



HAL
open science

A numerical study of 21-cm signal suppression and noise increase in direction-dependent calibration of LOFAR data

M. Mevius, F. Mertens, L.V.E. Koopmans, A.R. Offringa, S. Yatawatta, M.A. Brentjens, E. Chapman, B. Ciardi, H. Gan, B.K. Gehlot, et al.

► To cite this version:

M. Mevius, F. Mertens, L.V.E. Koopmans, A.R. Offringa, S. Yatawatta, et al.. A numerical study of 21-cm signal suppression and noise increase in direction-dependent calibration of LOFAR data. Monthly Notices of the Royal Astronomical Society, 2021, 509 (3), pp.3693-3702. 10.1093/mnras/stab3233. hal-03440881

HAL Id: hal-03440881

<https://hal.science/hal-03440881v1>

Submitted on 20 Apr 2023

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

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

A numerical study of 21-cm signal suppression and noise increase in direction-dependent calibration of LOFAR data

M. Mevius¹,^{*} F. Mertens^{1,2,3}, L. V. E. Koopmans², A. R. Offringa¹, S. Yatawatta¹, M. A. Brentjens¹, E. Chapman⁴, B. Ciardi⁵, H. Gan², B. K. Gehlot⁶, R. Ghara^{7,8}, A. Ghosh⁹, S. K. Giri¹⁰, I. T. Iliev¹¹, G. Mellema¹², V. N. Pandey¹ and S. Zaroubi^{7,2,8}

¹*Astron, Oude Hoogeveensedijk 4, NL-7991 PD Dwingeloo, the Netherlands*

²*Kapteyn Astronomical Institute, University of Groningen, PO Box 800, NL-9700 AV Groningen, the Netherlands*

³*LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Université, F-75014 Paris, France*

⁴*Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK*

⁵*Max-Planck Institute for Astrophysics, Karl-Schwarzschild-Straße 1, D-85748 Garching, Germany*

⁶*School of Earth and Space Exploration, Arizona State University, 781 Terrace Mall, Tempe, AZ 85287, USA*

⁷*Department of Natural Sciences, The Open University of Israel, 1 University Road, PO Box 808 Ra'anana 4353701, Israel*

⁸*Department of Physics, Technion, Haifa 32000, Israel*

⁹*Department of Physics, Banwarilal Bhalotia College, 713303 Asansol, India*

¹⁰*Institute for Computational Science, University of Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland*

¹¹*Astronomy Centre, Department of Physics and Astronomy, Pevensey II Building, University of Sussex, Brighton BN1 9QH, UK*

¹²*The Oskar Klein Centre, Department of Astronomy, Stockholm University, AlbaNova, SE-10691 Stockholm, Sweden*

Accepted 2021 November 3. Received 2021 November 2; in original form 2020 October 12

ABSTRACT

We investigate systematic effects in direction-dependent gain calibration in the context of the Low-Frequency Array (LOFAR) 21-cm Epoch of Reionization (EoR) experiment. The LOFAR EoR Key Science Project aims to detect the 21-cm signal of neutral hydrogen on interferometric baselines of 50–250 λ . We show that suppression of faint signals can effectively be avoided by calibrating these short baselines using only the longer baselines. However, this approach causes an excess variance on the short baselines due to small gain errors induced by overfitting during calibration. We apply a regularized expectation–maximization algorithm with consensus optimization (SAGECAL-CO) to real data with simulated signals to show that overfitting can be largely mitigated by penalising spectrally non-smooth gain solutions during calibration. This reduces the excess power with about a factor of 4 in the simulations. Our results agree with earlier theoretical analysis of this bias-variance trade off and support the gain-calibration approach to the LOFAR 21-cm signal data.

Key words: methods: numerical – methods: observational – techniques: interferometric – dark ages, reionization, first stars.

1 INTRODUCTION

Detection of the faint 21-cm signal of neutral hydrogen emitted at high redshifts is one of the hardest radio-astronomical programmes currently pursued. This is in particular so, since the contaminating foreground signals are many orders of magnitude stronger. To extract the 21-cm signal from the data, requires an exquisite understanding of the entire signal chain.

The Radio Interferometric Measurement Equation (RIME, Hamaker, Bregman & Sault 1996; Smirnov 2011) is a generalized framework to describe the propagation of the signal of radio sources from the source to the radio interferometer. It includes the full (polarized) signal path, including atmospheric and instrumental signal distortion, in a matrix formalism using Jones matrices. A model of the sky brightness distribution, updated iteratively during gain calibration, predicts the coherence matrix for a given baseline. In its most general form, for a dual-polarization instrument, all external

effects on the propagation of an electromagnetic wave, are merged into a single time (t) and frequency (f) dependent complex 2×2 gain matrix for both elements i, j of an interferometer, which is applied to the sky model coherence matrix to predict the corrupted visibilities V_{ij} of a single baseline:

$$\mathbf{V}_{ij}^{\text{pred}}(t, f) = \sum_{k=1}^K \mathbf{J}_i(t, f) \mathbf{C}_{ijk}(t, f) \mathbf{J}_j^{\text{H}}(t, f) + \mathbf{N}_{ij}, \quad (1)$$

where the sum is over K discrete sky model components. \mathbf{J}_i is the 2×2 complex Jones matrix for antenna i , \mathbf{C}_{ijk} is the coherency matrix giving the contribution of each sky model component k to the visibility of baselines i, j , and \mathbf{N}_{ij} is a 2×2 complex noise matrix.

A *solve* step performs a non-linear fit of the predicted visibilities $\mathbf{V}_{ij}^{\text{pred}}$ to the measurement to determine the complex gains of the Jones matrices. For current wide-field low-frequency telescopes, a single Jones matrix per antenna is often insufficient, since atmospheric effects and station beam errors vary over the field of view. These direction-dependent (DD) effects are in general taken into account by dividing the sky model into several patches, with scales over

* E-mail: mevius@astron.nl

which these effects are thought to vary little, each with its own Jones matrix. This means grouping the K discrete components in equation (1) into N ($N \ll K$) directions. In DD calibration, the number of free parameters increases linearly with the number of solved directions, enlarging the risk of overfitting. In practice, the sky model will always be incomplete, and during calibration this leads to gain errors whose impact are sometimes difficult to predict. Among these effects are the appearance of *ghost* sources in the images (Wijnholds, Grobler & Smirnov 2016), the increase of noise (Barry et al. 2016; Patil et al. 2016) and the suppression of the signal of the unmodelled sky (Mouri Sardarabadi & Koopmans 2019).

In this paper, we study the unwanted signal suppression and noise increase caused by the large number of parameters in DD calibration through simulation. We investigate this for the specific case, namely the study of the 21-cm signal power spectrum from the Epoch of Reionization (EoR) with the LOFAR radio telescope. This study serves two goals: First, it illustrates the motivation behind the calibration strategy used by Mertens et al. (2020), thereby verifying the soundness of their results in light of signal suppression. Secondly, it allows us to set a framework for investigating a generic context radio interferometric gain calibration using many directions.

The EoR refers to the period in the evolution of the Universe in which neutral hydrogen was ionised through irradiation by the first stars and quasars. It is one of the main research areas of modern low-frequency radio telescopes such as GMRT (Paciga et al. 2013), LOFAR (van Haarlem et al. 2013), MWA (Morales & Wyithe 2010) and HERA (DeBoer et al. 2017). Although direct detection of the 21-cm signal from the EoR requires instruments with higher sensitivity, the signals are predicted to be detectable through a statistical measurement, typically via its expected signature in a 21-cm power spectrum. This measurement is especially challenging since the power of the expected signal is orders of magnitudes smaller than that of the foreground radiation. It necessitates long integration times (typically thousands of hours) and foreground subtraction with unprecedented accuracy, such that the power of the background noise reaches that of thermal noise. This poses stringent requirements on the level of calibration accuracy. One of the main assumptions of the EoR-signal measurement is that the foregrounds as well as any instrumental or atmospheric effects are smooth in frequency, whereas the signal is not. As will be shown, this assumed smoothness in frequency is important to minimize signal suppression and noise increase in calibration.

The layout of this paper will be as follows. In Section 2, we outline the problem of signal suppression and noise enhancement in the framework of the LOFAR EoR analysis and we present our methods. In Section 3, we summarize the processing steps used in the LOFAR EoR analysis and in our simulation. Section 4 describes the simulation of the faint 21-cm EoR signals and the subsequent analysis of signal suppression and increased noise power through overfitting. In Section 5, we investigate whether we can reduce both signal suppression and added noise by enforcing smoothness of the gain parameters. Finally, conclusions and recommendations are given in Section 6.

2 METHOD

We base our analysis on typical data from the Low-Frequency ARray (LOFAR) (van Haarlem et al. 2013). The LOFAR array consists of 24 core stations, located roughly within a circle with a diameter of 3 km in the Eastern part of the Netherlands, and 14 remote stations with baselines up to ~ 100 km. A further 14 international stations are

not used for the EoR analysis. We use the high-band antenna data of LOFAR, covering frequencies roughly between 110 and 170 MHz.

Patil et al. (2017) presented the first upper limit on the 21-cm signal power spectra from LOFAR. They analysed the data of one of the main LOFAR EoR fields, the North Celestial Pole (NCP). In their analysis, gain calibration was performed without the baselines that were used for the 21-cm signal measurement (i.e. those shorter than 250λ). We will refer to this as ‘*applying a baseline cut*’ to the data. LOFAR has a sufficient number of longer baselines to determine the (multi-directional) gains per station. After calibration, the gain solutions are applied to all baselines, including the shorter baselines. The reasons for applying a baseline cut are twofold: first, the Galactic diffuse emission that is dominant on those short baselines is not included in the sky model; and secondly, it ensures that no signal suppression occurs in the calibration, since all data that contribute to the 21-cm signal analysis are excluded from the calibration. However, as was shown by Patil et al. (2016), overfitting increases the noise level on the baselines that are excluded from the calibration, resulting in an *excess variance* in the 21-cm signal power spectrum. A similar increase in noise is shown by Barry et al. (2016).

Therefore, we have produced a sky model that includes the diffuse emission in the NCP. Here, we shall investigate whether the new model allows us to remove the baseline cut to reduce the effect of overfitting. We will also, through simulation, study the effect of removing the baseline cut on signal suppression. Mouri Sardarabadi & Koopmans (2019) provide a theoretical framework for signal suppression. In this paper, we investigate signal suppression via numerical simulations, using the same codes that are used to analyse the LOFAR data.

2.1 Consensus optimization

Various algorithms have been developed that implement the RIME in equation (1), aimed at calibrating radio interferometric instruments. In this work, the software used for calibration is SAGECAL-CO (Yatawatta et al. 2013; Yatawatta 2015, 2016). SAGECAL-CO makes use of fast distributed systems that include graphics processing units (GPU), and performs consensus optimization to iteratively force the station-based and direction-(in)dependent gain solutions to approach a spectrally smooth function. Such smooth gain solutions have previously been shown to be crucial to mitigate signal suppression and noise increase (see Barry et al. 2016; Patil et al. 2016; Ewall-Wice et al. 2017; Mouri Sardarabadi & Koopmans 2019). To ensure a spectrally smooth behaviour of the gain solutions, an extra penalty function is added to the optimization problem at every iteration. This so called consensus optimization uses an augmented Lagrangian, with a regularization parameter to guide the solution to approach the chosen (smooth) regularization function, which is itself a function of several free parameters. Hence, if chosen wisely, the solutions will exactly match this functional form, reducing overfitting, although this would theoretically take an infinite number of iterations. In practice, a maximum number of iterations and regularization values need to be chosen carefully to ensure minimal deviations from the smooth regularization function. Yatawatta (2015) gives more details on the implementation. The Lagrangian (equation 14 in Yatawatta 2015) to be minimized in SAGECAL-CO is

$$\mathcal{L} = \sum_i g_{fi}(\mathbf{J}_{fi}) + \|\mathbf{Y}_{fi}^H(\mathbf{J}_{fi} - \mathbf{B}_{fi}\mathbf{Z})\| + \frac{\rho}{2} \|\mathbf{J}_{fi} - \mathbf{B}_{fi}\mathbf{Z}\|^2. \quad (2)$$

Here, $g(\mathbf{J})$ is the usual least-squares cost function resulting from the difference between the measured visibilities and the model, given Jones matrices \mathbf{J} ; \mathbf{Y} is a Lagrange multiplier; \mathbf{BZ} is the

smooth function over frequency, with \mathbf{B} a matrix composed of polynomial terms (in frequency) and \mathbf{Z} the fitted parameters; and ρ is a regularization parameter that determines the level of smoothness in frequency that is enforced on the individual gain solutions during each iteration. Fitting of the gains is done iteratively. The polynomial coefficients \mathbf{Z} are updated after each iteration. In practice, the number of iterations is limited by computation time. We use a Bernstein polynomial of rank three to regularize the frequency dependence of the gain parameters. Yatawatta (2015) uses simulated data to determine values for ρ , quoting typical values of $\rho \in [1, 10]$ over a bandwidth of about 60 MHz. The strongest regularization used in that work is $\rho = 50$. As we will show, the regularization parameter plays a key role in regard to overfitting. We show that increasing ρ to a substantially higher value over smaller bandwidths, greatly improves the result on real data. Possible reasons for this difference could be the additional noise power and the incomplete sky model which were not included in the simulations in Yatawatta (2015) and which slow down convergence for real data. In a slightly different simulation in Yatawatta (2016), an optimal value of $\rho \sim 200$ was found for a fixed number of 50 iterations.

3 DATA PROCESSING AND GAIN CALIBRATION

In this section, the LOFAR 21-cm signal data-processing steps as followed by Patil et al. (2017) and the impact of gain calibration on the 21-cm signal are discussed. Processing of the raw LOFAR visibilities follows several distinct steps, which we will summarize below. For more details, we refer to Patil et al. (2017) and Mertens et al. (2020). All analyses are done in the context of the NCP field.

3.1 Flagging and calibration

Step 1 – The raw visibilities have a 2 s time and 3.05-kHz spectral resolution per channel. The data sets are split into sub-bands, each with 64 channels. Each sub-band is cleaned of strong RFI using the package AOFLAGGER (Offringa et al. 2010; Offringa, van de Gronde & Roerdink 2012). The outer two channels on each side of the sub-band are discarded to avoid spurious signals caused by the poly-phase filter. Visibilities are subsequently averaged to three channels per sub-band, resulting in a 61-kHz spectral resolution.

Step 2 – After this initial filtering, flagging and averaging follows the initial gain calibration of the visibilities. A model of the brightest sources for the NCP field, consisting of ~ 1500 components each with an apparent flux > 35 mJy, is used during the initial calibration, and a model of the station beam is applied to the sky model. In direction-independent calibration, only full-Jones station-based complex gain solutions are solved for, in fact solving for the gains in a single effective direction. These gain solutions can then be used to correct the visibilities for station gains and first-order atmospheric effects. However, the relatively bright source 3C 61.1 is located near the first null of the station beam, causing an apparent flux that is strongly frequency and time-dependent. These gain variations can potentially degrade the accuracy of the initial calibration. To mitigate this effect, two separate gain-solutions are solved simultaneously, with one set of gain solutions for 3C 61.1 and one for the remainder of the field. The solution intervals over which the gains are assumed constant are 10 s and 3 channels (183 kHz). The data are subsequently corrected by applying the latter gain solutions. During initial calibration, only baselines $> 30\lambda$ are used to avoid significant diffuse emission affecting the gain solutions. SAGECAL-CO is used over the full 60-MHz bandwidth, using a third-order Bernstein polynomial to

regularize the gain solutions in the frequency direction, but with a relatively low level of regularization, determined from theoretical considerations in Yatawatta (2015). This level of regularization allows for fitting of the non-smooth spectral gain variations, such as cable reflections and filter effects, but at the same time reduces overfitting of the thermal noise and sky emission that is not part of the calibration model.

Step 3 – Direction-dependent (DD) calibration is done by solving for full-Jones matrices in 122 independent directions, clustered on a sky model consisting of 28 000 components. The time solution intervals vary between 4 and 20 min, depending on the apparent flux in the cluster associated with a direction, and the frequency solution interval is 3 channels (183 kHz). Before DD calibration, the corrected data of the previous step is averaged to a time resolution of 10 s. The frequency resolution remains 61 kHz.

SAGECAL-CO allows for regularization of the gain solutions by enforcing the gains to iteratively approach a smooth function in frequency. We make the following choices:

- (i) During DD calibration, a third-order Bernstein polynomial over the full 60-MHz frequency range is used for spectral regularization.
- (ii) The level of regularization for each iteration is specified by a single parameter per cluster of sources (i.e. a direction). The values of the regularization parameters ρ are based on the theoretically estimated required level of regularization for the centre of the field and are scaled with 0.1 and the relative apparent flux in the clusters. It is set to $\rho \approx 50$ for the brightest clusters and close to $\rho \approx 1.0$ for the faintest clusters.
- (iii) Each cluster of source components is subtracted from the visibilities after multiplication with their DD gains. This assumes that the gains are constant over the spatial extent of the cluster.

In Section 5, it is shown that this level of regularization still allows for small but significant frequency variations in the visibilities, which are caused by un-modelled sky emission (Barry et al. 2016). As the gain solutions have enough freedom to absorb part of the difference between the real sky and the sky model, this causes signal suppression. To avoid this, the baselines used during calibration and signal-extraction are fully separate in all of the current LOFAR 21-cm results (Patil et al. 2017; Mertens et al. 2020).

3.2 Baseline selection

Patil et al. (2017) set a lower limit of $> 250\lambda$ on the baselines used during DD calibration, under the assumption that the unmodelled diffuse emission, dominant on the shorter baselines, could result in signal suppression. In subsequent analysis, a new model, including diffuse galactic emission modelled at the shortest baselines ($< 250\lambda$) was introduced. In Section 4, we will investigate the effect of the DD calibration scheme with and without the baseline cut on a simulated 21-cm signal and on the noise power.

3.3 Power spectrum generation

After calibration and sky model subtraction, the residual visibilities are imaged, using WSCLEAN (Offringa et al. 2014). Imaging is performed with settings that are sufficiently accurate for 21-cm power spectra (Offringa et al. 2019). Each frequency channel is imaged individually, resulting in an image cube with dimensions l, m, f , with l, m the direction cosines and f the frequency. For the LOFAR-EoR analysis, only the baselines $< 250\lambda$ are taken into account. However, for inspection, we also make use of baselines up

to 500λ . The most common method to statistically detect the 21-cm signal is to look at the power spectrum of the data in cosmological units (Morales & Hewitt 2004). The power spectrum is generated by mapping the l, m, f coordinates to comoving distances and Fourier transforming to wavenumber (\mathbf{k}) space. We will present the results as a 2D power spectrum, distinguishing between the modes parallel (k_{\parallel}) and perpendicular (k_{\perp}) to the line of sight. This enables the possibility to specifically examine baseline (proportional to k_{\perp}) and frequency (proportional to k_{\parallel}) dependent effects.

4 SIMULATION OF FAINT SIGNALS

In order to test the level of suppression, we add mock 21-cm signals with varying amplitudes to simulated data. The simulated data represents the real data after initial calibration (Step 2 in Section 3), and includes realistic DD corruptions. The simulated data are formed by multiplying the 28 000 component model (without diffuse emission) with the DD gain solutions from real data calibration. Since these corruptions were taken from real data calibration, they were not forced to be spectrally smooth. In reality, we expect the gains, which are mainly the result of instrumental and ionospheric effects, to be spectrally smooth. Therefore, applying strong constraints on smoothness of the gain parameters during calibration, is not expected to introduce additional noise in real data, but it might add some noise in our simulations. This effect of this additional noise power in our simulations will be discussed in Section 5.1, where enforcing smooth gain solutions is studied.

A realistic level of complex Gaussian noise is added to the data. The variance is derived from the Stokes V noise in real data after DD calibration. The diffuse foreground emission is simulated using the model for Stokes I, Q and U that was introduced before for calibration purposes (see Section 3.2). This model consists of shapelets fitted to real data. For practical reasons, the diffuse emission is not corrupted with gain errors. However, in the calibration step, it is treated as one independent direction out of many, with its own gain solutions, which are expected to be close to the simulated gains of 1. Therefore, this simplification is not expected to have a large effect on the study presented here. In reality, the diffuse emission will be corrupted with DD errors, possibly varying over the extent of the emission. Solving for varying gains over the extended emission is currently not part of the LOFAR EoR calibration scheme, and therefore out of scope of this study. The effect of neglecting gain variations over the extended emission will likely add to the signal suppression in a similar way as an incomplete sky model does. The model that is used for calibration of the simulated data includes all simulated sources and the diffuse emission, but does of course not include the 21-cm signal. This implies that the sky model is complete, i.e. with no emission from unmodelled sources. In Mouri Sardarabadi & Koopmans (2019), it was shown that an incomplete sky model during calibration would add to the level of suppression. Therefore, we also performed one test in which we excluded one third of the simulated sources with the lowest flux from the sky model used during calibration. This resulted in no significant difference in the level of suppression with respect to the calibration using the complete sky model. Within our accuracy limits, we could therefore not validate that conclusion from Mouri Sardarabadi & Koopmans (2019). Although in reality the sky model will be incomplete, we are confident that the use of the complete sky model in simulations gives a reliable estimate of the minimal level of signal suppression during calibration.

The simulated 21-cm signal is generated from an image cube covering 115–200 MHz, with a frequency resolution of 0.5 MHz and a spatial resolution of 1.17 arcmin, that was originally pro-

duced by Jelić et al. (2008) using the 21CMFAST SIMULATION CODE (Mesinger, Furlanetto & Cen 2011). The image data are linearly interpolated to the LOFAR subband frequencies. Visibilities are predicted from the image cube using WSCLEAN (Offringa et al. 2014). Because a realistic 21-cm signal is orders of magnitude less bright than the noise in a single night of LOFAR observations, the simulated signal is artificially increased by three orders of magnitude. According to Mouri Sardarabadi & Koopmans (2019), the level of suppression is independent of the amplitude of the 21-cm signal, as long as the signal is faint enough not to significantly alter the gain solutions. To test the validity of this linearity, we add the signal with three different amplitudes with another factor of 4, 7, and 10, relative to the original multiplication factor of 1000 of the simulated 21-cm signal, still one to two orders of magnitude smaller than the simulated foreground signal. Just as the diffuse model, the simulated 21-cm signal is not corrupted with the SAGECAL gains. Since the residual visibilities are the result of subtracting the corrupted sky model, and not corrected with DD gain solutions, the fact that beam and atmospheric effects are ignored in the 21-cm signal is likely to have a negligible impact on this suppression study. Expanding equation (1) to specifically include the different components of our model, we obtain

$$\begin{aligned} \mathbf{V}_{ij}^{\text{pred}}(t, f) = & \sum_{k=1}^N \mathbf{J}_{ik}(t, f) \mathbf{C}_{ijk}(t, f) \mathbf{J}_{jk}^{\text{H}}(t, f) \\ & + \mathbf{G}_i(t, f) \sum_{s=1}^S \mathbf{D}\mathbf{G}_{ijs}(t, f) \mathbf{G}_j^{\text{H}}(t, f) \\ & + \mathbf{F} \cdot \mathbf{I}_i(t, f) \mathbf{21cm}_{ij}(t, f) \mathbf{I}_j^{\text{H}}(t, f) + \mathbf{N}_{ij}. \end{aligned} \quad (3)$$

Here, C_{ij} refers to the point source foreground model, where each cluster k is the sum over the individual components in the cluster. DG_{ij} is the modelled diffuse galactic emission, composed of S shapelet coefficients. 21cm_{ij} is the simulated 21-cm signal, multiplied with an artificial constant amplitude F . G and I are set to the identity matrices in our simulation and the J matrices are taken from earlier runs on real data. In the calibration, we solve for J and G .

4.1 Suppression measurement

We run SAGECAL-CO (Yatawatta 2016) with standard settings, but without the 250λ baseline cut, on the simulated data with and without an added 21-cm signal. The computationally intensive DD calibration step for a typical observation (~ 14 -h duration, 60MHz bandwidth) takes about 72 hours when processed using half of the dedicated high-performance GPU cluster (Pandey et al. 2020). This limits the possible number of independent simulation runs that can be executed in a reasonable time. We make sure that, by construction, a solution exists. It will, however, not be recovered precisely because noise has been added and the calibration is a non-linear process, which causes the gain solutions to be inexact. The latter point causes additional noise in the solutions, which we will refer to as the *solver noise*. Since SAGECAL-CO uses random parameter initialization, we can estimate the power of the solver noise by running SAGECAL-CO twice on the same simulated data and calculating the power spectrum of their difference. We assume the power of the solver noise to be consistent between runs. The power spectrum of the solver noise calculated this way is shown in the top left-hand panel of Fig. 1. The top right-hand panel of Fig. 1 shows the power spectrum of the simulated 21-cm signal before DD calibration. The bottom left-hand panel of Fig. 1 shows the power spectrum of the simulated 21-cm signal after calibration, and the




Figure 1. Top left-hand panel: Power spectrum of the solver noise. Top right-hand panel: simulated 21-cm signal (relative amplitude 10, see the text for explanation). Bottom left-hand panel: remaining 21-cm signal after DD calibration. Bottom right-hand panel: Ratio of 21-cm signal after and before calibration. DD calibration was performed without the 250λ cut on baseline length.



bottom right spectrum shows the ratio of the power spectra of the simulated signal after and before DD calibration. The solver noise is one to two orders of magnitude smaller than the added signal, but it does have an impact on the ratio in the bottom left-hand plot, albeit minor. In general, due to the solver noise, the power spectrum after SAGECAL will be enhanced and therefore the level of suppression will be underestimated. We compensate for this by subtracting the power of the solver noise from the signal power per bin. As can be seen from Fig. 1, the suppression of the signal is large, roughly one order of magnitude. Note that in this suppression measurement all baselines are included in DD calibration and the level of regularization is relatively low, in contrast to the analyses presented in Patil et al. (2017) and Mertens et al. (2020). The results with three different levels of input 21-cm signal are summarized in the top three rows of Table 1. The level of suppression is measured by averaging the power over all 750 k bins and taking the ratio with the same average of the input signal. They show that the level of suppression is more or less independent of the signal amplitude, as predicted by Mouri Sardarabadi & Koopmans (2019). There is a small correlation between signal strength and its measured suppression, leading to slightly more suppression on smaller input signals, but it is negligible compared to the overall level of suppression, and our current measurement is therefore a realistic indication of the level of suppression expected for signals with much smaller amplitudes, such as a true 21-cm signal.

4.2 Overfitting

From Section 4.1, it is clear that not using a 250λ baseline cut leads to unacceptable levels of signal suppression, even when diffuse emission is included in the model. Here, we test the effect on the signal and noise when a baseline cut is applied. As discussed in Patil et al. (2016) and Yatawatta (2015), excluding baselines from the calibration will increase the variance on these baselines. The significance of this effect is illustrated in Fig. 2, where the Stokes-V power spectrum ratios after and before SAGECAL are shown over a larger range of scales (between 50 and 500λ). For LOFAR, the Stokes-V correlations consist almost only of system noise, and Stokes V allows us therefore to analyse the effect of calibration on a noise-like signal. In the left-hand panel of Fig. 2, we show the power ratio of the settings as described in Section 4, using a diffuse model and all baselines included during calibration with SAGECAL-CO. We see a noise power ratio that is close to one. However, when applying the 250λ baseline cut, as shown in the right-hand panel of Fig. 2, a clear increase of the noise up to a factor of ~ 5 in power is observed on the largest scales. The turn up at 250λ , corresponding to k_{\perp} of 0.25 , is clearly visible. This turn up is a sign of overfitting of the data. Part of the noise is absorbed into the gain solutions and therefore transferred to the shortest baselines. Note that in our simulations the sky model is complete, so in this case, only the noise adds to overfitting of the data (Barry et al. 2016; Patil et al. 2016). In the following section, we will investigate whether adding more spectral constraints to the

Table 1. Remaining simulated 21-cm signal power after DD calibration with different settings. The calibration was performed with and without applying the 250λ cut and with a different regularization parameter and bandwidth.

| DD calibration scheme | Relative signal amplitude | Recovered power ratio | Solver noise fraction |
|--------------------------|---------------------------|-----------------------|-----------------------|
| No cut, low reg, 60 MHz | 4 | $8.3 \pm 0.2\%$ | 15% |
| No cut, low reg, 60 MHz | 7 | $8.7 \pm 0.2\%$ | 6% |
| No cut, low reg, 60 MHz | 10 | $9.1 \pm 0.2\%$ | 4% |
| Cut, low reg, 12 MHz | 10 | $103 \pm 3\%$ | 16% |
| No cut, high reg, 12 MHz | 10 | $65 \pm 2\%$ | 5% |
| Cut, high reg, 12 MHz | 10 | $100 \pm 2\%$ | 6% |

**Figure 2.** Ratio of the Stokes V power spectra after DD calibration and the input noise. Left-hand panel: all baselines used during calibration. Right-hand panel: calibrated with baselines larger than 250λ .**Figure 3.** Left-hand panel: Input amplitudes and amplitude solutions for a single timeslot, station and direction (≈ 3 degrees away from the field centre), using the four different calibration schemes described in 4 and table 1. Even though a spectral Bernstein polynomial constraint was applied during calibration, the solutions still fluctuate in frequency when the regularization parameter SAGECAL-CO is low. The large outliers correspond to subbands that had a lot of RFI (in real data) and therefore many of the data were flagged. The simulated gains were taken from real data calibration with low regularization and therefore have spectral structure. Middle panel: zoom of the left plot on the 12 MHz used in the latter 3 calibration schemes. A clear improvement in the smoothness of the solutions is observed when high regularization is applied. Right-hand panel: time-averaged delay spectrum of the complex gain solutions, same cluster and station. The reduction in noise power with higher regularization is obvious. Note the different binning due to the five times larger bandwidth of the first two lines.

gain solutions by increasing the regularization will reduce the level of signal suppression and/or the overfitting effect.

5 SPECTRAL REGULARIZATION

The calibration scheme discussed in Section 4 has also been applied on real data. As expected, removing the 250λ cut considerably reduces the noise in the Stokes I and V power spectra, compared to Patil et al. (2017). However, it also has a major effect on the 21-cm

signal itself. The solid black line in Fig. 3 left shows the DD gain variation for one timeslot and one of the 122 directions. From this figure, it is apparent that the final DD solutions are not spectrally smooth on MHz scales. From the physical processes involved (i.e. the ionosphere and the instrument response), one would expect the true gains to be spectrally smooth up to scales of 10 MHz or larger. The overfitting discussed in the previous section manifests itself in the irregularity of the solutions on much smaller frequency scales. As discussed by Mouri Sardarabadi & Koopmans (2019), smoother




Figure 4. Demonstration of the effect of the regularization parameter on the average gain variance, using the SAGECAL-CO solutions from a single cluster (real data).

solutions may mitigate the effect of signal suppression and the excess variance. We can enforce smoothness by increasing either the number of iterations or the value of the regularization parameter in SAGECAL-CO. A third-order Bernstein polynomial might also be too constraining to describe the expected frequency variation over the full 60-MHz bandwidth. Therefore, we reduce the bandwidth to 12 MHz, covering only 134–146 MHz, corresponding to the $z = 9.6$ – 8.7 redshift bin used by both Patil et al. (2017) and Mertens et al. (2020). By decreasing the bandwidth, a third-order polynomial provides enough freedom to capture the shape of the bandpass. Without any regularization, the number of gain parameters is reduced by a factor of 5 due to the limited bandwidth, but so is of course the number of input visibilities. Full regularization, forcing the solutions to lie on the Bernstein polynomial, would reduce the number of degrees of freedom by a factor of 20, since there are 62 subbands within this 12 MHz that have a individual gain solution. Since in our experiments full convergence is not always met, the number of degrees of freedom will only be partly be reduced by regularization. A more formal method to derive the remaining number of degrees of freedom is given in Yatawatta (2019).

For a few intervals of real data, we investigate the effect of the number of iterations and the regularization parameter on the smoothness of the solutions. From theory (Mouri Sardarabadi & Koopmans 2019), one expects that the solutions will be constrained to a smooth curve once convergence is reached. Selecting a regularization parameter that is too low may increase the number of iterations needed to reach convergence, whereas a high regularization parameter may force the solutions to a wrong minimum every iteration, thereby slowing convergence. We find that with the default SAGECAL-CO setting for the regularization parameter ρ , with values between 1 and 50, depending on the apparent cluster brightness, convergence to smooth gain solutions is not reached even after 400 iterations. Because the number of iterations is in practice constrained by compute power limitations, we decided to use 40 iterations and find the optimal regularization parameters for that setting.

A good measure of the smoothness of the gain solutions can be obtained by examining the differences of the gains of adjacent subbands. In Fig. 4, we present the variance of the subband differenced gain solutions for a typical cluster as a function of regularization parameter. To obtain this plot, we used real data. A clear minimum in differential gain variance is obtained near the regularization value of $\rho = 200$, which is 200 times higher than the original regularization parameter for this specific cluster. Similar optimal values were found

in Yatawatta (2016). In the same way, we find that for most of the 122 directions, the optimal regularization parameter is tens to hundreds times higher than the initial regularization. Only a few clusters that are close to the first null of the beam do not show such a clear minimum of differential gain variance. A possible explanation could be that a beam model is not applied to the sky model during DD calibration. This results in stronger frequency-dependent beam variations near the first null, causing spatial gain variations over the extent of the cluster not constant enough to be described by a single Jones matrix. Including the beam model is one of the envisioned improvements to our current calibration. The signal-to-noise ratio (S/N) of the gains using real data, defined as the absolute gain divided by the subband differenced gain variance, using the optimized regularization parameters, is shown in Fig. 5. Every point in this figure is a single component of the 28 000 sky model components. Clearly, clusters that share the same DD gain solution show the same S/N-ratio. The S/N ratio of the clusters shows a slight gradient to lower S/N values away from the phase centre, possibly related to the lower apparent flux at those positions. Some outliers with lower S/N values are visible near the first null of the beam.

5.1 Simulations

We repeat the tests in Section 4 on the same simulated data. This time, we only use a single realization of the simulated signal, corresponding to the brightest selected signal with a relative amplitude of 10. Additionally, we use the limited bandwidth and optimized regularization parameters, as described in the previous section. As before, we analyse the Stokes-V correlations to understand the effect of calibration on an almost noise-like signal. The Stokes-V power ratio plots with high regularization without and with the baseline cut are shown in the left- and right-hand panels of Fig. 6, respectively. Table 1 summarizes the results of the suppression and solver noise measurements for four different calibration schemes. The schemes have different baseline cuts and levels of regularization. The lowest stokes V power spectrum ratio is achieved when no baseline cut is applied, as expected. However, even with a high- gain regularization level, a considerable level of signal suppression is observed (35 per cent suppression), although it is much lower than without strong regularization. This remaining signal suppression is likely due to the fact that convergence is too slow for practical purposes, therefore maximum smoothness is not reached within the limited number of 40 iterations. This is in agreement with the theoretical considerations by Mouri Sardarabadi & Koopmans (2019). In Fig. 3, we show the amplitude solutions for all four calibration scenarios of a single time slot, station and direction. The reduction of spectral variations with increased regularization is visible and even more pronounced in the time-averaged delay power spectrum of the complex gains of the same station and direction in the same figure. From these results, we conclude that combining a baseline cut with a high level of regularization to enforce smooth gain solutions is optimal. When using these settings, the signal suppression on the baselines excluded from calibration is close to zero, as expected. Moreover, the extra noise power on the shorter baselines due to overfitting is reduced by about a factor four with respect to the test using low regularization. Also, the sharp break that can be observed in Fig. 2 at the k_{\perp} value corresponding to a baseline length of 250λ is hardly visible in the right-hand panel of Fig. 6. We conclude that this way an optimal bias-variance trade off can be reached.

Note that the way our simulated gains were constructed, namely from real data calibration, the input gains were not spectrally smooth. This is shown in Fig. 3, where the input gains in this example show

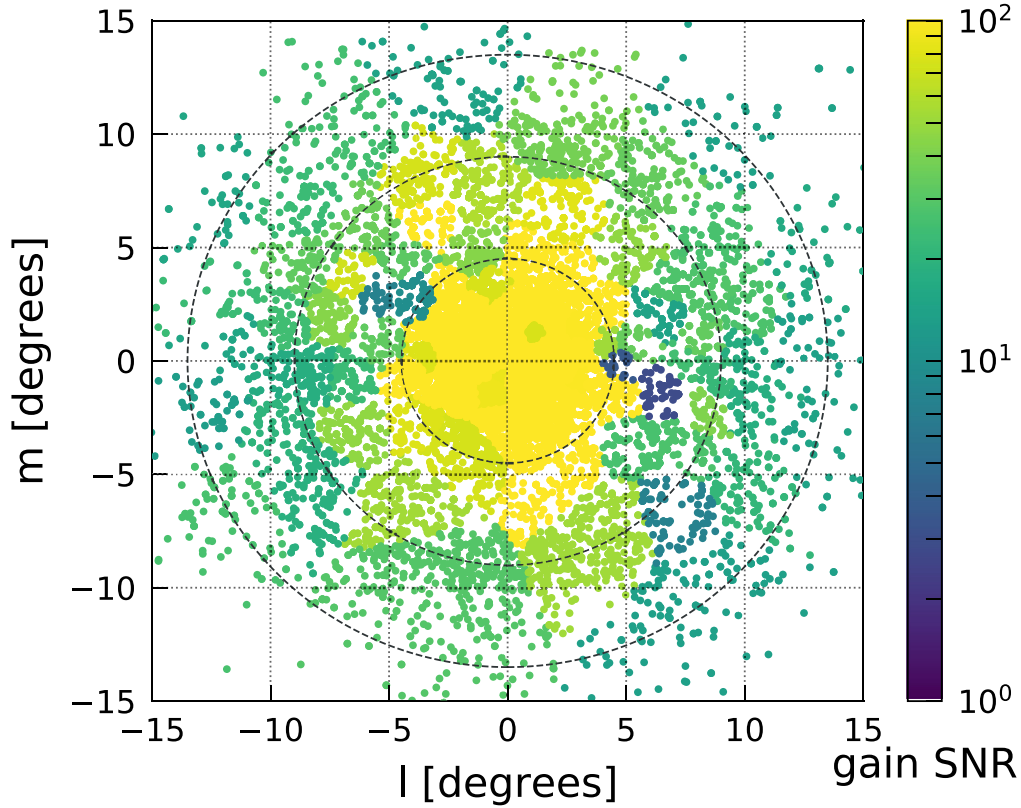


Figure 5. *S/N*-ratio of the subband differenced solutions for all $\sim 28,000$ sky components in real data. Clusters of components can be recognized since they share the same gain solutions. In total, there are 120 clusters in the image. The *S/N* ratio is similar for most clusters, but showing a slight gradient to lower *S/N* values away from the phase centre. Some outliers near the first null of the beam can be observed.




Figure 6. Ratio of the Stokes *V* power spectra after DD calibration with strong regularization over the input noise. Left-hand panel: No baseline cut during calibration. Right-hand panel: Excluding baselines shorter than 250λ from calibration.

similar structure to the gains of the initial test with low regularization. True spectral structure in the gains will introduce additional noise when enforcing maximal smoothness, an effect that needs to be carefully considered when working with real data. It is, however, expected that the true gains are spectrally smooth scales up to several MHz in real data. In our analysis, we are not able to directly detect this additional noise. In principle, it has been added to the power spectra in Fig. 6, such that our conclusion about the reduction of the overfitting

effect using strong regularization might be an underestimation. In fact, our result strengthens the conclusion that limiting the number of degrees of freedom in calibration is more important to reduce signal suppression and noise enhancement, than reaching the *True* gains.

6 DISCUSSION

We studied the effect of using DD calibration with a large number of free parameters (i.e. directions and spectral channels) on the

noise power as well as on a simulated 21-cm signal added to a typical LOFAR EoR data set pointing at the NCP. Although some parameters, such as the model components and regularization per direction were optimized for the chosen field and bandwidth, the main conclusion of our analysis is valid for any LOFAR EoR observation. We investigated the effects of excluding the shorter baselines that are used in the final power spectrum analysis ($<250\lambda$) from the calibration as well as the level of regularization of calibration solutions. The regularization constrains the gain solutions to be more smooth in frequency. We find that including all the baselines during calibration, even when using a model of the bright diffuse emission on those baselines, leads to an unacceptable level of signal suppression. Calibrating without those baselines that will be used in the final 21-cm signal power spectrum removes signal suppression. However, excluding short baselines from the calibration also increases the excess noise power on those same baselines. This is known as the bias-variance trade off. Forcing the solutions to be spectrally smooth, on the other hand, reduces signal suppression significantly, and limits the excess noise. The origin of signal suppression and noise enhancement due to calibration can be the lack of convergence to the *True* gains, the large number of degrees of freedom and sky model incompleteness. For the latter, the current analysis gives no definite answer, since most of the simulations were done with a perfect sky model, but it needs to be concerned in real data analysis. We found that convergence to smooth parameters could not be reached in practice on real data, even with a very large number of iterations. The simulated gains in our analysis contained spectral structure, but constraining the solutions to be spectrally smooth did not enhance the noise in our limited data set, strengthening the conclusion that the biggest impact comes from limiting the number of free parameters. Reducing the number of degrees of freedom by enforcing strong regularization clearly leads to lower signal loss and less noise enhancement.

We conclude that, for DD calibration of a typical LOFAR 21-cm data set, it is necessary to use a baseline cut, separating the baselines used for calibration from those used for signal extraction, as well as significantly increase the level of regularization. We also show that the conservative strategy chosen for the LOFAR EoR analysis by Mertens et al. (2020) does not lead to signal suppression during the calibration step, while the increased noise level due to overfitting is significantly reduced with respect to the strategy presented earlier by Patil et al. (2017). An incomplete sky model may further increase the level of noise due to overfitting or the level of signal suppression if no baseline cut is applied. In the presented study a perfect sky model (apart from the 21-cm signal) was assumed, therefore a lower limit on the level of suppression was established. Further testing is necessary to determine whether the baseline cut can be safely removed when using an improved model of the sky. However, as shown by the results in this work, the use of an optimized regularization parameter reduces the solver noise to a lower level (about a factor of 4 in power) and removing the cut is therefore less pressing.

Besides reducing the parameter space by enforcing smooth solutions in frequency space, we have additionally shown that, in real data, the gain parameters never converge to a polynomial curve for some of the clusters. This is specifically the case for clusters near the first null of the beam. Including a beam model during DD calibration could therefore improve results further. Since DD effects are expected to be smooth over the field of view, another way to decrease the degrees of freedom in the calibration is to constrain the solutions to be both spectrally and directionally smooth.

The final calibration scheme described here also does not yet include constrained solutions for the initial two-directional calibration.

The reason for this is that cable reflections and bandpass effects lead to DI gain variations on short frequency scales. These can not be taken into account with a third-order Bernstein polynomial. However, as was discussed by Barry et al. (2016) and Ewall-Wice et al. (2017), correcting the data with spurious rapid varying gain solutions, e.g. due to an incomplete sky model, could introduce artificial enhanced power on small frequency scales that cannot be reduced in the subsequent steps of the calibration. Contrary to solver noise, this model incompleteness induced power leakage is correlated and therefore, can not be integrated down by adding observations. Although in DD calibration the final visibilities are not multiplied with the calibration gains, the gain multiplied sky model is subtracted. Therefore, the correlated enhanced power at high k values, arising from incorrect spectral structure in the calibration gains related to model incompleteness, is expected to be less pronounced but could still be present if the DD gains are not forced to be spectrally smooth. In Mertens et al. (2020), it was shown that with the current settings, a large fraction of excess noise power (above the theoretical noise power) remains, which needs further investigation (Gan et al., in preparation). In future analyses, we will investigate whether adding extra constraints during our initial calibration or during the DD calibration will further reduce the excess noise power and is able to limit the leakage of power due to model incompleteness.

We note that, the results in this paper focus on the specific case of data from the LOFAR EoR project. However, the conclusions are also applicable to other experiments that try to recover faint signals using various forms of calibration with a large number of free parameters, and to the SKA, that has a design similar to that of LOFAR.

ACKNOWLEDGEMENTS

This result is part of a project that receives funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 884760). EC acknowledges the support of a Royal Society Dorothy Hodgkin Fellowship and a Royal Society Enhancement Award. ITI was supported by the Science and Technology Facilities Council [grant numbers ST/I000976/1 and ST/T000473/1] and the Southeast Physics Network (SEP-Net). AG would like to acknowledge IUCAA, Pune for providing support through the associateship programme.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Barry N., Hazelton B., Sullivan I., Morales M. F., Pober J. C., 2016, *MNRAS*, 461, 3135
- DeBoer D. R. et al., 2017, *Publ. Astron. Soc. Pac.*, 129, 045001
- Ewall-Wice A., Dillon J. S., Liu A., Hewitt J., 2017, *MNRAS*, 470, 1849
- Hamaker J. P., Bregman J. D., Sault R. J., 1996, *A&AS*, 117, 137
- Jelić V. et al., 2008, *MNRAS*, 389, 1319
- Mertens F. G. et al., 2020, *MNRAS*, 493, 1662
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Morales M. F., Hewitt J., 2004, *ApJ*, 615, 7
- Morales M. F., Wyithe J. S. B., 2010, *ARA&A*, 48, 127
- Mouri Sardarabadi A., Koopmans L. V. E., 2019, *MNRAS*, 483, 5480
- Offringa A. R. et al., 2014, *MNRAS*, 444, 606

- Offringa A. R., de Bruyn A. G., Biehl M., Zaroubi S., Bernardi G., Pandey V. N., 2010, *MNRAS*, 405, 155
- Offringa A. R., van de Gronde J. J., Roerdink J. B. T. M., 2012, *A&A*, 539, A95
- Offringa A. R., Mertens F., Van der Tol S., Veenboer B., Gehlot B. K., Koopmans L. V. E., Mevius M., 2019, *A&A*, 631, A12
- Paciga G. et al., 2013, *MNRAS*, 433, 639
- Pandey V., Koopmans L., Tiesinga E., Albers W., Koers H., 2020, in Pizzo R., Deul E., Mol J., Plaa H., Verkouter H., Williams R., eds, ASP Conf. Ser. Vol. 524, ADASS XXIV. Astron. Soc. Pac., San Francisco, p. 473
- Patil A. H. et al., 2016, *MNRAS*, 463, 4317
- Patil A. H. et al., 2017, *ApJ*, 838, 65
- Smirnov O. M., 2011, *A&A*, 527, A106
- van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Wijnholds S. J., Grobler T. L., Smirnov O. M., 2016, *MNRAS*, 457, 2331
- Yatawatta S. et al., 2013, *A&A*, 550, A136
- Yatawatta S., 2015, *MNRAS*, 449, 4506
- Yatawatta S., 2016, 24th European Signal Processing Conference (EU-SIPCO). IEEE, Budapest, Hungary, p. 265
- Yatawatta S., 2019, *MNRAS*, 486, 5646

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.