) Z_{i,j}(\mathbf{r} - \mathbf{r}_{\max});$$

3. form a vector $\mathbf{c}_{\text{raw}}$ of the entangled clean components $\lambda R_i(\mathbf{r}_{\max})$;
4. solve the equation $\mathbf{c} = \mathbf{M} \mathbf{c}_{\text{raw}}$ to disentangle the clean components;
5. store these for later cleaned-image reconstruction.

3.3. Orthogonality pros and cons

The Sault-Wieringa deconvolution includes a matrix inversion. If two or more beams from the set of N are identical, the matrix $\mathbf{M}$ is singular, i.e. uninvertible. If no two beams are exactly identical, but two or more are similar (have a normalized scalar product close to unity), $\mathbf{M}$ will be formally invertible but can be expected to be ill-conditioned (i.e., result in large errors in the final clean components). Hence approximate orthogonality between beams is desirable; but if this be established, little further reduction in error is to be gained by enforcing complete orthogonality.

There may however be other reasons to manipulate the beams, even to the point of transforming them into an orthonormal set, other than keeping the matrix $\mathbf{M}$ well-conditioned. Two such circumstances are described here. Firstly, the main aim of the Conway decomposition and subsequent vector cleaning is arguably to obtain a better image of the average sky brightness distribution. This image is identical to the image of the zero-order beam coefficient if and only if every other beam is zero-valued at its centre – because a non-zero value of $B_i(0)$ equates to a non-zero average of the i th basis function over the frequency and time range of the observation. Thus in order to simplify the post-production of an average-brightness image, one would want to modify the set of beams before cleaning by subtracting $B_i(0) \times B_0/B_0(0)$ from each B_i for $i > 0$.

The second circumstance concerns the practicalities of computing the deconvolution. The Sault-Wieringa algorithm requires that $N(N + 1)$ images be kept in memory during the iteration. For $N = 2$, as treated by Sault and Wieringa, this is a manageable number of images. However, some of the scenarios discussed in the present paper make use of values of N as high as 30. Maintenance within memory of close to 1000 arrays, each of which may be as large as 2048×2048 pixels, is likely to prove a strain for most machines at least in the lower echelons of present-day computing power. The load may be much reduced by orthogonalizing the beams before the start of the algorithm, i.e. by generating a set of modified beams B'_i such that $B'_i \cdot B'_j = 0$ for all $i \neq j$. The cross-correlation images $Z'_{i,j}$ generated from these will be such that their central values, thus the elements of $\mathbf{M}'$, will be 0 for all $i \neq j$. There is no need to store or even generate the “off-diagonal” images; the necessary inventory of images thus reduces to $2N$.

The transformation from the basis functions in Eq. (3) to beams in “sky-image” space does not preserve orthogonality. Because of this there seems little to be gained by requiring the basis functions to be orthogonal: one must work directly with the dirty beams themselves.

The Gram-Schmidt process of converting B_i to orthonormal B'_i can be expressed as the matrix equation

$$\mathbf{b} = \mathbf{G}\mathbf{b}', \quad (5)$$

where $\mathbf{b}$ is a vector of the N input dirty beams B_i , $\mathbf{b}'$ is a vector of the orthonormalized beams B'_i , and $\mathbf{G}$ is a lower-triangular matrix with elements defined as follows: $G_{i,i} = B'_i \cdot B'_i$ before B'_i is normalized, and $G_{i,j}$ for $j < i = B'_j \cdot B'_i$.

As Golub & Van Loan (1996) point out, the usual Gram-Schmidt procedure is known to be numerically unstable. These authors suggest an alternate, stabilized algorithm which amounts to solving Eq. (5) by backward substitution, the elements of $\mathbf{G}$ being calculated on the fly.

The Sault-Wieringa algorithm, working with orthonormalized beams, produces, for each image pixel (x, y) , a vector of clean components $\mathbf{c}'$ such that

$$D = \sum c'_i B'_i,$$

or, in vector notation,

$$D = \mathbf{c}'^T \mathbf{b}'. \quad (6)$$

How to reconstruct the clean components which give D in terms of the original beams? Inverting Eq. (5) and inserting (6) gives

$$D = \mathbf{c}'^T \mathbf{G}^{-1} \mathbf{b}.$$

Obviously the desired clean components $\mathbf{c}$ are given by

$$\mathbf{c}^T = \mathbf{c}'^T \mathbf{G}^{-1}.$$

3.4. Post-processing

The output of the Sault-Wieringa process is, as with Högbom CLEAN, a list of clean components; but here each component is no longer scalar but is itself an N -element vector or list. The question then arises of what to do with this information. Almost certainly one will want to make an image of the average sky brightness, since one of the motivations of the Conway decomposition has been to obtain a more accurate reconstruction of the sky, free from the spectral artifacts which one sees in Högbom-based deconvolution of wide-band observations. One might think that similar images should also be made from each of the higher-order elements of the clean components. Experience shows however that such images rapidly become too noisy to be of use as the order increases. In broad terms this is because the higher-order basis functions tend to vary more rapidly with displacement in the UV plane; which means they contain more power in higher spatial frequencies in this domain; which translates to a broader, less centrally peaked structure for the respective beam in the sky domain. This broadening and diluting of the beam profile for higher orders (visible in the example beam profiles of Conway et al. 1990) causes relatively more “false detections” at these orders to occur during the deconvolution, until at the end the real clean component values for this order may become lost among the false ones. The problem is exacerbated by the tendency of the true component values to decrease (with any sensible choice of basis function) as the order increases.

A possible remedy for this problem might be to choose smaller values of the loop gain λ as the order of the basis function increases. But we have not so far tested this conjecture.

For constructing a restored image of the average brightness, it seems clear that the best thing to do is construct an average-brightness beam, and fit the restoring beam to its centre, in analogy to the present standard practice when employing Högbom CLEAN. However, there is no such obvious source of a restoring beam profile for higher-order images, since their respective beams don't necessarily have a peak in the centre to fit to. Use of the average-derived restoring beam in these cases is probably acceptable though, since the underlying spatial resolution should be similar. Another possibility would be to use a Gaussian of the same full width half maximum as the beam.

4. Time-variable sources

4.1. Expansion in basis functions

It is a commonplace that radio continuum spectra tend to be smooth. In nearly all cases a continuum spectrum may be well approximated over an octave or so (a frequency range which the forthcoming wide-band interferometers are not expected to greatly exceed) by a polynomial of at most 2 or 3 orders. In this case there seems little to be gained by investigating the use of other types of basis function. Light curves, on the other hand, are much less well behaved: fluctuations can be expected over a wide range of timescales. Sault-Wieringa style parallel cleaning of observations which include time-variable sources may therefore require expansion of the nominal light curve to higher order than is usually necessary for spectra. This also makes it of more interest to compare different basis expansions in this case.

Note however that an observation which includes only a single variable source can be cleaned via the simple 2-beam technique described in Sect. 4.2.3.

All other aspects being equal, one would choose the basis set which resulted in the smallest residuals for a given N in Eq. (2), which is another way of saying that the sum over N “component” dirty images D_i should converge rapidly to the “true” dirty image D as N increases. And rapid convergence in the sky plane is surely linked to rapid convergence of the basis function expansion in the U-V plane. However, it is difficult to predict which set will best meet this criterion, for two reasons. Firstly the light curves encountered in practice may be very diverse: a basis set which converges rapidly to one sort may not do so for another. Secondly, there appears to be at present no analysis which predicts the image residuals. Thus trial and error, as with traditional CLEAN, probably remains the best guide here too.

In the present section we compare two sets of basis functions for the expansion of light curves, namely Fourier sinusoids and Chebyshev polynomials.

If a standard Fourier series is employed, the expansion will suffer from Gibbs phenomenon at the boundaries, unless the light curve is periodic continuous (in fact periodic analytic) over the chosen observation interval. Gibbs ringing can however be much reduced by the following treatment. Let the duration of the observation be denoted by T , and the light curve at a given sky direction by $f(t)$. Suppose one doubled the observation time, and filled the new interval with a function g defined such that

$$g(t) = \begin{cases} f(t), & 0 \leq t < T \\ f(T-t), & T \leq t < 2T. \end{cases}$$

The new function g is periodic continuous in the interval $[0, 2T]$, so there should be no zeroth-order Gibbs ringing in its Fourier expansion. g is alsosymmetric about zero, so only cosine terms

Fig. 1. Direct Fourier expansion of a simulated light curve compared to the result of a Sault-Wieringa clean using a set of basis functions which suppresses Gibbs oscillation at the boundaries.

will remain in the expansion. If the resulting Fourier basis functions are truncated back to $[0, T]$ they are seen to be given by

$$T_n(t) = \cos(\pi n t / T) \text{ for } n \in [0, \dots, N]. \quad (7)$$

Similar tricks can be used to enforce boundary differentiability to higher orders if desired.

If there are gaps in the time sequence of data values – patches of bad data perhaps, or even planned, periodic observations of a phase calibrator – such gaps might be expected to cause additional Gibbs-type problems. A normal Fourier expansion of a function which suddenly drops to zero at intervals would suffer Gibbs ringing at the discontinuities. However, Conway decomposition is not “normal expansion” in the basis functions. The actual basis onto which the dirty image is projected consists of the set of beams; and the construction of each beam from the corresponding basis function in the UV plane effectively disconnects it from any irregularities in the support of the basis functions. Orthonormal basis functions won’t necessarily produce orthonormal beams, and vice versa.

A final point to note is that this technique cannot disentangle or unwrap an observation which spans more than 24 h at the same frequency: in this case all that can be obtained is the net (cyclic superimposed) light curve.

4.2. Simulations

4.2.1. A single time-variable source

Figures 1 to 4 refer to a simulation constructed as follows. Artificial visibilities were calculated for a single time-varying source which was located at the phase centre. The array parameters were those of MERLIN; the chosen source declination was $+40^\circ$; integration time was 10 s, and 32 channels of 1 MHz width starting at 6 GHz were used. Perfect calibration was assumed and no instrumental noise was added. A gap was introduced into the data sequence between approximately 4 and 6.5 h postmeridian.

Figures 1–3 show the effect of different basis expansions on the Gibbs phenomenon; Fig. 4 concerns the relative efficiency of Högbom versus Sault-Wieringa cleans for time-variable sources.

Fig. 2. The same light curve as in Fig. 1 but now with residuals plotted. Shown are the residuals from the Fourier expansion of Fig. 1 (dashed curve) and the residuals from a Sault-Wieringa clean using the same set of Fourier basis functions (half-tone solid curve). For better clarity, values of small magnitude have not been plotted; and a different cutoff was used for each curve. Points falling within the data gap have also been omitted.

Fig. 3. This is similar to Fig. 2. The dashed curve again represents the Fourier residuals, whereas the dash-dot and dotted curves are residuals from a Sault-Wieringa clean using two different types of basis function. For the dash-dot curve, Chebyshev polynomials up to order 5 were employed to generate the Sault-Wieringa beams; the dotted curve is the result of using the half-frequency cosine functions (also to the order 5).

In Fig. 1, the light curve of the source (solid line) is compared firstly to a direct Fourier expansion of this light curve to order 10 (dashed line), and secondly to the light curve reconstructed after Conway decomposition into 6 beams (i.e. to order 5) from the half-frequency cosine basis functions prescribed by Eq. (7), followed by 1000 cycles of Sault-Wieringa cleaning with loop gain 0.1 (dotted line). The Gibbs ringing of the standard Fourier expansion at boundaries and other discontinuities is very obvious, as is the almost complete absence of such an effect in the Sault-Wieringa result.

Fig. 4. Both sides of this plot show rms image values as a function of radius from the source position. Values calculated from clean-component images appear in the left panel; those on the right come from residual images. The radial distribution of rms values in the dirty image (solid line) is included in both panels for purposes of comparison. The dashed lines represent the Högbom-cleaned data, whereas the dotted lines are from the Sault-Wieringa cleaning.

Figure 1 shows clearly that Gibbs phenomenon can be largely avoided in Sault-Wieringa cleaning of time-varying sources. It is of interest however to explore a little more deeply into the difference between boundary discontinuities and data gaps, and the relative efficacy of different basis functions. Figures 2 and 3 serve this purpose. For these figures, the same input data are used, but the plotting style is different to Fig. 1. Here are plotted, on a logarithmic scale, absolute values of the residuals of the curves of interest against the input light curve.

In Fig. 2 the direct Fourier expansion (dashed line) is compared to the result of a Sault-Wieringa clean (half-tone solid line) as in Fig. 1; but here the unmodified Fourier functions, i.e. the same functions used in the direct expansion, were chosen as the basis set (which generated 21 beams!). As expected, in both cases there is Gibbs oscillation at the boundary, but only in the direct expansion is there also Gibbs at the edges of the data gap. The conclusion here is that it is unnecessary to choose the basis set with any care in order to avoid Gibbs ringing at data gaps: the Sault-Wieringa technique by its nature avoids such troubles. The same is however obviously not true of the problem at the boundaries, which arises from the inability of the Fourier basis set to represent a function with a discontinuity at the boundary. In effect, the Sault-Wieringa deconvolution can interpolate over a gap, but cannot (without “outside help”) deal with a discontinuity.

Figure 3 again shows residuals from the direct Fourier expansion (dashed line) but compares them here to residuals from two different Sault-Wieringa deconvolutions: using firstly the half-frequency cosine basis of Eq. (7) (dotted line, as already displayed in Fig. 1), and secondly, Chebyshev polynomials (dot-dashed line). Both expansions were truncated at order 5. It is perhaps slightly unexpected that the cosine functions seem to fit the light curve better than the Chebyshevs.

Figure 4 compares a Högbom clean (1000 cycles at gain 0.1) of the same time-varying single source with Sault-Wieringa clean using the half-frequency cosine basis to order 5. (The Sault-Wieringa result is the same which generated the dotted

Fig. 5. Surfaces showing the frequency and time dependence of the 4 variable sources. The height of each surface represents the flux density of that source. Major contours occur at 1 jansky intervals. The non-variable source had a flat spectrum at 1 Jy. Note that the order of the sources has been slightly rearranged for better display of the surfaces.

curves in Figs. 1 and 3.) The data plotted in Fig. 4 come from images: the dirty image for the data set, and the clean and residual images from the Högbom and Sault-Wieringa cleans respectively. What has been done for each image is to calculate the rms of the image values in a polar coordinate system centred on the source at the phase centre. The RMS values have been plotted as a function of radius from this centre.

Use of the Sault-Wieringa procedure is seen to improve the dynamic range by about 2 orders of magnitude to almost 10^4 .

4.2.2. Several sources varying in frequency and time

The simulation described in the present section was devised to be an exacting test of the Conway/Sault-Wieringa technique as applied to sources which vary both in frequency and time. The input data were derived from a model containing five point sources, each of which had a power-law spectrum. The first source was to serve as a reference, and so had a flat spectrum with no time variation. For the remaining sources, both the average flux density and the spectral index were made to vary over time. The frequency and time dependences of the remaining sources are illustrated in Fig. 5.

The array parameters for this simulation were again those of MERLIN, with a phase centre declination of $+40^\circ$; however the bandwidth chosen this time spanned from 5 to 7 GHz, divided into 50 channels of width 40 MHz. The integration time was set to 60 s. These chosen values for channel width and integration time are much coarser than those usually associated with MERLIN, and would set severe limits to the breadth of field in a real observation. In the present case, the field of interest is narrow, so coarse values were chosen in order to reduce the computing load. The chosen values were the largest ones consistent with accurate gridding of the relatively small area of the field occupied by the simulated sources.

The simulated visibilities were gridded with uniform weighting and subjected separately to 2000 cycles of Högbom cleaning at a loop gain of 0.1 versus the same number of cycles, at the same gain, of Sault-Wieringa cleaning, using FT basis functions as described in Sect. 3.1, the F functions being Chebyshev polynomials (although to the small order used, these are scarcely distinguishable from Taylor series terms) and the T functions being the half-frequency Fourier cosine functions of Eq. (7).

Fig. 6. **a)** An image constructed from the average flux densities of the input sources to the simulation. **b)** The dirty image. **c)** The result of 2000 cycles of Högbom clean at a loop gain of 0.1. **d)** The result of a Sault-Wieringa clean (with the same niter and loop gain) as described in the text. The unsaturated brightness range for all four images was -0.02 to $+0.02$ Jy/beam. (For comparison purposes, the brightest pixel of the source image was 1.27 Jy/beam.) The grey scale is reversed. Black rectangles indicate the area chosen for the rms measurement as described in the text.

Table 1. RMS values from cleaned images.

Image name	rms (Jy)
Dirty image	$1.588\text{e-}2$
Högbom cleaned	$4.847\text{e-}3$
SW cleaned	$5.415\text{e-}4$

A number of images relating to this simulation are shown in Fig. 6.

An rms value was calculated for those pixels within a rectangular area which is shown on the plots by a black outline. These numbers are shown in Table 1. The Sault-Wieringa technique clearly offers a significant improvement in dynamic range.

As discussed in Sect. 3.4, the value of this technique lies more in its ability to remove artifacts from an image of average flux density, rather than as a way to estimate light curves or spectra. For wide-band observations in which there is no time variation over the observation, because of the typically slowly-varying nature of radio spectra, one may reasonably expect to extract low-order spectral information (such as spectral indices) from Conway decomposition. But because light curves may easily contain significant power at higher orders of the time basis functions, which are not so accurately recovered by the deconvolution, more caution is advisable in interpreting the time-dependent output of the deconvolution process. If an accurate light curve of a source is desired, it is probably always going to be preferable to just phase-rotate the array to that sky location.

4.2.3. Simpler method for a single variable source

A source which shows significant time variation within the course of a day must have a spatial dimension less than about a light-day across. When observed with, for example, the MERLIN array at 21 cm, any such source more distant than about a kiloparsec will be unresolved. Except in the case of masers, it is also unusual to find two such variable sources in close proximity. For many observations of time-variable radio

Fig. 7. **a)** An image constructed from the average flux densities of the input sources to the simulation. **b)** The dirty image. **c)** The result of 5000 cycles of Högbom clean at a gain of 0.01 . **d)** Shows the distribution of the zeroth component of a Sault-Wieringa clean, with similar gain and number of clean cycles, using the two-beam decomposition as described in the text. The unsaturated brightness scale was -0.005 to $+0.005$ for **a)** and **d)**, -0.05 to $+0.05$ for **b)** and **c)**. The grey scale is reversed for all images.

sources, therefore, we may expect the source to be point-like (unresolved) and the only source in the field which shows a significant time variation. In this case the light curve for the whole observed field, s_{total} , is just a simple sum of the source light curve s_{source} plus a time-invariant background flux density. Any point in the image can therefore be expressed as a sum of just two basis functions:

$$T_0(t) = 1$$

$$T_1(t) = s_{\text{total}} - \langle s_{\text{total}} \rangle.$$

Here the mean notation $\langle \rangle$ is used to indicate that beam 1 is constructed as follows. A raw beam 1 is first formed by gridding, weighting and transforming the visibilities from a source at the phase centre which has a light curve given by s_{total} . An amount of beam 0 is then subtracted from it such that the central value of the result is zero. Beam B_1 is adjusted in this way so that the average flux-density information over the field is contained entirely in the distribution of component 0.

A further simulation was constructed to test this technique. This simulation contained a time-variable point source located at the phase centre together with much fainter, extended emission (actually made up of 24 closely-spaced point sources) extending over about 0.2 arcsec (equal to 20 image pixels) either side of the central source. The average flux density of the central source was 1 Jy/beam whereas the extended emission ranged in brightness from about 3.5×10^{-3} to 1.5×10^{-3} Jy/beam. The light curve of the central source was the same as diagrammed in Fig. 1, but without the data gap. It brightens by 3.7 mag in the course of the observation.

A quartet of images (similar to Fig. 6) to exhibit the performance of this technique is shown in Fig. 7. For constructing the visibilities, MERLIN specifications were again used. The integration time was 5 s, and 32 channels of width 1 MHz, starting at 6 GHz, were specified.

It is easily seen that the Sault-Wieringa deconvolution recovers almost all the faint emission. The Högbom clean is unable to remove sidelobes at a level 10 times the flux density of the extended emission and thus is incapable of revealing it. Several

values of the Högbom gain and number of iterations were tried with no improvement on what is displayed here.

5. Conclusions

In this paper, a technique developed originally by Conway et al. (1990) and Sault & Wieringa (1994) to allow cleaning of multi-frequency-synthesis images has been generalized and shown to be applicable to earth-rotation-synthesis observations in which some of the sources vary significantly in brightness over the course of the observation. Sources which vary over the course of a day are not very common, but they do occur, and are sometimes (in the case for example of novae) sources which it is of the highest interest to map accurately. But in fact one does not have to look for natural variations in flux density to encounter this problem: any movement of the primary beam of the array on the sky during an observation will generate artificial fluctuations, not only in the average flux density of sources, but, due to the frequency-dependent size of the primary beam, also in their spectral indices. Since some kind of pointing error in a mechanically tracking dish is probably unavoidable, this is likely to be a problematic issue when performing any kind of wide-field or mosaiced observation (see e.g. Bhatnagar et al. 2008), particularly so in view of current hopes for improvements in dynamic range from the several wide-band, dish-antenna arrays, such as eVLA, eMERLIN and MeerKAT, which are currently under construction.

Deconvolution of time-varying sources reveals some issues which are not usually encountered in the frequency-synthesis case. The principal one of these is that choice of basis function is now of some importance. This issue was explored with some care, in particular the avoidance of Gibbs ringing when periodic basis functions are employed. It was also shown that, provided the time variation in the field is limited to a single, unresolved source, a particularly simple 2-beam technique can produce an almost perfect deconvolution.

Whereas the original parallel decomposition treatment employed at most 3 or 4 beams, some of the situations described in the present paper required as many as 30. Use of such large numbers of beams gives rise to a computational difficulty due to the requirement in the original Sault-Wieringa algorithm to store, for N input beams, of order N^2 cross-correlation images. It is shown here that this memory load can be much reduced if the set of beams is made orthogonal. An algorithm for doing this was proposed and it was shown how “de-orthogonalized” clean components can be recovered at the end of the cleaning process via a simple matrix inversion.

Acknowledgements. We thank the anonymous referee for some sound advice and for directing our attention to the compressive sensing papers.

References

Becker, R. H., Helfand, D. J., White, R. L., & Proctor, D. D. 2010, *AJ*, 140, 157
 Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *A&A*, 487, 419
 Bignall, H. E., Jauncey, D. L., Kedziora-Chudczer, L., et al. 2009, *ASP Conf. Ser.*, 402, 256

Bower, G. C., Saul, D., Bloom, J. S., et al. 2007, *ApJ*, 666, 346
 Candès, E. J., & Wakin, M. B. 2008, *IEEE Sig. Proc. Mag.*, 25, 21
 Clark, B. G. 1980, *A&A*, 89, 377
 Conway, J. E., Cornwell, T. J., & Wilkinson, P. N. 1990, *MNRAS*, 246, 490
 Cornwell, T. J. 1983, *A&A*, 121, 281
 Cornwell, T. J. 2008, *IEEE J. Sel. Topics Signal Proc.*, 2, 793
 Cornwell, T., Braun, R., & Briggs, D. 1999, *Synthesis Imaging in Radio Astronomy II*, *ASP Conf. Ser.*, 180, 151
 Dent, W. A. 1965, *Science*, 148, 1458
 De Vries, W. H., Becker, R. H., White, R. L., & Helfand, D. J. 2004, *AJ*, 127, 2565
 Eyres, S. P. S., Bode, M. F., O’Brien, T. J., Watson, S. K., & Davis, R. J. 2000, *MNRAS*, 318, 1086
 Eyres, S. P. S., Heywood, I., O’Brien, T. J., et al. 2005, *MNRAS*, 358, 1019
 Fender, R. P., Garrington, S. T., McKay, D. J., et al. 1999, *MNRAS*, 304, 865
 Fuhrmann, L., Krichbaum, T. P., Witzel, A., et al. 2008, *A&A*, 490, 1019
 Golub, G. H., & Van Loan, C. F. 1996, *Matrix Computations*, 3rd edition (The Johns Hopkins University Press)
 Hayama, K., Mohanty, S. D., & Nayak, K. R. 2006, 6th International LISA Symposium, *AIP Conf. Proc.*, 873, 465
 Haynes, R. F., Jauncey, D. L., Murdin, P. G., et al. 1978, *MNRAS*, 185, 661
 Heywood, I., & O’Brien, T. J. 2007, *MNRAS*, 379, 1453
 Hjellming, R. M. 1996, *ASP Conf. Ser.*, 93, 174
 Högbom, J. A. 1974, *A&AS*, 15, 417
 Jauncey, D. L., Kedziora-Chudczer, L., Lovell, J. E. J., et al. 2001, *Ap&SS*, 278, 87
 Kedziora-Chudczer, L. L., Jauncey, D. L., Wieringa, M. H., Tzioumis, A. K., & Reynolds, J. E. 2001, *MNRAS*, 325, 1411
 Kovalev, Y. Y., Kellermann, K. I., Lister, M. L., et al. 2005, *AJ*, 130, 2473
 Krauss, M., Sokolowski, J., Chomiuk, L., et al. 2011, *BAAS*, 43, presentation 304.07
 Lannes, A., Anterrieu, E., & Maréchal, P. 1997, *A&AS*, 123, 183
 Lazio, T. J. W., Ojha, R., Fey, A. L., et al. 2008, *ApJ*, 672, 115
 Li, F., Cornwell, T. J., & de Hoog, F. 2011, *A&A*, 528, A31
 Lovell, J. E. J., Rickett, B. J., Macquart, J.-P., et al. 2008, *ApJ*, 689, 108
 McCallum, J. N., Ellingsen, S. P., Lovell, J. E. J., Phillips, C. J., & Reynolds, J. E. 2009, *MNRAS*, 392, 1339
 O’Brien, T. J., Bode, M. F., Porcas, R. W., et al. 2006a, The 2006 explosion of the recurrent nova RS Ophiuchi, *Proceedings of the 8th European VLBI Network Symposium*
 O’Brien, T. J., Bode, M. F., Porcas, R. W., et al. 2006b, *Nature*, 442, L279
 Ojha, R., Fey, A. L., Lazio, T. J. W., et al. 2006, *ApJ*, 166, 37
 Quirrenbach, A., Witzel, A., Wagner, S., et al. 1991, *ApJ*, 372, L71
 Rushton, A., Spencer, R. E., Pooley, G., & Trushkin, S. 2010, *MNRAS*, 401, 2611
 Sault, R., & Wieringa, M. H. 1994, *A&AS*, 108, 585
 Schwab, F. R. 1984, *ApJ*, 89, 1076
 Schwartz, R. A. 2009, *BAAS*, 41, 839
 Steer, D. G., Dewdney, P. E., & Itoh, M. R. 1984, *A&AS*, 137, 159
 Stewart, I. M. 2008, Applications of multi-beam CLEAN, *Oxford Workshop on Imaging and Calibration Algorithms for EVLA, e-MERLIN and ALMA*, http://astrowiki.physics.ox.ac.uk/pub/Algorithms2008/WebHome/StewartOxford_nov_08.pdf
 Stewart, I. M. 2010, Parallel CLEAN: beyond the frequency domain, *ADASS XIX*, *ASP Conf. Ser.*, 434, 410
 Taylor, G. B., Carilli, C. L., & Perley, R. A. 1999, *Synthesis Imaging in Radio Astronomy II*, *ASP Conf. Ser.*, 180
 Tan, S. M. 1986, *MNRAS*, 220, 971
 Thompson, R. A., Moran, J. M., & Swenson, G. M. 2001, *Interferometry and Synthesis in Radio Astronomy*, 2nd edition (Wiley-VCH)
 Tudose, V., Fender, R. P., Tzioumis, A. K., Spencer, R. E., & van der Klis, M. 2008, *MNRAS*, 390, 447
 Tudose, V., Miller-Jones, J. C. A., Fender, R. P., et al. 2009, *MNRAS*, 401, 890
 Wagner, S. J., & Witzel, A. 1995, *ARA&A*, 33, 163
 Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., & Vanderghynst, P. 2009, *MNRAS*, 395, 1733