Self-Calibration of Arrays whose Elements are Strongly Polarized

Invited paper

SPIE Internatl. Symp. on Astronomical Telescopes and Instrumentation
Munich, March 27-30, 2000; SPIE Proceedings Vol. 4015.

Johan Hamaker

Netherlands Foundation for Research in Astronomy NFRA/ASTRON,
P.O. Box 2, 7990 AA Dwingeloo, Netherlands

ABSTRACT

Traditional calibration methods based on scalar self-calibration rely on a linearised *quasi-scalar* approximation. Several underlying assumptions will no longer hold for the phased-array telescopes being considered in this conference. The reason is that the element antennas will show a strong instrumental polarization, that is moreover strongly dependent on the pointing of the array. Scalar selfcal cannot be used for calibrating such an array.

This paper offers a rigourous description, based upon 2×2 matrices, that fully accounts for polarisation phenomena in an interferometer. It then translates the traditional selfcal algorithm into this matrix language by exploiting the analogies between scalar and matrix multiplications.

It is shown that the matrix algorithm does not yield a complete calibration: It only *aligns* all antenna-based errors by suppressing the scattering of radiation away from source components to places in the image that should be empty. In doing so, it satisfies the requirement for a high *dynamic range*.

However, it admits an unknown uniform in-place transformation, the *poldistortion*, of the matrix brightness. In terms of Stokes brightness parameters, the *polvector* (Q,U,V) in Stokes-vector space is rotated in an unknown way and there is an unknown mutual *polconversion* between Stokes I and the polvector. One must eliminate this poldistortion to make the image a faithful rendition of the source.

Unpolarized sources can be used as calibrators to suppress the polconversion effect, after which prior statistical knowledge about the orientations and ellipticities of the antenna elements serves to eliminate most of the polrotation, much in the same way as in the quasi-scalar method. For *homogeneous* arrays of identical elements, one supplementary phase measurement of some sort is required; for *heterogeneous* arrays this is not even necessary.

In the absence of unpolarized calibrators, polconversion must be eliminated by other means. There are no obvious ways of achieving this: Developing suitable calibration techniques will be a major challenge for the new generation of telescopes.

In the concluding section, the consequences of my findings for the design of phased-array aperture-synthesis telescopes are explored.

Keywords:

1. INTRODUCTION

The twenty-first century is likely to become the era of arrays built of mass-producible elements that have few of no moving parts. The transition from (arrays of) steerable paraboloids to this new type of telescope will bring about many changes, some of which are being explored in this conference. There are the more or less obvious changes in production, testing, maintenance and operational methods, the volume of the data and the complexity of its processing.

In addition to this, there is the very important question whether the usual approximations underlying our current thinking about radio interferometry will remain tenable. As I shall show, the answer is a categoric 'no': At a

E-mail: hamaker@nfra.nl

fundamental level, the polarization that is an unalienable property of both the electromagnetic field and our telescopes cannot adequately be represented by scalar equations. The approximations of the past century will no longer do. Instead, we need a theory that acknowledges the vector nature of the radiation field from the very start.

The theory to be presented here is based on 2×2 coherency and transfer ('Jones') matrices that replace the traditional scalars. It bears a strong formal resemblance to the familiar scalar one and may, to some extent, look deceptively simple. Clearly, both the noncommutativity of matrix multiplication and the fourfold content of the matrices introduce a number of new degrees of freedom which give rise to a range of important phenomena that have no scalar analogue; these will be discussed.

Apart from that, we have to revise our basic notions about how e.g. an antenna is to be described. It is customary to think to first order of the beam of a paraboloid as a scalar function of angular position. To second order, this description must be refined by introducing a (usually direction-dependent) instrumental polarization effect that is known as cross-polarization. The distinction between the two effects is a rational one because the latter is small and can for many purposes be neglected. As a matter of fact, our antennas conform so well to this description that we are easily led to believe that it is a universal one representing general antenna physics. This is not true, as can easily be demonstrated for a simple example.

Consider a phased array of crossed error-free short dipoles in a horizontal plane. If we assume ideal amplification, phasing and summation of their signals, the polarisation properties of the array beam are identical to those of the individual dipole pairs. Consider radiation incident in the vertical plane of the x dipole at a zenith angle ϕ. The component intercepted by the x dipole will be attenuated by the projection factor $\cos\phi$, that picked up by the y dipole will remain unattenuated. For this direction, the beam is strongly polarized and the scalar approximation makes little sense. The same is true for any other direction away from the zenith (except perhaps near the vertical planes at 45 degrees to both dipoles).

The only way to avoid this effect with static antenna elements would be to make their E- and H-plane responses closely identical (as one does with the feeds of parabolic dishes) over the direction range of interest. This may be possible in principle, but it would require elements considerably more complex than a simple dipole. Such complexity seems undesirable because it probably implies higher costs and certainly a greater sensitivity to production tolerances and other sources of error, such as mutual coupling of the elements. At any rate, feasible designs are still likely at best to represent crude approximations to the desired ideal, and to be beset with a degree of polarization that is large compared to our current standards.

The bottom line is that, with some exceptions, we must ban scalars from the theory of interferometry and put 2×2 matrices in their place. The 'new' effects that this leads to are not really new in the sense that they did not exist before. Matrix theory just brings them out more clearly and provides a full view of things that were obscured or hidden by the traditional quasi-scalar approach. Rather than complicate matters as it may seem to do at first, it helps us to see interferometers as they really are and to fully exploit their possibilities.

A number of important statements will be made and theorems stated without proof. A full account of the theory is to be found in a series of papers published elsewhere [1,2,3,4]. Most if not all of the proofs omitted here are provided in the latter of these papers.

1.1. Terminology and notation

Since I will be discussing scalar selfcal and its full-polarization analogue side by side, it is necessary to put a precise terminology in place.

The analogue of the scalar visibility is the *coherency*, a tensor that I represent in the form of a 2×2 matrix. It consists of four components that I shall occasionally refer to as *visibilities*.

The device that converts the electromagnetic field vector into a pair of voltages is called a *feed*; it consists of two *receptors* that are usually (but not necessarily) sensitive to nominally opposite polarizations. In a *homogeneous* array all feeds are nominally identical; in a *heterogeneous* array they differ.

The imaging process that I consider consists of the observations proper followed by a process of self-calibration. In the latter, *models* of the instrumental errors and the source brightness distribution are developed jointly in an iterative procedure. It is assumed to have converged when the models together correctly represent the observed coherencies within the noise. The *image* is the pictorial representation of the source model; the two words are almost

synonymous. The model itself may take various forms; its essential property is that it can be used to *predict* model coherency values that can be compared to those actually observed in order to estimate instrumental errors.

Vectors are denoted by bold lowercase symbols; bold uppercase represents matrices. Constant scalars and vectors are shown in roman, variables in italic font. A unit vector in the direction $\boldsymbol{x}$ will be denoted by $\mathbf{1}_{\boldsymbol{x}}$. The 'dagger' superscript † stands for hermitian transposition or conjugation, i.e. transposition plus complex conjugation.

Primes are generally used to distinguish *observed* or *fitted* values from *true* ones; occasionally they will also be used to distinguish input and output of a transformation or values of one variable under different conditions.

Quantities in the signal domain carry a single antenna subscript j or k; those in the coherency domain get an interferometer subscript jk. An additional subscript t will be used to indicate successive integration intervals or 'time slices'. The array consists of N antennas and an observation comprises M integration intervals.

2. COHERENCY-MATRIX FORMULATION OF INTERFEROMETRY

2.1. The scalar/matrix analogy

Scalar form	Matrix form	Section(s)
Arbitrary scalar a	Arbitrary 2×2 matrix $\boldsymbol{A}$	
Unity $= 1$	Identity 2×2 matrix $= \mathbf{I}$	
Phase factor $\exp i\alpha$	Unitary 2×2 matrix $\boldsymbol{X}$ Unimodular unitary 2×2 matrix $\boldsymbol{Y}$	
Positive real number $\lvert a \rvert$	Positive hermitian 2×2 matrix $\boldsymbol{G}$ Unimodular pos.-herm. 2×2 matrix $\boldsymbol{H}$	
Polar representation $a = a \exp i\alpha$	Polar representation $\boldsymbol{A} = a \exp i\alpha\, \boldsymbol{HY}$	6.1
Complex conjugation a^*	Hermitian transposition $\boldsymbol{A}^\dagger \equiv \boldsymbol{A}^{*\mathrm{T}}$	
Multiplication $c = ab = ba$	Multiplication $\boldsymbol{C} = \boldsymbol{AB} \neq \boldsymbol{BA}$	
Field or voltage transfer $e'_j = g_j e_j$ $g_j =$ (complex) gain	Field or voltage vector transfer $\boldsymbol{e}'_j = \boldsymbol{J}_j \boldsymbol{e}_j$ $\boldsymbol{J}_j =$ (complex) Jones matrix	2.3
Visibility $e_{jk} = < e_j e_k^* >$	Coherency $\boldsymbol{E}_{jk} = < \boldsymbol{e}_j \boldsymbol{e}_k^\dagger >$	2.2
Visibility transfer $e'_{jk} = g_j e_{jk} g_k^*$	Coherency transfer $\boldsymbol{E}'_{jk} = \boldsymbol{J}_j \boldsymbol{E}_{jk} \boldsymbol{J}_k^\dagger$	2.3

Table 1. Analogies between scalars and 2×2 matrices, their algebraic properties and their application in interferometry. Particulars are to be found in the sections listed.

The algebraic properties of scalars and matrices are very similar. Every elementary property of scalars has an immediate matrix counterpart, with one very important exception, viz. that matrix multiplication is non-commutative.

The analogy extends further. There is, for example, a matrix counterpart of the polar representation $a = |a| \exp i\alpha$ of complex numbers. An overview is given in Table 1.

2.2. The coherency matrix

The proper way to represent the electric field of radiation moving in the z direction of a cartesian coordnate frame is by a vector

$$\boldsymbol{e} = \begin{pmatrix} e_{\mathrm{x}} \\ e_{\mathrm{y}} \end{pmatrix}$$

The equivalent of the scalar visibility is the *coherency tensor*. It may be represented in various ways; the representation I use here is the *coherency matrix* [4,5]:

$$\boldsymbol{E}_{jk} = < \boldsymbol{e}_j \, \boldsymbol{e}_k^\dagger > = \begin{pmatrix} < e_{j\mathrm{x}} e_{k\mathrm{x}}^* > & < e_{j\mathrm{x}} e_{k\mathrm{y}}^* > \\ < e_{j\mathrm{y}} e_{k\mathrm{x}}^* > & < e_{j\mathrm{y}} e_{k\mathrm{y}}^* > \end{pmatrix} \tag{1}$$

2.3. The interferometer equation

The elements in the signal path in one antenna linearly transform the electric field or voltage vector. The transformations may be represented by equations of the form

$$\boldsymbol{w}_j = \boldsymbol{J}_j \boldsymbol{e}_j$$

where $\boldsymbol{J}_j$ is called a *Jones matrix*. It is then readily seen that an interferometer with Jones matrices $\boldsymbol{J}_j$ and $\boldsymbol{J}_k$ transforms the coherency matrix according to

$$\boldsymbol{W}_{jk} = \boldsymbol{J}_j \boldsymbol{E}_{jk} \boldsymbol{J}_k^\dagger \tag{2}$$

Coherency and Jones matrices having the same form, it should be clear from the context which is which, just as in the scalar domain. In addition, note that Jones matrices carry the single index of an antenna whereas the coherency matrices have a double, interferometer index. This difference will remain also when we later add another index t for sampling time.

I note in passing that in the particular case of a single dish, $j = k$ and Eq. (2) reduces to

$$\boldsymbol{W} = \boldsymbol{J} \boldsymbol{E} \boldsymbol{J}^\dagger \tag{3}$$

which is known as a *congruence transformation*. We will see in Sect. 5.1 that this same transformation describes a 'self-aligned' synthesis array.

3. APERTURE SYNTHESIS

3.1. The van Cittert-Zernike theorem

Aperture synthesis is based on the van Cittert-Zernike theorem [5,6,7], which (in its original form) states that there is a Fourier-transform relation between the scalar *brightness* $B(\boldsymbol{l})$, a function of sky direction $\boldsymbol{l}$, and the scalar *visibility* $v(\boldsymbol{r})$ which is a function of vector separation or *baseline* $\boldsymbol{r}$. * The visibility is defined as the spatial autocorrelation of the electric field:

$$v(\boldsymbol{r}) = < e(\boldsymbol{r}_0) \, e^*(\boldsymbol{r}_0 + \boldsymbol{r}) > \tag{4}$$

where $\boldsymbol{r}_0$ is an arbitrary reference position. We may approximate the van Cittert-Zernike relation by the sum

$$v(\boldsymbol{r}) = \sum w(\boldsymbol{r}, \boldsymbol{l}) \, B(\boldsymbol{l}) \tag{5}$$

where the $w(\boldsymbol{r}, \boldsymbol{l})$ are the values assumed by the Fourier kernel.

Rather than scanning B in the conventional way with an antenna beam, in aperture synthesis we sample the visibility by literally following the prescription of Eq. (4): Two antennas sample the field at positions separated by $\boldsymbol{r}$ and after amplification the samples are correlated and averaged in time. Collecting samples for many different values of $\boldsymbol{r}$, we may eventually reconstruct $v(\boldsymbol{r})$ and hence the brightness $\boldsymbol{B}(\boldsymbol{l})$.

The reconstruction problem is often an ill-conditioned one and we have to constrain the solution with prior knowledge to obtain a credible result. The most important constraint is that the sky is 'mostly empty', i.e. contains (like the visible sky) much empty space between isolated sources. I shall use this constraint but not discuss the way in which it is actually applied.

*One may note that it is the exact spatial analogon of the Wiener-Khinchin theorem that relates a signal's power spectrum and temporal autocorrelation.

3.2. Matrix brightness

The van Cittert-Zernike theorem can readily be generalised to show that each element of the coherency matrix is the Fourier transform of the corresponding element of a *brightness matrix*

$$\boldsymbol{E}(\boldsymbol{r}) \equiv \begin{pmatrix} e_{\rm xx}(\boldsymbol{r}) & e_{\rm xy}(\boldsymbol{r}) \\ e_{\rm yx}(\boldsymbol{r}) & e_{\rm yy}(\boldsymbol{r}) \end{pmatrix} \stackrel{\mathcal{FT}}{\longleftrightarrow} \boldsymbol{B}(\boldsymbol{l}) \equiv \begin{pmatrix} B_{\rm xx}(\boldsymbol{l}) & B_{\rm xy}(\boldsymbol{l}) \\ B_{\rm yx}(\boldsymbol{l}) & B_{\rm yy}(\boldsymbol{l}) \end{pmatrix} \tag{6}$$

In describing the observing process, I introduce a third index t to represent the successive integration intervals. Since the Earth is rotating, the successive intervals are associated for each interferometer with different baseline vectors $\boldsymbol{r}_{jkt}$, so the coherencies $\boldsymbol{E}_{jkt}$ that we measure are samples $\boldsymbol{E}(\boldsymbol{r}_{jkt})$ of the continuous coherency function $\boldsymbol{E}(\boldsymbol{r})$. In selfcal theory the integration intervals are assumed to coincide for all interferometers.

Each sample is an appropriate Fourier integral over $\boldsymbol{B}(\boldsymbol{l})$. Aanlogously to Eq. (5) I write

$$\boldsymbol{E}_{jkt} \;=\; \sum_{\boldsymbol{l}} w(\boldsymbol{r}_{jkt},\boldsymbol{l})\; \boldsymbol{B}\,(\boldsymbol{l}), \qquad j,k = 1,\ldots,N, \quad t = 1,\ldots,M \tag{7}$$

where the $w(\boldsymbol{r}_{jkt},\boldsymbol{l})$ are again values of the Fourier kernel.

3.3. Stokes brightness

The Stokes brightness [4,5,6] (I,Q,U,V) is also a function of $\boldsymbol{l}$, defined by the transformation

$$\boldsymbol{B} = \begin{pmatrix} I+Q & U-iV \\ U+iV & I-Q \end{pmatrix} = I\,\mathbf{I} + Q\,\mathbf{Q} + U\,\mathbf{U} + V\,\mathbf{V} \tag{8}$$

where

$$\mathbf{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \mathbf{U} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \mathbf{V} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \tag{9}$$

The matrices **I**, **Q**, **U** and **V** are known in physics as the *Pauli (spin) matrices.*

The Stokes parameter I is the total brightness or *intensity.* For (Q,U,V) a proper name is the 'polarized-brightness vector'; more conveniently, I call it the *polvector.* The dichotomy between I and the other Stokes parameters can be understood as a consequence of the Pauli matrix **I** being the identity matrix. The domain of the polvector is closely related to that of the Poincaré sphere [5].

It is convenient to introduce a symbolic shorthand for Eq. (8):

$$\boldsymbol{B}(\boldsymbol{l}) = [\,I(\boldsymbol{l}) + \boldsymbol{p}(\boldsymbol{l})\,] \tag{10}$$

where $\boldsymbol{p} \equiv (Q,U,V)$ is the polvector.

4. SELF-CALIBRATION

The invention of self-calibration [6,7] marks a watershed in the history of radio astronomy. It is a data processing technique that enabled us in a single stroke to turn inherently unstable interferometer arrays into superb imaging instruments whose dynamic range is measured in the hundred thousands. (And it is hoped that it can be stretched cosiderably further.)

To get the maximum out of the expensive steelware of N antennas, radio astronomers build correlator systems that implement most or all possible interferometers, their number totalling a maximum of $N(N-1)/2$. With modern electronics and digital techniques, these correlators are virtually error-free. It follows that all the gain and phase errors in the system are attributable to the individual antennas. Thus, their effect is present in the interferometers in a highly redundant form, albeit in combination with the actual visibilities of the source observed.

Self-calibration or *selfcal* is a (hitherto scalar) method of solving simultaneously for the source brightness $B(\boldsymbol{l})$ and the antennas' complex-gain errors. It is an iterative method alternating between two stages:

- Use the current estimate of the brightness distribution and the observed visibilities to improve the current estimate of the antenna errors.
- Use the current estimate of the antenna errors to correct the observed visibilities and improve the current estimate of the brightness.

It is not difficult to imagine that this kind of procedure will converge to some limiting image, but we have no formal proof that this limit is unique. Cases can be invented where this would clearly not be the case. The extensive experience built up over two decades shows, however, that the method is usually quite robust. Except for the overall brightness scale which selfcal leaves undefined, the astrophysically interesting features come out the same, independently of e.g. the way selfcal is initialised or twiddled through its steering parameters.

Scalar selfcal can be applied only with homogeneous arrays. In such arrays, the receptors form two independent subsets (X and Y or L and R) whose responses to the polarized radiation field can be described by scalar equations. The view that I shall take here of scalar selfcal is primarily aimed at providing for a generalisation to the matrix case. At the same time, it will provide another argument (but again: not a proof) for the uniqueness of the final solution. As such, it is interesting in its own right.

4.1. Scalar self-calibration

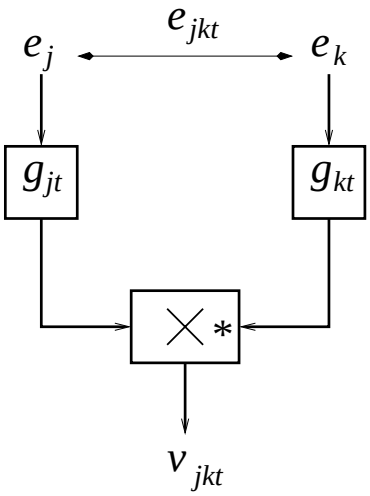

Figure 1. Model for a scalar interferometer. The complex gains $g_j,\ g_k$ may be in error; the correlator is assumed to be ideal.

Selfcal works on the basis of two assumptions:

- All instrumental effects are *antenna-based*; the correlator is error-free (fig. 2). Thus our observed visibility is given by

$$v(\boldsymbol{r}_{jkt}) = g_{jt}\, e(\boldsymbol{r}_{jkt})\, g^*_{kt} \tag{11}$$

- The sky is 'relatively empty': The source brightness is nonzero only in a minor fraction of the observed field, the source's *support*. (In practice, it turns out that the support need not be known a priori, but can be found and successively refined by inspection of provisional 'dirty' images.)

In the convergence limit, selfcal seeks to match the combination of complex antenna gains g'_{jt} and visibilities $e'(\boldsymbol{r}_{jkt})$ to the observations. Obviously, one possible match is provided by the *true* gains and visibilities. Admitting

the possibility of other solutions, let us allow the fitted gains g'_{jt} to differ from the true ones: $g'_{jt} = g_{jt}x_{jt}^{-1}$, and require the fitted visibilities to satisfy the equations

$$g_{jt}\, x_{jt}^{-1}\; e'(\boldsymbol{r}_{jkt})\; x_{kt}^{*-1}\, g_{kt}^{*} \quad = \quad v(\boldsymbol{r}_{jkt}) = g_{jt}\; e(\boldsymbol{r}_{jkt})\; g_{kt}^{*} \tag{12}$$

where $v(\boldsymbol{r}_{jkt})$ are the measured visibilities.

A solution is given by *any* set of values for the x_{jt} with a corresponding set of visibility samples

$$v'(\boldsymbol{r}_{jkt}) = x_{jt}\; v(\boldsymbol{r}_{jkt})\; x_{kt}^{*}, \tag{13}$$

These visibilities in turn correspond to a source model B' through Eq. (5):

$$x_{jt}\; v(\boldsymbol{r}_{jkt})\; x_{kt}^{*} = \sum_{\boldsymbol{l}} w(\boldsymbol{l}, \boldsymbol{r}_{jkt})\; B'\,(\boldsymbol{l}), \tag{14}$$

If the source support is limited as assumed, the sum contains only a limited number L of terms for which $B'(\boldsymbol{l})$ can differ from zero. For a properly conditioned observation, the number of visibility samples is much greater than that of unknowns: MN values of the x_{jt} plus L values of $B'(\boldsymbol{l})$.

The system is overdetermined; yet we have already seen that it admits of at least one solution. That solution is actually a member of a whole family: Indeed, if all the x_{jt} equal one value x, Eq. (14) can be rewritten as

$$x\; v(\boldsymbol{r}_{jkt})\; x^{*} = \sum_{\boldsymbol{l}} w(\boldsymbol{l}, \boldsymbol{r}_{jkt})\; x\; B(\boldsymbol{l})\; x^{*}$$

which defines a brightness solution

$$B'(\boldsymbol{l}) = x\; B(\boldsymbol{l})\; x^{*} \tag{15}$$

The fitted source model B' is identical to B except for a positive scale factor, and is obviously confined to the support of B. The selfcal method is indifferent to this scale factor and a practical algorithm will insert an arbitrary value. The proper scale must be determined by other means.

Other solutions are unlikely to exist. If we should allow the x_{jt} to take independent values, this would result in scattering of brightness away from the source to other parts of the image. It is reasonable to conjecture that it is impossible for any 'wild' combination of x_{jt} values to produce a false brightness image that nonetheless vanishes everywhere outside the source support. Practical experience of two decades supports this conjecture, — but I reiterate that a formal proof is lacking and the solution may not always be robust against the effect of noise.

4.2. Calibration versus alignment

The solution being non-unique, self-*calibration* is actually a misnomer: To really calibrate the image we need an independent fix of the brightness scale. What selfcal does achieve is to reduce all the errors x_{jt} in the individual visibility measurements to a single value x: It *lines up* the measurements, forcing them all to conform to one common unknown scale factor. As a result, extremely high dynamic ranges can be attained even though the absolute brightness scale is unknown.

Self-*alignment* describes more accurately what the procedure actually achieves. The distinction is a bit academic here, but becomes significant in the matrix version of 'selfcal'.

5. MATRIX SELF-ALIGNMENT

5.1. Self-alignment

The argument of the preceding section carries over in its entirety to the matrix domain. We may follow through the same steps, simply replacing all scalar gains and visibilities by the corresponding 2×2 matrices. We must now solve the matrix equivalent of Eq. (12):

$$\boldsymbol{J}_{jt}\, \boldsymbol{X}_{jt}^{-1}\; \boldsymbol{E}'(\boldsymbol{r}_{jkt})\; \boldsymbol{X}_{kt}^{*-1}\, \boldsymbol{J}_{k}^{*} = \boldsymbol{J}_{jt}\; \boldsymbol{E}(\boldsymbol{r}_{jkt})\; \boldsymbol{J}_{kt}^{*} \tag{16}$$

and the final result is the analogue of Eq. (15)

$$\boldsymbol{B}'(l) = \boldsymbol{X}\,\boldsymbol{B}(l)\,\boldsymbol{X}^{\dagger} \tag{17}$$

This equation is known as a *congruence transformation*. I shall follow mathematicians in referring to Eq. (17) as 'the congruence transformation $\boldsymbol{X}$', i.e. using the name of $\boldsymbol{X}$ as a synonym for the transformation that it effects: The source model $\boldsymbol{B}'$ is related to the true source $\boldsymbol{B}$ by an unknown congruence transformation $\boldsymbol{X}$.

Like Eq. (15) for scalar selfcal and with the same proviso, this is a basic relation that any matrix self-alignment solution must satisfy. And as for the scalar case, the appearance of this unknown transformation $\boldsymbol{X}$ is fundamental and unavoidable. I call it the *poldistortion*.

Although this result is formally the same as for scalar selfcal, it is worth some thought. As in scalar selfcal, all the, probably time-varying, errors in the observation have given way to a *single* poldistortion representing a set of unknown errors that is *constant* over the observation. In matrix self-alignment, the time-varying errors to be converted into unknown constants include

- Time-variable complex antenna gains.
- The variation in orientation of the antenna feeds relative to the object observed ('parallactic angle') that occurs in alt-az antennas. (Normally one would not rely on self-alignment to eliminate these, but rather 'predict' and correct for them on the basis of a geometric model.)
- Variations in the feed parameters (mainly the 'leakage' or 'D' terms in the quasi-scalar jargon)
- Variations in ionospheric Faraday rotation.

Moreover, all this is true not only for a single observation contiguous in time, but also for a set of observations spaced over a time interval in which the source does not change.

A schematic of the combined self-alignment and poldistortion elimination procedure is shown in fig. 1. As in scalar selfcal, we may first correct for the errors that we know of. My argument does not require this; however, the actual selfcal *algorithm* starts with an initial image, to be made from the raw observations, that must be good enough to extract a reasonable initial source model from it. If such a model can be obtained otherwise, e.g. from prior knowledge about the source, the initial corrections may just as well be omitted: They will automatically be subsumed in the corrections to be derived in self-alignment.

Whether or not we apply the prior corrections is likely to affect the poldistortion in the final solution, *but not our ignorance about its value*. No matter how we arrive at a self-aligned image, we *must* assume that it contains an unknown poldistortion and undertake to eliminate it. This problem will be discussed below.

As an aside, I observe that Fourier transforming Eq. (3) for a single dish, one obtains an equation of the same form as Eq. (17): The self-aligned array is equivalent to a single dish with unkown Jones matrix.

5.2. Scattering dynamic range and polarimetric fidelity

The solutions Eqs. (15) and (17) of the selfcal/self-alignment problem represent *in-place* transformations of the brightness: A simple uniform scaling in the scalar case, a uniform poldistortion in the matrix case. Their common characteristic is that they do not *scatter* radiation out of any point of the source into any other position in the image. In the scalar case, the scaling being the same for the entire image means that our image is a faithful (although scaled) replica of the source: *Dynamic range* is conserved, - and it is this feature that makes selfcal such an important and valuable tool.

In the matrix case, image fidelity and dynamic range are no longer synonymous. Self-alignment suppresses spatial scattering in the image; the residual scattering defines to what extent weak structures remain recognisable in the presence of strong features elsewhere in the source. The concept of dynamic range is appropriate to describe this effect. As in the scalar case, it is no more than a semi-quantitative measure of the quality of the image, and I do not bother to give a formal definition.

On top of the effect of scattering comes the poldistortion that is *independent* of it. Even if we were to produce a truly scatter-free image it would still misrepresent the source in an unknown way: As a complement to dynamic range we must consider the question of the *polarimetric fidelity* of the image.

5.3. Example: Faraday rotation

The concepts developed above and the benefits of the matrix approach can be neatly illustrated on the example of ionospheric Faraday rotation. This rotation changes the observed position angle of linear polarization. Over a synthesis observation, it varies and this results in scattering of linearly polarized brightness in the final image.

Corrections for the effect are generally based on external data: Ionosphere models and ground- and satellite-based measurements [6]. Often the results leave much to be desired. In some cases, the external correction can be improved upon by noting the apparent rotation of linear polarization during the observation. One might call this *Faraday self-aligment*. Obviously, it can only eliminate *variations* in the rotation; to find the true position angle of linear polarization one must determine the Faraday rotation's zero point by other means.

In the matrix approach, the adjustment of Faraday rotation is no longer a separate operation: It is subsumed in the overall self-alignment process. A scatter-free image results directly. The absolute rotation and position angle of linear polarization remain undetermined: This indeterminacy is now recognised as part of the poldistortion that is the unavoidable by-product of self-alignment. Obtaining a good dynamic range and correct rendition of the polarization appear as two distinct problems that must be independently addressed.

6. THE POLDISTORTION

6.1. Polrotation and polconversion

An arbitrary square matrix $\boldsymbol{X}$ can be subjected to a *polar decomposition*

$$\boldsymbol{X} = x\boldsymbol{H}\boldsymbol{Y} = x\boldsymbol{Y}\boldsymbol{H}'$$

where

- x is a complex constant;
- $\boldsymbol{Y}$ is a *unimodular unitary* matrix (i.e. $\boldsymbol{Y}\boldsymbol{Y}^\dagger = \mathbf{I}$ and $\det \boldsymbol{Y} = 1$);
- $\boldsymbol{H}$ is a *unimodular positive hermitian* matrix (i.e. $\boldsymbol{H} = \boldsymbol{H}^\dagger$, $\det \boldsymbol{H} = 1$ and $\mathrm{Tr}\ \boldsymbol{H} > 0$); so is $\boldsymbol{H}'$.

Applying the polar decomposition to Eq. (17) we get

$$\boldsymbol{B}' = xx^*\, \boldsymbol{H}\ (\boldsymbol{Y}\boldsymbol{B}\boldsymbol{Y}^\dagger)\ \boldsymbol{H}^\dagger \tag{18}$$

The positive scaling factor is the same as in scalar selfcal and we may ignore it for the same reasons. Apart from it, $\boldsymbol{B}'$ is derived from $\boldsymbol{B}$ through a succession of two specific transformations.

The first is the unitary transformation $\boldsymbol{Y}$. Its effect is to leave the intensity unchanged (naturally, since $\boldsymbol{Y}\mathbf{I}\boldsymbol{Y}^\dagger = \mathbf{I}$) and to *rotate the polvector* in its three-dimensional space. In the notation of Eq. (10)

$$[\, I' + \boldsymbol{p}'\,] = [\, I + \boldsymbol{R}\,\boldsymbol{p}]$$

where $\boldsymbol{R}$ is a rotation operator. I call $\boldsymbol{Y}$ the *polrotation* (transformation). It can be characterised by its *Gibbs polvector*

$$\mathbf{1}_{\boldsymbol{y}} \sin\eta,$$

where $\boldsymbol{y}$ is the direction of the rotation axis in polvector space and 2η the rotation angle. Obviously, any multiple of the Gibbs vector is an *eigenvector* of the transformation: $\boldsymbol{R}\,\mathbf{1}_{\boldsymbol{y}} = \mathbf{1}_{\boldsymbol{y}}$.

Like the polrotation, the positive hermitian transformation $\boldsymbol{H}$ is characterised by a Gibbs polvector

$$\mathbf{1}_{\boldsymbol{h}} \sinh\gamma$$

The transformation exchanges brightness between the intensity and that component of the polvector that is parallel to $\boldsymbol{h}$; any perpendicular component is not affected and is therefore an eigenvector. I call $\boldsymbol{H}$ the *polconversion*. We may summarise the above by stating that the self-aligned source model $\boldsymbol{B}'(\boldsymbol{l})$ is related to the true brightness $\boldsymbol{B}(\boldsymbol{l})$ by an in-place poldistortion transformation that is the product of a polrotation, a polconversion and a positive scale factor.

The poldistortion $\boldsymbol{X}$ is far more complicated than a simple scale error. As we have seen, polrotation and polconversion each contain three unknown parameters (the cartesian components of their Gibbs vectors) which, together with the scale factor, make a total of seven.

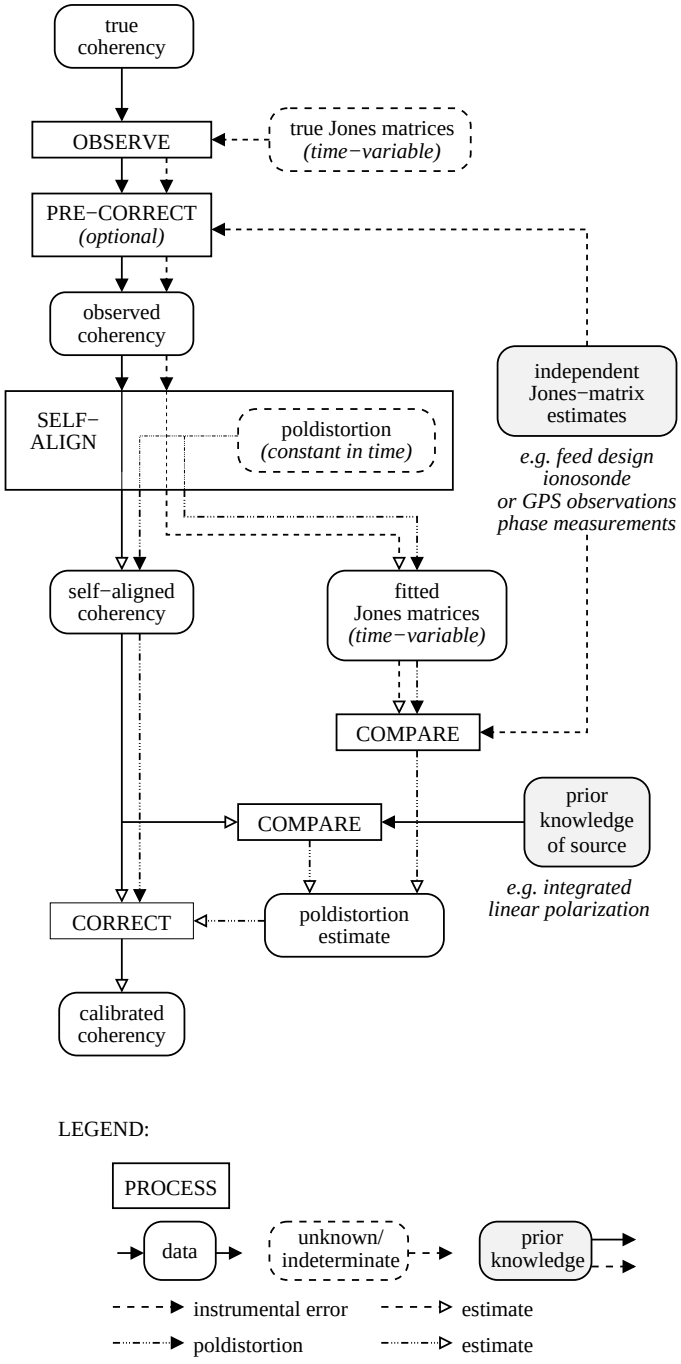

Figure 2. Flow diagram of the self-alignment and poldistortion calibration procedure. The crux of the matter resides in the self-alignment: It effects a split between the true coherencies and the time-variable part of the observing errors. The price to be paid is the insertion of an unknown constant error artefact, the poldistortion, in the estimated coherencies and its inverse in the estimated Jones matrices. To eliminate it, we must compare these estimates with a priori instrumental and astronomical information.

7. CONTROLLING THE POLDISTORTION

With self-alignment only the first half of the calibration job is done. To eliminate the poldistortion, we must bring other information to bear on our problem. It may take the form of either prior knowledge or additional measurements; both may be used either to constrain the self-alignment algorithm so as to (partly) suppress the poldistortion, or to remove the poldistortion afterwards.

For example, if we have reason to believe that our source field is unpolarized, we can impose this condition on our image and source model. Alternatively, we may allow self-alignment to produce a polarized image and determine and remove the polconversion afterwards, on the basis that in the proper image certain sources must be unpolarized.

Apart from the image $\boldsymbol{B}'$, self-alignment also yields estimates $\boldsymbol{J}'_j = \boldsymbol{J}_j\boldsymbol{X}^{-1}$ of the antenna Jones matrices. Comparing these with what we know about the true values gives us another handle on $\boldsymbol{X}$.

Current observational practice suggests that it is not necessary to control poldistortion to the same extreme that scattering must be controlled. Errors in the order of 0.1 to 1% in the measured relative polvector $\boldsymbol{p}/I$ are good enough for much meaningful astrophysics. This is fortunate because it may be very difficult to do better except in special cases.

7.1. Polrotation calibration

If one has access to sources known to be unpolarized, poldistortion is easily suppressed by imposing the same condition upon the corresponding images. Indeed, if the polvector vanishes in both source and image, conversion from unpolarized to polarized brightness cannot have taken place. Mathematically, this result is easily confirmed.

To deal with the polrotation, either prior knowledge or special measurements must be invoked. Usually, the orientations and ellipticities of the feeds are well known (from the geometries and/or lab measurements) and these can be used to eliminate two components of the polrotation. The remaining one corresponds to a receiver phase offset between the X and Y or L and R subsystems in conventional homogeneous arrays such as the ATCA, the VLA and the WSRT. The only way to fix it is through an additional measurement, either of a known polarized source or of a man-made pilot signal [2,4,6,7].

7.2. Heterogeneous arrays

An interesting new case arises when we relax the conventional assumption that the array is homogeneous. This condition is usually fulfilled for obvious engineering reasons and later became a conditio sine qua non for quasi-scalar self-calibration, but in the matrix formulation it can be dropped. It turns out that in a heterogeneous array the phase offset is tightly coupled to the feed properties; thus, constraining the latter to be close to their nominal values, one also constrains the receiver phases and no separate measurement is needed. The correctness of this assertion has been tested in simulations, and its practical usefulness is being explored in a special ad-hoc observation made with rotated feeds by the Westerbork Telescope.

This new result is not entirely unexpected: Indeed, over a quarter-century ago, Weiler [8] showed that an array of paraboloids with mostly N-S/E-W oriented linear dipole receptors, but with a few of them rotated over 45 degrees, can be calibrated completely (i.e. including receiver phases) to first order by observing an unpolarized calibrator alone. To support this finding in the more general form given here, several arguments can be advanced:

Most fundamentally, the nominal receptor characteristics in the heterogeneous case are arbitrarily distributed. There is no natural way to split the receptors into two disjoint subsets between which an asymmetry such as the phase difference might arise. The only way out is for the difference not to exist.

Another viewpoint is that the homogeneous case is, in physical terms, a *degenerate* one. The degeneracy results in a decoupling of the feed and phase characteristics which are normally coupled. This viewpoint is supported by the mathematics of the problem [4].

A third viewpoint is that a set of identical feeds defines a preferred direction in polvector space, viz. that of the polvectors to which the two receptors are matched (e.g. the V axis for circular feeds). This results in an asymmetry in the characteristics of the instrument that is directly related to the phase-difference problem. 'Randomised' feed characteristics destroy this preferential direction: The phase difference evapourates and all polarizations can be measured equally well, — or equally poorly.

7.3. General poldistortion calibration

What if we have no unpolarized calibrator sources? We have seen how to control polrotation once we have rid ourselves of poldistortion, but now we have no simple way to do so.

Frequently, the fields observed do contain background sources or parts of the target source itself that may be presumed to be unpolarized. In this case, the self-aligned image can be analysed to estimate and eliminate the poldistortion, after which the polrotation can be dealt with. Matrix theory provides a better framework for doing this and the Westerbork observation already referred to should provide a good test case.

In VLBI, the fields observed are small and hardly contain any background sources. At higher frequencies and finer resolutions, unpolarized sources become hard to find. Instead, observers are pioneering the use of the parallactic-angle variation in alt-az telescopes as a device to separate poldistorted intensity from true source polarization [9]. Considering this method from the viewpoint of matrix theory, I doubt its robustness, but this tack should also be pursued in the matrix framework [4].

8. IMPLICATIONS FOR PHASED-ARRAY INTERFEROMETER ARRAYS

It has already been argued that phased arrays of immovable antennas will require a radical departure from our past and present frame of mind. The notion of basically scalar fields and antennas in which polarization effects represent small departures must be abandoned and we must embrace the matrix formalism of this paper in its place.

As has been shown, this introduces a new degree of freedom in our designs by allowing for heterogeneous arrays. In the phased-array instruments that we consider, arraying occurs at two levels that we must consider separately.

8.1. Linearly or circularly polarized receptors?

The case of Faraday rotation discussed in Sect. 5.3 offers a good example of the benefits that accrue from the matrix approach. The rotation entails a mixing of the x and y fields that cannot be described in cartesian coordinates in terms of a linearising approximation. In the quasi-scalar approach, one introduces a transformation to circular coordinates [1,3] to diagonalise the rotation matrix, making the rotation fit in the quasi-scalar framework.

This is one of the reasons why many observers hold the view that circularly polarized feeds are superior. This mistaken view might lead to misguided design decisions. In the LOFAR [10] study group, the issue is being debated at length (J.D. Bregman, private communication). The view was advanced that, even if linear feeds would be selected for production/logistics reasons, it would be necessary to convert the signals afterwards to circular polarization in order that Faraday rotation can be accounted for. This would mean either expensive additional hardware or complications in the data processing.

In the matrix framework, this entire discussion becomes a red herring. The 'problem' of Faraday rotation is nothing more that an artefact of an inadequate mathematical model. From source to image, the signals can be handled in whichever coordinate frame one prefers. Probably the most attractive frame is the one in which the nominal Jones matrices of the feeds become identity matrices.

8.2. The necessarily homogeneous 'station' tied array

In international Square-Kiolmeter-Array (*SKA*) terminology, a *station* is the equivalent of a parabolic antenna in the existing telescopes. It is a large *tied array* consisting of receptors whose outputs are added ('tied') together to form the equivalent of the *primary beam* in traditional instruments. The pointing direction of this beam is determined by appropriate delay and phase control of the receptor signals prior to combining them; likewise, the shape of the beam is controlled by appropriate weighting or 'tapering'.

In this entire process, all receptors to be combined must be made identical in order to pick up the same component of the vector field. Failing to do so results in decorrelation and depolarization; their effect can be assessed and corrected for, but they entail an irrecoverable loss in signal-to-noise ratio. For this reason, *station arrays must be homogeneous*. (If they are not, appropriate matrix transformations must be applied on the signals before adding them; this eems to me to be an undesirable complication that would serve no purpose.)

As a consequence, the phasing/weighting/combining operation can be described in scalar terms; no 'new' mathematical theory is needed. In particular, the scalar adaptive sidelobe-nulling techniques and algorithms that are being explored and developed [11] remain valid for a homogeneous tied array. Only certain details will have to be worked out, e.g. whether the two (e.g. X and Y) receptor subsets should have two separate or one joint nulling algorithm.

8.3. The naturally heterogeneous interferometer array

The interferometers are formed from pairs of stations. Naturally, each station will be in its own cd .. local horizontal plane. It follows that when they all point at the same source, each will be pointing in a different direction in its local coordinate frame. The receptor's polarization responses to the same observed source will therefore be different, — the more so the farther the stations are apart.

In other words, an interferometer array built up from tied-array elements will naturally be heterogeneous. Clearly such heterogeneity could only be avoided at high cost, but there is no reason for trying to. Matrix self-alignment will handle it gracefully. In fact, it offers us the advantage of an inherent phase alignment. To take full advantage of this, we probably do well to deliberately make the array as heterogeneous as we can, e.g. by 'randomising' the stations' receptor orientations and by using both linearly and circularly (or veven elliptically) polarized receptor designs.

References

1. Hamaker J.P., Bregman J.D., Sault R.J. "Understanding Radio Polarimetry, I: Mathematical Foundations". *Astron. and Astrophys. Suppl.* 117, 137-1467, 1996.

2. Sault R.J., Hamaker J.P., Bregman J.D. "Understanding Radio Polarimetry, II: "Instrumental Calibration nof an Interferometer Array". *Astron. and Astrophys. Suppl.* 117, 149-159, 1996.

3. Hamaker J.P. "Understanding Radio Polarimetry, III: Interpreting the IAU/IEEE definitions of the Stokes parameters". emphAstron. and Astrophys. Suppl. 117, 161-165, 1996.

4. Hamaker, J.P. "Understanding Radio Polarimetry, IV: "The full-coherency analogue of scalar self-calibration: Self-alignment, dynamic range and polarimetric fidelity". *Astron. and Astrophys. Suppl.*, 2000, in press.

5. Born M., Wolf E. "Principles of Optics". Pergamon Press 1964.

6. Thompson A.R., Moran J.M., Swenson G.W. Jr. "Interferometry and Synthesis in Radio Astronomy". John Wiley & Sons, New York 1986.

7. Perley R.A., Schwab F.R., Bridle A.H. (Editors). "Synthesis Imaging in Radio Astronomy, a Collection of Lectures from the Third NRAO Synthesis Summer School". A.S.P. Conference Series Vol. 6, 1994.

8. Weiler, K.W. "The Synthesis Radio Telescope at Westerbork: Methods of Polarization Measurement". *Astron. and Astrophys.* 26, 403-407, 1973.

9. Leppänen K.J., Zensus J.A., Diamond P.J. "Linear Polarization Imaging with Very Long Baseline Interferometry at High Frequencies". *Astrophys. J.* 110, 2479-2491, 1995.

10. Kassim N.E., Lazio J., Weiler K.W., Gross C. "Low Frequency Array (LOFAR). *This Volume.*

11. Leshem A., van der Veen A. "Effect of Interference Mitigation on Radio-Astronomical Imaging". *This Volume.*