



Universiteit
Leiden
The Netherlands

A probabilistic approach to phase calibration - I. Effects of source structure on fringe-fitting

Natarajan, I.; Deane, R.; Bemmell, I. van; Langevelde, H.J. van; Small, D.; Kettenis, M.; ... ; Szomoru, A.

Citation

Natarajan, I., Deane, R., Bemmell, I. van, Langevelde, H. J. van, Small, D., Kettenis, M., ... Szomoru, A. (2020). A probabilistic approach to phase calibration - I. Effects of source structure on fringe-fitting. *Monthly Notices Of The Royal Astronomical Society*, 496(1), 801-813. doi:10.1093/mnras/staa1503

Version: Accepted Manuscript

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/3142198>

Note: To cite this publication please use the final published version (if applicable).

A probabilistic approach to phase calibration: I. Effects of source structure on fringe-fitting

Iniyan Natarajan¹★, Roger Deane^{2,1}, Ilse van Bemmél³, Huib Jan van Langevelde^{3,4}, Des Small³, Mark Kettenis³, Zsolt Paragi³, Oleg Smirnov^{1,5}, Arpad Szomoru³

¹Centre for Radio Astronomy Techniques and Technologies, Department of Physics and Electronics, Rhodes University, Grahamstown 6140, South Africa

²Department of Physics, University of Pretoria, Hatfield, Pretoria, 0028, South Africa

³Joint Institute for VLBI ERIC, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, Netherlands

⁴Sterrewacht Leiden, Leiden University, Postbus 9513, 2300 RA Leiden, Netherlands

⁵South African Radio Astronomy Observatory, Observatory 7925, Cape Town, South Africa

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We propose a probabilistic framework for performing simultaneous estimation of source structure and fringe-fitting parameters in Very Long Baseline Interferometry (VLBI) observations. As a first step, we demonstrate this technique through the analysis of synthetic short-duration Event Horizon Telescope (EHT) observations of various geometric source models at 230 GHz, in the presence of baseline-dependent thermal noise. We perform Bayesian parameter estimation and model selection between the different source models to obtain reliable uncertainty estimates and correlations between various source and fringe-fitting related model parameters. We also compare the Bayesian posteriors with those obtained using widely-used VLBI data reduction packages such as CASA and AIPS, by fringe-fitting 200 Monte Carlo simulations of each source model with different noise realisations, to obtain distributions of the Maximum A Posteriori (MAP) estimates. We find that, in the presence of resolved asymmetric source structure and a given array geometry, the traditional practice of fringe-fitting with a point source model yields appreciable offsets in the estimated phase residuals, potentially biasing or limiting the dynamic range of the starting model used for self-calibration. Simultaneously estimating the source structure earlier in the calibration process with formal uncertainties improves the precision and accuracy of fringe-fitting and establishes the potential of the available data especially when there is little prior information. We also note the potential applications of this method to astrometry and geodesy for specific science cases and the planned improvements to the computational performance and analyses of more complex source distributions.

Key words: techniques: interferometric – techniques: high angular resolution – methods: data analysis – methods: statistical

1 INTRODUCTION

Interferometry is concerned with estimating the spatial coherence function of electromagnetic fields measured between two different locations. A correlator must cross-correlate the signals from different stations while ensuring that they correspond to the same incoming wavefront instance. VLBI observations are similar to those made using connected-element interferometers insofar as the aim is to obtain this coherence function, but the measured visibility phases are less

well-behaved in VLBI due to a number of reasons. The atmospheric conditions at individual stations are different, leading to different atmospheric propagation delays that are uncorrelated (Thompson et al. 2017, Chapter 9). The use of independent frequency standards introduces another source of systematic uncertainty in determining the correct timestamps at which to cross-correlate the signals from different stations. Moreover, the rotation of the Earth causes the stations to move at different speeds with respect to the source, resulting in different Doppler shifts of the incoming wave, which introduces a time-variable delay (Taylor et al. 1999, Chapter 22). Finally, the presence of complex source struc-

★ E-mail: iniyannatarajan@gmail.com

ture also contributes to the variations in the phases of the measured visibilities.

A range of geometric, atmospheric, and instrumental effects are accounted for by the *correlator model*, a geometrical model that is applied to correct the visibility phases. This model is not perfect and contains residual errors in phases, that may vary appreciably with frequency and time. Various systematics such as errors in the source or antenna positions, offsets in the clock epoch, and errors in modelling the turbulence in the troposphere, especially at high frequency (Taylor et al. 1999, Chapter 28), contribute to these residual phases, and considerably reduce the time and frequency interval over which the measured visibilities can be averaged without a net loss of amplitude (*coherence time*), limiting the signal-to-noise ratio (SNR) of the data (Schwab & Cotton 1983). Correcting for these phase residuals and their slopes with frequency and time (known as *delay residuals* and *rate residuals* respectively) is known as *fringe-fitting* or *fringe-searching*, which allows one to average data over larger intervals of time and frequency, thereby increasing the coherence time and bandwidth, and improving the sensitivity.

Fringe-fitting is performed as one of the first steps after correlation and the resulting corrected visibilities are used as inputs for subsequent calibration procedures, such as self-calibration. Hence, studying any correlations between spatially-resolved source structure and fringe-fitting parameters and propagating reliable uncertainties further down to the data analysis steps where the astrophysical parameters are estimated is useful, especially since fringe-fitting is often performed with the assumption that the source is unresolved. Using the wrong source model during fringe-fitting can introduce false symmetry in the source structure, which may affect the results significantly when analysing strongly averaged data (e.g. Wielgus et al. 2019; Event Horizon Telescope Collaboration et al. 2019c). This may also be relevant to other science cases such as the asymmetric source structure observed in Global mm VLBI Array (GMVA) observations of Sgr A* (Issaoun et al. 2019), compact and extended structures observed in VLBI observations of gravitationally lensed radio sources (Spingola et al. 2019), and highly accurate source position estimates required for performing VLBI astrometry for the Bulge Asymmetries and Dynamic Evolution (BAaDE) targets (van Langevelde et al. 2019).

Incorporating a priori knowledge of the relevant parameters for self-calibration was first advocated by Cornwell & Wilkinson (1981) and a similar approach for fringe-fitting was advocated by Schwab & Cotton (1983). Reliable uncertainty propagation for calibration parameters has been discussed within the framework of Information Field Theory (IFT) by Enßlin et al. (2009, 2014); Enßlin (2018), and implemented in algorithms such as RESOLVE (Junklewitz et al. 2016). Lochner et al. (2015) discuss simultaneous estimation of source and instrumental parameters and model selection between partially-resolved source structures on Westerbork Synthesis Radio Telescope (WSRT) simulations within a Bayesian framework. In Natarajan et al. (2017), we undertook a full Bayesian analysis of European VLBI Network (EVN) observations of the blazar J0809+5341 to prove the existence of extended jet structure in the presence of residual station gain amplitudes after self-calibration.

In this first of an intended series of papers on the ap-

plication of probabilistic techniques to the wider problem of phase calibration in VLBI, we study the mutual effects of source structure and some fringe-fitting related parameters on each other, using short duration synthetic EHT observations with sparse uv-coverage¹. Of the parameters relevant to fringe-fitting, we simulate the phase and delay residuals; in actual observations, the phases also vary with time (*rate residuals*) and must be accounted for while fringe-fitting. However, in this proof-of-concept presentation of simultaneous source structure estimation and fringe-fitting, we perform these experiments only with simulated phase and delay residuals, since rate residuals are not significant for short duration snapshot observations such as these. We also assume that accurate a priori amplitude calibration has been performed on the data and do not consider gain amplitudes. Finally, we consider the presence of only one spectral window (subband) in this paper.

This work is organised as follows. In Section 2, we briefly discuss the theory behind fringe-fitting approaches used in current software packages. In Section 3, we present a probabilistic fringe-fitting formalism that simultaneously estimates the source-related parameters. In Sections 4 and 5, we apply this method to synthetic snapshot EHT observations of partially-and-fully-resolved symmetric and asymmetric geometric source models, and compare our results with those of other calibration packages such as CASA² and AIPS³. Finally, we outline ongoing efforts to update this framework to handle an expanded parameter space involving more complex source models and time-varying phase, delay, and rate residuals simultaneously, and the projects that we subsequently propose to undertake in Section 6.

2 FRINGE-FITTING

2.1 Terminology

There appears to be confusion in the relevant literature about the nomenclature used to describe the quantities that are estimated during fringe-fitting. Throughout this paper, we consistently refer to the first order phase residual terms (ψ_{0p}) as *phase residuals*; the derivative of the phase with respect to frequency (τ_p) as *delay residual* and the derivative of the phase with respect to time (r_p) as *rate residual*.

2.2 Background

The significant phase variations present in VLBI data lead to the loss of amplitude of the measured visibilities when averaged (*decoherence*), which fringe-fitting aims to minimise. These variations, to first order, manifest as residual phases, and phase slopes with frequency and time for ground-based VLBI arrays. Let $\tilde{V}_{pq}(t_m, \nu_n)$ be the measured visibility corresponding to baseline pq at time t_m and frequency ν_n . This is related to the true visibility $V_{pq}(t_m, \nu_n)$ as

$$\tilde{V}_{pq}(t_m, \nu_n) = g_p(t_m, \nu_n) \tilde{g}_q(t_m, \nu_n) V_{pq}(t_m, \nu_n) + \varepsilon_{pqmn} \quad , \quad (1)$$

¹ Throughout this paper, we use the terms *simulated* and *synthetic* interchangeably, as both are used in the relevant VLBI literature (e.g. Event Horizon Telescope Collaboration et al. 2019d).

² www.casa.nrao.edu.

³ www.aips.nrao.edu.

where ε_{pqmn} is the additive noise term arising due to thermal noise; g_p are complex-valued functions that incorporate antenna-based effects and can be expressed as

$$g_p(t, \nu) = |g_p(t, \nu)| e^{i\psi_p(t, \nu)} \quad (2)$$

Assuming that the amplitudes $|g_p|$ and $|V_{pq}|$ vary slowly enough that they are constant over the time and frequency averaging intervals, we can write to first-order (Schwab & Cotton 1983)

$$\begin{aligned} \tilde{V}_{pq}(t_m, \nu_n) \simeq & |g_p| |g_q| |V_{pq}(t_0, \nu_0)| \exp(i[(\psi_p - \psi_q)(t_0, \nu_0)]) \\ & \times \exp \left(i \left[\frac{\partial(\psi_p - \psi_q + \phi_{pq})}{\partial t} \Big|_{(t_0, \nu_0)} (t_m - t_0) \right. \right. \\ & \left. \left. + \frac{\partial(\psi_p - \psi_q + \phi_{pq})}{\partial \nu} \Big|_{(t_0, \nu_0)} (\nu_n - \nu_0) \right] \right) \quad (3) \end{aligned}$$

where $\phi_{pq} \equiv \arg V_{pq}$. Here t_0 and ν_0 are the time and frequency relative to which the calculations are carried out. It must be noted that using a fixed reference frequency across the band is valid only for moderate bandwidths as a first-order approximation, which will need to be revised for wide-band fringe-fitting. Here, we adopt this formalism from Schwab & Cotton (1983) which forms the basis of fringe-fitting in both CASA and AIPS. The derivative of phase with respect to time

$$r_{pq} \equiv \frac{\partial(\psi_p - \psi_q + \phi_{pq})}{\partial t} \Big|_{(t_0, \nu_0)} \quad (4)$$

is the expression for the rate residual, and the derivative of phase with respect to frequency

$$\tau_{pq} \equiv \frac{\partial(\psi_p - \psi_q + \phi_{pq})}{\partial \nu} \Big|_{(t_0, \nu_0)} \quad (5)$$

quantifies the delay residual, estimated relative to t_0 and ν_0 . These values are estimated by Fourier-transforming the per-baseline visibilities from the time and frequency domain to the delay and rate residual domain and finding its maximum (which would occur at r_{pq} and τ_{pq}). The position of this peak is used to define the phase centre of the observation, a fixed point in the sky relative to which the delay and rate residuals are estimated. The value of the function at this maximum (dependent on t and ν), gives the phase-corrected estimates of the visibility function in the time-frequency domain (Cotton 1995). It must be remembered that these estimates are only approximations to the fringe solutions, since the fringe-rate correction is constant over all frequencies. These can then be averaged coherently over time and frequency, to the extent that the first-order model in equation (3) is valid. It is important to note that the derivatives of the antenna phases ψ and those of the true visibility phases ϕ cannot be separated by this method.

Global fringe-fitting is performed to separate the antenna-based components in equation (3) and estimate them simultaneously for all stations using the data corresponding to all baselines. This is especially useful in cases where the SNR is low, which is often the case with VLBI observations. If the source is resolved, then a model of the source V_{pq}^M approximating the true visibilities V_{pq} is needed to enable this separation. If the source is (a) sufficiently compact so the 2-D FFT to the delay and rate residual domain has a well-defined peak and (b) the visibility function does not change significantly during the integration period (Alef

& Porcas 1986), then the default assumption of the source being point-like is made. For more elaborate source models, in theory, an initial image of the source can be used to enable the separation of baseline-based and antenna-based phase components (Schwab & Cotton 1983). Global fringe-fitting is related to phase self-calibration (Readhead & Wilkinson 1978; Readhead et al. 1980) in that both procedures aim to minimise the difference between the observed and model visibility phases. In addition to phase residuals, fringe-fitting also estimates their time and frequency slopes, and is always performed before averaging in order to preserve coherence.

Contemporary fringe-fitting algorithms such as the one employed by the CASA task `fringefit`⁴, perform baseline-based fringe-fitting before globalising the solutions to all baselines through a least-squares method (e.g. Janssen et al. 2019). This can be done, for instance, by performing baseline-based fringe-fitting with a subset of the more sensitive stations in the array, and letting the result serve as a starting point to a later global fringe-fitting step. A baseline-based approach that finds the maximum in a three-dimensional space of multi-band delay, single-band delay, and rate, is employed by the `fourfit` task (Cappallo 2017) in the Haystack Observatory Postprocessing System (HOPS)⁵ (Whitney et al. 2004). A pipeline based on HOPS that performs additional global fringe-fitting for delay and rate residuals (Blackburn et al. 2019b) is one of the primary fringe-fitting tools used to process data from the EHT, along with CASA and AIPS (Event Horizon Telescope Collaboration et al. 2019c), the tools we compare our results with in this work.

3 A RIME-BASED PROBABILISTIC APPROACH

We adopt a fully Bayesian approach to simultaneously estimate the source structure and the phase and delay residuals. This approach has the significant advantage that the prior assumptions (or biases) are quantified and propagated into the final visibility or image-plane analysis (e.g. Nataraajan et al. 2017). A natural outcome of this process is that we obtain reliable uncertainty estimates of and degeneracies between the various model parameters. A *model* or a *hypothesis* in our context, includes both source and instrumental parameters.

We use the Radio Interferometry Measurement Equation (RIME) formalism for modelling visibilities, which was originally developed for radio polarimetry by Hamaker et al. (1996) and extended to incorporate direction-dependent effects by Smirnov (2011). The RIME expresses the relationship between the source and instrumental characteristics in the form of 2×2 complex matrices known as *Jones* matrices (Jones 1941), which enables us to model the desired source properties and propagation path effects in parametric form. The generic RIME for a sky composed of multiple discrete sources is given by

$$V_{pq} = \sum_s J_{sp} B_s J_{sq}^H \quad (6)$$

⁴ https://casa.nrao.edu/casadocs-devel/stable/global-task-list/task_fringefit/about.

⁵ www.haystack.mit.edu/tech/vlbi/hops.html.

where V_{pq} denotes the visibility corresponding to baseline pq , J_{sp} represents the 2×2 cumulative Jones matrix corresponding to antenna p in the direction of source s , and B_s is the brightness matrix corresponding to the full polarisation output of the correlator. Each propagation path effect may be assigned its own Jones matrix, with the unknown or inseparable effects being subsumed into a generic Jones term.

The RIME for fringe-fitting is constructed as follows. In our analyses for this paper, we simulate only unpolarised sources with flat spectra. Hence, the brightness matrix for a source with flux density S_ν at frequency ν is given by

$$\mathbf{B} = \begin{pmatrix} S_\nu & 0 \\ 0 & S_\nu \end{pmatrix}. \quad (7)$$

This signal undergoes a linear transformation represented by the phase delay matrix K , representing the phase difference (κ_p) between the waves received by antenna p located at $\mathbf{u}_p = (u_p, v_p, w_p)$ relative to $\mathbf{u} = 0$:

$$K_p = e^{-i\kappa_p} \equiv e^{-i\kappa_p} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (8)$$

Since the K-Jones term exists even under ideal conditions with no propagation path effects, we define the *source coherency* matrix for convenience as (Smirnov 2011)

$$\mathbf{X}_{pq} = K_p \mathbf{B} K_q^H = \mathbf{B} e^{-i\kappa_{pq}}. \quad (9)$$

The E-Jones terms for primary beams are taken to be unity ($E_{sp} \equiv 1$) owing to the small FoVs often used for VLBI data analysis⁶. We construct a scalar phase-matrix G that affects both polarisation axes equally:

$$G_p = |g_p| \exp(i[\Delta\psi_p(t, \nu)]) \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (10)$$

where $\Delta\psi_p$ incorporates two independent terms that affect the visibility phases: the station-based phase residuals ψ_{0p} and delay residuals τ_p , relative to a reference antenna and frequency ν_{ref} :

$$\begin{aligned} \Delta\psi_p(t_m, \nu_n) &= \psi_{0p} + \frac{\partial\psi_p}{\partial\nu}(\nu_n - \nu_{\text{ref}}), \\ &= \psi_{0p} + \tau_p(\nu_n - \nu_{\text{ref}}). \end{aligned} \quad (11)$$

To keep the execution times (on a small compute server) manageable, for the purposes of this paper we restrict the number of model parameters relevant to fringe fitting, and simulate only the phase and delay residuals. More information on how we propose to expand this framework can be found in Section 6. Including the above terms, equation (6) becomes:

$$V_{pq}(t_m, \nu_n) = G_p(t_m, \nu_n) \left(\sum_s X_{spq}(t_m, \nu_n) \right) G_q^H(t_m, \nu_n). \quad (12)$$

For fringe-fitting, we are interested only in the phases of the G_p terms and set the gain amplitudes to unity ($|g_p| \equiv 1$). In practice, these terms are always present and uncertainties

in their estimates also affect the estimated source structure (e.g. Natarajan et al. 2017).

Having built this formalism for forward-modelling, we perform Bayesian inference on the synthetic data. Assuming a model (or hypothesis) H to be true, we estimate its parameters, Θ , by fitting them to the visibilities (data, D)⁷:

$$\mathcal{P}(\Theta|D, H) = \frac{\mathcal{P}(\Theta|H) \mathcal{P}(D|\Theta, H)}{\mathcal{P}(D|H)}, \quad (13)$$

where $\mathcal{P}(\Theta|H)$ is called the *prior* probability distribution, which encodes our beliefs about the parameters prior to the analysis of the data. $\mathcal{P}(\Theta|D, H)$ is the *posterior* probability distribution which describes how the data D modify our initial beliefs. $\mathcal{P}(D|\Theta, H) \equiv \mathcal{L}(\Theta|D, H)$ is the *likelihood*, which reflects how the uncertainties in the measurement are distributed. The denominator $\mathcal{P}(D|H)$ is a normalising constant called the *Bayesian evidence* or the *marginal likelihood* that becomes a valuable tool for ranking models while performing model comparison. Assuming Gaussian noise, which is a good approximation in the high SNR regime (e.g. Event Horizon Telescope Collaboration et al. 2019f), given the observed (V_D) and the modelled (V_M) visibilities (refer equation 12), and the uncertainties σ_k that vary with baseline, the likelihood function for parameter estimation for model H may be written as

$$\begin{aligned} \mathcal{L}(\Theta|V_D, H) &= \frac{1}{\prod_{k=1}^{2N_{\text{vis}}} \sqrt{2\pi\sigma_k^2}} \exp\left(-\frac{\chi^2}{2}\right), \\ \text{where } \chi^2 &= \sum_{k=1}^{2N_{\text{vis}}} \left(\frac{V_{M_k} - V_{D_k}}{\sigma_k} \right)^2, \end{aligned} \quad (14)$$

and N_{vis} is the total number of complex visibilities. The natural logarithm of $\mathcal{L}$, given by

$$\begin{aligned} \ln(\mathcal{L}) &= \sum_{k=1}^{2N_{\text{vis}}} \ln \left[(2\pi\sigma_k^2)^{-1/2} \right] - \frac{\chi^2}{2} \\ &= -\frac{1}{2} \sum_{k=1}^{2N_{\text{vis}}} \ln \left[2\pi\sigma_k^2 \right] - \frac{\chi^2}{2}, \end{aligned} \quad (15)$$

is often used in practice since it is more convenient to work with. The baseline-dependent per-visibility noise term for one polarisation is computed from the System Equivalent Flux Densities (SEFDs) of the individual stations using the radiometer equation (Thompson et al. 2017):

$$\sigma_{pq} = \frac{\text{SEFD}_{pq}}{\sqrt{2\delta\nu t_{pq}}}, \quad (16)$$

where $\text{SEFD}_{pq} = \sqrt{\text{SEFD}_p \times \text{SEFD}_q}$,

SEFD_p is the SEFD of station p , $\delta\nu$ is the channel bandwidth, and t_{pq} is the integration time for baseline pq .

The second level of inference we perform is model selection between different hypotheses. We use various geometric

⁶ For a discussion of the effects of E-Jones terms on EHT observations in the presence of antenna pointing errors, see Blecher et al. (2017).

⁷ A more detailed introduction to Bayesian inference may be found in Natarajan et al. (2017, and references therein). Here we give a short review relevant to our analysis.

Table 1. Criteria for model selection. B_{12} denotes the ratio of the evidences between hypotheses H_1 and H_2 (Kass & Raftery 1995), which measures the relative success of the two models at predicting the data.

$2\ln(B_{12})$	B_{12}	Evidence against H_2
0 to 2	1 to 3	Not worth more than a mention
2 to 6	3 to 20	Positive
6 to 10	20 to 150	Strong
> 10	> 150	Very strong

source models in our analyses, and always compare the results with those obtained with a point source model. Given hypothesis H , and a prior belief in the validity of H given by $\mathcal{P}(H|I)$, where I is any relevant background information, the model posterior probability may be computed using the evidence obtained from parameter estimation as

$$\mathcal{P}(H|D, I) \propto \mathcal{P}(D|H, I) \mathcal{P}(H|I) \quad (17)$$

Given two models H_1 and H_2 , we may define a model selection ratio between the posteriors of the two models as

$$\frac{\mathcal{P}(H_1|D, I)}{\mathcal{P}(H_2|D, I)} = \frac{\mathcal{Z}_1}{\mathcal{Z}_2} \frac{\mathcal{P}(H_1|I)}{\mathcal{P}(H_2|I)} = B_{12} \frac{\mathcal{P}(H_1|I)}{\mathcal{P}(H_2|I)}, \quad (18)$$

where $\mathcal{P}(H_1|I)/\mathcal{P}(H_2|I)$ is the ratio of the priors of the two models which we set to unity, indicating that there is no prior preference for one model over the other. The ratio of the evidences, B_{12} , known as the *Bayes factor* (Jeffreys 1961) then provides the odds in favour of H_1 ; the larger the value of B_{12} , the more is H_1 preferred over H_2 . Hence, Bayes factors between two models provide a more comprehensive metric for model comparison than traditional methods. Following Kass & Raftery (1995), we use twice the natural logarithm of this factor as a measure of how strongly one model is preferred over another (Table 1). An order of magnitude more support by the data in favour of H_1 is required to move up a level on this scale (Trotta 2008).

3.1 Software setup

We have developed a software package called ZAGROS⁸ to perform simultaneous source parameter estimation and fringe-fitting. ZAGROS uses CODEX-AFRICANUS⁹, a GPU-based forward-modelling software package (Perkins et al in prep), that speeds up model computation required for every iteration of the likelihood evaluation. ZAGROS can currently handle small datasets that can fit into the memory of an NVIDIA Tesla K20m GPU.

For sampling from the multi-dimensional posterior distribution we use DYPOLYCHORD, an implementation of the dynamic nested sampling algorithm (Higson et al. 2019), based on the publicly available POLYCHORD tool (Handley et al. 2015a,b). For generating synthetic EHT observations, we use MEQSILHOUETTE, a mm-VLBI synthetic data generation package capable of generating model visibilities from

parametric and non-parametric sky models and adding various propagation path effects such as tropospheric phase corruption, variable receiver gains, antenna pointing errors, and polarisation leakage (Blecher et al. 2017, Natarajan et al in prep).

4 FRINGE-FITTING AND SOURCE STRUCTURE

The EHT is a network of mm/sub-mm facilities spread across continents to create a telescope with high angular resolution ($\simeq 30\text{--}10\mu\text{as}$), with the longest baselines spanning the Earth’s diameter (Event Horizon Telescope Collaboration et al. 2019b). The primary goal of the EHT is to image the gravitationally-lensed photon ‘ring’ or *shadow* around the event horizons of the supermassive black holes at the centres of the Milky Way (Sgr A*) and the supergiant elliptical galaxy M87, which have the largest predicted apparent angular diameters that are resolvable by the EHT at 230 GHz (e.g. Broderick & Loeb 2009; Falcke & Markoff 2013). In the observing run of April 2017, the EHT imaged the black hole shadow of M87 within an asymmetric ring of size $42 \pm 3 \mu\text{as}$ (Event Horizon Telescope Collaboration et al. 2019d,f). At this frequency, the troposphere gives rise to a turbulent component to the delays, which is the major contributor to the decoherence of the visibilities. As the frequency increases, so does the tropospheric absorption, mainly due to the pressure-broadened transition lines of H_2O and O_2 (Carilli & Holdaway 1999). Unlike other chemical components, water vapour mixes poorly in the atmosphere, introducing rapid fluctuations in the measured visibility phases.

Under such difficult observing conditions, it is instructive to study the effects of partially or fully-resolved source structure on fringe-fitting. This is also relevant for self-calibration downstream, as the effects of the early decisions on the initial models will be impossible to remove further down the data analysis process. This is especially true for arrays with few stations such as EHT (given the small mutual visibility windows between stations, even fewer are observing simultaneously). The critical point in this paper is that the probabilistic approach has the advantage that it can simultaneously estimate the parameters related to source structure and fringe-fitting, and reveal whether they are degenerate, as well as enabling model selection via the Bayesian evidence. This not only minimises the effects that arise from using incorrect source models, but also enables the propagation of phase uncertainties further down in the calibration process. This may not always yield an appreciable difference in the final calibrated data, but may prove highly useful for high-value science targets such as the ones mentioned in Section 1.

With the recent introduction of the fringe-fitting task *fringeFit*, CASA is being adopted for VLBI data processing (e.g. Janssen et al. 2019; van Bemmelen et al. 2019) and hence serves as a good tool to compare our results with. We also perform some comparisons with the AIPS task *FRINGE*¹⁰

⁸ <https://github.com/saiyanprince/zagros>.

⁹ <https://github.com/ska-sa/codex-africanus>.

¹⁰ <http://www.aips.nrao.edu/cgi-bin/ZXHLP2.PL?FRINGE>.

Table 2. EHT stations participating in the mock-observations.

Station	Diameter (m)	Nominal SEFD (Jy)
SMA (SM)	25	6000
SMT (AZ)	30	1300
LMT (LM)	32	560
ALMA (AA)	25	220
JCMT (JC)	15	5000
SPT (SP)	10	1600
APEX (AP)	12	4500

(for the elliptical Gaussian model), which has traditionally been used for fringe-fitting. Since Bayesian methods are intrinsically concerned with probability distributions, we want to compare them to *distributions* of parameter estimates output by conventional methods; since the latter generate only MAP estimates (albeit with bounds), we run them over a set of 200 Monte Carlo simulations of each sky model in the study, each with a different noise realisation, and thus obtain these distributions.

4.1 Simulations of geometric source models

We consider the following source models in our simulations: a point source at the centre (PT), a circular Gaussian source (CIRC), an elliptical Gaussian source (ELLIP), and two point sources (2PT). To each source model, phase and delay residuals and baseline-dependent Gaussian thermal noise are introduced. The EHT stations used for these simulations are the Submillimeter Array (SMA) and James Clerk Maxwell Telescope (JCMT) in Hawai'i, Submillimeter Telescope (SMT) in Arizona, Large Millimeter Telescope (LMT) in Mexico, the Atacama Large Millimeter Array (ALMA) and the Atacama Pathfinder Experiment (APEX) in Chile, and the South Pole Telescope (SPT) (Table 2). The phase centre of the observations coincides with the coordinates of Sgr A*, $\alpha_{J2000} = 17^{\text{h}}45^{\text{m}}40^{\text{s}}.04088$, $\delta_{J2000} = -29^{\circ}0'28''.118$. The mock-observations were conducted for 3 minutes with a 2 s integration time, at a frequency of 230 GHz, with a 2.56 GHz bandwidth divided into 32 channels. The start time of the observation was selected for maximum mutual visibility between the stations.

The parametrisation of each model evaluated is shown in Table 3 and the hyperparameters used in the analyses are shown in Table 4. In all the models, the position (l, m) of the central source was fixed at the phase centre. The position of the secondary source in 2PT (l_2, m_2) is given a uniform prior overlapping with the position of the primary source. CIRC and ELLIP use the same parameters to describe the shape, with CIRC imposing the additional constraints $e_{\min} = e_{\text{maj}}$ and $\text{PA} = 0$. The position angle PA is set to vary over a range of 180 degrees, to avoid degeneracies between position angles oriented in opposite directions. LM was used as the reference station and hence its phase and delay residuals were set to zero. The centre frequency of the band was chosen as the reference frequency.

Table 3. Models evaluated in this study. Each model is composed of source-related parameters and the station-based phase and delay residuals. The degrees of freedom (DoF) indicates the number of parameters that are free to vary independently (excludes parameters with delta priors, such as the phase and delay residuals corresponding to reference stations; refer text).

Model	Parameters/DoF	Parametrisation
CIRC	24/14	Flux Density (S_{ν})
		Position (l, m)
		Shape ($e_{\text{maj}}, e_{\text{min}}, \text{PA}$)
		Phase residuals (ψ_{0p})
		Delay residuals (τ_p)
ELLIP	24/16	Flux Density (S_{ν})
		Position (l, m)
		Shape ($e_{\text{maj}}, e_{\text{min}}, \text{PA}$)
		Phase residuals (ψ_{0p})
		Delay residuals (τ_p)
2PT	24/16	Flux Density ($S_{\nu 1}, S_{\nu 2}$)
		Position ($(l_1, m_1), (l_2, m_2)$)
		Phase residuals (ψ_{0p})
		Delay residuals (τ_p)

Table 4. Prior distributions for all the model parameters. All parameters were set uniform priors with the range indicated by the values in the square brackets. For parameters with delta priors, refer text.

Parameter (Units)	Prior distribution
S_{ν} (Jy)	[0, 2]
l_2 (μas)	[-50, 50]
m_2 (μas)	[-10, 110]
e_{maj} (μas)	[0, 40]
e_{min} (μas)	[0, 40]
PA ($^{\circ}$)	[0, 180]
ψ_{0p} ($^{\circ}$), where $p \neq \text{LM}$	[0, 360]
τ_p (ps), where $p \neq \text{LM}$	[-200, 200]

4.1.1 Point Source and Circular Gaussian (PT & CIRC)

We simulate 21 datasets, each with a central circular Gaussian of total flux density 1 Jy, with increasing HPBW from 0 to 20 μas , up to about half the size of the point spread function (PSF) of the array. In all ZAGROS runs, we use 400 live points and 100 points for the initial exploratory run of DY-POLYCHORD, via the `nlive_const` and `ninit` parameters respectively (Higson et al. 2019). To perform Bayesian model selection on each dataset, we compare the evidences for a circular Gaussian model (CIRC) with that of a point source model (PT). In each case, the correct model is favoured (PT for the 0 μas source and CIRC for the rest) with *very strong* Bayesian evidence (Table 1). As the source gets more resolved, the odds ratio in favour of CIRC increases from 23 : 1 (for 1 μas Gaussian) to 10^{12} : 1 (for 20 μas Gaussian), indicating that the correct model is very strongly preferred. The error in relative natural-logarithmic evidence (the quantity defined in Table 1) in each case is ± 0.7 . It must be noted that, in actual observations, the minimum resolvable source size depends on the SNR, array configuration and calibration

uncertainties, and hence additional calibration steps will be necessary (Martí-Vidal et al. 2012; Natarajan et al. 2017).

Fig. 1 shows the results of ZAGROS fringe-fitting on the synthetic data with a 20 μas circular Gaussian source at the centre. We see from the 1-D marginalised posterior plots on the main diagonal, that all the model parameters are estimated to a high precision and accuracy. The 2-D joint posteriors show that there are no significant degeneracies between the source and instrumental parameters in this case, apart from a slight correlation between the delay residual estimates for AA with AP, likely due to their proximity.

As the radius of the simulated circular Gaussian source increases, the peak flux density goes down which results in decreased SNR of the data. This reduction in SNR contributes to the widening of the posteriors for larger source sizes. Figs. 2 and 3 demonstrate this effect for the estimated phase and delay residuals respectively for synthetic data generated with increasingly resolved CIRC models. In addition, the larger sources are resolved by the longer baselines and the corresponding visibility amplitudes decrease. This effect is demonstrated in Fig. 4 which plots the posteriors of phase and delay residuals for all antennas for the 20 μas CIRC model. The posteriors corresponding to SM, JC, and SP that contribute to the longest baselines are wider than those of the other stations whose baselines are relatively shorter. AA has the narrowest posteriors due to its high sensitivity.

To obtain the distribution of CASA fringe-fit results for comparison with ZAGROS, we simulate 200 instances of each dataset with independent noise realisations, with the same rms σ_{pq} in equation (16). Following this, we fringe-fit each dataset twice, with and without incorporating the exact source model in the process, to obtain the best and worst possible estimates from CASA. Figs. 5 and 6 show the results of this process. Only those solutions with $\text{SNR} \geq 3$ are included for global least-squares minimisation via the `minsnr` parameter. The phase residuals obtained with `fringe-fit` are re-referenced to the central frequency of the band. For this symmetric source structure, we see that the point source assumption during fringe-fitting is sufficient and the histograms of the `fringe-fit` results coincide well with the Bayesian posteriors.

4.1.2 Elliptical Gaussian (ELLIP)

We next perform fringe-fitting on elliptical Gaussian models (ELLIP) of size $25 \times 5 \mu\text{as}$ oriented at a position angle of 60° . Model selection between PT and ELLIP on these data results in a Bayes factor on the order of 10^{12} (with the error in relative evidence being ± 0.66), indicating a very high preference for the elliptical Gaussian.

The posterior distributions of the parameters corresponding to ELLIP are shown in Fig. 7. Here too, we see that the Bayesian framework estimates the source parameters and the phase and delay residuals accurately, as shown by the vertical green lines indicating the true values. There are no significant degeneracies between the source and instrumental parameters, except for the slight correlation between the AA and AP delay estimates as with the CIRC model. There are no appreciable systematic offsets in the delay posteriors.

Figs. 8 and 9 respectively show the relative widths of the phase and delay estimates output by CASA `fringe-fit`

for 200 Monte Carlo simulations of the ELLIP source with different noise realisations. As with CIRC, the phase and delay histograms coincide with the Bayesian posteriors. Neglecting the source structure for this centrally-located source distribution with two axes of symmetry does not affect the accuracy of the fringe-fitting.

For the ELLIP case, we also compare the CASA estimates with the results of fringe-fitting this dataset using the AIPS task FRING. We choose this model for the AIPS comparison since the 2PT model which has more than one source cannot be input to the FRING task in AIPS. As with `fringe-fit`, the fringe-fitting is performed by inputting the exact source model to FRING. A comparison of the `fringe-fit` and FRING results is shown in Figs. 10 and 11). The results from both AIPS and CASA are re-referenced to the same frequency (the central frequency of the band) used by ZAGROS. There is almost perfect correlation between the AIPS and CASA delay residual estimates, while the phase residual estimates also follow the same trend, but are more loosely correlated. This correspondence between CASA, AIPS, and ZAGROS results indicates that our software is consistent with different implementations of the fringe-fitting algorithm.

4.1.3 Two point sources (2PT)

The next class of source models we test, 2PT, are asymmetric source models with a primary point source of flux density 1 Jy at the phase centre and a secondary point source of 0.3 Jy located away from the central source at different distances. While the flux density ratios and source separations may change, the overall structure in this source model is more typical of VLBI sources and therefore is an important test to perform.

We simulate 9 datasets in which the secondary source is located at varying distances from the central source, from 20 μas to 100 μas in Declination ($\Delta\alpha = 0$), in steps of 10 μas . The flux densities of both the sources and the location of the secondary source are allowed to vary with uniform priors with the hyperparameters shown in Table 4. Crucially, the position prior of the secondary source is allowed to overlap with that of the primary source. In each case, the correct model (2PT) was favoured with a Bayes factor on the order of 10^{12} . The error in relative evidence is ± 0.7 .

Fig. 12 shows the posterior distributions of the parameters of 2PT, when the secondary source is located 100 μas away from the centre. This asymmetric source structure reveals stronger correlations between the source structure and phase residuals. The phase residuals of the different stations are also correlated with each other. The delay residuals are independent of the source structure, as with the previous models.

The ground-truth values for almost all fringe-fitting parameters are located within a credible interval of 1σ , except for the delay residuals of AZ and SP. For these two antennas, the ground-truth values lie at the 2σ level of the posteriors. This could be due to contributions from various factors such as the source orientation with respect to the PSF of the array, which in turn is a function of the baseline orientation. To test this, we generate synthetic data with a 2PT source model rotated by 90° in the sky ($\Delta\alpha = -100 \mu\text{as}$, $\Delta\delta = 0$) and perform fringe-fitting with ZAGROS. Fig. 13 shows the comparison between the delay posteriors of AZ and SP obtained




Figure 1. The 1-D and 2-D posteriors of the parameters of CIRC of size $20\ \mu\text{as}$. The principal diagonal gives the 1-D marginalised posterior distributions of the estimated flux density (S_ν), the HPBW of the Gaussian profile ($e_{\text{maj}} = e_{\text{min}}$), and the station delay residuals τ_p shifted to zero mean. The 68 (1σ) and 95 (2σ) per cent credible regions are indicated by the dark-red and light-red shaded regions respectively. The vertical green lines indicate the ground-truth values used in the simulation. Parameters with delta priors are not shown.




Figure 2. Posteriors of the phase residuals estimated by ZAGROS for synthetic data generated using CIRC models of three different sizes: 4, 12 and $20\ \mu\text{as}$. The ground-truth values are shifted to zero.




Figure 3. Same as Fig. 2, but for delay residuals.

for this dataset and that shown in Fig. 12. We see that the systematic offsets in the posterior peaks for both AZ and




Figure 4. The posteriors of the phase and delay residuals from Fig. 1 (20 μ s CIRC model) for all antennas. The ground-truth values are shifted to zero.




Figure 5. ZAGROS posteriors (green) from Fig. 1 (20 μ s CIRC) shown alongside the histograms of the phase residuals obtained for 200 Monte Carlo simulations with different noise realisations using CASA *fringefit*, with PT (blue) and CIRC (red) provided as input source models independently. The ground-truth values are shifted to zero.




Figure 6. Same as Fig. 5, but for delay residuals.

SP are reduced. Fig. 14 also shows the relative widths and the magnitude of the spread of the posteriors from Fig. 12 around the ground-truth values for all antennas plotted together. All the posteriors have comparable widths, with AA being the narrowest due to its high sensitivity.

Fig. 15 compares the ZAGROS posteriors with CASA estimates. Here, the point source assumption introduces significant offsets of up to $\pm 15^\circ$ in the estimated phase residuals. Incorporating the exact source structure in CASA results in the estimated phases coinciding with the ZAGROS estimates. Fig. 16 shows the comparison between the delay estimates obtained with both ZAGROS and *fringefit*. In this case, not accounting for the source structure results in offsets of up to ± 4 ps in the estimated delay residuals. Figure 17 shows the differences in phases between the corrected visibilities, after fringe-fitting with and without incorporating the correct source model. We may expect these offsets to be much larger in the low SNR regime. A crucial advantage of ZAGROS is that it can capture the uncertainties in and degeneracies between the source parameters and the fringe-fitting related parameters completely, allowing one to draw more robust scientific inferences in the more difficult cases.

5 DISCUSSION

As the source structure becomes more complex, it must be incorporated in the fringe-fitting process to avoid significant offsets in the corrected visibility phases. We find that our probabilistic approach to fringe-fitting performs at least as well as other widely-used software packages used for fringe-fitting, if not better. This is especially true in cases where a good source model from previous observations is unavailable or is difficult to obtain from the data, as is normally the case when the *uv*-coverage is sparse. It also possesses the unique advantage of returning full posterior distributions for the model parameters and thus reliable propagation of uncertainties and the use of more informative priors further down in the calibration process where iterative calibration is necessary.

We find that for resolved Gaussian source morphologies at the phase centre, the effects of not incorporating the source model is not as significant as for an asymmetric source model such as two point sources. While for short-duration observations, the sparse sampling will ultimately constrain the amount of information that can be obtained on source structure, especially if the source structure is complex, the availability of posteriors on super-resolved source parameters provides a mechanism for establishing the veracity of the results obtained using conventional methods. For simple symmetric sources, the conventional methods generally yield results that are close to the truth which is borne out by the correspondence with the Bayesian posteriors. For asymmetric source distributions, any unmodelled structure not accounted for during fringe-fitting introduces systematic offsets in the estimates of the phase and delay residuals, as evident from the point source assumption using CASA *fringefit* and the relative Bayesian evidences that are orders of magnitude (up to $10^{12}:1$) in favour of the correct source model in ZAGROS.

Furthermore, the unmodelled flux could be subsumed into self-calibration solutions, causing artefacts and sup-




Figure 7. Same as Fig. 1, but for the model ELLIP of size $25 \times 5 \mu\text{as}$.




Figure 8. ZAGROS posteriors (green) from Fig. 7 ($25 \times 5 \mu\text{as}$ ELLIP) shown alongside the histograms of the phase residuals obtained for 200 Monte Carlo simulations with different noise realisations using CASA fringeFit, with PT (blue) and ELLIP (red) provided as input source models independently. The ground-truth values are shifted to zero.




Figure 9. Same as Fig. 8 but for delay residuals.

pressing real sources, even in the case of connected-element interferometers with a large number of stations (Grobler et al. 2014; Sob et al. 2019). VLBI observations can be especially affected by this, given the sparsity of uv -coverage. Calibration and imaging is an ill-posed problem, with the so-




Figure 10. A comparison of AIPS and CASA estimates of the phase residuals, with the true values shifted to zero, for the 200 synthetic datasets of the $25 \times 5 \mu\text{as}$ ELLIP source with the exact source model input during fringe-fitting.




Figure 11. Same as Fig. 10, but for delay residuals.

lutions selected by self-calibration strongly affected by initial conditions. Obtaining joint posteriors of source parameters while performing fringe-fitting allows for the starting conditions to be closer to the ground-truth. In cases where the source structure is comparable to or smaller than the size of the PSF, it is difficult to obtain a model as a result of a previous imaging step. In such cases, using a priori knowledge of the source structure, modelled even as simple geometric models, selected based on Bayesian evidences obtained by a relatively quick run of ZAGROS, provides a good starting point for calibration; this would be especially useful in the low SNR regime.

Application of this technique to synthetic data with low SNR will help set an upper limit to the effect that the systematic offsets of the posteriors from the ground-truth will have on subsequent calibration of actual VLBI data. The correlation between source parameters and phase and delay residuals may be relevant to specialist experiments mentioned in Section 1.

Apart from imaging experiments, this technique will find special applications in high precision astrometry and

geodesy. For instance, simultaneous estimation of source positions in the presence of phase residuals on simulated European VLBI Network (EVN) data at 6.7 GHz, yield highly accurate estimates of the source positions (van Langevelde et al. 2019). Reversing this assumption and treating the source positions as known and the antenna locations as unknown, we can measure geodetic quantities more accurately (Taylor et al. 1999, Chapter 23). The crucial advantage to geodetic experiments here is the capability to estimate the frequency-dependent core-shifts of sources simultaneously with the antenna positions and study the degeneracies between them, since it is straightforward to account for frequency-dependence in the source structure in our current framework.

6 COMPUTATIONAL CONSIDERATIONS AND OUTLOOK

The Bayesian analyses performed for this paper take a maximum of about 8 hours to complete on a machine with an NVIDIA Tesla K20m GPU, compared to the ~ 30 minutes required for obtaining the histograms by fringe-fitting 200 datasets using CASA fringe-fit (excluding the ~ 2 hours required to simulate 200 datasets using MEQSILHOUTTE) on the same machine. In order to systematically explore a larger parameter space with ZAGROS in a reasonable amount of time (e.g. more complex source models, time and frequency variable complex gains, rate residuals, etc.), faster model computation and sampling techniques become necessary. For larger visibility data set sizes, the forward-modelling step must be distributed over multiple CPU and/or GPU nodes due to increased memory requirements, which therefore necessitates processing on High Performance Computing (HPC) clusters. A distributed version of CODEX-AFRICANUS (Perkins et al., in prep.) and the corresponding ZAGROS version which can distribute model computation between multiple HPC nodes using the DASK¹¹ library are currently under development, and will be available for production in the near future. Including complex-valued, time and frequency varying delay and rate residuals is a straightforward extension of the current framework that would benefit from increased computational speed.

The geometric models explored in this paper have relatively simple structures, modelled using point sources and Gaussians. While the asymmetric source models with secondary point sources placed at different locations is typical of some VLBI sources and are useful for studying the limitations of traditional fringe-fitting, the primary EHT targets have much more complex structure (e.g. Event Horizon Telescope Collaboration et al. 2019a,d). We plan to generate synthetic data with geometric ring and crescent models with jet components and fringe-fit them in two different ways: (i) using multiple Gaussian components to capture the source structure in the RIME and (ii) constructing the source-related terms of the RIME by performing Fast Fourier Transform (FFT) on parametrised ring and jet structures input in FITS format. Comparing the fringe-fitting

¹¹ <https://dask.org>.




Figure 12. Same as Fig. 1, but for the model 2PT with the second source located at $\Delta\delta = 100\mu\text{s}$.




Figure 13. Comparison of the delay posteriors of AZ and SP from Fig. 12 (left) with those obtained for synthetic data generated with the same sky model rotated by 90° ($\Delta\alpha = -100\mu\text{s}$, $\Delta\delta = 0\mu\text{s}$).

sults of the two methods will reveal the limitations of using simple geometric models to model rings and crescents.

Knowledge of these systematics will enable us to choose between these methods to fringe-fit synthetic observations of general relativistic magnetohydrodynamical (GRMHD) models (e.g. Mościbrodzka et al. 2016) of the photon ring surrounding the shadow of a black hole (e.g. Event Hori-




Figure 14. The posteriors of the phase and delay residuals from Fig. 12 ($\Delta\alpha = 0\mu\text{s}$, $\Delta\delta = 100\mu\text{s}$) for all antennas. The ground-truth values are shifted to zero.

zon Telescope Collaboration et al. 2019a,e). Using the FFT method, important physical parameters necessary to build GRMHD models can be directly sampled in image space and posteriors obtained at the end of a ZAGROS run. A systematic study of these fringe-fitting approaches, followed by a suite of self-calibration procedures to measure differences in




Figure 15. ZAGROS posteriors (green) from Fig. 12 shown alongside the histograms of the phase residuals obtained for 200 Monte Carlo simulations with different noise realisations using CASA fringe-fit, with PT (blue) and 2PT (red) provided as input source models independently. The ground-truth values are shifted to zero.




Figure 16. Same as Fig. 15, but for delay residuals.

the final maps is a natural extension of this and will be the subject of a future paper.

Blackburn et al. (2020) have shown that the usage of closure quantities is equivalent to numerical marginalisation over unconstrained gains. Under the RIME framework, scalar closure relations generalise to matrix closure relations and apply under certain assumptions of source symmetry (Smirnov 2011). Marginalising over uniform priors on instrumental phases and using these estimates to initialise gains for the full inference problem may serve as a good computational improvement and will be explored. We also plan to extend these analyses to multiple spectral windows so that this framework can be used to estimate multi-band delays. The ability to include both spatially and frequency-resolved source structure in fringe-fitting will prove useful especially when fractional bandwidths are large, as will be the case for arrays such as the next generation Very Large Array (ngVLA) (Carilli et al. 2015) and the next generation EHT (ngEHT) (Blackburn et al. 2019a).




Figure 17. Difference between the corrected phases for the 2PT simulation, after fringe-fitting with and without incorporating the source model, shown as a function of baseline length. The data are coloured by baseline and correspond to all channels and time.

ACKNOWLEDGEMENTS

We thank Simon Perkins for helpful discussions on using CODEX-AFRICANUS; Edward Higson, Will Handley, and Mike Hobson for discussions on POLYCHORD and DYPOLYCHORD. We thank Lindy Blackburn, Bill Cotton, and Landman Bester for useful discussions. We also thank the anonymous referee for their helpful comments.

IN and RPD are grateful for the support from the New Scientific Frontiers with Precision Radio Interferometry Fellowship awarded by the South African Radio Astronomy Observatory (SARAO), which is a facility of the National Research Foundation (NRF), an agency of the Department of Science and Technology (DST) of South Africa. OS is supported by the South African Research Chairs Initiative of the Department of Science and Technology and the National Research Foundation. The authors acknowledge funding from the European Union’s Horizon 2020 research and innovation programme JUMPING JIVE (grant agreement No 730884), the National Research Foundation (grant No 107324), and the Dutch Science Organization, NWO (contract No 629.003.009). All computation was carried out on the servers located at the Rhodes University Centre for Radio Astronomy Techniques and Technologies (RATT).

REFERENCES

- Alef W., Porcas R. W., 1986, *A&A*, **168**, 365
- Blackburn L., et al., 2019a, arXiv e-prints, p. [arXiv:1909.01411](https://arxiv.org/abs/1909.01411)
- Blackburn L., et al., 2019b, *ApJ*, **882**, 23
- Blackburn L., Pesce D. W., Johnson M. D., Wielgus M., Chael A. A., Christian P., Doeleman S. S., 2020, *ApJ*, **894**, 31
- Blecher T., Deane R., Bernardi G., Smirnov O., 2017, *Monthly Notices of the Royal Astronomical Society*, **464**, 143
- Broderick A. E., Loeb A., 2009, *The Astrophysical Journal*, **697**, 1164
- Cappallo R., 2017, FOURFIT user’s manual. MIT Haystack Observatory
- Carilli C. L., Holdaway M. A., 1999, *Radio Science*, **34**, 817
- Carilli C. L., et al., 2015, arXiv e-prints, p. [arXiv:1510.06438](https://arxiv.org/abs/1510.06438)

- Cornwell T. J., Wilkinson P. N., 1981, *Monthly Notices of the Royal Astronomical Society*, **196**, 1067
- Cotton W. D., 1995, in Zensus J. A., Diamond P. J., Napier P. J., eds, *Astronomical Society of the Pacific Conference Series Vol. 82, Very Long Baseline Interferometry and the VLBA*. p. 189
- Enßlin T. A., 2018, arXiv e-prints, p. [arXiv:1804.03350](https://arxiv.org/abs/1804.03350)
- Enßlin T. A., Frommert M., Kitaura F. S., 2009, *Phys. Rev. D*, **80**, 105005
- Enßlin T. A., Junklewitz H., Winderling L., Greiner M., Selig M., 2014, *Physical Review E*, **90**, 043301
- Event Horizon Telescope Collaboration et al., 2019a, *ApJ*, **875**, L1
- Event Horizon Telescope Collaboration et al., 2019b, *ApJ*, **875**, L2
- Event Horizon Telescope Collaboration et al., 2019c, *ApJ*, **875**, L3
- Event Horizon Telescope Collaboration et al., 2019d, *ApJ*, **875**, L4
- Event Horizon Telescope Collaboration et al., 2019e, *ApJ*, **875**, L5
- Event Horizon Telescope Collaboration et al., 2019f, *ApJ*, **875**, L6
- Falcke H., Markoff S. B., 2013, *Classical and Quantum Gravity*, **30**, 244003
- Grobler T. L., Nunhokee C. D., Smirnov O. M., van Zyl A. J., de Bruyn A. G., 2014, *MNRAS*, **439**, 4030
- Hamaker J. P., Bregman J. D., Sault R. J., 1996, *Astronomy & Astrophysics*, **117**, 137
- Handley W. J., Hobson M. P., Lasenby A. N., 2015a, *Monthly Notices of the Royal Astronomical Society*, **450**, L61
- Handley W. J., Hobson M. P., Lasenby A. N., 2015b, *Monthly Notices of the Royal Astronomical Society*, **453**, 4384
- Higson E., Handley W., Hobson M., Lasenby A., 2019, *Statistics and Computing*, **29**, 891
- Issaoun S., et al., 2019, *ApJ*, **871**, 30
- Janssen M., et al., 2019, *A&A*, **626**, A75
- Jeffreys H., 1961, *Theory of Probability*, 3rd edn. Clarendon Press
- Jones R. C., 1941, *Journal of the Optical Society of America*
- Junklewitz H., Bell M. R., Selig M., Enßlin T. A., 2016, *Astronomy & Astrophysics*, **586**, A76
- Kass R. E., Raftery A. E., 1995, *Journal of the American Statistical Association*, **90**, 773
- Lochner M., Natarajan I., Zwart J., Smirnov O., Bassett B., Oozeer N., Kunz M., 2015, *Monthly Notices of the Royal Astronomical Society*, **450**, 1308
- Martí-Vidal I., Pérez-Torres M. A., Lobanov A. P., 2012, *Astronomy & Astrophysics*, **541**, A135
- Mościbrodzka M., Falcke H., Shiokawa H., 2016, *A&A*, **586**, A38
- Natarajan I., Paragi Z., Zwart J., Perkins S., Smirnov O., van der Heyden K., 2017, *Monthly Notices of the Royal Astronomical Society*, **464**, 4306
- Readhead A. C. S., Wilkinson P. N., 1978, *The Astrophysical Journal*, **223**, 25
- Readhead A. C. S., Walker R. C., Pearson T. J., Cohen M. H., 1980, *Nature*, **285**, 137
- Schwab F. R., Cotton W. D., 1983, *The Astronomical Journal*, **88**, 688
- Smirnov O. M., 2011, *Astronomy & Astrophysics*, **527**, A106
- Sob U., Landman Bester H., Smirnov O., Kenyon J., Grobler T., 2019, arXiv e-prints, p. [arXiv:1910.08136](https://arxiv.org/abs/1910.08136)
- Spingola C., McKean J. P., Lee M., Deller A., Moldon J., 2019, *MNRAS*, **483**, 2125
- Taylor G. B., Carilli C. L., Perley R. A., eds, 1999, *Synthesis Imaging in Radio Astronomy II*, 1st edn. Astronomical Society of the Pacific Conference Series Vol. 180, Astronomical Society of the Pacific Conference Series
- Thompson A. R., Moran J. M., George W. Swenson J., 2017, *Interferometry and Synthesis in Radio Astronomy*, 3rd edn. Springer Verlag
- Trotta R., 2008, *Contemporary Physics*, **49**, 71
- Whitney A. R., et al., 2004, *Radio Science*, **39**
- Wielgus M., Blackburn L., Issaoun S., Janssen M., Johnson M., Koay J., 2019, EHT Memo Series, 2019-CE-02
- van Bemmell I., Small D., Kettenis M., Szomoru A., Moellenbrock G., Janssen M., 2019, arXiv e-prints, p. [arXiv:1904.11747](https://arxiv.org/abs/1904.11747)
- van Langevelde H., et al., 2019, arXiv e-prints, p. [arXiv:1901.07804](https://arxiv.org/abs/1901.07804)

This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.