HAL
open science

The variance of radio interferometric calibration solutions

Etienne Bonnassieux, Cyril Tasse, Oleg Smirnov, Philippe Zarka

To cite this version:

Etienne Bonnassieux, Cyril Tasse, Oleg Smirnov, Philippe Zarka. The variance of radio interferometric calibration solutions. Astronomy and Astrophysics - A&A, 2018, 615, pp.A66. 10.1051/0004-6361/201732190 . insu-01969808

HAL Id: insu-01969808

https://insu.hal.science/insu-01969808v1

Submitted on 20 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A&A 615, A66 (2018)
https://doi.org/10.1051/0004-6361/201732190
© ESO 2018

Astronomy & Astrophysics

The variance of radio interferometric calibration solutions

Quality-based weighting schemes

Etienne Bonnassieux[1,2], Cyril Tasse[3], Oleg Smirnov[2,5], and Philippe Zarka[1,4]

[1] LESIA, Observatoire de Paris, PSL, CNRS, Sorbonne Université, UPMC Univ. Paris 06, Univ. Paris Diderot, Sorbonne Paris Cité, 5 place Jules Janssen, 92195 Meudon, France
e-mail: etienne.bonnassieux@obspm.fr
[2] Department of Physics & Electronics, Rhodes University, PO Box 94, Grahamstown 6140, South Africa
[3] GEPI, Observatoire de Paris, 5 place Jules Janssen, 92190 Meudon, France
[4] Station de Radioastronomie de Nançay (SRN, Observatoire de Paris, CNRS USN, PSL Research University, Université d'Orléans, 18330 Nançay, France
[5] South Africa Radio Astronomy Observatory, 3rd floor, The Park, Park Roads, 7405 Pinelands, South Africa

Received 27 October 2017 / Accepted 6 February 2018

ABSTRACT

This paper investigates the possibility of improving radio interferometric images using an algorithm inspired by an optical method known as "lucky imaging", which would give more weight to the best-calibrated visibilities used to make a given image. A fundamental relationship between the statistics of interferometric calibration solution residuals and those of the image-plane pixels is derived in this paper. This relationship allows us to understand and describe the statistical properties of the residual image. In this framework, the noise map can be described as the Fourier transform of the covariance between residual visibilities in a new differential Fourier plane. Image-plane artefacts can be seen as one realisation of the pixel covariance distribution, which can be estimated from the antenna gain statistics. Based on this relationship, we propose a means of improving images made with calibrated visibilities using weighting schemes. This improvement would occur after calibration, but before imaging; it is thus ideally used between major iterations of self-calibration loops. Applying the weighting scheme to simulated data improves the noise level in the final image at negligible computational cost.

Key words. instrumentation: interferometers – instrumentation: adaptive optics – methods: analytical – methods: statistical – techniques: interferometric – radio continuum: general

1. Introduction

Interferometers sample Fourier modes of the sky brightness distribution corrupted by instrumental and atmospheric effects rather than measuring the sky brightness directly. This introduces two problems for astronomers to invert: calibration and imaging. Both of these problems are ill-conditioned.

The problem of imaging consists of correcting for the incomplete *uv*-coverage of any given interferometer by deconvolving the instrument's point-spread function (PSF) from images. Its poor conditioning comes from our limited a priori knowledge of the sky brightness distribution, combined with large gaps in our *uv*-coverage, which prevents us from placing strong constraints on image deconvolution. It can be better-conditioned in different ways, including through the use of weighting schemes (see Briggs 1995; Yatawatta 2014, and references therein) to improve image fidelity at the start of deconvolution. When inverting the imaging problem, we often assume that the sky is stable within the domain (i.e. is constant in time and frequency). There are exceptions, such as wide-band deconvolution algorithms (e.g. Rau & Cornwell 2011) that explicitly take into account the sky's frequency dependence, but still assume that the sky brightness distribution does not vary with time.

The problem of calibration is what concerns us in this paper. It consists of estimating and correcting for instrumental errors (which include effects such as antenna pointing errors, but also the phase delays caused by ionospheric activity, troposphere, and other factors). Calibration consists of solving for gain estimates, where a gain models the relationship between the electromagnetic field of an astrophysical source and the voltage that an antenna measures for this source. Because measurements are noisy, calibration often involves some fine-tuning of solution intervals to ensure that the solutions are well-constrained while the solution intervals stay as small as signal-to-noise allows. The calibration inverse problem involves three competing statistical effects: thermal noise in the measurements, true gain variability, and sky model incompleteness. If gain solutions are sought individually for each measurement, then calibration estimates will be dominated by thermal noise and will not adequately describe the actual gains. Similarly, if a single gain estimate is fitted to too many measurements, the intrinsic gain variability will be averaged out; for example, a choice of time and frequency interval that is too large will cause the solver to estimate a constant gain while the underlying function varies quickly, thereby missing much of the gain structure. This will introduce error, which will be correlated in time and frequency. This occurs, for example, when solving for ionospheric phase delays: in the most extreme case, where the solution interval is significantly larger than the scale of ionospheric fluctuations, its varying phase can average out to zero over the interval in time and frequency. Finally, if the model being fitted is incomplete, un-modelled physical flux will likely be absorbed unpredictably into both the gain solutions

Open Access article, published by EDP Sciences, under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

and the residual visibilities: this absorption of physical flux into gain solutions is known as source suppression (see Grobler et al. 2014; Kazemi & Yatawatta 2013, and references therein).

In practice, it is reasonable to assume that gain variation is generally slower than some given scale: we can then reduce the noise of our gain estimates by finding a single gain solution for a small number of measurements, assuming that the underlying gain variation is very small and stable over short intervals. This is generally a valid hypothesis, but the specific value for the variation scale can be contentious. Indeed, while the noise level can often be treated as constant throughout an observation, the gain variability itself is generally not constant: there will be time periods where the gains will tend to remain constant for longer, and others where variability will be very quick. This means that, for any choice of calibration interval, some gain estimates will be better than others, and almost all could have been improved (at a cost to others) by a different choice of time (and frequency) intervals.

Since we have measurements that are better-calibrated than others (in that better estimates for their gains were obtained through chance alone), we could, in principle, take inspiration from "lucky imaging" (an optical-domain method for making good images: for more details, see Fried 1978, and references therein) to weigh our visibilities according to their calibration quality. Those weights would in effect be an improvement of currently existing methods such as clipping noisy residual visibilities: in the extreme case where all visibilities are equally-well calibrated except a few which are extremely noisy, it should be equivalent to clipping. Otherwise, the weights should show at least a slight improvement over clipping.

The key finding of the present paper is a fundamental relationship between the covariance matrices of residual visibilities and the map of the covariance in the image plane: the Cov-Cov relationship between visibility covariance to image-plane covariance. We show that the pixel statistics in the image plane are determined by a noise PSF, convolved with each source in the sky (modelled or not). This noise PSF is the product of the Fourier transforms of the gain covariance matrix with each cell mapped not from uv space to lm coordinates but rather between their respective covariance spaces: from a new differential Fourier plane (henceforth ($\delta u \delta v$)-plane) to the image-plane covariance space $\delta l \delta m$. This image-plane covariance space describes the variance in each pixel and the covariance between pixels[1]. It describes the expected calibration artefacts and thermal noise around each source, does not vary as a function of direction, and is convolved with each source in the field to yield the final error map. Because all unwanted (in our case, unphysical) signal can be thought of as noise, we will refer to the pixel variance map as the noise map.

The notion of a ($\delta u \delta v$)-plane arises organically from the framework of radio interferometry: we are associating a correlation between visibilities to coordinates in covariance space, just as we associate the visibilities themselves to the uv-domain. The ($\delta u \delta v$)-plane is the natural domain of these correlations. As previously stated, even if all sources in the field are perfectly known and modelled, a poor choice of calibration interval can introduce correlated noise in the residuals, which would then introduce larger variance near sources in the noise map. Conversely, if calibration is perfect, the noise map should be completely flat (i.e. same variance for all pixels), as there would be no noise correlation between pixels.

The main result of this paper consists in describing a new adaptive, quality-based weighting scheme based on this insight. Using the Cov-Cov relationship, we can create a new weighting scheme by estimating the residual visibility covariance matrix in a given observation. By weighting visibilities so as to change their covariance matrix, one can change the shape of the noise PSF and thus improve the final image: this manifests as either decreased noise or decreased calibration artefacts. This weighting is applied after calibration, but before image deconvolution; applying it will therefore not only improve the residual noise in the image, and thus the sensitivity achievable with a given pipeline, but will also improve deconvolution by minimising calibration artefacts in the field. It should thus effectively remove spurious, un-physical emission from final data products. Estimating the covariance matrix is the main difficulty of our framework; we do not know the underlying covariance matrix, and the conditioning of our estimation thereof is limited by the number of measurements within each solution interval. As such, we have no guarantee that our estimate of the corrected visibility covariance matrix is accurate. This problem can be alleviated, for example, by estimating the covariance matrix for the antenna gains themselves, and using this estimate to build the visibility covariance matrix: this effectively improves conditioning (cf. Sect. 4).

This paper is split into four main sections. In Sect. 2, we derive the Cov-Cov relationship. With its newfound insights, we propose quality-based weighting schemes with which to improve radio interferometric images in Sect. 3. We follow in Sect. 4 by showing how to estimate, from real data, the covariance matrix from which the quality-based weights are derived. Our approach seems to give good results. Finally, we close the paper on a discussion of the applicability and limitations of the quality-based weighting scheme.

2. Building the noise map

In this section, we derive our first fundamental result: the Cov-Cov relationship, Eq. (25), which describes how the statistics of residual visibilities (and thus the antenna calibration solutions, henceforth gains) relate to the statistics of the image plane, that is, of images made using the associated visibilities. The dimensions of the matrices (denoted by boldface capital letters) and vectors (denoted by boldface lowercase letters) used in this paper are given in Table 1, along with the scalar numbers used to denote specific dimensions. All other variables are scalars.

2.1. The Cov-Cov relationship in the $\delta u \delta v$ plane

Let us begin by defining visibility gains. Using the Radio Interferometry Measurement Equation formalism for a sky consisting of a single point source (Hamaker et al. 1996; Smirnov 2011; and companion papers), we can write the following relation between the sky and the signal as measured by a single baseline at time t and frequency ν:

$$\boldsymbol{V}^{t\nu}_{pq} = \sum_d \boldsymbol{K}^d_{p,t\nu} \boldsymbol{J}^d_{p,t\nu} \boldsymbol{B}^d_\nu \left(\boldsymbol{J}^d_{q,t\nu}\right)^H \left(\boldsymbol{K}^d_{q,t\nu}\right)^H + \boldsymbol{N}. \tag{1}$$

All the quantities above are 2×2 matrices. Equation (1) implies a linear relationship between the coherency matrix $\boldsymbol{B}^d_\nu$ and the visibilities recorded by a given baseline ($\boldsymbol{V}^{t\nu}_{pq}$), with the addition of a thermal noise matrix $\boldsymbol{N}$, which is also of shape

[1] The noise PSF also relates δw to δn, as shown in the matrix formalism, but this is not explicitly referenced in the text since visibility space is usually referred to as "the UV-plane" in literature, rather than "the UVW-space".

Table 1. Meaning and dimensions of vectors and matrices used in Sect. 2.1.

Scalars	
$n_{\rm pix}$	Total number of pixels in image plane & number of cells in uv-grid
$n_{\rm ant}$	Total number of antennas in the array
n_b	Number of visibilities
b	Index for a single visibility. Equivalent to $(pq, t\nu)$
τ	Equivalent to (t, ν)
Vectors	
$\tilde{\boldsymbol{y}}$	Residual image vector, size $n_{\rm pix}$
$\boldsymbol{\epsilon}$	Vector of ϵ, size $n_{\rm pix}$
$\widetilde{\boldsymbol{\gamma}}$	Contains gain products, size n_b. see Eq. (12)
$\mathbf{1}$	Vector containing 1 in every cell, size n_b
$\boldsymbol{\delta u}_{bb'}$	Vector of coordinates in differential Fourier plane, of length 3.
$\boldsymbol{l}_d$	Vector of sky coordinates, of length 3.
Matrices	
$\boldsymbol{V}^{t\nu}_{pq}$	Visibility seen by a baseline pq at time and frequency t, ν. Size 2×2.
$\boldsymbol{K}^d_{p,t\nu}$	Fourier kernel for direction d, antenna p and one (t, ν) pair. Size 2×2.
$\boldsymbol{J}_{p,t\nu}$	Jones matrix for antenna p for one (t, ν) pair. Contains the gains information. Size 2×2
$\boldsymbol{B}$	Sky brightness distribution matrix, of size 2×2.
$\boldsymbol{N}$	Noise matrix, of size 2×2. Contains a single realisation n of the thermal noise in each cell.
$\boldsymbol{\mathcal{F}}$	Fourier transform matrix, of size $n_{\rm pix} \times n_{\rm pix}$.
$\boldsymbol{S}_b$	Baseline selection matrix, which picks out 1 visibility out of the full set. Size $n_{\rm pix} \times n_b$
$\boldsymbol{C}_b$	$n_{\rm pix} \times n_{\rm pix}$ convolution kernel that defines the PSF.
$\boldsymbol{\mathcal{F}}_{bb'}$	Convolution matrix mapping one $\delta u \delta v$ to $\delta l \delta m$. The set of all $\boldsymbol{\mathcal{F}}_{bb'}$ determines the noise-PSF. Size $n_{\rm pix} \times n_{\rm pix}$.

Notes. Only scalars which give matrix dimensions or indices are given here.

2×2 and contains different complex-valued realisations of the noise in each cell. Since electric fields are additive, the sky coherency matrix can be described as the sum of the contributions from individual sources in directions d in the sky. We also assume that the sky does not vary over time, that is, that $\boldsymbol{B}^d_\nu$ is not a function of time. The Jones matrices ($\boldsymbol{J}^d_{\ldots,t\nu}$) contain the antenna gain information in matrix form, while $\boldsymbol{K}^d_{\ldots,t\nu}$ is the Fourier kernel. Let us limit ourselves to the scalar case, which corresponds to assuming that emission is unpolarised. We assume that $\boldsymbol{B}^d = s_d \mathbf{I}$, where s is the flux of our single point source and $\mathbf{I}$ is the 2×2 identity matrix. We also assume that $\boldsymbol{J}_{p,t\nu} = g_{p,t\nu}\mathbf{I}$, where $g^{t\nu}_p$ is the complex-valued gains of antenna p at time t and frequency ν. This means that we assume that the gains are direction-independent, and so $\boldsymbol{J}^d_{\ldots,t\nu}$ becomes $\boldsymbol{J}_{\ldots,t\nu}$. Similarly, $\boldsymbol{K}^d_{p,t\nu} = k^d_{p,t\nu}\mathbf{I}$, the Fourier kernel in the direction of the source, d. The noise matrix $\boldsymbol{N}$ has one realisation of ϵ in each cell, where

$$\epsilon \sim \mathcal{N}(0, \sigma) + i\mathcal{N}(0, \sigma), \tag{2}$$

where σ is the variance of the thermal noise. Let us denote each (t, ν) pair by τ, and ignore the sky's frequency dependence. The following scalar formulation is then equivalent to Eq. (1):

$$V^\tau_{pq} = \left(\sum_d s_d k^d_{p,\tau}\overline{k^d_{q,\tau}}\right) g^\tau_p \overline{g^\tau_q} + \epsilon, \tag{3}$$

$$k^d_{p,\tau} = \exp\left(2\pi i\left(u_{p,\tau}l_d + v_{p,\tau}m_d + w_{p,\tau}\left(n_p - 1\right)\right)\right). \tag{4}$$

Calibration is the process of finding an accurate estimate of g^τ_p for all antennas p, at all times t and frequencies ν. Since we are in a direction-independent regime, the quality of our calibration then determines the statistical properties of the residual visibilities (and the image-plane equivalent, the residual image). The residual visibilities associated with calibration solutions are defined as our measured visibilities minus the gain-corrupted model visibilities. The variable $\widehat{g^\tau_p}$ then denotes our calibration estimate for g^τ_p. We now begin to limit the generality of our framework by assuming that all sufficiently bright sources have been modelled and subtracted: unmodelled flux is then negligible. We can then write the residual visibilities as

$$\tilde{r}^\tau_{pq} = \sum_d s_d \left(k^d_{p,\tau}\overline{k^d_{q,\tau}}\right)\left(g^\tau_p\overline{g^\tau_q} - \widehat{g^\tau_p}\overline{\widehat{g^\tau_q}}\right) + \epsilon. \tag{5}$$

The flux values in the image-plane pixels[2] are the Fourier transform of the visibility values mapped onto each pixel. This can be written as

$$\tilde{\boldsymbol{y}} = \begin{pmatrix} \vdots \\ \sum_{pq} I^r_{pq,lm} \\ \vdots \end{pmatrix} \tag{6}$$

$$I^r_{pq,lm} = \sum_\tau \omega_{pq,\tau}\tilde{r}^\tau_{pq}k^{lm}_{pq,\tau}, \tag{7}$$

where lm are the directional cosine positions of a given pixel, and $k^{lm}_{pq,\tau} = k^d_{p,\tau}\overline{k^d_{q,\tau}}$ the Fourier coefficient mapping a point in Fourier space to a point on the image-plane. The scalar $\omega_{pq,\tau}$ is the weight associated to a given visibility. Let us now write this using a matrix formalism. The contribution of a single visibility $b = (pq, \tau)$ to the image-plane residuals can be written as

$$\tilde{\boldsymbol{y}}_b = \boldsymbol{\mathcal{F}}^H \boldsymbol{S}_b \omega_b \left(\kappa_b \tilde{\boldsymbol{\gamma}} + \boldsymbol{\epsilon}\right), \tag{8}$$

$$\tilde{\boldsymbol{y}} = \sum_b \tilde{\boldsymbol{y}}_b, \tag{9}$$

where $\boldsymbol{\epsilon}$ is a vector of the ϵ of Eq. (2) and $\tilde{\boldsymbol{y}}$ is a vector of size $n_{\rm pix}$, with

$$\kappa_b = \sum_d s_d \kappa^d_b, \tag{10}$$

$$\kappa^d_b = k^d_{p,\tau}\overline{k^d_{q,\tau}} = k^{lm}_{pq,\tau}, \tag{11}$$

$$\widetilde{\gamma}_b = g^\tau_p\overline{g^\tau_q} - \widehat{g^\tau_p}\overline{\widehat{g^\tau_p}}, \tag{12}$$

$$\tilde{\boldsymbol{\gamma}} = \begin{pmatrix} \vdots \\ \widetilde{\gamma}_b \\ \vdots \end{pmatrix}, \tag{13}$$

[2] As opposed to the Fourier-plane pixels, which are the elements of the grid onto which the measured visibilities are mapped for imaging.

and w_b is the scalar weight associated with each visibility. By default, $w_b = \frac{1}{n_b}$: all visibilities then have the same weight, and $\tilde{\boldsymbol{y}}$ then becomes the average of all $\tilde{\boldsymbol{y}}_b$. Here, $\tilde{\boldsymbol{\gamma}}$ is a vector of all $\tilde{\gamma}_b$, and thus of size n_b. Similarly, $\boldsymbol{\mathcal{F}}$ is the Fourier kernel, of size $n_{\rm pix} \times n_{\rm pix}$. Finally, $\boldsymbol{S}_b$ is a matrix of size $n_{\rm pix} \times n_b$: its purpose is to encode the *uv*-coverage. Each $\boldsymbol{S}_b$ contains only a single non-zero cell, different for different $\boldsymbol{S}_b$. The height (number of rows) of $\boldsymbol{S}_b$ is determined by the size of the *uv*-grid, and its length (number of columns) by the number of visibilities.

The order of operations is thus the following: each residual visibility $(\kappa_b \tilde{\boldsymbol{\gamma}} + \boldsymbol{n})$ is assigned some weight w_b and its *uv*-coordinates are set by $\boldsymbol{\mathcal{S}}_b$. The inverse Fourier transform $(\boldsymbol{\mathcal{F}}^H)$ is then applied to this grid, and so we recover its image-plane fringe. By averaging over all fringes, we recover the dirty image. The residual image will thus depend on three quantities: the residual gains, the flux in the image, and the weighting scheme. Let us consider the relationship between the statistics of residual visibilities and the variance at a given point in the corresponding residual image.

2.2. Statistical analysis

In the following analysis, we treat our gain solutions and thermal noise as random variables in order to compute the covariance matrix of our residual image, $\mathrm{Cov}\{\tilde{\boldsymbol{y}}\}$. The diagonal of this matrix gives the variance for each pixel, while the wings give the covariance between pixels. Using the property that $\mathrm{Cov}\{\boldsymbol{A}\boldsymbol{x}\} = \boldsymbol{A}\mathrm{Cov}\{\boldsymbol{x}\}\boldsymbol{A}^H$, we can apply the Cov{} operator to Eq. (9) to write

$$\mathrm{Cov}\{\tilde{\boldsymbol{y}}\} = \sum_{bb'} \boldsymbol{\mathcal{F}}^H \underbrace{w_b w_{b'} \kappa_b \overline{\kappa_{b'}}}_{\stackrel{\rm def}{=} \phi_{bb'}} \boldsymbol{S}_b \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\} \boldsymbol{S}_{b'}^T \boldsymbol{\mathcal{F}} \tag{14}$$

$$+ \sum_b w_b^2 \sigma^2 \underbrace{\boldsymbol{\mathcal{F}}^H \boldsymbol{S}_b \mathbf{I} \boldsymbol{S}_b^T \boldsymbol{\mathcal{F}}}_{\stackrel{\rm def}{=} \boldsymbol{C}_b} \tag{15}$$

$$= \sum_{bb'} \phi_{bb'} \boldsymbol{\mathcal{F}}^H \boldsymbol{S}_b \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\} \boldsymbol{S}_{b'}^T \boldsymbol{\mathcal{F}} + \sum_b w_b^2 \sigma^2 \boldsymbol{C}_b. \tag{16}$$

So far, we have only applied definitions. The net effect of $\boldsymbol{S}_b \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\} \boldsymbol{S}_{b'}^T$ (dimensions of $n_{\rm pix} \times n_{\rm pix}$) is to encode where a given baseline samples the *uv*-plane, and map one cell at matrix coordinates (b, b') from the correlation matrix $\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}$ onto the visibility grid. The matrix $\boldsymbol{\mathcal{S}}_b$ is not the gridding kernel, but rather the sampling matrix, which determines where we have measurements and where we do not. We can thus write that $\boldsymbol{S}_b \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\} \boldsymbol{S}_{b'}^T = [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'} \boldsymbol{S}_b \mathbf{1}\mathbf{1}^T \boldsymbol{S}_{b'}^T$, where $[\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'}$ is the value from the appropriate cell and $\mathbf{1}$ is the vector-of-ones of appropriate length (here, n_b). This allows us to write

$$\mathrm{Cov}\{\tilde{\boldsymbol{y}}\} = \sum_{bb'} \phi_{bb'} [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'} \boldsymbol{\mathcal{F}}_{bb'} + \sum_b w_b^2 \sigma^2 \boldsymbol{C}_b, \tag{17}$$

$$\text{with} \qquad \boldsymbol{\mathcal{F}}_{bb'} = (\boldsymbol{\mathcal{F}}_b)^H \underbrace{\boldsymbol{\mathcal{F}}_{b'}}_{\stackrel{\rm def}{=} \mathbf{1}^T \boldsymbol{S}_{b'}^T \boldsymbol{\mathcal{F}}}. \tag{18}$$

Here, $\boldsymbol{C}_b$ is a Toeplitz matrix, that is, a convolution matrix, associated with baseline b. The set of all $\boldsymbol{C}_b$ defines the convolution kernel that characterises the point-spread function (henceforth PSF) associated with a given *uv*-coverage, of size $n_{\rm pix} \times n_{\rm pix}$. The matrix $\boldsymbol{\mathcal{F}}_{bb'}$, meanwhile, is not generally Toeplitz. Its cells can be written as

$$\boldsymbol{\mathcal{F}}_{bb'}[d, d'] = e^{2i\pi(u_b l_d - u_{b'} l_{d'} + v_b m_d - v_{b'} m_{d'} + (n_d - 1) w_b - (n_{d'} - 1) w_{b'})}. \tag{19}$$

Let us investigate how the sky brightness distribution (i.e. *d*-dependence) affects the noise map. We can write the sum over bb' as two sums: one over $b, b' = b$ and one over $b, b' \neq b$. Thus

$$\mathrm{Cov}\{\tilde{\boldsymbol{y}}\} = \sum_b \left(\phi_{bb} [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb} + w_b^2 \sigma^2\right) \boldsymbol{C}_b + \sum_{b, b' \neq b} \phi_{bb'} [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'} \boldsymbol{\mathcal{F}}_{bb'}. \tag{20}$$

The only direction-dependent terms in the above are s_d and κ_b^d, which are both inside $\phi_{bb'}$ (for both $b = b'$ and $b \neq b'$). By making the approximation that the Fourier kernels of different sources are orthogonal (i.e. that $\kappa_b^d \kappa_{b'}^{d'} = (\kappa_b^d)^2 \delta_{bb'}$)[3] we can write

$$\phi_{bb'} = w_b w_{b'} \kappa_b \overline{\kappa_{b'}} \tag{21}$$

$$= w_b w_{b'} \left(\sum_d s_d \kappa_b^d\right) \left(\sum_{d'} s_{d'} \overline{\kappa_{b'}^{d'}}\right) \tag{22}$$

$$\approx \sum_d w_b w_{b'} s_d^2 \kappa_b^d \overline{\kappa_{b'}^d} \tag{23}$$

$$\phi_{bb'} \approx \sum_d \phi_{bb'}^d. \tag{24}$$

We note that $\boldsymbol{\mathcal{F}}_{bb} = \boldsymbol{C}_b$, since those are the coordinates along the diagonal; for these values, the matrix-of-ones at the centre of $\boldsymbol{\mathcal{F}}_{bb'}$ becomes the identity matrix. We note also that $\phi_{bb}^d = w_b^2 s_d^2$, since $\kappa_b^d \overline{\kappa_b^d} = 1$. We can then write Eq. (20) as

$$\mathrm{Cov}\{\tilde{\boldsymbol{y}}\} = \sum_d \left(\sum_b \phi_{bb}^d \left([\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb} + \frac{w_b^2 \sigma^2}{\phi_{bb}} \right) \boldsymbol{C}_b + \sum_{b, b' \neq b} \phi_{bb'}^d [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'} \boldsymbol{\mathcal{F}}_{bb'} \right), \tag{25}$$

where we have now limited our formalism to the case where the sky is dominated by distant point-like sources. This is our fundamental result: assuming unpolarised emission coming from distant point sources and normally-distributed thermal noise, it gives a direct relationship between the covariance of the residual visibilities and the covariance of the residual image-pixel values. We thus call it the Cov-Cov relationship. It describes the statistical properties of the image plane as the result of a convolution process changing an average noise level at different points in the image plane, allowing us to describe the behaviour of variance and covariance in the image. Let us focus on the first. By applying the Diag {} operator (which returns the diagonal of an input matrix as a vector) to both sides of Eq. (25), we can find an expression for the variance map in the image plane:

$$\mathrm{Var}\{\tilde{\boldsymbol{y}}\} = \mathrm{Diag}\,\{\mathrm{Cov}\{\tilde{\boldsymbol{y}}\}\} \tag{26}$$

$$= \sum_d \left(\sum_b \left(\phi_{bb}^d [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb} + w_b^2 \sigma^2\right) \underbrace{\mathrm{Diag}\,\{\boldsymbol{C}_b\}}_{=\mathbf{1}} + \sum_{b, b' \neq b} \phi_{bb'}^d [\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{bb'} \mathrm{Diag}\,\{\boldsymbol{\mathcal{F}}_{bb'}\} \right), \tag{27}$$

$$\text{where} \qquad \boldsymbol{l}_d = (l_d, m_d, (n_d - 1)), \tag{28}$$

$$\boldsymbol{\delta u}_{bb'} = (\delta u_{bb'}, \delta v_{bb'}, \delta w_{bb'}) \tag{29}$$

$$= (u_b - u_{b'}, v_b - v_{b'}, w_b - w_{b'}). \tag{30}$$

[3] This hypothesis is equivalent to assuming that the sky is dominated by distant point sources, where "distant" means that the sources are multiple PSF full-width half-maximum apart

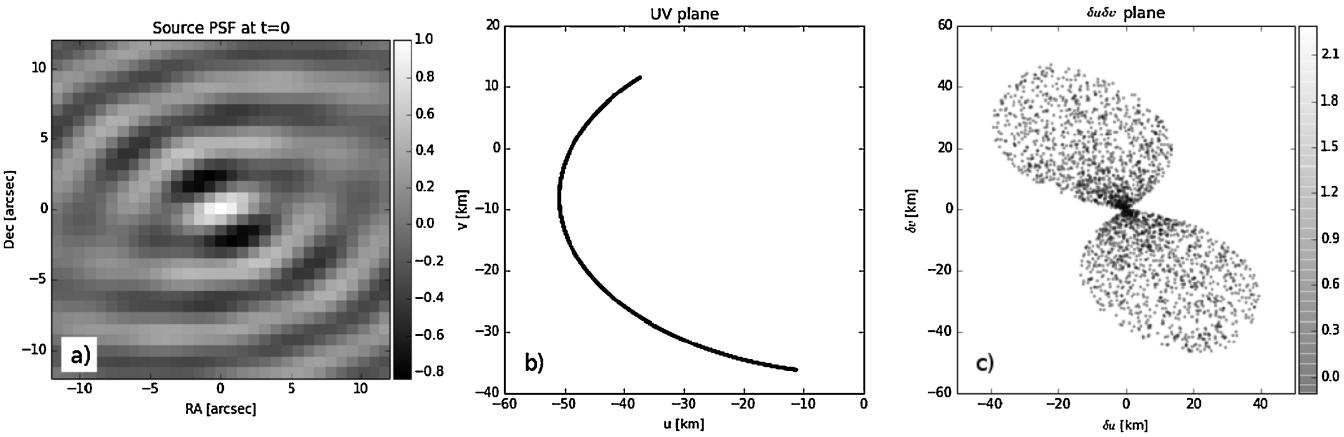

Fig. 1. *Panel a*: PSF image of a simulated 1Jy source at phase centre. Colour-bar units are in Jansky. *Panel b*: associated *UV* track, *panel c*: the corresponding $(\delta u, \delta v)$ tracks. The $\delta u \delta v$ plane does not have a homogeneous point density, but is denser near its origin: here, this is shown by plotting only one random point in 10 000.

In Eq. (27), we have

$$\mathrm{Diag}\left\{\mathcal{F}_{bb'}\right\}[d] = e^{2i\pi \boldsymbol{l}_d \cdot \boldsymbol{\delta u}_{bb'}}. \tag{31}$$

For $b \neq b'$, the diagonals of $\mathcal{F}_{bb'}$ are the Fourier kernels mapping $\delta u \delta v$ space to $\delta l \delta m$. Applying $\mathcal{F}_{bb'}$ can then be thought of as performing a Fourier transform. It is not a diagonal matrix. It behaves as a covariance fringe, allowing us to extend standard interferometric ideas to covariance space. Each fringe can be thought of as a single "spatial filter" applied to the pixel covariance matrix. Just as a given baseline has coordinates in *uv*-space, a given correlation between baseline residual errors has coordinates in *uv* correlation space, which we will henceforth refer to as $\delta u \delta v$-space.

This $\delta u \delta v$ space warrants further discussion. Figure 1 shows, for a given *uv*-track (Fig. 1b), both the corresponding $\delta u \delta v$ domain (Fig. 1c) and point-spread function (Fig. 1a). The symmetric, negative *uv*-track is treated as a separate track, and thus ignored. This means that we do not fully constrain the noise PSF (since the covariance matrix of the symmetric track is simply the Hermitian of the first), but we do not seek to constrain it in this section, but rather to show that our results hold. We can see that the $\delta u \delta v$-tracks are symmetrical about the origin. The $\delta u \delta v$ space corresponding to a given *uv*-track can thus be most concisely described as a "filled *uv*-track", with its boundaries defined by the ends of the *uv*-track. The set of $\mathcal{F}_{bb'}$, each of which maps one value of the covariance matrix to a fringe in the image plane, would then characterise a PSF equivalent for the noise distribution, which we refer to as the noise PSF. In our formalism, the only source of statistical effects in the field are calibration errors and thermal noise. The average variance in all pixels will be given by the diagonal of the covariance matrix and the thermal noise, provided that it truly follows a normal distribution. The only effects which will cause the variance in the image plane to vary from one pixel to the next are those mapped onto the covariance fringes; such position-dependent variance fluctuations will be caused by correlated gain errors, which are spurious signal introduced by erroneous gain estimates. Assuming all sources in the field are point-like and distant, then these variance fluctuations will follow a specific distribution, convolved to every source in the field. Since the variance fluctuations act as tracers for calibration artefacts, artefacts in the image can be understood as one realisation of the variance map, which is characterised by an average level determined by the variance in the gains and thermal noise, and a noise PSF convolved with the sky brightness distribution. The actual artefacts in the image will still be noisy, as a realisation of the true variance map. For the same reason, in the absence of correlated gain errors, $[\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}]_{b \neq b'}$ are all zero and $\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}$ is a diagonal matrix. We then recover a "flat" noise map: the variance will be the same in all pixels, as the noise PSF is absent. In the ideal case, were we to recover the true value of the gains for all times and frequencies, this would become pure thermal noise.

2.3. Noise-map simulations

We have shown in Eq. (25) that there exists an analytical relationship between residual visibility statistics and image-plane residual statistics. This section gives details of simulations we have performed to support our claims about this "Cov-Cov" relationship. Specifically, we simulate residual visibilities for a single baseline by generating a set of correlated random numbers with zero mean and a distribution following a specified covariance matrix $\boldsymbol{C}$. It contains a periodic function of period T along the diagonal, which is then convolved with a Gaussian of width σ_τ corresponding to the characteristic scale of correlation. The values of these parameters are chosen arbitrarily. A small constant term is added on the diagonal, the net value of which is strictly positive. This simulates a low thermal noise. Finally, singular value decomposition is used to ensure that this matrix is Hermitian positive semi-definite. The net effect is a non-stationary correlation: some residuals are correlated with their neighbours, and uncorrelated with others. An example of this covariance matrix for arbitrary parameter values is shown in Fig. 2. We see that, for any given point, correlation is stronger with some neighbours than others (as determined by σ_τ). Samples are drawn as follows: $\boldsymbol{C}$ is built following user specifications as described above. We then find its matrix square root $\boldsymbol{C}_0$ so as to apply it to random numbers generated from a normal distribution. We generate 2000 realisations i of our random variables $\boldsymbol{r}_i$:

$$\tilde{\boldsymbol{y}}_i = \sum_b \mathcal{F}_b \boldsymbol{r}_i \quad \text{with} \quad \boldsymbol{r}_i = \boldsymbol{C}_0 \boldsymbol{x}, \tag{32}$$

$$\text{and} \quad \boldsymbol{x} \leftarrow \mathcal{N}(0,1), \tag{33}$$

$$\tilde{\boldsymbol{y}} = \left(\ldots \tilde{\boldsymbol{y}}_i \ldots\right), \tag{34}$$

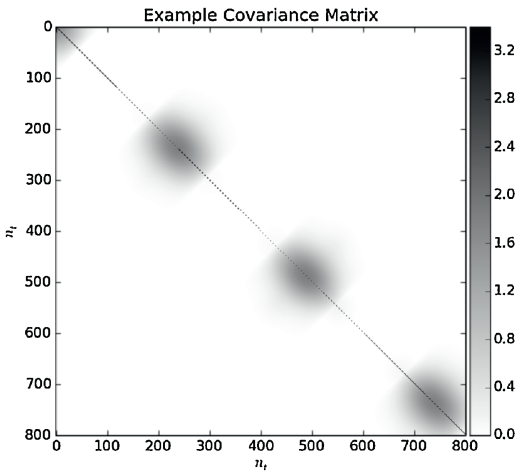

Fig. 2. Example of a non-stationary covariance matrix, which can be used to simulate Cov$\{\tilde{\boldsymbol{\gamma}}\}$. The colour bar units are Jy^2. The correlation scale σ_τ is 40 cells, and the variability period is 500 cells. The matrix is made positive semi-definite (and therefore a covariance matrix) through SVD decomposition. The maximum size of the "bubbles" is determined by σ_τ.

where $\tilde{\boldsymbol{y}}$ is a matrix of dimensions $n_{\rm realis} \times n_b$. Since $\boldsymbol{x}$ follows a normal distribution, Cov$\{x\} = \mathbf{I}$ and the covariance matrix of each $\tilde{\boldsymbol{y}}_i$ is, by construction, $\boldsymbol{C} = \boldsymbol{C}_0\boldsymbol{C}_0^H$. The covariance matrix of $\tilde{\boldsymbol{y}}$ is thus also $\boldsymbol{C}$.

As for the *uv*-track, our simulations read a single one from a specified dataset. In this case, we read an eight-hour *uv*-track for a baseline between two arbitrary LOFAR (LOw Frequency Array) stations (specifically, CS001HBA0 and RS310HBA) in an observation of the Bootes extragalactic field. The effective baseline length varies between 37.9 km and 51.8 km. The dataset includes 20 channels, each with a spectral width of 97.7 kHz; the central observing frequency is 139 MHz. The temporal resolution is one measurement per second.

We compare the measured variance map $\boldsymbol{V}_y^m$, built by measuring the variance across realisations at each pixel in the image-plane, with the predicted variance map Var$\{\tilde{\boldsymbol{y}}\}$, built using the Cov-Cov relationship (Eq. (25)). Since we are only interested in the variance map, rather than the covariance between pixels, we compute only the diagonal terms,

$$\boldsymbol{V}_{\tilde{y}}^m = \mathrm{Diag}\left\{\tilde{\boldsymbol{y}}\tilde{\boldsymbol{y}}^H\right\}, \tag{35}$$

$$\boldsymbol{V}_{\tilde{y}}^{pr} = \sum_b [\boldsymbol{C}]_{bb}\mathbf{I} + \sum_{b,b'\neq b} \mathrm{Diag}\left\{\mathcal{F}_{bb'}\right\}[\boldsymbol{C}]_{bb'}, \tag{36}$$

where the thermal noise is already incorporated into the diagonal of $\boldsymbol{C}$ and Eq. (36) is merely the diagonal operator applied to the Cov-Cov relationship.

2.3.1. Simulation with a single point source

We model our sky as containing a single 1 Jy point source at phase centre: we thus have $\phi_{bb} = w_b^2$. The source as seen through the set of *uv*-tracks used in our simulation, along with their corresponding $(\delta u, \delta v)$ space, are shown in Fig. 1. The simulated noise map is calculated by drawing a large sample ($n_{\rm realis} = 2000$) of random numbers from the correlated distribution, thereby creating 2000 sets of residual visibilities. By Fourier-transforming the visibilities to the image plane and taking the variance of the values for each image pixel (i.e. each l, m pair) as per Eq. (35), we can estimate Var$\{\tilde{\boldsymbol{y}}\}$. The predicted noise map, meanwhile, was found by assigning each cell of $\boldsymbol{C}$ to the appropriate point in the $(\delta u, \delta v)$ plane and Fourier transforming from this plane into the image plane, as per Eq. (36). We compare the outcome of simulating a large number $n_{\rm realis}$ of realisations and taking the variance across these realisations for each pixel with mapping the covariance matrix onto the $\delta u \delta v$-plane and using the Cov-Cov relationship.The results of our simulations are shown side-by-side in Fig. 4: the predicted and simulated noise PSFs match. The peak-normalised predicted noise PSF is less noisy, as shown in Fig. 3 for different correlation scales. This is expected, since it is calculated directly from the underlying distribution, rather than an estimate thereof. As $n_{\rm realis} \to \infty$, we expect the two methods to fully converge. As the maximum characteristic correlation length σ_τ increases, the variance becomes ever more sharply peaked.

Since our simulated sky consists of a 1Jy source at phase centre, there is only one noise PSF to modulate the average noise map, and it lies at phase centre. Let us test our formalism further by considering a model with multiple point sources.

2.3.2. Simulation with three point sources

We wish to test our prediction that the noise map can be described as a convolutional process modulating an average noise level. We thus perform another simulation, this time with three 1 Jy point sources. The associated dirty image is shown in Fig. 5d.

This dirty image simply consists of performing a direct Fourier transform (i.e. without using a Fast Fourier Transform algorithm) on simulated visibilities corresponding to these three point sources. We now perform a similar test as above on this field. Firstly, we "apply" gain errors to these visibilities by multiplying our model with our residual gain errors. This allows us to find 2000 realisations of residual visibilities, and find the variance for each pixel across these realisations. This gives us the simulated noise map, shown in Fig. 5a. Secondly, we perform a Direct Fourier Transform (DFT) from the differential Fourier plane to the (l, m) plane as before, assigning one cell of Cov$\{\tilde{\boldsymbol{\gamma}}\}$ to each point of the differential Fourier plane. This time, however, $\phi_{bb'}^d$ is not simply unity for all points in the differential Fourier plane. Instead, it is calculated for the three-point-source model, and applied for each point. This gives us the predicted noise map in Fig. 5b. Finally, Fig. 5c shows the absolute value of the difference between the two images. We see that there is some structure present in these residuals: this is expected, as the PSF of the sources in the dirty images clearly overlap. We are thus not quite in the regime where emission is fully spatially incoherent. Nevertheless, our predictions hold to better than 5% accuracy.

It bears repeating that, for correlated noise, this map can be understood as a distribution map for calibration artefacts: the amount of spurious correlated emission seen by each baseline will determine the noise map, and the true image-plane artefacts will then be one set of realisations of this underlying distribution.

3. Adaptive quality-based weighting schemes

As discussed in Sect. 1, some intervals of an observation will have lower gain variability. These will show up in the gain covariance matrix as intervals with lower variance. Similarly, those with larger intrinsic gain variability will have greater error in their gain estimate. By giving greater weights to the former, and lower weights to the latter, we expect to be able to improve image reconstruction. We thus talk of adaptive quality-based

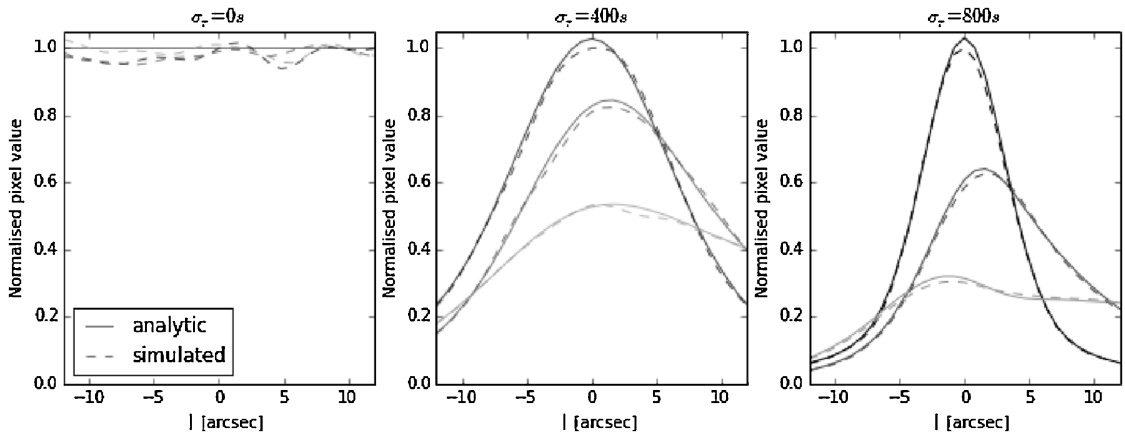

Fig. 3. Three lines in each figure corresponding to three horizontal cross-sections from images in Fig. 4. The units on the y-axis are dimensionless $[Jy^2/Jy^2]$. σ_τ is the maximum characteristic error correlation length. In decreasing intensity, they correspond to $m = 0''$, $m = 4''$, and $m = 8''$. The dashed lines correspond to the variance measured with 2000 realisations for each pixel, while the solid line corresponds to the predicted value at that pixel. There are 31 pixels. We do not show cross-sections for negative m due to image symmetry about the origin.

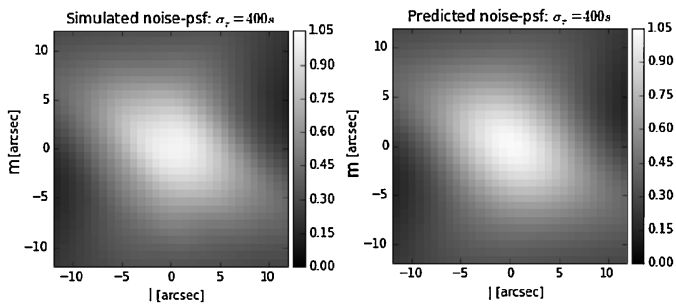

Fig. 4. Simulated noise maps, compared with theoretical prediction. The pixel values are normalised by the average value of the covariance matrix: the units of the colour bar are thus dimensionless (Jy^2/Jy^2). These are on the same angular scale as the source shown in Fig. 1.

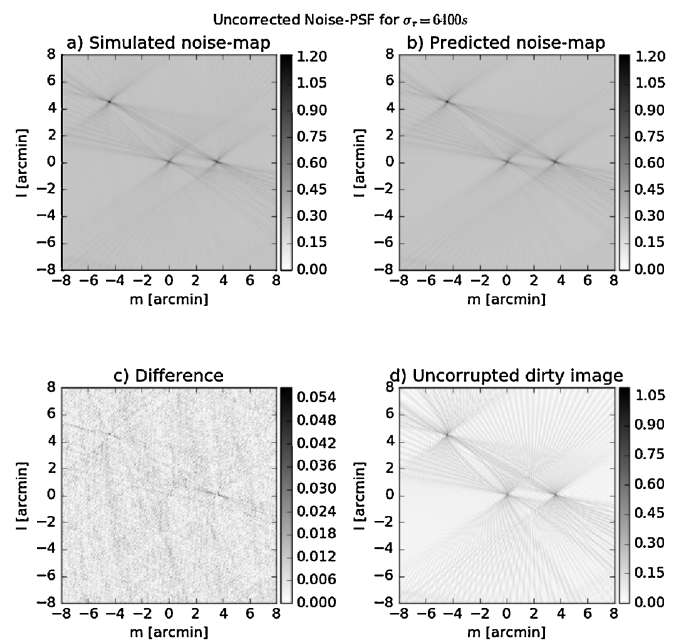

Fig. 5. Noise map of sky with correlated gain errors and three point sources. The colour bars of (*a*), (*b*), and (*c*) have dimensionless units, while that of (*d*) is in Jansky. We note the presence of structure in the residuals (*c*): these show the limits of our hypothesis that sources are spatially incoherent.

weighting, as the weights will adapt based on the calibration quality.

The pixel variance is determined by the visibility covariance matrix, as shown in Eq. (27). The diagonal of the visibility covariance matrix will add a flat noise to all pixels, while its wings will determine the calibration artefact distribution, which will be convolved to the sky brightness distribution. We thus have two sources of variance in the image plane. Minimising the far-field noise (i.e. the variance far from sources) in an image would involve down-weighting noisier calibration intervals while up-weighting the more quiescent ones, without taking noise correlation between visibilities into account. This is because the far-field noise will be dominated by the diagonal component of the covariance matrix (cf. Eq. (27)). By the same token, minimising calibration artefacts would involve down-weighting measurements with strongly-correlated noise, and up-weighting the less-correlated. This would not, however, minimise the diagonal component: in fact, it will likely exaggerate its up-weighting and down-weighting. As such, it will increase the constant level of the noise map, but flatten the noise-PSF's contribution. There are thus two competing types of noise that we seek to minimise: uncorrelated noise, which corresponds to $\delta u \delta v = 0$ (i.e. the diagonal components of the gain covariance matrix), and correlated noise, which corresponds to $\delta u \delta v \neq 0$ (i.e. its wings). Minimising the first will minimise far-field noise without optimally reducing artefacts, while minimising the last will minimise noise near sources at a cost to far-field noise. In the following sub-sections, we will discuss weighting schemes used to accomplish this.

3.1. Optimising sensitivity

The Cov-Cov relationship (Eq. (25)) tells us that, far from any sources, the variance map (Eq. (27)) is dominated by a constant term: the contribution from thermal noise and the diagonal of the residual visibility covariance matrix. Maximising sensitivity far from sources therefore implies minimising Diag $\{\mathrm{Cov}\{\tilde{\gamma}\}\}$. This is equivalent to assigning visibilities weights inversely proportional to their variance:

$$w_b = \frac{1}{\mathrm{Var}\{\tilde{\gamma}_b\}}. \tag{37}$$

For each baseline, those times with larger variance in the residuals will be down-weighted, and those with smaller variance will be up-weighted; this scheme does not require information about the underlying gains, only the error on our solutions. Since we are treating σ_n^2 as a constant for all antennas and all times, those times where our gains estimate is closer to the true gains will be up-weighted, and those moments where they are farther from the actual gains will be down-weighted: hence the term "adaptive quality-based weighting". We note that the diagonal of the weighted residuals' covariance matrix should therefore become constant: this weighting scheme explicitly brings the residuals closer to what is expected in the case of perfect calibration, assuming uncorrelated noise. For the remainder of this paper, we will refer to these weights as sensitivity-optimal weighting.

3.2. Minimising calibration artefacts

Minimising calibration artefacts - optimising the sensitivity near bright sources - means flattening the noise map. Since the noise map can be understood as a noise PSF convolved with all the modelled sources in the sky modulating the background variance level, it will be flattest when its peak is minimised. From the Cov-Cov relationship (Eq. (25)), we can see that, at the peak of the noise PSF (which would be the variance at the pixel where $l = m = 0$), the Fourier kernel is unity: the variance for that pixel is thus the sum of all the cells in the covariance matrix. By accounting for normalisation, we can write the variance at the centre of the noise-PSF as

$$V(\boldsymbol{w}) = \frac{\boldsymbol{w}^T \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}\boldsymbol{w}}{\boldsymbol{w}^T \mathbf{1}\mathbf{1}^T \boldsymbol{w}}. \tag{38}$$

Our optimality condition is then, after some algebra,

$$0 = \frac{\partial}{\partial \boldsymbol{w}}(V), \tag{39}$$

$$\leftrightarrow \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}\boldsymbol{w} = \mathbf{1}\mathbf{1}^T\boldsymbol{w}\left(\boldsymbol{w}^T\mathbf{1}\mathbf{1}^T\boldsymbol{w}\right)^{-1}\boldsymbol{w}^T\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}\boldsymbol{w}. \tag{40}$$

We find that one $\boldsymbol{w}$ which satisfies the above is

$$\boldsymbol{w} = \mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}^{-1}\mathbf{1}, \tag{41}$$

where $\mathbf{1}$ is a vector of ones. These weights depend only on calibration quality: badly-calibrated cells will include spurious time-correlated signal introduced by trying to fit the noise n on visibilities. Down-weighting these cells helps suppress artefacts in the field, at the cost of far-field sensitivity. This weighting scheme is thus only a function of the relative quality of calibration at different times, boosting better-calibrated visibilities and suppressing poorly-calibrated visibilities. For the remainder of this paper, we will refer to these weights as artefact-optimal weighting.

4. Estimating the covariance matrix

In our simulations, we have worked from a known covariance matrix and shown that our predictions for the residual image's behaviour hold. With real data, however, we do not have access to this underlying covariance matrix. Since our weights are extracted from said matrix, estimating it as accurately as possible remains a challenge; this is in turn limited by the number of samples that can be used for each cell.

Each cell in the covariance matrix is built by averaging a number of measurements or samples. The more samples available, the better our estimate becomes. Once we have more samples than degrees of freedom, we say that our estimation is well-conditioned. Otherwise, it is poorly-conditioned. In this section, we will discuss ways in which we can improve the conditioning of the covariance matrix estimation.

4.1. Baseline-based estimation

One way to improve the conditioning of our covariance matrix estimation is to make the same hypothesis as the calibration algorithm: we can treat the underlying gains as constant within each calibration interval. Provided this interval is known, this allows us to find a single estimate for each interval block of the covariance matrix, turning a $n_b \times n_b$ matrix into a smaller $n_{\mathrm{intervals}} \times n_{\mathrm{intervals}}$ equivalent, where $n_{\mathrm{intervals}}$ is the number of solution intervals used to find the gain solutions. We then improve our conditioning by a factor of n_{int}, which is the number of samples in a calibration interval. The estimate $\widehat{\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}}$of the covariance matrix $\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}$ is built by applying the covariance operator

$$\widehat{\mathrm{Cov}\{\tilde{\boldsymbol{\gamma}}\}} \stackrel{\mathrm{def}}{=} \widehat{\boldsymbol{C}}_{\tilde{\gamma}} = \frac{1}{n_{\mathrm{int}}}\sum_{i\in n_{\mathrm{int}}}(\tilde{\boldsymbol{\gamma}}_i - \langle\tilde{\boldsymbol{\gamma}}\rangle)(\tilde{\boldsymbol{\gamma}}_i - \langle\tilde{\boldsymbol{\gamma}}\rangle)^H, \tag{42}$$

where the $\langle\ldots\rangle$ operator denotes taking the average over the full vector. If the calibration solver's gain estimates are unbiased (i.e. $\mathrm{E}\{\hat{\mathbf{g}}\} = \mathbf{g}$) and the model of the sky is sufficiently complete, this quantity should be zero. Having created $\widehat{\boldsymbol{C}}_{\tilde{\gamma}}$, which will be of size $n_b \times n_b$, its cells can now be averaged over blocks of $n_{\mathrm{int}} \times n_{\mathrm{int}}$. This allows us to estimate the weights for each baseline and each time.

Mathematically, we retrace the steps of Sect. 2. In the absence of direction-dependent effects, we define the residual visibilities as before, and use them to define the normalised residual visibilities ρ_b:

$$r_b = w_b\kappa_b\tilde{\gamma}_b + \epsilon, \tag{43}$$

$$\rho_b = \frac{r_b}{k_b}. \tag{44}$$

We then organise the residuals in cells:

$$\boldsymbol{r}_C = \begin{pmatrix} \vdots \\ \rho_{b\in C} \\ \vdots \end{pmatrix}, \tag{45}$$

$$\boldsymbol{R} = \begin{pmatrix} \ldots & \boldsymbol{r}_C & \ldots \end{pmatrix}. \tag{46}$$

The matrix $\boldsymbol{R}$ contains all the residual visibilities within one calibration cell C, that is, for $b \in C$ where $\mathbf{g}_C = \mathrm{const}$. It is therefore of size $n_{\mathrm{intervals}} \times n_C$, where $n_{\mathrm{intervals}}$ is the number of calibration intervals in the observation. Normalising the residual visibilities by k_b allows us to recover the underlying covariance matrix by multiplying the residual visibility matrix $\boldsymbol{R}$ with its Hermitian conjugate,

$$\widehat{\boldsymbol{C}}_{\tilde{\gamma}}[b \in C, b' \in C'] = \left(\boldsymbol{R}^H\boldsymbol{R}\right)[C, C']. \tag{47}$$

We have divided the noise term by the flux model S_b, which can be very small in some cells. As such, care must be taken not to cause the relative thermal noise contribution to explode. Those cells where this would occur are dominated by thermal noise and

information on the covariance matrix cannot be recovered from them.

In this framework, we simply treat the index C as containing all the times and frequencies, for individual baselines, corresponding to a single calibration interval. The matrix $\widehat{\boldsymbol{C}}_{\tilde{\gamma}}$ is then an estimate of the residual visibility covariance matrix.

4.2. Antenna-based estimation

In the sub-section above, we assumed that finding one solution per interval would give us strong enough constraints to make the problem of estimating the covariance matrix well-conditioned: this may not be true in all cases. Conditioning may then need to be improved further: in this sub-section, we show one way in which this can be done. There are others, for example using the rank of the matrix itself to find better-conditioned estimates of the covariance matrix at a lower resolution (i.e. a single estimate for a greater number of cells). They will not be presented in this paper, but are a possible avenue for future work on this topic.

In estimating the covariance matrix for each baseline and each calibration cell, we are severely limited by the small number of samples in each cell. One way to overcome this problem is to find estimates for the variance of antenna gains and use these to return to the baseline variances. In this formalism, we extend C to include all visibilities pointing at a single antenna at a given time. Let us begin by writing an expression for the gain vector, which contains the gains for all antennas and all calibration cells,

$$\hat{\mathbf{g}}_c = \begin{pmatrix} \vdots \\ \hat{g}_p^{\tau\in c} \\ \vdots \end{pmatrix}, \tag{48}$$

$$\hat{\mathbf{G}} = \begin{pmatrix} \dots & \hat{\mathbf{g}}_c & \dots \end{pmatrix}, \tag{49}$$

and the variance on each antenna in each calibration cell is then

$$= \mathrm{E}\{\hat{\mathbf{g}}_c\hat{\mathbf{g}}_c^H\} - \underbrace{\mathrm{E}\{\hat{\mathbf{g}}_c\}\mathrm{E}\{\hat{\mathbf{g}}_c\}^H}_{=\mathbf{g}_c\mathbf{g}_c^H}. \tag{50}$$

As we can see, Eq. (50) is simply a vector form of Eq. (12). The residual gains of Eq. (12) can now be understood as random samples of the covariance between the gains for antennas p and q at a given time, assuming complete sky-model subtraction. We can thus define the variance sample matrix as an estimate of the variance matrix:

$$\widehat{\boldsymbol{V}_C} = \widehat{\mathrm{Var}\{\mathbf{g}_c\}} \tag{51}$$

$$= \sum_{\tau\in C}\left(\hat{\mathbf{g}}_\tau\hat{\mathbf{g}}_\tau^H - \mathbf{g}_\tau\mathbf{g}_\tau^H\right). \tag{52}$$

We define the residual matrix as

$$\boldsymbol{r}_\tau = \sum_d s_d\boldsymbol{K}_{d,\tau}\left(\hat{\mathbf{g}}_\tau\hat{\mathbf{g}}_\tau^H - \mathbf{g}_\tau\mathbf{g}_\tau^H\right)\boldsymbol{K}_{d,\tau}^H + \epsilon, \tag{53}$$

where we explicitly place ourselves in the limits of our formalism, that is, we do not have direction-dependent gains. We now see that at the core of Eq. (53) lies $\widehat{\boldsymbol{V}}_{\hat{\mathbf{g}}_\tau}$, where $\sum_{\tau\in C}\widehat{\boldsymbol{V}}_{\hat{\mathbf{g}}_\tau} = \widehat{\boldsymbol{V}_C}$. The K-matrix is defined as

$$\boldsymbol{K}_{d,\tau} = \begin{pmatrix} k_{p,\tau}^d & & 0 \\ & k_{q,\tau}^d & \\ 0 & & \ddots \end{pmatrix}. \tag{54}$$

Since the residual matrix depends on the gains, we define the residual visibility vectors as

$$\boldsymbol{r}_C = \begin{pmatrix} \vdots \\ \boldsymbol{r}_{\tau\in C} \\ \vdots \end{pmatrix}, \tag{55}$$

$$\boldsymbol{R} = \begin{pmatrix} \dots & \boldsymbol{r}_c & \dots \end{pmatrix}. \tag{56}$$

The matrix $\boldsymbol{r}_C$ contains all the residual visibilities within one calibration cell C, that is, for $\tau \in C$ where $\mathbf{g}_C = \text{const}$. Let us define n_C as the number of elements in each calibration cell. The residual variance sample matrix can now be built by multiplying the residual visibility matrix with its Hermitian conjugate:

$$\mathbb{V} = \boldsymbol{R}^H\boldsymbol{R}. \tag{57}$$

We do this because it allows us to turn a single noise realisation ϵ into a statistical quantity σ. We can relate $\mathbb{V}$ to the variance of individual antenna gains:

$$\mathbb{V} = \sum_{\tau\in C}\left(\sum_{d,d'} s_d\boldsymbol{K}_{d,\tau}\left(\widehat{\boldsymbol{V}_C}\right)^H\boldsymbol{K}_{d,\tau}^H s_{d'}\boldsymbol{K}_{d',\tau}\left(\widehat{\boldsymbol{V}_C}\right)\boldsymbol{K}_{d',\tau}^H + \mathbf{I}\sigma^2\right). \tag{58}$$

To reach this point, in Eq. (23), we made the hypothesis that the sky brightness distribution is dominated by spatially incoherent emission. Applying this hypothesis again here, we can make the approximation that the cross terms in the sum over d, d' average to zero: $\sum_{d,d'\neq d} \approx 0$. We then have:

$$\mathbb{V} \approx \sum_\tau\left(\sum_d s_d^2\boldsymbol{K}_{d,\tau}\left(\widehat{\boldsymbol{V}_C}\right)^H\left(\widehat{\boldsymbol{V}_C}\right)\boldsymbol{K}_{d,\tau}^H + \mathbf{I}\sigma^2\right) \tag{59}$$

$$= \left(\widehat{\boldsymbol{V}_C}\right)^2 \circ \left(\underbrace{\sum_\tau\sum_d s_d^2\boldsymbol{k}_{d,\tau}\boldsymbol{k}_{d,\tau}^H}_{\stackrel{\text{def}}{=}\mathcal{S}}\right) + n_C\mathbf{I}\sigma^2, \tag{60}$$

$$\widehat{\boldsymbol{V}_C} = \sqrt{\mathcal{S}^{\circ-1}\left(\mathbb{V} - n_C\mathbf{I}\sigma^2\right)}, \tag{61}$$

where $\circ$ denotes the Hadamard or entrywise product and $\boldsymbol{k} = \mathrm{Diag}\{\boldsymbol{K}\}$. Thus, $\mathbb{V}$ allows us to estimate the variance of each antenna and for each calibration cell by using all the visibilities pointing to that antenna within that calibration cell. With this information, we can then rebuild the baseline-dependent matrix, having improved our sampling by a factor of n_{ant}.

5. Applying the correction to simulated data

In this section, we show the impact of our weighting schemes on a noise map made from arbitrarily strongly-correlated residuals. Here, we assume that our sky contains only a single point source at phase centre: there is thus only a single instance of the noise PSF, placed at phase centre, to modulate the average variance level. We sample this instance by taking a cross-section from $(l, m) = 0$ to a large l, keeping m constant. The only difference between these cross-sections is the weighting scheme applied: unit weights for all visibilities ("Uncorrected", blue), sensitivity-optimal weights (green), and artefact-optimal weights (red). We plot both the result predicted by the Cov-Cov relationship (solid line) and the variance estimated across 2000

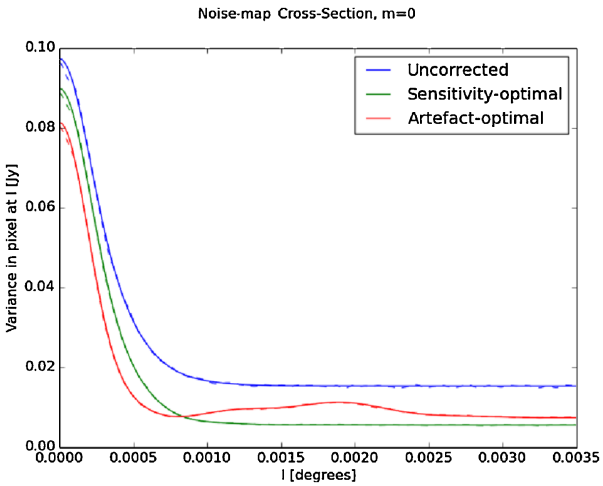

Fig. 6. Sensitivity-optimal (green) and artefact-optimal (red) weights both give improvements over the un-weighted noise map (blue).

realisations (dashed line): the result is shown in Fig. 6. The two remain in such agreement throughout the cross-section as to be nearly indistinguishable.

There are a few significant points to note on this figure. Firstly, most of the power in the matrix lies along the diagonal; both weighting schemes thus give good improvements in variance across the map. The artefact-optimal weights, while decreasing the peak further, as expected, also increase the noise far from sources. This is due to the fact that the artefact-optimal weights are in a sense more "selective" than the sensitivity-optimal weights; they up-weigh and down-weigh more severely, and will only result in a constant covariance matrix if this matrix is zero everywhere outside of the diagonal. In effect, the noise map becomes flatter, but much broader.

6. Applying the correction to real data

In this section, we show the effect of adaptive quality-based weighting on real data. The dataset used in this section is a single sub-band from an eight-hour LOFAR observation centred on the Extended Groth Strip (α = 14:19:17.84, δ = 52:49:26.49). The observation was performed on August 28, 2014. The sub-band includes eight channels of width 24.414 kHz each, for a total bandwidth ranging from 150.2 to 150.5 MHz. The data have been averaged in time to one data point per second. The data were calibrated using Wirtinger calibration (see Tasse 2014; Smirnov & Tasse 2015, and references therein) and a sky model consisting only of a nearby calibrator source, 3C295. A reference image (a cutout of which is shown in Fig. 7a) was made by calibrating the data according to best practice for LOFAR survey data: one calibration solution per eight seconds and per four channels. The residual data was then corrected by the gain solutions and imaged using Briggs weighting (robust = 0), pixel size of 1.5″, and deconvolved using the default deconvolution algorithm in DDFacet (see Tasse et al. 2018).

The data was then time-averaged to create a new, 2.4 GB dataset with one data point per eight seconds. Deliberately poor calibration was then performed on this dataset, solving for one calibration solution per two minutes (caeteris paribus). The resulting corrected residual data was imaged using the same imaging parameters as the reference image, and a cutout of the result is shown in Fig. 7b. As expected, the very long calibration intervals introduce calibration artefacts into the image. The brightest sources are still visible, but much of the fainter emission is buried under these artefacts. We then have the case where our residual visibilities are dominated by calibration error rather than sky model incompleteness.

Weights were then calculated based on the badly-calibrated residual visibilities. Figure 7c was made using the same visibilities as Fig. 7b and applying baseline-based, sensitivity-optimal weight. Similarly, Fig. 7d used the poorly-calibrated residual visibilities with the application of baseline-based, artefact-optimal weighting. These weights are likely to be poorly-conditioned. In both cases, all other parameters were conserved.

Applying antenna-based sensitivity-optimal weighting to the badly-calibrated data (not shown here) allows us to recover the reference image with only a very small increase in rms (increased by a factor of 1.14). Further testing on complex field simulations will be required to ascertain the usefulness of artefact-optimal weighting; it is likely that it fails to correct the image fully due to the poor conditioning of the covariance matrix used here.

The pixel histograms show us that the weights do not completely mitigate the poor calibration interval choice, but certainly give a dramatic improvement over the un-weighted, poorly-calibrated residuals. This is compatible with our statement that the weights give similar residuals in the image with a dramatic improvement in time at some cost to sensitivity. It is interesting to note that while Fig. 7d looks noisier than Fig. 7c, its residual flux histogram is actually closer to that of Fig. 7a.

As for performance, the weights used for Fig. 7d took eight hours of computing time on a single core[4], working on a 29 GB dataset, which is not particularly large for LOFAR data. Since the problem is massively parallel, this cost can be alleviated. The main bottleneck is likely due to very poor code optimization. As for the weights used for Fig. 8b, they are computed in 1min6s on the same single core.

7. Discussion

This paper began by investigating the use of an algorithm inspired by "lucky imaging" to improve images made using radio interferometric data. By investigating the statistics of residual visibilities, we have made the following findings:

- A relationship between the statistics of residual visibilities and residual pixel values (the "Cov-Cov relationship").
- A description of the noise map in the image plane as a constant variance level modulated by a noise PSF convolved with the sources in the field. This gives the variance in the flux of the image as a function of distance from the sources in the sky for a given calibration.
- An adaptive quality-based weighting scheme, which reduces the noise in the image (and the presence of calibration artefacts) by minimising either the constant noise term or the noise PSF.

While our results are not a panacea for poor calibration, they show that we can not only improve images made with well-calibrated data, but also mitigate the worst effects of poorly-calibrated visibilities in otherwise well-calibrated datasets. Provided that the gain variability timescale is long enough at certain points of the observation, we can effectively get images of similar quality using both the "standard" best-practice calibration interval for LOFAR survey data (calibration solution interval of eight seconds) and a significantly larger solution interval

[4] Core type: Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz.

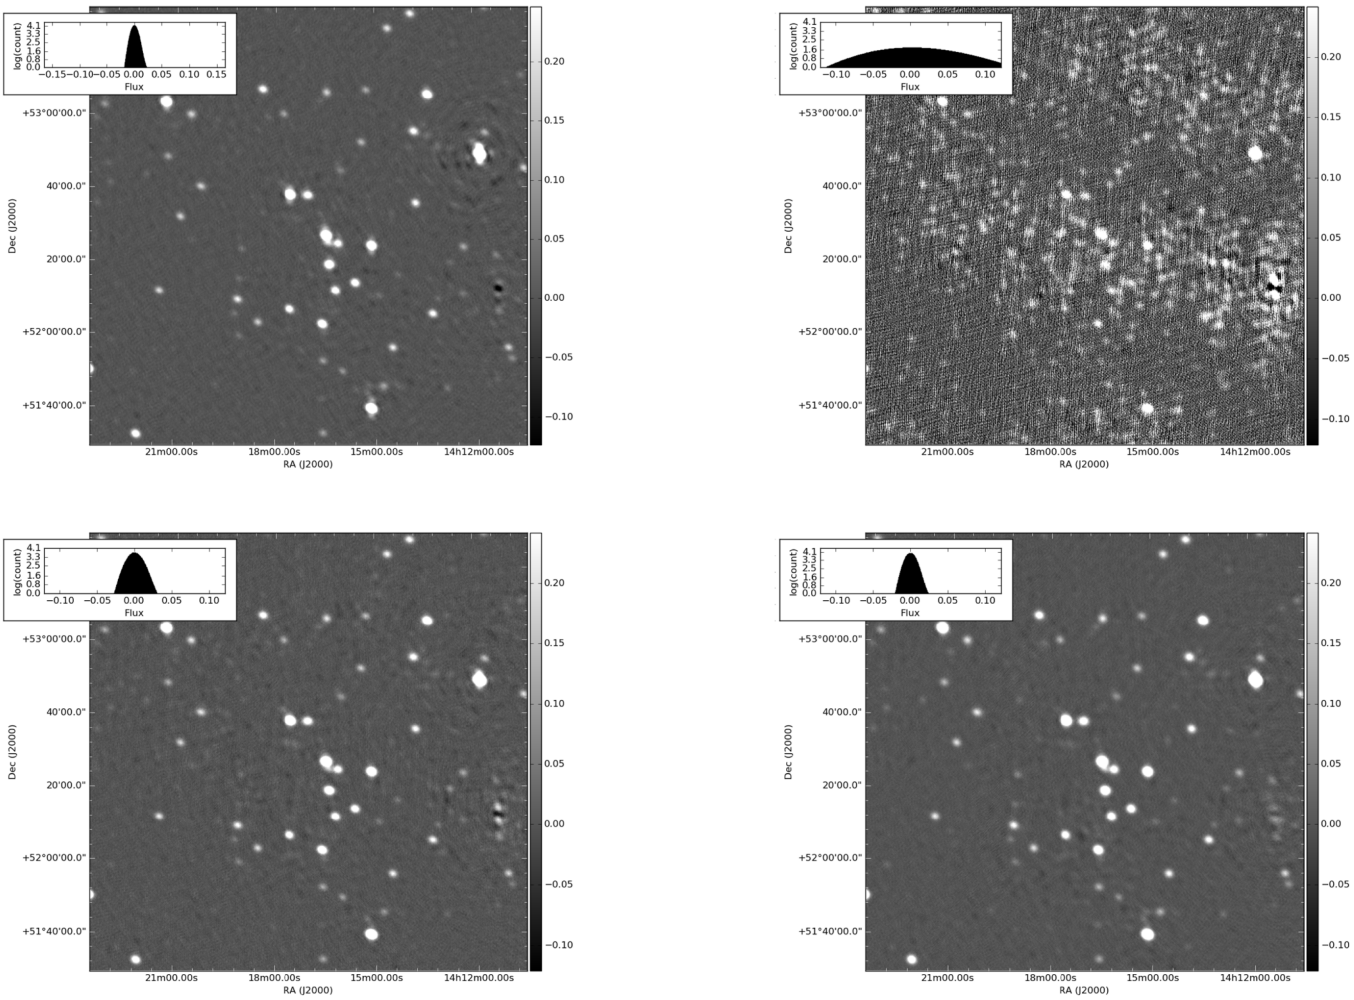

Fig. 7. Restored images of the centre of the Extended Groth Strip, as seen with an eight-hour observation using the full LOFAR array. *Panel a*: shows an image of the field made with good calibration intervals; well-calibrated, un-weighted restored image of the sky near the centre of the Extended Groth Strip. Used for comparison with the other images. Units of colour bar are Janskys. This image was made with data calibrated following best practice (solution intervals of 8 s, half the bandwidth). rms = 5.87 mJy beam^{-1}. *Panel b*: shows an image of the field made with poor calibration intervals; poorly-calibrated, un-weighted restored image of the sky near the centre of the Extended Groth Strip. Units of colour bar are Janskys. This image was made with the same data as for *panel a*, but averaged in time and calibrated using larger gain solution intervals: two minutes and half the bandwidth. rms = 86.4 mJy beam^{-1}. *Panel c*: shows image made with the same visibilities and imaging parameters, but with the application of the sensitivity-optimal weighting scheme; image made using the same imaging parameters and corrected visibilities as *panel b*, using sensitivity-optimal weighting. Units of colour bar are Janskys. rms = 9.69 mJy beam^{-1}. *Panel d*: similarly, differs from *panel c* only in that artefact-optimal weights, rather than sensitivity-optimal weights, were used; image made using the same imaging parameters and corrected visibilities as *panel b*, using artefact-optimal weighting. Units of colour bar are Janskys. rms = 15.8 mJy beam^{-1}. The histograms of pixel values in each image have 1000 flux bins ranging from –0.16 Jy to 0.16 Jy. Their ordinates are in log scale. Pixel size is 1.5″ in all images.

of two minutes (frequency interval unchanged). Of course, if no such stable interval exists, there will be no good intervals to up-weigh, and we will be left only with equally-poor data chunks. This means that, in the right conditions, net pipeline time can be sped up by a factor of nearly three, at a slight cost to sensitivity. This increase will be greater than what could be achieved with existing comparable methods such as "clipping".

We emphasise that the adaptive quality-based weighting schemes work because the noise map describes the underlying noise distribution, of which calibration artefacts are one single realisation. To fully characterise the artefacts, the correlation between different pixels (i.e. off-diagonal elements of $\mathrm{Cov}\{\tilde{\boldsymbol{y}}\}$) must be computed; this has not been done in this paper. Nevertheless, lesser constraints on the spatial distribution of artefacts can be found using only the diagonal elements of $\mathrm{Cov}\{\tilde{\boldsymbol{y}}\}$. The weighting schemes merely seek to minimise this spatial distribution as much as possible; the end result is fewer artefacts, which can be distributed across a much larger area. This is the source of the dramatic improvement from Figs. 7b and c. We have simply down-weighted those visibilities where spurious signal was introduced by the calibration solutions, and up-weighted those visibilities where such signal was lesser. Since this spurious signal is the source of calibration artefacts, down-weighting the associated visibilities reduces it dramatically. The poor improvement from Figs. 7b and d is likely due to limits in the conditioning of our estimation of the covariance matrix.

The work presented here can be improved upon, notably by working on improving the conditioning of our covariance matrix estimate; for real observations, it is impossible to have more than

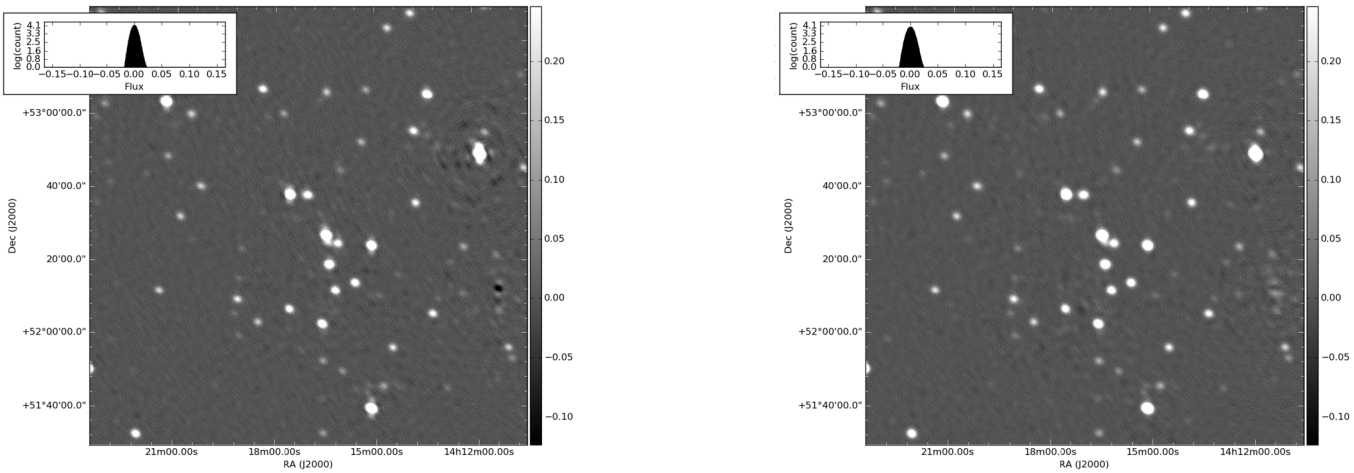

Fig. 8. Comparison between the well-calibrated image (i.e. the same image as Fig 7a) and antenna-based sensitivity-optimal weights. Units of both colour bars are Jansky. *Panel a*: well-calibrated, un-weighted restored image of the sky near the centre of the Extended Groth Strip. Used for comparison with the other images. Calibration solution intervals used were eight seconds, half the bandwidth. rms = 5.87 mJy beam^{-1}. *Panel b*: image made using the same imaging parameters and corrected visibilities as Fig. 7b, with the application of antenna-based, sensitivity-optimal weighting. Solution interval of two minutes, half the bandwidth. rms = 6.69 mJy beam^{-1}. We see that we recover a very similar image, despite the fact that the data used for the weighted image are averaged by a factor of 8 compared to those used for the un-weighted image.

one realization of each gain value for all antennas. By treating each visibility within a calibration interval as a realization of the true distribution, we can better estimate the covariance matrix per baseline, and thus reach a better estimate of the variance in the image plane. Of course, in practice, we can never access the true, underlying time-covariance matrix for each baseline. Significant hurdles remain:

- The impact of sky model incompleteness (since calibration requires a sky model) is ignored in this paper; we start by assuming that we have a complete sky model. In practice, of course, acquiring a complete sky model is often a key science goal in and of itself. The impact of this hypothesis therefore ought to be investigated in future work.
- The conditioning of our covariance matrix estimation remains a concern. By using an antenna-based approach, we can improve conditioning by a factor of $n_{\rm ant}$, but this is only one approach among many. Further work is needed to investigate which method, if any, proves optimal.

Acknowledgements. This work is based upon research supported by the South African Research Chairs Initiative of the Department of Science and Technology and National Research Foundation. The authors thank the LOFAR Surveys KSP Primary Investigators for allowing us to use their data for our tests and images in this paper. They would also like to thank M. Atemkeng and T. Gobler for fruitful discussions over the course of this work. They helped shape this work through their insightful comments on both style and substance while proofreading.

References

Briggs, D. S. 1995, in Amer. Astron. Soc. Meet. Abstr., BAAS, 27, 1444
Fried, D. L. 1978, J. Opt. Soc. Am. (1917-1983), 68, 1651
Grobler, T. L., Nunhokee, C. D., Smirnov, O. M., van Zyl, A. J., & de Bruyn A. G. 2014, MNRAS, 439, 4030
Hamaker, J. P., Bregman, J. D., & Sault, R. J. 1996, A&AS, 117, 137
Kazemi, S., & Yatawatta, S. 2013, MNRAS, 435, 597
Rau, U., & Cornwell, T. J. 2011, A&A, 532, A71
Smirnov, O. M. 2011, A&A, 527, A106
Smirnov, O. M., & Tasse, C. 2015, MNRAS, 449, 2668
Tasse, C. 2014, ArXiv e-prints [arXiv: 1410.8706]
Tasse, C., Hugo, B., Mirmont, M., et al. 2018, A&A, 611, A87
Yatawatta, S. 2014, MNRAS, 444, 790