Ionospheric calibration of low frequency radio interferometric observations using the peeling scheme

I. Method description and first results

H. T. Intema¹, S. van der Tol², W. D. Cotton³, A. S. Cohen⁴, I. M. van Bemmelen¹, and H. J. A. Röttgering¹

¹ Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, The Netherlands
e-mail: intema@strw.leidenuniv.nl

² Electr. Engineering, Math. and Computer Science, Delft University of Technology, Delft, The Netherlands

³ National Radio Astronomy Observatory, Charlottesville, VA, USA

⁴ Naval Research Laboratory, Washington DC, USA

Received 6 October 2008 / Accepted 16 April 2009

ABSTRACT

Calibration of radio interferometric observations becomes increasingly difficult towards lower frequencies. Below ~ 300 MHz, spatially variant refractions and propagation delays of radio waves traveling through the ionosphere cause phase rotations that can vary significantly with time, viewing direction and antenna location. In this article we present a description and first results of SPAM (Source Peeling and Atmospheric Modeling), a new calibration method that attempts to iteratively solve and correct for ionospheric phase errors. To model the ionosphere, we construct a time-variant, 2-dimensional phase screen at fixed height above the Earth's surface. Spatial variations are described by a truncated set of discrete Karhunen-Loève base functions, optimized for an assumed power-law spectral density of free electrons density fluctuations, and a given configuration of calibrator sources and antenna locations. The model is constrained using antenna-based gain phases from individual self-calibrations on the available bright sources in the field-of-view. Application of SPAM on three test cases, a simulated visibility data set and two selected 74 MHz VLA data sets, yields significant improvements in image background noise (5–75 percent reduction) and source peak fluxes (up to 25 percent increase) as compared to the existing self-calibration and field-based calibration methods, which indicates a significant improvement in ionospheric phase calibration accuracy.

Key words. atmospheric effects – methods: numerical – techniques: interferometric

1. Introduction

Radio waves of cosmic origin are influenced by the Earth's atmosphere before detection at ground level. At low frequencies (LF, $\lesssim 300$ MHz), the dominant effects are refraction, propagation delay and Faraday rotation caused by the ionosphere (e.g., Thompson et al. 2001, TMS2001 from here on). For a ground-based interferometer (array from here on) observing a LF cosmic source, the ionosphere is the main source of phase errors in the visibilities. Amplitude errors may also arise under severe ionospheric conditions due to diffraction or focussing (e.g., Jacobson & Erickson 1992).

The ionosphere causes propagation delay differences between array elements, resulting in phase errors in the visibilities. The delay per array element (antenna from here on) depends on the line-of-sight (LoS) through the ionosphere, and therefore on antenna position and viewing direction. The calibration of LF observations requires phase corrections that vary over the field-of-view (FoV) of each antenna. Calibration methods that determine just one phase correction for the full FoV of each antenna (like self-calibration; e.g., see Pearson & Readhead 1984) are therefore insufficient.

Ionospheric effects on LF interferometric observations have usually been ignored for several reasons: (i) the resolution and sensitivity of the existing arrays were generally too poor

to be affected; (ii) existing calibration algorithms (e.g., self-calibration) appeared to give reasonable results most of the time; and (iii) a lack of computing power made the needed calculations prohibitively expensive. During the last 15 years, two large and more sensitive LF arrays have become operational: the VLA at 74 MHz (Kassim et al. 2007) and the GMRT at 153 and 235 MHz (Swarup 1991). Observations with these arrays have demonstrated that ionospheric phase errors are one of the main limiting factors for reaching the theoretical image noise level.

For optimal performance of these and future large arrays with LF capabilities (such as LOFAR, LWA and SKA), it is crucial to use calibration algorithms that can properly model and remove ionospheric contributions from the visibilities. Field-based calibration (Cotton et al. 2004) is the single existing ionospheric calibration & imaging method that incorporates direction-dependent phase calibration. This technique has been successfully applied to many VLA 74 MHz data sets, but is limited by design for use with relatively compact arrays.

In Sect. 2, we discuss ionospheric calibration in more detail. In Sect. 3, we present a detailed description of SPAM, a new ionospheric calibration method that is applicable to LF observations with relatively larger arrays. In Sect. 4, we present the first results of SPAM calibration on simulated and real VLA 74 MHz observations and compare these with results from

self-calibration and field-based calibration. A discussion and conclusions are presented in Sect. 5.

2. Ionosphere and calibration

In this section, we describe some physical properties of the ionosphere, the phase effects on radio interferometric observations and requirements for ionospheric phase calibration.

2.1. The Ionosphere

The ionosphere is a partially ionised layer of gas between ~ 50 and 1000 km altitude over the Earth's surface (e.g., Davies 1990). It is a dynamic, inhomogeneous medium, with electron density varying as a function of position and time. The state of ionization is mainly influenced by the Sun through photo-ionization at UV and short X-ray wavelengths and through injection of charged particles from the solar wind. Ionization during the day is balanced by recombination at night. The peak of the free electron density is located at a height around 300 km. The free electron column density along a LoS through the ionosphere is generally referred to as *total electron content*, or TEC. The TEC unit (TECU) is 10^{16} m^{-2} which is a typically observed value at zenith during nighttime.

The refraction and propagation delay are caused by a varying refractive index n of the ionospheric plasma along the wave trajectory. For a cold, collisionless plasma without magnetic field, n is a function of the free electron density n_e and is defined by (e.g., TMS2001)

$$n^2 = 1 - \frac{\nu_p^2}{\nu^2}, \quad (1)$$

with ν the radio frequency and ν_p the plasma frequency, given by

$$\nu_p = \frac{e}{2\pi} \sqrt{\frac{n_e}{\epsilon_0 m}}, \quad (2)$$

with e the electron charge, m the electron mass, ϵ_0 the vacuum permittivity. Typically, for the ionosphere, ν_p ranges from 1 – 10 MHz, but may locally rise up to ~ 200 MHz in the presence of sporadic E-layers (clouds of unusually high free electron density). Cosmic radio waves with frequencies below the plasma frequency are reflected by the ionosphere and do not reach the Earth's surface. For higher frequencies, the spatial variations in electron density cause local refractions of the wave (Snell's Law) as it travels through the ionosphere, thereby modifying the wave's trajectory. The total propagation delay, integrated along the LoS, results in a phase rotation given by

$$\phi^{\text{ion}} = -\frac{2\pi\nu}{c} \int (n - 1) dl, \quad (3)$$

with c the speed of light in vacuum. For frequencies $\nu \gg \nu_p$, this can be approximated by

$$\phi^{\text{ion}} \approx \frac{\pi}{c\nu} \int \nu_p^2 dl = \frac{e^2}{4\pi\epsilon_0 m c \nu} \int n_e dl, \quad (4)$$

where the integral over n_e on the right is the TEC along the LoS. Note that this integral depends on the wave's trajectory, and therefore on local refraction. Because the refractive index is frequency-dependent, the wave's trajectory changes with frequency. As a consequence, the apparent scaling relation $\phi^{\text{ion}} \propto \nu^{-1}$ from Eq. (4) is only valid to first order in frequency.

Although bulk changes in the large scale TEC (e.g., a factor of 10 increase during sunrise) have the largest amplitudes, the fluctuations on relatively small spatial scales and short temporal scales are most troublesome for LF interferometric observations. Most prominent are the traveling ionospheric disturbances (TIDs), a response to acoustic-gravity waves in the neutral atmosphere (e.g., van Velthoven 1990). Typically, medium-scale TIDs are observed at heights between 200 and 400 km, have wavelengths between 250 and 400 km, travel with near-horizontal velocities between 300 and 700 km h^{-1} in any direction and cause 1 – 5 percent variations in TEC (TMS2001).

The physics behind fluctuations on the shortest spatial and temporal scales is less well understood. Temporal and spatial behaviour may be coupled through quasi-frozen patterns that move over the area of interest with a certain velocity and direction (Jacobson & Erickson 1992). Typical variations in TEC are on the order of 0.1 percent, observed on spatial scales of tens of kilometers down to a few km, and time scales of minutes down to a few tens of seconds. The statistical behaviour of radio waves passing through this medium suggests the presence of a turbulent layer with a power-law spectral density of free electron density fluctuations $P_{n_e}(q) \propto q^{-\alpha}$ (e.g., TMS2001), with $q \equiv |q|$ the magnitude of the 3-dimensional spatial frequency. $P_{n_e}(q)$ is defined in units of electron density squared per spatial frequency. The related 2-dimensional structure function of the phase rotation ϕ of emerging radio waves from a turbulent ionospheric layer is given by

$$D_\phi = \langle [\phi(\mathbf{x}) - \phi(\mathbf{x} + \mathbf{r})]^2 \rangle \propto r^\gamma, \quad (5)$$

where $\mathbf{x}$ and $\mathbf{x} + \mathbf{r}$ are Earth positions, $r \equiv |\mathbf{r}|$ is the horizontal distance between these two points, $\langle \dots \rangle$ denotes the expected value and $\gamma = \alpha - 2$. For pure Kolmogorov turbulence, $\alpha = 11/3$, therefore $\gamma = 5/3$.

Using differential Doppler-shift measurements of satellite signals, van Velthoven (1990) found a power-law relation between spectral amplitude of small-scale ionospheric fluctuations and latitudinal wave-number with exponent $\alpha/2 = 3/2$. Combining with radio interferometric observations of apparent cosmic source shifts, van Velthoven derived a mean height for the ionospheric perturbations of 200 – 250 km. Through analysis of differential apparent movement of pairs of cosmic sources in the VLSS, Cohen & Röttgering (2008) find typical values for $\gamma/2$ of 0.50 during nighttime and 0.69 during daytime. Direct measurement of phase structure functions from different GPS satellites (van der Tol, *unpublished*) shows a wide distribution of values for γ that peaks at ~ 1.5 . On average, these results indicate the presence of a turbulent layer below the peak in the free electron density that has more power in the smaller scale fluctuations than in the case of pure Kolmogorov turbulence. Note that for individual observing times and locations, the behaviour of small-scale ionospheric fluctuations may differ significantly from this average.

2.2. Image plane effects

Interferometry uses the phase differences as measured on baselines to determine the angle of incident waves, and is therefore only sensitive to TEC differences. A baseline is sensitive to TEC fluctuations with linear sizes that are comparable to or smaller than the baseline length. At 75 MHz, a 0.01 TECU difference on a baseline causes a ~ 1 radian visibility phase error (Eq. (4)). Because the observed TEC varies with time, antenna position and viewing direction, visibility phases are distorted by time-varying differential ionospheric phase rotations.

An instantaneous spatial phase gradient over the array in the direction of a source causes an apparent position shift in the image plane (e.g., Cohen & Röttgering 2008), but no source deformation. If the spatial phase behaviour deviates from a gradient, this will also distort the apparent shape of the source. Combining visibilities with different time labels while imaging causes the image plane effects to be time-averaged. A non-zero time average of the phase gradient results in a source shift in the final image. Both a zero-mean time variable phase gradient and higher order phase effects cause smearing and deformation of the source image, and consequently a reduction of the source peak flux (see Cotton & Condon 2002 for an example). In the latter case, if the combined phase errors behave like Gaussian random variables, a point source in the resulting image experiences an increase of the source width and reduction of the source peak flux, but the total flux (the integral under the source shape) is conserved.

For unresolved sources, the *Strehl ratio* is defined as the ratio of observed peak flux over true peak flux. In case of Gaussian random phase errors, the Strehl ratio R is related to the rms phase error σ_ϕ by (Cotton et al. 2004)

$$R = \exp\left(-\frac{\sigma_\phi^2}{2}\right). \quad (6)$$

A larger peak flux is equivalent to a smaller rms phase error. This statement is more generally true, because all phase errors cause scattering of source power into sidelobes.

A change in the apparent source shape due to ionospheric phase errors leads to an increase in residual sidelobes after deconvolution. Deconvolution subtracts a time-averaged source image model from the visibility data at all time stamps. In the presence of time-variable phase errors, the mean source model deviates from the apparent, instantaneous sky emission and subtraction is incomplete. Residual sidelobes increase the rms background noise level and, due to its non-Gaussian character, introduce structure into the image that mimics real sky emission. In LF observations, due to the scaling relation of the dirty beam with frequency (width $\propto \nu^{-1}$), residual sidelobes around bright sources can be visible at significant distances from the source.

2.3. Ionospheric phase calibration

Lonsdale (2005) discussed four different regimes for (instantaneous) ionospheric phase calibration, depending on the different linear spatial scales involved. These scales are the array size A , the scale size S of ionospheric phase fluctuations and the projected size V of the field-of-view (FoV) at a typical ionospheric height. We use the term *compact* array when $A \ll S$ and *extended* array when $A \gtrsim S$. Note that these definitions change with ionospheric conditions, so there is no fixed linear scale that defines the difference between compact and extended. A schematic overview of the different regimes is given in Fig. 1.

The combination AV/S^2 is a measure of the complexity of ionospheric phase calibration. Both S and V depend on the observing frequency ν . For a power-law spectral density of free electron density fluctuations (see Sect. 2.1) S scales with ν , and for a fixed circular antenna aperture V scales with ν^{-1} . Therefore, AV/S^2 scales with ν^{-3} , signalling a rapid increase in calibration problems towards low frequencies.

Under *isoplanatic* conditions ($V \ll S$), the ionospheric phase error per antenna does not vary with viewing direction within the FoV, for both compact and large arrays (Lonsdale regimes 1 and 2, respectively). Phase-only self-calibration on

short enough time-scales is sufficient to remove the ionospheric phase errors from the visibilities.

Under *anisoplanatic* conditions ($V \gtrsim S$), the ionospheric phase error varies over the FoV of each antenna. A single phase correction per antenna is no longer sufficient. Self-calibration may still converge, but the resulting phase correction per antenna is a flux-weighted average of ionospheric phases across the FoV (see Sect. 3.1). Accurate self-calibration and imaging of individual very bright and relatively compact sources is therefore possible, even with extended arrays (see Gizani, Cohen & Kassim 2005 for an example). For a compact array (Lonsdale regime 3), the FoV of different antennas effectively overlap at ionospheric height. The LoS of different antennas towards one source run close and parallel through the ionosphere. For an extended array (Lonsdale regime 4), the FoV of different antennas may partially overlap at ionospheric height, but not necessarily. Individual LoS from widespread antennas to one source may trace very different paths through the ionosphere.

In regime 3, ionospheric phases behave as a spatial gradient over the array that varies with viewing direction. This causes the apparent position of sources to change with time and viewing direction, but no source deformation takes place. The 3-dimensional phase structure of the ionosphere can be effectively reduced to a 2-dimensional phase screen, by integrating the free electron density along the LoS (Eq. (4)). Radio waves that pass the virtual screen experience an instantaneous ionospheric phase rotation depending on the *pierce point* position (where the LoS pierces the phase screen). When assuming a fixed number of required ionospheric parameters per unit area of phase screen, calibration of a compact array requires a minimal number of parameters because each antenna illuminates the same part of the phase screen.

In regime 4, the dependence of ionospheric phase on antenna position and viewing direction is more complex. This causes source position shifts and source shape deformations that both vary with time and viewing direction. A 2-dimensional phase screen model may still be used, but only when the dominant phase fluctuations originate from a restricted height range $\Delta h \ll S$ in the ionosphere. The concept of a thin layer at a given height is attractive, because it reduces the complexity of the calibration problem drastically. When using an airmass function to incorporate a zenith angle dependence, the spatial phase function is in effect reduced to 2 spatial dimensions. Generally, a phase screen in regime 4 requires a larger number of model parameters than in regime 3, because the phase screen area illuminated by the total array is larger.

It is currently unclear under which conditions a 2-dimensional phase screen model becomes too inaccurate to model the ionosphere in regime 4. For very long baselines or very severe ionospheric conditions, a full 3-dimensional ionospheric phase model may be required, where ionospheric phase corrections need to be determined by ray-tracing. Such a model is likely to require many more parameters than can be extracted from radio observations alone. To first order, it may be sufficient to extend the phase screen model with some form of height-dependence. Examples of such extensions are the use of several phase screens at different heights (Anderson 2006) or introducing smoothly varying partial derivatives of TEC or phase as a function of zenith angle (Noordam et al. *in preparation*).

Calibration needs to determine corrections on sufficiently short time scales to track the ionospheric phase changes. The phase rate of change depends on the intrinsic time variability of the TEC along a given LoS and on the speed of the LoS from the array antennas through the ionosphere while tracking

Fig. 1. Schematic overview of the different calibration regimes as discussed by Lonsdale (2005). For clarity, only two spatial dimensions and one calibration time interval are considered. In this overview, the array is represented by three antennas at ground level, looking through the ionospheric electron density structure (grey bubbles) with individual fields-of-view (red, green and blue areas). Due to the relatively narrow primary beam patterns in regimes 1 and 2 (*top left* and *top right*, respectively), each individual antenna “sees” an approximately constant TEC across the FoV. The relatively wide primary beam patterns in regimes 3 and 4 (*bottom left* and *bottom right*, respectively) causes the antennas to “see” TEC variations across the FoV. For the relatively compact array configurations in regimes 1 and 3, the TEC variation across the array for a single viewing direction within the FoV is approximately a gradient. For the relatively extended array configurations in regimes 2 and 4, the TEC variation across the array for a single viewing direction differs significantly from a gradient. The consequences for calibration of the array are discussed in the text.

a cosmic source. The latter may range up to $\sim 100 \text{ km h}^{-1}$ at 200 km height. The exact requirements on the time resolution of the calibration are yet to be determined. In principle, the time-variable ionospheric phase distortions needs to be sampled at least at the Nyquist frequency. However, during phase variations of large amplitude ($\gg 1$ radian), 2π radian phase winding introduces periodicity on much shorter time scales. To successfully unwrap phase winds, at least two corrections per 2π radian phase change are required.

2.4. Proposed and existing ionospheric calibration schemes

Schwab (1984) and Subrahmanya (1991) have proposed modifications to the self-calibration algorithm to support direction-dependent phase calibration. Both methods discuss the use of a spatial grid of interpolation nodes (additional free parameters) to characterize the spatial variability of the ionospheric phase error. Schwab suggests to use a different set of nodes per antenna, while Subrahmanya suggests to combine these sets by positioning them in a quasi-physical layer at fixed height above the Earth’s surface (this to reduce the number of required nodes when the FoVs from different antennas overlap at

ionospheric height). Neither of both proposed methods have been implemented.

Designed to operate in Lonsdale regime 3, field-based calibration by Cotton et al. (2004) is the single existing implementation of a direction-dependent ionospheric phase calibration algorithm. Typically, for each time interval of 1–2 min of VLA 74 MHz data, the method measures and converts the apparent position shift of 5–10 detectable bright sources within the FoV into ionospheric phase gradients over the array. To predict phase gradients in arbitrary viewing directions for imaging of the full FoV, an independent phase screen per time interval is fitted to the measured phase gradients. The phase screen is described by a 5 term basis of Zernike polynomials (up to second order, excluding the constant zero order).

Field-based calibration has been used to calibrate 74 MHz VLA observations, mostly in B-configuration (e.g., Cohen et al. 2007) but also several in A-configuration (e.g., Cohen et al. 2003, 2004). Image plane comparison of field-based calibration against self-calibration shows an overall increase of source peak fluxes (in some cases up to a factor of two) and reduction of residual sidelobes around bright sources, a clear indication of improved phase calibration over the FoV

(Cotton & Condon 2002). The improved overall calibration performance sometimes compromises the calibration towards the brightest source.

Zernike polynomials are often used to describe aberrations in optical systems, because lower order terms match well with several different types of wavefront distortions, and the functions are an orthogonal set on the circular domain of the telescope pupil. Using Zernike polynomials to describe an ionospheric phase screen may be less suitable, because they are not orthogonal on the discrete domain of pierce points, diverge when moving away from the field center and have no relation to ionospheric image aberrations (except for first order, which can model a large scale TEC gradient). Non-orthogonality leads to interdependence between model parameters, while divergence is clearly non-physical and leads to undesirable extrapolation properties.

For extended LF arrays or more severe ionospheric conditions, the ionospheric phase behaviour over the array for a given viewing direction is no longer a simple gradient. Under these conditions, performance of field-based calibration degrades. For the 74 MHz VLA Low-frequency Sky Survey (VLSS; Cohen et al. 2007), field-based calibration was unable to calibrate the VLA in B-configuration for about 10–20% of the observing time due to severe ionospheric conditions. Observing at 74 MHz with the ~3 times larger VLA A-configuration leads to a relative increase in the failure rate of field-based calibration. This is to be expected, as the larger array size results in an increased probability for the observations to reside in Lonsdale regime 4.

The presence of higher order phase structure over the array in the direction of a calibrator requires an antenna-based phase calibration rather than a source position shift to measure ionospheric phases. The calibration methods proposed by Schwab and Subrahmanya (see above) do allow for higher order phase corrections over the array and could, in principle, handle more severe ionospheric conditions. An alternative approach is to use the *peeling* technique (Noordam 2004), which consists of sequential self-calibrations on individual bright sources in the FoV. This yields per source a set of time-variable antenna-based phase corrections and a source model. Because the peeling corrections are applicable to a limited set of viewing directions, they need to be interpolated in some intelligent way to arbitrary viewing directions while imaging the full FoV. Peeling is described in more detail in Sect. 3.3.

Noordam (2004) has proposed a “generalized” self-calibration method for LOFAR (e.g., Röttgering et al. 2006) that includes calibration of higher order ionospheric phase distortions. Similar to ‘classical’ self-calibration, instrumental and environmental (including ionospheric) parameters are estimated by calibration against a sky brightness model. Sky model and calibration parameters are iteratively updated to converge to some final result. Uniqueness of the calibration solution is controlled by putting restrictions on the time-, space- and frequency behaviour of the fitted parameters. The effects of the ionosphere are modeled in a Minimum Ionospheric Model (MIM, Noordam *in preparation*), which is yet to be defined in detail. The philosophy of the MIM is to use a minimal number of physical assumptions and free parameters to accurately reproduce the observed effects of the ionosphere on the visibilities for a wide-as-possible range of ionospheric conditions. The initial MIM is to be constrained using peeling corrections.

3. Method

SPAM, an abbreviation of “Source Peeling and Atmospheric Modeling”, is the implementation of a new ionospheric

calibration method, combining several concepts from proposed and existing calibration methods. SPAM is designed to operate in Lonsdale regime 4 and can therefore also operate in regimes 1 to 3. It uses the calibration phases from peeling sources in the FoV to constrain an ionospheric phase screen model. The phase screen mimics a thin turbulent layer at a fixed height above the Earth’s surface, in concordance with the observations of ionospheric small-scale structure (Sect. 2.1). The main motivation for this work was to test several aspects of ionospheric calibration on existing VLA and GMRT data sets on viability and qualitative performance, and thereby support the development of more advanced calibration algorithms for future instruments such as LOFAR.

Generally, the instantaneous ionosphere can only be sparsely sampled, due to the non-uniform sky distribution of a limited number of suitable calibrators and an array layout that is optimized for UV-coverage rather than ionospheric calibration. To minimize the error while interpolating to unsampled regions, an optimal choice of base functions for the description of the phase screen is of great importance. Based on the work by van der Tol & van der Veen (2007), we use the discrete Karhunen-Loève (KL) transform to determine an optimal set of base ‘functions’ to describe our phase screen. For a given pierce point layout and an assumed power-law slope for the spatial structure function of ionospheric phase fluctuations (see Sect. 2.1), the KL transform yields a set of base vectors with several important properties: (i) the vectors are orthogonal on the pierce point domain; (ii) truncation of the set (reduction of the model order) gives a minimal loss of information; (iii) interpolation to arbitrary pierce point locations obeys the phase structure function; and (iv) spatial phase variability scales with pierce point density, i.e., most phase screen structure is present in the vicinity of pierce points, while it converges to zero at infinite distance. More detail on this phase screen model is given in Sect. 3.4.

Because the required calibration time resolution is still an open issue, and the SPAM model does not incorporate any restrictions on temporal behaviour, independent phase screens are determined at the highest possible time resolution (which is the visibility integration time resolution).

SPAM calibration can be separated in a number of functional steps, each of which is discussed in detail in the sections to follow. The required input is a spectral-mode visibility data set that has flux calibration and bandpass calibration applied, and radio frequency interference (RFI) excised (see Lazio, Kassim & Perley 2005 or Cohen et al. 2007 for details). The SPAM recipe consists of the following steps:

1. obtain and apply instrumental calibration corrections for phase (Sect. 3.1);
2. obtain an initial model of the apparent sky, together with an initial ionospheric phase calibration (Sect. 3.2);
3. subtract the sky model from the visibility data while applying the phase calibration. Peel apparently bright sources (Sect. 3.3);
4. fit an ionospheric phase screen model to the peeling solutions (Sect. 3.4);
5. apply the model phases on a facet-to-facet basis during re-imaging of the apparent sky (Sect. 3.5).

Steps 3 to 5 define the SPAM calibration cycle, as the image produced in step 5 can serve as an improved model of the apparent sky in step 3.

The scope of applications for SPAM is limited by a number of assumptions that were made to simplify the current implementation:

- the ionospheric inhomogeneities that cause significant phase distortions are located in a single, relatively narrow height range;
- there exists a finitely small angular patch size, which can be much smaller than the FoV of an individual antenna, over which the ionospheric phase contribution is effectively constant. Moving from one patch to neighbouring patches results in small phase transitions ($\ll 1$ radian);
- there exists a finitely small time range, larger than the integration time interval of an observation, over which the apparent ionospheric phase change for any of the array antennas along any line-of-sight is much smaller than a radian;
- the bandwidth of the observations is small enough to be effectively monochromatic, so that the ionospheric dispersion of waves within the frequency band is negligible;
- within the given limitations on bandwidth and integration time, the array is sensitive enough to detect at least a few (≥ 5) sources within the target FoV that may serve as phase calibrators;
- the ionospheric conditions during the observing run are such that self-calibration is able to produce a good enough initial calibration and sky model to allow for peeling of multiple sources. This might not work under very bad ionospheric conditions, but for the applications presented in this article it proved to be sufficient;
- after each calibration cycle (steps 3 to 5), the calibration and sky model are equally or more accurate than the previous. This implies convergence to a best achievable image;
- the instrumental amplitude and phase contributions to the visibilities, including the antenna power patterns projected onto the sky towards the target source, are constant over the duration of the observing run.

SPAM does not attempt to model the effects of ionospheric Faraday rotation on polarization products, and is therefore only applicable to intensity measurements (stokes I).

In our implementation we have focussed on functionality rather than processing speed. In its current form, SPAM is capable of processing quite large offline data sets, but is not suitable for real-time processing as is required for LOFAR calibration. SPAM relies heavily on functionality available in NRAO's Astronomical Image Processing System (AIPS; e.g., Bridle & Greisen 1994). It consists of a collection of Python scripts that accesses AIPS tasks, files and tables using the ParselTongue interface (Kettenis et al. 2006). Two main reasons to use AIPS are its familiarity and proven robustness while serving a large group of users over a 30 year lifetime, and the quite natural way by which the ionospheric calibration method is combined with polyhedron imaging (Perley 1989a; Cornwell & Perley 1992). SPAM uses a number of 3rd party Python libraries, like *scipy*, *numpy* and *matplotlib* for math and matrix operations and plotting. For non-linear least squares fitting of ionospheric phase models, we have adopted a Levenberg-Marquardt solver (LM; e.g., Press et al. 1992) based on IDL's MPFIT package (Markwardt 2008).

3.1. Instrumental phase calibration

Each antenna in the array adds an instrumental phase offset to the recorded signal before correlation. At low frequencies, changes in the instrumental signal path length (e.g., due to temperature

induced cable length differences) are very small compared to the wavelength, therefore instrumental phase offsets are generally stable over long time periods (hours to days). SPAM requires removal of the instrumental phase offsets from the visibilities prior to ionospheric calibration.

Instead of directly measuring the sky intensity $I(l, m)$ as a function of viewing direction cosines (l, m) , an interferometer measures an approximate Fourier transform of the sky intensity. For a baseline consisting of antennas i and j , the perfect response to all visible sky emission for a single time instance and frequency is given by the measurement equation (ME) for visibilities (e.g., TMS2001):

$$V_{ij} = \iint I(l, m) e^{-2\pi J[u_{ij}l + v_{ij}m + w_{ij}(n-1)]} \frac{dl dm}{n}, \quad (7)$$

where $J = \sqrt{-1}$, $n = \sqrt{1 - l^2 - m^2}$, u_{ij} and v_{ij} are baseline coordinates in the UV plane (expressed in wavelengths) parallel to l and m , respectively, and w_{ij} is the perpendicular baseline coordinate along the LoS towards the chosen celestial *phase tracking center* at $(l, m) = (0, 0)$. In practice, these measurements are modified with predominantly antenna-based complex gain factors a_i that may vary with time, frequency, antenna position and viewing direction. This modifies the ME into

$$\hat{V}_{ij} = \iint a_i(l, m) a_j^*(l, m) I(l, m) e^{-2\pi J[u_{ij}l + v_{ij}m + w_{ij}(n-1)]} \frac{dl dm}{n}. \quad (8)$$

Determination of the gain factors is generally referred to as *calibration*. When known, only gain factors that do not depend on viewing direction can be removed from the visibility data prior to image reconstruction by applying the calibration:

$$V_{ij} = (a_i a_j^*)^{-1} \hat{V}_{ij} \quad (9)$$

This operation is generally not possible for gain factors that do depend on viewing direction, because these gain factors cannot be moved in front of the integral in Eq. (8). One may still choose to apply gain corrections for a single viewing direction (e.g. to image a particular source), but the accuracy of imaging and deconvolution of other visible sources will degrade when moving away from the selected viewing direction. A solution for wide-field imaging and deconvolving in the presence of direction-dependent gain factors is discussed in Sect. 3.5.

The standard approach for instrumental phase calibration at higher frequencies is to repeatedly observe a bright (mostly unresolved) source during an observing run. Antenna-based gain phase corrections $g_i \approx a_i^{-1}$ are estimated by minimizing the weighted difference sum S between observed visibilities $\hat{V}_{ij}$ and source model visibilities $V_{ij}^{\text{model}} \approx V_{ij}$ (e.g., TMS2001; implemented in AIPS task CALIB):

$$S = \sum_i \sum_{j>i} W_{ij} \|V_{ij}^{\text{model}} - g_i g_j^* \hat{V}_{ij}\|^p, \quad (10)$$

with W_{ij} the visibility weight (reciprocal of the uncertainty in the visibility measurement), $g_i = e^{J\phi_i^{\text{cal}}}$ and p the power of the norm (typically 1 or 2). The source model visibilities V_{ij}^{model} are calculated using Eq. (7) with $I(l, m) = I^{\text{model}}(l, m)$. The phase corrections ϕ_i^{cal} consist of an instrumental and an atmospheric part. The corrections are interpolated in time and applied to the

target field visibilities, under the assumptions that the instrumental and atmospheric phase offsets vary slowly in time, and that the atmospheric phase offsets in the direction of the target are equal to those in the direction of the calibrator.

At low frequencies, there are two complicating factors for the standard approach: (i) the FoV around the calibrator source is large and includes many other sources; and (ii) the ionospheric phase offset per antenna changes significantly with time and viewing direction. The former can be overcome by choosing a very bright calibrator source with a flux that dominates over the combined flux of all other visible sources on all baselines. For the VLSS (Cohen et al. 2007), the 17 000 Jy of Cygnus A was more than sufficient to dominate over the total apparent flux of 400–500 Jy in a typical VLSS field. The latter requires filtering of the phase corrections to extract only the instrumental part, which is then applied to the target field visibilities.

For SPAM, we have adopted an instrumental phase calibration method that is very similar to the procedure used for field-based calibration (Cotton et al. 2004). Antenna-based phase corrections are obtained on the highest possible time resolution by calibration on a very bright source k using the robust L1 norm (Eq. (10) with $p = 1$; Schwab 1981). A phase correction ϕ_{ikn}^{cal} for antenna i at time interval n consist of several contributions:

$$\phi_{ikn}^{\text{cal}} = \phi_i^{\text{instr}} + \phi_{ikn}^{\text{ion}} - \phi_{rkn} - \phi_{ikn}^{\text{ambig}}, \quad (11)$$

where the instrumental and ionospheric phase corrections, ϕ_i^{instr} and ϕ_{ikn}^{ion} respectively, are assumed to be constant resp. vary with time and antenna position over the observing run. The other right-hand terms are the phase offset $\phi_{rkn} = \phi_r^{\text{instr}} + \phi_{rkn}^{\text{ion}}$ of an arbitrarily chosen reference antenna $r \in \{i\}$, and the phase ambiguity term $\phi_{ikn}^{\text{ambig}} = 2\pi N_{ikn}$ with integer N_{ikn} that maps ϕ_{ikn}^{cal} into the $[0, 2\pi)$ domain.

The antenna-based phase corrections are split into instrumental and ionospheric parts on the basis of their temporal and spatial behaviour. The phase corrections are filtered by iterative estimation of invariant instrumental phases (together with the phase ambiguities) and time- and space-variant ionospheric phases. The instrumental phases are estimated by robust averaging ($+3\sigma$ rejection) over all time intervals n :

$$\tilde{\phi}_i^{\text{instr}} = \langle (\phi_{ikn}^{\text{cal}} - \tilde{\phi}_i^{\text{ion}}) \bmod 2\pi \rangle_n. \quad (12)$$

The phase ambiguity estimates follow from

$$\tilde{\phi}_{ikn}^{\text{ambig}} = 2\pi \text{round} \left(\left[\tilde{\phi}_i^{\text{instr}} + \tilde{\phi}_i^{\text{ion}} - \phi_{ikn}^{\text{cal}} \right] / 2\pi \right), \quad (13)$$

where the round() operator rounds a number to the nearest integer value. The instrumental phase offset of the reference antenna is arbitrarily set to zero. The ionospheric phases are constrained by fitting a time-varying spatial gradient $\mathbf{G}_{kn}$ to the phases over the array. The gradient fit consists of an initial estimate directly from the calibration phase corrections, followed by a refined fit by using the LM solver to minimize

$$\chi_{kn}^2 = \sum_i \left[\left(\phi_{ikn}^{\text{cal}} - \tilde{\phi}_i^{\text{instr}} + \tilde{\phi}_{ikn}^{\text{ambig}} \right) - \underbrace{\mathbf{G}_{kn} \cdot (\mathbf{x}_i - \mathbf{x}_r)}_{\phi_{ikn}^{\text{ion}}} \right]^2, \quad (14)$$

where $\mathbf{x}_i$ is the position of antenna i . The ionospheric phase offset of the reference antenna is arbitrarily set to zero, which makes it a pivot point over which the phase gradient rotates. Higher order ionospheric effects are assumed to average to zero in Eq. (12).

3.2. Initial phase calibration and initial sky model

The instrumental phase calibration method described in Sect. 3.1 assumes that the time-averaged ionospheric phase gradient over the array in the direction of the bright phase calibrator is zero. Any non-zero average is absorbed into the instrumental phase estimates, causing a position shift of the whole target field and thereby invalidating the astrometry. Before entering the calibration cycle (Sects. 3.3–3.5), SPAM requires restoration of the astrometry and determination of an initial sky model and initial ionospheric calibration.

To restore the astrometry, the instrumentally corrected target field data from Sect. 3.1 is phase calibrated against an apparent sky model (AIPS task CALIB). The default is a point source model, using NVSS catalog positions (Condon et al. 1994, 1998), power-law interpolated fluxes from NVSS and WENSS/WISH catalogs (Rengelink et al. 1997) and a given primary beam model. To preserve the instrumental phase calibration as obtained in Sect. 3.1 during further processing, time-variable phase corrections resulting from calibration steps in this and the following sections are stored in a table (AIPS SN table) rather than applied directly to the visibility data. The sky model calibration is followed by wide-field imaging (AIPS task IMAGR) and several rounds of phase-only self-calibration (CALIB and IMAGR) at the highest possible time resolution, yielding the initial sky model and initial phase calibration.

For wide-field imaging with non-coplanar arrays, the standard imaging assumptions that the relevant sky area is approximately flat and the third baseline coordinate (w -term in Eq. (7)) is constant across the FoV are no longer valid. To overcome this, SPAM uses the polyhedron method (Perley 1989a; Cornwell & Perley 1992) that divides the large FoV into a hexagonal grid of small, partially overlapping *facets* that individually do satisfy the assumptions above (AIPS task SETFC). Additional facets are centered on relatively bright sources inside and outside the primary beam area to reduce image artefacts due to pixellation (Perley 1989b; Briggs & Cornwell 1992; Briggs 1995; Voronkov & Wieringa 2004; Cotton & Uson 2008).

The Cotton-Schwab algorithm (Schwab 1984; Cotton 1999; Cornwell et al. 1999) is a variant of CLEAN deconvolution (Högbom 1974; Clark 1980) that allows for simultaneous deconvolution of multiple facets, using a different dirty beam for each facet. *Boxes* are used to restrict CLEANing to real sky emission, making sure that sources are deconvolved in the nearest facet only (CLEAN model components are stored in facet-based AIPS CC tables). After deconvolution, the CLEAN model is restored to the relevant residual facets (AIPS task CCRES) using a CLEAN beam, and the facets are combined to form a single image of the full FoV (AIPS task FLATN).

3.3. Peeling

To construct a model of ionospheric phase rotations in arbitrary viewing directions within the FoV, SPAM requires measurements in as many directions as possible. When no external sources of ionospheric information are available, the target field visibilities themselves need to be utilized. Calibration on individual bright sources in the FoV can supply the required information, even in the presence of higher order phase structure over the array. After instrumental phase offsets are removed, phase calibration corrections are an relative measure of ionospheric phase:

$$\phi_{ikn}^{\text{cal}} = \phi_{ikn}^{\text{ion}} - \phi_{rkn}^{\text{ion}} - \phi_{ikn}^{\text{ambig}}, \quad (15)$$

where we used Eq. (11) with $\phi_i^{\text{instr}} = \phi_r^{\text{instr}} = 0$.

SPAM uses the peeling technique (Noordam 2004) to obtain phase corrections in different viewing directions. Peeling consists of self-calibration on individual sources, yielding per source a set of time-variable antenna-based phase corrections and a source model, after which the source model is subtracted from the visibility data set while temporarily applying the phase corrections (AIPS tasks SPLIT, UVSUB and CLINV/SPLIT).

For peeling to converge, the source needs to be the dominant contributor of flux to the visibilities on all baselines. Especially at low frequencies, the presence of many other sources in the large FoV adds considerable noise to the peeling phase corrections. To suppress this effect, the following steps are performed: (i) The best available model of the apparent sky is subtracted from the visibility data while temporarily applying the associated phase calibration(s). The initial best available model and associated phase calibration is the self-calibration output of Sect. 3.2. Individual source models are added back before peeling; (ii) Sources are peeled in decreasing flux order to suppress the effect of brighter sources on the peeling of fainter sources; (iii) Calibration only uses visibilities with projected baseline lengths longer than a certain threshold. This excludes the high “noise” in the visibilities near zero-length baselines from the coherent flux contribution of imperfectly subtracted sources.

The radio sky can be approximated by a discrete number of isolated, invariant sources of finite angular extent. Visibilities in the ME (Eq. (7)) for a single integration time n can therefore be split into a linear combination of contributions from individual sources k :

$$V_{ijn} = \sum_k V_{ijkn} = \sum_k \int \int I_k(l, m) \times e^{-2\pi J[u_{ijn}l + v_{ijn}m + w_{ijn}(n-1)]} \frac{dl dm}{n}. \quad (16)$$

The subtraction of all but the peeling source k' from the measured visibilities in step (i) above can be described as

$$\hat{V}_{ijk'n} \approx \hat{V}_{ijn} - \sum_{k \neq k'} (g_{ikn} g_{jkn}^\dagger)^{-1} V_{ijkn}^{\text{model}}, \quad (17)$$

with $g_{ikn} = g_i(l_k, m_k, t_n) = e^{J\phi_{ikn}^{\text{cal}}}$ the best available calibration in the viewing direction of source k , and V_{ijkn}^{model} the visibilities that are derived from the best available model I_{ijk}^{model} of source k . The peeling itself consists of iterative calibration and imaging steps of the peeling source k' . The calibration (Eq. (10) with $p = 1$) updates the antenna gain corrections g_{ikn} by minimizing

$$S_n = \sum_i \sum_{j>i} w_{ijn} \|V_{ijk'n}^{\text{model}} - g_{in} g_{jn}^\dagger \hat{V}_{ijk'n}\|, \quad (18)$$

while the imaging step updates $I_{ijk'}^{\text{model}}$ and therefore $V_{ijk'n}^{\text{model}}$.

In practice, due to incompleteness of the sky model and inaccuracies in the phase calibration, there will always remain some contaminating source flux in the visibilities while peeling. Complemented with system noise, sky noise, residual RFI and other possible sources of noise, the noise in the visibilities propagates into the phase corrections from the peeling process.

Absolute astrometry is not conserved during peeling, because self-calibration allows antenna-based phase corrections to vary without constraint. In subsequent peeling cycles, small non-zero phase gradients in the phase residuals after calibration can cause the source model to wander away from its true position. In SPAM, astrometry errors are minimized by re-centering the

source model to its true (catalog) position before calibration in each self-calibration loop. By default, SPAM re-centers the peak of the model flux to the nearest bright point source position in the NVSS catalog (Condon et al. 1994, 1998). It is recommended to visually check the final peeling source images for possible mismatches with the catalog (e.g., in case of double sources or sources with a spatially varying spectral index).

While peeling, SPAM attempts to calibrate sources on the highest possible time resolution, which is the visibility time grid. The noise in the resulting phase corrections depends on the signal-to-noise ratio (SNR) of the source flux in the visibilities. To increase the number of peeling sources and limit the phase noise in case of insufficient SNR, SPAM is allowed to increase the calibration time interval beyond the visibility integration time up to an arbitrary limit. Through image plane analysis, SPAM estimates the required calibration time-interval per source:

$$n_t = \left(\frac{\sigma_L}{\alpha S_p} \right)^2 N_t, \quad (19)$$

where n_t is the required number of integration times in a calibration interval, N_t is the total number of integration times within the observation, α is the minimum required SNR per integration time (a tweakable parameter that sets the balance between the SNR and the time resolution of the peeling phase corrections), and S_p and σ_L are the measured source peak flux and local background noise level in the image. For a fixed upper limit on the calibration time interval, an increase in α results in a decrease in the number of peeling sources. For $n_t < 1$, phase corrections are determined on the visibility time grid. For $n_t > 1$, a spline is used to resample the phase corrections per antenna in time onto the visibility time grid.

Apart from SNR issues, the number of sources that can be peeled is fundamentally limited by the available number of independent visibility measurements. When peeling N_s sources, self-calibration fits $N_s(N_a - 1)$ phase solutions per calibration time interval to the visibility data, where N_a is the number of antennas. For self-calibration to converge to a unique combination of phase solutions and source model, this number needs to be much smaller than the number of independent visibility measurements. The maximum of visibilities measurements that is available in one calibration time interval is given by $N_c \langle n_t \rangle N_a (N_a - 1)/2$, with N_c the number of frequency channels and $\langle n_t \rangle$ the average number of visibility integration times in a calibration interval. In the ideal case, when we assume that each visibility is an independent measurement, the determination of antenna-based phase corrections for all peeling sources is well constrained if

$$N_s \ll \frac{N_a N_c \langle n_t \rangle}{2}. \quad (20)$$

The applications presented in this article do satisfy this minimal condition (see Sect. 4).

Equation (20) is equivalent to stating that the number of degrees-of-freedom (DoF; the difference between the number of independent measurements and the number of model parameters) should remain a large positive number. Correlation between visibilities over frequency and time may reduce the number of independent measurements drastically, thereby also reducing the number of DoFs. The exact number of DoFs for any data set is hard to quantify. When this number becomes too low, the data is “over-fitted” (e.g., Bhatnagar et al. 2008), which could result in an artificial reduction of both the image background noise level and source flux that is not represented in the self-calibration

model (Wieringa 1992). Although we have found no evidence of this effect occurring in the applications presented in this article, the SPAM user should be cautious not to peel too many sources. In case of a high number of available peeling sources, one can choose a subset with a sufficiently dense spatial distribution over the FoV (e.g., one source per isoplanatic patch; see Sect. 3.5).

3.4. Ionospheric phase screen model

The phase corrections that are obtained by peeling several bright sources in the FoV (Sect. 3.3) are only valid for ionospheric calibration in a limited patch of sky around each source. To correct for ionospheric phase errors over the full FoV during wide-field imaging and deconvolution, SPAM requires a model that predicts the phase correction per antenna for arbitrary viewing directions.

SPAM constructs a quasi-physical phase screen model that attempts to accurately reproduce and interpolate the measured ionospheric phase rotations (or more accurately: the peeling phase corrections). The phase screen is determined independently for each visibility time stamp, therefore we drop the n -subscript in the description below. Figure 2 is a schematic overview of the geometry of ionospheric phase modeling in SPAM. The ionosphere is represented by a curved phase screen at a fixed height h above the Earth's surface, compliant to the WGS84 standard (NIMA 1997). The total phase rotation experienced by a ray of radio emission traveling along a LoS through the ionosphere is represented by an instantaneous phase rotation $\phi^{\text{ion}}(\mathbf{p}, \zeta)$ on passage through the phase screen that is a function of pierce point position $\mathbf{p}$ and zenith angle ζ . For a thin layer ($\Delta h \ll S$; see Sect. 2.3), the dependence of ϕ^{ion} on ζ can be represented by a simple airmass function, so that

$$\phi^{\text{ion}}(\mathbf{p}, \zeta) = \frac{\phi^{\text{ion}}(\mathbf{p})}{\cos(\zeta)}. \quad (21)$$

SPAM uses an angular local longitude/latitude coordinate system to specify $\mathbf{p}$, relative to the central pierce point from array center to field center. For the applications presented in this article, the angular distances between pierce points over the relevant ionospheric domain are all < 5 degrees, which effectively makes the pierce point vector $\mathbf{p}$ a 2-dimensional cartesian vector.

The 2-dimensional phase screen $\phi^{\text{ion}}(\mathbf{p})$ is defined on a set of KL base vectors, generated from the instantaneous pierce point configuration $\{\mathbf{p}_{ik}\}$ and an assumed power-law shape for the phase structure function (Sect. 2.1). The KL base vector generation and interpolation is based on the work by van der Tol & van der Veen (2007) and is described in detail in Appendix A. The phase screen model requires one free parameter per KL base vector. The initial complete set of KL base vectors is arbitrarily reduced in order by selecting a subset based on statistical relevance (principle component analysis). This reduces the effect of noise in the peeling solutions on the model accuracy and simultaneously limits the number of model parameters. However, the subset should still be large enough to accurately reproduce the peeling phase corrections. Per visibility time stamp, the KL base vectors are stored for later use during imaging (for this purpose, we mis-use the AIPS OB table). As an example, the first six interpolated KL base vectors for a single configuration of ionospheric pierce points are plotted in Fig. 3.

The peeling phase corrections ϕ_{ik}^{cal} are interpreted to be relative measurements of the absolute ionospheric phase screen model $\phi^{\text{ion}}(\mathbf{p}, \zeta)$ which may be determined up to a constant. The model parameters are determined by minimizing the differences between the observed and the model phases using the

Fig. 2. Schematic overview of the SPAM thin ionospheric phase screen model geometry. For clarity, only two spatial dimensions and one calibration time interval are considered. In this overview, five ground-based array antennas (labelled 1 to 5) observe three calibrator sources (colored red/green/blue and labelled A to C) within the FoV. The (colored) LoSs from the array towards the sources run parallel for each source and pierce the phase screen at fixed height h (colored circles). The LoS from antenna i at Earth location $\mathbf{a}_i$ towards a peeling source k at local sky position $\hat{\mathbf{s}}_{ik}$ intersects the phase screen at a single pierce point $\mathbf{p}_{ik}$ under a zenith angle ζ_{ik} . For a single LoS from antenna 1 towards source A, we have indicated how the pierce point position $\mathbf{p}_{1A}$ and zenith angle $\zeta_{1A} = \zeta_{1A}$ relate to the antenna position $\mathbf{a}_1 = \mathbf{a}_1$ and the local sky position $\hat{\mathbf{s}}_{1A}$ of the source. For some LoSs the pierce points may overlap (or nearly overlap), as is the case for 1C & 4A and 2C & 5A in our example. The total (integrated) phase rotation along any LoS through the ionosphere is modeled by an instantaneous phase rotation ϕ_{ik}^{ion} at the phase screen height. For example, radio waves traveling along LoSs from source A towards antennas 1 to 5 experience an instantaneous phase rotation ϕ_{1A}^{ion} to ϕ_{5A}^{ion} , respectively, while passing the screen at their related pierce points $\mathbf{p}_{1A}$ to $\mathbf{p}_{5A}$, respectively. Peeling the three calibrator sources yields measurements of the ionospheric phases ϕ_{ik}^{ion} , relative to a common reference antenna (in this example antenna 3; encircled).

LM non-linear least-squares solver, for which a χ^2 sum needs to be defined. From Eq. (15) it follows that

$$\phi_{ik}^{\text{cal}} = \phi^{\text{ion}}(\mathbf{p}_{ik}, \zeta_{ik}) - \phi^{\text{ion}}(\mathbf{p}_{rk}, \zeta_{rk}) - \phi_{ik}^{\text{ambig}}. \quad (22)$$

Consequently, the phase correction in the direction of source k for a baseline consisting of antennas i and j is

$$\phi_{ik}^{\text{cal}} - \phi_{jk}^{\text{cal}} = \left[\phi^{\text{ion}}(\mathbf{p}_{ik}, \zeta_{ik}) - \phi^{\text{ion}}(\mathbf{p}_{jk}, \zeta_{jk}) \right] - \left[\phi_{ik}^{\text{ambig}} - \phi_{jk}^{\text{ambig}} \right]. \quad (23)$$

The χ^2 sum is defined as:

$$\chi^2 = \sum_k \sum_i \sum_{j>i} \left[\left(\phi_{ik}^{\text{cal}} - \phi_{jk}^{\text{cal}} \right) - \left(\phi^{\text{ion}}(\mathbf{p}_{ik}, \zeta_{ik}) - \phi^{\text{ion}}(\mathbf{p}_{jk}, \zeta_{jk}) \right) \bmod 2\pi \right]^2. \quad (24)$$

Fig. 3. Plots of the interpolations of the first six KL base vectors, derived for an artificial but realistic configuration of ionospheric pierce points. In this example, the pierce points (black crosses) are calculated for a single time instance during a 74 MHz VLA-B observation with 13 available calibrator sources in the ~ 10 degree FoV, adopting a phase screen height $h = 200$ km and a structure function power-law slope $\gamma = 5/3$. The horizontal and vertical axes represent angular distances in East-West and North-South directions, respectively, as seen from the center of the Earth, relative to the phase screen's pierce point along the line-of-sight from array center to pointing center, with East- and Northward offsets being positive. At this height, a 0.1 degree angular offset represents a physical horizontal offset of ~ 11.5 km. The direction-dependent phase for each interpolated KL base vector is color-coded and scaled to an arbitrary amplitude range.

This definition has several properties: (i) By remapping the χ^2 terms into the $[0, 2\pi)$ domain, the phase ambiguity terms do not have to be fitted explicitly; (ii) the χ^2 terms of all calibrator sources are weighted equally, so the model is not biased towards the brightest source (as is the case for self-calibration); and (iii) using χ^2 terms from all possible antenna pairs prevents a bias towards the reference antenna.

Using Eq. (24), the LM solver yields a set of model parameters per visibility time stamp. These are stored for later use during imaging (AIPS NI table). The square root of the average of the χ^2 terms equals the average rms phase residual between peeling and model phases. Time intervals that have a bad fit are identified and removed by means of an upper limit ($+2.5\sigma$ rejection) on the distribution of rms phase residuals over time.

Convergence of the LM solver is troubled by 2π phase ambiguities, because these introduce local minima in χ^2 space. A good initial guess of the model parameters greatly helps to overcome this problem. For this purpose, SPAM estimates the global phase gradient over all the pierce points directly from the phase corrections ϕ_{ik}^{cal} and projects it onto the KL base vectors before invoking the LM solver.

Figure 4 shows an example of an ionospheric phase screen that was constructed as described above. The pierce point layout consists of multiple projections of the array onto the phase screen. The low density of calibrators causes a minimal overlap between array projections. Figure 5 shows a comparison

between time-sequences of phase corrections from self-calibration, peeling and model fitting. Because the self-calibration corrections are a flux-weighted average for the full FoV, they are biased towards the brightest source. They look somewhat similar to the peeling solutions of the brightest source, but the latter contains additional fluctuations that vary on a relatively short timescale. The model phases appear similar to the peeling phases, but vary more smoothly. Their values fall somewhere in between the self-calibration phases and the peeling phases. The difference between the peeling phases and model phases are mainly caused by the constraints on the spatial variability of the phase screen model.

3.5. Imaging

With an ionospheric phase screen model available for a given visibility data set, antenna-based phase corrections for any direction in the wide FoV can be calculated (Eq. (22)). Because each visibility consist of contributions from visible sources in different viewing directions, there is no simple operation that removes the ionospheric phase rotations from a visibility data set prior to imaging. Instead, SPAM requires an algorithm that calculates and applies the appropriate model phase corrections during imaging and deconvolving for different parts of the FoV.

SPAM works under the assumption that there exists a fixed angular *isoplanatic patch* size on the sky, with a projected size

Fig. 4. Example of an ionospheric phase screen model fit. The color map represents an ionospheric phase screen at 200 km height that was fitted to the peeling phase solutions of 8 calibrator sources at time-interval $n = 206$ of 10 s during a VLSS observing run of the 74 MHz VLA in BnA-configuration (see Sect. 4, the J1300-208 data set). The plot layout is similar to Fig. 3. The overall phase gradient (depicted in the bottom-left corner) was removed to make the higher order terms more clearly visible. The collection of pierce points from all array antennas to all peeling sources are depicted as small circles. The color in the circle represents the measured peeling phase (the reference antenna VLA N36 was set to match the phase screen value). The size of the circle scales with the magnitude of the estimated phase residual after model correction. The overall rms phase residual $\sigma_{\text{phase}} = 21.799$ degrees (averaged over all pierce points) was one of the better fitting results during this particular observing run.

at ionospheric height smaller than the scale size of ionospheric phase fluctuations, over which variations in ionospheric phase rotation are negligible. Each isoplanatic patch requires at least one phase correction per antenna per visibility time interval. For the VLA at 74 MHz, the isoplanatic patch size is estimated to be 2–4 degrees (Cotton & Condon 2002).

The facet-based polyhedron method for wide-field imaging (see Sect. 3.2) allows for a relatively simple implementation of ionospheric phase correction (Schwab 1984). By choosing a facet size smaller than the isoplanatic patch size, a set of model phase corrections calculated for the center of a facet are assumed to be accurate for the whole facet area. Ionospheric phase model corrections are calculated and stored (AIPS SN tables) for each facet center in the FoV prior to imaging and deconvolution. For the additional facets centered on bright sources (see Sect. 3.2), model phase corrections are optionally replaced by peeling phase corrections to allow for optimized calibration towards these sources.

The SPAM imaging and deconvolution procedure is similar to the procedure used for the field-based calibration method by Cotton et al. (2004), which differs from the standard Cotton-Schwab algorithm by the temporary application of the facet-based phase corrections (AIPS tasks SPLIT and CLINV/SPLIT) to the visibility data for the duration of major CLEAN cycles on individual facets (AIPS tasks IMAGR and UVSUB). After deconvolution, facets are combined to form a single image of the full FoV (AIPS task FLATN). Because antenna-based phase corrections change very little between adjacent facets, the complete

Fig. 5. Example of phase corrections from different steps in the ionospheric calibration process, resulting from processing a VLSS data set with SPAM (see Sect. 4, the real J0900+398 data set). The antenna under consideration is VLA E28, with W20 being the reference antenna (an 5.7 km east-west baseline). The plots represent 25 min of observing time, using a 10 second time resolution. *Top:* Antenna-based phase corrections resulting from self-calibration on the whole FoV. *Middle:* Phase corrections resulting from peeling the brightest (30 Jy) source. *Bottom:* corrections resulting from ionospheric phase modeling in the direction of the (same) brightest source.

set of partly overlapping facet images combine into a continuous image of the FoV.

4. Applications

To demonstrate the capabilities of SPAM, we have defined three test cases based on observations with the VLA at 74 MHz (Kassim et al. 2007). In each test case, SPAM is used for ionospheric phase calibration and imaging of a VLSS visibility data set (Cohen et al. 2007), following the steps described in Sect. 3. In the first test case, SPAM was applied to simulated data to validate basic functionality in a controlled environment. In the next two test cases, SPAM was applied to visibility data from real observations under varying ionospheric conditions. We compare SPAM performance against self-calibration (SC) and field-based calibration (FBC) by analyzing the resulting images. The setup and results of these test cases are described in detail in the following sections.

4.1. Data selection, preparation and processing

In this section we describe how the visibility data sets for the three test cases were selected/constructed. Furthermore, we present details on how these data sets were processed by SPAM into calibrated images of the FoV.

Two VLSS observations, at pointing centers J0900+398 and J1300-208, respectively, have been picked from more than 500 available VLSS observations on the following criteria: (i) both fields contain a relatively large number of bright sources that

Table 1. Overview of the processing parameters for the three data sets that are handled with SPAM as defined in the test cases.

Field name	VLSS J0900+398 (simulated)	VLSS J0900+398 (real)	VLSS J1300-208 (real)
Pixel size ^a	18.9''	18.9''	11.1''
Number of facets	347	243	576
Facet separation	1.18°	1.18°	0.62°
SPAM calibration cycles ^b	1	1	3
Peeling sources	10 ^c	11	9
KL model height	1000 km ^d	200 km	200 km
Fitted KL model terms	15	15	20 ^e
Rejected time intervals	0 / 464	25 / 464	86 / 484
Model fit phase rms	3.0° ± 0.8°	21.3° ± 2.4°	23.2° ± 3.2°
Peeling corrections applied directly	no	yes	yes

^a The pixel size for all field-based calibration images is 20''. ^b Adding more cycles did not significantly improve the image quality. ^c Arbitrarily limited to mimic a more realistic scenario. ^d Increased to improve match with FBC phase screen. ^e In this case, 15 terms proved to be insufficient.

can serve as calibrators and (ii) the ionospheric conditions during the observations appear to be relatively good (J0900+398) and relatively bad (J1300-208). The presence of more than 5 bright sources of at least 5 Jy compensates for the relatively poor efficiency of the VLA 74 MHz receiving system (Kassim et al. 2007). The ionospheric conditions were derived from the apparent smearing of point sources in the images, due to residual phase errors after applying FBC. From experience, we adopted the qualification “good” when the mean width of apparent point sources was at most 5'' larger than the intrinsic 80'' resolution, while for “bad” conditions the mean point source width was larger by at least 15''. In terms of Strehl ratio R (Eq. (6)), “good” and “bad” conditions correspond with $R > 0.996$ and $R < 0.966$, respectively. Additionally, candidate fields were visually inspected for evidence of residual phase errors by the presence or absence of image artefacts near bright sources, which lead to the final selection of the two fields mentioned above.

The difference in observed ionospheric conditions between the two real data sets may be the result of the difference in array size and elevation of the target field. From the VLA site at +34 degrees declination, the J0900+398 field was observed in B-configuration (up to 11 km baselines) at relatively high elevation, while the J1300-208 field was observed in BnA-configuration (up to 23 km baselines) at relatively low elevation. For the J1300-208 observation, the array observed through the ionosphere at larger separations and along longer path lengths than for the J0900+398 observation, which is expected to result in both larger and less coherent phase errors over the array.

Because both real data sets have been previously calibrated and imaged with FBC, the data sets were already partly reduced at the start of SPAM processing. Instrumental calibration was applied (including instrumental phase calibration, similar to Sect. 3.1), most RFI-contaminated data was flagged and the spectral resolution was reduced (see Cohen et al. 2007 for details), but no FBC has been applied yet. For the simulated data set, which is based on the real J0900+398 observations, the measured visibilities were replaced by noiseless model visibilities of an idealized sky, consisting of 91 bright point sources with peak fluxes (larger than 1 Jy) and positions as measured in the J0900+398 FBC image. For each point source, the corresponding model visibility phases were corrupted using the direction-dependent ionospheric phase model that was obtained with FBC to correct the real J0900+398 data.

FBC images of the two real data sets were available in the VLSS archive. For the simulated J0900+398 data set, an “undisturbed” image was made before applying the ionospheric phase

corruptions. All three VLSS data sets have been processed with SPAM, yielding both an SC image and an ionosphere-corrected SPAM image. Relevant details on the processing can be found in Table 1. For SC and SPAM imaging, we adopted most of the imaging-specific settings from FBC (like uniform weighting). Noticeable differences are the use of CLEAN boxes, a smaller pixel size and a different facet configuration.

By choosing a minimum SNR per time interval of 15 and a maximum peeling time interval of 4 min (see Eq. (19)), SPAM was able to peel ~10 sources in each of the real data sets. Lowering the SNR resulted in a much larger scatter in the peeling phases over time, or prevented peeling from converging at all. The peeling time upper limit was chosen to roughly match the spatial density of calibrator sources used in FBC. Determining phase corrections on a 4 min time scale could result in undersampling the time evolution of ionospheric phase errors. Note that this only applies to the faintest of the calibrator sources. The limitations on spatial and temporal sampling of the ionosphere are dictated by the given sensitivity of the VLA.

Because of the high SNR, all 91 sources in the simulated J0900+398 data set qualified for peeling at the highest time resolution of 10 seconds. To mimic a more realistic scenario for further SPAM processing, the number of calibrators was arbitrarily limited to 10. Generally, for all data sets, the images of peeling sources showed larger peak fluxes and less background structure than their counterparts in the SC image, although the contrast became less apparent for weaker and extended (mostly doubles) peeling sources.

As stated in Sect. 3.3, the number of peeling sources is fundamentally limited by the requirement for a large positive number of degrees-of-freedom in the available visibility data. The minimal requirement is given in Eq. (20). Typically, for the VLSS data sets, there were 25 active antennas, 12 frequency channels and 6 visibility intervals (of 10 seconds) in an average peeling interval of 1 min. In our test cases, we typically peel 10 sources, which is much less than $25 \times 12 \times 6 / 2 = 900$, thereby satisfying the minimal requirement.

Due to the uncertainty in their optimal values, it is left to the SPAM user to specify the phase screen model order (the number of KL base vectors), the height h of the phase screen and the power-law exponent γ of the phase structure function. For the applications presented here, we used $h = 200$ km and $\gamma = 5/3$, which is compliant to the measured values given in Sect. 2.1 given the uncertainty in these values. For the simulated data set, we chose instead $h = 1000$ km to better match the corrupting FBC ionospheric phase model that is attached to the sky plane at infinite height. These values gave satisfactory results

for the test applications presented here, but can be further optimized. The optimal model order was found to lie in the range of 15–20 terms, which is 1.5–2 times the number of available peeling sources. Increasing or decreasing the model order caused the model fit to be less accurate or more problematic in terms of convergence.

For both the simulated and real J0900+398 data sets, no improvement in background noise was observed by adding a second calibration cycle after the first. This indicates fast convergence of the SPAM calibration method for quiet ionospheric conditions, where the initial self-calibration is already close to the best achievable calibration of SPAM. For the real J1300-208 data set, adding up to third calibration cycle did improve over the previous cycles.

4.2. Phase calibration accuracy

For the simulated J0900+398 data set, the absolute accuracy of ionospheric calibration can be determined by a direct comparison between the corrupting FBC phase screen and the correcting SPAM phase screen. For this purpose, phase corruptions and corrections were calculated from the models for a hexagonal grid of 342 viewing directions within the FoV. Per viewing direction, the rms phase error was calculated by differencing of the phases from both models and averaging over all time stamps and baselines. The result is depicted in Fig. 6.

For areas near the calibrators and in the center of the field in general, there is a relatively good match between the input and output model, with typical rms phase errors $\lesssim 5$ degrees. The absence of calibrator sources south-west of the field center still results in relatively accurate predictions by the SPAM model. In the direction of peeling sources, the measured rms phase error can be split into a contribution from inaccuracies in the peeling process and a contribution from imperfect model fitting. The latter is approximately 3 degrees (Table 1), therefore the rms phase error introduced by peeling is $\lesssim 4$ degrees. Considering the model setup, the only possible source of error is contamination from other sources while peeling (which appears to happen despite the initial subtraction of the SC model).

Overall, the change in model base from the corrupting FBC model (5 Zernike polynomials) to the correcting SPAM model (15 KL vectors) has a constant accuracy over large parts of the FoV. Towards some parts of the edge of the field the phase errors are substantially larger, up to 20–25 degrees at worst. This agrees with the different asymptotic behaviour towards large radii of the Zernike model (diverge to infinity) and the KL model (converge to zero) in the absence of calibrators. The presence of calibrator sources near the edge (like the one on the North-East edge of the field) leads to a better local match between corrupting ionosphere and correcting model.

For the real observations, in the absence of external sources of information (e.g., GPS measurements), it is not possible to derive the absolute accuracy of ionospheric calibration from the observations themselves. Instead, the residual rms phase error of the model fit to the peeling phases is used as an relative indicator for calibration accuracy over time. For both the real J0900+398 and J1300-208 data sets, the residual rms phase error of ~ 22 degrees is much larger than for the simulated data. This already excludes rejected time stamps with exceptionally large rms values. By inspecting model fits on individual time stamps, we found that there are often a few pierce point phases that deviate significantly more from the fitted model than most neighbouring points. These errors do not appear to be antenna-based instrumental errors, because peeling solutions for the same antenna

Fig. 6. For the simulated J0900+398 data set, the grayscale map represents the residual phase rms between the distorting and correcting ionospheric phase models across the primary beam area, averaged over baselines and time. The phase rms was calculated for a hexagonal grid of viewing directions across the FoV. Each viewing direction is depicted by a small circular area. Overplotted is a contour map of the point sources as seen in the SPAM image (which extends slightly beyond the grid of circles). The 10 peeling sources are marked by circles. The correspondence between the models is largest near the calibrator sources and over a large part of the inner primary beam. The discrepancy is largest near the South-East and North-West borders, away from the calibrators.

towards other calibrator sources do not deviate in the same manner. Typically, these deviating points persist for a few time stamps before disappearing. The ionosphere may be responsible for these very small scale deviations. Another possibility is that the peeling solutions are (sometimes) noisy due to limitations in source SNR.

4.3. Background noise

In this and the next sections, we revert to analyzing image properties for an indirect, relative comparison between the different calibration techniques. In the presence of residual phase errors, part of the image background noise level consists of residual sidelobes after CLEANing. The local sidelobe noise increases with both the rms phase error and the local source flux. When measured over a large image area, the mean sidelobe noise depends mainly on rms phase error. For all relevant output images, the mean image noise σ was determined by fitting a Gaussian to the histogram of image pixel values from the inner quarter radius of the FoV (AIPS task IMEAN). Note that these images have not been corrected for primary beam attenuation. The results are given in Table 2.

Because no noise was added to the simulated J0900+398 data set, the resulting image noise of $3.0 \text{ mJy beam}^{-1}$ in the

Table 2. Overview of results from calibrating and imaging three test case data sets with no ionosphere (Undisturbed), self-calibration (SC), field-based calibration (FBC) and SPAM.

Field name	VLSS J0900+398 (simulated)	VLSS J0900+398 (real)	VLSS J1300-208 (real)
Mean background noise σ [mJy beam ⁻¹]:			
Undisturbed	3.0	–	–
SC	10.2	71	92
FBC	–	87	118
SPAM	6.7	67	68
Number of sources with a peak flux larger than 5σ :			
Undisturbed	91	–	–
SC	91	393	374
FBC	–	310	285
SPAM	91	372	392
5σ source fraction with an NVSS counterpart within $80''$:			
Undisturbed	1.	–	–
SC	1.	0.83	0.60
FBC	–	0.86	0.74
SPAM	1.	0.97	0.97

undisturbed image is caused by incomplete UV coverage and inaccuracies in the imaging process (see Sect. 3.2), limiting the dynamic range to $\sim 10^4$. The local noise is highest near the sources, but significantly less near the brightest 10 sources with dedicated facets centered on their peak position. The SC and SPAM images from this data set were created using the same facet configuration. The SC image noise of $10.2 \text{ mJy beam}^{-1}$ is 3.4 times as high as the undisturbed image noise, therefore dominated by phase error induced sidelobe noise. The SPAM image noise of $6.7 \text{ mJy beam}^{-1}$ is a significant improvement over the SC image, but still 2.2 times as high as in the undisturbed image. The local noise in the SC and SPAM images has increased most apparently near bright sources as compared to the undisturbed image, which confirms the presence of residual phase errors after calibration.

For the real J0900+398 data set, both the SC and SPAM images have an image noise of $\sim 70 \text{ mJy beam}^{-1}$. The SPAM image noise is slightly lower than SC. The local noise in the SC image is higher near bright sources. This is not the case in the SPAM image, which must be a direct result of an improved calibration accuracy near these sources. The FBC image noise for this data set is ~ 20 percent higher, a combination of a higher average noise over the FoV and higher local noise near bright sources.

For the real J1300-208 data set, the SPAM image has the same image noise as for the real J0900+398 data set, with no apparent increase near bright sources. At the same time, the noise levels in the SC and FBC images have increased with 30 and 35 percent, respectively. The noise in the SC image is highest near the bright sources. The FBC noise is highest near the brightest source and remains high in the rest of the image. The significant increase of the average FBC noise level indicates a dependence on ionospheric conditions, and therefore on calibration accuracy. The SPAM image noise appears to have little or no dependence on varying ionospheric conditions (Fig. 7).

4.4. Source properties

The presence of residual phase errors changes the apparent distribution of flux of a source (see Sect. 2.2). In the time-averaged image, sources may appear offset from their intrinsic position, may suffer from smearing or deformation, and sidelobes may be misidentified as sources. Comparing the properties of the same sources in differently calibrated images allows for a

relative comparison of the performance of the different calibration techniques.

To allow for comparison of source properties, we applied the source extraction tool BDSM (Mohan et al. 2008) on all relevant images. BDSM performed a multiple 2-dimensional Gaussian fit on islands of adjacent pixels with amplitudes above a specified threshold based on the *local* image noise σ_L in the image. Multiple overlapping Gaussians were grouped together into single sources. We applied BDSM to all images, using the default extraction criteria, except for the following: a source detection requires at least 4 adjacent pixel values above $2.5\sigma_L$, with at least one pixel value above $4\sigma_L$.

4.4.1. Source counts

Due to the non-Gaussian character of the phase-induced sidelobe noise, the source catalogs will contain spurious detections. To suppress these, we removed sources with a peak flux smaller than 5σ from the catalogs. The remaining number of catalog entries are listed in Table 2. Additionally, each catalog was cross-associated against the NVSS catalog, which has a slightly higher resolution ($45''$). For an average spectral index of -0.8 , the NVSS detection limit is at least 10 times lower than for the VLSS. At the risk of missing an incidental ultra-steep spectrum source, we determined the source fraction that has an NVSS counterpart within an $80''$ radius (one VLSS beamsize), which are also listed in Table 2.

For the simulated J0900+398 data set, all 91 input sources are detected and matched against NVSS counterparts, regardless of the calibration method. Due to the low noise levels and the lower limit of 1 Jy on the input source catalog, all sources are effectively $\geq 100\sigma$ detections. None of the sources had more than one Gaussian fitted to it, despite the freedom to do so.

For the real J0900+398 data set, the higher σ in the FBC image is reflected in a smaller number of source detections as compared to SC and SPAM. SC detects slightly more sources than SPAM, despite the slightly higher σ . However, there is a very large fraction of sources in the SPAM catalog that has an NVSS counterpart, significantly larger than for both the SC and FBC catalogs. This suggests that the SPAM catalog is much less contaminated by false detections than the SC and FBC catalogs, resulting in a larger absolute number of true detections.

Fig. 7. Greyscale plots of a 3.5×3.5 square degree area in the VLSS J1300-208 field centered on the bright (40 Jy) point source 3C 283. All three images have contours (black lines) overplotted at [0.15, 0.48, 0.83, 1.16, 1.50] Jy. *Left:* image after self-calibration, *middle:* image after field-based calibration, and *right:* image after SPAM calibration.

This is further strengthened by the results from the real J1300-208 data set. For this test case, the SPAM image has the largest number of source detections. Again, the SPAM catalog has the largest fraction of associations with the NVSS catalog, the same fraction as with the J0900+398 data set. In contrast, the fraction of NVSS counterparts for SC and VLSS have both gone down. This is best explained by an increase in (non-Gaussian) sidelobe noise in the image background due to calibration errors, which corresponds with the observed increase in σ .

4.4.2. Source peak fluxes

The presence of residual phase errors after calibration can cause an unresolved source shape to deviate from a point source shape. The source flux is redistributed over a larger area and the peak flux of the source drops. At $80''$ resolution, most sources in a VLSS field are unresolved. Therefore, a mean increase of source widths over the point source width is a direct measure of ionospheric conditions. This argument was used in the pre-selection of data sets for our test cases.

For significant source deformations or low SNR sources, determination of the shape of individual sources is subject to large uncertainties (e.g., Condon et al. 1997). Because determination of peak fluxes is much more robust, we use these for a relative comparison of calibration accuracy. Starting with the original catalogs as produced by BDSM, we associate sources between the undisturbed, FBC, SC and SPAM catalogs that lie within $80''$ of the same NVSS source and has a peak flux larger than 5σ in at least one of the two catalogs.

For the simulated J0900+398 data set, the true peak fluxes of all 91 sources are known. A comparison between peak fluxes from the undisturbed image and the input catalog identifies a small (<1 percent) CLEAN bias of $3.6 \text{ mJy beam}^{-1}$ (e.g., Condon et al. 1994, 1998; Becker et al. 1995). Ignoring the image noise dependency of CLEAN bias, we applied this small correction to the peak fluxes in the undisturbed, SC and SPAM source catalogs before proceeding. Figure 8 shows a comparison of the measured-to-input peak flux ratios for sources in the SC and SPAM images. The mean peak flux ratio for both images is approximately equal and just slightly smaller than one. The larger scatter in the SC peak fluxes is consistent with a higher σ . Using Eq. (6), the random part of the mean rms phase error for

Fig. 8. Peak flux ratios of point sources in the simulated J0900+398 field. Peak fluxes were measured in the self-calibration image and the SPAM image, corrected for a small CLEAN bias and divided by the input model peak fluxes. The size of each dot scales with the input model peak flux, ranging from 1.02 to 26.7 Jy. Ideally (without phase errors), the peak flux ratios would be scattered around *one* (solid lines) due to image noise dependent errors in the peak flux determination. Instead, the peak flux ratio distributions along the *x*- and *y*-axis are centered around 0.995 and 0.996, respectively (dotted lines), which is a direct result of the residual phase errors. The smaller and larger scatter in distribution of SPAM- and self-calibration peak flux ratios is consistent with peak flux determination inaccuracies due image background noise levels.

both SC and SPAM is estimated at 5–6 degrees. This value is comparable to the observed rms phase error over large parts of the SPAM image (Sect. 4.2).

To study the nature of residual rms phase errors after application of SPAM, we plot the rms phase errors at the source positions from Fig. 6 against SPAM-to-input peak flux ratios (Fig. 9). For Gaussian random phase errors, the peak flux ratio is

Fig. 9. Peak flux ratios in the simulated J0900+398 field: *Left:* Peak flux ratios of the 91 extracted sources from the SPAM image as compared to the input model sources, plotted as a function of the residual rms phase error after SPAM calibration. Overplotted is the theoretical Strehl ratio (solid line) as given in Eq. (6). For larger rms phase errors, the measured peak flux ratios do not follow the theoretical strehl ratio curve. This indicates that systematic phase errors dominate the larger rms phase errors. *Right:* same peak flux ratios plotted as a function of absolute position offset between extracted sources in the SPAM image and the input model (see Fig. 12). The presence of a strong correlation indicates that residual phase gradients dominate the larger rms phase errors.

Fig. 10. Peak fluxes in the real J0900+398 field *Left:* peak flux comparison for 367 sources detected in both the self-calibration and SPAM images. The straight diagonal line represents equality, the dashed lines represent $3\sigma_c$ deviations (where σ_c is the combined noise level from both images), and the dotted lines indicate the 5σ detection limit. For bright sources (peak fluxes ≥ 1 Jy beam $^{-1}$), the average peak flux ratio is 1.00. *Middle:* same for 329 sources in the field-based calibration (VLSS) and SPAM images. The average bright peak flux ratio of SPAM over field-based calibration is 1.10. *Right:* same for 313 sources in the self-calibration and field-based calibration (VLSS) images. The average bright peak flux ratio of self-calibration over field-based calibration is 1.10. In all plots, the image noise causes a larger scatter in the peak flux determinations of faint sources (≤ 1 Jy beam $^{-1}$) and consequently, a selection bias towards positively enhanced peak fluxes that increases with image noise.

expected to decrease with increased rms phase error as described in Eq. (6). However, the discrepancy between the data points and Eq. (6) indicates that for larger rms values the phase errors are predominantly systematic rather than random.

For the real J0900+398 data set, Fig. 10 shows a comparison of peak fluxes for associated sources in the SC, FBC and SPAM catalogs. There is a good match between peak fluxes measured in the SC and SPAM catalogs. For high SNR sources with a peak flux above 1 Jy, the SPAM peak fluxes match on average within 1 percent with the SC peak fluxes. Similarly, SC and SPAM peak fluxes are on average 10 percent higher than FBC peak fluxes. The systematic increase of peak fluxes for SC and SPAM as compared to FBC for many more than the calibrator

sources denotes a more accurate calibration over large parts of the FoV. Towards the low flux end, source detections are slightly biased towards the image with the highest noise level, which is the FBC image.

Figure 11 shows the same comparison of peak fluxes for the real J1300-208 data set. For high SNR sources with a peak flux above 1 Jy, the SC peak fluxes are by far the smallest, while FBC and SPAM peak fluxes are on average higher by 15 and 24 percent, respectively. The relative loss of peak flux in the SC image is a clear indication of the break-down of the assumption of isoplanaticity across the FoV. Under the conditions that clearly need direction-dependent corrections, the SPAM peak fluxes are on average 7 percent higher than the FBC peak fluxes.

Fig. 11. Peak fluxes in the (real) J1300-208 field: *Left*: peak flux comparison for 247 sources detected in both the self-calibration and SPAM images. For bright sources (peak fluxes ≥ 1 Jy beam $^{-1}$), the average peak flux ratio of SPAM over SC is 1.24. *Middle*: Same for 278 sources in the field-based calibration (VLSS) and SPAM images. The average bright peak flux ratio of SPAM over field-based calibration is 1.07. *Right*: same for 202 sources in the self-calibration and field-based calibration (VLSS) images. The average bright peak flux ratio of field-based calibration over self-calibration is 1.15.

Fig. 12. Position offsets in the simulated J0900+398 field: *Left*: offsets between the measured source positions in the self-calibration image as compared to the input model. *Right*: same for the SPAM image. In both cases, the distribution around the origin is non-Gaussian. For the SPAM image, the tail of points extending roughly northwards indicates the presence of persistent phase gradients in local parts of the SPAM image. All source position offsets fall well within the size of the 80'' restoring beam (dotted line).

4.4.3. Astrometry

When the time-average of residual phase errors towards a source contains a non-zero spatial gradient, the source will appear to have shifted its position in the final image (see Sect. 2.2). This gradient may indicate a limitation of the calibration model to reproduce the ionospheric phase corruptions (e.g., in the absence of nearby calibrators), but may also be introduced by the peeling process. The latter occurs when a peeling source is re-centered to the wrong catalog position (see Sect. 3.3). Because such an error propagates into the calibration model, many sources in the vicinity of the peeling source may also suffer from a systematic astrometric error.

For the simulated data set, the peak positions of sources as determined by BDSM were compared against the positions of counterparts in the input model. For the real data sets, we compared against the NVSS catalog instead. When comparing against NVSS positions, apparently large position offsets may occur due to resolution differences and spectral variation across

the source. Averaged over a large number of sources, these offsets should have no preferential orientation. In contrast, a residual phase gradient in a certain viewing direction is expected to cause systematic offsets for groups of sources in a certain preferential direction.

For the simulated J0900+398 data set, Fig. 12 shows that the positions for both SC and SPAM are accurate to within $\sim 10''$, except for a small tail of ~ 15 SPAM sources that have somewhat larger offsets. These sources are all positioned near the edge of the FoV, where the rms phase error is large (Fig. 6). Figure 9 also confirms this by the clear correlation between rms phase error and absolute position offsets.

For the real J0900+398 data set, the source position offsets for SC, FBC and SPAM relative to NVSS catalog positions are plotted in Fig. 13. The larger scatter as compared to the simulated J0900+398 data set can be the (combined) result of less accurate position measurements due to higher image noise, resolution and spectral differences between the observations and the NVSS catalog or larger residual rms phase errors after

Fig. 13. Position offsets in the real J0900+398 field *Left*: offsets between the measured source positions in the self-calibration image as compared to the NVSS catalog. *Middle*: same for the field-based calibration (VLSS) image. *Right*: Same for the SPAM image.

calibration. The observed scatter for SC is centered around a point that is offset from the origin by $\sim 5''$, which is either caused by inaccuracies in the initial sky model or during the self-calibration process (Sect. 3.2). The scatter of both FBC and SPAM offsets is centered close to the origin. The rms of the scatter around the mean position offset is $10.5''$ for both FBC and SPAM (despite the apparently larger scatter for SPAM, which is due to a larger number of data points), both smaller than the $11.9''$ for SC.

For the real J1300-208 data set, the source position offsets for SC, FBC and SPAM relative to NVSS catalog positions are plotted in Fig. 14. The position scatter for all three methods is significantly larger than for the real J0900+398 data set, and all suffer from systematic position offsets in varying degrees of severity. The position offsets in the SC image have a seriously distorted distribution, which includes a large tail of points that extends roughly southwards. This indicates the presence of varying systematic source offsets over the whole FoV. The distribution of position offsets in the FBC image is more compact but also asymmetric, and is approximately centered around a point that is $\sim 10''$ offset in northward direction from the origin. A large number of the SPAM position offsets are clustered near the origin, similar to the real J0900+398 data set, but there is an additional tail of points that runs roughly northwards. The rms of the scatter around the mean position offset is $20.7''$, $16.5''$ and $14.8''$ for SC, FBC and SPAM, respectively, which confirms the apparently strongest clustering of points in the SPAM position offset plot.

Systematic position offsets in the images can be reduced by distortion and regridding of the images. For this purpose, Cohen et al. (2007) fit a fourth order Zernike polynomial to the (time constant) position offsets of typically more than 100 sources in the FBC images of the VLSS. They estimate that, after correction, the final residual position error in the full VLSS catalog due to the ionosphere is $\lesssim 3''$ in both RA and Dec.

5. Discussion and conclusions

The SPAM method for ionospheric calibration has been successfully tested on one simulated and two carefully selected visibility data sets of 74 MHz observations with the VLA (taken from the

VLSS; Cohen et al. 2007). From the results of these test cases, we draw the following conclusions:

- (i) A proof-of-concept is given for several different techniques that were incorporated in SPAM calibration. The peeling technique (Noordam 2004) was successful in providing relative measurements of ionospheric phase errors in the direction of several bright sources in the FoV. The Karhunen-Loève phase screen (van der Tol & van der Veen 2007) at fixed height was able to combine these measurements into a consistent model per time stamp. For relatively bad ionospheric conditions, it was demonstrated that the ionospheric calibration cycle (repeated ionospheric calibration and subsequent imaging; Noordam 2004) converges within a few iterations to a calibration of similar accuracy as under relatively good ionospheric conditions (for which one iteration was sufficient).
- (ii) Ionospheric calibration with SPAM is more accurate than the existing self-calibration (e.g., Pearson & Readhead 1984) and field-based calibration (Cotton et al. 2004) techniques. Even for relatively compact array configuration like VLA-B and BnA, significant improvements in image quality are obtained by allowing for higher-order (i.e., more than a gradient) spatial phase corrections over the array in any viewing direction. In the resulting images, we obtained dynamic range improvements of 5–45% and 70–80% under relatively good and bad ionospheric conditions, respectively.
- (iii) Although the mean astrometric accuracy of source positions in SPAM images is similar to or better than for self-calibration and field-based calibration, systematically larger astrometric errors are present in regions of the output images of all calibration methods. This is caused by a shortage of available calibrators in these regions and positional inaccuracies in the reference source catalog used for calibration.

The 65 mJy beam^{-1} noise levels in the SPAM images match the lowest noise levels of the more than 500 images that define the VLSS survey. A potential reduction of the average noise level from $100 \text{ mJy beam}^{-1}$ to 65 mJy beam^{-1} for the full VLSS survey would significantly increase the number of source detections from $\sim 70\,000$ to about $120\,000$ (an increase of $\sim 75\%$), but also it would greatly enhance virtually every science goal. For example, using the radio luminosity function for high-luminosity radio

Fig. 14. Position offsets in the real J1300-208 field: *Left*: offsets between the measured source positions in the self-calibration image as compared to the NVSS catalog. *Middle*: same for the field-based calibration (VLSS) image. *Right*: same for the SPAM image.

galaxies from Jarvis et al. (2001), the estimated number of detectable HzRGs in the VLSS would increase by 65%, but also the maximum redshift would increase. For a luminous radio galaxy with luminosity of $2 \times 10^{28} \text{ W Hz}^{-1} \text{ sr}^{-1}$ at 74 MHz, the redshift limit would rise from $z = 5.7$ to $z = 6.8$. Another example is the detection and study of cluster radio halos. Using available halo population models (Enßlin & Röttgering 2002; Cassano et al. 2006), the anticipated noise reduction would roughly double the number of detectable halo systems.

For the VLSS, the estimated theoretical thermal noise level of 35 mJy beam^{-1} is still a factor of two lower than the average background noise level of $\sim 65 \text{ mJy beam}^{-1}$ in the SPAM images. From inspection of the SPAM images we cannot identify an obvious single cause for this. Therefore, similar to Cohen et al. 2007, we expect the remaining excess noise to be the combined result of several different causes, including residual ionospheric phase errors after SPAM calibration, but also residual RFI, collective sidelobe noise from many non-deconvolved sources (too faint or outside the FoV) and variable source amplitude errors (e.g., due to pointing errors and non-circular antenna beam patterns; see Bhatnagar et al. 2008).

The SPAM test results indicate that the ionospheric calibration accuracy may be further improved. The typical model fit rms phase error per antenna of $\sim 20\text{--}30$ degrees for real data sets is much larger than the 3 degrees for the noiseless simulated data set. There are several possible sources of error, either in the peeling phase corrections or the ionospheric phase model. Noise in the visibilities (either thermal or non-thermal), contamination from other sources, inaccuracies in the peeling source model and undersampling of the fastest phase fluctuations are factors that degrade the accuracy of peeling. Also, the ionospheric phase screen model may be a poor representation of reality, either because it is incomplete (e.g., absence of vertical structure) or the fixed model parameters are chosen poorly (e.g., screen height, spectral index of phase fluctuations). Several of these issues will be addressed in future work (Sect. 6).

The potential problems with the peeling technique raises the question whether one should use alternative methods. Apart from the precautions described in Sect. 3.3, we have found little means to improve the accuracy of the peeling process for single sources any further. One unexplored option is to peel sources in groups, e.g. identify isoplanatic patches of sky with a large enough total flux from multiple sources. Two possible alternative approaches to peeling are: (i) simultaneous self-calibration towards multiple sources in the FoV; or (ii)

fitting the ionosphere model directly to the visibilities rather than using peeling as an intermediate step. Although these alternative approaches have not been tested by us in practice, we anticipate little improvement over our current accuracy. Van der Tol, Jeffs & van der Veen (2007) show that, theoretically, iterative peeling converges to the same solution as simultaneous self-calibration. A direct fit of the ionosphere model to the visibilities is, similar to self-calibration, biased towards accurate solving in the direction of the apparently strongest source in the FoV. Although not conclusive for this approach, tests with SPAM show that using even a moderate flux-based weighting into the ionospheric phase model fitting of peeling phase corrections introduces a strong bias towards the brightest source, while calibration accuracy towards other peeled sources degrades severely.

For the existing and future large low-frequency radio interferometer arrays like VLA-A, GMRT, LOFAR, LWA and SKA, the need for a direction-dependent ionospheric calibration method is evident. Based on the results presented in this paper, it is difficult to draw quantitative conclusions on the achievable calibration accuracy for these arrays. If a SPAM-like calibration algorithm is to be used in a very high signal-to-noise observing regime under quiet to moderate ionospheric conditions, it seems likely that residual rms phase errors in the order of a few degrees could be achieved, comparable to the SPAM results on the simulated VLSS data set.

When relying on the array itself to provide the necessary measurements to constrain ionospheric correction models, ionospheric calibration requires an array layout and sensitivity that allows for sampling the ionosphere over the array at the relevant spatial scales and time resolution. The spatial sampling is determined by the instantaneous pierce point distribution (or more general, the distribution of lines-of-sight through the ionosphere), which depends on the array layout and the detectable calibrator constellation. For future design of low-frequency arrays, it is recommended to optimize the array layout not just for scientific arguments (in general, centrally dense and sparse outside for good UV coverage), but also for ionospheric calibrability (in general, both uniform and randomized).

6. Future work

To test the robustness and limitations of the method, it is necessary to apply SPAM calibration on a wide variety of data sets at different (low) frequencies, obtained with different arrays under

different ionospheric conditions. Our highest priority is to test SPAM on observations from the largest existing LF arrays; the VLA in A-configuration and the GMRT. Data for these tests have been obtained and tests are currently in progress. One important possible limitation is the use of a 2-dimensional phase screen to represent the ionosphere. We plan to expand the SPAM model by including multiple screens at different heights and compare the resulting image properties against the current single screen model.

Another limitation of the current implementation is the absence of restrictions on the time behaviour of the model. Antenna-based peeling phases clearly show a coherent temporal behaviour, which is likely to exist for physical reasons. This could be used to reduce the number of required model parameters and suppress the noise propagation from the peeling solutions. We are currently investigating the possibilities of forcing the SPAM model to be continuous in time.

Several of the authors of this article are currently involved in setting up a simulation framework in which one has full control over the sky emission, ionospheric behaviour and array characteristics when generating artificial low frequency observations. Like in the test case on simulated data presented in Sect. 4, this allows for direct and quantitative comparison between the distorting ionosphere model and the recovered ionospheric phase model by SPAM. We plan to use this setup to further test optimize SPAM calibration for a broad range of ionospheric conditions.

Acknowledgements. The authors would like to thank Rudolf Le Poole, Reinout van Weeren, Sridharan Rengaswamy, Amitesh Omar, Mamta Pandey, Oleksandr Usov, James Anderson, Ger de Bruyn, Jan Noordam, Maaijke Mevius, Dharam Vir Lal, C.H. Ishwara-Chandra, A. Pramesh Rao, Jim Condon and Juan Uson for useful discussions. Special thanks to Mark Kettenis for fast bug fixing and implementation of new functionality in ParselTongue, and Niruj Ramanujam Mohan for similar performance on BDSM. The authors also thank the anonymous referee for useful comments and suggestions. HTI acknowledges a grant from the Netherlands Research School for Astronomy (NOVA). SvdT acknowledges NWO-STW grant number DTC.5893. This publication made use of data from the Very Large Array, operated by the National Radio astronomy Observatory. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

Appendix A: Derivation and interpolation of the KL base vectors

This section contains an outline of the derivation and interpolation of the Karhunen-Loève (KL) base vectors that are used to describe the ionospheric phase screen in SPAM. The KL technique is adopted from the work by van der Tol & van der Veen (2007). For a given stochastic model of spatial electron density fluctuations in the ionosphere, the differential phase rotation on rays of passing radio waves can be described by a (zero mean) isotropic phase screen $\phi(\mathbf{p})$ with a given spatial covariance function

$$C_{\phi\phi}(r) = \langle \phi(\mathbf{p})\phi(\mathbf{p} + \mathbf{r}) \rangle, \quad (\text{A.1})$$

where $\langle \dots \rangle$ denotes the expected value and $r = |\mathbf{r}|$ is the length of vector $\mathbf{r}$. In Kolmogorov turbulence theory, a phase structure function is defined as (Eq. (5))

$$D_{\phi\phi}(\mathbf{r}) = \langle [\phi(\mathbf{p}) - \phi(\mathbf{p} + \mathbf{r})]^2 \rangle. \quad (\text{A.2})$$

The structure function and the covariance function are related through

$$\begin{aligned} D_{\phi\phi}(r) &= 2 [C_{\phi\phi}(0) - C_{\phi\phi}(r)] \Leftrightarrow \\ C_{\phi\phi}(r) &= C_{\phi\phi}(0) - \frac{1}{2} D_{\phi\phi}(r), \end{aligned} \quad (\text{A.3})$$

where $C_{\phi\phi}(0) \equiv \sigma_\phi^2$ is the phase variance. The phase structure function of a plane wave that passed through a turbulent layer is found to behave as a power-law over a large range of spatial scales (see Sect. 2.1)

$$D_{\phi\phi}(r) = (r/r_0)^\gamma, \quad (\text{A.4})$$

where r_0 is a measure of the scale size of phase fluctuations and γ is the power-law slope.

Because the domain of $\mathbf{p}$ is a limited set of P discrete ionospheric pierce points $\mathbf{p}$, we switch to matrix notation. The elements of the $(P \times P)$ phase covariance matrix are given by

$$C_{\phi\phi}[i, j] = C_{\phi\phi}(r_{ij}), \quad (\text{A.5})$$

with $\mathbf{p}_i, \mathbf{p}_j \in \{\mathbf{p}\}$ and $r_{ij} = |\mathbf{p}_i - \mathbf{p}_j|$. The symmetric covariance matrix can be decomposed into

$$\mathbf{C}_{\phi\phi} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T, \quad (\text{A.6})$$

where the columns of $(P \times P)$ matrix $\mathbf{U}$ contain the orthonormal eigenvectors of $\mathbf{C}_{\phi\phi}$, $\mathbf{U}^T$ is the transpose of $\mathbf{U}$ and the $(P \times P)$ diagonal matrix $\mathbf{\Lambda}$ contains the eigenvalues. The eigenvectors in the columns of $\mathbf{U}$ are a suitable set of base vectors to describe the phase screen at the pierce points $\{\mathbf{p}\}$. The eigenvalues on the diagonal of $\mathbf{\Lambda}$ are a measure of the variance of the eigenvector coefficients. Reducing the fitting order is an arbitrary, but necessary step. An optimal subset of eigenvectors is determined by selecting only those with the largest eigenvalues, as the coefficients with the highest variances are the most significant in the modeling problem.

When the fitting order is reduced from P to $\tilde{P} < P$, we are left with a subset of eigenvectors in the columns of $(P \times \tilde{P})$ matrix $\tilde{\mathbf{U}}$, and eigenvalues on the diagonal of $(\tilde{P} \times \tilde{P})$ matrix $\tilde{\mathbf{\Lambda}}$, for which

$$\mathbf{C}_{\phi\phi} \approx \tilde{\mathbf{U}}\tilde{\mathbf{\Lambda}}\tilde{\mathbf{U}}^T. \quad (\text{A.7})$$

The phase screen at the pierce point locations can be approximated by

$$\Phi \approx \tilde{\mathbf{U}}\mathbf{q}, \quad (\text{A.8})$$

where we have denoted the phases and eigenvector coefficients in matrix notation as Φ ($P \times 1$) and $\mathbf{q}$ ($\tilde{P} \times 1$), respectively. $\mathbf{q}$ is unknown and needs to be solved for by use of a non-linear least squares method using Eq. (24). When $\mathbf{q}$ is determined, the phase screen can be evaluated at arbitrary pierce point locations $\{\tilde{\mathbf{p}}\}$ through Kriging interpolation (Matheron 1973):

$$\hat{\Phi} = \mathbf{C}_{\hat{\phi}\phi} \mathbf{C}_{\phi\phi}^{-1} \tilde{\mathbf{U}}\mathbf{q} \quad (\text{A.9})$$

where $\mathbf{C}_{\hat{\phi}\phi}$ is the $(\tilde{P} \times P)$ covariance matrix between $\{\tilde{\mathbf{p}}\}$ and $\{\mathbf{p}\}$. Inversion of the full $\mathbf{C}_{\phi\phi}$ can be approximated using

$$\mathbf{C}_{\phi\phi}^{-1} \approx \tilde{\mathbf{U}}\tilde{\mathbf{\Lambda}}^{-1}\tilde{\mathbf{U}}^T. \quad (\text{A.10})$$

The elements of the covariance matrix can be calculated using Eqs. (A.3) and (A.4). For our application, the absolute value of r_0 is not relevant because we only require the relative eigenvalues

for the order reduction. The elements of the $(P \times P)$ structure matrix are given by

$$D_{\phi\phi}[i, j] = D_{\phi\phi}(r_{ij}). \quad (\text{A.11})$$

The relation between the structure matrix and the covariance matrix is given by

$$C_{\phi\phi} = \sigma_{\phi}^2 \mathbf{1}\mathbf{1}^T - \frac{1}{2} D_{\phi\phi}, \quad (\text{A.12})$$

with $\mathbf{1}$ a $(P \times 1)$ vector containing ones. The phase variance terms σ_{ϕ}^2 are removed from the equations by explicitly removing the mean phase from the individual phases through the substitution

$$\begin{aligned} \phi &\rightarrow \phi' = \phi - \langle \phi \rangle \Rightarrow \\ \Phi &\rightarrow \Phi' = \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) \Phi, \end{aligned} \quad (\text{A.13})$$

where $\mathbb{I}$ is the $(P \times P)$ identity matrix. Applying this substitution to the covariance matrix yields

$$\begin{aligned} C_{\phi'\phi'} &= \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) C_{\phi\phi} \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) \\ &= \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) \left[-\frac{1}{2} D_{\phi\phi} \right] \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right). \end{aligned} \quad (\text{A.14})$$

The σ_{ϕ}^2 terms have dropped because of the properties

$$\left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) \mathbf{1} = \mathbf{0}, \quad \mathbf{1}^T \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) = \mathbf{0}. \quad (\text{A.15})$$

For Kriging interpolation, a similar substitution is performed as in Eq. (A.13):

$$\hat{\Phi} \rightarrow \hat{\Phi}' = \hat{\Phi} - \left(\frac{1}{P} \hat{\mathbf{1}}\hat{\mathbf{1}}^T \right) \Phi, \quad (\text{A.16})$$

with $\hat{\mathbf{1}}$ a $(\hat{P} \times 1)$ vector containing ones. The covariance matrix between $\{\hat{p}\}$ and $\{p\}$ (Eq. (A.9) now becomes

$$C_{\hat{p}p} = \left(\left[-\frac{1}{2} D_{\hat{p}\hat{p}} \right] - \left(\frac{1}{P} \hat{\mathbf{1}}\hat{\mathbf{1}}^T \right) \left[-\frac{1}{2} D_{\phi\phi} \right] \right) \times \left(\mathbb{I} - \frac{1}{P} \mathbf{1}\mathbf{1}^T \right) \quad (\text{A.17})$$

where $D_{\hat{p}\hat{p}}$ is the $(\hat{P} \times P)$ structure matrix between $\{\hat{p}\}$ and $\{p\}$, calculated using Eqs. (A.4) and (A.11).

References

- Anderson, J. M. 2006, Evaluation of Ionospheric Correction Method for the European VLBI Network, poster presentation at the URSI General Assembly, Delhi, India
- Becker, R. H., White, R. L., & Helfand, D. J. 1995, *ApJ*, 450, 559
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *A&A*, 487, 419

- Bridle, A. H., & Greisen, E. W. 1994, *AIPS Memo*, 87
- Briggs, D. S. 1995, Ph.D. Thesis, New Mexico Institute of Mining Technology, Socorro, New Mexico, USA
- Briggs, D. S., & Cornwell, T. J. 1992, *ASPC*, 25, 170
- Cassano, R., Brunetti, G., & Setti, G. 2006, *MNRAS*, 369, 1577
- Clark, B. G. 1980, *A&A*, 89, 377
- Cohen, A. S., & Röttgering, H. J. A. 2008, Probing Fine-Scale Ionospheric Structure with the Very Large Array Radio Telescope, to be submitted
- Cohen, A. S., Röttgering, H. J. A., Kassim, N. E., et al. 2003, *ApJ*, 591, 640
- Cohen, A. S., Röttgering, H. J. A., Jarvis, M. J., Kassim, N. E., & Lazio, T. J. W. 2004, *ApJS*, 150, 417
- Cohen, A. S., Lane, W. M., Cotton, W. D., et al. 2007, *ApJ*, 134, 1245
- Condon, J. J. 1997, *PASP*, 109, 166
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1994, *ASPC*, 61, 155
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, 115, 1693
- Cornwell, T. J., & Perley, R. A. 1992, *A&A*, 261, 353
- Cornwell, T. J., Braun, R., & Briggs, D. S. 1999, *ASPC*, 180, 151
- Cotton, W. D., 1999, *ASPC*, 180, 357
- Cotton, W. D., & Condon, J. J. 2002, Proceedings of URSI General Assembly, paper 0944
- Cotton, W. D., & Uson, J. M. 2008, *A&A*, 490, 455
- Cotton, W. D., Condon, J. J., Perley, R. A., et al., 2004, *SPIE*, 5489, 180
- Davies, K., *Ionospheric Radio* 1990, Institution of Engineering and Technology, Stevenage, UK
- Enßlin, T. A., & Röttgering, H. J. A. 2002, *A&A*, 396, 83
- Gizani, N. A. B., Cohen, A. S., & Kassim, N. E., *MNRAS*, 358, 1061
- Högbom, J. A. 1974, *A&AS*, 15, 417
- Jacobson, A. M., & Erickson, W. C. 1992, *Planet. Space Sci.*, 40, 447
- Jarvis, M. J., Rawlings, S., Willott, C. J., et al. 2001, *MNRAS*, 327, 907
- Kassim, N. E., Lazio, T., Joseph, W., et al. 2007, *ApJS*, 172, 686
- Kettenis, M., van Langevelde, H. J., Reynolds, C., & Cotton, B. 2006, *ASPC*, 351, 497
- Lazio, T. J. W., Kassim, N. E., & Perley, R. A. 2005, *Low-Frequency Data Reduction at the VLA: A Tutorial for New Users*, version 1.13
- Lonsdale, C. J. 2005, *ASPC*, 345, 399
- Markwardt, C. B. 2008 [[arXiv:0902.2850](http://arxiv.org/abs/0902.2850)]
- Matheron, G. 1973, *Adv. Appl. Prob.*, 5, 439
- Mohan, N. R. 2008, ANAAMIKA manual, version 2.1, available online through http://www.strw.leidenuniv.nl/~mohan/anaamika_manual.pdf
- NIMA Technical Report TR8350.2, 3rd Edition
- Noordam, J. E. 2004, *SPIE*, 5489, 817
- Noordam, J. E., et al. 2008, The MIM principle, version 2.0, in preparation
- Pearson, T. J., & Readhead, A. C. S. 1984, *ARA&A*, 22, 97
- Perley, R. A. 1989, *ASPC*, 6, 259
- Perley, R. A. 1989, *ASPC*, 6, 287
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C*, 2nd Edn (Cambridge, UK: Cambridge University Press)
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, *A&AS*, 124, 259
- Röttgering, H. J. A., et al. 2006 [[arXiv:astro-ph/0610596](http://arxiv.org/abs/astro-ph/0610596)]
- Schwab, F. R. 1981, *VLA Scientific Memo* 136
- Schwab, F. R. 1984, *AJ*, 89, 1076
- Subrahmanya, C. R. 1991, *ASPC*, 19, 218
- Swarup, G. 1991, *ASPC*, 19, 376
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2001, *Interferometry and Synthesis in Radio Astronomy*, Second Edition, Wiley, New York, USA
- van der Tol, S., & van der Veen, A. J. 2007, Proceedings of the IEEE International Symposium Signals, Circuits, Systems, Iasi, Romania
- van der Tol, S., Jeffs, B. D., & van der Veen, A. J. 2007, *IEEE Transactions on Signal Processing*, 55, 9, 4497
- van Velthoven, P. F. J. 1990, Ph.D. Thesis, Eindhoven University of Technology, Eindhoven, The Netherlands
- Voronkov, M. A., & Wieringa, M. H. 2004, *Exper. Astron.*, 18, 13
- Wieringa, M. H. 1992, *Exper. Astron.*, 2, 203