



HAL
open science

Spectral modelling of Cygnus A between 110 and 250 MHz. Impact on the LOFAR 21-cm signal power spectrum

E Ceccotti, A.R Offringa, L.V.E Koopmans, F.G Mertens, M Mevius, A Acharya, S.A Brackenhoff, B Ciardi, B.K Gehlot, R Ghara, et al.

► To cite this version:

E Ceccotti, A.R Offringa, L.V.E Koopmans, F.G Mertens, M Mevius, et al.. Spectral modelling of Cygnus A between 110 and 250 MHz. Impact on the LOFAR 21-cm signal power spectrum. *Astronomy & Astrophysics* - A&A, 2025, 696, pp.A56. <10.1051/0004-6361/202453106>. <hal-04989439>

HAL Id: hal-04989439

<https://hal.science/hal-04989439v1>

Submitted on 28 Apr 2025

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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Spectral modelling of Cygnus A between 110 and 250 MHz

Impact on the LOFAR 21-cm signal power spectrum

E. Ceccotti^{1,2,*}, A. R. Offringa^{3,1}, L. V. E. Koopmans¹, F. G. Mertens⁴, M. Mevius^{3,1}, A. Acharya⁵, S. A. Brackenhoff¹, B. Ciardi⁵, B. K. Gehlot¹, R. Ghara⁶, J. K. Chege¹, S. Ghosh¹, C. Höfer¹, I. Hothi^{4,7}, I. T. Iliev⁸, J. P. McKean^{1,9,10}, S. Munshi¹, and S. Zaroubi^{11,1,5}

¹ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands

² INAF – Istituto di Radioastronomia, Via P. Gobetti 101, 40129 Bologna, Italy

³ ASTRON, PO Box 2, 7990 AA Dwingeloo, The Netherlands

⁴ LERMA, Observatoire de Paris, PSL Research University, CNRS, Sorbonne Université, F-75014 Paris, France

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

⁶ Department of Physical Sciences, Indian Institute of Science Education and Research Kolkata, Mohanpur, WB 741246, India

⁷ Laboratoire de Physique de l'ENS, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, 75005 Paris, France

⁸ Astronomy Centre, Department of Physics & Astronomy, Pevensey III Building, University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom

⁹ South African Radio Astronomy Observatory (SARAO), PO Box 443, Krugersdorp 1740, South Africa

¹⁰ Department of Physics, University of Pretoria, Lynnwood Road, Hatfield, Pretoria 0083, South Africa

¹¹ Department of Natural Sciences, The Open University of Israel, 1 University Road, PO Box 808, Ra'anana 4353701, Israel

Received 21 November 2024 / Accepted 24 February 2025

ABSTRACT

Studying the redshifted 21-cm signal from the neutral hydrogen during the Epoch of Reionisation and Cosmic Dawn is fundamental for understanding the physics of the early universe. One of the challenges that 21-cm experiments face is the contamination by bright foreground sources, such as Cygnus A, for which accurate spatial and spectral models are needed to minimise the residual contamination after their removal. In this work, we develop a new, high-resolution model of Cygnus A using Low Frequency Array (LOFAR) observations in the 110–250 MHz range, improving upon previous models by incorporating physical spectral information through the forced-spectrum method during multi-frequency deconvolution. This approach addresses the limitations of earlier models by providing a more accurate representation of the complex structure and spectral behaviour of Cygnus A, including the spectral turnover in its brightest hotspots. The impact of this new model on the LOFAR 21-cm signal power spectrum is assessed by comparing it with both simulated and observed North Celestial Pole datasets. Significant improvements are observed in the cylindrical power spectrum along the Cygnus A direction, highlighting the importance of having spectrally accurate models of the brightest foreground sources. However, this improvement is washed out in the spherical power spectrum, where we measure differences of a few hundred mK at $k < 0.63 h \text{ cMpc}^{-1}$, but not statistically significant. The results suggest that other systematic effects must be mitigated before a substantial impact on 21-cm power spectrum can be achieved.

Key words. methods: data analysis – techniques: interferometric – cosmology: observations – dark ages, reionization, first stars – radio continuum: galaxies

1. Introduction

If we look at the sky at 150 MHz, we expect to find more than one source per deg^2 with a flux density lower than 1 Jy (Wilman et al. 2008; Intema et al. 2017). Radio sources become increasingly rare as their brightness increases (e.g. Franzen et al. 2019), with only a few having flux densities higher than 100 Jy. The brightest radio sources in constellations are named after their constellation with an ‘A’ suffix, forming the so-called ‘A-team’. The most powerful A-team source at 150 MHz is currently Cygnus A (Cyg A; 3C 405, RA 19^h59^m28^s.36, Dec +40°44′02″.10 in J2000; Baumgartner et al. 2013; Charlot et al. 2020; Vinyaikin 2014), an FR-II radio galaxy with a flux density of about 11 kJy at 150 MHz, that, together with Cassiopeia A (Cas A), Taurus A (Tau A), and Virgo A (Vir A), forms the northern A-team (e.g. de Gasperin et al. 2020).

The high brightness of the A-team sources is a double-edged sword for radio interferometric observations. On the one hand, these sources are so easy to measure even for low sensitivity instruments that their fluxes are well known, making them good (amplitude) calibrators, especially when unresolved. On the other hand, they can affect observations of a given target field even when they are tens of degrees away. This poses a challenge for wide-field instruments, such as LOFAR¹ (van Haarlem et al. 2013) and MWA² (Lonsdale et al. 2009), which have strong primary beam sidelobes. Different techniques have been developed to reduce such a contribution in the targetted observations. Examples are fast flagging of the data where the predicted contamination is above a certain threshold (Shimwell et al. 2017),

¹ Low-Frequency Array, <https://www.astron.nl/telescopes/lofar>

² Murchison Widefield Array, <http://www.mwatelescope.org>

* Corresponding author; emilio.ceccotti@inaf.it

more computationally demanding direction-dependent (DD) calibration (Gan et al. 2023), demixing (Van der Tol 2009) when the A-team source is more than 30° away, or peeling (Noordam 2004; Mitchell et al. 2008) when it is closer. In all of these cases, good spatial and spectral models of the A-team sources are required to minimise the residual contamination after the removal.

Minimising the A-team contamination is particularly important for those experiments that aim to probe the neutral hydrogen (HI) during the Epoch of Reionisation (EoR) and Cosmic Dawn (CD) via the redshifted 21-cm line (see e.g. Furlanetto et al. 2006; Pritchard & Loeb 2012; Liu & Shaw 2020, for some reviews on 21-cm cosmology). Because the current generation of instruments does not have the sensitivity to make direct images of the 21-cm signal, they mainly focus on estimating the spatial fluctuations of the 21-cm signal by measuring its power spectrum. Increasingly more stringent upper limits have been set on the 21-cm power spectrum from the EoR by radio interferometers such as LOFAR (Mertens et al. 2020) and MWA (Trott et al. 2020; Kolopanis et al. 2023) with the highest frequency band, GMRT (Paciga et al. 2013), PAPER³ (Parsons et al. 2010; Kolopanis et al. 2019), and HERA⁴ (DeBoer et al. 2017; Abdurashidova et al. 2022), and from the CD by instruments such as LOFAR (Gehlot et al. 2019) and MWA (Yoshiura et al. 2021) with the lowest frequency band, OVRO-LWA⁵ (Greenhill & Bernardi 2012; Eastwood et al. 2019), NenuFAR⁶ (Zarka et al. 2012; Munshi et al. 2024), and AARTFAAC⁷ (Gehlot et al. 2020). The next generation of instruments, such as HERA and SKA-Low⁸ (Dewdney et al. 2009; Koopmans et al. 2015) are expected to have the sensitivity to detect the cosmological signal via power spectra, and SKA-Low has the potential to directly make tomographic images (e.g. Giri et al. 2018a,b; Bianco et al. 2024).

In all of the aforementioned cases, the extra-galactic and Galactic foreground emission must be mitigated, being up to five orders of magnitude brighter than the cosmological 21-cm signal (e.g. Jelić et al. 2008; Bowman et al. 2009; Bernardi et al. 2009, 2010). In principle, a separation between the two signals can be made based on the spectral smoothness. The foreground emission, being dominated by synchrotron radiation, is expected to be smooth over tens of MHz, while the 21-cm signal, being a hyperfine spectral line, fluctuates rapidly over MHz-scales (Shaver et al. 1999; Morales & Hewitt 2004; Santos et al. 2005; Morales & Wyithe 2010). However, instrumental systematics and incorrect calibration can cause spatial and spectral fluctuations in the observed data and mix the signals, making it harder to separate them (e.g. Datta et al. 2010; Vedantham et al. 2012; Barry et al. 2016; Ewall-Wice et al. 2017; Mazumder et al. 2022; Gorce et al. 2023). For this reason, having a precise model of the sky emission, with correct spectral information, is vital to minimise such errors.

Foreground avoidance methods, such as the ones used by the HERA collaboration, are less dependent on sky modelling, because they mitigate the foreground by using the foreground avoidance technique, which discards the Fourier modes dominated by the foregrounds (e.g. Parsons et al. 2012; Thyagarajan et al. 2015). However, they might still need accurate sky models to improve the calibration (e.g. Li et al. 2018; Kern et al. 2020) or to access the most sensitive small Fourier modes (e.g. Kerrigan et al. 2018; Ghosh et al. 2020). On the other hand, sky-modelling subtraction methods, such as the ones used by the LOFAR team, require an accurate sky model (Mertens et al. 2020; Gan et al. 2022). The latest upper limits published by the LOFAR-EoR Key Science Project (KSP) showed an excess variance that, on the larger angular scales, reaches 22 times the thermal noise in the residual power of the North Celestial Pole (NCP) field after foreground removal (Mertens et al. 2020). There might be multiple causes for this excess power, such as polarisation leakage (Jelić et al. 2015; Asad et al. 2015, 2016), model incompleteness (Barry et al. 2016; Patil et al. 2016; Ewall-Wice et al. 2017), low-level radio frequency interference (RFI; Wilensky et al. 2019; Offringa et al. 2019a), and uncorrected ionospheric effects (Vedantham & Koopmans 2016). However, Gan et al. (2022) concluded that the excess seen in LOFAR data is likely dominated by residuals of bright sources outside the main lobe of the primary beam. The sources that mainly affect the NCP field are Cas A and Cyg A, which are approximately 30° and 50° away from the field centre, respectively. At that distance, the chromatic effects of the primary beam are stronger and can introduce rapid spectral fluctuations to the sources, mixing their emission with other foregrounds and the 21-cm signal (Pofer et al. 2016). If the beam sidelobes have large spatial and spectral gradients and the source is extended enough, the beam could also imprint different spectral structures to different parts of such a source, making it even harder to calibrate for it (Cook et al. 2022).

Previous efforts to model bright and extended galaxies have been made, for example, by Procopio et al. (2017) for MWA observations. They demonstrated that when such sources are incorrectly modelled as point sources, their subtraction leaves strong residuals. As a result, an improved wide-field model was made by Lynch et al. (2021) for use in the MWA processing pipeline. Efforts to improve models of bright sources were also made by Line et al. (2020) on the Fornax A radio galaxy. In this case, shapelets (Refregier 2003) were employed to model the extended emission. Shapelets allow a source to be modelled with fewer components than when using a collection of point sources and 2D Gaussian functions, thereby reducing the computational time required for calibration. However, incorporating spatially varying spectral information into shapelets remains challenging and has not yet been tested in real scenarios. Line et al. (2020) showed that, while a theoretical improvement in the 21-cm power spectrum was achieved with the shapelet model, the improvement was negligible in real observations due to other systematics. Similar results were observed by Gehlot et al. (2022) when modelling diffuse Galactic emission in AARTFAAC observations. Shapelets have also been applied to LOFAR observations by Yatawatta et al. (2013), who made a model of 3C 61.1, the brightest source in the NCP field. This model was subsequently used by Mertens et al. (2020). However, the fixed spectrum of the shapelet components proved to be a limitation, and Mertens et al. (2025) showed that using a multi-scale model of 3C 61.1 with a spatially varying spectral index led to improved upper limits. McKinley et al. (2022) showed that combining point sources and Gaussian components can still produce

³ Precision Array for Probing the Epoch of Reionization, <http://eor.berkeley.edu>

⁴ Hydrogen Epoch of Reionization Array, <https://reionization.org/>

⁵ Owens Valley Radio Observatory Long Wavelength Array, <https://leo.phys.unm.edu/~lwa>

⁶ New Extension in Nançay Upgrading LOFAR, <https://nenufar.obs-nancay.fr>

⁷ LOFAR Amsterdam-ASTRON Radio Transients Facility and Analysis Centre, <http://aartfaac.org>

⁸ Square Kilometre Array, <https://www.skao.int/en/explore/telescopes>

an accurate, high-resolution model of a large source such as Centaurus A. All these efforts suggest that having a spatially and spectrally realistic model of A-team sources is a fundamental requirement to further investigate the cause of the excess power.

In this paper, we will focus on Cyg A, which is the brightest A-team source at 150 MHz. It is a complex source, with both extended and compact emission. At low-frequencies, its spectral energy distribution deviates from a pure power law. Using several measurements between 12 MHz and 15 GHz, [McKean et al. \(2016\)](#); henceforth cited as MK16) showed that the total flux density S at frequency ν is well described by a third-order ($n = 4$) logarithmic polynomial function of the form

$$S(\nu) = S(\nu_0) \prod_{i=1}^{n-1} 10^{c_i \log_{10}^i(\nu/\nu_0)}, \quad (1)$$

where c_1 corresponds to the spectral index α and c_2 to the spectral curvature β (see [MK16](#) for the coefficients values). This total spectral shape does not describe the spatial variation of the Cyg A emission, whose spectral indices go from $\alpha \approx -0.65$ in the lobes to $\alpha \approx -1.25$ in the plume region, with a rapid turnover in the two brightest hotspots, the western of which is labelled A and the eastern D ([Leahy et al. 1989](#); [MK16](#)). In particular, the spectral index of hotspot A, measured at 150 MHz, is $\alpha_A = +0.18 \pm 0.01$, with the turnover around 140 MHz, while for hotspot D it is $\alpha_D = +0.36 \pm 0.02$, with the turnover at 160 MHz ([MK16](#)).

The main goal of this work is to create an accurate low-frequency spatial and spectral model of Cyg A and assess its impact on the LOFAR 21-cm power spectrum. Both [Mertens et al. \(2020\)](#) and [Gan et al. \(2022\)](#) used models of Cyg A with a spatially fixed spectral index $\alpha = -0.8$ at a resolution that is not enough to capture the finest structures of the radio galaxy. By using a better spatial and spectral model of Cyg A, we can investigate whether an improved model of a bright off-axis source reduces the observed excess power in the 21-cm signal power spectra. In order to embed physical spectral information into the Cyg A model, we take advantage of the forced-spectrum method during a multi-frequency deconvolution [Ceccotti et al. \(2023\)](#). This method uses an input spectral index map to constrain the spectral index of each component by fitting a logarithmic polynomial function during the deconvolution. Without this regularisation of the spectral behaviour, logarithmic (or ordinary) polynomial fitting can generate unrealistic spectral indices, especially when the frequency bandwidth is limited or systematics in data are high ([Offringa et al. 2016](#)). In this work, we used an updated version of the method that supports higher order terms, such as the curvature, as long as maps of the higher order terms are also provided. The spectrally varying model of Cyg A that we get from the forced-spectrum method is the first of its kind in the analysis of the LOFAR power spectrum⁹. Another important goal of this work is to provide a strategy to make a model with physical spectral information by using the forced-spectrum method, thereby making it easier to make accurate source models. The final aim is given by combining the two goals: assessing the importance of improved models of off-axis sources, including spectral information, and the role they play in the excess power that is observed in the 21-cm power spectra.

In Sect. 2, we describe how we obtained the new Cyg A model, and compare it with the model used in the current

LOFAR-EoR processing pipeline. In Sect. 3, we introduce both the simulated and observed NCP datasets that we use to test the impact of the new model onto the 21-cm upper limits. In Sect. 4, we describe the NCP data processing, from calibration to power spectrum estimation. The results on the two datasets are presented in Sect. 5. Finally, in Sect. 6, we summarise and discuss the results, drawing the conclusions.

2. Cygnus A modelling with LOFAR HBA

The current LOFAR-EoR processing pipeline uses the Cyg A model obtained from low-frequency VLA observations¹⁰ (see Appendix A for the catalogue of the model components), which we will refer to as the ‘old’ Cyg A model. This model has been used by standard LOFAR processing pipelines, such as LINC¹¹ (LOFAR Initial Calibration pipeline; [de Gasperin et al. 2019](#)), for calibration or demixing. Here we identify three major issues of this model when applied to HBA data:

1. *Spectral extrapolation*: the flux density of each component has been derived from a clean-based model of VLA observations at 73.8 MHz, and it is then extrapolated to higher frequencies using a given spectral index α ;
2. *Low resolution*: the model consists of four point and eight Gaussian components, five of which have a zero-length minor axis, that are unable to model the complex structures of the source visible at higher frequencies (see Sect. 2.3.4);
3. *Spatially constant spectral index*: each component has an artificial, constant spectral index $\alpha = -0.8$, even though the source has a strongly varying spectral index.

The total flux density given by the old Cyg A model at 73.8 MHz is 15.9 kJy, while the results of [MK16](#) indicate it is 16.3 ± 2.2 kJy when extrapolated to that frequency using Eq. (1). Whereas this offset is acceptable within the 1σ error, issues (1) and (3) lead to a much larger difference at higher frequencies. Extrapolating the flux density of the old model to 150 MHz by using a first-order ($n = 2$) logarithmic polynomial with $\alpha = -0.8$ results in a total flux of 9.0 kJy, which is 16% lower than the expected 10.7 kJy ([MK16](#)). The inconsistent flux density scale is another reason for improving the model of Cyg A. We will describe the new model in the next section.

2.1. The new Cygnus A model

LOFAR has the capabilities to operate between 110 and 250 MHz, by selecting two sub-ranges of the high band antenna (HBA) system, HBA-low in the 110–190 MHz range and HBA-high in the 210–250 MHz range. For each observation, a total of 488 sub-bands (SB) can be recorded, each with 195.3 kHz of bandwidth, which are further divided into 64 frequency channels. LOFAR currently consists of 24 core stations (CS) near the centre of the array, distributed within a diameter of ≈ 4 km and 14 remote stations (RS) across the Netherlands, with a maximum baseline of ≈ 120 km, forming what is called the Dutch array. Another 14 (international) stations are spread across Europe, providing baselines up to ≈ 2000 km. Each CS includes two HBA sub-stations (HBA0 and HBA1) that can either function independently (‘HBA Dual’ mode), improving the uv -coverage, or be individually selected (‘HBA Zero/One’ mode).

Given the large frequency coverage of the HBA system and the high spatial resolution provided by the Dutch array

⁹ Main field sources have different spectral indices, but they are only described by single point-like components, so there is no spectral variation within a single source.

¹⁰ <https://github.com/lofar-astron/prefactor/tree/master/skymodels>

¹¹ <https://linc.readthedocs.io/en/latest/>

Table 1. Observational details of the Cyg A dataset after flagging, averaging, and calibration.

Parameter	Value
Telescope	LOFAR HBA
Project code	LC1_013
Antenna configuration	HBA One
Number of stations	36 (CS + RS) ^(a)
Phase centre (J2000):	
Right Ascension	19 ^h 59 ^m 28 ^s 00
Declination	+40°44′02″00
Obs. start time (UTC):	
HBA-low	2014 May 13; 00:11
HBA-high	2014 May 14; 00:00
Frequency range:	
HBA-low	110.7–181.4 MHz
HBA-high	210.0–248.6 MHz
Duration of observation	12 h
Time resolution	6.0 s
Frequency resolution	195.3 kHz

Notes. ^(a)CS006 and RS208 did not participate in this observation.

(approximately 3 arcsec), our goals with the improvement of the model for Cyg A are:

1. *Obtain an accurate model at 120–160 MHz:* this is the frequency range targeted by the LOFAR-EoR KSP and, therefore, we want to use HBA observations;
2. *Increase resolution to 5 arcsec:* hotspots are compact sources and we need a sufficient high resolution to not mix their emission with the surrounding lobes: 5 arcsec is the highest resolution provided by LOFAR HBA to get a model that is accurate for all Dutch baselines;
3. *Capture the spatially varying spectral index:* Cyg A is a morphologically complex radio galaxy, where different regions show very different spectral behaviours that have to be taken into account.

To satisfy goal (1), we use a combination of 12 h LOFAR HBA-low and HBA-high observations taken in 2014. The data, previously processed and calibrated (see Appendix B), consist of 412 SBs spanning 111–181 and 210–249 MHz, with a frequency resolution of 195.3 kHz (i.e. one channel per SB) and an integration time of 6 s. Other observational details are reported in Table 1. All modelling and data processing described in this paper has been performed on the ‘Dawn’ high-performance computing cluster (Pandey et al. 2020) at the University of Groningen.

In the next sections, we describe how we achieved goals (2) and (3) to get a new, more complete model of Cyg A. We used the WSCLEAN imager software (Offringa et al. 2014) to make images of the data. We start by describing the deconvolution methods that we employed.

2.2. MS-MF deconvolution and forced-spectrum fitting

Modelling the spectral structures of a radio source requires fitting some smooth function to the data from the frequency channels. This can be done during a multi-frequency (MF) clean approach, which models the spectral behaviour during deconvolution (Sault & Wieringa 1994). In the joined-channel deconvolution described by Offringa & Smirnov (2017) and implemented in WSCLEAN, a wide frequency band is divided into narrower output channels which are jointly imaged and decon-

volved, and combined into a high dynamic range, frequency-integrated continuum image. The integrated image is used by the cleaning process to locate the peak of each model component. The frequency-varying brightness is measured across the images at each output frequency channel. This MF approach allows the components to capture the spectral variations of the source. Fitting the measured spectrum of each found component to an ordinary or logarithmic polynomial function ensures spectral smoothness of the model, which can be expected for sources dominated by synchrotron emission (e.g. Reich & Reich 1988; Tegmark et al. 2000; Jackson 2005; Jelić et al. 2008).

However, a clean-component spectral model might contain incorrect spectral information because, for instance, of the limited frequency bandwidth or when systematics in the data are large (e.g. Rau et al. 2016; Offringa et al. 2016). As a result, such spectral indices cannot be used for extrapolation. The forced-spectrum method (Ceccotti et al. 2023) solves this problem by using a spectral index map to assign to each clean-component a specific spectral index value based on the position of the component, fitting a logarithmic polynomial function of the form of Eq. (1). A modified weighted linear least-squares is then used to only determine the magnitude $S(\nu_0)$ of each clean-component. Ceccotti et al. (2023) used the forced-spectrum method to produce a model with only spectral indices. In this work, we additionally constrain higher order terms, such as the spectral curvature. This is necessary to model the spectral turnover of the Cyg A hotspots. When using higher order terms, it is important that the input spectral maps have been calculated at the same reference frequency at which WSCLEAN performs the MF spectral fitting. In Sect. 2.3 we will show a method to make spectral index and curvature maps from in-band data.

The forced-spectrum method, and in general the MF deconvolution, works well in combination with a multi-scale (MS) method (Cornwell 2008; Rich et al. 2008; Rau & Cornwell 2011). The MS deconvolution can incorporate diffuse emission by making a cleaned model as a summation of delta (i.e. point-like components) and basis functions (e.g. Gaussian or tapered quadratic). In this work, we used the MS deconvolution implemented in WSCLEAN (Offringa & Smirnov 2017) to model the extended emission with (circular) Gaussian functions that have an analytically defined Fourier transform, unlike tapered quadratic functions used by the algorithm described by Cornwell (2008). The combination of the MS and forced-spectrum methods result in a compactly described model, which is essential for calibration.

2.3. Imaging and spectral modelling of Cygnus A

To build the high-resolution spectral model of Cyg A, our imaging process was divided into four key steps: initial high-resolution imaging, imaging for accurate spectral modelling, spectral index extraction, and forced-spectrum fitting. This approach was designed to first capture the spatial details of the source, and then to accurately translate these observations into a comprehensive spectral model. The first step, the initial imaging, is optional and was done to check the data quality and the spectral features of hotspots A and D at the highest resolution allowed by the dataset.

2.3.1. High-resolution imaging and spectra of the hotspots

Our Cyg A dataset spans a wide bandwidth, from approximately 110 to 250 MHz, allowing us to create images with a synthesised beam (i.e. the point-spread function) of about 2 arcsec in

size at the highest frequency. Although we cannot take advantage of this for reliable spectral index extraction, where maintaining a consistent beam size across the bandwidth is essential to ensure each frequency samples the same spatial scale, we can utilise it to generate high-resolution images for data quality assessment. However, if we clean deep enough so that the residual is mostly free of emission from the source, we can set the restoring beam to a fixed value and still get good spectral information from the restored images. This is the approach used in this section to extract the spectral behaviour of the hotspots.

For this initial step, we used the MF joined-channel deconvolution, splitting the full bandwidth into 100 independent (i.e. without spectral fitting) output channels. We did not use MF weighting, ensuring a better image quality for each output frequency by gridding the uv -plane of each channel on a separate grid. The output channels were weighted with a Briggs weighting scheme with a robust parameter of -2 , resulting in an integrated synthesised beam with a full width at half-maximum (FWHM) of 3×2 arcsec. We used the MS deconvolution with point and Gaussian components of size 5, 10, 15, 20, 40, and 80 arcsec. We performed the cleaning to an initial threshold of 5σ using auto-masking and a final threshold of 1σ , where $\sigma = 8.3$ mJy/beam for the frequency integrated image, with a dynamic range of 1:19 540. We found that using a 1σ cleaning threshold, the residuals of the 100 output channels are noise-like at the hotspots location, but not completely on the rest of the source. This produced 10 times stronger residuals at the edges of the source at low frequencies and in the lobes at high frequencies. These are probably deconvolution artefacts caused by to the morphological complexity of the source. Because the condition of having the residual free of source emission is met for the brightest structures of Cyg A, we can still set a constant restoring beam to extract the spectra of hotspots A and D. We used a restoring circular beam with a FWHM of 3 arcsec, slightly larger than the frequency-integrated synthesised beam. Other imaging parameters are listed in the ‘High resolution’ column of Table 2.

The resulting 111–249-MHz frequency-integrated image is shown in the left panel of Fig. 1. At 3 arcsec resolution, the two brighter hotspots, A and D, are well resolved and we can extract their spectra from the 100 output images. We fitted their peak brightness using a non-linear least-squares method with a second-order ($n = 3$) logarithmic polynomial of the form

$$S(\nu) = S(\nu_0) \left(\frac{\nu}{\nu_0} \right)^\alpha 10^{\beta \log_{10}^2(\nu/\nu_0)}, \quad (2)$$

which is equivalent to Eq. (1) when using three terms, where α and β are, respectively, the spectral index and the curvature terms evaluated at the reference frequency ν_0 . We used 150 MHz as reference frequency to compare our spectral values with MK16. The panels on the right-hand side of Fig. 1 show these fits with a blue dashed line. Here, the black solid line represents the peak brightness as extracted from each output image, to which we have associated 1σ errors (grey shaded area). These are calculated by adding in quadrature the standard deviation of the noise of each image and a standard calibration error of 2% of the flux density (MK16), where the latter dominates the uncertainties. We find that the best fit values for the hotspot A are $\alpha_A = +0.1025 \pm 0.0002$ and $\beta_A = -1.541 \pm 0.009$, and for the hotspots D are $\alpha_D = +0.3081 \pm 0.0002$ and $\beta_D = -1.819 \pm 0.009$, confirming the presence of a turnover in both of them. Our values differ from those found by MK16 for a couple of reasons: they only had HBA-low data in the range 109–183 MHz, and they fitted a first-order logarithmic polynomial. Restricting our

Table 2. WSCLEAN parameters used for the initial high-resolution imaging and the forced-spectrum method, which produces the final model of Cyg A.

Parameter	High resolution	Forced spectrum
Image size	1024 × 1024	256 × 256
Pixel scale	0.25 arcsec	1 arcsec
Weighting	Briggs -2	Briggs -2
Gridder	W-gridder	W-gridder
Minor gain	0.01	0.01
Major gain	0.8	0.8
Spatial mask	–	1 Jy/beam [–]
Auto-mask ^(a)	5σ	3σ [6σ]
Auto-threshold ^(b)	1σ	1σ
MS scales (arcsec)	0–80	0–99 (default)
MS gain ^(c)	0.2	0.05
Output channels	100 (joined)	41 (joined)
Spectral fitting	–	Log. pol. $n = 3$ [–]
uv -range (k λ)	–	0.16–45.5
Beam size	3 arcsec	5 arcsec

Notes. Most of the parameters from the forced-spectrum method were also applied when generating the images used to extract the spectral index and curvature maps. The specific parameters for this imaging process are provided in square brackets. ^(a)Initial threshold at which a mask is constructed by finding clean-components, based on the noise standard deviation. ^(b)Stopping threshold inside the mask, based on the noise standard deviation. ^(c)Size of step made in the MS sub-minor loop.

analysis in the same frequency range and fitting a simple power law without the curvature term (red dot-dashed line in Fig. 1), we find $\alpha_A = +0.1720 \pm 0.0004$ and $\alpha_D = +0.3917 \pm 0.0004$, which are consistent with the values reported by MK16, namely, $\alpha_A = +0.18 \pm 0.01$ and $\alpha_D = +0.36 \pm 0.02$. The 1σ uncertainties on α and β assume uncorrelated errors, whereas systematics such as calibration and deconvolution errors are typically correlated, and are therefore not captured by these uncertainties.

2.3.2. Spectral index extraction

The second step towards constructing an accurate spectral model of Cyg A was to generate images that could be used for a pixel-by-pixel spectral index extraction. As mentioned in the previous section, using the 100 high-resolution images was not ideal due to the $\approx 10\sigma$ residual emission, which could affect the spectral indices of fainter emission. While the clean model would be convolved with a consistent restoring beam across the bandwidth, the residual images retain the synthesised beam of their respective output channels. This might result in a mix of components at different spatial scales from low to high frequencies, preventing the extraction of a reliable source spectrum. To address this, we performed a separate imaging step with appropriate settings for an accurate full-source spectral index extraction.

Because we could not be certain that the residual images were noise-like across the entire bandwidth, we ensured a consistent synthesised beam (i.e. the same uv -coverage) by applying a strict uv -cut. We selected only the (projected) baselines in the range 160–45 500 λ , where the lower value corresponds to the shortest baseline at 249 MHz and the upper value to the longest baseline at 111 MHz. By removing baselines shorter than 160 λ , we lost sensitivity to structures larger than ≈ 20 arcmin, which was not a concern since our focus was only on Cyg A, which has a maximum extent of 2.3 arcmin. Conversely, the upper

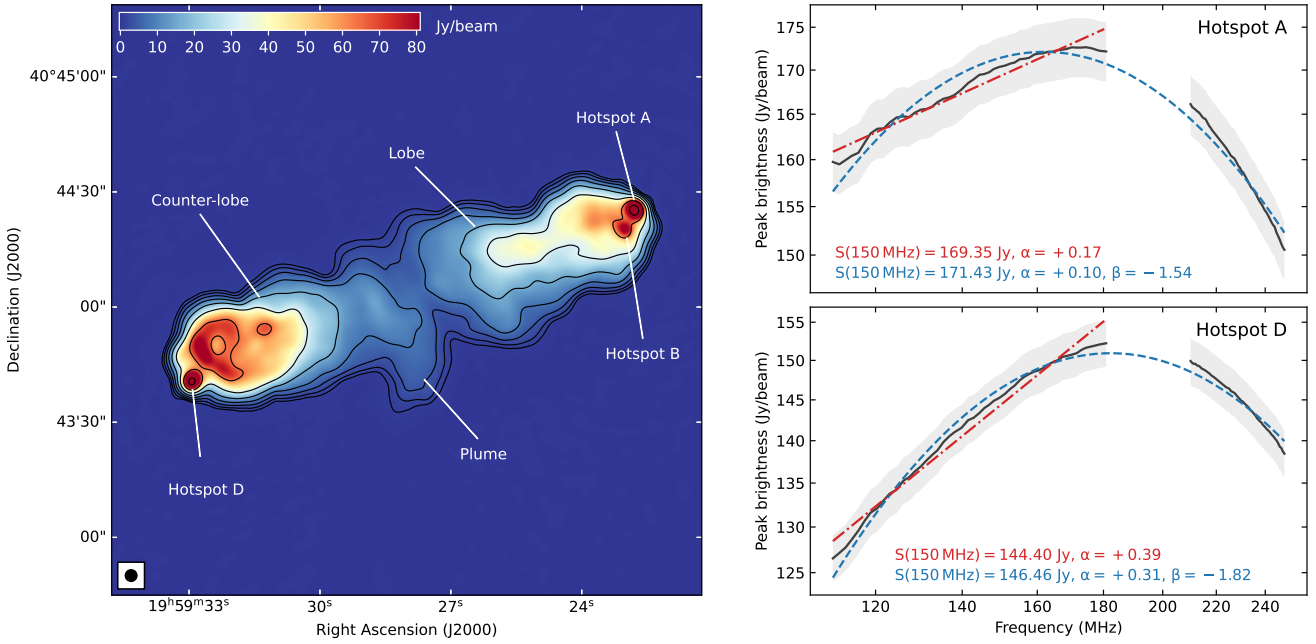


Fig. 1. High-resolution, frequency-integrated image of Cyg A between 111 and 249 MHz (left), and peak brightness of the hotspots A (top right) and D (bottom right), extracted from the high-resolution images of the 100 output channels (black solid line, with 1σ uncertainties shown as a grey shaded area). The fits using a second-order (blue dashed line) and a first-order logarithmic polynomial function (red dot-dashed line) are also shown, with the best fit parameters reported in the same colours. The latter has been fitted only between 111 and 181 MHz for a consistent comparison with MK16. The image on the left has a noise level of $\sigma = 10$ mJy/beam, with contours starting at 1 Jy/beam and increasing by a factor of 2.

uv -limit reduced the resolution from 3 to 4.5 arcsec, leading us to fix the restoring beam to a circular Gaussian with a FWHM of 5 arcsec. Consequently, we increased the pixel scale from 0.25 arcsec (used for high-resolution imaging) to 1 arcsec, producing images of 256×256 pixels. We also reduced the number of output channels to 41 to get a better noise per output image. This number was still high enough to minimise any bias that would occur during the later fitting process. We measured a noise level of $\sigma = 12.8$ mJy/beam for the frequency-integrated image (dynamic range of 1:20 550). With these settings, we found that the MS deconvolution gave better results with the default spatial scales of 12, 25, 49, and 99 arcsec. The other WSCLEAN parameters were similar to those used in the high-resolution imaging. As we will apply similar settings in the next forced-spectrum fitting step, all parameters are listed in the ‘Forced spectrum’ column of Table 2, with specific parameters for this step provided in square brackets.

The resulting 41 output images were used to extract spectral index and curvature maps. The brightness of each pixel was fitted using Eq. (2) at the reference frequency of 190.18 MHz, which is the same frequency at which WSCLEAN extracted the model clean-components for our specific settings¹². In this operation, we considered only pixels brighter than 10 Jy/beam in the frequency-integrated image (this roughly corresponds to 3.6 Jy/beam in the frequency-integrated, high-resolution image of Fig. 1). The extracted spectral index and curvature maps are shown in Fig. 2. Although the image resolution is lower than in the preliminary imaging step, the large LOFAR HBA bandwidth and high image dynamic range still allow us to identify the main spectral features across the entirety of Cyg A. Hotspots A and D exhibit the flattest spectral indices and the lowest cur-

vature, as expected due to their turnover. Additionally, in line with synchrotron spectral ageing models (e.g. Carilli et al. 1991; MK16), the spectral index map shows a steepening from the hotspots, where relativistic particles are continuously accelerated by the jets, to the outer regions of the lobe and counter-lobe, where the primary energy loss mechanism is synchrotron radiation. More energetic plasma, which emits at higher frequencies, loses energy faster than less energetic plasma, creating a break frequency that shifts to lower frequencies as the plasma ages (unless re-acceleration occurs). Because the plume of Cyg A has the steepest spectral index ($\alpha < -1$), it consists of the oldest plasma in the source, indicating that this plasma underwent its last acceleration much earlier than in any other region.

2.3.3. Forced-spectrum fitting and spectral rescaling

In both Sects. 2.3.1 and 2.3.2, the imaging was performed without any constraints on the spectral behaviour. In the next processing step, the extracted spectral index and curvature maps were used as inputs to the forced-spectrum method, acting as an imaging regularisation. We mostly applied the same WSCLEAN settings as in the previous step, as listed in the ‘Forced spectrum’ column of Table 2. The same uv -range was used to ensure that the sampled spatial scales were consistent with the spectral index map. While we could have used higher resolution to obtain a more detailed spatial model of Cyg A, the 5 arcsec resolution of the spectral index map meant that this would not have provided any additional benefit on the spectral modelling side. Furthermore, we preferred to maintain the same uv -cut and resolution to limit the number of clean-components in the final model, which consisted of 2082 components (1320 point sources and 762 Gaussians).

Due to the spectral constraints of the forced-spectrum method, the residual images may exhibit higher residuals, resulting from

¹² The reference frequency needs to be checked beforehand each time, as different weighting and uv -coverage can shift it with the same dataset.




Fig. 2. Forced-spectrum input spectral index (left) and curvature maps (right) at 190.18 MHz, extracted from the 5 arcsec resolution images. Contours in both figures are the same of the left panel of Fig. 1. The colour scale of the spectral index map spans the range $-1.27 \leq \alpha \leq -0.38$, while the curvature is plotted in the range $-1.04 \leq \beta \leq 1.04$.

the clean-component overlap described in Ceccotti et al. (2023). This occurs because the spectral index is assigned based on the clean-component position, but for Gaussians, the wings overlap with other components in regions where the spectral index differs, causing a bias in the final model image and, consequently, in the residual image. For spectrally complex sources such as Cyg A, this bias can be significant, particularly in the extended regions where many Gaussians are used during cleaning. To mitigate this, we constrained the cleaning process using a spatial mask that included only pixels brighter than 1 Jy/beam, based on the previous frequency-integrated image (corresponding to approximately 0.4 Jy/beam in the high-resolution image of Fig. 1). This approach allowed us to push the initial cleaning threshold deeper, reducing residual emission and preventing artefacts from being cleaned beyond the source extension. As a result, clean components can still be selected between the spatial mask boundary and the spectral index map extension, which was defined down to 10 Jy/beam. To minimise the impact of the overlap bias in the spectrum of the final model within the 10 Jy/beam contour, we assigned values of $\alpha = 0$ and $\beta = 0$ to the remaining areas of the spectral maps.

The output of our forced-spectrum imaging step is a new, high-resolution model of Cyg A, whose total flux density matches, between 110 and 250 MHz, the spectrum found by MK16 as shown in the top panel of Fig. 3. However, when extrapolated outside the LOFAR HBA frequency range, the spectral energy distribution diverges from MK16, likely because a sum of components with varying spectra would not necessarily follow the flux scale to which the data were calibrated. Moreover, the overlap bias might play a role. While extrapolation always requires caution, this new Cyg A model could be useful for applications at lower frequencies, such as calibrating or demixing LOFAR low band antenna (LBA) data, between 10 and 80 MHz. To address this, we evaluated and corrected the total flux density of our model at 10 frequencies between 10 and 300 MHz to match MK16. Each clean component was then adjusted by the found correction factors and fitted to a logarithmic polynomial to restore spectral index and curvature values at the model reference frequency. The rescaled model agrees with MK16 as shown by the black line in Fig. 3. The middle and bottom panels show the flux density of hotspots A and D, integrated within a circle of 5 arcsec diameter. Compared to Fig. 3, the hotspot spectra appear flatter, with the turnover shifted to lower frequencies. This is due to the lower resolution used for

both the spectral maps and the final model, allowing more lobe emission to contaminate the hotspots. Such flattening is accentuated by the rescaling, although most effects impact frequencies outside the LOFAR HBA range.

2.3.4. Old versus new Cygnus A model

The new, rescaled model of Cyg A and the old model are shown, respectively, in the top and bottom panels of Fig. 4, both rendered at 150 MHz. It is immediately clear that the new model is at a higher resolution than the old one. Hotspots A, B and D are modelled with a dominant point component, while the diffuse lobe and counter-lobe emission consists of a collection of Gaussian components of different sizes. Because WSCLEAN uses only circular Gaussians during the MS deconvolution, point components are required to shape these Gaussians into real source structures, such as the bow-shaped structure in the counter-lobe. This explains the presence of components with negative fluxes, that have no physical meaning but they are still needed to get models resembling real data. The final result is a model where also the fine structures of Cyg A are visible and easily distinguishable, fulfilling goal (2), mentioned in Sect. 2.1. On the other hand, we can hardly distinguish the different structures in the old model. Assuming that the two outermost Gaussian components model the hotspots A and D emission, a shift in the coordinates towards the north-west is present in the old model, the cause of which is unknown. This issue would be easily solved by shifting back each components to the correct location in the sky, but it is hard to find a reference point, such as hotspot A, in this model. When the old model is used on its own for calibration or demixing, such an offset, which is not too large, is absorbed into the gain phases.

The new model also outperforms the old one in terms of spectral modelling. The forced-spectrum method transferred the pixel-by-pixel spectral information from the spectral index and curvature maps into the new model, whose components have different spectra depending on their location, as we saw in the middle and bottom panels of Fig. 3. Thus, the resulting model of Cyg A also satisfies goal (3). Plotting the old model in the top panel, extrapolated from 74.8 MHz to the HBA frequency range, shows the large difference in flux scale between the expected values, as already discussed at the beginning of Sect. 2.

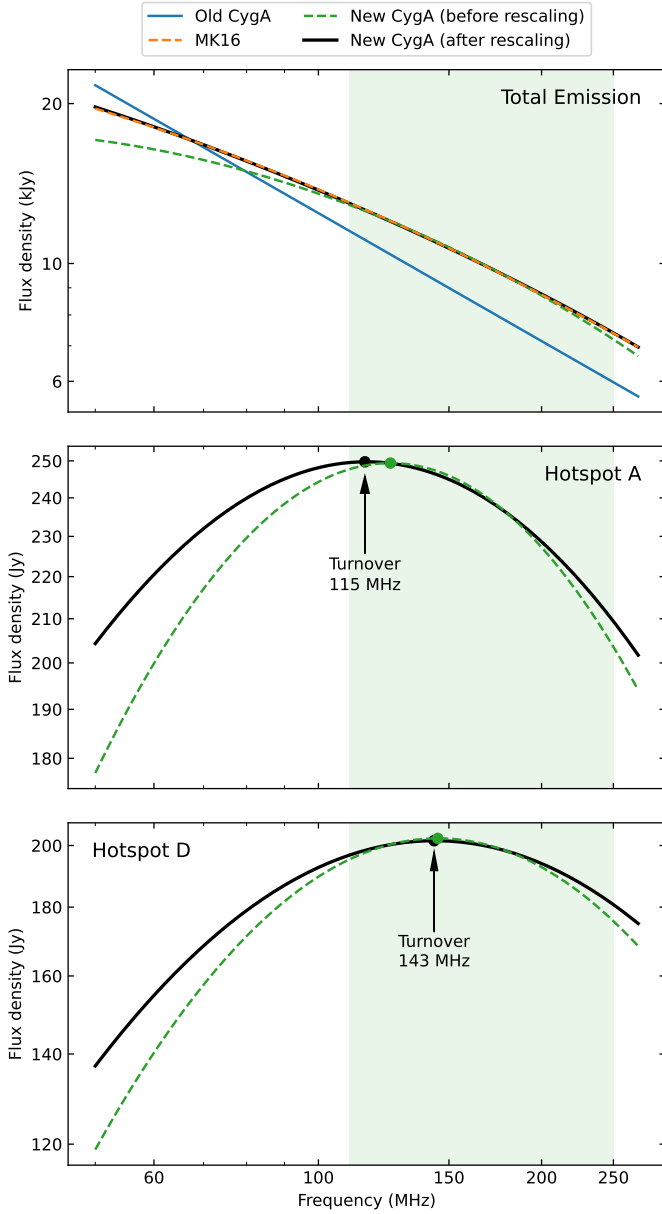


Fig. 3. Flux density between 50 and 270 MHz measured from the Cyg A models for the total emission (top), hotspot A (middle), and hotspot D (bottom). The top panel shows the old model (solid blue line), the new model before (dashed green line) and after spectral rescaling (solid black line), and the third-order logarithmic polynomial used by MK16 (dashed orange line). In the middle and bottom panels, only the flux densities from the new Cyg A models before and after rescaling are shown. These values were extracted from a circular region with a 5 arcsec diameter, centred on the peak of each hotspot. The filled points indicate the spectral peaks, and the turnover frequency after rescaling is also reported. The green shaded area in all panels marks the LOFAR HBA range used to extract the new Cyg A model.

3. NCP datasets

To test the performances of the new model of Cyg A against the old one, we used it within the LOFAR-EoR processing pipeline, comparing the results both from a simulated dataset and from a real observation of the NCP field. This allows us to assess the impact of this better model on the 21-cm power spectrum measurements, which is one of the main motivations of this work. We selected an average-quality night from Mertens et al.

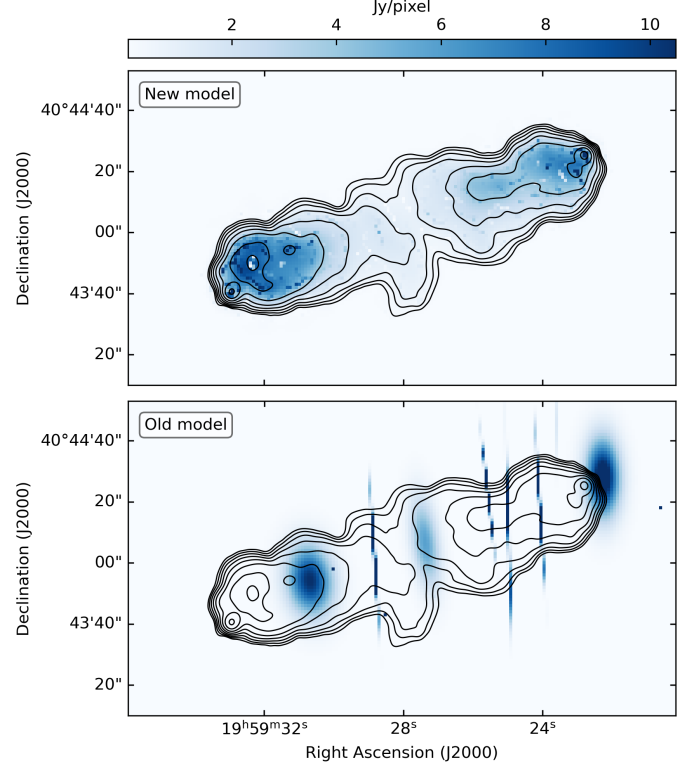


Fig. 4. New (top) and old model (bottom) of Cyg A, rendered at 150 MHz. Contours are the same of Fig. 1. The zero-length axes of the old model Gaussians have been replaced with 1 arcsec axes for imaging purposes, resulting in the appearance of vertical spikes.

Table 3. Observational details of the NCP dataset used for testing the Cyg A model.

Parameter	Value
Telescope	LOFAR HBA
Project code	LC2_019
Observation ID	L246309
Antenna configuration	HBA dual inner
Number of stations	62 (CS + RS)
Phase centre (J2000):	NCP (0 ^h , +90°)
Obs. start time (UTC):	2014 Oct. 16; 17:11:36
Frequency range:	134.1–147.1 MHz
Duration of observation	12.6 h [6 h]
Time resolution	2.0 s
Frequency resolution	61.0 kHz

Notes. The duration value within square brackets refers to the simulation dataset.

(2020) where Cyg A is above an elevation of $\approx 10^\circ$ for most of the observing duration. The observational details of the selected dataset are reported in Table 3.

To produce the simulated dataset, we made a copy of the same measurement set and forward predicted (i.e. Fourier transformed from image space to visibility space) a sky model into it. When predicting visibilities, the calculation of the primary beam is computationally expensive. The computing time increases with the number of model components, because the primary beam has to be calculated at each component coordinate. Because the NCP sky model has more than 28 700 components (Yatawatta et al. 2013; Mertens et al. 2020), we used a

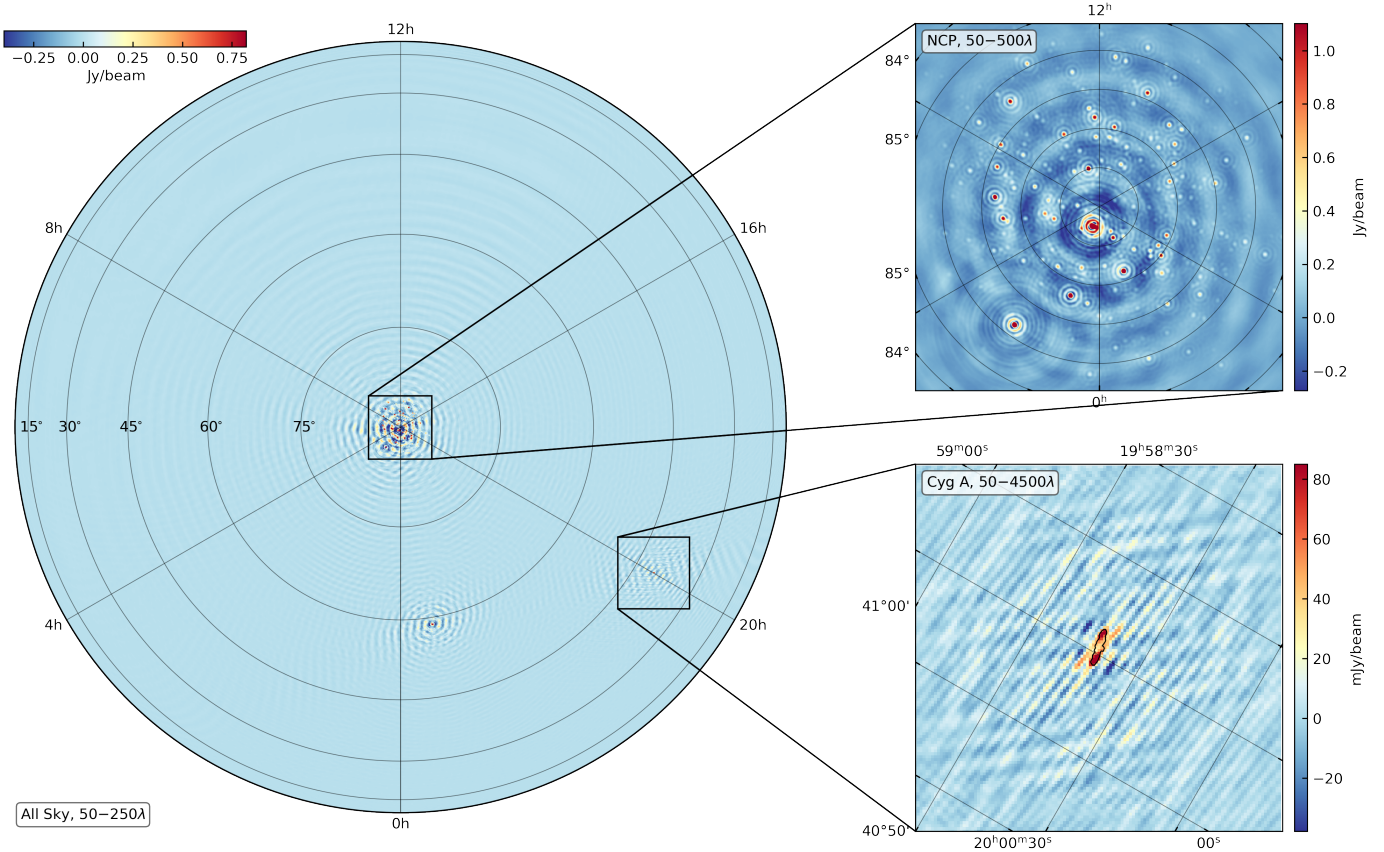


Fig. 5. All sky dirty image of the simulated NCP dataset (left), with zoom-in to the NCP main field (top right) and the Cyg A direction (bottom right). The all sky image has been obtained using a baseline range of 50–250 λ , the NCP zoom-in with 50–500 λ , while the Cyg A zoom-in with 50–4500 λ . A natural weighting scheme has been used for the all sky and the NCP zoom-in images, whereas a uniform weighting has been chosen for the Cyg A zoom-in to make the source structures more visible. The contour in the bottom right panel is the 1 Jy/beam level from Fig. 1 to show the position and extension of Cyg A.

simpler NCP model with only 684 point components with a flat spectrum, presented in Brackenhoff et al. (2024). This model extends up to 6° from the NCP and consists of 14 source clusters (i.e. directions), and its sources dominate the total power spectrum of the sky, making it sufficient for the purpose of this work. More details are reported in Appendix C. We added a low resolution model of Cas A, consisting of 24 Gaussian components with $\alpha = -0.8$ constant across the source and $\nu_0 = 74$ MHz, similar to the old Cyg A model. Finally, we add our new Cyg A model. Assuming that this new model is the true sky, we will be able to measure the theoretical excess power that might be caused by not having a Cyg A model with correct spectral indices.

The combined sky model was predicted on to the visibilities of the previously selected night with SAGECAL¹³ (Yatawatta 2016), which allows using GPUs to speed up the process (see Sect. 4). The primary beam was applied during the prediction, and we finally added a per-visibility noise estimated from the radiometer equation (e.g. Thompson et al. 2017):

$$\sigma_{\text{vis}} = \frac{\text{SEFD}}{\sqrt{2\Delta t \Delta \nu}}, \quad (3)$$

where SEFD is the system equivalent flux density, $\Delta t = 2$ s is the visibility integration time and $\Delta \nu = 61$ kHz is the frequency

channel width. The SEFD of CS and RS is expected to be different because RS have double collecting area (van Haarlem et al. 2013). However, our observations were carried with the ‘Inner’ setup, which sets the RS collecting area to the CS one by switching off the outer tiles of the station. For this reason, we assigned to every visibility the same SEFD of 4253 Jy as estimated by Mertens et al. (2020). To reduce the computing time, we kept only the first 6 h of the dataset, where the elevation of Cyg A peaks (77.9° after 35 min from the beginning) and stays above 35°.

The full-sky image of the NCP simulations is shown in Fig. 5, with a zoom-in of the NCP main field (top-right panel) and into the Cyg A field (bottom-right panel). Cyg A is 49.3° away from the NCP, and, because of that, the primary beam attenuates its flux by a factor of $\approx 10^{-4}$. However, its sidelobes are still strong and extend across the whole image, reaching the NCP main field and affecting its residuals if the source is not well modelled and subtracted. The same holds for Cas A, the effect of which could be even larger because it is closer to the NCP (31.2° away) and has a higher apparent brightness.

4. NCP data processing

Both datasets have been processed through an updated LOFAR-EoR production pipeline, which steps are extensively described by Mertens et al. (2020, 2025). The pipeline consists essentially of (1) pre-processing and RFI removal, (2) initial calibration, (3) DD solving and DD sky-model subtraction, (4)

¹³ <https://sagecal.sourceforge.net/>

imaging, (5) removal of residual foregrounds, and (6) power spectrum estimation. In this work, we skipped step (1) because we did not include any RFI into the simulated measurement set and we start with the already pre-processed observed data from Mertens et al. (2020). Step (5), which is performed with the Gaussian process regression (GPR) method (Mertens et al. 2018, 2024; Hothi et al. 2021; Acharya et al. 2024) was also skipped for both simulated and observed datasets, because this step does not use a sky model and we therefore do not need it to make the final comparison. In the following, we briefly describe the steps used for this work.

4.1. Initial calibration

The main goal of the initial calibration is to set the correct flux scale and correcting the sky-averaged phase offset per station. For the NCP, this is done in an unconventional way, by obtaining calibration solutions for two directions. One direction covers 3C 61.1, which is the brightest source in the main field with a flux density of ≈ 38 Jy at 150 MHz (Baldwin et al. 1985) and close to the first null of the primary beam, making it necessary to solve the varying DD effects of the primary beam and imperfections in the LOFAR primary beam model. The other direction includes the rest of the field, where only sources brighter than 35 mJy are selected to reduce the computing time while keeping a high enough signal-to-noise ratio. This reduces the NCP sky model used for our real observations from 28 773 to 1416 sources, thereby reducing the computational cost. The remaining sources account for only 1% of the total flux model, and this has therefore only a small effect. For consistency, we performed a similar procedure for the sky model used for the simulations, but such a flux cut has a relatively small effect here, because the model has fewer sources (i.e. from 684 to 666 components). During the initial calibration, we only use baselines in the range $50\text{--}5000\lambda$. The lower limit was applied in real observations to avoid poor uv -coverage and the diffuse Galactic emission which is not present in our model and would bias the calibration (Patil et al. 2017), while the upper limit minimises ionospheric effects and avoids resolving compact sources (Patil et al. 2016; Mevius et al. 2022).

The initial calibration was performed with SAGECAL, that uses the space alternating generalised expectation maximisation (SAGE) algorithm (Kazemi et al. 2011). It distributes the processing across multiple nodes using a consensus optimisation (Boyd et al. 2011). This method applies a frequency smoothness regularisation with a third-order Bernstein polynomial (Farouki & Rajan 1988; Yatawatta 2019), using an alternating direction method of multipliers (ADMM; Yatawatta 2015, 2016). The level of frequency smoothness for each iteration is determined by the regularisation parameter ρ , whose value depends on the apparent sky brightness for each source cluster and is therefore direction dependent (Mevius et al. 2022). More details about SAGECAL and its capabilities can be found in Yatawatta (2015, 2016, 2018), Yatawatta et al. (2017), and Gan et al. (2023). This smooth gain calibration was performed with a frequency and time solution interval of 195 kHz (i.e. one solution per SB) and 30 s, respectively, to address relatively fast direction-independent ionospheric effects on the longer baselines. We note, though, that we did not include any ionospheric effects in our simulations. Recent studies have shown that the effect of the ionosphere is currently still expected to be below the thermal noise level (Vedantham & Koopmans 2015, 2016; Brackenhoff et al. 2024). The data were corrected using the solutions from the main field.

After this high-temporal and smooth-spectral resolution calibration, we performed a low-temporal and high-spectral resolution bandpass calibration with the same sky model to solve for the fine frequency response of the instrument. This was done using a single solution per SB over a time interval of 3 h to mitigate the introduction of frequency structures due to an incomplete or inaccurate model as much as possible. We still used SAGECAL, but each SB was solved independently and no smoothness was applied to the solutions. As before, only the main field solutions were applied to the data.

After the bandpass calibration, all visibilities with amplitude higher than 70 Jy were flagged to prevent SAGECAL from diverging from the correct solutions in the subsequent DD calibration step (Mertens et al. 2020). We also averaged the data to a time resolution of 10 s to speed up the next step, which is the most demanding step of the pipeline.

4.2. Direction-dependent solving and source subtraction

The wide field of view of LOFAR introduces DD effects due to, for example, the time-varying primary beam and the ionospheric effects, that have to be taken into account to perform a good foreground subtraction. This is done by clustering the sky model components into a number of directions and ensuring a spectral smoothness using the consensus optimisation of SAGECAL. These clusters typically are $1\text{--}2^\circ$ in size to reach a sufficient signal-to-noise with the sources inside the cluster and to minimise spatially varying DD effects. For the simulations, the NCP main field was divided into 14 clusters (see Table C.1), while for real observations it currently consists of 108 clusters, because the sky model extends up to 15° from the NCP. Two extra clusters were added for Cas A and Cyg A, because, similarly to the 3C 61.1 case, they are so bright and so far away that they need separate solutions. No flux limit was applied to the sky model during this DD step. The gains were solved for all clusters simultaneously using a frequency interval of 195 kHz. We set solution time intervals between 2.5 and 20 min for the main field directions and an interval of 2.5 min for Cas A and Cyg A to capture faster primary beam changes. All the sources in the sky models are included in this step.

To avoid over-fitting, a uv -cut at 250λ was introduced between baselines used for calibration and 21-cm signal extraction, where only baselines in the range $250\text{--}5000\lambda$ were used during the DD solving. This cut causes excess noise on the uncalibrated baselines, which is reduced by enforcing spectrally smooth solutions (Yatawatta 2016; Barry et al. 2016). Using this approach, it has been shown that there is no signal loss on the baselines of interest and very limited excess power (Mouri Sardarabadi & Koopmans 2019; Mevius et al. 2022).

In contrast to the initial calibration, after the DD gain solving we did not correct the data with the solutions. Instead, we used the gains when subtracting the sky model from the visibilities, to obtain the residuals that were finally used for creating power spectra.

In order to compare the performance of the new Cyg A model with the old one, we performed the DD calibration steps for both Cyg A models. In addition, we did a third run in which we modelled Cyg A as a point source with a spectral energy distribution following a third-order logarithmic polynomial with the coefficients given in MK16. This extreme case will demonstrate whether having a spatially resolved model is important, even when the source is not resolved on the baselines used for calibration. In fact, calibration can still be affected by the source structure if the signal-to-noise ratio is high enough, as in the case of

Cyg A. The results from the simulations (see Sect. 5.2) that used the new model during the DD steps are only a benchmark for the calibration efficiency within the LOFAR-EoR pipeline, and do not say anything about how good our new model is. After the DD calibration, we considered only the first 6 h of the observed data, in order to better compare the results with the simulated dataset.

4.3. Imaging and conversion to Kelvin

After calibration and foreground subtraction, we gridded and imaged each SB of the residual visibilities using WSCLEAN, making an (l, m, ν) image cube. During the gridding, each visibility received equal weight and we applied a Kaiser-Bessel anti-aliasing filter with a kernel extent of 15 uv -cells, an oversampling factor of 4096, and 32 w -layers. These parameters ensure that any error arising from the gridding process is confined significantly below the expected 21-cm signal (Offringa et al. 2019b,a; Mertens et al. 2020).

We produced Stokes I and V images with a natural weighting scheme. Additionally, we generated even and odd time-step Stokes V images, the difference of which is used to estimate the thermal noise variance. We used a baseline range of 0–600 λ during the imaging, setting a pixel scale of 0.5 arcmin and an image size of 512 \times 512 pixels, which results in a field of view of 4.3 \times 4.3°. After gridding, only the gridded uv -values between 50 and 250 λ were used for the power spectrum.

The image cube, which has units of Jy/beam, has to be converted into units of brightness temperature, which is Kelvin. This is done by Fourier transforming the images into a gridded (u, v, ν) visibility cube and then using the conversion technique detailed by Offringa et al. (2019a). Making the visibility cube, we applied a spatial Tukey function with a 4° diameter to focus on the centre of the primary beam, which has a FWHM of $\approx 4.1^\circ$ at 140 MHz, while also mitigating its uncertainties farther from the centre. The gridded visibilities (in units of Kelvin) and the number of visibilities in each (u, v, ν) -cell were used to estimate the power spectrum.

4.4. Power spectrum estimation

If $T(\mathbf{x})$ is the brightness temperature of a signal at the physical coordinate \mathbf{x} , its power spectrum as a function of the wavenumber \mathbf{k} (units of $h \text{ cMpc}^{-1}$) is defined as

$$P(\mathbf{k}) = V |\tilde{T}(\mathbf{k})|^2, \quad (4)$$

where V is the observed comoving volume and \tilde{T} is the discrete Fourier transform of the brightness temperature. This 3D power spectrum is usually in units of $\text{K}^2 h^{-3} \text{ cMpc}^3$. The components of the wavenumber \mathbf{k} perpendicular and along the line of sight are respectively defined as (Morales & Hewitt 2004; Vedantham et al. 2012):

$$\mathbf{k}_\perp = \frac{2\pi}{D_M(z)} \mathbf{u}, \quad k_\parallel = \frac{2\pi H_0 \nu_{21} E(z)}{c(1+z)^2} \eta, \quad (5)$$

where $D_M(z)$ is the transverse comoving distance at redshift z , $\nu_{21} = 1420 \text{ MHz}$, H_0 is the Hubble constant, $E(z)$ is the dimensionless Hubble parameter, c is the speed of light, and η is the Fourier dual of the frequency ν (often referred as delay).

The power $P(\mathbf{k})$ can be averaged in cylindrical shells of radius $k_\perp = |\mathbf{k}_\perp|$ to get the cylindrically averaged (2D) power spectrum:

$$P(k_\perp, k_\parallel) = \frac{\sum_{\mathbf{k} \in (k_\perp, k_\parallel)} P(\mathbf{k})}{N_{(k_\perp, k_\parallel)}}, \quad (6)$$

where $N_{(k_\perp, k_\parallel)}$ is the number of \mathbf{k} voxels falling into the (k_\perp, k_\parallel) -bin. If we choose spherical shells of radius $k = |\mathbf{k}| = \sqrt{k_\perp^2 + k_\parallel^2}$, a spherically averaged (1D) power spectrum can also be obtained:

$$\Delta^2(k) = \frac{k^3}{2\pi^2} \frac{\sum_{\mathbf{k} \in k} P(\mathbf{k})}{N_k}, \quad (7)$$

where N_k is the number of \mathbf{k} voxels falling into the k -bin. This is the power spectrum estimator used to set 21-cm upper limits because it is in units of mK^2 . Both cylindrical and spherical power spectra are weighted by the gridded visibility thermal noise (see Mertens et al. 2020, for more details).

Because k_\parallel is essentially the Fourier dual of the frequency, a spectrally smooth signal is confined to low k_\parallel values. This is the case for Galactic and extra-galactic foreground emission, which is expected to be dominated by synchrotron radiation at low frequencies and spectrally smooth over tens of MHz. On the other hand, the 21-cm signal is fluctuating rapidly in frequency, and its power extends over all k_\parallel (e.g. Santos et al. 2005). However, every interferometer is intrinsically chromatic because a single uv -cell is crossed by different baselines at different frequencies, and this effect becomes stronger at longer baselines, corresponding to larger k_\perp (Morales et al. 2012, 2019). Because of this, the foreground power is spread at higher k_\parallel , generating a characteristic shape known as the ‘foreground wedge’ in the cylindrically averaged power spectrum (Datta et al. 2010; Liu et al. 2014a,b). The maximum extent of the wedge in k_\parallel is given by the horizon line, that can be accurately calculated in a non-flat sky formalism knowing the horizon delay (Munshi et al. 2025). In the same way, we could find the delay lines for every direction in the sky (Munshi et al. 2025). We then expect that all the sky emission lies in the wedge below the horizon line, leaving a foreground-free region called ‘EoR window’. However, the finite bandwidth of the data causes the Fourier transformation to leak some power above the wedge, requiring some frequency window function to reduce such a leakage (e.g. Vedantham et al. 2012). Also calibration errors might introduce rapid frequency fluctuations that contaminate the EoR window. This effect can be limited by using frequency smoothed gains, as we did in the first calibration step and during the DD solving.

In this work, we estimated the power spectra using the power spectrum pipeline PPIPE¹⁴. We used only gridded visibilities in the 50–250 λ baseline range and applied a Hann window function to the frequency bandwidth. The reported power spectrum uncertainties were estimated from the sample variance, as described by Mertens et al. (2020).

5. Results

In this section, we show the results from both the simulated and observed datasets, comparing the residuals after the DD subtraction and their power spectra. We also compare results from the three Cyg A models described previously: the old, new, and point source models. All the cylindrical power spectra are shown as a ratio between the Stokes I and the thermal noise power spectrum, where the thermal noise was evaluated from time differences of Stokes V images. This ratio removes most of the effects that are caused by the sensitivity response of the instrument, allowing us to focus on the differences between the different Cyg A models.

In the next section, we sketch how errors in the sky model could propagate into the foreground-subtracted visibilities after the DD subtraction. This is an important aspect to better understand the following results.

¹⁴ <https://gitlab.com/flomertens/pspipe>

5.1. DD gain errors due to model incompleteness

In the following sections, the effect of model incompleteness for Cyg A, after its subtraction, is illustrated both in the image domain via residual images and in the power spectrum domain. The latter is particularly relevant for 21-cm cosmology experiments on baselines where the 21-cm signal is expected and illustrates that the effect in the image domain can be extremely subtle because continuum images contain all baselines and are averaged over frequency, thereby averaging out frequency-dependent fluctuations. In the cylindrical power spectra (Sects. 5.2.1 and 5.3.1), however, the effects as a function of baseline (proportional to k_{\perp}) and frequency (or its Fourier dual, ‘delay’, proportional to k_{\parallel}) become much more evident.

In general, when determining LOFAR station-based gains, especially their DD gains, the incompleteness of the sky model can have a significant impact on the quality of the gains as a function of baselines and frequency (Patil et al. 2016; Ewall-Wice et al. 2017; Mouri Sardarabadi & Koopmans 2019; Mevius et al. 2022). When solving for multiple directions, the gains can modify the point spread function (PSF) used for subtraction in such a way that its sidelobes absorb unmodelled structures in the sky, thereby removing diffuse foregrounds and 21-cm emission on larger scales (see Patil et al. 2016 for a demonstration). This is why spectrally smooth gain regularisation was introduced (e.g. Yatawatta 2015; Gan et al. 2023), as it suppresses the ability of gains to remove unmodelled source structures.

Similarly, very bright sources far from the phase centre (e.g. Cas A and Cyg A) can have a dominant impact via their PSF sidelobes, leaking power into the main target field (e.g. NCP or 3C 196), as demonstrated by Gan et al. (2022) and Munshi et al. (2024). Hence, an incomplete sky model for these bright sources might lead to incorrect DD solutions in their direction. When subtracting them with their DD gains applied, the PSF-convolved source model will be incorrectly subtracted, leaving residuals in the main field of interest (NCP in our case). For fainter sources, this effect is much smaller and averages out when there are many sources in various directions and at different distances. However, Cyg A and Cas A are two extremely dominant sources, and their effect does not average away (Gan et al. 2022; Munshi et al. 2024).

The errors on these source models can then enter the power spectrum in several different and complex ways: (i) even if the DD solutions in the direction of a bright source were known perfectly, applying them to an incorrect source model would leave residuals in the main field via the sidelobes, as discussed above; (ii) in practice, however, the incomplete sky model (not just the bright source model) also leads to complex gain errors, even in the case of strong gain regularisation as a function of frequency (see e.g. Yatawatta 2015). The latter effect can redistribute power across baselines depending on which baselines exhibit the largest differences between data and model, as well as on which baselines have the most visibilities (both contributing to the minimisation of the penalty function for the gain solutions).

In the following sections, we illustrate these effects for the case of an incomplete model for Cyg A.

5.2. Results from simulations

After the DD subtraction step, we imaged the time and frequency averaged residuals at the Cyg A position in the simulated dataset for all the three tested models. In the top row of Fig. 6, from second to fourth panel, we show these residuals divided by the

apparent peak brightness of Cyg A in the initially calibrated data, shown in the first panel. This gives an estimate of the relative errors using different models during the DD solving. In the second image, the model used for prediction and calibration was the same. The residuals therefore do not reflect the quality of the model, but the effectivity of calibration. The third and fourth panels show the extra residual power caused by (i) not modelling spectral indices, and (ii) assuming Cyg A is a point source. In the bottom row we also report the differences between the three residual images, to better highlight the structures affected by the different Cyg A models.

Although the old model is incomplete and spectrally incorrect, it leaves residual similar to the new model, both with a maximum ratio of approximately 0.13 close to the source centre. The point source model shows higher residuals, with a maximum of 0.17. The removal of all the three models generated ripples far from the sources and directed toward the NCP, eventually affecting the power spectrum. The intensity of these ripples are slightly larger for the extreme point source model, which also shows stronger small-scale fluctuations in the perpendicular direction. It is interesting to note that residuals for the new model are not fully consistent with the noise, even though the same model was used for both prediction and DD calibration. This highlights that even in the perfect scenario, calibration errors still affect the resulting residuals. This can be partly attributed to the added noise and the non-linear behaviour of calibration. However, in our experience, sources in the centre of the primary beam, are easier to subtract and leave lower residuals. A possible explanation for the relatively large residuals with a perfect model might be that the source passes through nulls in the primary beam and that not all stations have sufficient signal-to-noise at all times to accurately solve the source.

5.2.1. Cylindrical power spectra and ratios

The cylindrical power spectra $P(k_{\perp}, k_{\parallel})$, centred on the NCP, were estimated both after initial calibration and DD subtraction, and show that the DD steps removed most of the foreground emission within the wedge. The top row of Fig. 7 shows the Stokes I power spectra divided by the thermal noise, estimated from time-differenced Stokes V . The thick line represents the horizon limit, above which the EoR window should be foreground-free. The other two pairs of lines, the solid and the dashed ones, show the range where we can maximally find Cyg A and Cas A, respectively. Even with error-free simulations, the residual power within the wedge is quite high after the DD subtraction, up to 10 times the thermal noise. This might be due to calibration errors occurring during the DD solving or inaccuracies in the primary beam prediction, which also leaks power to higher k_{\parallel} , into the EoR window. Most of the residual power is along the Cyg A direction, while Cas A seems better subtracted. The maximum residual brightness of Cas A is 1% of the peak of Cas A after the initial calibration and it is the same irrespective of which Cyg A model was used in the calibration. This shows that using spatially and spectrally different models of Cyg A during the DD calibration step does not considerably affect the gain calibration of other bright sources such as Cas A. The bottom-left panel of Fig. 7 shows that the DD subtraction removes most of the power within the horizon. The residual power is usually taken into account with GPR, a step skipped in this work.

Ratios of DD subtracted power spectra for the three Cyg A models are shown in the bottom row of Fig. 7. Overall, the new model leaves the least residual power compared to the other two models. This is expected because the new model

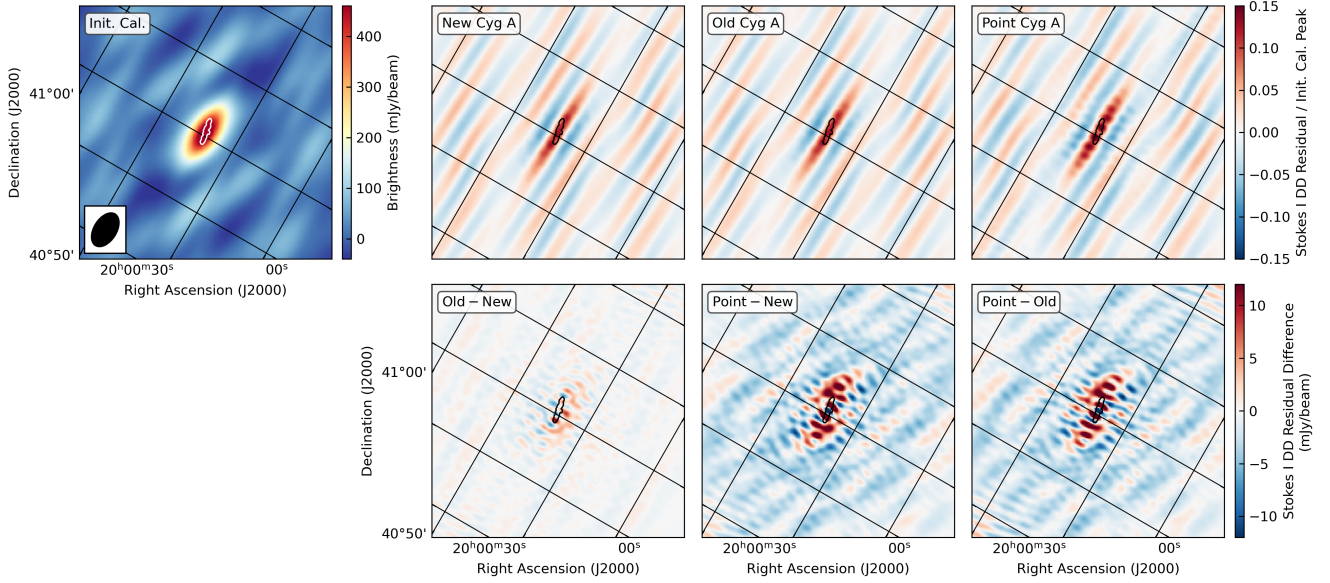


Fig. 6. Dirty images of the simulated NCP dataset in the direction of Cyg A. The top row shows, from left to right: (i) the image after initial calibration, and residual images after DD subtraction using the (ii) new, (iii) old and (iv) point source model, divided by the peak brightness of the initially calibrated image. The bottom row shows the differences between residual images obtained after DD subtraction for the three Cyg A models: Old – New (left), Point – New (middle), and Point – Old (right). All the images have been obtained with a natural weighting scheme and using a baseline range of 50–4500 λ , resulting in the synthesised beam shown in the bottom-left corner of the top-left panel. The contour in every panel is the 1 Jy/beam level from Fig. 1 to show the position and extension of Cyg A.

is the same model used during the prediction step. Most differences are visible around the Cyg A lines (thin solid lines), where the new model outperforms the old one, particularly at $k_{\perp} \lesssim 0.12 \text{ h cMpc}^{-1}$. The point source model, however, leaves the highest residuals across all $k_{\perp} - k_{\parallel}$ values within the Cyg A lines.

The top panel of Fig. 8 shows the 1D power spectrum $P(k)$ for the three models, extracted from the cylindrical power spectra (i.e. top row of Fig. 7, but not divided by the thermal noise) along the Cyg A direction and averaged within k -bins of size 0.07 h cMpc^{-1} .¹⁵ The middle and bottom panels show the ratios of these $P(k)$ pairs (white-faced markers), alongside the ratio values from bottom row of Fig. 7 for each $(k_{\perp}, k_{\parallel})$ -cell within the Cyg A delay lines (filled markers). We plotted the ratio $P_{\text{new}}(k)/P_{\text{old}}(k)$ separately in the middle panel to make it easier to see where the new model has improved residuals compared to the old model. The 1σ uncertainties associated with $P(k)$ have been calculated as

$$P_{\text{err}}(k) = \sqrt{\frac{\sum_{(k_{\perp}, k_{\parallel}) \in k} P_{\text{err}}^2(k_{\perp}, k_{\parallel})}{N_k^2}}, \quad (8)$$

where N_k is the number of $(k_{\perp}, k_{\parallel})$ -cells falling in k -bin and $P_{\text{err}}(k_{\perp}, k_{\parallel})$ is the error on the cylindrical power spectrum, estimated from the sample variance. For the ratio $P_i(k)/P_j(k)$, the sample variance is not a reliable uncertainty estimator, as all three datasets were drawn from the same sample and have the same noise. Therefore, the uncertainties of $P_i(k)/P_j(k)$ were set to the standard deviation over all $(k_{\perp}, k_{\parallel})$ -cells of the ratios shown in the bottom row of Fig. 7. The New/Old ratio reaches a minimum value of $P_{\text{new}}(k)/P_{\text{old}}(k) = 0.80 \pm 0.03$ at $k =$

0.30 h cMpc^{-1} and is significantly less than one for all $k < 0.4 \text{ h cMpc}^{-1}$. At higher k values, the ratio is close to one, suggesting that the new model does not offer improvement over the old model at smaller frequency and spatial scales. On the other hand, both $P_{\text{new}}(k)/P_{\text{point}}(k)$ and $P_{\text{old}}(k)/P_{\text{point}}(k)$ are consistently lower than one. The inverse-variance weighted averages across all k -bins are $\langle P_{\text{new}}(k)/P_{\text{old}}(k) \rangle = 0.95 \pm 0.05$, $\langle P_{\text{new}}(k)/P_{\text{point}}(k) \rangle \approx \langle P_{\text{old}}(k)/P_{\text{point}}(k) \rangle = 0.7 \pm 0.1$, indicating that the point source model performs worse than both the new and old models. This highlights that, although the point source model has correct spectral information, even a rough spatial characterisation of the source is necessary to improve DD calibration results.

5.2.2. Spherical power spectra and differences

Although the power spectra $P(k)$ along the Cyg A direction give a robust estimation of the impact that the three different models have on the residual power, we saw in the third and fourth bottom panels of Fig. 7 that differences are present also in different regions of the cylindrical power spectrum. We can then use Eq. (7) to estimate the spherical power spectrum $\Delta^2(k)$ in k -bins spanning the whole $(k_{\perp}, k_{\parallel})$ -space and assess the overall impact of the three models on the 21-cm upper limits.

The total Stokes I spherical power spectra are shown in the top panel of Fig. 9, evaluated within ten logarithmically spaced k -bins in the range $0.08 \leq k \leq 1.15 \text{ h cMpc}^{-1}$, with a bin size of $dk/k \approx 0.3$. The differences between model pairs $i - j$ and the associated 1σ uncertainties have been calculated as

$$D_{ij} = \Delta_i^2 - \Delta_j^2, \quad D_{ij, \text{err}} = \sqrt{(\Delta_{i, \text{err}}^2)^2 + (\Delta_{j, \text{err}}^2)^2}, \quad (9)$$

where $\Delta_{\text{err}}^2(k)$ is estimated from the sample variance of the power spectrum (see P_{err} in Eq. 8). In this case, the sample variance is used also for the uncertainties of the differences, because D_{ij}

¹⁵ Our definition of $P(k)$ differs from $\Delta^2(k)$ due to how the averages are done. In fact, $P(k)$ is calculated as a simple average of $P(k_{\perp}, k_{\parallel})$ without considering the weights, which are instead used to calculate $\Delta^2(k)$ from $P(k)$ (see Eq. 7).

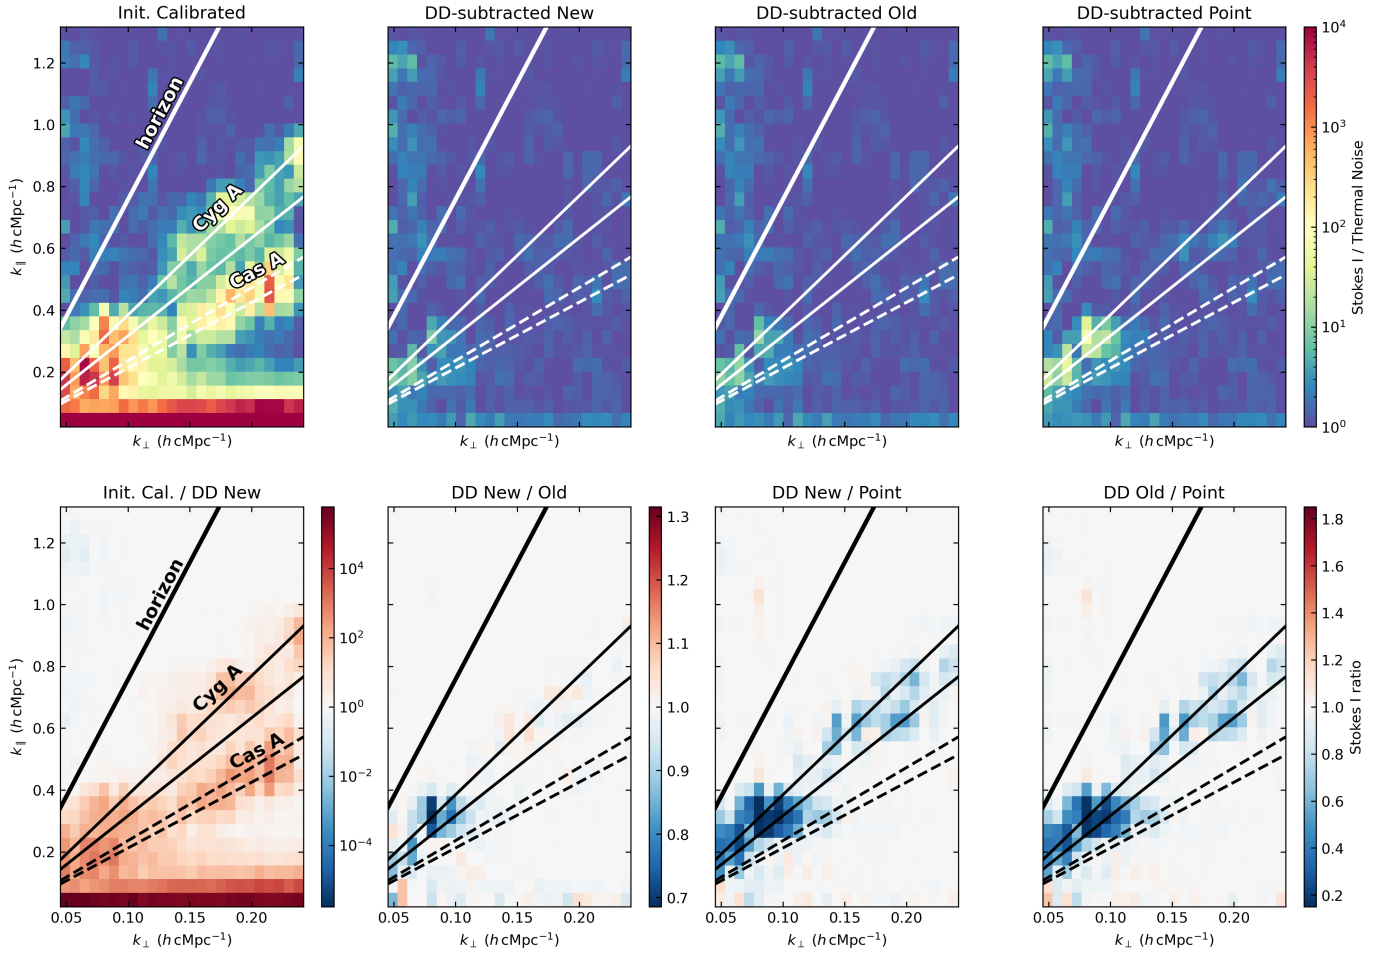


Fig. 7. Stokes I cylindrical power spectra divided by the thermal noise power spectrum at different stages of the processing of the simulated NCP dataset. Top row, from left to right: power spectrum ratio after initial calibration, after DD subtraction using the new, old, and point source model of Cyg A. The bottom row shows ratios of results from processing with different models. From left to right: the initial calibration results over new model results, new model results over old model results, new model results over point source model results, and old model results over point source model results. ‘DD New/Point’ and ‘DD Old/Point’ share the same colour bar. In all the panels, the thick line indicates the horizon delay line (foreground wedge), whereas the thin solid and dashed line pairs indicate the delay ranges where we expect most of the power of Cyg A and Cas A, respectively.

is calculated directly from $\Delta^2(k)$. The values are reported in Table D.1, with also the inverse-variance weighted mean of the differences calculated over two representative k -bin ranges. The difference $\Delta_{\text{old}}^2 - \Delta_{\text{new}}^2$ (middle panel of Fig. 9) is larger than zero across most of the k -bins, with just a few exceptions. However, all these values are consistent with zero when the sample variance error is taken into account, except at $k = 0.35 \text{ h cMpc}^{-1}$ where $\Delta_{\text{old}}^2 > \Delta_{\text{new}}^2$ significantly. We observe the strongest differences at this k -bin for all the model pairs, consistent with Fig. 8. Looking at the differences with the point source model (bottom panel of Fig. 9), we find that Δ_{point}^2 is higher than Δ_{new}^2 and Δ_{old}^2 across all k , with significant values at $0.19 \leq k \leq 0.85 \text{ h cMpc}^{-1}$.

5.3. Results from real observations

For the actual observations, we follow a similar route compared to the simulated data. In the top row of Fig. 10, from second to fourth panel, we show the residuals at the Cyg A position for all the three tested models, divided by the apparent peak brightness of the initial calibrated Cyg A, shown in the first panel. Differences between the residual images are shown in the bottom row. The residuals left from the new and old model are similar,

as we saw in the simulations, with a maximum ratio within the Cyg A region of 0.09 and 0.11, respectively. The point source model shows the highest residuals, with a maximum ratio of 0.23. Although the measured residuals for the new and old models are smaller than those in the simulations, the fluctuations originated from the subtraction of these models are stronger even far from the source and more complex. The subtraction of a single point component leaves the strongest sidelobes and ripples, both close and far from the source location.

5.3.1. Cylindrical power spectra

The top row of Fig. 11 shows the Stokes I cylindrical power spectra after the initial calibration and the DD subtraction, all divided by the thermal noise power spectrum. In this real case scenario, the DD subtraction is not able to remove some foreground emission at low k_{\parallel} . The power spectrum reaches 10^4 times the thermal noise level, because of systematics such as calibration errors and sky model incompleteness, including diffuse Galactic emission not present in the model. However, the decrease in power after the DD subtraction step is a factor of 10^3 at low k_{\parallel} . Some residual power is also observed along the Cas A




Fig. 8. Power spectra along the Cyg A direction after DD subtraction (top panel) and their ratios (middle and bottom panels) for the simulated NCP dataset. The top panel shows Stokes I $P(k)$ after the DD subtraction step using the new (blue), old (orange), and point source (green) models, along with the thermal noise level (dashed grey line). The shaded areas represent the 1σ uncertainties. The middle panel shows the ratio New/Old (black dots), while the bottom panel shows the ratios New/Point (pink squares), and Old/Point (purple diamonds). Filled markers represent the ratio of the cylindrical power spectra $P(k_{\perp}, k_{\parallel})$ for each $(k_{\perp}, k_{\parallel})$ -cell within the Cyg A delay lines, while the white-faced markers indicate the ratios of the power spectra $P(k)$. For the latter, the horizontal error bars indicate the k -bin extension, while the vertical error bars indicate the 1σ standard deviation over all $(k_{\perp}, k_{\parallel})$ -cells of the ratios shown in Fig. 7. A small offset in k has been added to the white-faced pink squares and purple diamonds in the bottom panel to avoid overlapping of error bars.

direction, which might be the cause of some of the ripples visible in Fig. 10.

In the bottom row of Fig. 11 (second to fourth panel), we show the ratios of the DD subtracted power spectra for the three Cyg A models. In all three cases, most differences are confined along the Cyg A direction. Although residuals from Cas A are observed in the top panels, the three different Cyg A models do not affect them differently, and the power along the Cas A direction cancels out in the ratio.

The 1D power spectra $P(k)$ along the Cyg A direction for the three models are shown in the top panel of Fig. 12. The ratio $P_{\text{new}}(k)/P_{\text{old}}(k)$ is shown in the middle panel, while

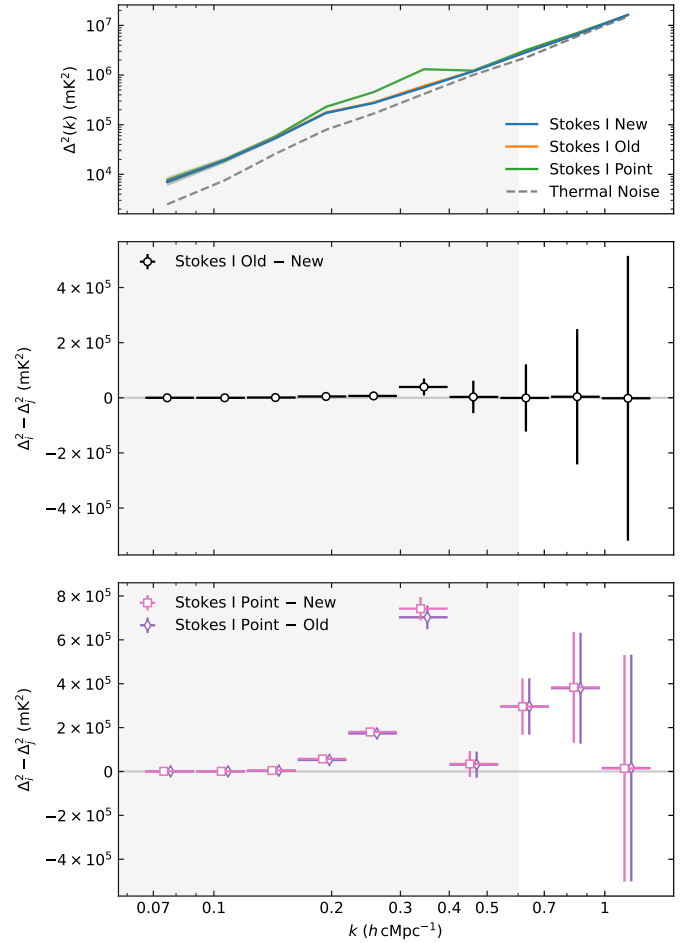


Fig. 9. Spherical power spectra after DD subtraction (top panel) and their difference (middle and bottom panels) for the simulated NCP dataset. The top panel shows Stokes I $\Delta^2(k)$ after the DD subtraction step using the new (blue), old (orange), and point source (green) Cyg A model, and the thermal noise level (dashed grey line). The shaded areas represent the 1σ uncertainties. The middle panel shows the difference $\Delta_{\text{old}}^2 - \Delta_{\text{new}}^2$ (black dots), while the bottom panel shows the differences $\Delta_{\text{point}}^2 - \Delta_{\text{new}}^2$ (pink squares), and $\Delta_{\text{point}}^2 - \Delta_{\text{old}}^2$ (purple diamonds), with the associated 1σ uncertainties (values reported in Table D.1). The horizontal error bars indicate the k -bin extension. We added a small offset in k to the pink and purple markers in the bottom panel to avoid overlapping of error bars. The grey shaded area delimits the range in k where Δ^2 upper limits are usually extracted from LOFAR data, being the most sensitive range for LOFAR.

$P_{\text{new}}(k)/P_{\text{point}}(k)$ and $P_{\text{old}}(k)/P_{\text{point}}(k)$ are shown in the bottom panel. In this real case, the 1σ uncertainties associated with ratio $P_i(k)/P_j(k)$ have been estimated from the sample variance as

$$\left[\frac{P_i}{P_j} \right]_{\text{err}} = \frac{P_i}{P_j} \sqrt{\left(\frac{P_{i,\text{err}}}{P_i} \right)^2 + \left(\frac{P_{j,\text{err}}}{P_j} \right)^2}, \quad (10)$$

where P_{err} is from Eq. (8). In real observations, the new model results in a residual power spectrum similar to the old model at low k . The ratio $P_{\text{new}}(k)/P_{\text{old}}(k)$ reaches a minimum of 0.8 ± 0.1 at $k = 0.89 h \text{ cMpc}^{-1}$, with an inverse-variance weighted average of 0.91 ± 0.05 over the entire k range. In contrast, the point source model shows higher power at all scales, with $\langle P_{\text{new}}(k)/P_{\text{point}}(k) \rangle = 0.67 \pm 0.03$ and $\langle P_{\text{old}}(k)/P_{\text{point}}(k) \rangle = 0.74 \pm 0.04$. These values are consistent with those measured from the




Fig. 10. Dirty images of the real observed NCP dataset in the direction of Cyg A. The top row shows, from left to right: (i) the image after initial calibration, and residual images after DD subtraction using the (ii) new, (iii) old and (iv) point source model, divided by the peak brightness of the initially calibrated image. The bottom row shows the differences between residual images obtained after DD subtraction for the three Cyg A models: Old – New (left), Point – New (middle), and Point – Old (right). All the images have been obtained with a natural weighting scheme and using a baseline range of 50–4500 λ , resulting in the synthesised beam shown in the bottom-left corner of the top-left panel. The contour in every panel is the 1 Jy/beam level from Fig. 1 to show the position and extension of Cyg A.

simulations. Interestingly, the new model leaves lower residuals than the old one at high k , whereas in the simulations we observed the opposite trend, with larger differences at small k . This discrepancy results from the DD calibration, which successfully removed most of the Cyg A power at $k_{\perp} < 0.12 \text{ h cMpc}^{-1}$ in the observed dataset (see Fig. 11), but not in the simulated one, where the cylindrical power spectra were noise-like at higher k_{\perp} (see Fig. 7).

5.3.2. Spherical power spectra and differences

The total Stokes I spherical power spectra $\Delta^2(k)$ are shown in the top panel of Fig. 13. As already seen in the top panel of Fig. 11, the power is a few orders of magnitude higher than the thermal noise at $k < 0.3 \text{ h cMpc}^{-1}$, which corresponds to approximately $k_{\parallel} < 0.3 \text{ h cMpc}^{-1}$ in the cylindrical power spectrum. Here, the foreground emission within and close to the main lobe of the primary beam has not been completely removed, because of either calibration errors or an incomplete sky model (e.g. extended Galactic emission is not removed and may affect low k_{\perp}). At higher k , Δ^2 is closer to the thermal noise, even if some excess power is present because of poorly subtracted off-axis sources, mainly Cyg A and Cas A. In fact, the difference between model pairs, shown in the middle and bottom panels of Fig. 13 and evaluated with Eq. (9), increases in these k -bins, reaching peaks of $\approx (1600 \text{ mK})^2$ for $\Delta_{\text{old}}^2 - \Delta_{\text{new}}^2$ and $\approx (4500 \text{ mK})^2$ for both $\Delta_{\text{point}}^2 - \Delta_{\text{new}}^2$ and $\Delta_{\text{point}}^2 - \Delta_{\text{old}}^2$ at $k = 0.85 \text{ h cMpc}^{-1}$. The real observed data do not prefer one model at $k \leq 0.35 \text{ h cMpc}^{-1}$, where all the differences have large uncertainties and are consistent with zero. This is in contrast with the simulations, where differences and uncertainties were smaller. This happens because in the observed data there are other systematics, such as model incompleteness, calibration errors, and inaccuracies in the primary beam modelling, that play a bigger role during the DD subtraction. Given the current situation of the systematics, we do not

expect to see larger differences by using different Cyg A models. However, an overall improvement is achievable by using the new model instead of the old one in the range $0.63\text{--}1.15 \text{ h cMpc}^{-1}$, where $\langle \Delta_{\text{old}}^2 - \Delta_{\text{new}}^2 \rangle = (1259 \pm 742 \text{ mK})^2$. At lower k , the average difference is consistent with zero, but the expected value is $\approx (257 \text{ mK})^2$ in the range $0.08\text{--}0.63 \text{ h cMpc}^{-1}$. This shows that, in principle, the new Cyg A model is capable of reducing the excess power by hundreds of mK, even if Δ_{new}^2 is still 2–3 times higher than the thermal noise level.

6. Discussion and conclusions

The latest upper limits on the 21-cm power spectrum from the LOFAR-EoR KSP show an excess power in $\Delta^2(k)$ with respect to thermal noise (Mertens et al. 2020). Gan et al. (2022) analysed various causes of such an excess, concluding that it might be dominated by residuals of bright sources, especially those in the far field. Their power spectra, even after GPR foreground removal, show non-negligible power in the Cas A and Cyg A direction, hinting that the modelling and subtraction of these A-team sources has been performed poorly. The models of these sources, also used by the current LOFAR-EoR processing pipeline, are low resolution, with no spectral variation and incorrect spectral indices. These models have been obtained from VLA data at $\approx 74 \text{ MHz}$ and assigned a fixed spectral index of -0.8 to each component, which is used to extrapolate values to the LOFAR HBA band where the 21-cm power spectrum is calculated.

In this work, we focused on Cyg A, the brightest of the A-team at 150 MHz, and used it as a test case to assess the impact of a better spatial and physical model on the 21-cm power spectrum. In Sect. 2, we introduced what we called the ‘old’ Cyg A model and showed that extrapolating it to the HBA frequency range results in 16% missing flux. Not including any spectral variation across the source is incorrect, because we

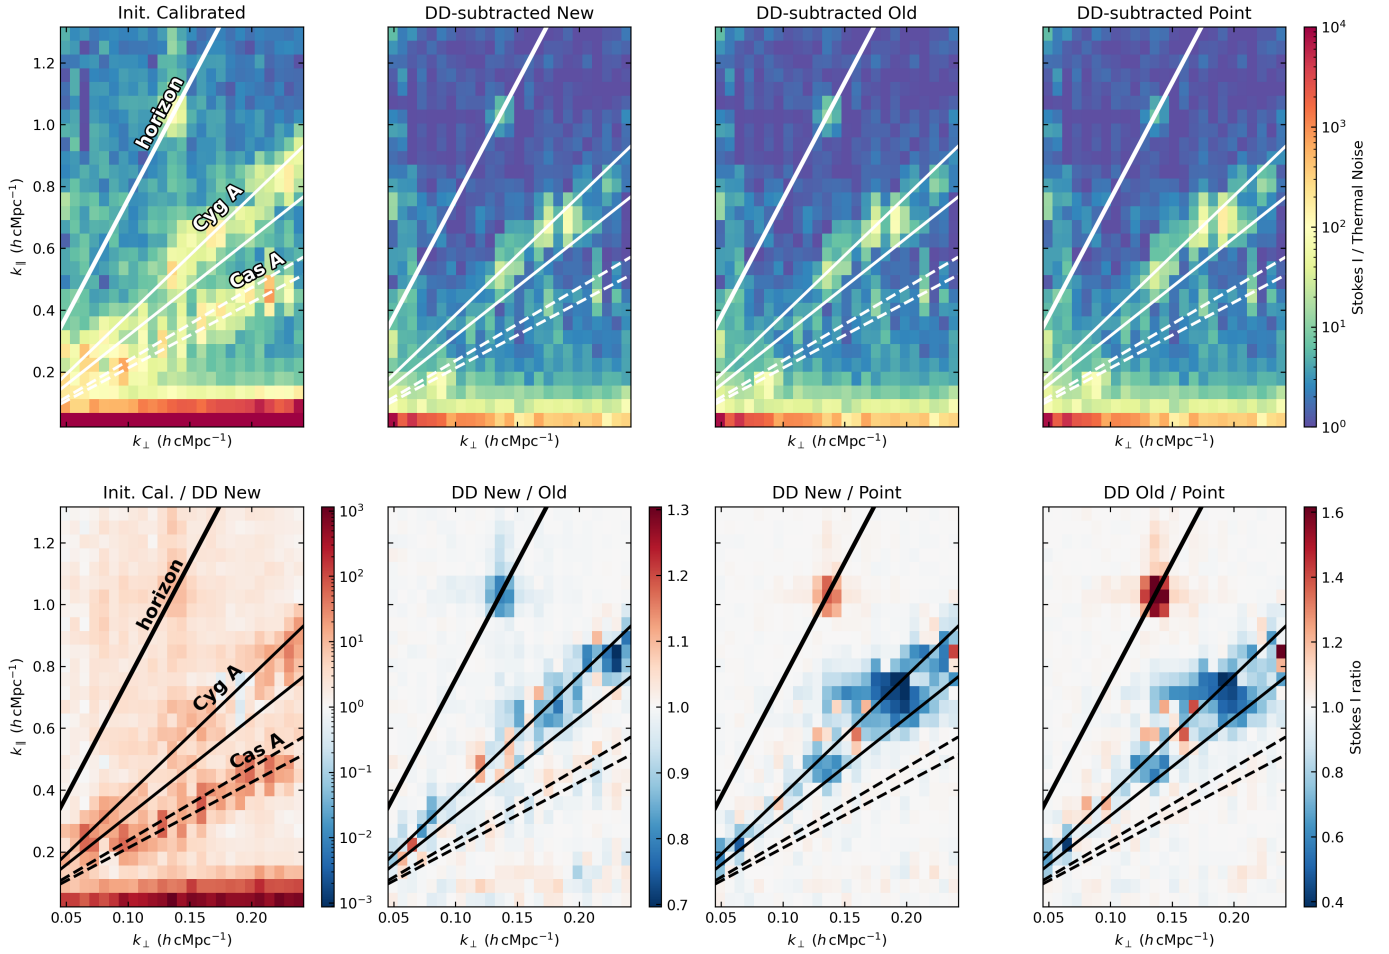


Fig. 11. Stokes I cylindrical power spectra divided by the thermal noise power spectrum at different stages of the processing of the real observed NCP dataset. Top row, from left to right: power spectrum ratio after initial calibration, after DD subtraction using the new, old, and point source model of Cyg A. The bottom row shows ratios of results from processing with different models. From left to right: the initial calibration results over new model results, new model results over old model results, new model results over point source model results, and old model results over point source model results. ‘DD New/Point’ and ‘DD Old/Point’ share the same colour bar. In all the panels, the thick line indicates the horizon delay line (foreground wedge), whereas the thin solid and dashed line pairs indicate the delay ranges where we expect most of the power of Cyg A and Cas A, respectively.

know from MK16 that the brightest hotspots (A and D) show a frequency turnover at around 140–150 MHz. To solve these issues, we made a new model directly from HBA observations (110–250 MHz), taking advantage of the high spatial resolution allowed by the Dutch LOFAR configuration (i.e. ≈ 5 arcsec) and using a second-order logarithmic polynomial to model the spectral behaviour of each component of the source (Sect. 2.1). We used the forced-spectrum method of WSCLEAN to transfer this spectral information into the model during a multi-scale and multi-frequency deconvolution (Ceccotti et al. 2023). This makes our new Cyg A model the first test case of such a method for improving 21-cm power spectrum results.

In Sect. 5, we compared the performances of the old model with our new one within the LOFAR-EoR processing pipeline (Sect. 4), using both simulated and observed data of the NCP field. Additionally, we considered an extreme Cyg A model, constituted by a single point source at the Cyg A location but with correct spectral information. On the simulated data, the most significant differences appear in $P(k)$ along the Cyg A direction at $k < 0.4 h cMpc^{-1}$, where the new model performs better than both the old and point source models. These differences are washed out in $\Delta^2(k)$, where all $(k_{\perp}, k_{\parallel})$ -cells are con-

sidered, resulting in $\Delta_{\text{new}} \approx \Delta_{\text{old}}$ across most k bins. For the real observed data, differences between the three models in both $P(k)$ along the Cyg A direction and $\Delta^2(k)$ are less pronounced at $k < 0.4 h cMpc^{-1}$, where uncertainties make their power spectra consistent with one another. Larger differences are observed at higher k , where the point source model results in significantly higher power in both the power spectrum estimators. The lowest residuals are achieved with the new model at $k > 0.6 h cMpc^{-1}$, where the old and point source models show an average excess of more than $(1250 \text{ mK})^2$ and $(2300 \text{ mK})^2$, respectively.

As we did not inject any systematics in the simulations, apart from adding the thermal noise, we expect all the differences to come from the different Cyg A models. In the real observations, where other systematics are present, these systematics have a larger impact on the residuals, hiding the effect of the different models, especially at the lower k where upper limits on the 21-cm power spectrum are usually extracted. We know that Cyg A is the strongest source at 150 MHz, but not in apparent flux for NCP observations. There are many other bright sources outside the main lobe of the primary beam that could leave subtraction residuals strong enough to outweigh those left from Cyg A, most notably Cas A, which has an apparent flux approximately three

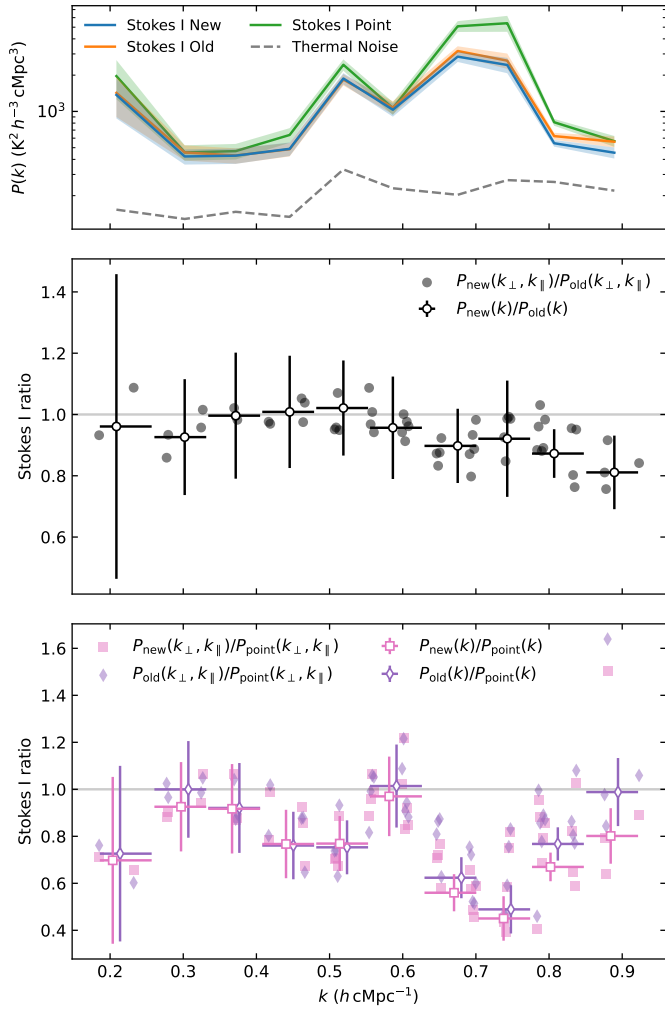


Fig. 12. Power spectra along the Cyg A direction after DD subtraction (top panel) and their ratios (middle and bottom panels) for the real observed NCP dataset. The top panel shows Stokes I $P(k)$ after the DD subtraction step using the new (blue), old (orange), and point source (green) models, along with the thermal noise level (dashed grey line). The shaded areas represent the 1σ uncertainties. The middle panel shows the ratio New/Point (pink squares), and Old/Point (purple diamonds). Filled markers represent the ratio of the cylindrical power spectra $P(k_{\perp}, k_{\parallel})$ for each $(k_{\perp}, k_{\parallel})$ -cell within the Cyg A delay lines, while the white-faced markers indicate the ratios of the power spectra $P(k)$. For the latter, the horizontal error bars indicate the k -bin extension, while the vertical error bars indicate the 1σ uncertainties of the ratio. A small offset in k has been added to the white-faced pink squares and purple diamonds in the bottom panel to avoid overlapping of error bars.

times higher than Cyg A. The DD subtracted plots in Fig. 11 show that some residuals associated with Cas A are present, highlighting the need of a better model for this source as well. However, as long as we focus our analysis at $k < 0.6 h cMpc^{-1}$, as in the analysis of real data, we do not expect significant improvements until the effects of other systematics are lowered.

In conclusion, we have demonstrated that the forced-spectrum method can generate spectrally accurate models of complex sources such as Cyg A, provided that physical spectral index maps are available or can be extracted from the data. This approach ensures that robust spectral information is embedded in a high-resolution model, which is fundamental for using such a

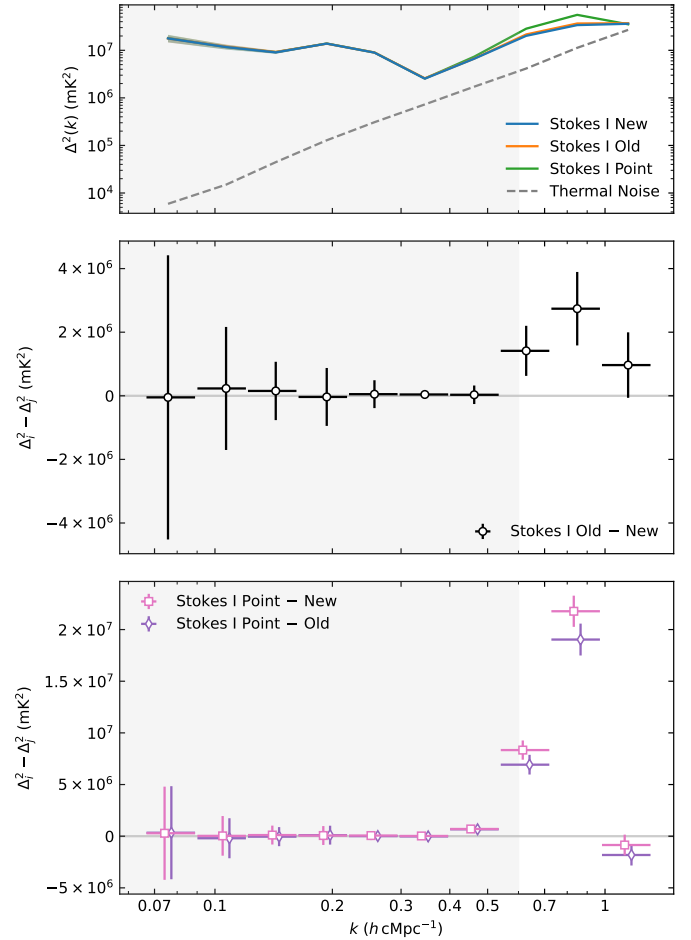


Fig. 13. Spherical power spectra after DD subtraction (top panel) and their difference (middle and bottom panels) for the real observed NCP dataset. The top panel shows Stokes I $\Delta^2(k)$ after the DD subtraction step using the new (blue), old (orange), and point source (green) Cyg A model, and the thermal noise level (dashed grey line). The shaded areas represent the 1σ uncertainties. The middle panel shows the difference $\Delta^2_{old} - \Delta^2_{new}$ (black dots), while the bottom panel shows the differences $\Delta^2_{point} - \Delta^2_{new}$ (pink squares), and $\Delta^2_{point} - \Delta^2_{old}$ (purple diamonds), with the associated 1σ uncertainties (values reported in Table D.1). The horizontal error bars indicate the k -bin extension. We added a small offset in k to the pink and purple markers in the bottom panel to avoid overlapping of error bars. The grey shaded area delimits the range in k where Δ^2 upper limits are usually extracted from LOFAR data, being the most sensitive range for LOFAR.

model at different frequencies from those at which the model was originally extracted. For example, we showed that while the old model has a correct flux density at 74 MHz, its incorrect spectral index results in a much lower extrapolated flux density at 150 MHz than the expected value. Moreover, the new model is described as a standard sky model, with a low number of point and Gaussian components with spectral information, and can therefore directly be used in calibration. The new Cyg A model presented in this work, and more generally, the forced-spectrum fitting method, could greatly benefit high-resolution and wide-field interferometers operating at low frequencies, such as the upcoming SKA-Low. This advancement will be beneficial for enhancing the accuracy and reliability of astronomical observations at low frequencies.

Because no GPR analysis was applied before power spectrum estimation, all the improvements that we measured in the 21-cm power spectrum come from the DD calibration, the only step where we used Cyg A. Alongside with assessing the impact of a better model on the power spectrum, we demonstrated that the calibration itself can benefit from a forced-spectrum model. We cannot make a one-to-one comparison with Mertens et al. (2020) about the excess power in the LOFAR power spectrum, because we did not apply a GPR technique. However, we can conclude that improving the modelling of Cyg A can reduce the excess power up to $(250 \text{ mK})^2$ over the $k = 0.08\text{--}0.63 \text{ h cMpc}^{-1}$ range, as measured in the real observations (see Table D.1). This reduction can be smaller after GPR, because some of the extra power that we observed would be absorbed in the foreground mode-mixing kernel, making the difference between the three models less strong. Our simulation scenario can be expected to be comparable to the results of a GPR analysis, because we have not included the systematics that GPR removes. Here, indeed, the measured differences are smaller, with $\langle \Delta_{\text{old}}^2 - \Delta_{\text{new}}^2 \rangle \approx (17 \text{ mK})^2$ over the $k = 0.08\text{--}0.63 \text{ h cMpc}^{-1}$ range. What we showed is that using both spatially and spectrally accurate models of bright off-axis sources is important to lower the upper limits on the 21-cm power spectrum.

In order to see a significant impact of a better spectral modelling in real data, other systematics must be reduced first. Because we flagged data after the DD subtraction, we expect even low-level RFI to be removed and they would therefore not be a dominant source of excess power. Another possible cause of the excess power might be ionospheric effects that have not been completely solved for in the initial calibration step. Although, Gan et al. (2022) and Brackenhoff et al. (2024) concluded that the ionosphere is currently not the dominant source of the excess power. In the analysis of Brackenhoff et al. (2024), it is shown that rapid spatial and temporal variations in the primary beam model could, however, result in incorrect DD calibration solutions. This is especially true for directions far from the field centre, where primary beam sidelobes move faster in time and have steep edges with deep nulls. Whereas we do not see incorrect DD gains, we see bright residuals associated with Cyg A even in the simulations where the calibration model is the same as the one used in the prediction step (first panel of Fig. 6). This could mean that the DD calibration gains are not able to capture the finest spatial, temporal, and spectral variations of the source due to the primary beam sidelobes. Understanding and characterising the time and frequency behaviour of the LOFAR HBA primary beam appears to be the next important step in trying to reduce the excess power in the NCP field. Processing and analysing a different sky region, such as the 3C 196 field, could be helpful as well in understanding the source of the systematics (Ceccotti et al., in prep.).

Data availability

The component list of the new model of Cygnus A is available at <https://doi.org/10.5281/zenodo.14863290>. The Cygnus A dataset (Project code: LC1_013) and both the observed and simulated NCP datasets (Observation ID: L246309) underlying this article will be shared upon reasonable request to the corresponding author.

Acknowledgements. EC, ARO and LVEK would like to acknowledge support from the Centre for Data Science and Systems Complexity (DSSC), Faculty of Science and Engineering at the University of Groningen. EC acknowledges support from the Ministry of Universities and Research (MUR) through the

PRIN project ‘Optimal inference from radio images of the epoch of reionization’. SAB, SG, LVEK and SM acknowledge the financial support from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 884760, ‘CoDEX’). FGM acknowledges support from a PSL Fellowship. RG acknowledges support from SERB, DST Ramanujan Fellowship no. RJF/2022/000141. The post-doctoral contract of IH was funded by Sorbonne Université in the framework of the Initiative Physique des Infinis (IDEX SUPER). LOFAR, the Low Frequency Array designed and constructed by ASTRON, has facilities in several countries, that are owned by various parties (each with their own funding sources), and that are collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. In this work, we made use of the KVIS (Gooch 1996) and DS9 (Joye & Mandel 2003) FITS file image viewers, and the ASTROPY (Astropy Collaboration 2022), MATPLOTLIB (Hunter 2007), NUMPY (Harris et al. 2020), PANDAS (McKinney 2010), SCIPY (Virtanen et al. 2020) PYTHON packages.

References

- Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *ApJ*, 925, 221
 Acharya, A., Mertens, F., Ciardi, B., et al. 2024, *MNRAS*, 527, 7835
 Asad, K. M. B., Koopmans, L. V. E., Jelić, V., et al. 2015, *MNRAS*, 451, 3709
 Asad, K. M. B., Koopmans, L. V. E., Jelić, V., et al. 2016, *MNRAS*, 462, 4482
 Astropy Collaboration (Price-Whelan, A. M., et al.) 2022, *ApJ*, 935, 167
 Baldwin, J. E., Boysen, R. C., Hales, S. E. G., et al. 1985, *MNRAS*, 217, 717
 Barry, N., Hazelton, B., Sullivan, I., Morales, M. F., & Pober, J. C. 2016, *MNRAS*, 461, 3135
 Baumgartner, W. H., Tueller, J., Markwardt, C. B., et al. 2013, *ApJS*, 207, 19
 Bernardi, G., de Bruyn, A. G., Brentjens, M. A., et al. 2009, *A&A*, 500, 965
 Bernardi, G., de Bruyn, A. G., Harker, G., et al. 2010, *A&A*, 522, A67
 Bianco, M., Giri, S. K., Prelogović, D., et al. 2024, *MNRAS*, 528, 5212
 Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2009, *ApJ*, 695, 183
 Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. 2011, *Foundations and Trends in Machine Learning*, 3, 1
 Brackenhoff, S. A., Mevius, M., Koopmans, L. V. E., et al. 2024, *MNRAS*, 533, 632
 Carilli, C. L., Perley, R. A., Dreher, J. W., & Leahy, J. P. 1991, *ApJ*, 383, 554
 Ceccotti, E., Offringa, A. R., Koopmans, L. V. E., et al. 2023, *MNRAS*, 525, 3946
 Charlot, P., Jacobs, C. S., Gordon, D., et al. 2020, *A&A*, 644, A159
 Cook, J. H., Trott, C. M., & Line, J. L. B. 2022, *MNRAS*, 514, 790
 Cornwell, T. J. 2008, *IEEE Journal of Selected Topics in Signal Processing*, 2, 793
 Datta, A., Bowman, J. D., & Carilli, C. L. 2010, *ApJ*, 724, 526
 DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, 129, 045001
 de Gasperin, F., Dijkema, T. J., Drabent, A., et al. 2019, *A&A*, 622, A5
 de Gasperin, F., Vink, J., McKean, J. P., et al. 2020, *A&A*, 635, A150
 Dewdney, P. E., Hall, P. J., Schilizzi, R. T., & Lazio, T. J. L. W. 2009, *IEEE Proc.*, 97, 1482
 Eastwood, M. W., Anderson, M. M., Monroe, R. M., et al. 2019, *AJ*, 158, 84
 Ewall-Wice, A., Dillon, J. S., Liu, A., & Hewitt, J. 2017, *MNRAS*, 470, 1849
 Farouki, R., & Rajan, V. 1988, *Computer Aided Geometric Design*, 5, 1
 Franzen, T. M. O., Vernstrom, T., Jackson, C. A., et al. 2019, *PASA*, 36, e004
 Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, 433, 181
 Gan, H., Koopmans, L. V. E., Mertens, F. G., et al. 2022, *A&A*, 663, A9
 Gan, H., Mertens, F. G., Koopmans, L. V. E., et al. 2023, *A&A*, 669, A20
 Gehlot, B. K., Mertens, F. G., Koopmans, L. V. E., et al. 2019, *MNRAS*, 488, 4271
 Gehlot, B. K., Mertens, F. G., Koopmans, L. V. E., et al. 2020, *MNRAS*, 499, 4158
 Gehlot, B. K., Koopmans, L. V. E., Offringa, A. R., et al. 2022, *A&A*, 662, A97
 Ghosh, A., Mertens, F., Bernardi, G., et al. 2020, *MNRAS*, 495, 2813
 Giri, S. K., Mellema, G., Dixon, K. L., & Iliev, I. T. 2018a, *MNRAS*, 473, 2949
 Giri, S. K., Mellema, G., & Ghara, R. 2018b, *MNRAS*, 479, 5596
 Gooch, R. 1996, *ASP Conf. Ser.*, 101, 80
 Gorce, A., Ganjam, S., Liu, A., et al. 2023, *MNRAS*, 520, 375
 Greenhill, L. J., & Bernardi, G. 2012, ArXiv e-prints [arXiv:1201.1700]
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357
 Hothi, I., Chapman, E., Pritchard, J. R., et al. 2021, *MNRAS*, 500, 2264
 Hunter, J. D. 2007, *Comput. Sci. Eng.*, 9, 90
 Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, 598, A78
 Jackson, C. 2005, *PASA*, 22, 36
 Jelić, V., Zaroubi, S., Labropoulos, P., et al. 2008, *MNRAS*, 389, 1319
 Jelić, V., de Bruyn, A. G., Pandey, V. N., et al. 2015, *A&A*, 583, A137
 Joye, W. A., & Mandel, E. 2003, *ASP Conf. Ser.*, 295, 489
 Kazemi, S., Yatawatta, S., Zaroubi, S., et al. 2011, *MNRAS*, 414, 1656

- Kern, N. S., Dillon, J. S., Parsons, A. R., et al. 2020, *ApJ*, **890**, 122
- Kerrigan, J. R., Pober, J. C., Ali, Z. S., et al. 2018, *ApJ*, **864**, 131
- Kolopanis, M., Jacobs, D. C., Cheng, C., et al. 2019, *ApJ*, **883**, 133
- Kolopanis, M., Pober, J. C., Jacobs, D. C., & McGraw, S. 2023, *MNRAS*, **521**, 5120
- Koopmans, L., Pritchard, J., Mellema, G., et al. 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 1
- Leahy, J. P., Muxlow, T. W. B., & Stephens, P. W. 1989, *MNRAS*, **239**, 401
- Li, W., Pober, J. C., Hazelton, B. J., et al. 2018, *ApJ*, **863**, 170
- Line, J. L. B., Mitchell, D. A., Pindor, B., et al. 2020, *PASA*, **37**, e027
- Liu, A., & Shaw, J. R. 2020, *PASP*, **132**, 062001
- Liu, A., Parsons, A. R., & Trott, C. M. 2014a, *Phys. Rev. D*, **90**, 023018
- Liu, A., Parsons, A. R., & Trott, C. M. 2014b, *Phys. Rev. D*, **90**, 023019
- Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. 2009, *IEEE Proc.*, **97**, 1497
- Lynch, C. R., Galvin, T. J., Line, J. L. B., et al. 2021, *PASA*, **38**, e057
- Mazumder, A., Datta, A., Chakraborty, A., & Majumdar, S. 2022, *MNRAS*, **515**, 4020
- McKean, J. P., Godfrey, L. E. H., Vegetti, S., et al. 2016, *MNRAS*, **463**, 3143
- McKinley, B., Tingay, S. J., Gaspari, M., et al. 2022, *Nat. Astron.*, **6**, 109
- McKinney, W. 2010, in *Proceedings of the 9th Python in Science Conference*, eds. S. van der Walt, & J. Millman, 56
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, *ASP Conf. Ser.*, **376**, 127
- Mertens, F. G., Ghosh, A., & Koopmans, L. V. E. 2018, *MNRAS*, **478**, 3640
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2020, *MNRAS*, **493**, 1662
- Mertens, F. G., Bobin, J., & Carucci, I. P. 2024, *MNRAS*, **527**, 3517
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2025, *A&A*, submitted [arXiv:2503.05576]
- Mevius, M., Mertens, F., Koopmans, L. V. E., et al. 2022, *MNRAS*, **509**, 3693
- Mitchell, D. A., Greenhill, L. J., Wayth, R. B., et al. 2008, *IEEE Journal of Selected Topics in Signal Processing*, **2**, 707
- Morales, M. F., & Hewitt, J. 2004, *ApJ*, **615**, 7
- Morales, M. F., & Wyithe, J. S. B. 2010, *ARA&A*, **48**, 127
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, **752**, 137
- Morales, M. F., Beardsley, A., Pober, J., et al. 2019, *MNRAS*, **483**, 2207
- Mouri Sardarabadi, A., & Koopmans, L. V. E. 2019, *MNRAS*, **483**, 5480
- Munshi, S., Mertens, F. G., Koopmans, L. V. E., et al. 2024, *A&A*, **681**, A62
- Munshi, S., Mertens, F. G., Koopmans, L. V. E., et al. 2025, *A&A*, **693**, A276
- Noordam, J. E. 2004, *SPIE Conf. Ser.*, **5489**, 817
- Offringa, A. R., & Smirnov, O. M. 2017, *MNRAS*, **471**, 301
- Offringa, A. R., McKinley, B., Hurley-Walker, N., et al. 2014, *MNRAS*, **444**, 606
- Offringa, A. R., Trott, C. M., Hurley-Walker, N., et al. 2016, *MNRAS*, **458**, 1057
- Offringa, A. R., Mertens, F., van der Tol, S., et al. 2019a, *A&A*, **631**, A12
- Offringa, A. R., Mertens, F., & Koopmans, L. V. E. 2019b, *MNRAS*, **484**, 2866
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, **433**, 639
- Pandey, V. N., Koopmans, L. V. E., Tiesinga, E., Albers, W., & Koers, H. U. A. 2020, *ASP Conf. Ser.*, **527**, 473
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *AJ*, **139**, 1468
- Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012, *ApJ*, **756**, 165
- Patil, A. H., Yatawatta, S., Zaroubi, S., et al. 2016, *MNRAS*, **463**, 4317
- Patil, A. H., Yatawatta, S., Koopmans, L. V. E., et al. 2017, *ApJ*, **838**, 65
- Pober, J. C., Hazelton, B. J., Beardsley, A. P., et al. 2016, *ApJ*, **819**, 8
- Pritchard, J. R., & Loeb, A. 2012, *Rep. Progr. Phys.*, **75**, 086901
- Procopio, P., Wayth, R. B., Line, J., et al. 2017, *PASA*, **34**, e033
- Rau, U., & Cornwell, T. J. 2011, *A&A*, **532**, A71
- Rau, U., Bhatnagar, S., & Owen, F. N. 2016, *AJ*, **152**, 124
- Refregier, A. 2003, *MNRAS*, **338**, 35
- Reich, P., & Reich, W. 1988, *A&AS*, **74**, 7
- Rich, J. W., de Blok, W. J. G., Cornwell, T. J., et al. 2008, *AJ*, **136**, 2897
- Santos, M. G., Cooray, A., & Knox, L. 2005, *ApJ*, **625**, 575
- Sault, R. J., & Wieringa, M. H. 1994, *A&AS*, **108**, 585
- Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. 1999, *A&A*, **345**, 380
- Shimwell, T. W., Röttgering, H. J. A., Best, P. N., et al. 2017, *A&A*, **598**, A104
- Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. 2000, *ApJ*, **530**, 133
- Thompson, A. R., Moran, J. M., Swenson, G. W. J. 2017, *Interferometry and 1480Synthesis in Radio Astronomy*, 3rd edn. (Springer)
- Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015, *ApJ*, **804**, 14
- Trott, C. M., Jordan, C. H., Midgley, S., et al. 2020, *MNRAS*, **493**, 4711
- Van der Tol, S. 2009, Ph.D. Thesis, Technical University of Delft, The Netherlands
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, **556**, A2
- Vedantham, H. K., & Koopmans, L. V. E. 2015, *MNRAS*, **453**, 925
- Vedantham, H. K., & Koopmans, L. V. E. 2016, *MNRAS*, **458**, 3099
- Vedantham, H., Udaya Shankar, N., & Subrahmanyam, R. 2012, *ApJ*, **745**, 176
- Vinyaikin, E. N. 2014, *Astron. Rep.*, **58**, 626
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Meth.*, **17**, 261
- Wilensky, M. J., Morales, M. F., Hazelton, B. J., et al. 2019, *PASP*, **131**, 114507
- Wilman, R. J., Miller, L., Jarvis, M. J., et al. 2008, *MNRAS*, **388**, 1335
- Yatawatta, S. 2015, *MNRAS*, **449**, 4506
- Yatawatta, S. 2016, ArXiv e-prints [arXiv:1605.09219]
- Yatawatta, S. 2018, ArXiv e-prints [arXiv:1805.00265]
- Yatawatta, S. 2019, *MNRAS*, **486**, 5646
- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, *A&A*, **550**, A136
- Yatawatta, S., Diblen, F., & Spreuw, H. 2017, *2017 IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 1
- Yoshiura, S., Pindor, B., Line, J. L. B., et al. 2021, *MNRAS*, **505**, 4775
- Zarka, P., Girard, J. N., Tagger, M., & Denis, L. 2012, in *SF2A-2012: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*, eds. S. Boissier, P. de Laverny, N. Nardetto, et al., 687

Appendix A: Old Cygnus A model components

In Table A.1 we report the components of the old Cyg A model used by the current LOFAR-EoR processing pipeline. The model catalogue is available at https://raw.githubusercontent.com/lofar-astron/prefactor/master/skymodels/Ateam_LBA_CC.skymodel.

Table A.1. Components of the old Cygnus A model obtained from VLA observations.

Name	Type	RA	Dec	I (kJy)	α	MajorAxis (arcsec)	MinorAxis (arcsec)	Orientation (deg)
CygA_4_1GG	GAUSSIAN	19 ^h 59 ^m 30 ^s :691	+40°43′54″:226	4.263	−0.8	18.21	11.69	69.14
CygA_4_2GG	GAUSSIAN	19 ^h 59 ^m 22 ^s :263	+40°44′27″:859	4.164	−0.8	20.15	87.32	103.82
CygA_4_0GG	GAUSSIAN	19 ^h 59 ^m 27 ^s :371	+40°44′06″:504	1.768	−0.8	23.13	69.14	78.24
CygA_4_5GG	GAUSSIAN	19 ^h 59 ^m 28 ^s :833	+40°44′02″:168	1.394	−0.8	23.60	0	36.24
CygA_4_7GG	POINT	19 ^h 59 ^m 30 ^s :024	+40°43′58″:388	1.309	−0.8	–	–	–
CygA_4_3GG	GAUSSIAN	19 ^h 59 ^m 24 ^s :073	+40°44′18″:226	1.226	−0.8	25.73	0	35.31
CygA_4_4GG	GAUSSIAN	19 ^h 59 ^m 24 ^s :980	+40°44′09″:455	0.865	−0.8	34.47	0	158.91
CygA_4_6GG	GAUSSIAN	19 ^h 59 ^m 25 ^s :564	+40°44′21″:868	0.700	−0.8	21.55	0	66.46
CygA_4_11GG	POINT	19 ^h 59 ^m 28 ^s :553	+40°43′43″:492	0.066	−0.8	–	–	–
CygA_4_8GG	POINT	19 ^h 59 ^m 20 ^s :591	+40°44′17″:513	0.057	−0.8	–	–	–
CygA_4_9GG	GAUSSIAN	19 ^h 59 ^m 23 ^s :597	+40°44′44″:562	0.051	−0.8	31.08	0	112.37
CygA_4_10GG	POINT	19 ^h 59 ^m 27 ^s :150	+40°43′49″:356	0.001	−0.8	–	–	–

Notes. The flux density I and the spectral index α are estimated at a reference frequency of 73.8 MHz.

Appendix B: Processing of the Cygnus A dataset

The initial dataset consisted of 399 and 203 SBs for HBA-low and HBA-high, respectively, with frequency and time resolutions of 48.8 kHz (i.e. 4 channels per SB) and 2 s. In the following, we will briefly describe the flagging, averaging, and calibration that have been previously performed using CASA (Common Astronomy Software Applications; McMullin et al. 2007).

Firstly, the dataset has been averaged to a single channel per SB (i.e. 195.3 kHz of frequency resolution) and to an integration time of 2 s. The SB at the edges of both bands are flagged because of strong RFI. Then, the calibration of the HBA-low data has been performed using the best model from McKean et al. (2016), solving both phase and amplitude for each channel (i.e. one solution every 195.3 kHz). The calibrated data has been imaged with the CASA multiscale multifrequency synthesis deconvolution, decomposing the spectral variation with two terms of the logarithmic polynomial function (see Eq. 1). This model is used to calibrate the HBA-high data, which are then concatenated to the HBA-low data.

The full dataset has been self-calibrated with a two Taylor term model obtained from the combined bands and finally flagged, resulting in a total of 412 SBs spanning 111–181 and 210–249 MHz.

Appendix C: Simple NCP sky model

In Table C.1, we report the central position and total flux density of the each of the 14 clusters that constitute the simple NCP sky model. We used this model for the simulated dataset.

Table C.1. Details of the NCP model used in the simulation step.

Cluster	RA	Dec	Flux density (Jy)
1	02 ^h 59 ^m 32 ^s .94	+86°25′28″.75	40.66
2	12 ^h 57 ^m 54 ^s .30	+86°38′13″.36	33.73
3	19 ^h 06 ^m 55 ^s .94	+86°42′41″.86	30.11
4	16 ^h 23 ^m 41 ^s .99	+86°30′09″.31	28.94
5	07 ^h 32 ^m 55 ^s .82	+86°36′33″.91	27.96
6	10 ^h 00 ^m 37 ^s .52	+86°23′21″.35	27.40
7	01 ^h 49 ^m 35 ^s .26	+88°25′47″.16	26.49
8	21 ^h 17 ^m 24 ^s .94	+86°20′27″.92	25.32
9	14 ^h 35 ^m 08 ^s .05	+88°22′37″.94	25.06
10	20 ^h 00 ^m 36 ^s .09	+88°47′30″.65	23.62
11	05 ^h 00 ^m 54 ^s .66	+87°05′35″.73	19.10
12	08 ^h 13 ^m 21 ^s .06	+88°40′41″.76	18.82
13	23 ^h 06 ^m 49 ^s .67	+87°04′30″.64	9.19
14	00 ^h 44 ^m 22 ^s .68	+86°22′36″.81	6.55

Appendix D: Final power spectrum differences

In Table D.1, we report the differences $D_{ij} = \Delta_i^2 - \Delta_j^2$ and the associated 1σ uncertainties evaluated using Eq. (9). Both the results from simulations (Sect. 5.2.2) and real observations (Sect. 5.3.2) are shown. In the last two rows, we also report the inverse-variance weighted mean of the differences D_{ij} over the k -bins ranges and the associated 1σ uncertainties, evaluated as

$$\langle D_{ij} \rangle = \frac{\sum_k (D_{ij}(k) \cdot D_{ij, \text{err}}^{-2}(k))}{\sum_k D_{ij, \text{err}}^{-2}(k)}, \quad \langle D_{ij} \rangle_{\text{err}} = \sqrt{\frac{1}{\sum_k D_{ij, \text{err}}^{-2}(k)}}. \quad (\text{D.1})$$

Table D.1. Differences of spherical power spectrum pairs for simulated and observed NCP data, with the associated 1σ error.

k (h cMpc ⁻¹)	Simulations				Observations	
	$\Delta_{\text{old}}^2 - \Delta_{\text{new}}^2$ (mK ²)	$\Delta_{\text{point}}^2 - \Delta_{\text{new}}^2$ (mK ²)	$\Delta_{\text{point}}^2 - \Delta_{\text{old}}^2$ (mK ²)	$\Delta_{\text{old}}^2 - \Delta_{\text{new}}^2$ (mK ²)	$\Delta_{\text{point}}^2 - \Delta_{\text{new}}^2$ (mK ²)	$\Delta_{\text{point}}^2 - \Delta_{\text{old}}^2$ (mK ²)
0.08	+(7.10) ² ± (41.99) ²	+(26.02) ² ± (42.94) ²	+(25.03) ² ± (43.01) ²	-(223.23) ² ± (2113.76) ²	+(535.38) ² ± (2123.78) ²	+(580.06) ² ± (2122.31) ²
0.11	-(5.17) ² ± (57.06) ²	+(23.31) ² ± (57.48) ²	+(23.88) ² ± (57.46) ²	+(480.41) ² ± (1390.86) ²	+(180.77) ² ± (1384.98) ²	-(445.10) ² ± (1391.81) ²
0.14	+(30.91) ² ± (79.87) ²	+(67.85) ² ± (81.24) ²	+(60.40) ² ± (81.57) ²	+(390.49) ² ± (957.66) ²	+(331.09) ² ± (956.53) ²	-(207.02) ² ± (960.52) ²
0.19	+(69.71) ² ± (115.49) ²	+(239.41) ² ± (124.51) ²	+(229.03) ² ± (125.15) ²	-(189.67) ² ± (954.30) ²	+(256.63) ² ± (956.06) ²	+(319.12) ² ± (955.44) ²
0.26	+(82.28) ² ± (122.58) ²	+(424.47) ² ± (142.62) ²	+(416.42) ² ± (143.09) ²	+(225.44) ² ± (661.87) ²	+(248.76) ² ± (662.08) ²	+(105.15) ² ± (663.02) ²
0.35	+(198.16) ² ± (175.76) ²	+(861.36) ² ± (230.58) ²	+(838.25) ² ± (231.88) ²	+(203.28) ² ± (346.58) ²	+(150.94) ² ± (345.95) ²	-(136.16) ² ± (347.35) ²
0.46	+(55.75) ² ± (242.08) ²	+(185.30) ² ± (243.67) ²	+(176.71) ² ± (243.82) ²	+(179.24) ² ± (539.60) ²	+(835.65) ² ± (553.24) ²	+(816.20) ² ± (553.84) ²
0.63	-(24.35) ² ± (349.32) ²	+(543.64) ² ± (358.42) ²	+(544.19) ² ± (358.40) ²	+(1188.90) ² ± (887.23) ²	+(2887.20) ² ± (964.72) ²	+(2631.06) ² ± (976.20) ²
0.85	+(61.80) ² ± (495.40) ²	+(619.05) ² ± (502.52) ²	+(615.96) ² ± (502.59) ²	+(1654.94) ² ± (1074.88) ²	+(4666.05) ² ± (1228.38) ²	+(4362.70) ² ± (1242.11) ²
1.15	-(41.87) ² ± (718.76) ²	+(118.13) ² ± (718.93) ²	+(125.33) ² ± (718.91) ²	+(982.77) ² ± (1014.69) ²	-(923.45) ² ± (1001.99) ²	-(1348.55) ² ± (1008.87) ²
0.08–0.63	+(17.38) ² ± (38.57) ²	+(56.30) ² ± (39.38) ²	+(54.20) ² ± (39.43) ²	+(256.98) ² ± (324.09) ²	+(463.29) ² ± (325.06) ²	+(391.35) ² ± (326.23) ²
0.63–1.15	+(13.94) ² ± (326.95) ²	+(547.32) ² ± (334.35) ²	+(547.11) ² ± (334.35) ²	+(1259.39) ² ± (741.59) ²	+(2663.22) ² ± (788.59) ²	+(2361.20) ² ± (796.34) ²

Notes. Differences were estimated using Eq. (9) at given k -bins. In the last two rows, we report the inverse-variance weighted mean of the differences over the indicated k -bins ranges, with the associated standard error, estimated using Eq. (D.1).