

Calibration artefacts in radio interferometry.

I. Ghost sources in WSRT data

T. L. Grobler^{12*}, C. D. Nunhokee¹, O. M. Smirnov¹², A. J. van Zyl³, A. G. de Bruyn⁴⁵

¹*Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown, 6140, South Africa*

²*SKA South Africa, 3rd Floor, The Park, Park Road, Pinelands, 7405, South Africa*

³*Department of Mathematics and Applied Mathematics, University of Pretoria, Private bag X20, Hatfield, Pretoria 0028, South Africa*

⁴*ASTRON, P.O. Box 2, Dwingeloo, 7900AA, The Netherlands*

⁵*Kapteyn Astronomical Institute, University of Groningen, 9700 AV, Groningen, The Netherlands*

ABSTRACT

This work investigates a particular class of artefacts, or ghost sources, in radio interferometric images. Earlier observations with (and simulations of) the Westerbork Synthesis Radio Telescope (WSRT) suggested that these were due to calibration with incomplete sky models. A theoretical framework is derived that validates this suggestion, and provides predictions of ghost formation in a two-source scenario. The predictions are found to accurately match the result of simulations, and qualitatively reproduce the ghosts previously seen in observational data. The theory also provides explanations for many previously puzzling features of these artefacts (regular geometry, PSF-like sidelobes, seeming independence on model flux), and shows that the observed phenomenon of flux suppression affecting unmodelled sources is due to the same mechanism. We demonstrate that this ghost formation mechanism is a fundamental feature of calibration, and exhibits a particularly strong and localized signature due to array redundancy. To some extent this mechanism will affect all observations (including those with non-redundant arrays), though in most cases the ghosts remain hidden below the noise or masked by other instrumental artefacts. The implications of such errors on future deep observations are discussed.

Key words: Instrumentation: interferometers, Methods: analytical, Methods: numerical, Techniques: interferometric

1 INTRODUCTION

In the context of radio interferometry, the term *calibration* refers to the estimation and correction of instrumental errors (which are traditionally taken to also include effects of the troposphere and ionosphere) on the observed visibilities. Current calibration approaches boil down to a joint fit to the observations of a sky model and an instrumental model, such as that provided by the radio interferometer measurement equation (RIME; Hamaker et al. 1996; Smirnov 2011). A typical observing strategy will include intermittent *calibrator scans* of a known calibrator field, for which an accurate prior sky model is available; the obtained instrumental solutions can then be interpolated onto scans of the target field. These can be further refined through a process known as *self-calibration* or *selfcal* (Cornwell & Wilkinson 1981). Selfcal is an iterative approach (the sky model is refined at each iteration) that minimizes the error between predicted

visibilities corrupted by the instrumental model (the free parameters) and the observed visibilities in a least squares sense during each iteration. An initial sky model for selfcal can be obtained by imaging visibilities that have been corrected by the interpolated calibrator solutions. Where a reasonable initial sky model for the target field is available, it can even provide the starting point for selfcal without the need for an external calibrator.

Traditional selfcal assumes an instrumental model where all effects are direction-independent. The increased field of view of modern radio interferometers implies that direction dependent effects can no longer be ignored during calibration. Incorporating direction dependent effects into calibration solutions (third-generation calibration, or 3GC) has become a major research field over the past few years (Intema et al. 2009; Smirnov 2011a; Kazemi et al. 2011, 2013; Wijnholds & van der Veen 2009). van der Veen et al. (2004); Rau et al. (2009) have conducted good literature reviews on calibration.

* E-mail: t.grobler@ru.ac.za

It is well-established that calibration can lead to imper-

fect images, even to the generation of spurious source components, elimination or suppression of real components, and the deformation of the structure of extended sources (Linfield 1986; Wilkinson et al. 1988; Taylor et al. 1999; Martí-Vidal et al. 2010; Martí-Vidal & Marcaide 2008). Kazemi & Yatawatta (2013) have proposed a novel calibration technique that is meant to minimize the amount of source suppression that occurs if the sky model is incomplete. This approach uses a t -distribution to model the residual noise. Martí-Vidal & Marcaide (2008) has shown that spurious point sources can form when performing calibration on data containing nothing but white noise.

One of the more striking sightings of calibration artefacts occurred in 2004, in a 92cm Westerbork Synthesis Radio Telescope (WSRT) observation of J1819+3845 by de Bruyn (Fig. 1). After self-calibration, the map exhibited a string of “ghosts” – point-source-like objects, mostly of negative flux, arranged along a line linking the brightest object in the field with Cyg A, which was about 20° away (i.e. in a distant sidelobe of the primary beam, and extremely attenuated by WSRT’s extremely low sidelobe response). The pattern was highly peculiar as the ghosts were arranged with some regularity, and their positions did not vary with frequency. No other observations at the time were known to exhibit such features, and the problem remained open until a series of 21cm WSRT observations in 2010, which were done as part of the “Quality Monitoring Committee” project (Smirnov 2011b). In these observations, a large pointing error was deliberately introduced, and the resulting residual images (post self-cal) exhibited similar artefacts. The ghosts were fainter, but there were several strings of them, all associated with the brightest objects in the field (Fig. 2). The problem was then investigated empirically, through the use of simulations (Smirnov 2010), and this revealed a number of features:

- The ghosts were associated with sky model errors (i.e. missing or incorrect flux in the sky model, or direction-dependent errors towards the brightest sources). In the QMC case, this was due to the large pointing error; in the J1819+3845 case this was due to insufficiently accurate modelling of Cyg A. Correcting for these errors (by solving for differential gains towards the brighter sources in the QMC case, and towards Cyg A in the J1819+3845 case) made the ghosts disappear.

- A simple simulation of a two-source (1 Jy and 1 mJy) field, where only the 1 Jy source was included in the calibration model, while the second source played the role of “contaminator”, produced a similar ghost pattern in the residual visibilities. The peak intensity of the pattern was roughly at μ Jy level, and appeared to be proportional to the flux of the contaminator source (but did not depend on the flux of the model source!) This suggested that ghosts should **always** arise in the presence of incomplete sky models, but would generally be buried in the thermal noise, unless the observations were very sensitive, or the missing model sources were sufficiently bright.

- The ghosts always arranged themselves along a line (or lines) passing through the unmodelled or poorly modelled source(s), and the dominant source(s) in the sky model. The positions of the ghosts corresponded to some (but not all) rational fractions of the interval between the sources (i.e.

$1/2, 1/3, 2/3, 1/5$, etc.), with significant variation in intensity. The positions did not depend on frequency.

- The ghosts exhibited sidelobes that were similar, but not identical to, the PSF of the telescope.
- Similar simulations with other telescopes (VLA) showed a far less regular artefact pattern.

These facts strongly suggested that the regularity of the ghost patterns in Figs. 1 and 2 was somehow related to the highly redundant geometry of the WSRT, but the mechanism by which they arose was not clear. Another conclusion of that study was that ghosts could be minimized and eventually driven below the noise by laboriously improving the sky model and/or applying direction-dependent solutions.

The problem of calibration artefacts is becoming more important with the advent of new observational techniques and new radio telescopes such as the Low Frequency Array (LOFAR), the upgraded Jansky Very Large Array (JVLA), MeerKAT, etc., as well as the upcoming Square Kilometer Array (SKA). Not only do these telescopes promise (and in some cases already deliver) unprecedented sensitivity, they also increase the data rates substantially. The increased sensitivity means that fainter artefacts cannot be ignored, while the data rates require that calibration become largely automated, with careful and laborious manual data reduction no longer an option. Observations of the diffuse HI in the $0.5 < z < 20$ in order to probe galaxy formation (e.g. Furlanetto et al. 2006) or the nature of dark energy (Chang et al. 2010) need to face exquisite calibration in order to subtract foreground that are spatially and spectrally orders of magnitude brighter (i.e. Bernardi et al. 2009; Switzer et al. 2013; Pober et al. 2013). Similar calibration requirements are needed for future stacking techniques for HI detection (e.g. Delhaize et al. 2013) and continuum surveys aimed at revealing the nJy population (Norris et al. 2013). On the other hand, transient source detection pipelines require very fast on-the-fly calibration, which necessitates the use of shallow and incomplete sky models (Wijnholds 2013, priv. comm.) It is clear that a deeper and theoretical understanding of calibration artefacts, including those buried in the noise, is required if these new instruments and techniques are to achieve their scientific goals.

We have been investigating two classes of calibration artefacts, namely spurious point sources (ghosts) and source suppression (i.e. the reduction in observed flux of sources not included in the sky model). This work establishes that both are manifestation of the same underlying mechanism, and aims to provide a theoretical understanding of this. The objective is to extend the results of previous papers (Linfield 1986; Wilkinson et al. 1988; Martí-Vidal & Marcaide 2008; Taylor et al. 1999), by studying the underlying theoretical principles which are responsible for ghost formation and source suppression (as opposed to conducting an empirical survey of the different types of calibration artefacts that have been identified).

This paper is the first in a series; ongoing work is concentrating on ghost formation and source suppression in non-redundant interferometers, and on direction-dependent calibration. Ghosts have already been spotted in LOFAR data (de Bruyn priv. comm.), and particularly in the presence of transient sources (Fender priv. comm.), and early




Figure 1. Ghost sources in a 92cm WSRT observation of J1819+3845. The target field is in the lower-image corner of the image, and Cyg A is to the upper left, just outside the image.




Figure 2. Ghost sources in a 21cm WSRT observation of the QMC2 field. Note that this is a residual dirty map, i.e. the sky model sources have been subtracted, and the visibly brighter sources here are in fact relatively faint. The positions of the brightest subtracted sources are indicated by red circles. Multiple strings of ghosts are visible, here highlighted by the ellipses. Note how the strings are firmly associated with the brightest sources.

indications are that the same mechanism is responsible. Future papers in the series will cover these phenomena.

2 PROBLEM OVERVIEW AND DEFINITIONS

In this paper, we will concentrate on the WSRT example, since its highly redundant East-West geometry makes for prominent and regular ghosts. The results can be extended to other telescopes, which the follow-up work on source sup-

pression will also deal with. We make a number of further simplifications:

- We consider a case where the true sky consists of two discrete point sources with fluxes A_1 and A_2 , the former at the phase centre, and the calibration model consists of just the central source A_1 .
- The sources are unpolarized, and we consider only a single frequency channel.
- Only direction-independent calibration (i.e. regular self-cal) is performed.

Multiple sources, multiple frequencies, polarization and studying the effects of direction-dependent solutions will be the subject of future work.

This section will anticipate some results of the following sections, in order to provide a logical outline that the rest of the paper will fill in.

2.1 Calibration

In its general form, (unpolarized) calibration entails finding a diagonal antenna gain matrix $\mathbf{G} = \text{diag}(\mathbf{g}) = \text{diag}([g_1, g_2, \dots, g_n]^T)$ such that

$$\|\mathcal{R} - \mathbf{G}\mathcal{M}\mathbf{G}^H\| \quad (1)$$

is minimized at each observational time-step. The superscript notation $()^H$ in Eq. 1 denotes the Hermitian transpose. The Hermitian matrix \mathcal{R} is the observed unpolarized visibility matrix, where element r_{pq} of \mathcal{R} is the visibility measured by the baseline formed by antennas p and q . The matrix \mathcal{M} is the corresponding visibility matrix generated from the calibration sky model. Eq. 1 can then be restated as

$$\|\mathcal{R} - \mathbf{g}\mathbf{g}^H \odot \mathcal{M}\| = \|\mathcal{R} - \mathcal{G} \odot \mathcal{M}\|, \quad (2)$$

where “ \odot ” represents element-by-element multiplication (Hadamard product), and $\mathcal{G} = \mathbf{g}\mathbf{g}^H$ is the matrix product of the gain solution vector with its own Hermitian transpose. The elements of \mathcal{G} will be denoted by g_{pq} . Crucially, \mathcal{G} is a rank one matrix by construction, and conversely, any rank one matrix can be decomposed into a product of the form $\mathbf{g}\mathbf{g}^H$.

Conventional approaches to radio interferometric calibration ignore the autocorrelations (i.e. the diagonal of the visibility matrix), since these are subject to a high additional self-noise term, and employ non-linear optimization techniques such as Levenberg-Marquardt (Levenberg 1944; Marquardt 1963) to find a maximum likelihood (ML) solution for the off-diagonal terms of Eq. 1. When a Gaussian noise model is assumed, this becomes equivalent to least squares (LS) minimization:

$$\min_{\mathbf{g}} \sum_{p \neq q} (r_{pq} - g_p m_{pq} \bar{g}_q)^2, \quad (3)$$

where \bar{g}_q denotes the complex conjugate of g_q .

We will use the term LS calibration to refer to an LS solution of Eq. 3. Most reduction packages in current use perform some sort of LS calibration. Kazemi & Yatawatta (2013) propose an alternative approach called *robust calibration*, where a ML solution is obtained under the assumption of a t -distribution for the noise. Since implementations of

robust calibration are not yet publicly available, we do not study it in this work. Where required, we make use of the MeqTrees package (Noordam & Smirnov 2010) to do LS calibration.

If autocorrelations are included in the optimization problem, an approach called Alternating Least Squares gain estimation (ALS; Boonstra & van der Veen 2003; Wijnholds & van der Veen 2009) can be used to obtain a solution to Eq. 2. It is not clear whether ALS provides a practical advantage over traditional LS without autocorrelations, since implementations of ALS compatible with conventional radio interferometric data do not exist. The issue deserves to be investigated in a separate study. For our purposes, ALS turns out to provide a vital theoretical framework in which ghost formation can be understood analytically.

2.2 Ghost formation

In a nutshell, ghost sources are produced when Eq. 1 is solved for with an incomplete or incorrect model \mathcal{M} . Consider the simple case where the observed visibilities \mathcal{R} correspond to two point sources, and the calibration model consists of a single point source at centre, $\mathcal{M} = \mathbf{1}$, where $\mathbf{1}$ (boldface 1) represents a matrix of all ones (not a unity matrix!) If we then assume a perfect instrument with unity gains, the actual solutions for \mathbf{G} will not be quite equal to unity, as they will attempt to fit for the difference between \mathcal{M} and \mathcal{R} . Qualitatively, this process can be understood as follows: calibration attempts to move some “real flux” from the model \mathcal{M} to compensate for the unmodelled flux of the second point source. When these solutions are applied to the data, the resulting corrected visibilities

$$\mathcal{R}^{(c)} = \mathbf{G}^{-1} \mathcal{R} \mathbf{G}^{-H}, \quad (4)$$

will contain ghost sources in addition to real sources. In Eq. 4, $()^{-1}$ denotes standard matrix inversion, while $()^{-H}$ designates $(())^H)^{-1}$. The rest of this paper analyses the mechanism by which this comes about. Note that we do not consider the effects of noise in our analysis; earlier empirical work (Smirnov 2010) has shown that the same pattern arises with or without noise.

In this context, LS calibration has proven to be very difficult to study theoretically. By contrast, the ALS formulation does yield the necessary insights. In this paper we therefore approach the problem of ghosts from several directions:

- We develop a theoretical framework based on ALS that predicts ghost formation;
- We empirically compare the results of ALS and LS calibration, and show that they yield similar ghost patterns (with minor differences that are explained);
- We provide empirical results for ghost formation using ALS and LS, and show that these match the theoretical predictions;
- We show that all of the above match observed ghost patterns in real data, such as those seen in Fig. 1.

These results suggest that the theoretical insights gained from the ALS framework are valid for the LS approaches, while the last point demonstrates that our simplified assumptions provide a good fit to real observations.

2.3 Distillation

Since ghost sources are relatively faint (as we'll show below), they can be difficult to detect over the thermal noise and the PSF sidelobes of actual sources. In hindsight, this probably explains why the phenomena was not spotted earlier. A straightforward way to detect the ghosts in simulations is to “distill” them into residual visibilities as follows:

- (i) We form predicted visibilities from a “true” sky (\mathcal{M}_0), and an incomplete calibration sky model (\mathcal{M}).
- (ii) The “observed visibilities” \mathcal{R} then correspond to \mathcal{M}_0 .
- (iii) We obtain calibration solutions \mathcal{G} by solving Eq. 1 using \mathcal{R} and \mathcal{M} .
- (iv) We apply the solutions to \mathcal{R} (Eq. 4), yielding corrected visibilities $\mathcal{R}^{(c)}$.
- (v) We image the residuals $\mathcal{R}^\Delta = \mathcal{R}^{(c)} - \mathcal{R}$. The real sources then (mostly) cancel out, the noise term, if any, also (mostly) cancels out, and the resulting image yields the “distilled” ghost sources.

Note that in real-life observations, actual gains are never unity, and the residuals $\mathcal{R}^{(c)} - \mathcal{R}$ would not reveal much since real sources would not cancel out. However, in our perfect telescope simulation, the gain solutions account for sky model incompleteness and nothing more, and the ghosts are easily visible in images of the residuals.

In the two-source, noise-free case considered here, the true sky \mathcal{M}_0 is equal to

$$\mathcal{R} = \mathcal{M}_0 = A_1 \mathbf{1} + A_2 \mathbf{K},$$

where \mathbf{K} is a Fourier kernel matrix of complex phase terms corresponding to the offset of the second source w.r.t. the phase centre. The residuals then correspond to

$$\mathcal{R}^\Delta = A_1 \mathbf{G}^{-1} \mathbf{1} \mathbf{G}^{-H} + A_2 \mathbf{G}^{-1} \mathbf{K} \mathbf{G}^{-H} - A_1 \mathbf{1} - A_2 \mathbf{K}$$

By defining the matrix $\mathcal{G}^\top = \mathbf{G}^{-1} \mathbf{1} \mathbf{G}^{-H} = \{g_{pq}^{-1}\} = \left\{ \frac{1}{g_{pq}} \right\}$ (i.e. the element-by-element inverse or the Hadamard inverse of \mathcal{G}), we can rewrite this as

$$\mathcal{R}^\Delta = A_1 (\mathcal{G}^\top - \mathbf{1}) + A_2 (\mathcal{G}^\top - \mathbf{1}) \odot \mathbf{K}. \quad (5)$$

The matrix $\mathcal{G}^\top - \mathbf{1}$ is in some sense fundamental. As will be shown below, it yields the basic ghost pattern corresponding to one source. From the equation above, we can see that the residuals will contain a superposition of two ghost patterns, scaled by A_1 and A_2 , with the second pattern shifted to the position of the second source. In the general case, the residuals will correspond to a convolution of the true sky with $\mathcal{G}^\top - \mathbf{1}$. Since in practice $A_2 \ll A_1$ (i.e. the missing flux in the model is usually considerably less than the flux accounted for), the first realization of the pattern is dominant (moreover, as will also be shown below, in the WSRT case the positions of the ghosts in the two patterns fall on top of one another). We shall refer to $\mathcal{G}^\top - \mathbf{1}$ as the *distilled ghost pattern*.

3 THEORETICAL DERIVATION

In this section analytic expressions for the elements of \mathcal{G}^\top are derived. In the image domain each element of \mathcal{G} represents a different ghost pattern. The ghost patterns that

are associated with \mathcal{G} form due to a loss of information. Since \mathcal{G} is of a lower rank than \mathcal{R} (assuming a two source sky and a single source in the model) some information is lost when \mathcal{G} is computed. The inadequate rank one model \mathcal{G} leads to a significant change in the Fourier characteristics of the original matrix \mathcal{R} . The change in Fourier characteristics manifest as ghost patterns when \mathcal{G} is imaged. When \mathcal{G}^\top is calculated the ghost patterns remain the same (the fluxes of the sources do however change). When the antenna gain solutions are applied to \mathcal{R} the ghost patterns of \mathcal{G}^\top get convolved with the true sky, which implies that $\mathcal{R}^{(c)}$ will contain ghost sources.

A brief introduction to the Appendix is given in Sect. 3.1, since it is crucial to the theory in this section. The mathematical definition of a regularly-spaced array is given in Sect. 3.2, while Sect. 3.3 gives a better description of the experimental test case that is considered. Analytic expressions for the elements of \mathcal{G}^\top are then derived in Sect. 3.5. Section 3.6 describes how this results in ghost patterns in the dirty images, while Sect. 3.7 analyzes the effect of \mathcal{G}^\top on the corrected visibilities.

The derivations in this section are highly mathematical; the crucial result is Eq. 13, which shows that the calibrated visibilities on each baseline, in case of a two-source sky and one-source model, are sampled from a periodic one-dimensional uv -distribution, which in turn corresponds to a string of delta functions in the image plane. The reader wishing to skip the heavier mathematics is encouraged to take Eq. 13 at face value, and skip to Sect. 3.6, which explains how the “strings” corresponding to each baseline combine to form ghosts in the final image.

3.1 Introduction to Appendix

The brief introductory explanation from above will be expanded upon in the rest of Sect. 3. The Appendix will be one of the main tools we will use to accomplish this. A brief introduction to the Appendix is therefore needed. The appendix contains lemmas and propositions. The propositions are the main results that are used to derive the theoretical results in this section. The lemmas are the dependencies that are required by these propositions. The relation between the lemmas and propositions are discussed in greater detail in the Appendix itself. The Appendix proves certain properties of $\mathcal{R}(\mathbf{b})$, $\mathcal{G}(\mathbf{b})$ and $\mathcal{G}^\top(\mathbf{b})$, which are the extrapolated counterparts of \mathcal{R} , \mathcal{G} and \mathcal{G}^\top (see Definition 1.3 and Definition 1.4). These properties turn out to be essential in deriving the distilled ghost pattern. The following propositions are proven in the Appendix:

(i) Proposition 1.5, the rank of $\mathcal{R}(\mathbf{b})$ is rank two. This proposition quantifies the amount of information that is being lost during the computation of $\mathcal{G}(\mathbf{b})$.

(ii) Proposition 1.6, the elements of the function-valued matrix $\mathcal{G}(\mathbf{b})$ are periodic, effectively one-dimensional, differentiable, Hermitian functions.

(iii) Proposition 1.7, it follows from Proposition 1.6 that the elements of $\mathcal{G}(\mathbf{b})$ can be written as an effectively one-dimensional Fourier-series (which ultimately leads to the formation of ghosts).

(iv) Proposition 1.8, the elements of $\mathcal{G}^\top(\mathbf{b})$ are also periodic, effectively-one dimensional, differentiable, Hermitian

functions and therefore by Proposition 1.7 can also be expressed as an effectively one-dimensional Fourier-series.

3.2 Regular and redundant array geometries

Since the geometric regularity of the WSRT layout will turn out to have an important effect on ghost formation, let's provide a formal mathematical definition here.

Definition 1.1. (Regularly-spaced array) Let us pick a coordinate system with origin at the first antenna position $\mathbf{u}_1 = 0$. We shall call a set of antenna positions $\{\mathbf{u}_p\}$ *regularly-spaced* if there exists a *common quotient baseline* (CQB) \mathbf{b}_0 such that each antenna position is an integer multiple of \mathbf{b}_0 , i.e. that $\mathbf{u}_p = \phi_p \mathbf{b}_0$, with ϕ_p being a whole number. We will also require that \mathbf{b}_0 is the largest such baseline (equivalently, the greatest common divisor of $\{\phi_p\}$ is 1).

Definition 1.2. (Array geometry matrix) The array geometry matrix Φ is an $n \times n$ integer matrix with elements $\phi_{pq} = \phi_q - \phi_p$.

Obviously, a regularly-spaced array defined in this way is necessarily one-dimensional. Note that regularity is preserved under rotation (but only in an East-West array). Note also that \mathbf{b}_0 does not necessarily correspond to a real baseline. Most commonly-used configurations of the WSRT are regularly-spaced: the 10 fixed antennas have a CQB of 144m, while the CQB of the array as a whole is determined by the positions of the movable antennas RTA to RTD, with typical CQB lengths of 6 or 12m¹. A *redundant* array will have many identical entries in Φ . A regularly-spaced array is not necessarily redundant, but WSRT itself is highly redundant.

The matrix Φ has a few interesting mathematical properties, which will be fully derived in the Appendix. Note that the actual uv -coordinates of each baseline are given by $\mathbf{b}_0 \Phi$. The matrix Φ can be thought of as representing a whole number scaling relationship between all the uv -tracks of the interferometer, and the *reference track* given by the CQB $\mathbf{b}_0(t)$, which is a function of time due to the Earth's rotation.

3.3 The two source problem

Let us consider a sky composed of two unpolarized point sources of flux A_1 and A_2 , and a calibration sky model consisting of just the primary source A_1 . Since the solutions to the calibration equation (Eq. 1) are invariant with respect to amplitude rescaling and positional shifts that are applied to both the sky and the model, we may, without loss of generality, restrict ourselves to the case where the primary source has unity flux and is located at the phase centre. The “true sky” as a function of position $\mathbf{s} = (l, m)$ (where l and m are the direction cosines) is then equal to $I_{\mathcal{R}}(\mathbf{s}) = A_1 \delta(\mathbf{s}) + A_2 \delta(\mathbf{s} - \mathbf{s}_0)$, and the “model sky” to $I_{\mathcal{M}}(\mathbf{s}) = A_1 \delta(\mathbf{s})$, where $A_1 = 1$, $\mathbf{s}_0 = (l_0, m_0) \neq 0$ is the position of the secondary source, and δ is the Kronecker delta function. Let us further assume a perfect interferometer with

¹ Since the WSRT movable antennas can in principle be placed at any position along a continuum, a non-regularly-spaced configuration is technically possible, but never used in practice.

unity gains, a monochromatic observation, and integration intervals sufficiently short to make smearing negligible. The “observed visibility” corresponding to the true sky $I_{\mathcal{R}}$ can be written as

$$r(\mathbf{u}) = A_1 + A_2 e^{-2\pi i \mathbf{u} \cdot \mathbf{s}_0}, \quad (6)$$

where $\mathbf{u} \cdot \mathbf{s}_0$ is a dot product. If the array is regularly-spaced as defined above, then the visibility observed by baseline pq at uv -coordinates $\mathbf{u}_{pq} = \phi_{pq} \mathbf{b}_0$ is

$$V_{pq} = r(\mathbf{u}_{pq}) = r(\phi_{pq} \mathbf{b}_0) = A_1 + A_2 e^{-2\pi i \phi_{pq} \mathbf{b}_0 \cdot \mathbf{s}_0}, \quad (7)$$

where $\mathbf{b}_0 = \mathbf{b}_0(t)$ is the CQB. The “model visibilities” corresponding to the model sky $I_{\mathcal{M}}$ above, are trivially all unity.

3.4 The extrapolated visibility matrix

The observed visibilities for each observational time step can be packed into a two dimensional matrix

$$\mathcal{R} = \begin{bmatrix} V_{11} & V_{12} & \cdots & V_{1n} \\ V_{21} & V_{22} & \cdots & V_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ V_{n1} & V_{n2} & \cdots & V_{nn} \end{bmatrix}. \quad (8)$$

The elements of \mathcal{R} are functions that depend on time. For a regularly-spaced array, Eq. 7 can be utilized to rewrite all the elements of \mathcal{R} as functions of \mathbf{b}_0 . We can express this formally via the following definition:

Definition 1.3. (Extrapolated visibility matrix) Let $\mathcal{R}(\mathbf{b}) : \mathbf{R}^2 \rightarrow \mathbf{C}^{n \times n}$ be an $n \times n$ Hermitian function-valued matrix with entries

$$r_{pq}(\mathbf{b}) = r(\phi_{pq} \mathbf{b}), \quad (9)$$

where r is given by Eq. 6, ϕ_{pq} is given by the array geometry matrix Φ , $\mathbf{b} = (u, v)$ and $\mathbf{s}_0 = (l_0, m_0) \neq 0$ are real two-vectors, $A_1 = 1$, and $0 < A_2 < 1$.

This allows us to formally define $\mathcal{R}(\mathbf{b})$ over the entire uv -plane, i.e. for any value \mathbf{b} . The actual observed visibilities \mathcal{R} at time t are given by $\mathcal{R}(\mathbf{b}_0(t))$. For any given baseline pq , $r_{pq}(\mathbf{b}_0)$ corresponds to the visibilities measured by that baseline. Since $\mathbf{b}_0(t)$ follows an elliptical track, our actual “measurements” on baseline pq (i.e. the values subject to calibration) are restricted to that series of uv -points. However, by replacing \mathbf{b}_0 by the free variable \mathbf{b} in Eq. 9, we automatically define an “extrapolated” visibility function over the entire uv -plane. By definition, the values of r_{pq} over the track $\mathbf{b}_0(t)$ are equal to visibilities measured by baseline pq over the track $\phi_{pq} \mathbf{b}_0(t)$.

Note also that Eq. 9 can also be seen as a coordinate scaling relationship between the observed visibility distribution $r(\mathbf{u})$ and any given $r_{pq}(\mathbf{b})$. To emphasize this, we use the variable \mathbf{u} to represent coordinates in the “observed” uv -plane (where r lives), and \mathbf{b} for coordinates in the “scaled” uv -planes (where the r_{pq} 's live). This also implies that the “sky” corresponding to any r_{pq} (i.e. the inverse Fourier transform of r_{pq}) is a scaled and stretched version of the true sky.

Finally and most crucially (as we'll see in the discussion of ALS below), the $\mathcal{R}(\mathbf{b})$ matrix for any $\mathbf{b} \neq 0$ can be shown to have rank two (see Proposition 1.5 in the Appendix).




Figure 3. The functions $r_{12}(\mathbf{b})$ and $g_{12}(\mathbf{b})$.

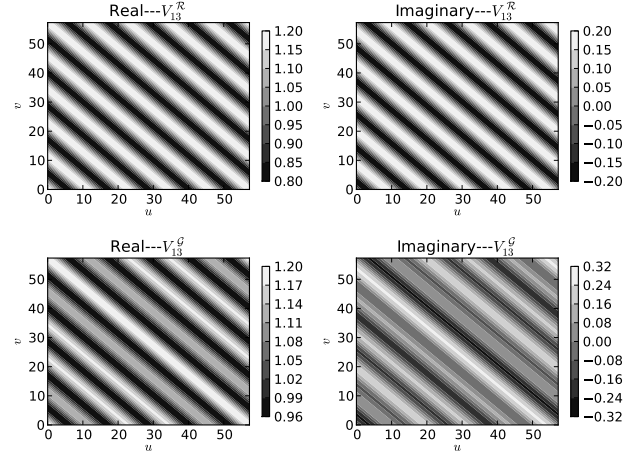


Figure 4. The functions $r_{13}(\mathbf{b})$ and $g_{13}(\mathbf{b})$.

3.5 The calibration matrix

Since our model visibilities are all unity, the calibration process (Eq. 2) entails finding some kind of “best fit” rank one matrix \mathcal{G} , given \mathcal{R} . In effect, the calibration process results in a mapping $\mathcal{R} \rightarrow \mathcal{G}$; by extension, this also defines a mapping $\mathcal{R}(\mathbf{b}) \rightarrow \mathcal{G}(\mathbf{b})$ for any \mathbf{b} . For LS calibration, the best fit is given by Eq. 3. This has proven difficult to explore analytically, so we will consider ALS calibration instead (and later empirically show that it yields similar results).

In a nutshell, ALS calibration obtains a \mathcal{G} by “de-ranking” \mathcal{R} , i.e. keeping just its largest eigenvalue. More precisely:

Definition 1.4. (ALS calibration matrix) Let $\mathcal{G}(\mathbf{b}) = \lambda(\mathbf{b})\mathbf{x}(\mathbf{b})\mathbf{x}^H(\mathbf{b})$, where $\lambda(\mathbf{b})$ is the largest eigenvalue of $\mathcal{R}(\mathbf{b})$, and $\mathbf{x}(\mathbf{b})$ is its associated normalized eigenvector.

We will designate the elements of \mathcal{G} as $g_{pq}(\mathbf{b})$.

To provide a specific example of the above, let us create a theoretical three-element interferometer with a geometry matrix of

$$\Phi = \begin{bmatrix} 0 & 3 & 5 \\ -3 & 0 & 2 \\ -5 & -2 & 0 \end{bmatrix}, \quad (10)$$

and place a secondary source of flux $A_2 = 0.2\text{Jy}$ at $l_0 = 1^\circ, m_0 = 1^\circ$. Fig. 3, Fig. 4 and Fig. 5 graphically display the resulting $\mathcal{R}(\mathbf{b})$ and $\mathcal{G}(\mathbf{b})$ matrices. The following observations can also be made. The functions $r_{pq}(\mathbf{b})$ (top row of each figure) are trivial phase gradients, and are thus continuous, differentiable, Hermitian and periodic in the u and v direction with periods of $\frac{1}{\phi_{pq}|l_0|}$ and $\frac{1}{\phi_{pq}|m_0|}$. They are effectively one-dimensional, i.e. constant along each line $v = -\frac{l_0}{m_0}u + c$ for any c . The $\mathcal{G}(\mathbf{b})$ functions (bottom row of each figure) have a more interesting structure, but are also differentiable, Hermitian, one-dimensional and periodic, with periods along u and v of $\frac{1}{|l_0|}$ and $\frac{1}{|m_0|}$. Moreover, this holds for any ALS calibration matrix as defined above (see Proposition 1.6). The difference between the g_{pq} ’s is that each of them has a different secondary harmonic which is determined by ϕ_{12}, ϕ_{13} and ϕ_{23} .




Figure 5. The functions $r_{23}(\mathbf{b})$ and $g_{23}(\mathbf{b})$.

Any periodic, one-dimensional and differentiable function can be written out in terms of a one-dimensional discrete Fourier transform. We can therefore decompose each g_{pq} as follows:

$$g_{pq}(\mathbf{b}) = \sum_{j=-\infty}^{\infty} c_{j,pq}^{\mathcal{G}} e^{-2\pi i j \mathbf{b} \cdot \mathbf{s}_0} \quad (11)$$

Proposition 1.7 derives this result formally. The coefficients $c_{j,pq}^{\mathcal{G}}$ (real, since g_{pq} is Hermitian) have a very non-trivial structure, but they can be calculated using Eq. A13.

Now, since $g_{pq}(\mathbf{b})$ represents the predicted corrupted visibilities (given that we have a unity model), it is fair to ask, what image-plane distribution corresponds to the visibility distribution $g_{pq}(\mathbf{b})$? Doing an inverse 2D Fourier transform, we obtain:

$$\mathcal{F}^{-1}\{g_{pq}\}(\mathbf{s}) = \sum_{j=-\infty}^{\infty} c_{j,pq}^{\mathcal{G}} \delta(\mathbf{s} - j\mathbf{s}_0), \quad (12)$$

i.e. a sum of delta-functions whose locations are integer multiples of $\mathbf{s}_0 = (l_0, m_0)$.

Let us now define something we'll call the “ \mathcal{G} -sky of baseline pq ” as follows:

$$\begin{aligned} I_{pq}^{\mathcal{G}}(\mathbf{s}) &= \mathcal{F}^{-1} \left\{ g_{pq} \left(\frac{\mathbf{b}}{\phi_{pq}} \right) \right\} \\ &= \sum_{j=-\infty}^{\infty} c_{j,pq}^{\mathcal{G}} \delta \left(\mathbf{s} - \frac{j\mathbf{s}_0}{\phi_{pq}} \right). \end{aligned} \quad (13)$$

The physical meaning of $I_{pq}^{\mathcal{G}}$ is as follows: it is a sky distribution whose Fourier transform yields a visibility distribution that, along the uv -track given by $\phi_{pq}\mathbf{b}_0(t)$, is consistent with the *predicted corrupted visibilities* g_{pq} along the track given by $\mathbf{b}_0(t)$ (note how the scaling relationship of Eq. 9 enters into Eq. 13). In other words, after the best-fitting calibration gains have been applied, the resulting predicted visibilities for each baseline pq will be consistent with a sky of delta functions spaced at intervals of s_0/ϕ_{pq} , with intensities given by $\{c_{j,pq}^{\mathcal{G}}\}$.

These delta functions are the fundamental ingredients of the ghosts observed in Fig. 1. We will shortly show that the corrected visibilities exhibit a similar structure, but first let us consider what happens to the visibilities given by g_{pq} during imaging.

3.5.1 More on extrapolation

It has been our experience that the mathematical construct of extrapolated visibility functions, which is key to the above arguments, is particularly difficult to explain or justify clearly. In this section we attempt to reformulate the argument again in general terms.

In order to understand ghost formation, we need to understand the behaviour of the best-fitting visibilities produced by the calibration process. The actual visibilities per each baseline are sampled along an elliptical track in the uv -plane. Analysing the mathematical properties of functions defined along a specific uv -track proved to be a difficult problem, analysing continuous functions defined over the entire uv -plane proved more fruitful. We therefore proceed by finding a unique mapping from the former (specific) problem to the latter (general) problem, and back.

More specifically, the extrapolation operation defined above (Eq. 9) provides a formal recipe for mapping sets of per-baseline visibilities onto functions defined over the entire uv -plane (giving us a per-baseline “virtual uv -plane” that is consistent with the visibilities over the one specific track of that baseline). We then define the ALS calibration process in terms of operations on such virtual uv -planes. The virtual uv -planes corresponding to the best-fit visibilities ($g_{pq}(\mathbf{b})$, bottom row of Figs. 3-5), of which the actual visibilities are a subset (given by the baseline’s uv -track) turn out to have certain mathematical properties: they are Hermitian, one-dimensional and periodic (with the same period across all baselines), and therefore correspond to a string of delta-functions in the image domain. Note how these properties are straightforward to establish for functions defined over the entire uv -plane, but are a lot less obvious if one only considers a subset of the uv -plane along a track. (This observation is the main justification for the extrapolation formalism.) This establishes that the best-fit visibilities per

each baseline are consistent with a string of delta functions. Finally, the geometric scaling relationship implicit in Eq. 9 causes the spacing of the delta-functions to be inversely proportional to baseline length.

3.6 Imaging

In the situation above, each baseline’s predicted corrupted visibilities correspond to its own apparent sky I_{pq} . During conventional interferometric imaging, the per-baseline visibilities are interpolated onto a “common” uv -plane using convolutional gridding, and the result is Fourier transformed back into an estimate of the sky (the so-called “dirty image”). Mathematically, this can be described as follows:

$$I_D = \mathcal{F}^{-1} \left\{ \sum_{pq} S_{pq} \mathcal{F}\{I_{pq}\} \right\}. \quad (14)$$

Here, S_{pq} is the *sampling function* of baseline pq . The sampling function is only non-zero in the neighbourhood of the track described by \mathbf{u}_{pq} , and accounts for both the imaging weights and the interpolation coefficients of the gridding process. This can be rewritten as

$$I_D = \sum_{pq} P_{pq} \circ I_{pq}, \quad (15)$$

where “ \circ ” denotes convolution, and $P_{pq} = \mathcal{F}^{-1}\{S_{pq}\}$ is the (unnormalized) PSF associated with baseline pq . Note that in the case of each baseline seeing a common sky I , the above becomes

$$I_D = \left(\sum_{pq} P_{pq} \right) \circ I, \quad (16)$$

which is the familiar result that the dirty image I_D is the convolution of the true sky I by the point spread function of the array P given by

$$P = \sum_{pq} P_{pq}. \quad (17)$$

Now, recall that Eq. 13 describes a string of delta functions spaced at intervals of s_0/ϕ_{pq} . If we define ϕ_0 as the least common multiple of all ϕ_{pq} , we can rewrite the equation as a sequence of delta functions spaced at intervals of s_0/ϕ_0 , some of them possibly of zero amplitude:

$$I_{pq}^{\mathcal{G}}(\mathbf{s}) = \sum_{k=-\infty}^{\infty} d_{k,pq}^{\mathcal{G}} \delta \left(\mathbf{s} - \frac{k\mathbf{s}_0}{\phi_0} \right), \quad (18)$$

where $d_{k,pq}^{\mathcal{G}} = c_{j,pq}^{\mathcal{G}}$ if there is an integer j such that $k\phi_{pq} = j\phi_0$, and zero otherwise. To simplify further equations, we'll use the δ_k as shorthand for the k -th delta function above:

$$\delta_k(\mathbf{s}) = \delta \left(\mathbf{s} - \frac{k\mathbf{s}_0}{\phi_0} \right).$$

Substituting this into Eq. 15, we get

$$I_D^{\mathcal{G}} = \sum_{pq} P_{pq} \circ \left(\sum_{k=-\infty}^{\infty} d_{k,pq}^{\mathcal{G}} \delta_k \right) \quad (19)$$

$$= \sum_{k=-\infty}^{\infty} \left(\sum_{pq} d_{k,pq}^{\mathcal{G}} P_{pq} \right) \circ \delta_k. \quad (20)$$

Physically, this can be interpreted as follows. The dirty image $I_D^{\mathcal{G}}$ which we get as a result of imaging the predicted corrupted visibilities consists of a string of delta functions at regularly-spaced locations $k\mathbf{s}_0/\phi_0$, each one convolved with its own *ghost spread function* (GSF) $P_k^{\mathcal{G}}$:

$$P_k^{\mathcal{G}} = \sum_{pq} d_{k,pq}^{\mathcal{G}} P_{pq}. \quad (21)$$

Comparing this to Eq. 17, we can now understand the previously puzzling observation that the ghost sources in Fig. 1 appear to be convolved with differing point spread functions, similar but not identical to the nominal PSF of the WSRT.

Furthermore, ghost positions do not depend on frequency (only on array and source geometry) – though the GSF of course does. This is also consistent with previous observations.

3.7 Corrected visibilities

In real life, one would typically be imaging the *corrected visibilities* (Sect. 2.2) given by

$$\mathcal{R}^{(c)} = \mathbf{G}^{-1} \mathcal{R} \mathbf{G}^{-H} = \mathcal{G}^{\top} \odot \mathcal{R}, \quad (22)$$

and our real goal is to understand the effect of \mathcal{G} on the *corrected sky* $I^{(c)} = \mathcal{F}^{-1}\{\mathcal{R}^{(c)}\}$. To get there, we need to take an intermediate step. First, let us define a “ \mathcal{G}^{\top} -sky” whose Fourier transform is consistent with the visibility distribution given by g_{pq}^{-1} . Proposition 1.8 shows² that the visibility distribution $g_{pq}^{-1}(\mathbf{b})$ can also be decomposed into a Fourier series:

$$g_{pq}^{-1}(\mathbf{b}) = \sum_{j=-\infty}^{\infty} c_{j,pq}^{\top} e^{2\pi i j \mathbf{b} \cdot \mathbf{s}_0}, \quad (23)$$

which implies that the corresponding “ \mathcal{G}^{\top} -sky” has a similar form to Eq. 18, but with a different set of coefficients:

$$I_{pq}^{\mathcal{G}^{\top}} = \sum_{k=-\infty}^{\infty} d_{k,pq}^{\top} \delta_k. \quad (24)$$

Now, consider the matrix $\mathcal{G}^{\top}(\mathbf{b}) \odot \mathcal{R}(\mathbf{b})$. We’ll designate its elements as $r_{pq}^{\top}(\mathbf{b})$. The inverse Fourier transform of each element is then

$$\mathcal{F}^{-1}\{r_{pq}^{\top}\} = \mathcal{F}^{-1}\{g_{pq}^{-1}\} \circ \mathcal{F}^{-1}\{r_{pq}\}, \quad (25)$$

and the inverse Fourier transforms of both components have already been derived above. This means that the “corrected sky” corresponding to the corrected visibilities of baseline pq is given by

$$I_{pq}^{(c)} = I_{pq}^{\mathcal{G}^{\top}} \circ I^{\mathcal{R}}, \quad (26)$$

² Note that this proposition implicitly assumes $g_{pq} \neq 0$, i.e. that the ALS calibration solutions are not null. Intuition suggests that this is a safe assumption: a null gain solution would yield null predicted visibilities, which could hardly be a “best fit” to the calibration equation in any sense. However, obtaining a rigorous proof of this has been surprisingly difficult, so we will let the assumption stand as is.

i.e. is simply a convolution of the real sky $I^{\mathcal{R}}$ with the “ghost pattern” of delta functions given by $I_{pq}^{\mathcal{G}^{\top}}$ above. In other words, the “corrected sky” will contain multiple instances of the fundamental ghost pattern (what we call the distilled ghost pattern), centered on each source, and scaled by the flux of that source. It can be seen that the sky corresponding to the residuals \mathcal{R}^{Δ} is given by

$$I_{pq}^{\Delta} = I_{pq}^{\mathcal{G}^{\top-1}} \circ I^{\mathcal{R}} = (I_{pq}^{\mathcal{G}^{\top}} - \delta) \circ I^{\mathcal{R}}. \quad (27)$$

The $I_{pq}^{\mathcal{G}^{\top-1}}$ term is the per-baseline sky associated with the distilled ghost pattern $\mathcal{G}^{\top} - 1$. By analogy with Eq. 20, we can derive an expression for the full dirty image:

$$I_D^{\mathcal{G}^{\top-1}} = \sum_{k=-\infty}^{\infty} \left(\sum_{pq} \hat{d}_{k,pq}^{\top} P_{pq} \right) \circ \delta_k, \quad (28)$$

where $\hat{d}_{k,pq}^{\top} = d_{k,pq}^{\top} - 1_{\{k=0\}}$, and the notation $1_{\{k=0\}}$ represents a series whose coefficients are 1 at $k = 0$ and 0 elsewhere. Note the physical meaning of this operation: the $d_{0,pq}^{\top}$ coefficient corresponds to the position of the real source A_1 in the image, and is close to unity, since the corrected visibilities contain the real source as well as all the ghosts. Subtracting 1 from this corresponds to taking the residual visibilities.

In our simple case the real sky consists of two sources, and the resulting corrected sky is a superposition of two patterns given by $I_{pq}^{\mathcal{G}^{\top}}$, scaled by A_1 and A_2 , and centered on origin and on \mathbf{s}_0 , respectively. Since each pattern yields ghosts at discrete intervals of \mathbf{s}_0/ϕ_0 , the two sets of positions align, and we can work out the amplitudes of the resulting superposed ghost sources by summing up the corresponding coefficients. By analogy with Eq. 20, we can derive the following equation for the dirty image formed from corrected visibilities:

$$I_D^{(c)} = \sum_{k=-\infty}^{\infty} \left(\sum_{pq} (A_1 d_{k,pq}^{\top} + A_2 d_{k-\phi_0,pq}^{\top}) P_{pq} \right) \circ \delta_k. \quad (29)$$

From Eq. 27, it follows that the dirty image corresponding to the residuals can be obtained by subtracting unity from the corresponding d coefficient:

$$I_D^{\Delta} = \sum_{k=-\infty}^{\infty} \left(\sum_{pq} (A_1 \hat{d}_{k,pq}^{\top} + A_2 \hat{d}_{k-\phi_0,pq}^{\top}) P_{pq} \right) \circ \delta_k, \quad (30)$$

where again $\hat{d}_{k,pq}^{\top} = d_{k,pq}^{\top} - 1_{\{k=0\}}$.

Equations 20, 28, 29 and 30 summarize the formation of ghost patterns in the predicted corrupted, corrected and residual visibilities.

For our purposes, it is important to derive a theoretical prediction for the amplitudes of individual ghost sources as a fraction of the “missing flux” A_2 . Consider the k -th ghost source located at $k\mathbf{s}_0/\phi_0$. From Eq. 28, it follows that the amplitude of the ghost source is given by the weighted sum

$$\sum_{pq} \hat{d}_{k,pq}^{\top} P_{pq}(0),$$

where the per-baseline weights $P_{pq}(0)$ are ultimately determined by the imaging weights. For simplicity, let us consider the case of natural weighting, in which case the sum becomes unweighted. We can then define the quantity

$$\zeta_k = \langle \hat{d}_{k,pq}^\top \rangle_{pq} \quad (31)$$

where $\langle \cdot \rangle_{pq}$ denotes averaging over all baselines pq . This gives us the amplitude of the k -th ghost source in the distilled pattern (assuming natural weighting), and can be computed analytically from the results above. Likewise, the k -th source in the corrected residuals is a superposition of two appropriately scaled sources from the distilled pattern, and its amplitude is given by (assuming natural weighting):

$$\zeta_k^\Delta = A_1 \zeta_k + A_2 \zeta_{k-\phi_0} \quad (32)$$

Of particular interest is the quantity $\zeta_{\phi_0}^\Delta/A_2$, as this gives the relative amplitude of the “flux suppression ghost” sitting on top of source A_2 . Indirectly, this one ghost has been observable since the invention of selfcal, since it corresponds to the previously noted phenomenon of flux suppression in unmodelled sources. The theoretical derivation given here provides an explanation for this.

4 RESULTS

Section 3.5 provides a theoretical framework for understanding ghost formation, as well as a mechanism for predicting the distribution and amplitudes of ghosts in the two-source case. In this section, we apply the mechanism to predict ghost formation for a specific observational scenario, and compare the results with simulations.

As discussed above, the ghost pattern is highly dependent on the array configuration. The results in this section were all generated with a traditional (36,108,1332,1404m) WSRT configuration. Unless specified otherwise, $A_2 = 0.2$ Jy, $A_1 = 1$, Jy, $l_0 = 1^\circ$ and $m_0 = 0^\circ$, and we assumed monochromatic observations at a frequency of 1.45 GHz. To verify the theory developed above, we compare the distribution of ghost sources obtained by three methods:

- A theoretically predicted distribution, using the framework of Section 3.5.
- ALS calibration of simulated data (using a custom-made implementation).
- LS calibration of simulated data using the MeqTrees (Noordam & Smirnov 2010) package.

Fig. 6 displays the theoretically determined distilled ghost patterns for a selection of baselines (9A: 36m, 01 and 12: 144m, 05 and 16: 720m, 0D: 2.7 km). We also obtain a set of simulated distilled ghost patterns for the same set of baselines (Figs. 7, 8) as follows:

- We run ALS or LS calibration on a set of simulated visibilities, and derive the calibrated visibilities;
- We image the calibrated visibilities for each baseline (using the *lwimager* program – an FFT-based imager derived from the CASA libraries and functionally equivalent to the CASA imager). The resulting dirty maps are given in Fig. 9;
- We measure fluxes at the ghost source positions in the resulting dirty images, resulting in Figs. 7 and 8.

As predicted by Sect. 3, short baselines yield a few coarsely-spaced ghosts (e.g. 9A), while long baselines (e.g. 0D) yield many finely-spaced ghosts.

Comparing Figs. 6, 7 and 8, the following general observations (and subsequent conclusions) can be made:

- The bright ghost sources in Fig. 6 and the bright sources in Figs. 7–8 show up at the same lm coordinates. This validates Eq. 24.

- The weaker sources given by the theoretical ghost patterns (Fig. 6) are not visible in Fig. 7. Furthermore, there are small differences in flux between the theoretical ghost patterns and the measured fluxes of the corresponding ghost sources in Fig. 7. Note, however, that the dirty images are dominated by the sidelobes of the brighter ghost sources – which in general (for $n > 2$) are not amenable to normal deconvolution, since each ghost spread function (Eq. 21) is different. This both masks the fainter ghost sources, and distorts the flux measurements.

- The ghost patterns yielded by ALS and LS calibration (Fig. 7 and Fig. 8) are qualitatively similar, but show different amplitudes. This is understandable, as they are products of slightly different optimization problems, and therefore yield slightly different calibration solutions. In particular, there are negative ghosts at 0° in Fig. 7, while there are none in Fig. 8. This implies that ALS tends to also suppress the flux of the modelled source, while LS doesn’t. This can be explained by the following argument. The total flux of the sky and the calibration model is given by the diagonal terms (autocorrelations) of \mathcal{R} and \mathcal{M} , while the total power in the off-diagonal terms is zero. When the autocorrelation constraints are ignored (as in LS), there is no restriction on the total flux in the model, which leaves the gain solutions $g_p \bar{g}_q$ in Eq. 3 more freedom to fit the mean amplitude of r_{pq} over time. If the autocorrelations are also fitted (as is the case in ALS, Eq. 2), then the gain solutions must also account for the total flux of the sky ($A_1 + A_2$) using a model containing a total flux of only A_1 . This yields mean gain-amplitudes of slightly above unity in \mathcal{G} , and thus below unity in \mathcal{G}^\top . This results in a negative ghost source at the phase centre in the distilled ghost pattern $I_{pq}^{\mathcal{G}^\top-1}$, i.e. flux suppression of the primary source.

Figs. 10–11 display the ghost patterns that are obtained for ALS and LS calibration using all antennas during imaging. In Fig. 10, the secondary source is at $l = 1^\circ$, while in Fig. 11 it is placed at $l = 20^\circ$ (thus qualitatively reproducing the observational scenario of Fig. 1). In Fig. 11 only the “inner ghosts” (the ghosts between the primary and secondary source) are visible, while some “outer ghosts” are emerging in Fig. 10.

Fig. 12 displays the theoretically determined distilled ghost pattern for the full WSRT array, as percentage of A_2 flux (for the $A_2 = 0.2$ at 1° case). Compare these to the per-baseline patterns in Fig. 6. The pattern exhibits a number of interesting features:

- Most (though not all) ghosts have negative amplitudes; the positive ghosts tend to be fewer and much weaker.
- The strongest response is the “flux suppression” ghost at the A_2 position (1° , or $k = \phi_0$). At about 13%, it is perfectly consistent with the amount of flux suppression normally observed when calibrating WSRT data with LS. As discussed in Sect. 3.6, the $k = n\phi_0$ positions are shared by the ghost patterns of all baselines, and thus favour the formation of strong ghosts. It is not surprising that the $k = \phi_0$ position shows the strongest response overall, as that is where the missing flux that the calibration process is trying to fit is located. The next-brightest ghost ($\sim 6\%$),




Figure 6. Theoretical ghost pattern for baselines 9A (36m), 01 and 12 (144m), 05 and 16 (720m), 0D (2.7km).




Figure 7. ALS ghost pattern for baselines 9A (36m), 01 and 12 (144m), 05 and 16 (720m), 0D (2.7km).




Figure 8. LS ghost pattern for baselines 9A (36m), 01 and 12 (144m), 05 and 16 (720m), 0D (2.7km).

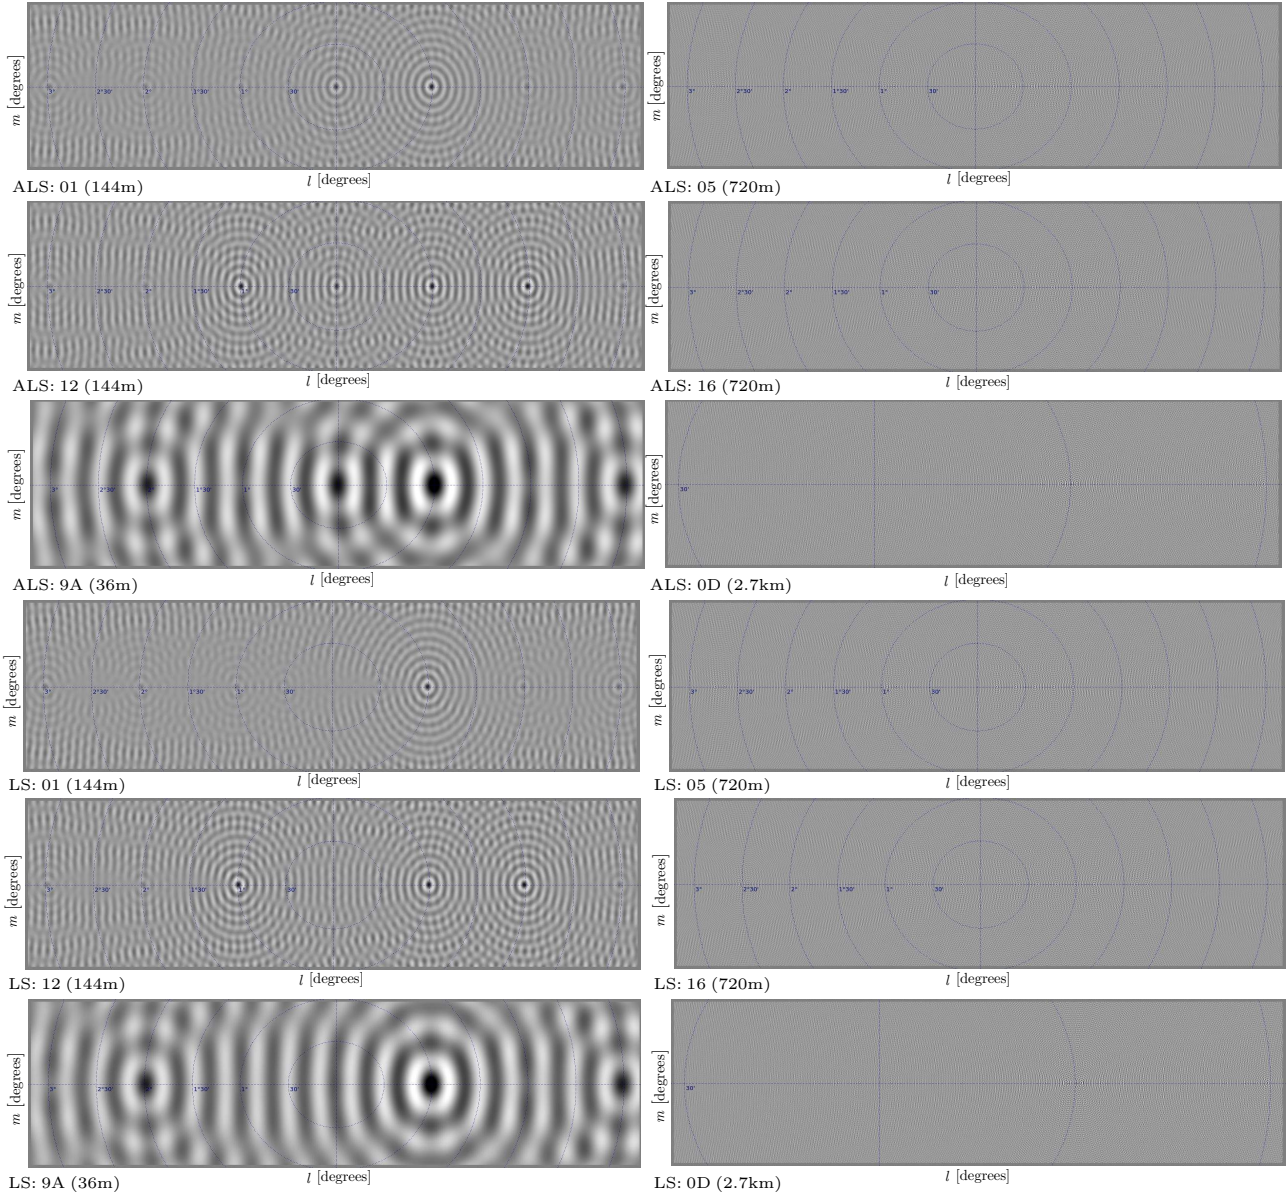


Figure 9. Dirty images of ALS and LS ghost patterns for baselines 9A (36m), 01 and 12 (144m), 05 and 16 (720m), and 0D (2.7km).




Figure 10. ALS (left) and LS (right) distilled ghost pattern for the full WSRT array, with the A_2 source at 1° .

is at 0° . As discussed above, this particular ghost is specific to ALS.

- Curiously, the other “favoured” positions (-2° , -1° , -2° , etc.) show a much diminished response – only about $1 \sim 2\%$ – i.e. weaker than the strongest of the inner ghosts

(see e.g. the “halfway ghost” at $1/2^\circ$, with 2.5%). Comparing this to Fig. 6, we can partially understand how this comes about: different baselines show a mix of positive and negative responses at these positions, whereas the 0° and 1° responses




Figure 11. ALS (left) and LS (right) distilled ghost pattern for the full WSRT array, with the A_2 source at 20° .

are consistently negative. The terms in the sum of Eq. 21 therefore average down at the other positions.

- Most of the strongest remaining ghosts are the inner ones between the two sources ($0^\circ - 1^\circ$). However, the “outer ghosts” seem to extend indefinitely at the $0.1 \sim 0.2\%$ level.

The latter two points are especially puzzling, and there should be some fundamental mathematical reason for why this should be so, but it escapes us at present.

4.1 Dependence on flux ratio

Now that Eq. 24 has been validated, the natural question arises of how l_0, m_0, A_1 and A_2 influence the amplitudes of the ghost pattern. Proposition 1.7 implies that the position of the secondary source l_0, m_0 has no influence on the amplitudes – it only stretches or shrinks the ghost patterns, and determines their orientation. (This also explains why it was sufficient to verify the validity of Eq. 24 at only one position of the secondary source.) The source fluxes, obviously, do have an effect. As discussed in Section 3.3, the matrix \mathcal{G} is determined by A_2 (which is equivalent to the flux ratio, since we’ve been assuming $A_1 = 1$ throughout), which implies that the ghost amplitudes given by the various $d_{k,pq}$ coefficients are dependent on A_2 .

The actual ghost amplitudes do not have a simple analytic representation as they are ultimately determined by the interaction between the largest eigenvalue of \mathcal{R} and its associated eigenvector (Eq. A13). We can, however, empirically show an approximately linear dependence on A_2 . Let us postulate this dependence:

$$\zeta_k \approx K_k A_2, \quad (33)$$

and find an estimate for each K_k over a range of A_2 values using least-squares. The relative magnitude of the error of the fit:

$$\epsilon_k = \left| \frac{\zeta_k - K_k A_2}{\zeta_k} \right|, \quad (34)$$

as a function of A_2 for the thirteen brightest ghosts is plotted in Fig. 13. This shows that most ghosts vary linearly with A_2 to within 10%. Curiously, the flux suppression ghost (1°) is linear to within 1%, but the ghosts at 2° and 3° are the least linear of all.³ Approximate linear models for the amplitudes of all the ghosts in the distilled ghost pattern can be derived in this manner.

³ Yet another mathematical puzzle raised by the ghost phenomenon. We have no theoretical explanation for this at present!

Consider now the dependence on A_1 (for which we’ve used a fixed value of 1 Jy until now). It is obvious from the calibration equation that rescaling the true sky and the model sky by the same factor will have no effect on the solutions matrix \mathcal{G} , which completely determines the ζ_k coefficients, i.e. the distilled ghost amplitudes. This means that a calibration problem with fluxes of $A_1 = A'_1 \neq 1$, $A_2 = A'_2$ will produce the same coefficients as one with $A_1 = 1$, $A_2 = A'_2/A'_1$. In other words, *the distilled ghost pattern is determined by the flux ratio of the two sources rather than their absolute fluxes.*

The ghost pattern in the corrected or residual visibilities, on the other hand, is a convolution of the distilled pattern with the sky, and therefore *will* scale with absolute flux. To be more precise, for fluxes A'_1, A'_2 , the resulting residual ghost amplitudes (Eq. 32) will be:

$$\zeta_k^\Delta \approx A'_1 K_k \frac{A'_2}{A'_1} + A'_2 K_{k-\phi_0} \frac{A'_2}{A'_1} = K_k A'_2 + K_{k-\phi_0} \frac{A'^2_2}{A'_1}. \quad (35)$$

When $A'_2 \ll A'_1$, the first term in the sum dominates, which makes the ghost patterns in the residual image *nearly independent of A'_1* . This explains the behaviour observed by Smirnov (2010) and discussed in the introduction.

Fig. 14 shows the theoretically-derived relative ghost source amplitudes ζ_k/A_2 for the 13 strongest⁴ ghosts of the distilled ghost pattern, as a function of A_2 . A true linear dependence would have yielded constant horizontal lines; deviation from horizontal indicates deviation from linearity. As we saw above, the 2° and 3° ghosts appear to be the least linear.

Fig. 15 shows the same amplitudes for the residual ghost pattern. Since the residual pattern is a superposition of two scaled fundamental patterns, the dependence is different due to the additional linear component given by the second term of Eq. 35 (since we’re plotting relative amplitudes, the equation should be divided by A'_2). As expected, this component becomes negligible for $A_2 \ll A_1$. For larger A_2 the linear component can actually come to dominate – note how the 2° ghost becomes stronger than the 1° ghost for larger values of A_2 (which is not surprising, since it contains the 1° component from the distilled ghost pattern, scaled by A_2).

⁴ Strongest at $A_2 = 0.5$, to be precise. As the plots show, the relative ranking of the ghosts can actually change as a function of A_2 .




Figure 12. Theoretical residual ghost pattern within 3° (left plot) and 20° (right plot) of phase centre, for the full WSRT array. The A_2 source was at 1° . Amplitudes are given as a percentage of A_2 flux. The 3° plot shows all the ghosts, while in the 20° plot, only ghosts with amplitudes in excess of $\approx 0.009\%$ of A_2 are shown, and the y axis is cut off just below the $1/2^\circ$ ghost – the 0° and 1° ghost response thus extends well below the plot limits.

4.2 The role of array redundancy

WSRT’s highly redundant configuration plays a very important role in ghost formation. Theoretically, this is explained by Eq. 18. The set of all possible ghost positions is discrete, and given by $\{k\mathbf{s}_0/\phi_0\}$. Each baseline pq yields ghosts at a specific subset of these positions, i.e. at intervals (in k) of ϕ_0/ϕ_{pq} . For short baselines, ϕ_{pq} is small, and few ghosts are produced, and vice versa for long baselines. Positions corresponding to redundant baselines, or more generally to common integer factors of multiple ϕ_{pq} ’s, will then host stronger ghosts due to a contribution from multiple baselines.

This effect is vividly illustrated by Figs. 6–8. The shortest baseline (9A, 36m, circle symbol) produces the most widely-spaced ghosts, at intervals of 1° . The 144m baselines (01 and 12, up/down triangles) produce ghosts at 0.25° , the 720m baselines (05 and 16, left/right triangles) produce ghosts at 0.05° , and the longest baseline (0D, 2.7km, horizontal tick marks) produces the most finely spaced ghosts. Groups of redundant baselines (01 and 12, 05 and 16) yield ghosts at exactly the same positions, but with different amplitudes (sometimes even of different sign). The difference is explained by the fact that the antennas constituting redundant spacings form slightly different sets of baselines to other antennas, and are thus subject to different calibration constraints.

The positions corresponding to $\{k = n\phi_0\}$ (in this case, multiples of 1 degree) will have contributions from all baselines, and indeed (as we’ve shown above), the 0 and 1 positions yield the strongest ghosts in the combined pattern. Likewise, the next-strongest ghost appears at the halfway point ($k = \phi_0/2$), since many ϕ_{pq} ’s are even in the (36,108,1332,1404m) WSRT configuration. Other prominent ghosts may be expected at other rational fractions of \mathbf{s}_0 , which fully explains earlier observations.

Equation 18 also provides us with a qualitative understanding of ghost patterns for a less regularly-spaced East-West array. As a mental experiment, we may pick a length for \mathbf{b}_0 (say, 1m), and imagine moving the WSRT antennas to

new positions such that the spacings are still integer multiples of \mathbf{b}_0 , but are mutually prime. The least common multiple ϕ_0 would then be the product of all spacings, and would be very large (and most entries of the geometry matrix Φ would be very large). Furthermore, no two baselines would yield ghosts at any common position apart from $\{k = n\phi_0\}$. The resulting pattern would then consist of very many finely spaced and weak ghosts, with a lot of interaction between the GSF sidelobes, and would therefore be a lot more noise-like. Further decreasing \mathbf{b}_0 would increase ϕ_0 even more, thus spacing the ghosts even finer and further washing out the overall response. Of course, to within some fraction of the dish size, any conceivable array layout can be considered regularly-spaced (with a very large ϕ_0), so we can only properly talk about arrays that are more regular (WSRT, small ϕ_0) or less regular (large ϕ_0). The argument above suggests that highly regular array layouts result in more widely separated and stronger ghosts. This is a hitherto unforeseen disadvantage to redundancy, and should be investigated and quantified in light of current arguments promoting redundancy in future telescope designs (Noorishad 2013) so as to exploit the redundancy calibration technique (Noordam & de Bruyn 1982).

Similar considerations apply to fully 2D/3D arrays such as LOFAR and the JVLA, where we may expect the overall ghost response to be a lot more scattered and noise-like. An upcoming Paper II (Grobler et al., in prep) will study this subject in more detail. Here we will just note that an exception to the above considerations are ghosts occupying the $\{k = n\phi_0\}$ positions, which the theory shows must be yielded by all baselines, regardless of redundancy. The two strongest ones, at positions 1 and 0 – the “flux suppression ghost” and the primary source ghost – sit on top of actual sources, and are therefore not easy to detect as separate artefacts. The other $\{k = n\phi_0\}$ positions seem to yield much weaker ghosts in practice (see above). However, ghosts at the -1 position have recently been spotted in LOFAR data in two independent instances (de Bruyn, priv. comm., Fender, priv. comm.) The latter in particular was associated with a bright




Figure 13. The relative error magnitude of the linear fit (Eq. 34) for the thirteen brightest ghosts. The legends are the ghost locations, with source A_2 being at 1° .

transient source. This suggests that the -1 ghost can yield a strong response under some conditions. We will investigate this phenomenon in Paper II.

5 CONCLUSIONS AND FUTURE WORK

In this work, we have demonstrated that the ghost source phenomenon is a fundamental aspect of the selfcal process, and will invariably arise during calibration with an incomplete sky model. It is perhaps surprising that these features were not explicitly noted before 2004, but radio interferometrists are accustomed to esoteric instrumental artefacts showing up in poorly-calibrated maps, so something that is faint to begin with and goes away with a more accurate calibration would not necessarily have attracted attention. The strongest ghost – the one sitting on top of the missing model source – has been indirectly observed as a matter of course, under the guise of “flux suppression”, but being masked by the source itself, it did not give enough clues to the true extent of the phenomenon. Consider also that the spacing between ghosts (given by s_0/ϕ_0) decreases as s_0 becomes small, while the size of the ghosts remains the same (as it is similar to the PSF size). There is therefore some critical s_0 beyond which the ghosts will begin to blend together into a single “spoke” connecting the sources. In the authors’ experience, poorly calibrated WSRT maps will often exhibit spokes connecting nearby brighter sources, which tend to go away once the sky model is improved. It is quite plausible that we’ve been seeing “blended ghosts” all along. Since the primary beam makes it more likely for apparently bright sources to be near rather than distant, such blends may be the most frequent manifestation of the ghost mechanism.

Although we have only studied the two-source case theoretically, empirical results (as well as the observations of Fig. 2) suggest that with multiple missing sources (which would be a far more typical case), many individual “ghost strings” are produced at different orientations. This would cause the artefacts and their positive/negative sidelobes to overlap and, to an extent, average out, thus washing out the regular geometric pattern and further contributing to their

non-detection. It took a special set of circumstances – i.e. a single relatively strong and distant unmodelled source, high sensitivity, and WSRT’s regular and highly redundant geometry – to make the ghosts stand out (Fig. 1) as a regular and peculiar geometric feature.

Given the increased sensitivity of current and upcoming instruments, and the emergence of statistical detection techniques, as argued in the introduction, a good understanding of the ghost phenomenon is vital. We have developed a theoretical framework for explaining ghost formation in the two-source WSRT case, which has yielded predictions that closely match simulations, qualitatively match actual observations (which are of course always more complex than just two sources), and explain all the previously puzzling features of the ghost phenomena that were empirically observed earlier (positioning at rational fractions along a line connecting two sources, a “ghost spread function” that differs from the PSF, independence – to first order – on the flux of the modelled source). We have also established that the well-known phenomenon of flux suppression is actually just one manifestation of the same underlying mechanism.

A particularly intriguing avenue of further research is ghost formation in the presence of direction-dependent effects. Recall that in the original QMC result (Fig. 2), the ghosts were found to be associated with uncorrected DDEs, and went away once direction-dependent solutions had been applied. In this case, the unmodelled flux responsible for ghost formation must have been due to errors in the voltage beam, caused by the artificially large pointing error. We may postulate that any unmodelled flux, whether due to errors in the sky model, or unaccounted-for DDEs, will lead to ghosts on some level.

More work is required to describe the ghost phenomenon, both empirically and theoretically, for 2D arrays, for multiple and extended unmodelled sources, and for various forms of direction-dependent calibration. In particular, flux suppression, as the strongest manifestation of the ghost mechanism, needs to be studied in more detail (especially in light of upcoming deep, blind surveys). We must also investigate alternative calibration approaches. In particular, robust calibration (Kazemi & Yatawatta 2013) has been shown to result in less flux suppression, and must necessarily exhibit different ghost behaviour.

Ultimately, we need to develop a theoretical and/or numerical mechanism for answering the following fundamental questions. Given an observational scenario, how deep/accurate does a sky model need to be in order to suppress ghosts to a given level? And, what then are the statistical properties and signatures (power spectrum, etc.) of the remaining ghost artefacts? Building on the theory and numerical tools developed in this work, prospects are good that we can eventually provide rigorous answers to these questions.

ACKNOWLEDGMENTS

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.

Jan Noordam lovingly maintained a picture of the 2004 ghosts on a wall of his office, thus keeping the puzzle alive

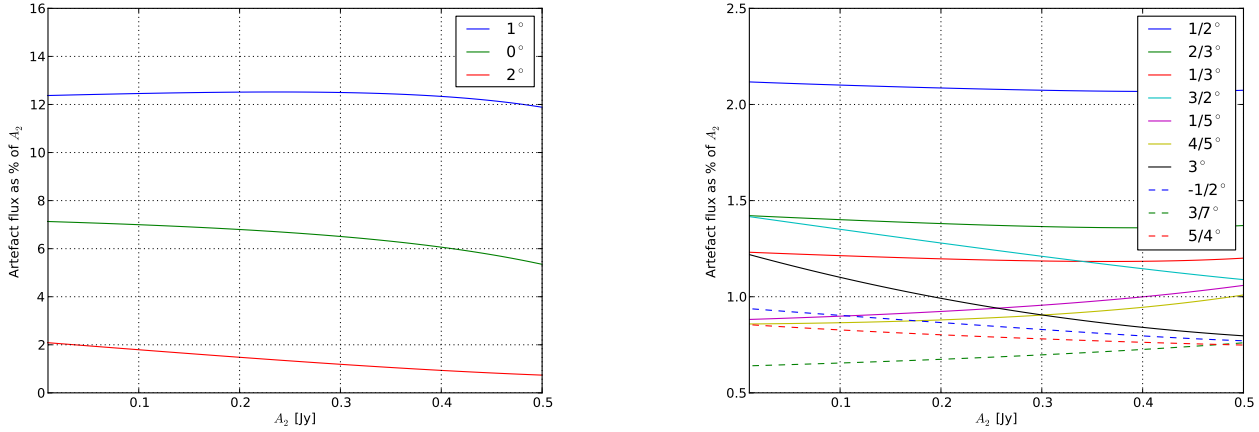


Figure 14. The relative amplitude $|\zeta_k/A_2|$ of the top 13 (top 3 on the left, 4–13 on the right) ghosts in the distilled ghost pattern, as a function of A_2 . Ranking is by ghost amplitude at $A_2 = 0.5$. The ghost positions are indicated by the legend, with source A_2 being at 1° .

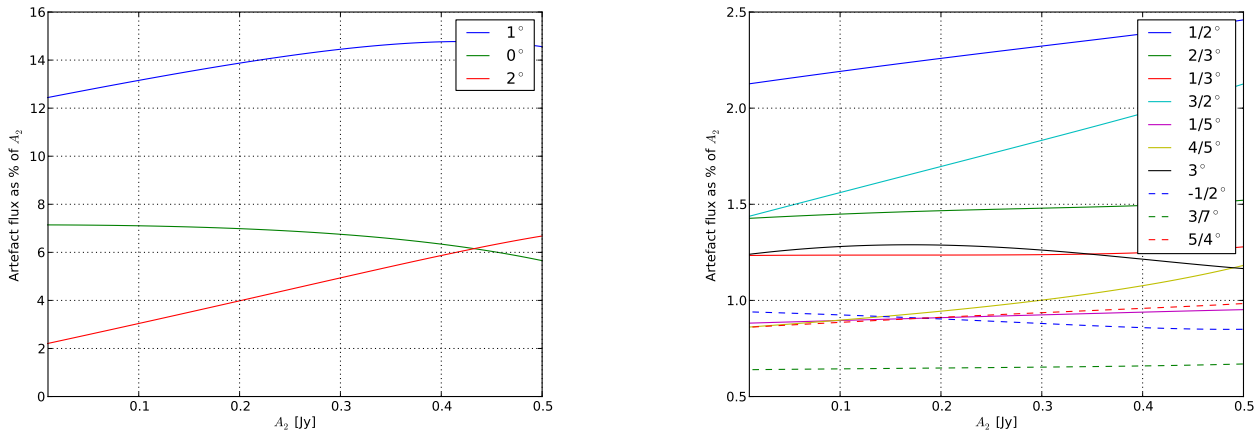


Figure 15. The relative amplitude $|\zeta_k^\Delta/A_2|$ of the top 13 (top 3 on the left, 4–13 on the right) ghosts in the residual ghost pattern, as a function of A_2 . Ranking is by ghost amplitude at $A_2 = 0.5$. The ghost positions are indicated by the legend, with source A_2 being at 1° .

until the next clue dropped in 2010, and contributed many challenging theories subsequently. The link to ALS was made thanks to fruitful discussions with Tobia Carozzi and Griffin Foster. We thank Gianni Bernardi and Cyril Tasse for critical comments on earlier drafts.

We would like to thank the anonymous referee for a thorough review, and very helpful suggestions that improved the paper.

REFERENCES

- Bernardi G., de Bruyn A. G., Brentjens M. A., Ciardi B., Harker G., Jelić V., Koopmans L. V. E., Labropoulos P., Offringa A., Pandey V. N., Schaye J., Thomas R. M., Yatawatta S., Zaroubi S., 2009, *A&A*, 500, 965
- Blinn J., 1996, *IEEE Comput. Graphics and App.*, 16, 82
- Boonstra A., van der Veen A., 2003, *IEEE Trans. Sig. Proc.*, 51, 25
- Chang T.-C., Pen U.-L., Bandura K., Peterson J. B., 2010, *Nature*, 466, 463
- Cornwell T. J., Wilkinson P. N., 1981, *MNRAS*, 196, 1067
- Delhaize J., Meyer M. J., Staveley-Smith L., Boyle B. J., 2013, *MNRAS*, 433, 1398
- Furlanetto S. R., Oh S. P., Briggs F. H., 2006, *Phys. Rep.*, 433, 181
- Hamaker J. P., Bregman J. D., Sault R. J., 1996, *A&AS*, 117, 137
- Ikramov K. D., 2009, *J. of Math. Sci.*, 157, 695
- Intema H., Van der Tol S., Cotton W., Cohen A., Van Bemmelen I., Röttgering H., 2009, *A&A*, 501, 1185
- Kazemi S., Yatawatta S., 2013, preprint (arXiv:1307.5040)
- Kazemi S., Yatawatta S., Zaroubi S., 2013, *MNRAS*, 430, 1457
- Kazemi S., Yatawatta S., Zaroubi S., Lampropoulos P., de Bruyn A., Koopmans L., Noordam J., et al., 2011, *MNRAS*, 414, 1656
- Kılıç M., Mengi E., Yıldırım E. A., 2011, preprint (arXiv:1109.2080)
- Kopp J., 2008, *Int. J. of Modern Phys. C*, 19, 523
- Lax P., 1996, *Linear Algebra and its applications*. John

- Wiley & sons
- Levenberg K., 1944, The Quarterly J. of App. Math., 2, 164
- Linfield R., 1986, AJ, 92, 213
- Marquardt D., 1963, SIAM J. of App. Math., 11, 431
- Martí-Vidal I., Guirado J., Jiménez-Monferrer S., Marcaide J., 2010, A&A, 517
- Martí-Vidal I., Marcaide J., 2008, A&A, 480, 289
- Noordam J. E., de Bruyn A. G., 1982, Nature, 299, 597
- Noordam J. E., Smirnov O. M., 2010, A&A, 524, A61
- Noorishad P., 2013, PhD thesis, Univ. Groningen
- Norris R. P., Afonso J., Bacon D., Beck R., Bell M. e. a., 2013, PASA, 30, 20
- Perlis S., 1952, Theory of matrices. Dover Publications
- Pober J. C., Parsons A. R., Aguirre J. E., Ali Z., Bradley R. F., Carilli C. L., DeBoer D., Dexter M., Gugliucci N. E., Jacobs D. C., Klima P. J., MacMahon D., Manley J., Moore D. F., Stefan I. I., Walbrugh W. P., 2013, ApJ, 768, L36
- Rau U., Bhatnagar S., Voronkov M., Cornwell T., 2009, Proc. of the IEEE, 97, 1472
- Reed M., Simon B., 1978, Analysis of operators. Methods of modern mathematical physics IV. New York: Academic
- Smirnov O. M., 2010, Ghostbusters: The Unknown Unknowns Of Selfcal, presentation at CALIM2010 conference (Dwingeloo 2010), http://www.astron.nl/calim2010/presentations/14_Ghostbusters_Smirnov.pdf
- Smirnov O. M., 2011, A&A, 527, A106
- Smirnov O. M., 2011a, A&A, 527
- Smirnov O. M., 2011b, Solving For Primary Beams, Pointing Errors And The Westerbork Wobble, presentation at CALIM2011 conference (Manchester 2011), <https://indico.skatelescope.org/getFile.py/access?contribId=20&sessionId=9&resId=0&materialId=0&confId=171>
- Switzer E. R., Masui K. W., Bandura K., Calin L.-M., Chang T.-C., Chen X.-L., Li Y.-C., Liao Y.-W., Natarajan A., Pen U.-L., Peterson J. B., Shaw J. R., Voytek T. C., 2013, MNRAS, 434, L46
- Taylor G., Carilli C., Perley R., 1999, in Synthesis Imaging in Radio Astronomy II Vol. 180, Synthesis imaging in radio astronomy II
- van der Veen A., Leshem A., Boonstra A., 2004, in Sensor Array and Multichannel Signal Processing Workshop Proceedings, 2004 Signal processing for radio astronomical arrays. pp 1–10
- Wijnholds S., van der Veen A., 2009, IEEE Trans. Sig. Proc., 57, 3512
- Wilkinson P., Conway J., Biretta J., 1988, in The Impact of VLBI on Astrophysics and Geophysics Vol. 129, Developments in continuum imaging. pp 509–518

APPENDIX A: LEMMAS AND PROPOSITIONS

This appendix contains formal mathematical derivations of the propositions used in Sect. 3.

Proposition 1.5. If the function-valued matrix $\mathcal{R}(\mathbf{b}): \mathbf{R}^2 \rightarrow \mathbf{C}^{n \times n}$ is defined as stated in Definition 1.3 its

rank does not exceed two and its eigenvalues are either equal to zero or

$$\frac{n(A_1 + A_2)}{2} \pm h, \quad (\text{A1})$$

where $h = \frac{1}{2} \sqrt{[n^2 - 4 \binom{n}{2}][A_1 + A_2]^2 + \kappa}$ and $\kappa = 4 \sum_{p < q} (A_1^2 + A_2^2 + 2A_1 A_2 \cos(2\pi \phi_{pq} \mathbf{b} \cdot \mathbf{s}_0))$.

Proof. The fact that the rank of $\mathcal{R}(\mathbf{b})$ does not exceed two follows trivially from Lemma 1.9–1.11. Since the rank of $\mathcal{R}(\mathbf{b})$ is at most two its characteristic equation is equal to (Ikramov 2009; Blinn 1996)

$$\left(\lambda^2 - \text{tr}(\mathcal{R}(\mathbf{b}))\lambda + \sum_{p < q} \begin{vmatrix} r_{pp} & r_{pq} \\ r_{qp} & r_{qq} \end{vmatrix} \right) \lambda^{n-2} = 0. \quad (\text{A2})$$

Solving for λ in Eq. A2 produces the result. \square

Proposition 1.5 states that the rank of $\mathcal{R}(\mathbf{b})$ is two and gives an analytic expression of its largest eigenvalue, $\lambda(\mathbf{b})$ (Eq. A1). The expression of $\lambda(\mathbf{b})$ is also used in Proposition 1.6 and Lemma 1.17. Proposition 1.5 have three direct dependencies namely, Lemma 1.9, Lemma 1.10 and Lemma 1.11. Lemma 1.9 gives the properties that a matrix must have so that its rank does not exceed $k \in \mathbf{N}$. Lemma 1.10 and Lemma 1.11 show that $\mathcal{R}(\mathbf{b})$ has the required properties so that its rank does not exceed two. The validity of Lemma 1.10 and Lemma 1.11 follows from Lemma 1.14, which gives the mathematical properties of Φ .

Proposition 1.6. The entries $g_{pq}(\mathbf{b})$ of $\mathcal{G}(\mathbf{b})$ are differentiable Hermitian functions. Moreover, $g_{pq}(u, v) = g_{pq}(u + \frac{j}{l_0}, v + \frac{k}{l_0})$ and $g_{pq}(u, -\frac{l_0}{m_0}u + c) = g_{pq}(0, c) \forall j, k \in \mathbf{Z}$ and $\forall u, v, c \in \mathbf{R}$.

Proof. By Lemma 1.15 and Eq. A1

$$\begin{aligned} \mathcal{G}\left(u + \frac{j}{l_0}, v + \frac{k}{l_0}\right) &= \lambda(u, v) \mathbf{x}(u, v) \mathbf{x}^H(u, v) \\ &= \mathcal{G}(u, v), \end{aligned} \quad (\text{A3})$$

for all $j, k \in \mathbf{Z}$. Eq. A3 implies that $g_{pq}(u, v) = g_{pq}(u + \frac{j}{l_0}, v + \frac{k}{l_0})$, $\forall j, k \in \mathbf{Z}$. Similarly, $g_{pq}(u, -\frac{l_0}{m_0}u + c) = g_{pq}(0, c)$, $\forall u, c \in \mathbf{R}$ (by Lemma 1.16). The fact that $g_{pq}(\mathbf{b})$ is a differentiable function is established by Lemma 1.17.

The function $g_{pq}(u, v)$ is the best possible (in a least squares sense) fit of $r_{pq}(u, v)$ (see Eq. 2). From this observation and the fact that \mathcal{R} and \mathcal{G} are Hermitian matrices (as well as the fact that $\mathcal{R}(-\mathbf{b}) = \overline{\mathcal{R}(\mathbf{b})} \Rightarrow \mathcal{G}(-\mathbf{b}) = \overline{\mathcal{G}(\mathbf{b})}$) the following statements logically follow

- (i) The best possible fit of $r_{pq}(-u, -v) = \overline{r_{pq}(u, v)}$ is $g_{pq}(-u, -v)$.
- (ii) The best possible fit of $r_{qp}(u, v) = r_{pq}(-u, -v) = \overline{r_{pq}(u, v)}$ is $g_{qp}(u, v) = \overline{g_{pq}(u, v)}$.

The above statements imply that $g_{pq}(-u, -v) = \overline{g_{pq}(u, v)}$. \square

Proposition 1.6 shows that the elements $g_{pq}(\mathbf{b})$ of $\mathcal{G}(\mathbf{b})$ are periodic, effectively one-dimensional, differentiable, Hermitian functions. The properties of $g_{pq}(\mathbf{b})$ follow from Lemma 1.15 (periodicity), Lemma 1.16 (one-dimensionality) and Lemma 1.17 (differentiability). Lemma 1.17 is a consequence of Rellich's theorem.

Proposition 1.7. Each element $g_{pq}(\mathbf{b})$ of $\mathcal{G}(\mathbf{b})$ can be written as the following sum

$$g_{pq}(\mathbf{b}) = \sum_{j=-\infty}^{\infty} c_j e^{2\pi i j \mathbf{b} \cdot \mathbf{s}_0}, \quad (\text{A4})$$

that is

$$g_{pq}(u, v) = \sum_{j=-\infty}^{\infty} c_j e^{2\pi i j (ul_0 + vm_0)}, \quad (\text{A5})$$

where

$$c_j = \mu \int_{-\frac{1}{2|m_0|}}^{\frac{1}{2|m_0|}} \int_{-\frac{1}{2|l_0|}}^{\frac{1}{2|l_0|}} g_{pq}(u, v) e^{-2\pi i j (ul_0 + vm_0)} dudv, \quad (\text{A6})$$

with $\mu = |l_0||m_0|$ and $c_j \in \mathbf{R}$.

Proof. Since g_{pq} is a differentiable periodic function in \mathbf{R}^2 (Proposition 1.6), consider the standard Fourier series expansion

$$g_{pq}(u, v) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} c_{jk} e^{2\pi i (j l_0 u + k m_0 v)}, \quad (\text{A7})$$

with

$$c_{jk} = \mu \int_{-\frac{1}{2|m_0|}}^{\frac{1}{2|m_0|}} \int_{-\frac{1}{2|l_0|}}^{\frac{1}{2|l_0|}} g_{pq}(u, v) e^{-2\pi i (j u l_0 + k v m_0)} dudv \quad (\text{A8})$$

and $\mu = |l_0||m_0|$.

Since g_{pq} is Hermitian (Proposition 1.6), the coefficients (c_{jk}) are real numbers. Fix $c \in \mathbf{R}$. Note that $g_{pq}(u, -\frac{l_0}{m_0}u + c) = g_{pq}(0, c)$, $\forall u \in \mathbf{R}$ (Proposition 1.6). We'll denote this constant $g_{pq}(0, c)$ by $\alpha \in \mathbf{C}$.

Evaluating the ‘‘diagonal’’ of the series in Eq. A7, at $(u, v) = (u, -\frac{l_0}{m_0}u + c)$, results in another constant (i.e. independent of u), say $\beta \in \mathbf{C}$.

Thus

$$h(u) \equiv \sum_{j=-\infty}^{\infty} \sum_{\substack{k=-\infty \\ k \neq j}}^{\infty} c_{jk} e^{2\pi i (j-k)l_0 u + k m_0 c} = \alpha - \beta. \quad (\text{A9})$$

So, setting $n = j - k$, we get

$$\alpha - \beta = \sum_{j=-\infty}^{\infty} \sum_{\substack{k=-\infty \\ k \neq j}}^{\infty} c_{jk} e^{2\pi i k m_0 c} \cdot e^{2\pi i (j-k)l_0 u}, \quad (\text{A10})$$

$$= \sum_{j=-\infty}^{\infty} \sum_{\substack{n=-\infty \\ n \neq 0}}^{\infty} c_{j, j-n} e^{2\pi i (j-n)m_0 c} \cdot e^{2\pi i n l_0 u}, \quad (\text{A11})$$

$$= \sum_{\substack{n=-\infty \\ n \neq 0}}^{\infty} d_n e^{2\pi i n l_0 u}, \quad (\text{A12})$$

where $d_n = \sum_{j=-\infty}^{\infty} c_{j, j-n} e^{2\pi i (j-n)m_0 c}$. Thus $h(u)$, which is a one dimensional Fourier series without a constant term, is a constant $\alpha - \beta$. This is only possible if $\alpha - \beta = 0$ and $d_n = 0$ whenever $n \neq 0$. This is again only possible if each $c_{j, j-n} = 0$ whenever $n \neq 0$. Thus $c_{jk} = 0$ whenever $j \neq k$.

Therefore, in the two dimensional Fourier series expansion of g_{pq} , only the terms with $j = k$ contribute. \square

Proposition 1.7 states that $g_{pq}(\mathbf{b})$ can be expressed as an effectively one-dimensional Fourier-series and follows from Proposition 1.6.

Note that Proposition 1.7 can also be stated using $e^{-2\pi i j \mathbf{b} \cdot \mathbf{s}_0}$ instead of $e^{2\pi i j \mathbf{b} \cdot \mathbf{s}_0}$ in which case Eq. A6 becomes

$$c_j = \mu \int_{-\frac{1}{2|m_0|}}^{\frac{1}{2|m_0|}} \int_{-\frac{1}{2|l_0|}}^{\frac{1}{2|l_0|}} g_{pq}(u, v) e^{2\pi i j (ul_0 + vm_0)} dudv, \quad (\text{A13})$$

with $\mu = |l_0||m_0|$ and $c_j \in \mathbf{R}$. It is also important to note that Proposition 1.7 assumes that $l_0 \neq 0$ and $m_0 \neq 0$. When either l_0 or m_0 is zero the derivation simplifies and becomes one dimensional. To avoid cluttering the derivation of the one dimensional case is not repeated here.

Proposition 1.8. Let $h(u, v) = \frac{1}{g_{pq}(u, v)}$, then $h(u, v)$ will be a differentiable Hermitian function if $g_{pq}(u, v) \neq 0, \forall u, v \in \mathbf{R}$. Moreover, $h(u + \frac{j}{l_0}, v + \frac{k}{m_0}) = h(u, v)$ and $h(u, -\frac{l_0}{m_0}u + c) = h(0, c), \forall j, k \in \mathbf{Z}$ and $u, v, c \in \mathbf{R}$.

Proof. To see that h has the same period as g_{pq} , notice that for any $j, k \in \mathbf{Z}$ we have $h(u + j\frac{1}{l_0}, v + k\frac{1}{m_0}) = (g_{pq}(u + j\frac{1}{l_0}, v + k\frac{1}{m_0}))^{-1} = g_{pq}(u, v)^{-1} = h(u, v)$. Similarly $h(u, u - \frac{l_0}{m_0}u + c) = h(0, c) \forall u, c \in \mathbf{R}$. To see that h is Hermitian, recall that complex conjugation satisfies $\overline{\frac{1}{z}} = \frac{1}{\overline{z}}$. Thus one computes $h(-u, -v) = \frac{1}{g_{pq}(-u, -v)} = \frac{1}{\overline{g_{pq}(u, v)}} = \overline{\left(\frac{1}{g_{pq}(u, v)}\right)} = \overline{h(u, v)}$. Finally, $h(u, v)$ is also differentiable since

$$\frac{\partial h(u, v)}{\partial u} = -\frac{\frac{\partial g_{pq}(u, v)}{\partial u}}{g_{pq}^2(u, v)}, \quad (\text{A14})$$

$$\frac{\partial h(u, v)}{\partial v} = -\frac{\frac{\partial g_{pq}(u, v)}{\partial v}}{g_{pq}^2(u, v)}, \quad (\text{A15})$$

$$(\text{A16})$$

exist ($g_{pq}(u, v) \neq 0$ by assumption). \square

Proposition 1.8 shows that the elements of $\mathcal{G}^\top(\mathbf{b})$ are also periodic, effectively one-dimensional, differentiable, Hermitian functions. Proposition 1.7 and Proposition 1.8 therefore implies that the elements of $\mathcal{G}^\top(\mathbf{b})$ also have a one-dimensional Fourier-series representation.

Lemma 1.9. Let \mathbf{A} be symmetric or Hermitian. If all principal submatrices having $k+1$ rows or $k+2$ rows are singular, the rank of \mathbf{A} does not exceed k (Perlis 1952).

Lemma 1.10. All 3×3 function-valued principal submatrices of $\mathcal{R}(\mathbf{b})$ are singular.

Proof. Due to the construction of $\mathcal{R}(\mathbf{b})$ all 3×3 function-valued principal submatrices of $\mathcal{R}(\mathbf{b})$ have the following form (see Lemma 1.14)

$$\begin{bmatrix} A_1 + A_2 & A_1 + A_2 e^{-2\pi i a \mathbf{b} \cdot \mathbf{s}_0} & A_1 + A_2 e^{-2\pi i A \mathbf{b} \cdot \mathbf{s}_0} \\ A_1 + A_2 e^{2\pi i a \mathbf{b} \cdot \mathbf{s}_0} & A_1 + A_2 & A_1 + A_2 e^{-2\pi i b \mathbf{b} \cdot \mathbf{s}_0} \\ A_1 + A_2 e^{2\pi i A \mathbf{b} \cdot \mathbf{s}_0} & A_1 + A_2 e^{2\pi i b \mathbf{b} \cdot \mathbf{s}_0} & A_1 + A_2 \end{bmatrix}, \quad (\text{A17})$$

where $A = a + b$ and $a, b \in \mathbf{N}$. The determinant of the matrix in Equation A17 is equal to zero (Kopp 2008). \square

Lemma 1.11. All 4×4 function-valued principal submatrices of $\mathcal{R}(\mathbf{b})$ are singular.

Proof. Due to the construction of $\mathcal{R}(\mathbf{b})$ all 4×4 function-valued principal submatrices of $\mathcal{R}(\mathbf{b})$ have the following

form (see Lemma 1.14)

$$\begin{bmatrix} A_1 + A_2 & A_1 + A_2e^{-kA} & A_1 + A_2e^{-kA} & A_1 + A_2e^{-kC} \\ A_1 + A_2e^{kA} & A_1 + A_2 & 1 + Ae^{-kb} & A_1 + A_2e^{-kB} \\ A_1 + A_2e^{kA} & A_1 + A_2e^{kb} & A_1 + A_2 & A_1 + A_2e^{-kC} \\ A_1 + A_2e^{kC} & A_1 + A_2e^{kB} & A_1 + A_2e^{kC} & A_1 + A_2 \end{bmatrix}, \quad (\text{A18})$$

where $k = 2\pi i \mathbf{b} \cdot \mathbf{s}_0$, $A = a + b$, $B = b + c$, $C = a + b + c$ and $a, b, c \in \mathbf{N}$. The determinant of the matrix in Equation A18 is equal to zero. \square

Definition 1.12. Let $|\mathbf{A}|_{d+1}$, where $\mathbf{A} \in \mathbf{Z}^{k \times k}$, be defined as $\sum_{p=1}^{n-1} a_{pp+1}$.

Definition 1.13. Let \mathcal{A} denote the set of all $m \times m$ principal sub-matrices of Φ , with $m = n - 1$. Let \mathcal{B} denote the set of all $k \times k$ principal sub-matrices of Φ , with $2 \leq k \leq n$.

Lemma 1.14. The array geometry matrix Φ has the following properties:

- (i) $\phi_{pp} = 0$ (diagonal),
- (ii) $\phi_{pq} \neq 0$ (non-diagonal),
- (iii) $\phi_{pq} > 0$; $\forall q > p$,
- (iv) $\phi_{pq} = -\phi_{qp}$,
- (v) $\gcd(\{\phi_{pq}\}_{q>p}) = 1$.
- (vi) $|\Phi|_{d+1} = \phi_{1n}$.
- (vii) $|\mathbf{B}|_{d+1} = b_{1k}$, $\forall \mathbf{B} \in \mathcal{B}$.

Proof. Property (i) is true since $\phi_{pp} = \phi_p - \phi_p = 0$. Properties (ii)–(iv) follow trivially from the assumption that the antenna positions $\{\mathbf{u}_p\}$ satisfy $\|\mathbf{u}_q\|_2 > \|\mathbf{u}_p\|_2 \forall q > p$. Property (v) is true since $\gcd(\{\phi_{pq}\}_{q>p}) = \gcd(\gcd(\{\phi_{1q}\}), \{\phi_{dq}\}_{q>d, d>1}) = \gcd(\gcd(\{\phi_q\}), \{\phi_{dq}\}_{q>d, d>1}) = \gcd(1, \{\phi_{dq}\}_{q>d, d>1}) = 1$. Property (vi) is true since $|\Phi|_{d+1} = \sum_{p=1}^{n-1} \phi_{pp+1} = \sum_{p=1}^{n-1} \phi_{p+1} - \phi_p = \phi_{1n}$.

Property (vii) can be proven using the following argument. Assume that $\mathbf{A}^j \in \mathcal{A}$ is obtained from Φ by deleting the j -th row and column from Φ (where j was chosen arbitrarily). When calculating $|\mathbf{A}^j|_{d+1}$ three separate cases arise,

- $1 < j < n$: $|\mathbf{A}^j|_{d+1} = \sum_{i=1}^{m-1} a_{ii+1}^j = \sum_{\substack{p=1 \\ p \neq j, j+1}}^{n-1} \phi_{pp+1} + \phi_{j,j+2} = \sum_{\substack{p=1 \\ p \neq j, j+1}}^{n-1} \phi_{pp+1} + \phi_{jj+1} + \phi_{j+1,j+2} = \phi_{1n} = a_{1m}$,
- $j = n$: $|\mathbf{A}^j|_{d+1} = \phi_{1n-1} = a_{1m}$,
- $j = 1$: $|\mathbf{A}^j|_{d+1} = \phi_{2n} = a_{1m}$.

The above shows that $|\mathbf{A}|_{d+1} = a_{1m} \forall \mathbf{A} \in \mathcal{A}$ (since j was chosen arbitrarily). Expanding the above derivation by using $1 < t \leq n - 2$ arbitrary deletions yields the required result. \square

Lemma 1.15. Let $\lambda(u, v)$ denote the largest eigenvalue of $\mathcal{R}(u, v)$ and $\mathbf{x}(u, v)$ its associated normalized eigenvector, then $\mathbf{x}(u, v) = \mathbf{x}(u + \frac{j}{l_0}, v + \frac{k}{m_0})$, $\forall j, k \in \mathbf{Z}$.

Proof. Notice that $\forall j, k \in \mathbf{Z}$

$$\mathcal{R}\left(u + \frac{j}{l_0}, v + \frac{k}{m_0}\right) = \mathcal{R}(u, v), \quad (\text{A19})$$

$$\lambda\left(u + \frac{j}{l_0}, v + \frac{k}{m_0}\right) = \lambda(u, v); \quad (\text{A20})$$

implying that $\mathbf{x}(u, v) = \mathbf{x}(u + \frac{j}{l_0}, v + \frac{k}{m_0})$, $\forall j, k \in \mathbf{Z}$. \square

Lemma 1.16. Let $\lambda(u, v)$ denote the largest eigenvalue of $\mathcal{R}(u, v)$ and $\mathbf{x}(u, v)$ its associated normalized eigenvector, then $\mathbf{x}(u, -\frac{l_0}{m_0}u + c) = \mathbf{x}(0, c)$, $\forall u, c \in \mathbf{R}$.

Proof. Notice that $\forall u, c \in \mathbf{R}$

$$\mathcal{R}\left(u, -\frac{l_0}{m_0}u + c\right) = \mathcal{R}(0, c), \quad (\text{A21})$$

$$\lambda\left(u, -\frac{l_0}{m_0}u + c\right) = \lambda(0, c), \quad (\text{A22})$$

implying that $\mathbf{x}(u, -\frac{l_0}{m_0}u + c) = \mathbf{x}(0, c)$, $\forall u, c \in \mathbf{R}$. \square

Lemma 1.17. Let $\lambda(u, v)$ denote the largest eigenvalue of $\mathcal{R}(u, v)$ and $\mathbf{x}(u, v)$ its associated normalized eigenvector. The real function $\lambda(u, v)$ and the function-valued vector $\mathbf{x}(u, v)$ are differentiable.

Proof. The parameter dimension of $\mathcal{R}(\mathbf{b})$ is effectively one, i.e. $\mathcal{R}_t(t) := \mathcal{R}(u, v)$, $\lambda_t(t) := \lambda(u, v)$ and $\mathbf{x}_t(t) := \mathbf{x}(u, v)$ with $t(u, v) := \mathbf{b} \cdot \mathbf{s}_0$. The entries of $\mathcal{R}_t(t)$ are analytic functions depending on $t \in \mathbf{R}$. Therefore, Rellich's theorem (Lemma 1.18) implies that $\lambda(t)$ and $\mathbf{x}(t)$ are analytic (the largest eigenvalue is simple-Eq. A1) and therefore also differentiable. We can therefore calculate

$$\frac{\partial \lambda(u, v)}{\partial u} = \frac{d\lambda_t(t(u, v))}{dt} \frac{\partial t(u, v)}{\partial u}, \quad (\text{A23})$$

$$\frac{\partial \lambda(u, v)}{\partial v} = \frac{d\lambda_t(t(u, v))}{dt} \frac{\partial t(u, v)}{\partial v}. \quad (\text{A24})$$

The above equations imply that the real function $\lambda(u, v)$ is differentiable. A similar argument can be used to prove that $\mathbf{x}(u, v)$ is also differentiable. \square

Lemma 1.18. (Rellich's Theorem) Let $\mathbf{A}(t) : \mathbf{R} \rightarrow \mathbf{C}^{n \times n}$ be a Hermitian function-valued matrix that depends on t analytically.

(i) The n roots of the characteristic polynomial of $\mathbf{A}(t)$ can be arranged so that each root $\lambda_j(t)$ for $j = 1, \dots, n$ is an analytic function of t .

(ii) There exists an eigenvector $\mathbf{x}_j(t)$ associated with $\lambda_j(t)$ for $j = 1, \dots, n$ satisfying

$$(a) \|\mathbf{x}_j(t)\|_2 = 1 \forall t \in \mathbf{R},$$

(b) $\mathbf{x}_j(t)$ is an analytic function-valued vector of t (Reed & Simon 1978; Lax 1996; Kılıç et al. 2011).

This paper has been typeset from a \LaTeX / \LaTeX file prepared by the author.