


Author: J.E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

LOFAR Calibration framework

J.E. Noordam
(*jnoordam@astron.nl*)

Dwingeloo
15 Oct 2006

Reviewed:			
Name	Signature	Date	Rev.nr.
LOFAR Calibration Stuurgroep		15 Oct 2006	1.0

Accepted:		
Work Package Manager	System Engineering Manager	Program Manager
J.E. Noordam	K. van der Schaaf	J. Reitsma
..... <i>date:</i> <i>date:</i> <i>date:</i>

©ASTRON 2006
All rights are reserved. Reproduction in whole or in part is prohibited without the written consent of the copyright owner.

Distribution list:

Group:	For Information:
LOFAR Calibration Review Committee	O.M. Smirnov
LOFAR Calibration Stuurgroep:	S.B. Yatawatta
H. Falcke	M. Mevius
K. van der Schaaf	G. van Diepen
W.N. Brouw	J. van Swieten
A.G. de Bruyn	M.A. Brentjens
T.A. Oosterloo	S.J. Wijnholds
R.J. Nijboer	J. D. Bregman
M.P. van Haarlem	J. P. Hamaker

Document revision:


Revision	Date	Section	Page(s)	Modification

Abstract


This document contains a step-by-step outline of the LOFAR calibration framework. The *basic procedure* can be adapted for the various observing modes of the Key Science Groups and other applications. The main difference between LOFAR and existing instruments is the prevalence of large *image-plane effects*, i.e. instrumental effects that vary over the field of view. This requires a generalisation of traditional self-calibration, which has been developed. Unfortunately, the extra sophistication requires a rather large increase in processing. On the positive side, two rather fundamental *calibratability conditions* are satisfied for LOFAR. First of all, there are enough bright calibrator sources, and thus sufficient information available to calibrate LOFAR. Secondly, there are enough equations to solve for the parameters of the LOFAR Measurement Equation(s). Finally, it is very important to realise that, although we have made a promising start, the full development of LOFAR calibration will take time, and will only happen if we create the right conditions for it.

Contents

1	Introduction	4
2	The LOFAR Measurement Equation (M.E.)	6
3	The LOFAR calibration strategy	9
4	Step-by-step calibration procedure	10
4.1	Station calibration	10

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

4.2	Preliminaries	12
4.3	Creating suitable conditions for the Major Cycle: The Rough	12
4.4	The Major Cycle: The Smooth	13
4.5	Delivering the Goods	14
4.6	Further processing by the user: The Empty	14
5	Estimating M.E. parameters	15
5.1	Estimating primary instrumental parameters: peeling	15
5.1.1	The effects of peeling contamination	15
5.1.2	The A-team	16
5.2	Estimating Cat I source parameters themselves	16
5.3	Estimating secondary parameters from primary ones	16
6	Ionosphere calibration	17
6.1	Estimating MIM parameters (IJones)	17
6.2	Phase locking: finding the ambiguity numbers of the MIM	18
6.3	Estimating ionospheric Faraday rotation (FJones)	19
7	Equations and unknowns	19
7.1	Solving for instrumental parameters	20
7.2	Solving for source parameters	21
8	Cat II subtraction	21
9	Dynamic range	22
10	Engineering requirements	22
11	Conclusions	23
A	Appendices	26
B	LOFAR vital statistics	26
C	Some operational choices	26
D	Special techniques	32
E	Some words on efficiency	32
E.1	Cutting corners (suggested optimisations)	33
E.2	Look for alternative methods	33
F	The role of simulation	33

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

1 Introduction

LOFAR calibration will be a challenge, because of a pathological ionosphere, crowded fields, very bright sources, extended sources, unstable station beams, high station beam sidelobes (all-sky imaging), and unsmooth bandpasses. It will be more difficult than calibrating existing radio aperture synthesis telescopes, partly because of less favourable conditions, and partly because the higher sensitivity requires a higher dynamic range. Therefore, new calibration methods and algorithms are needed for LOFAR (and SKA)¹.

We define calibration as the **capability to subtract 'foreground' sources**, as illustrated in fig 1. The goal is to produce residual images that are *as close to gaussian noise as possible*. This is possible only if the Measurement Equation (M.E., see section 2) is correct, and the values of its parameters are accurately known. The problem is that the LOFAR M.E. is rather complicated. The biggest difference with earlier instruments is the prevalence of (large) *image-plane effects*: ionosphere and station voltage beamshapes cause instrumental errors that vary over the field. Not only does this increase the number of instrumental parameters that have to be estimated, but it also increases the processing requirement by a large factor.

Since uv-data can only be corrected for a single point in the sky, calibrated uv-data do not exist in the presence of image-plane effects. Therefore, as many sources as possible must be subtracted (or filtered) from uv-data, i.e. not from images. For the brightest sources this is already standard practice in most reduction packages. In the case of LOFAR this principle must be extended to all sources in the Local Sky Model (LSM). Since there will be thousands of such sources per field, this will be one of the bottlenecks of LOFAR data processing. We define the following categories of sources:

	brightness	#/field	subtract from	remarks
Cat I	$SNR_{sample} > 3$	20 – 30	uv-data	are used for parameter estimation
Cat II	$\sigma_{image} > 10$	$10^3 - 10^4$	uv-data	are subtracted in groups (patches)
Cat III	$\sigma_{image} < 10$	many	image	will be convolved with PSF(\vec{l})


Note that Cat I sources are treated ('peeled', see below) individually, and are there for subtracted with maximum accuracy, using their own 'private' parameters. Cat II sources are predicted for subtraction by interpolating smooth functions for beamshapes and ionosphere. PSF(\vec{l}) is a point-spread function² that depends on source position \vec{l} . The dividing lines between the source categories are dynamic. As multiple observations of the same field are added, σ_{image} will decrease, so more Cat III sources will be turned into Cat II sources. And a bright Cat II source that 'causes trouble', can always be promoted to Cat I status, for individual treatment³. A somewhat special case is the so-called 'A-team', i.e. the small number of very bright sources (Cas A, Cyg A, Tau A, Vir A), at least two of which are visible in *every* LOFAR observation, due to the relatively high station sidelobes⁴. They are treated as Cat I sources, but they are only used for beamshape estimation if they are in the main lobe.

¹Obviously, the highly productive existing instruments (WSRT, VLA, ATCA, VLBI, etc) should also be able to profit from the new calibration software. With the existing packages (NEWSTAR, AIPS, MIRIAD, AIPS++), only the WSRT occasionally reaches the thermal noise in all polarisations, over the entire field of view. This is due to the virtual absence of closure errors, on-axis receivers, equatorial mounts, and NEWSTAR.

²We strongly encourage the use of PSF i.s.o. (synthesized) beam. The latter is confusing.

³Unproven thesis: Any source that is bright enough to cause trouble, is bright enough to be dealt with. This means that, if its residuals after being subtracted as a Cat II source are too large, it must be bright enough to be promoted to Cat I status, and dealt with individually.

⁴LOFAR imaging is all-sky imaging' (Jaap Bregman)

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	


The LOFAR calibration strategy makes the following assumptions:

1. The ionospheric phase will vary substantially over the field, and in time and frequency (see fig 1). It will not be possible to predict the ionospheric phase by external means (PIM, GPS) with sufficient accuracy, i.e. better than a degree. This means that it must be measured continuously during the observations, using bright (Cat I) sources in the field. Based on work done in Cambridge and at the VLA, we have adopted a (uniform) rate of 0.1 rad/s as the maximum that LOFAR calibration should be able to handle.
2. The instrumental polarization is determined by the projected dipole angles, and voltage beamshapes. It can only be approached by means of a proper (matrix) M.E. Its absolute calibration requires sources with substantial polarization, of which there appear to be few at LOFAR frequencies. It is complicated by ionospheric Faraday rotation. It will take time to master this aspect of calibration, but in the meantime it will not be a show-stopper.
3. All station voltage beams will be different, and will vary substantially in time and frequency. It will not be possible to control or predict their shapes with sufficient accuracy. This means that the shapes (at least of the inner parts, where sources must be subtracted) must be measured continuously during the observations, using bright (Cat I) sources in the field.
4. Because the LOFAR stations are horizontal, the voltage beamshapes will change considerably while tracking a field for a few hours (see fig 6). It is possible to derive an analytical expression for a position-dependent 'error-PSF', using the estimated station beamshapes. This can be used to deconvolve Cat III sources. Because these are faint ($S < 10\sigma$), only the inner part of the PSF is needed.
5. It is not possible to determine the shape of the far sidelobes of the station voltage beams with sufficient accuracy to contemplate any subtraction of (Cat II) sources in that area. See fig 1). Therefore all possible measures must be taken to design the instrument in such a way that the effects of such faraway sources on uv-data are negligible. See section 10.
6. There are very few accurate sky models at LOFAR frequencies. Therefore, solving for source parameters must be an integral part of the calibration procedure. For Cat I sources, it is possible (with some limitations) to do so simultaneously with instrumental parameters. The parameters of Cat II sources must be derived from residual images.
7. Some level of source subtraction remnants is unavoidable. The signature of such remnants should be fully understood, so that they can be distinguished from astrophysical phenomena.

It will be clear that the calibratability of LOFAR (see also section 7) depends critically on the following factors:

- The availability of enough bright calibrator sources. As long as there is enough information available, the problem can be solved in principle, and will be eventually⁵. Fig 3 shows that there are indeed sufficient calibrators, i.e. 20-30 sources per field that give $SNR > 3$ per uv-sample (10 s).
- The number of Measurement Equation (M.E.) parameters should be much smaller than the number of independent data-samples, and they must be sufficiently 'distinguishable' from each other. This is investigated in section 7.

⁵Note that this condition is not met for optical interferometry.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

The centre piece of this document is the outline of the *calibration strategy (section 3)*, followed by a step-by-step description of the *basic procedure*⁴. The remaining sections provide some more detail. It is highly recommended to read the figure captions carefully, since they tend to contain some of the real 'meat' for LOFAR calibration cognoscenti. For the same group, an appendix has been added that touches on some more practical aspects. The LOFAR vital statistics relevant to calibration are summarized in appendix B.

LOFAR calibration will be implemented in the so-called Black Board system (BBS) as described in an accompanying document [2]. The existence of a working prototype (implemented as MeqTrees) will certainly help in this process. The present document only contains recommendations, and does not bind the implementation team in any way.

2 The LOFAR Measurement Equation (M.E.)

The Measurement Equation describes the relationship between the true brightness distribution, and the measured visibility data \vec{V}_{ij} . For many years, scalar M.E.'s were used, until the full-polarisation matrix form was introduced [3], and extended to include image-plane effects [4].

In the matrix M.E., the flux of a source is represented as vector of 4 Stokes parameters $\vec{I}_k(f, t) = [I, Q, U, V]$, which define it in full polarisation. The visibility measured by an interferometer between stations i and j is a vector of 4 correlations $V_{ij}(f, t) = [XX, XY, YX, YY]$, which are the results of correlating all 4 combinations of the signals from the two sets of station dipoles. The most general form is:

$$\vec{V}_{ij} = \vec{A}_{ij} + M_{ij}(J_i \otimes J_j^*) \sum_k (J_{ik} \otimes J_{jk}^*) S \vec{I}_k + \vec{N}_{ij} \quad (1)$$

Note that a visibility sample is the sum of the contributions from the various sources in the field. For extended sources, or groups of sources, these are integrals over a small area (patch) of sky. The 4×4 Stokes matrix S is constant, and depends on the chosen polarisation representation. The 2×2 instrumental *Jones matrices* are station-based. The symbol \otimes designates the Kronecker product. If we ignore the noise N_{ij} , and assume that there are no additive (A_{ij}) or multiplicative (M_{ij}) interferometer-based effects, we can write:


$$\vec{V}_{ij}(f, t) = (J_i \otimes J_j^*) \sum_k \int dl \int dm (J_{ik} \otimes J_{jk}^*) S \vec{I}_k \quad (2)$$

The contributions of extended sources must of course be integrated over the sky coordinates (for clarity, the form of equ 1 is for point sources only). The Jones matrices before the Σ in equ 2 represent *uv-plane effects*, which are valid for the entire FOV, but may vary from beam to beam:

$$J_i(f, t) = B_i G_i [T_i] \quad (3)$$

while the ones after the Σ represent *image-plane effects*, which depend of source position (\vec{l}):

$$J_{ik}(f, t, l, m) = E_{ik} P_{ik} I_{ik} F_{ik} K_{ik} \quad (4)$$

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

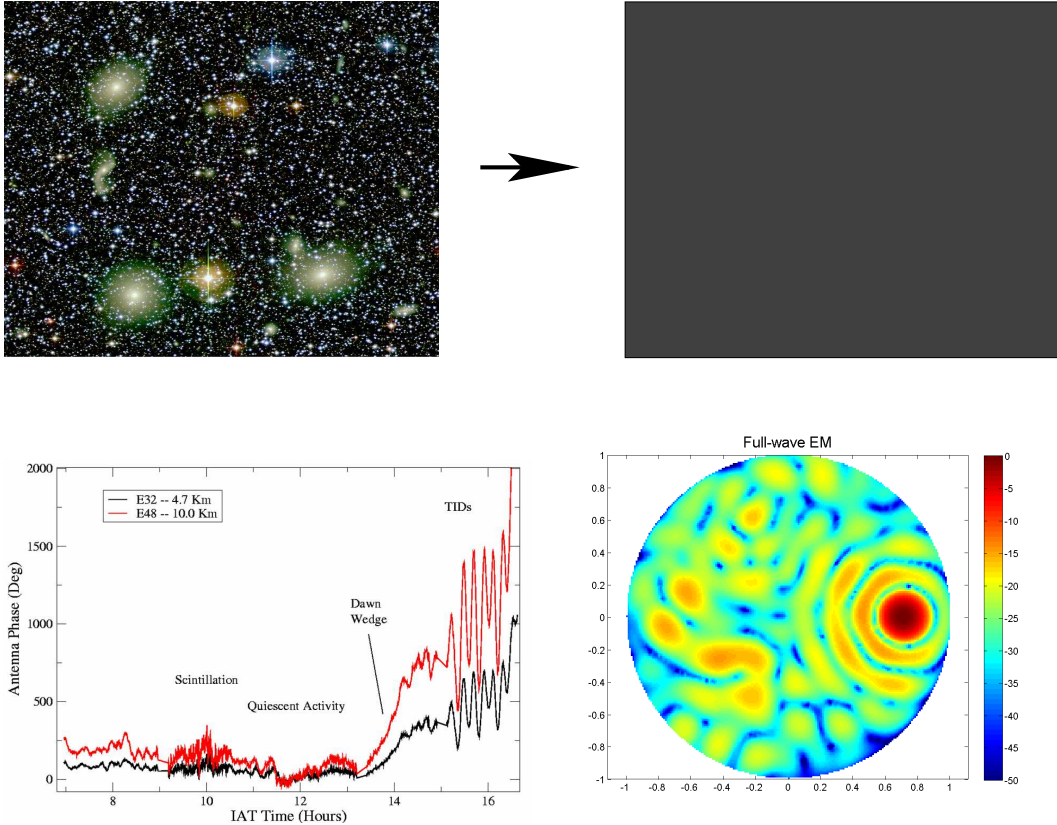



Figure 1: *The LOFAR calibration challenge. Our definition of calibration is the capability to subtract foreground sources. This is only possible if the Measurement Equation is correct, and its parameters are accurately known. The (optical) input image in the top-left panel illustrates some of the salient features of a LOFAR field. It has 10-20 bright Category I sources, which are used for estimating instrumental parameters in their direction, and to subtract the sources themselves with maximum accuracy. In addition, the instrumental parameters are interpolated to subtract the thousands of fainter Category II sources from the uv-data. The output residual image in the top-right panel is grey to illustrate that, ideally, it should only contain gaussian noise. In practice, residual images will also contain many faint Category III sources, which will be convolved with a position-dependent error PSF. They will also contain some remnants of incompletely subtracted Cat I/II sources.*

LOFAR calibration will be complicated by a number of factors. First of all, we have to observe through a 'pathological' ionosphere, especially at frequencies smaller than 100 MHz. This is illustrated by the phases in the bottom-left panel, which were measured by the VLA at 74 MHz, by Perley and Bust. Note that the maximum phase rate for the 10 km baseline is about 0.5 deg/s, well within our adopted limit of 0.1 rad/s.

In addition, the LOFAR station beams (bottom-right) will be less stable than those of the traditional parabolic dishes, and sparse stations will have higher sidelobes.


Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

Each of these 2×2 Jones matrices represents a specific instrumental effect. Each matrix element is a mathematical expression, which usually has parameters. Note that Jones matrices do not always commute with each other, so their order is important. This is discussed in more detail in section 3.

IF stands for Intermediate Frequency, and traditionally indicates the signal channel from a single dipole. For LOFAR it is the sum of 96 dipoles of a station.

- **$\mathbf{BJones}(\mathbf{f}, \mathbf{t}, \mathbf{beam})$** : Diagonal matrix. The raw IF bandpass will be 'granular', i.e. it will vary rapidly with frequency. This is caused by the sub-band filters, and by the way the sub-bands are joined. It is assumed to be known *a priori* through station calibration, and divided out from the uv-data in an early stage. Any further variations are assumed to be slow, in time and frequency, and are absorbed into EJones.
- **$\mathbf{IJones}(\mathbf{f}, \mathbf{t}, \vec{l}, \vec{x})$** : Scalar. The ionospheric phase is modelled in the form of a Minimum Ionospheric Model (MIM, explained in the caption of fig 5), which is a *large-sky* model with a minimum number of parameters. The frequency-dependence ($\phi \propto \lambda$) is assumed. The MIM coefficients are derived from the selfcal phase solutions in the direction of one or more bright (Cat I) calibrator sources.
 - NB: If significant ionospheric phase structure exists at scales smaller than a few km, there will also be amplitude effect. This manifests itself as amplitude scintillation. The latter becomes important when diffraction dominates refraction. This occurs when $r_{diff} < r_{Fresnel} = \sqrt{\lambda D}$, where r_{diff} is the linear scale over which the phase changes by a radian. It should be possible to derive the magnitude of amplitude scintillations from the rapidity of the phase variations. For the moment, it will be assumed that we will not observe under such conditions. However, if necessary, the MIM can be extended to include (smooth) amplitude effects as well.
 - The decorrelation due to a uniform change of 1 radian over an integration interval of 10 s (the LOFAR design spec) would cause a 4% decrease in amplitude. Obviously, such a large gain effect must be taken into account. As long as the variations are smooth, this can easily be done during the prediction of visibilities, in the same way as time/freq smearing.
- **$\mathbf{GJones}(\mathbf{f}, \mathbf{t}, \mathbf{beam})$** : Diagonal matrix. The IF complex gain is a uv-plane effect, which includes IF electronics, the more rapid atmospheric phase variations, crosstalk, etc.
 - [**$\mathbf{TJones}(\mathbf{f}, \mathbf{t}, \mathbf{beam})$**]: Scalar. Optionally, we may use an explicit model for the atmospheric phase, using the atmospheric pressure, temperature etc as external parameters. TJones is only weakly dependent on frequency.
- **$\mathbf{EJones}(\mathbf{f}, \mathbf{t}, \vec{l}, \mathbf{beam})$** : Has four non-zero (complex) elements, so it does not commute easily with other matrices. It models the main lobe, and perhaps the inner sidelobes, of the station voltage beam. See fig 6. They are modelled by smooth functions, with coefficients(f, t) that are estimated from Cat I selfcal (see section 6.1)⁶.
 - **$\mathbf{PJones}(\vec{l}, \mathbf{beam})$** : Rotation matrix. The effective projection of the dipoles on the sky. Could be combined with EJones, but that might make FJones calibration more difficult. Deterministic in principle, but it may be necessary to solve for some of its parameters initially, in a special program of observations.
- **$\mathbf{FJones}(\mathbf{f}, \mathbf{t}, \vec{l}, \vec{x})$** : Rotation matrix. The ionospheric Faraday rotation is related to IJones, but depends on the angle between source direction and the local Earth magnetic field. It is treated separately (see sections 3 and 6.3). Like IJones, FJones is modelled by a single *large-sky* (MIM-like) model.

⁶In the WSRT, EJones is a diagonal matrix, whose two elements represent the two separate voltage beams associated with the X and Y dipoles of an antenna. For LOFAR, with its 4 non-zero matrix elements, we have to think in terms of a single voltage beam per station.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

- **KJones($\mathbf{f}, \mathbf{t}, \vec{l}, \vec{x}$)**: Scalar. The Fourier Transform phase 'kernel' depends on station position (\vec{x}) and relative source direction (\vec{l}). It is deterministic in general, but could be used to solve for station positions.

Traditional selfcal only deals with uv-plane effects, usually just the complex gain (GJones) per station, and per frequency-channel. Sometimes DJones, the on-axis polarisation 'leakage' is also included. Station voltage beamshapes are assumed to be all identical, and constant in time, so they can be represented by a multiplicative *power* beam in the image plane. None of these assumptions are valid for LOFAR. All this is different for LOFAR.

Note that there will be more than one version of 'the' LOFAR M.E., depending on the kind of observation, the observing conditions and the required dynamic range. Although they will all have the same general structure of equ 1, they may differ in the mathematical expressions, and thus the parametrisation, of its matrix elements.

3 The LOFAR calibration strategy


This section outlines the recommended LOFAR calibration strategy, and the reasoning behind it. In an important sense it is the core of this document. The strategy has three stages, which will be labelled the **Rough**, the **Smooth** and the **Empty**⁷, as illustrated in fig 2. The first stage creates suitable conditions for the much more processing-intensive second stage by taking care the 'wilder' instrumental effects. In particular, it tracks the large ionospheric phase variations (IJones), and removes the instrumental effects that vary rapidly with frequency (BJones) and time (GJones). The second stage is the Major Cycle (MC), which iteratively determines the slowly varying shape of the station voltage beams (EJones) and the ionospheric Faraday rotation (FJones). It also improves the Local Sky Model (LSM) by adding new sources to it, and estimating better source parameters. The third stage deals with the residuals, images and/or uv-data, from which all LSM sources have been subtracted. These are not really empty, of course, but contain noise, subtraction remnants, and faint sources. The first two stages are an integral part of LOFAR operations, while the third stage is up to the user.

The following table summarises the relevant properties of the *groups of M.E. parameters* that we have to solve for in the first two stages.

parm group	matrix	stage	freq	time	\vec{x}	\vec{l}	
BJones	diagonal	Rough	per channel	hours	*	-	IF bandpass
IJones	scalar	Rough	$\propto \lambda$	10 min	*	*	ionospheric phase
GJones	diagonal	Rough	-	per timeslot	*	-	IF complex gain, incl TJones
EJones	'rotation'	Smooth	smooth	10-100 min	*	*	voltage beam shape, incl PJones
FJones	rotation	Smooth	$\propto \lambda^2$	10 min	*	*	ionospheric Faraday rotation
LSM sources	-	Smooth	smooth	'constant'	-	*	Local Sky Model

The four columns in the centre indicate dependencies in four dimensions. Differences between groups determine whether they can (and should) be solved for separately. A 'smooth' frequency dependence indicates a smooth spectrum, e.g. a spectral index with only a few terms. Effects that do not depend on sky position \vec{l} are called 'uv-plane effects', the others are 'image-plane effects'.

⁷The resemblance to a famous film title is completely accidental. But it does make it easier to remember the scheme.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Kind of issue: Public	File: <i>lofar/</i>	

Obviously, the calibration strategy should not violate matrix commutation. In principle, instrumental effects should be corrected in their (reverse) order along the signal path, unless their Jones matrices commute. We are rather fortunate that I_{Jones} is a scalar (multiplied by a 2×2 unit matrix), which commutes with anything. Therefore, it can be moved to the other side of E_{Jones} . (Note that this is not possible for the related F_{Jones}). Diagonal matrices do commute with each other, and with scalars.

In the Rough stage, the uv-data are corrected for the uv-plane effects B_{Jones} and G_{Jones} , and perhaps for the I_{Jones} phase for the centre of the field. At the very least, this makes it easier to visualize the data. But the main goal is to create optimal conditions for the iterative Smooth stage. First of all, this means a minimum number of (E_{Jones}) parameters to be solved for. It is particularly important that the E_{Jones} frequency dependence should be smooth. Secondly, the I_{Jones} phase is used to shift the data to the *apparent* positions of Cat I sources, and the centre of patches of Cat II sources. This smoothes the visibility function of the source/patch of interest, thereby limiting the amount of processing.

The question is of course whether we are allowed to do this. Note that any errors made in the determination of I_{Jones} and G_{Jones} have to be absorbed in E_{Jones} . Our contention is that this is no problem since I_{Jones} is smooth at the scale of a station, and E_{Jones} has more degrees of freedom at that scale. Thus, I_{Jones} errors are simply absorbed as semi-linear phase gradients over the beams. (NB: As explained in section 6.1, the main source of I_{Jones} errors is contamination by E_{Jones} differences anyhow, so it might be argued that the effect is absorbed by the correct Jones matrix in a circuitous way). Any errors in G_{Jones} translate into a multiplicative factor of the entire station beam.


Finally, we cannot over-emphasize the importance of using a large-sky model like the MIM for the ionosphere. Thanks to the strong constraints it imposes, the large ionospheric effects can be separated from other instrumental effects, thereby almost reducing LOFAR to a 'normal' telescope.

4 Step-by-step calibration procedure

The flow diagram for the basic LOFAR calibration procedure is shown in fig 2. It is expected that different LOFAR observation modes will use parts or all of this procedure, in some form.

4.1 Station calibration

1. **Station beam-forming:** Fiddling the beam-former coefficients (open-loop!) to achieve one or more of the following effects:
 - (a) Reducing the side-lobe level.
 - (b) Approximating a constant shape of the main lobe.
 - (c) Adaptive RFI nulling (not recommended!).
2. Beamshape estimation: A first-order approximation (10%), to reduce the number of iterations downstream. Not vital.
3. Bandpass estimation: Needed to divide out B_{Jones} .

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

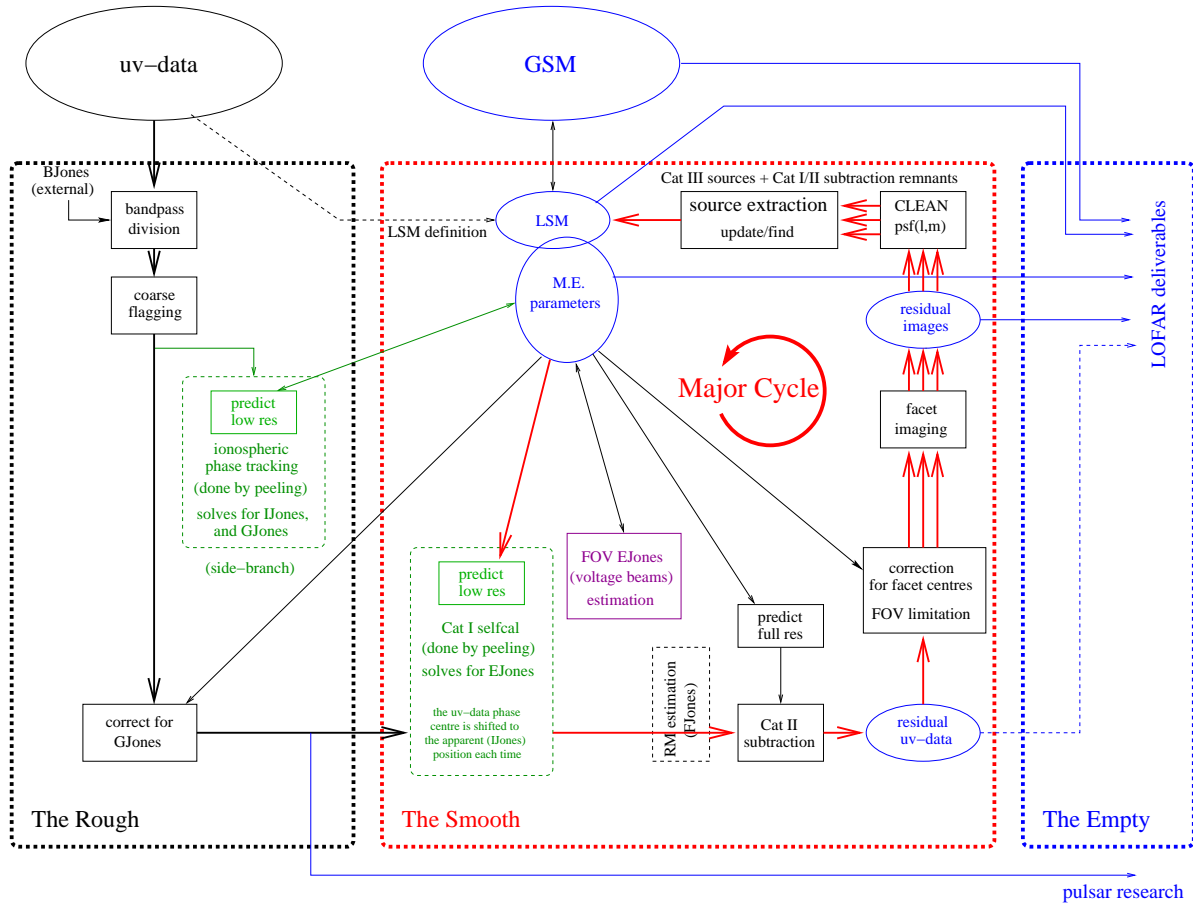



Figure 2: *The LOFAR Calibration Flow Diagram is divided into three stages: The Rough, the Smooth and the Empty. The first 'tames' the large ionospheric phase variations, and removes 'rough' uv-plane effects. The second stage is also called the Major Cycle⁹ (MC). It is iterative, and deals with smooth phenomena like voltage beams. In addition, increasingly fainter LSM sources are found or improved in each iteration, and subtracted in the uv-plane. The third stage deals with the 'empty' residual images or uv-data, and is up to the user.*

Note that the M.E. parameters overlap with the Local Sky Model, indicating that the LSM source parameters are regular M.E. parameters, just like the instrumental ones. Two-way arrows mean that parameters are not only used for prediction, but are also solved for.

The uv-data are corrected for all uv-plane effects (BJones, GJones) before entering the MC. The ionospheric phase (IJones) is used to shift the phase centre of the data to the apparent position of a Cat I source or a patch. The Cat I selfcal solves for the elements of the EJones (voltage beam) matrices that are associated with each bright calibrator source. These are then used to solve for the coefficients of station EJones matrices.

The bright Cat I sources are subtracted with their 'private' parameters, for maximum accuracy. Cat II sources are subtracted in groups, using interpolated values of the station EJones, or the ionospheric Faraday rotation.

The residual uv-data are transformed into residual images, and deconvolved with a position-dependent error PSF. The resulting CLEAN components are used to update existing LSM sources, and finding new ones. The MC is repeated until the residual images have a specified dynamic range.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

4.2 Preliminaries


For each observation, a suitable processing 'tree' must be generated.

1. **Generate a Local Sky Model (LSM):** Use the instrumental 'windows' of the observation (primary beam, spectral window) to select relevant sources from the Global Sky Model (GSM). The LSM includes all the known sources in the main lobe(s) of the primary beam, and a small number of very bright sources (the A-team) all over the sky.
2. **Generate a suitable 'tree'** to process this particular observation, with the correct number (and order!) of Cat I selfcal stages, followed by Cat II patch subtraction stages. The LSM plays an important role in this. Upon request, it returns a list of sources in reverse order of *apparent* (...) brightness, which can be used to select Cat I peeling sources. It will also tile the FOV with 'prediction patches' of a suitable size, to be used in Cat II subtraction.
3. **Transfer** any externally measured parameter values, to be used as initial parameter values. This reduces the number of iterations.

4.3 Creating suitable conditions for the Major Cycle: The Rough

The uv-data is processed in 'snippets', i.e. chunks up to several minutes, containing up to 4000 frequency channels. See also fig 4. The following operations are done to the snippets only once, before feeding them into the Major Cycle.

1. **Bandpass division (BJones):** Divide out the high-granularity (channel-by-channel) IF bandpasses caused by the electronics (filters etc). They may be determined at station level by injecting a suitable test-signal.
2. **Coarse flagging:** Only the RFI that can be clearly distinguished from the signal by relatively cheap detection algorithms. NB: FLAGS AFFECT THE WEIGHTS OF THE UV-DATA SAMPLES, AND THEIR UV-COORDINATES....
3. **Ionospheric phase (IJones) tracking:** Use one or more bright calibrators to update the parameters of the Minimum Ionospheric Model (MIM). See section 6.
 - (a) As a by-product, the GJones are measured as well. They are distinguished from IJones by the fact that they are uv-plane effects (i.e. independent of direction (\vec{l})), and by their frequency dependence.
 - (b) At the start of an observation, several minutes may be needed *to acquire ionospheric phase-lock*, i.e. to carry out a program of trial and error to eliminate 2π ambiguities. See section 6.
 - (c) The simplest approach is to use calibrator sources in the FOV. However, this will not always be sufficient to constrain the MIM over a large enough part of the sky. *Therefore, it should definitely also be possible to switch beams rapidly in order to point directly at calibrator sources.* Obviously, this will complicate the observation and calibration schemes somewhat, but once it exists it will open the way to a whole family of new observational modes.
4. **Correct the data for uv-plane effects**, i.e. GJones. Note that the data are NOT corrected for IJones, although one might consider correcting them for the MIM phase of the field centre, e.g. for visualisation or monitoring. The MIM that is used downstream must be adjusted for this, of course.


Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

- The uv-data snippets are now ready for one or more passes through the Major Cycle. They are corrected for uv-plane effects, and no sources have been subtracted (IJones determination is done in a side-branch).

4.4 The Major Cycle: The Smooth

The uv-data snippets are passed through one or more loops of the Major Cycle. Each loop results in better values for the M.E. parameters (which include both instrumental parameters *and* LSM source parameters!), and more LSM sources.

- [Optional: Regenerate the forest]**. It may happen that a Cat II source is dynamically 'promoted' to be a Cat I source, because it is bright enough to 'cause trouble'.
- Cat I selfcal (peeling, see section 5.1)**: Solve for voltage beam (EJones) parameters in the direction of bright (Cat I) sources. This is done in order of decreasing apparent brightness. The phase-centre of the uv-data is moved to the *apparent* position of each source, i.e. corrected for ionospheric phase errors (IJones). The contribution of each Cat I source is subtracted (peeled off) from the data before moving to the next one.
 - Integration time issues...
 - Contamination issues...
 - Some of the Cat I sources will be the very bright sources of the so-called A-team (Cas A, Cyg A, Tau A, Sun etc). Because of the high side-lobes of the LOFAR station beams, they will always be visible, so their effects must always be subtracted. Any special treatment due to their large distance from the phase centre is baked into the trees.
- Estimation of station voltage beams (EJones)**: The parameters of the 'private' EJones matrices associated with individual Cat I sources are used to estimate the the elements of the station EJones matrices.
 - A potential complication is that the size and shape of the beams will vary considerably with elevation. This has to be taken into account when selecting the Cat I sources to be used.
 - It may be desirable to use longer integration on a number of fainter calibrator sources on the field.
- [Optional: Estimation of ionospheric Faraday Rotation (FJones)]**: See section 6.3.
- Progressive flagging**: As more bright sources are peeled off, we may flag the residuals for increasingly subtle RFI, while still using relatively cheap detection algorithms.
- Subtraction of Cat II sources**: Done in groups (patches). See section 8. This could well be a major bottleneck, so we should look for alternatives.
- [Optional: Estimation of interferometer-based errors]**: Multiplicative and/or additive. They violate the so-called 'selfcal-condition', which requires that all instrumental effects should be station-based. Only then is the number of independent parameters much smaller than the number of data. Therefore, in a well-designed system they should be negligible. But since even the WSRT has them (very small ones) we should assume that LOFAR will also suffer from this affliction to a certain degree. In practice, we can only do something about them if we have *a priori* information, e.g. that they are constant over the entire observation.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

8. **Facet imaging.** The size of a facet is determined by ionospheric conditions, and perhaps by the w-coordinate. The residual ionospheric phase errors over a facet should be less than one radian. For each facet, the residual uv-data are corrected for the facet centre, after which the phase centre is shifted to that position, and the FOV is reduced by integrating over freq and time. The results are gridded and Fourier Transformed to a set of 4D residual images (which may overlap a little). A position-dependent error-PSF is generated as well.
9. **Cat III deconvolution.** Apart from gaussian noise, the 4D residual images contain Cat III sources and the remnants of incompletely subtracted Cat I/II sources. Deconvolution is somewhat complicated because the error PSF depends on the position (\vec{l}) of a source. See section ??.
10. **Source extraction** from images of CLEAN components is done in two modes:
 - (a) Updating the parameters of existing LSM sources. The LSM is used to inspect the areas where Cat I/II sources have been subtracted to see whether there is anything left. Any residuals are used to solve for improvements of the source parameters, using the source trees in the LSM.
 - (b) Finding new LSM sources. The most difficult problem is to decide which groups of CLEAN components represent a new source, and how this source should be parametrized (if at all). Extended sources may be modeled in terms of base functions like shapelets (see fig 8) or pixons. Really pathological sources may be stored in the LSM as 4D images.

The updated LSM will now be used to predict/subtract all sources from the uv-data. The Major Cycle is repeated until some calibration quality criterion is met. It is expected that at least two cycles will be needed to reach a dynamic range of $1 : 10^4$, as required in the survey mode (see section 9).


4.5 Delivering the Goods

After exiting the Major Cycle, the results have to be disposed of:

1. Update the Global Sky Model (GSM) from the LSM.
2. Make LSM images, i.e. images from the LSM sources, to be used by Surveys KSG.
3. Deliver the deliverables:
 - (a) residual images
 - (b) [LSM] or images
 - (c) [GSM]
 - (d) [MeqParm tables]
 - (e) [metadata]

4.6 Further processing by the user: The Empty

For some applications, the deliverables will be sufficient. For others, it will be necessary to do more processing to extract the desired astrophysical information from the 'empty' residual images (or residual uv-data).

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

5 Estimating M.E. parameters

Instrumental M.E. parameters are estimated with the help of a smallish number of Cat I sources, i.e. bright calibrator sources. We distinguish 'primary' instrumental parameters that are associated with individual Cat I sources. They are estimated by peeling. Usually, the Cat I sources are assumed to be known, but since their parameters are regular M.E. parameters, it is possible to solve for them also. There are also 'secondary' parameters, which are derived from primary ones. These are the coefficients of smooth functions like the ionospheric MIM, or the shapes of station voltage beams.

5.1 Estimating primary instrumental parameters: peeling

It will be assumed here that Cat I parameters are estimated by peeling¹⁰. This means that the Cat I sources are treated one at a time, in order of decreasing *apparent* brightness. In practice, this is achieved by defining a chain of 'peeling stages', each with its own solver(s). At each stage, the phase centre of the uv-data is shifted to the apparent position of the peeling source. This is done by adding the ionospheric phase (IJones) to the nominal phase-shift factor. This makes the visibility function of the peeling source as smooth as possible, allowing a large reduction in the number of snippet domain cells. The contribution of each peeling source is subtracted from the uv-data before moving to the next one.

Of course it is possible to solve for the parameters of all kinds of Jones matrices, in all kinds of Measurement Equations. However, in the calibration scheme adopted here, only the station voltage beams (EJones) are estimated. After all, the uv-data have already been corrected for the uv-plane effects BJones and GJones, while IJones is used to shift them to the apparent position of the peeling source. Since the latter is in the phase centre, no KJones is needed. Thus the M.E. used for predicting the visibility function $M_{ijk}(f, t)$ of the peeling source k reduces to:

$$\vec{M}_{ijk}(f, t) = \int dl \int dm ((P_{ik} E_{ik}) \otimes (E_{jk}^* P_{jk}^*)) S \vec{I}_k \quad (5)$$


where PJones is a deterministic projection matrix. It should be emphasized that, since the voltage beams are much less well-behaved than those of the WSRT, all 4 EJones elements will be non-zero. Therefore, we should solve for their real and imaginary parts, i.e. 8 independent real parameters per station. If we decide to solve for some parameters of the peeling source itself, this number may increase by 1-10. For LOFAR, this is well within the maximum of N real parameters per station (see section 7).

Since we are not refining the uv-plane effects or IJones in the Major Cycle, any imperfections in the parameters of those Jones matrices, including time variations in BJones, are absorbed into EJones.

5.1.1 The effects of peeling contamination

When assuming that the brightest source is the only one in the sky, we are ignoring the contamination from other (fainter) sources. The latter will be present in the measured uv-data, but not in the predicted visibility. The result will be that the selfcal solution is distorted, which will lead to higher source subtraction residuals down-stream.

¹⁰The alternative to peeling is a simultaneous solution for multiple (or all) Cat I sources. This approach is considered too expensive, and will not be considered here.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

First of all, it is possible to reduce the contamination by taking some of the other sources into account in the predict (while only solving for the parameters associated with the peeling source). However, since this is expensive, we should attempt to minimize the contamination by various averaging mechanisms. It also has its limits, see section 7.

Fig 4 illustrates how a snippet (or even a domain cell) may cover multiple 'corrugations' of the visibility function of an off-axis source. We should only use interferometers whose snippets cover as many corrugations as possible, i.e. use the largest possible frequency band, and avoiding certain baseline orientations. Obviously, this has its limitations, especially for sources close to the phase centre (where the peeling source is). However, there are other mechanisms that work in or favour: The selfcal error will average out over longer periods, so if we may assume that instrumental errors vary slowly, we may smooth the solution in time and thus reduce the effects of contamination. This is equivalent to enlarging the snippets in the time direction.

We may also observe that the *vector sum* of the visibility contributions of a large number of evenly distributed sources will tend to zero. This limits the number of sources that we have to worry about the the 1-10 brightest ones, and then only the ones that are relatively close to the peeling source.

Summarizing: There are many things we can do about peeling contamination. The matter has been analyzed in more detail in [6].

5.1.2 The A-team

There is a small number of very bright sources that will be visible in any LOFAR observation. Therefore, they always have to be included in the list of Cat I sources. However, unless they are in the field-of-view, they are not used for the estimation of beamshapes or MIM. Only their apparent flux has to be determined, so that they can be predicted and subtracted. The latter does not have to be done very accurately.


5.2 Estimating Cat I source parameters themselves

During the first year(s) of LOFAR operation, the models of the Cat I sources will not be known with sufficient accuracy. A dedicated campaign to redress this will have its own problems, and even then many observations will have some Cat I sources whose parameters need to be solved for to get the best results. Give numbers of Cat I sources

Since source parameters are regular M.E. parameters, they can be solved for, in any combination with other M.E. parameters. The trick is to make sure that this process converges to ever better models for LSM/GSM sources.

5.3 Estimating secondary parameters from primary ones

The ionospheric MIM (IJones), the IF complex gains (GJones) and the station voltage beamshapes (EJones) are derived from phases and gains measured in the direction of individual peeling sources. In this section, we will concentrate on EJones (station voltage beam) estimation. The case of the ionosphere is dealt with in the next section.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

The primary parameters are the 8 real parameters in the 4 complex EJones matrix elements $e_{ik}^{ab}(f, t)$ per station i , per Cat I source k , estimated in section 5.1. From these, we estimate the secondary parameters, i.e. the coefficients of the smooth functions that describe the inner parts of the station voltage beams. For the moment we will assume that each of the 4 matrix elements of the overall station EJones matrix has its own function $e_i^{ab}(f, t, l, m)$, with 10-20 real parameters $p_m(f, t)$. We equate these functions with their measured values in the direction of the Cat I sources:

$$e_{ik}^{ab}(f, t) = e_i^{ab}(f, t, l_k, m_k) = e_i^{ab}(l_k, m_k, p_0, p_1, p_2, \dots, p_m) \quad (6)$$

and solve for the p_m (p_m^{ab} really, because they are independent per matrix element ab and per station i). Obviously, the number of parameters p_m must be less than the number of Cat I sources, and the latter should be evenly distributed over the field-of-view. Fewer Cat I sources will be needed if the 4 functions per station share some of the parameters, which will almost certainly be the case.

In this stage, we will not endeavour to write down suitable expressions for the functions e_i^{ab} . The radial part might be something like a sinc. The lateral part will be a little more tricky, especially if the inner sidelobe(s) are to be included. The way to proceed is to use simulation and measurements to make the best possible model, and to use physical considerations to impose constraints on the range of values that its parameters can take. The latter will require the inclusion of a temporary rest-function, which will gradually be reduced to zero as our understanding of the instrument deepens.

6 Ionosphere calibration


Two aspects of the ionosphere are modelled in two separate sets of Jones matrices: IJones deals with the ionospheric phase, and FJones deals with ionospheric Faraday rotation. Since the ionosphere plays such a large role in LOFAR calibration, it merits its own section.

6.1 Estimating MIM parameters (IJones)

The Minimum Ionospheric Model (MIM) is described briefly in the caption of fig 5, and more elaborately in [10]. The MIM parameters are secondary M.E. parameters, which are estimated every 10 s from the individual selfcal phase solution(s) of one or more bright calibrator sources. As fig 5 shows, more than one calibrator is needed in most cases, so we will assume that. Thus, we have equations of the form:

$$m_{ijk} = \psi_{ijk} + (\Delta\phi_{ik} - \Delta\phi_{jk}) - 2\pi(a_{ik} - a_{jk}) + (\Delta g_{ik} - \Delta g_{jk}) + \Delta e_{ij} + c_{ijk} + n_{ijk} \quad (7)$$

where m_{ijk} is the estimated phase for interferometer ij in the direction of calibrator k , $\psi(u_{ij}, v_{ij}, l_k, m_k)$ is the source model phase, $\Delta\phi_{ik}$ is the incremental MIM phase for station i in the direction of calibrator k , Δg_{ik} is the incremental GJones phase, Δe_{ij} is the difference between voltage beam phases, a_{ik} is an *ambiguity number*, c_{ijk} is contamination caused by other sources in the field, and n_{ijk} is (non-gaussian) noise. Ignoring noise, contamination and voltage beam effects (see below), and transferring the 'known' terms to the left-hand side, the right-hand side contains only quantities that we wish to solve for:

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

$$m_{ijk} - \psi_{ijk} + 2\pi(a_{ik} - a_{jk}) = (\Delta\phi_{ik} - \Delta\phi_{jk}) + (\Delta g_{ik} - \Delta g_{jk}) \quad (8)$$

After replacing the $\Delta\phi$ with the MIM expression, we solve simultaneously for the incremental MIM parameters Δp and the incremental GJones parameters Δg . As usual, we solve for *incremental* values of the various M.E. parameters, because we use the best available values as starting point for the solving process. Thus, the input data have been corrected for the extrapolated values of IJones and GJones.

It will be necessary to correct the uv-data for the 'random' part of the electronic (GJones) phase before a solution for the much smoother ionosphere can be contemplated. This cannot be estimated separately, because it would be dominated by the ionospheric phase. Therefore, we propose to do it all simultaneously. The GJones will absorb those phases that do not vary smoothly with station position, are not proportional to λ , and do not change with source direction. It is possible that some of the phases will be attributed to the wrong Jones matrix this way, but the question is whether this makes any difference downstream.

Another question is whether it is allowed to ignore terms in equ 7. The noise is probably OK because the calibrators are very bright. For the contamination c_{ijk} , we should take the same precautions as in Cat I selfcal above. The term Δe_{ij} is the difference between the two voltage beam phases. If the beams were virtually identical, as in the WSRT, the phases in the direction of the same source would be identical also. However, this will almost certainly not be the case for LOFAR beams. Fortunately, any error that this causes in the MIM parameters will be absorbed as a smooth phase function in the voltage beam EJones estimation in the Major Cycle (making this issue nicely circular).


6.2 Phase locking: finding the ambiguity numbers of the MIM

Because of the large-sky nature of the MIM, we will only get a consistent solution if we have a 'suitable' set of values for the *ambiguity numbers* a_{ik} . This is done by means of a trial and error algorithm. The proposed procedure is to start with two inner stations, and then to include the others one by one, steadily increasing the distance from the centre. After that, the values of a_{ik} are modified by one, looking for the minimum χ^2 of the MIM solution.

There are many sets of ambiguity numbers that will lead to a consistent solution. We may go one step further by imposing the condition that the ionospheric phase $\phi \propto \lambda$ and $\phi_{f=0} = 0$. The MIM now predicts the *absolute* ionospheric phase (except for errors made in the estimation process).

The process of obtaining a set of suitable ambiguity numbers (acquiring *ionospheric phase-lock*) from scratch will take at least a few minutes at the start of a new observation. After that, the array must be kept in phase-lock by using these numbers each time. But even then, it will probably be necessary to adjust them whenever the χ^2 of the MIM solution suddenly increases.

Finally, the value of χ^2 can also be used to check whether the MIM has the optimal numbers of terms for the prevailing ionospheric conditions. Whenever it increases, an extra term may be added (dynamically!), until it drops below an acceptable level. In the same way, the number of terms may be tentatively decreased when the conditions improve.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

6.3 Estimating ionospheric Faraday rotation (FJones)

Ionospheric Faraday rotation is related to the MIM phase, since it is also proportional to the integrated TEC. However, it is not easily derived from the phase, since it requires the local Earth magnetic field. In addition, because it is a differential effect, measuring the Faraday rotation requires much less accuracy than the IJones phase. Therefore, it is not part of the MIM, but estimated separately. However, the result is also an large-sky function $F(\vec{x}, \vec{l}, f, t)$, which gives the elements of an FJones matrix for any direction (\vec{l}) in the sky, as seen from any station-position (\vec{x}). There are several ways to approach this, and the strategy is to pursue them in the following order:

- Use GPS (or Galileo or GLONASS) measurements. Unlike the phase, this is accurate enough, especially for the higher LOFAR frequencies.
- Use the extended foreground polarization discovered with the WSRT. However, this may not be polarized at < 150 MHz, and certainly not < 75 MHz. In addition, it is not yet clear how exactly to use this information.
- Estimate FJones in the direction of sources with known linear polarization. Unfortunately, there appear to be not too many of those, and their polarized flux tends to be rather weak, usually < 30 mJy. In addition, we need long baselines to minimise beam depolarisation.

7 Equations and unknowns


Given an array of N stations, we have $4N(N - 1)/2$ complex uv-samples per timeslot per freq channel, which allows us to form $4N(N - 1)$ real selfcal equations. For an observation with n_t timeslots and n_f freq channels, the number of *independent* equations is:

$$neq = 4N(N - 1) \times (1 + \alpha(n_t - 1)) \times (1 + \beta(n_f - 1)) \quad (9)$$

This reduces to the familiar $neq = 4N(N - 1)$ for $n_t = n_f = 1$, and neq always increases (albeit slowly) for more data samples. Two samples are independent if they are separated in the uv-plane by more than the station diameter D , i.e. half of the station autocorrelation 'footprint'. The sampling in freq and time will usually be denser than that, especially for the short baselines (as illustrated in fig 4). This is reflected by the factors $\alpha < 1$ and $\beta < 1$. Since the number long baselines in a LOFAR configuration is relatively small, we will tentatively use $\alpha = 0.1$ and $\beta = 0.1$.

Eventually, we need a more complete analysis, which includes the channel width, the integration time and the uv-coverage. The resulting expression should of course yield a maximum $neq_{max} = 4\pi(L_{max}/D)^2$, corresponding to a fully filled, critically sampled uv-plane¹¹. For LOFAR, $D = 50m$. For a maximum baseline $L_{max} = 100km$, we have $neq_{max} \approx 6 \times 10^6$. But for $L_{max} = 3km$ (LOFAR core only), we have $neq_{max} = 5000(!)$. Obviously, these numbers are per LOFAR beam.

¹¹Note that, because the LOFAR stations are horizontal, i.e. in the same plane as the array. Thus, their footprints on the uv-plane are foreshortened with elevation in the same way. Thus, neq_{max} is independent of declination. The factor 0.5 indicates that we only sample half the uv-plane.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

The maximum number n_p of real numbers that we can solve for (in any possible manner!) is equal to the number of independent equations. But, allowing for noise, non-orthogonality and other limitations to the solving process, we will conservatively use $n_p = 0.5 \times neq$. Note that, since in our definition M.E. parameter values are represented by smooth functions like polynomials in time, n_p is the total number of *coefficients* of these functions.

7.1 Solving for instrumental parameters

For simplicity, we will assume that, on average, the frequency dependence of all M.E. parameters can be captured in 4 coefficients, e.g. a 3rd order polynomial. Since we will always have at least 4 independent samples in a LOFAR frequency band, the factor $(1 + \beta(n_f - 1))$ can be replaced with 4, and we can remove the frequency dependence entirely. Thus, the maximum number of real numbers that we can solve for in the time direction is:


$$n_p = 2N(N - 1) \times (1 + \alpha(n_t - 1)) \quad (10)$$

For the important case of a single timeslot ($n_t = 1$) we can solve for $n_p = 2N$ real parameters *per station*. Note that this implicitly assumes that all samples in a timeslot are independent. So, if instrumental parameters would vary so rapidly that a separate value would be required for each timeslot, we could solve for $2N \approx 60$ parameters per station for the LOFAR core. If we take 4 for GJones, and 2 for the MIM, and 40 for a full-polarisation station voltage beam (EJones), there is some room left for the estimation of source parameters (see below). In reality the situation is more favourable, since only the atmospheric phase (2 parameters/station) would vary that rapidly, while the others require only one parameter per 10-100 timeslots.

The number n_p is a *fundamental* limit, independent of the path that is taken. For instance, the station beamshapes are derived from gains and phases measured in the direction of individual Cat I sources. These can be estimated simultaneously, or sequentially by peeling. In the first case, the number of simultaneous parameters per solution is large, but within the limits of n_p provided we solve over the longer time interval allowed by the slow variation of beam parameters¹². This is not very practical, and that is one of the reasons for peeling. Since there we deal with one Cat I source at a time, the number of parameters per solution is much smaller, and we can choose any interval that is convenient. The disadvantage of this approach is that the peeling solutions will be influenced by the 'contamination' caused by the other sources in the field. It has been argued that this can always be reduced to arbitrarily low levels by taking increasing numbers of these contamination sources into account in the prediction process (see section 5.1). However, this would imply that the instrumental effects in the direction of these contaminating sources are known. This can only be true up to the limit imposed by n_p , whichever path we choose. *Thus, we have a consistent story, but a potential problem!* Note that the same limitation would apply to the simultaneous solution.

The solution of this problem lies in the fact that contamination is non-linear, and that the $\log N - \log S$ curve is steep enough. The latter means that the next brightest source will usually have less than half the flux of the brightest one. Therefore, only the contamination of a very small number of sources will have an appreciable effect on the selfcal solution. It is very important to realise that the brightest sources have a much greater effect on the dynamic range than the smaller ones. Therefore, we should concentrate on removing them, within the limits set by n_p .

¹²With a simultaneous solution, it would also be possible to solve directly for the beam parameters, rather than use the 2-step process via the parameters associated with Cat I sources.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

7.2 Solving for source parameters

Source parameters are also M.E. parameters, just like instrumental ones. Therefore, their determination is subject to the same limits of available information, irrespective of whether they are determined by selfcal or NEWSTAR 'updating' in the uv-plane, or by source extraction in the image plane.

Since the vast majority of LOFAR sources will be unpolarized point sources, let us assume that they have an average of 5 M.E. parameters like RA, Dec, I, [Q, U, V, RM, shape]. These will have an average of 1 coefficient in the time direction (i.e. they are largely constant in time), and an average of 2 coefficients in the frequency direction. So we have to determine an average of 10 real numbers per LSM source. Thus, in the absence of instrumental errors, we can solve for a maximum of $n_p/5 = 0.1 \times neq$ sources per beam (FOV).

For the full LOFAR ($neq = 6 \times 10^6$), this is 600.000 sources per beam, which seems ample. However, for the LOFAR core alone ($neq = 5000$), this would be only 500 sources per beam, or less if we take instrumental errors into account.

Obviously, this kind of reasoning is closely related to, and must be consistent with, traditional confusion limits.

8 Cat II subtraction


The (very) bright Cat I sources are subtracted in the peeling stages, with maximum accuracy. The (thousands of) fainter Cat II sources are subtracted in patches, i.e. groups that are in a smallish area on the sky. Only Cat II sources in the main lobe and the inner sidelobes will be subtracted.

The contribution of patches of Cat II sources to the visibility is done by means of so-called *uv-bricks* (see appendix C). These have the important property that they allow the application of image-plane effects, which are different for different interferometers. On the other hand, they cause patch 'tiling' in the residual images, because the quality of the prediction, and thus the subtraction, decreases towards the edge of the patch.

Cat II subtraction will be very expensive. However, we should realise that we are doing too much, because for many observations we are not interested in the positions and fluxes of the Cat II sources: we just want to remove them. Up to this point, there are two possible alternatives to pursue, perhaps to be used in combination:

- Subspace decomposition (see the caption of fig 9). It is not possible to target a specific kind of sources, but it is possible to use only the longer baselines, and subtract the result from all baselines. This would preserve the EoR signature, which is known to be extended.
- Time/freq differencing of uv-data. This very effectively removes sources around the phase centre, the extended ones more than the point sources. Thus, it could be used to target specific kinds of sources, either to subtract or to preserve them.

There might be other possibilities.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

9 Dynamic range

The dynamic range of LOFAR images is limited by residual instrumental errors because they cause imperfections in the subtraction of sources. Systematic errors are worse than rapidly varying ones. Therefore, in order to assess their impact on a residual image, we should look at their *correlation footprint* in the uv-plane. For instance, ionospheric phase errors are shared by all interferometers, but change relatively rapidly in time. Station beamshape errors only affect a subset of the interferometers, but vary much more slowly. RFI often looks terrible in the data, but since it varies very rapidly, it may not cause serious effects in the image.

We also have to study the impact of instrumental effects that are specific to LOFAR. Some examples are:

- The bandpass sawtooth ripple (BSR) effect, which is caused by the fact that the station beamformers use phases rather than time-delays. This causes a position-dependent gain effect that looks like a sawtooth in the frequency direction. It may cause an appreciable effect in the FOV, due to bright sources outside it.
- Two different operations will give rise to a 'tiling' structure in the residual image, with different periods. They are caused by the fact that the prediction of Cat II sources will be less accurate towards the edge of an LSM patch, and by the fact that residual instrumental errors increase towards the edge of a residual image (facet).


NB: I was planning to develop the theoretical framework in which we can study and compare the propagation of all kinds of effects into the image a little further. Unfortunately, there has not been time before the deadline of this document. But it remains an important subject, which will have to be addressed in the not too distant future.

In the meantime, we refer to the DR requirements of a typical LOFAR observation, as given by [7]. Remarkably, the number turns out to be $\approx 10^4$ for both wavelength ranges, and for the core (2 km) as well as the lon-baseline (75 km) array. For a typical observation, it should be possible to achieve this in two passes through the Major Cycle. This presupposes that the bright Cat I sources are already in the LSM, with approximately the correct parameters. This allows the Cat II sources to be found in a single source extraction operation, and subtracted from the uv-data.

10 Engineering requirements

The needs of calibration impose requirements on LOFAR engineering. The most important ones are:

- The station sensitivity must be sufficient to get $SNR_{sample} > 3$ for a sufficient number (20-30) of Cat I calibrator sources in a typical LOFAR field.
- Smoothness (f, t, \vec{l}) of instrumental effects, so that they can be modelled by the smallest number of parameters.
 - If time-discontinuities cannot be avoided, like in the pointing of the HBA racks, they should all happen together, at well-known moments.
 - If frequency roughness cannot be avoided, like in BJones, it should be divided out.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

- Station configuration: Since we cannot subtract sources in the far sidelobes (except the A-team), the stations must be designed in such a way that the response to sources in the far sidelobes averages out as much as possible over all *interferometers*. This affects selfcal as well as imaging! Moreover, since the effects of mutual coupling between antennas on the spectral response are unknown, the configuration should be chosen in such a way that mutual coupling is minimised.

Finally, instrumental errors that cannot be avoided should be as un-systematic as possible, i.e. they should have a minimum uv-plane correlation footprint. See section 9 .

11 Conclusions


As stated in the introduction, LOFAR calibratability depends crucially on two factors. We are confident about the first, the availability of enough bright calibrator sources. The second, the balance between equations and parameters, is a difficult topic. A framework for understanding it is outlined in section 7. The numbers are encouraging for the proposed Calibration Strategy (section 3). But the issue must be laid to rest by experience. The cautious conclusion is that we are optimistic about the calibratability of LOFAR.

In the last few years, the LOFAR Calibration Studies group has made very substantial progress in finding ways to deal with the problems presented by the next generation of large radio telescopes. *Generalised selfcal* is based on an arbitrary Measurement Equation, and solves for arbitrary subsets of its parameters. The latter include source parameters as well as instrumental parameters. All M.E. parameters are assumed to be smooth functions of frequency and time, which allows us to make maximum use of known continuities in those dimensions. We have a Minimum Ionosphere Model (MIM), which requires a remarkably small number of parameters. We have new ways of representing extended sources (e.g. with shapelets), and to apply image-plane effects when predicting their visibilities. We can also solve for arbitrary source parameters, either from uv-data or from residual images. We have a way of generating position-dependent error PSF's, so that we can deconvolve sources that move through the station beams during observation.

We also have (the beginnings of) some new frameworks for understanding the fundamental limits of selfcal (section 7), and the propagation of instrumental effects into the final image (section ??).

Very importantly, we have a *working prototype* (implemented as MeqTrees) of the kind of software that is needed to implement innovations listed above. This is an invaluable help in implementing the actual LOFAR processing system, and to guide its further development.

However, all this new sophistication has a price, in memory use and processing cycles. Bigger computers are only a part of the solution. New ideas are necessary, to do things in different, more efficient ways. Some of these have already been identified. The so-called *peeling* technique offers considerable savings in processing by shifting the phase-centre from source to source. Appendix E.1 lists a number of ways to minimise processing dynamically, i.e. by allowing the software to take its own decisions based on the data situation. In addition, we have identified some promising new avenues to be explored. For instance, the subtraction of Cat II sources from uv-data is expected to be a major bottleneck. Techniques like subspace decomposition (see fig 9) might offer a way to 'filter' them out without having to know their details. Similar or related methods may be used to filter out RFI, or instrumental effects with a particular signature, like the Bandpass Sawtooth Ripple effect[11]. Appendix D lists some more examples of special techniques that are developed in other packages, or other fields, and which should be part of the LOFAR calibration toolbox.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

In summary, although we have made a promising start in creating a LOFAR calibration system, this is only the beginning of a longish development process. The framework described here will be sufficient to make discoveries with the early LOFAR, but it is unreasonable to expect it to keep pace with future extensions and rising expectations. This development process has taken a long time with earlier instruments, and there is no reason why it should be different with LOFAR. Therefore, it is vital to *create the conditions* in which new calibration ideas will continue to be generated, and can be quickly implemented. Some suggestions:

- A critical mass of clever and passionate people should be involved with LOFAR operations for a long time (years), and kept motivated somehow.
- They should have a system that offers many 'windows' on what is actually going on. Visualization is a most efficient generator of ideas and understanding.
- They should have a system that allows rapid experimentation. The existing instruments have performed significantly below their real capabilities because of the difficulties in implementing new calibration ideas.

Note that it is not suggested that the LOFAR 'workhorse' system should be burdened with this much versatility. In view of the expected data-volumes, this will not be practical. The necessary experimentation should be carried out with relatively small amounts of data, using a parallel system. However, it is still highly recommended to minimize the time and effort required for generating (and debugging!) new ideas in the workhorse system, by using something like the Tree Definition Language (TDL) used by MeqTrees.


But first, we must build LOFAR in such a way that its calibratability is maximised. This means that the engineering requirements summarized in section 10 must be taken very seriously. Calibrating LOFAR will be difficult enough without avoidable complications.

Acknowledgements


I gratefully acknowledge the contributions (and pressure) of the members and guests of the LOFAR Calibratie Stuurgroep in shaping this document. In no particular order: Ger de Bruyn, Wim Brouw, Jaap Bregman, Stefan Wijnholds, Johan Hamaker, Kjeld van der Schaaf, Heino Falcke, Ronald Nijboer, Tom Oosterloo and Michiel van Haarlem. Of course, the development of LOFAR calibration has taken some years. During that time, Ger de Bruyn, Jaap Bregman, Johan Hamaker and Wim Brouw have always been available with their vast experience, to challenge me on every issue. I also enjoyed working with the Haystack/MIT group. Sanjay Bhatnagar and Tim Cornwell from NRAO have influenced our thinking about the application of image-plane effects. Finally, it is a continuing privilege and source of inspiration to be part of the MeqTree development team: Oleg Smirnov, Sarod Yatawatta, Ronald Nijboer, Maaijke Mevius, Tony Willis and, in an earlier stage, Ger van Diepen.

References

- [1] J.E. Noordam: *LOFAR Calibration Strategy* (CDR-I, LOFAR-ASTRON-ADD-014, November 2005)
- [2] K. van der Schaaf: *LOFAR Selfcal Implementation* (LOFAR-ASTRON-SDD-050, Oct 2006)
- [3] J.P. Hamaker, J.D. Bregman, R.J. Sault: *Understanding radio polarimetry I* (Astron. Astrophys. Suppl Ser 117 137-147, 1996)

Author: J.E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

- [4] J.E. Noordam: *The Measurement Equation of a generic radio telescope* (AIPS++ implementation note 185, 1996)
- [5] S van der Tol, B.D. Jeffs, A.J. van der Veen: *Self Calibration for the LOFAR Radio Astronomical Array*, (June 2006), Subm. IEEE Transaction on Signal Processing
- [6] J.E. Noordam: *Peeling the Visibility Onion, the optimal way of doing self-calibration* (LOFAR-ASTRON-MEM-212, 2004)
- [7] A.G. de Bruyn, J.E. Noordam: *LOFAR dynamic range requirements and snapshot imaging* (LOFAR-ASTRON-MEM-213, 2006)
- [8] S.J. Wijnholds: *Study on LOFAR Core Calibratability Based on Source Statistics*(LOFAR-ASTRON-MEM-205, 2006)
- [9] S.J. Wijnholds: *LOFAR station calibration: Pitfalls and Possibilities* (LOFAR-ASTRON-ADD-217, Oct 2006)
- [10] J.E. Noordam: *A Minimum Ionospheric Model (MIM)* (LOFAR-ASTRON-MEM-214, 2005)
- [11] J.E. Noordam: *The Bandpass Sawtooth Ripple (BSR) effect* (LOFAR-ASTRON-MEM-218, 2004)

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

A Appendices

The following appendices are not directly relevant for the LOFAR calibration framework. But they may answer some of the questions that will occur to the reader, about the practical aspects of such an ambitious scheme.

B LOFAR vital statistics

Here is a brief summary of the essential LOFAR numbers that are relevant for calibration. The instrument has been described more fully elsewhere.

LOFAR is a distributed sensor network covering an area of 100 km in diameter¹³, centered on Exloo in the Netherlands. The main sensor types are antennas used for radio astronomy observations¹⁴. The sensors are grouped in 77 stations, 32 of which are in a compact core with a diameter of ~3 km, and 45 are distributed along 5 spiral arms. Each station forms a phased array, which can form up to 8 independent beams. For each beam (pointing direction), the signals of all stations are combined (correlated) centrally to form an aperture synthesis telescope. The maximum bandwidth is 32 MHz, divided in sub-bands, which do not have to be contiguous. Each sub-band has 256 channels of 0.76 kHz each, giving a total number of spectral channels of 42240. It is possible to trade bandwidth against number of beams. The basic integration time is 1 s.

LOFAR has two frequency ranges, each of which has a separate set of antennas at each station.

- The LBA units are sensitive to 20-80 MHz, and consist of individual feeds of two dipoles each. The 96 units of a station are placed randomly over a circular area with a diameter of 60(?) m, with an increasing density towards the centre. Each station has a different random configuration, thus minimising the rms sidelobe pattern of an interferometer.
- The HBA units are sensitive to 115-240 MHz, and consists of sub-arrays (racks or tiles) of 4 × 4 feeds in a rectangular pattern. Each rack has its own analog beamformer, with 5-bit phase-shifters. This means that the rack beam has to be moved discontinuously every 10 min or so. It is not yet clear how the 96 HBA units in a station will be configured.

The total nr of signal paths is 7700. It is not (yet) possible to observe simultaneously in the two bands.


C Some operational choices

In addition to the calibration principles above, a number of operational choices have been made for MeqTrees, and are highly recommended for BBS. They are one level above implementation choices.

1. A Measurement Equation (M.E.) is represented as a 'forest' of parallel *trees* (graphs, really), one for each interferometer. A tree is built up from software nodes of various types, each of which represents a smallish

¹³LOFAR will be extended to a size of several hundred kilometers eventually.

¹⁴This document deals exclusively with the radio astronomical use of LOFAR.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

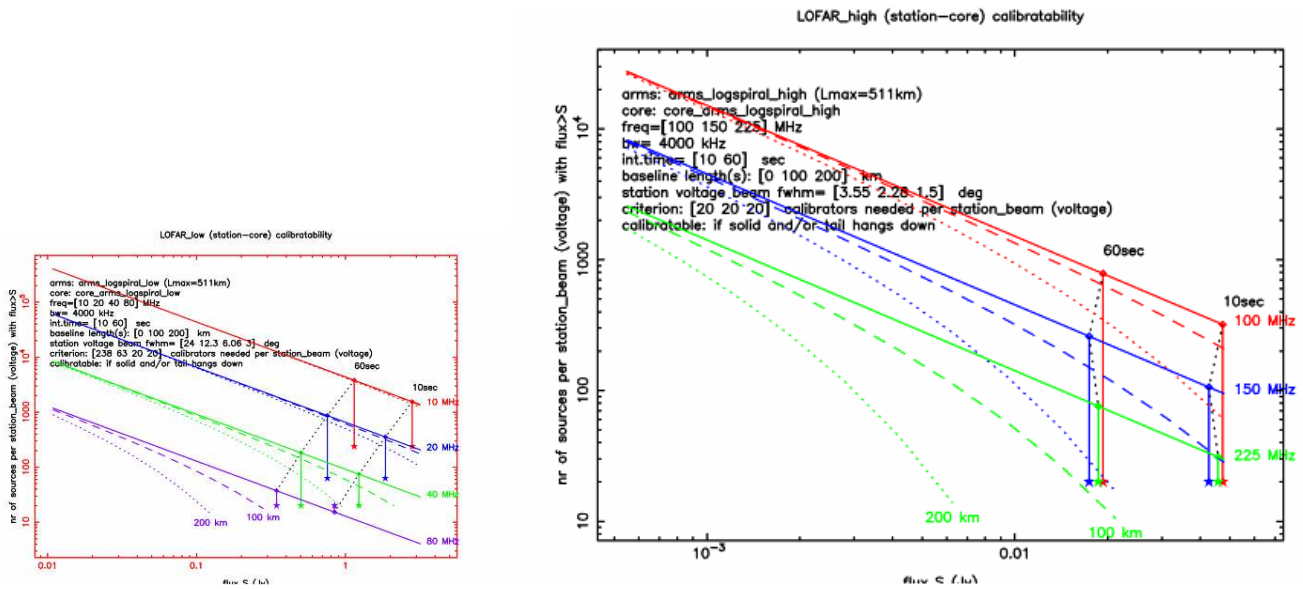



Figure 3: The availability of enough bright (Cat I) sources in the field is the cornerstone of LOFAR calibratability. The right one is larger to increase readability somewhat (see also below). The solid lines give the number of available sources in the (main lobe of the) station voltage beam, given the instrumental parameters in the top left corner. The stars indicate the numbers that are required for calibratability. Thus, things are alright as long as the 'tail' hangs down.

The solid lines are valid for short baselines, or for point sources. Since many of the LOFAR calibrator sources will be slightly extended, they will become less visible to longer baselines. This is indicated with the broken lines, for baselines of 100 and 200 km. The tentative conclusion is that any LOFAR station must be closer than about 50 km to another station. This will become an important consideration when the LOFAR array is be extended in a few years time.

NB: IT SHOULD BE NOTED THAT THESE FIGURES WERE MADE FOR THE PDR, A FEW YEARS AGO. WE WILL ENDEAVOUR TO REDRAW THEM BEFORE THE CDR. HOWEVER, THE OLD PDR FIGURES CONTAIN RELEVANT INFORMATION, EVEN THOUGH THERE ARE SOME DIFFERENCES. THE LBA STATIONS WILL HAVE A DIAMETER OF 60 M RATHER THAN 100 M, AND THE HBA STATIONS 50 M RATHER THAN 73 M. THE TOTAL NUMBER OF ELEMENTS ANTENNA UNITS WILL BE 7600 RATHER THAN 10400. THESE TWO POINTS ROUGHLY CANCEL EACH OTHER IN THE QUESTION OF AVAILABLE CALIBRATOR SOURCES. THE BOTTOM LINE IS THAT THERE ARE ENOUGH, ALSO BECAUSE WE HAVE SINCE THEN REALISED THAT OUR INITIAL REQUIREMENTS WERE TOO SEVERE. ESPECIALLY THE MIM NEEDS VERY MUCH FEWER CALIBRATORS THAN INDICATED HERE. THE CHARACTERIZATION OF THE STATION VOLTAGE BEAMS STILL REQUIRES UP TO 20-30 CALIBRATORS, BUT THEY MAY BE LESS BRIGHT BECAUSE WE CAN INTEGRATE LONGER. THUS, IT IS NO LONGER NECESSARY TO USE THE LOFAR CORE AS A SUPER-STATION, WHICH MAKES THE CALIBRATION SCHEME MUCH SIMPLER.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	





Figure 4: The uv -data are processed by domains, or snippets, i.e. rectangles in freq-time space, subdivided in cells. A snippet usually covers only a few time-slots (a few minutes), and one or more spectral windows. Mapped on the uv -plane, snippets tend to be longer in the **radial** (freq) direction than in the **lateral** (time) direction, especially if the fractional bandwidth is relatively large (10% here). Note that the area of a snippet is proportional to the distance to the origin of the uv -plane, i.e. the projected baseline length. Thus, longer baselines cover more uv -plane than shorter ones. This is an important consideration in Multi Frequency Synthesis (MFS).

Also indicated in the figures are the ridges of the cosine corrugation pattern (red) that represents the visibility function of an off-axis source. The period of the corrugation is inversely proportional to the distance of the source to the phase centre ($l = 0, m = 0$). In various calibration issues (peeling contamination, time/freq smearing), the number of visibility corrugations across a snippet, or even a cell, plays an important role. The figure on the right illustrates how this number depends on source position, baseline orientation, and snippet size in the time and freq direction. The most practical way to maximize the snippet size is to increase the fractional bandwidth.

Another important issue is the number of cells in a snippet. At full uv -data resolution (e.g. 1 s, 10 kHz), this can easily exceed 1000. Some operations, like shifting the phase centre, or source subtraction, have to be performed at full resolution. But if the visibility function of interest varies only slowly over the snippet domain, and we need to solve for only a few coefficients in the freq direction, we may resample to a (much) smaller number of larger cells. This represents a considerable saving in processing and memory use, especially when we consider that we also need to calculate derivatives w.r.t. to each solvable coefficient, for each cell and for each solver iteration. This is one of the main reasons for shifting the phase centre of the uv -data to the apparent position of the peeling source.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

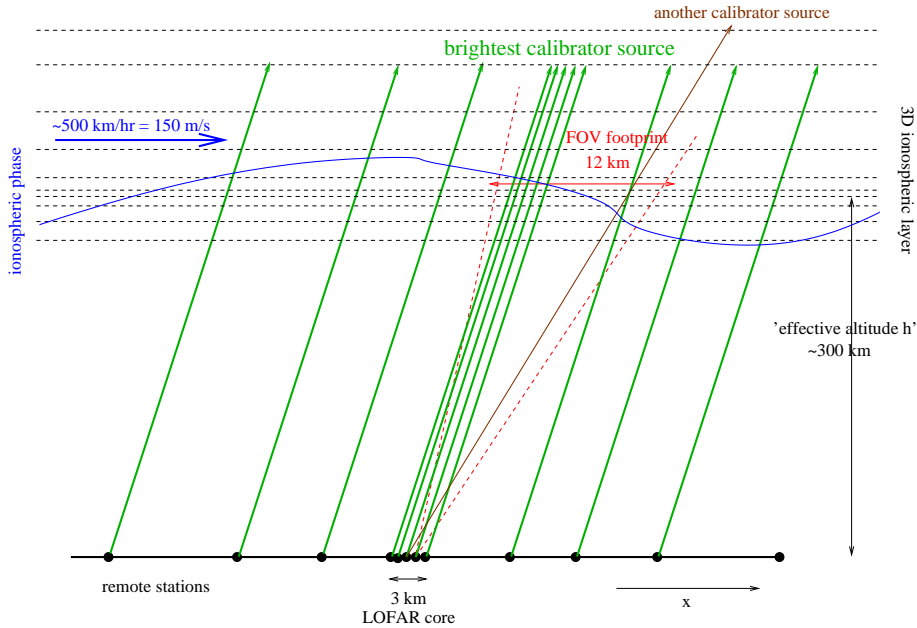



Figure 5: The Minimum Ionospheric Model (MIM) is not concerned with the internal structure of the 3D ionosphere, but only with the phenomenological phase $\phi(x, y, l, m, f, t)$, as seen from the position (x_i, y_i) of a particular station (i) , in a particular viewing direction (l, m) . The procedure is to postulate the simplest possible function $\phi = \phi(x, y, l, m, f, t)$, consisting of a minimum number of smooth parametrized base functions, e.g. low-order polynomials. For instance, in one dimension, we might postulate $\phi_0(x) = \lambda(p_0 + p_1x + p_2x^2 + \dots)$ for the ionospheric phase in the zenith direction. Note that this includes the a priori knowledge that $\phi \propto \lambda$, so the parameters p are only functions of time (i.e. not freq). A reasonable starting value would be $p_0 = -25TEC$ rad, with the integrated Total Electron Content in TEC units ($10^{16}m^{-2}$). Thus, for $\lambda = 3m$ and a typical night-time value of $TEC = 5$, the excess phase would be -375 rad.

For zenith angle z , the excess path will be longer, and the ionosphere will be 'pierced' at a different position. This can be expressed as $\phi(x, z) = \phi_0(x - h \tan(z)) \cdot S(x, z)$, where h is a 'coupling parameter with the dimension of an effective altitude (e.g. $h=300km$), and $S(x, z)$ incorporates the ionospheric charge profile and the curvature of the Earth surface. See [10] for a more thorough discussion.

The MIM parameters p_k (including h !) are secondary M.E. parameters, i.e. they are estimated every 10 s from the individual selfcal phase solution(s) of one or more bright calibrator sources. From the figure it can be gleaned that a single calibrator in the FOV would suffice to constrain the MIM, provided the LOFAR array is larger than the 'footprint' of the FOV at 'the' effective altitude of the ionosphere (~ 300 km), or if it may be assumed to be a simple linear phase wedge. However, with only using the LOFAR core (< 5 km), it is safer to use more calibrators, either inside the FOV, or outside. In the latter case, it could be advantageous to rapidly switch the LOFAR beam direction, so that it points directly at the calibrator for a few seconds. More calibrators will be needed as the ionospheric conditions deteriorate, i.e. as the structure size decreases.

Once its parameters are known, the MIM is able to calculate 'absolute' ionospheric phases for any station (\vec{x}) , and any viewing direction (\vec{l}) . These phases are used to shift the phase-centre of the uv -data to the apparent source/patch position when doing Cat I selfcal, or when predicting Cat II sources for subtraction.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	




Figure 6: *LOFAR station voltage beams are much less well-behaved than we are used to. For parabolic dishes, they are assumed to be constant in time, and roughly equal for all stations. The latter has the additional advantage that their phases largely cancel, because they are the same for the same source position in the two station beams of an interferometer. So, if the beam is roughly circular, or does not rotate on the sky (equatorial mount, WSRT), the response of 'the' primary (power) beam may be removed by division in the image-plane. Non-circular effects like instrumental polarization are assumed to average out when they rotate over the sky during a 12 hr observation (alt-az mount, VLA/ATCA).*


All this will change with LOFAR. Because the dipole arrays are fixed horizontally on the ground, the station beamshapes will become elongated at lower elevations, and these ellipses will rotate w.r.t. the sky as a function of azimuth. This is illustrated in the three images above, which show a highly schematic voltage beam for three (az, el) directions during a single observation. It is schematic because this is the voltage beam that would be produced by a filled aperture station. The dotted lines indicate the positions where the gain is zero. The phase increases with the distance from the centre. The central ellipse is the main lobe. An actual voltage beam will have more lateral structure, because of the discrete dipoles (see fig 1).

The source near the field centre stays in the main lobe, but will be subject to substantial gain (and some differential phase) variations. The other source wanders through different sidelobes. This will make it more difficult to deconvolve Cat III sources in residual images, because they will be convolved with an error PSF that depends rather strongly on their position (\vec{l}). The solution is to derive an analytical function for this error $PSF(\vec{l})$, and use that for deconvolution. Since Cat III sources are close to the noise ($< 10\sigma$), it is sufficient to use only the inner part. Another approach is to circularize the station beam by beamforming, at the cost of some sensitivity. A combination of these two solutions is also possible.

A station voltage beam is described by its 2×2 EJones matrix. For LOFAR, it will have 4 non-zero complex elements, giving substantial instrumental polarisation. Also, because of low phase-shift resolution of the analog beamformer of the HBA, its beams will jump to a new position every 10 min or so, which will have an effect on the station beamshape (although this appears to be quite small in the observations with the WHAT test station!). And finally, because of temperature variations and cheap electronics, the beamshapes will vary individually, in unpredictable ways.


For all these reasons, LOFAR station voltage beamshapes must be measured continuously, using the brighter sources in the field. The procedure is to obtain complex gains in the direction of these bright sources (by peeling), and to use the results to estimate the parameters (f, t) of a suitable beamshape function.

Finally, since it will not be practical to subtract Cat II sources in the far sidelobes, every available method must be used to attenuate them as much as possible.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

subset of mathematical operations. The main function of a node is to return values upon request, for the cells of a given domain (see below). It just passes the request to its children, and combines their results according to its specific function. The end-points of a tree are 'leaf' nodes, which have no children, but use other sources of information to return the requested values. Examples are data nodes and parameter nodes. A different M.E. requires a different forest of trees. It should be possible to generate new forests quickly and easily, for rapid experimentation.

2. All M.E. *parameters* are functions of freq and time, and sometimes of other dimensions as well (e.g. station beamshapes). Their values are stored in tables, in the form of zero or more *funklets*, i.e. arrays of coefficients of suitable base-functions. Each funklet has its own validity domain. The simplest and most common funklet is the $\text{polc}(f,t)$, i.e. a 2D array of coefficients of a 2D time-freq polynomial. When solving for an M.E. parameter, we actually solve for funklet coefficients.
3. Operations on uv-data are done per *domain*, i.e. a rectangle in freq-time space. Each domain is subdivided in cells, for which the trees generate values. The number of cells, i.e. the domain resolution, is determined by the available data (e.g. 1 s, 10 kHz), but can be changed by resampling. Mapped on the uv-plane, domains tend to be radial, i.e. longer in the freq direction than the time direction. See fig 4.
4. Information about all the relevant sources for a particular observation resides in a Local Sky Model (LSM). This is an object with three interfaces: With the Global Sky Model (GSM), the uv-data processing kernel, and residual images. See fig 7. The LSM sources are grouped in *punits* (prediction units), which may represent individual sources, or small areas (patches) of the sky. Very importantly, the LSM contains some kind of *predisol* mechanism for each source. The latter defines the relationship between its four image manifestations (I,Q,U,V) and its parametrisation, which is necessary to predict their visibilities. It also allows solving for source parameters.
5. There is no fundamental distinction between instrumental parameters and LSM source parameters (the ionosphere is regarded as part of the instrument). They are all parameters of the M.E. It must be possible to solve for any subset of them. The system must be able to measure the ortogonality of such subsets, and to take graceful remedial action in case of inter-dependence. There are multiple solvers in the system, each for its own subset of M.E. parameters.
6. There are various methods for predicting corrupted source visibilities. Especially the application of position-dependent (image-plane) effects requires attention. All methods allow solving for source parameters.
 - (a) Point sources present no problem. Their true visibilities may be calculated for each interferometer, and then corrupted with instrumental errors valid for their position. Sources that are just resolved may be treated in the same way, assuming that the errors are constant over their extent.
 - (b) Compact extended sources may be modelled as *shapelets* (see fig 8). Calculating visibilities is efficient, since the Fourier Transform of shapelet base functions is relatively cheap. Image-plane effects are applied by using a (small) image of the source, made from its shapelets. For each interferometer, a corrupted image is made by multiplication, and decomposed into a new set of corrupted shapelets. These are then transformed into corrupted visibilities.
 - (c) Visibilities of (patch) images are predicted by means of *uv-bricks* [?]. These are 4D cubes of gridded visibilities, over the entire uv-plane. They may be interpolated to give values for a particular interferometer, for a requested domain(f,t). This is relatively cheap if the uv-grid is coarse, which is the case for smallish patches near the phase centre. Image-plane effects are applied by adding a few terms, different for each interferometer, to the interpolation function.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	

7. Peeling. The punits (sources/patches) are dealt with one by one, and subtracted (peeled) from the uv-data before moving on to the next one. An *essential feature* of peeling is that the phase-centre of the uv-data is moved to the apparent position of each punit, thereby smoothing its visibility function over the requested domain. This makes it possible to reduce the number of cells, since we require that the visibility function of interest is approximately linear over a cell.

- (a) In the case of Cat I selfcal, the main advantage of peeling is efficiency. Partly because of smaller solution matrices, since we deal with one source at a time. But mostly because selfcal prediction (which is dominated by the calculation of derivatives!) is needed for very much fewer domain cells. A potential problem with peeling is the selfcal error caused by *contamination* of other (fainter) sources in the field.
- (b) In the case of Cat II subtraction, uv-bricks are used to predict these fainter sources in groups (patches). This requires the phase-centre to be moved to the centre of the patch.


D Special techniques

There are a number of special techniques available in various existing packages that have been designed bring out certain astrophysical features from the data. Most of them can be used in calibration also, so they should be available in the LOFAR toolbox. At the very least, they should not be designed out. Therefore, they are just listed here, without specifying an area of application. *The important thing is to be able to include such techniques in the processing, rapidly, efficiently and naturally.* Some examples:

- **Freq differencing:** Subtract adjacent freq channels from each other. Used for EoR detection and Cat II/III subtraction
- **Time differencing:** Subtract adjacent time-slots from each other. Used for transient detection and Cat II/III subtraction
- **Continuum subtraction:** Subtract a low-order polynomial from uv-data spectra. This removes the continuum from sources close to the phase centre, in a messy sort of way. But it does enhance the contrast for spectral features. Called UVLIN in AIPS (pioneered by Cotton and Van Langevelde, and elaborated in Miriad by Sault et al).
- **RM synthesis:** Vary the Rotation Measure and look for peaks in Q/U.
- Apply **large-scale constraints**, like $V=0$.
- **Fringe fitting:** From VLBI. Vary the phase and look for peaks in I.
- **Subspace decomposition:** See fig 9. Used for Cat II/III subtraction and RFI removal.
- **Using Bayesian learning of interference signals for effective flagging.**

E Some words on efficiency

This subject falls outside the scope of this document, which just deals with the principles of LOFAR calibration, and not its implementation.

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
Kind of issue: Public		Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

E.1 Cutting corners (suggested optimisations)

LOFAR calibration will be very expensive in terms of processing power and memory use. Potential bottlenecks are Cat II subtraction, number of solver iterations, number of passes through the main cycle, etc... Therefore, we must try to cut all possible corners in processing. Local intelligence in the various parts of the calibration process can reduce processing considerably, especially if it is done *dynamically*. It must also be done judiciously, of course, since it may have effects on the result. Some suggestions:

- Minimize solver iterations by using solver metrics, and continuous solutions.
- Use relatively simple flagging algorithms, at multiple levels.
- Minimize the subset of uv-data used for selfcal (e.g. only the longer baselines).
- Minimize the number of MIM terms, depending on the state of the ionosphere.
- Minimize the number of contamination sources to be taken into account for selfcal, i.e. only whenever $(ul_k + vm_k)$ is small. Here (l_k, m_k) is the position of the contaminating source w.r.t. to the phase-centre (i.e. the position of the peeling source).
- Use uv-track crossing points (weak redundancy) as solver constraints.
- Subtract only the minimum number of Cat II sources (depends on changing size of main lobe)
- Do *incremental* subtraction of Cat II sources: keep the uv-data residuals, and subtract only those Cat II sources that have changed in this cycle (obviously, this is impossible if any instrumental errors have changed!)
- etc..


Finally, it is good to realise that each major cycle will require more processing than the one before. The multiplication factor could be 2 or even more, i.e. more than half the processing could be in the last cycle.

E.2 Look for alternative methods

E.g. subspace decomposition, etc. See section 8.

F The role of simulation

The development of the various LOFAR calibration strategies within this framework will require extensive testing on simulated data. The selfcal stage, with its prediction of corrupted uv-data values, represents a natural simulation capability. However, it is obviously less desirable if simulation and calibration are done with the same system. Fortunately, LOFAR has two systems: one for development and experimentation (MeqTrees), and a sleek operational version (BBS). The MeqTree system has already demonstrated its capabilities in this area, and is poised to play an important role in SKA simulations. At this moment, most elements of LOFAR calibration are only available in their MeqTree implementation. As soon as their BBS counterparts are ready, they can be tested with simulated data generated by the other system.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	





Figure 7: The Local Sky Model (LSM) contains information about all the sources that are relevant for a particular observation. The LSM contains a list of parametrised sources and 4D images, and some database of parameter values for the parametrised sources. In addition it contains a list of punits (prediction units) in descending order of brightness. A punit can be an individual source, or a smallish image (patch) that contains multiple sources or an image. Finally, it has some objects to help it make decisions. Its 'obswin' contains information about the spectral window, and an idealised primary beam. Its 'obsres' contains information about spectral and spatial resolution. An LSM has three interfaces:

1. **With the Global Sky Model:** The GSM contains all the sources that LOFAR can see. It uses its obswin to select a subset that is relevant for the present observation. Afterward processing, the GSM may be updated from the modified LSM.
2. **With the uv-data processing kernel:** The punit list obtained from the LSM is used to generate a forest of suitable processing trees with relevant Cat I selfcal, and Cat II subtraction stages. For each source, the 'predisol' mechanism (e.g. subtree, see below) in the LSM is connected to the relevant part of the trees.
3. **With the residual images:** The LSM plays an important role in source extraction. Firstly, the (deconvolved) 4D images are inspected at the positions where Cat I/II sources have been subtracted. Any remains are used to solve for incremental improvements in the source parameters, using the LSM predisol mechanism. Secondly, new sources may be identified and created in the LSM. One of the most interesting problems will be to automatically choose a parametrisation for such sources.

Note that residual images are derived from imaging facets, and are quite unrelated to punit patches. Finally, a very important feature of the LSM is the 'predisol' mechanism for each source. It implements the mathematical relationship between the 4 image manifestations (I, Q, U, V) of a source, and its (arbitrary) parametrisation. It is able to predict IQUV, and solve for its parameters.

Author: E. Noordam	Date of issue: 15 Oct 2006 Kind of issue: Public	Scope: Project Documentation Doc.nr LOFAR-ASTRON-ADD-015	
	Status: Final Revision nr.: 1.0	File: <i>lofar/</i>	





Figure 8: *Shapelet decomposition, pioneered for us by Sarod Yatawatta, is expected to play an important role. At LOFAR frequencies and resolution, most sources will be more or less extended. They can be represented elegantly and efficiently by a (surprisingly) small number of shapelet coefficients. Shapelets are a set of orthogonal base functions (e.g. Hermite polynomials, multiplied by gaussians), which have the desirable property that they are their own Fourier Transform. There are Cartesian and Polar shapelets. Above, we show an input image, its shapelet coefficients and its reconstruction, and the difference image. The number of significant coefficients may be minimised by carefully choosing the origin (which is clearly not optimal here). Shapelets will be especially useful with the bright Cat I sources that are used for selfcal. Not only can they deal with arbitrary source shapes efficiently, but it is also possible to apply ifr-dependent image-plane effects, and to solve for both source parameters and instrumental parameters. Finally, shapelets provide an elegant mechanism for source extraction from residual images.*

Author: E. Noordam	Date of issue: 15 Oct 2006	Scope: Project Documentation	
	Kind of issue: Public	Doc.nr: LOFAR-ASTRON-ADD-015	
	Status: Final Revision: 1.0 nr.:	File: <i>lofar/</i>	

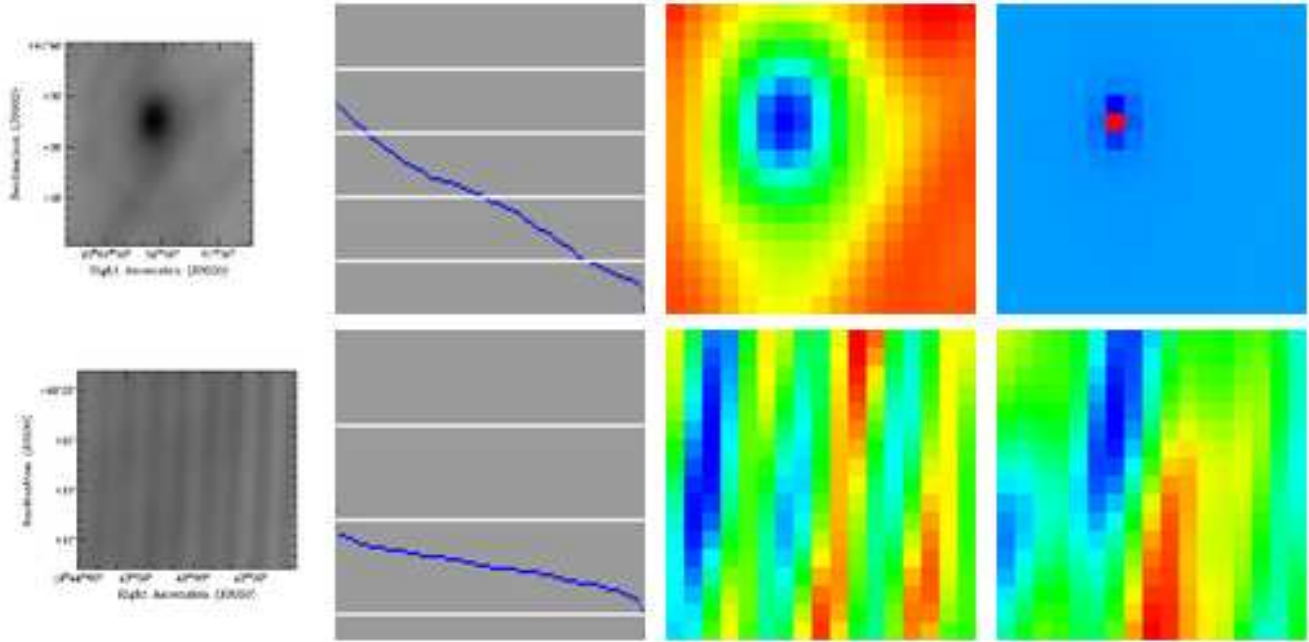


Figure 9: *Subspace decomposition is a promising technique to separate different kinds of structure in an image, a spectrum or a time-series. The method is being pioneered for us by Sarod Yatawatta, but it has been a staple of the LOFAR RFI group for some time. In the illustration above, the two rows of four images each show the first frequency plane of an input image cube, the eigenvalues of the associated autocorrelation in descending order, and the eigenmode reconstructions associated with the two largest eigenvalues. The larger eigenvalues are associated with structure, while the smaller ones are associated with noise. The eigenvalue plots are logarithmic, the maximum (left) being 1.0.*

In general, the inputs must be column vectors that 'look at' the same thing in different ways. These are correlated, and subjected to singular value decomposition (SVD). In this particular case, the input vectors are the frequency planes of an image-cube. They can also be the rows or columns of a single image, or 'parallel' time-series of visibilities for different baselines.

Possible applications are the removal of Cat II sources from uv -data, without knowing their details. This is important, because the regular predict/subtract method will be a major bottleneck. A potential problem is that this would also remove the Cat III sources, including the EoR signature etc. This may be avoided when we understand the technique better.

Another application is RFI detection in uv -data.