

Improving pulsar-timing solutions through dynamic pulse fitting

Rowina S. Nathan^{1,2}★ Matthew T. Miles^{1,2,3} Gregory Ashton^{1,4} Paul D. Lasky,^{1,2} Eric Thrane^{1,2} Daniel J. Reardon^{1,2,3} Ryan M. Shannon^{1,2,3} and Andrew D. Cameron^{1,2,3}

¹*School of Physics and Astronomy, Monash University, Melbourne, VIC 3800, Australia*

²*The ARC Center of Excellence for Gravitational Wave Discovery – OzGrav, Hawthorn, VIC 3122, Australia*

³*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia*

⁴*Department of Physics, Royal Holloway, University of London, London TW20 0EX, UK*

Accepted 2023 May 31. Received 2023 May 30; in original form 2023 April 5

ABSTRACT

Precision pulsar timing is integral to the detection of the nanohertz stochastic gravitational-wave background as well as understanding the physics of neutron stars. Conventional pulsar timing often uses fixed time and frequency-averaged templates to determine the pulse times of arrival, which can lead to reduced accuracy when the pulse profile evolves over time. We illustrate a dynamic timing method that fits each observing epoch using basis functions. By fitting each epoch separately, we allow for the evolution of the pulse shape epoch to epoch. We apply our method to PSR J1103–5403 and find evidence that it undergoes mode changing, making it the fourth millisecond pulsar to exhibit such behaviour. Our method, which is able to identify and time a single mode, yields a timing solution with a root-mean-square error of 1.343 μs, a factor of 1.78 improvement over template fitting on both modes. In addition, the white-noise amplitude is reduced 4.3 times, suggesting that fitting the full data set causes the mode changing to be incorrectly classified as white noise. This reduction in white noise boosts the signal-to-noise ratio of a gravitational-wave background signal for this particular pulsar by 32 per cent. We discuss the possible applications for this method of timing to study pulsar magnetospheres and further improve the sensitivity of searches for nanohertz gravitational waves.

Key words: methods: data analysis – stars: neutron – pulsars: general – pulsars: individual: J1103–5403.

1 INTRODUCTION

Due to their remarkable regularity, pulsars are exceptionally accurate clocks. Pulsar timing is therefore useful across many areas of physics. Timing many stable pulsars across the sky in a pulsar timing array can detect low-frequency nanohertz gravitational waves (Hellings & Downs 1983) from sources such as supermassive binary black holes (Sesana et al. 2004; Kocsis & Sesana 2011; Taylor, Simon & Sampson 2017; Burke-Spolaor et al. 2019) and phase transitions in the early Universe (e.g. Starobinsky 1980; Grishchuk 2005; Lasky et al. 2016). The most recent data sets from three of the world’s major pulsar timing arrays, the European Pulsar Timing Array (Kramer & Champion 2013), the North American Nanohertz Observatory for Gravitational Waves (McLaughlin 2013), and the Parkes Pulsar Timing Array (Manchester et al. 2013), have shown evidence of a common, red noise process in the residuals of their pulsar arrival times (Arzoumanian et al. 2020; Chen et al. 2021; Goncharov et al. 2021). Common red noise was also found when these data were combined by the International Pulsar Timing Array (IPTA; Antoniadis et al. 2022). Common red noise can arise due to the stochastic gravitational-wave background. However, in order to make an unambiguous detection, one must observe an angular correlation function consistent with the Hellings and Downs curve

(Hellings & Downs 1983), a quadrupolar correlation between timing delays across the sky. With new data releases imminent, the detection of nanohertz gravitational waves may be around the corner.

Current pulsar-timing methods employ a single matched-filter template to calculate pulse times of arrival.¹ Often, the template is created by averaging together many of the observed pulses. This single, static template is then used to time all the pulses. However, all pulsars exhibit at least some degree of pulse-to-pulse shape variation (Parthasarathy et al. 2021), which conventional pulsar-timing methods are not able to account for.

There are a number of phenomena that are known to cause obvious changes in pulse shape. For example, giant pulses are extraordinarily bright and narrow outbursts of radio flux (Staelin & Reifenstein 1968; Geyer et al. 2021; Caleb et al. 2022). Pulses can be temporarily broadened or lensed by the interstellar medium (Rickett 1970; Shannon & Cordes 2017; Bilous, Ransom & Demorest 2019; Lin et al. 2021; Askew, Reardon & Shannon 2023). Geodetic precession, where the strong gravitational field of the pulsar results in spin–orbit coupling, causes pulse shapes to evolve over time (Kramer 1998; Fonseca, Stairs & Thorsett 2014; Desvignes et al. 2019; Venkatraman

* E-mail: rowina.nathan@gmail.com

Krishnan et al. 2019; Noutsos et al. 2020). A pulsar can undergo a period of mode changing, in which it switches between two or more characteristic pulse shapes (Bartel et al. 1982; Wang, Manchester & Johnston 2007; Miles et al. 2022). On occasion, pulses have been seen to cease completely, in a process known as nulling (Backer 1970; Gajjar, Joshi & Kramer 2012).

There are also more subtle pulse-shape changes, such as stochastic changes known as ‘jitter’ (Shannon et al. 2014; Parthasarathy et al. 2021) and slow evolution in pulse shape due to sub-pulse drifting (Drake & Craft 1968; Backer 1973; Chen et al. 2023). A single pulse profile template does not capture the pulse-shape variation from these and other effects. Since the average template may not be a good match for some pulses, the estimated pulse time of arrival can be significantly wrong, leading to increased errors in the timing solution. Pulse-shape variation therefore is a source of noise for pulsar timing arrays, reducing our ability to detect nanohertz gravitational waves.

While our primary goal is to reduce the timing solution residuals for gravitational-wave searches, modelling the evolution of pulse shape is interesting in its own right. By studying how pulse shape varies over time, it may be possible to gain a better understanding of the pulsar magnetosphere and/or pulsar emission (Rankin 1986; Cairns, Johnston & Das 2004; Janagal et al. 2022). Moreover, pulsar timing allows for tests of general relativity in the strong field limit (Kramer et al. 2006b, 2021a, b; Freire et al. 2012; Archibald et al. 2018; Voisin et al. 2020) and it provides us with an improved understanding of the neutron star equation of state (Demorest et al. 2010; Antoniadis et al. 2013; Fonseca et al. 2021; Miller et al. 2021; Riley et al. 2021) and thus the behaviour of matter at extreme densities (Oppenheimer & Volkoff 1939; Kurkela et al. 2014; Özel et al. 2016).

All pulsars exhibit pulse-shape variations. However, large shape-change events are less common in millisecond pulsars; they are mostly stable, which makes them good candidates for long-term pulsar timing. There are currently three millisecond pulsars known to mode-change: PSR B1957+20 (Mahajan et al. 2018), PSR J0621+1002 (Wang et al. 2021), and PSR J1909–3744 (Miles et al. 2022). PSR J1103–5403 is a pulsar observed by the MeerKAT telescope as part of the MeerKAT Pulsar Timing Array (MPTA; Miles et al. 2023). We show that this pulsar exhibits the characteristics of a mode-changing pulsar, as it has a group of early arriving outliers in the timing solution residuals. This millisecond pulsar has a period of only ~ 3.4 ms and is a good candidate for a timing array pulsar. However, the mode changing severely restricts its timing accuracy. This makes it an ideal test case for an alternate timing method that is better able to constrain pulse-shape variability.

Lentati, Alexander & Hobson (2015) developed a profile-domain timing method where individual time-averaged epochs were fit, allowing for simultaneous estimation of the timing model, dispersion measure variations, and pulse profile evolution. When this was implemented with broad-band frequency evolution, an improvement of up to 40 per cent was seen in pulsar-timing precision (Lentati et al. 2017). Pulse-to-pulse timing has already been shown to measure the glitch rise time of the Vela pulsar (Ashton et al. 2019b). However, pulse-to-pulse observations are only available for the brightest millisecond pulsars. It is therefore common to use epochs for timing, where an observing period is folded and summed over many pulse periods to increase the brightness of the pulse. We implement here an epoch-to-epoch fitting method, in order to determine whether the flexibility of this method provides insights on time-averaged data.

In this paper, we present a pulsar-timing method that allows for and is sensitive to pulse-shape variation, using PSR J1103–5403 as a case study. We are able to confidently determine outliers in

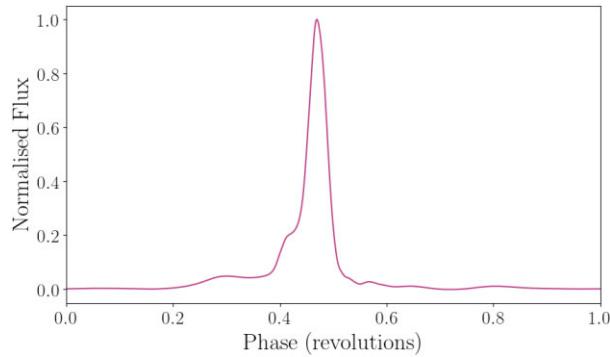


Figure 1. The smoothed total intensity profile of PSR J1103–5403. This is the average of 52 observing epochs from across 3 yr, which is then smoothed to emphasize the main pulse features.

pulse shape, the removal of which from the total data results in a reduction of the timing solution root-mean-square error (RMS; this error arises from the difference between pulse times of arrival predicted by the timing solution and observations). We reduce the RMS of this pulsar by a factor of 1.78 consequently improving the sensitivity of PSR J1103–5403 to the gravitational-wave background by 32 per cent. We describe how our method can be more broadly applied to other pulsars. The remainder of this paper is organized as follows. In Section 2, we describe our mathematical formalism and individual epoch fitting. In Section 3, we present the results of our analysis of PSR J1103–5403 and compare them to the matched-filter template method. We then use the parameter fits to each pulse to characterize the shapes present in the mode changing, and produce a single mode timing solution. Finally, in Section 4 we discuss the implications of our results and consider avenues for future work.

2 INDIVIDUAL EPOCH FITTING

In order to fit the flux profile of each epoch, we fit a sum of basis functions (similar to Kramer et al. 1994; Lentati et al. 2015; Padmanabha et al. 2021; Cameron et al. 2022). The parameters of these basis functions are fit independently for each epoch. We employ a Bayesian approach, using nested sampling (Skilling 2004, 2009) to explore the parameter space.

In Fig. 1, we show the averaged pulse profile for PSR J1103–5403, which has been smoothed to reduce noise using `psrsmooth` (Demorest et al. 2013). By visually inspecting the pulse shape, we identify three main pulse features: the shoulder at ~ 0.30 in phase, the leading edge of the pulse at ~ 0.42 , and the brightest portion of the pulse at ~ 0.48 .² Based on this morphology, we choose a sum of three Gaussians as the basis functions to model the pulse profile

$$F(\phi) = \sum_{i=0}^2 C_i e^{-(\phi - \phi_i)^2 / \beta_i^2} + \sum_{j=0}^2 B_j (\phi - 0.5)^j, \quad (1)$$

where $F(\phi)$ is the flux at phase ϕ , ϕ_i is the phase centring of each Gaussian, C_i is the amplitude, and β_i is the width. The second summation is included to model the baseline flux as a sum of polynomials of order j . The scale of the baseline flux for each polynomial is modelled by B_j with a subtraction of 0.5 in phase to ensure centring in the middle of the pulse.

²Following conventions in pulsar-timing literature, we measure phase in units of revolutions (from zero to one).

We use Fig. 1 to inform our priors, summarized in Table A1. A priori we know that the width of the Gaussians should be small compared to the pulse period motivating priors that are uniform in $\log_{10}(\beta_i)$. We place a minimum on each β_i that is five times the width of the time bins to prevent the model from fitting small noise fluctuations. The maximum on β_i prevents the Gaussian fitting the shoulder of the pulse (Gaussian 1) from interfering with the main pulse (which we use to determine the time of arrival as discussed below). The prior on both the Gaussian means ϕ_i and amplitudes C_i is a bounded uniform prior, forcing $C_0 < C_1 < C_2$ and $\phi_0 < \phi_1 < \phi_2$. This prior on the $\{\phi_i\}$ ensures that the Gaussians do not have multimodal posteriors due to swapping order. The constraints on the amplitudes are motivated by the average pulse profile and ensure that the model does not fit noise features to the right of the main pulse.

We also explore other potential models. For example, we fit a model with two Gaussians, but find that the Bayesian evidence prefers a three-component model. We fit a four-Gaussian model and find that the additional Gaussian increases variation in the phase centring of the largest amplitude Gaussian, increasing uncertainty in the timing solution. We attempt to model the epoch profile with a single higher order shapelet, as Lentati et al. (2017) (see Refregier 2003, for a definition of shapelets) or sums of multiple higher order shapelets, but again find for these models that the data prefer the three-Gaussian model, or that the computing expense is too great. As the pulse profile differs greatly between pulsars, this style of investigation is likely required for dynamic pulse fitting on all pulsars.

The fitting is done with KOOKABURRA (Ashton & Nathan 2020), an open-source PYTHON library. This library allows for easy definition of the priors and basis functions, and uses BILBY (Ashton et al. 2019a) for the fits themselves. KOOKABURRA fits pulses using shapelets (Gaussians are zeroth-order shapelets), but users may define other basis functions as well. KOOKABURRA outputs the posterior of the parameters, as well as a Bayes factor comparing the signal hypothesis (that the pulses consist of linear combination of basis functions) with the noise hypothesis that no pulse is present.

We fit our model to the data using nested sampling. KOOKABURRA and BILBY allow for the use of a number of sampling algorithms; we use the nested sampling algorithm pymultinest (Feroz, Hobson & Bridges 2009; Buchner et al. 2014). We individually fit 52 de-dispersed, frequency- and polarization-averaged, time-folded observing epochs from PSR J1103–5403. We use observations taken using the MeerKAT radio telescope L -band receiver, collecting between 856 and 1712 MHz (Bailes et al. 2020). The observations have a nominal cadence of two weeks (Miles et al. 2023), from 2019 August 2 to 2021 September 25.

We show a fit to an example epoch (observed on 2020 March 12) in Fig. 2. In both panels, the black curve shows the maximum-likelihood model fit and the grey curve shows the flux data. The top panel shows the 90 and 99 per cent confidence intervals from the posterior distribution of the sampling in pink and the bottom panel shows the three Gaussians making up the maximum-likelihood fit. The pink Gaussian characterizes the shoulder, the orange characterizes the leading edge of the main pulse, and the teal Gaussian fits the brightest portion of the pulse. Fig. 3 shows the posterior distribution for the mean of each Gaussian.³

After fitting for the pulse profile in individual epochs, we construct a timing solution for PSR J1103–5403. There are numerous ways one could define the time of arrival. We choose to identify the time of

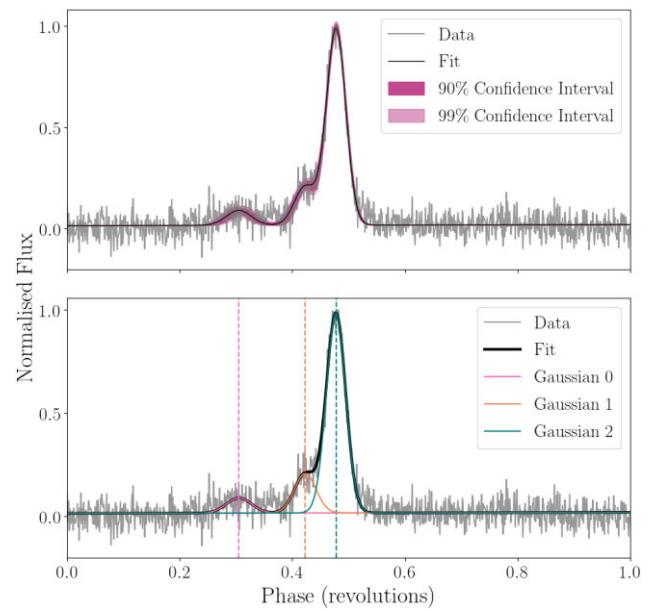


Figure 2. An example of one of the 52 individual pulse fits (this is the observing epoch from 2020 March 12). The top panel shows the maximum-likelihood fit in black, with 90 and 99 per cent confidence intervals from the posterior distribution of the fit shown in pink. The bottom panel shows the three Gaussians (shown in orange, teal, and pink) summed for the maximum-likelihood fit (black). The pulse data are shown in grey.

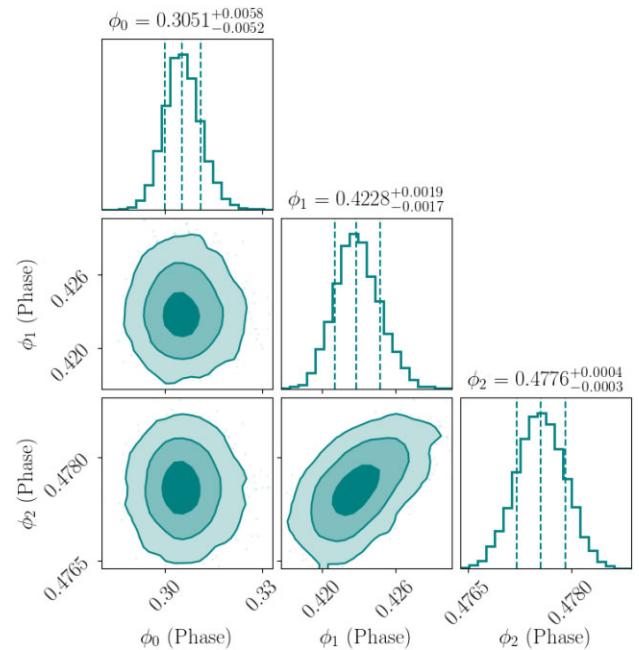


Figure 3. Posterior distributions for the centres of each Gaussian ϕ_i fit to the pulse in Fig. 2.

arrival as the maximum-likelihood estimate for ϕ_2 , the mean of the third and most prominent Gaussian peak (denoted by the dashed teal line in Fig. 2). Our rationale is that this prominent feature of the pulse profile is likely the most stable over long periods of time. We convert ϕ_2 , which is measured in phase, into a time of arrival in Modified Julian Day (MJD) in order to fit a timing solution. We record the 1σ credible interval for ϕ_2 ; this is an estimate for our uncertainty on the

³The corresponding corner plot showing the posteriors for all parameters is shown in Appendix B in Fig. B1.

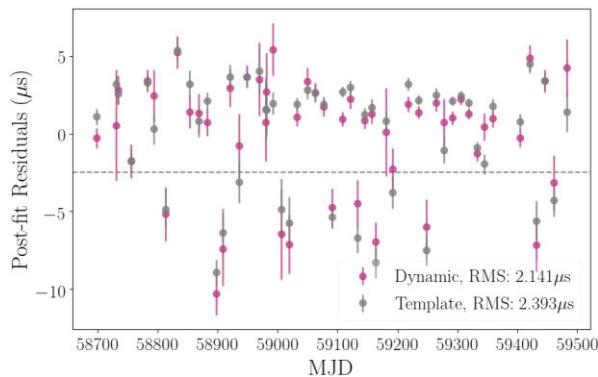


Figure 4. A comparison of the timing solutions produced by template fitting (grey) and dynamic basis-function fitting with KOOKABURRA (pink). Our method is flexible to changes in pulse shape while maintaining timing accuracy. Our timing solution improves on matched-filter template methods, with a smaller RMS of $2.141\,\mu\text{s}$. Both timing solutions show evidence of mode changing, highlighted by the grey line dividing the two groups of residuals.

pulse arrival time.⁴ The list of times of arrival (and their associated uncertainties) is passed to TEMPO2 (Edwards, Hobbs & Manchester 2006; Hobbs, Edwards & Manchester 2006), which finds the best-fitting pulsar model. This is achieved by a chi-squared minimization of the residuals (see Hobbs et al. 2006).

3 TIMING SOLUTIONS

We show our timing solution in Fig. 4. The pink points represent the post-fit residuals obtained with our three-Gaussian fit. We compare our dynamic method to matched-filter template times of arrival obtained by PSRCHIVE (Hotan, van Straten & Manchester 2004). We fit the template using the default PSRCHIVE fitting method (Fourier phase gradient; see Taylor 1992). This fit is represented by the grey times of arrival in Fig. 4. The two methods yield qualitatively similar timing solutions. However, our three-Gaussian fit yields a somewhat lower RMS: $2.141\,\mu\text{s}$ down from $2.393\,\mu\text{s}$. We attribute this reduction in RMS to the flexibility of our pulse profile fits, which we posit yield a more accurate fit on average.

The full power of our pipeline is yet to be demonstrated as additional improvement in the RMS is possible using analysis of the pulse profile shape. In Fig. 4, we see that there is evidence of mode changing in PSR J1103–5403, which is evident as a tendency for the post-fit residuals to cluster around two distinct values: one at zero and one at $\approx -7.5\,\mu\text{s}$. The dashed line on Fig. 4 differentiates between the two modes.

To investigate the shapes of the two modes, we plot the average profile of the two groups in Fig. 5. The difference in shapes exhibited in Fig. 5 suggests that the Gaussian fit to the leading edge of the pulse (Gaussian 1) may be able to differentiate between the modes. To further examine this, in the top panel of Fig. 6 we show a histogram of C_1/C_2 , the ratio of the height of the first Gaussian C_1 to the height of the second Gaussian C_2 . The total flux varies between epochs, thus this ratio allows comparison between epochs. This bimodal plot suggests that there could be two distinct sub-populations of pulse profiles, delineated by the vertical grey dashed line. Plotting C_1/C_2 versus residual in the bottom panel of Fig. 6, we see that the different

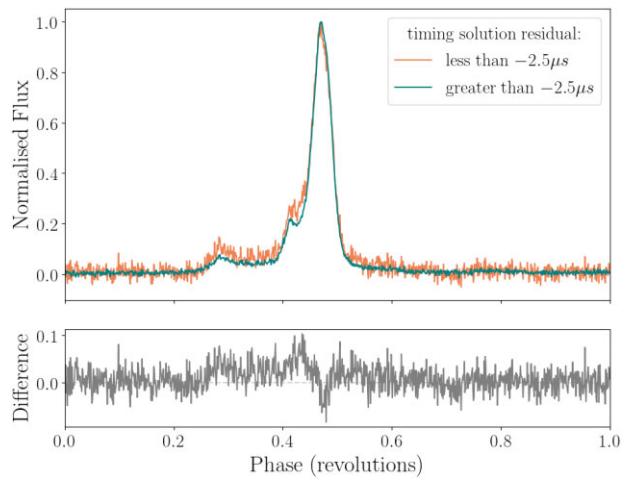


Figure 5. The average profile shapes of the two proposed modes, separated by the dashed line in Fig. 4. The top panel contains a teal curve comprised of epochs with a timing solution residual greater than $-2.5\,\mu\text{s}$ (Mode A), whereas the orange curve comprises of epochs with timing solution residuals less than $-2.5\,\mu\text{s}$ (Mode B). The bottom panel shows the difference between these two profiles in grey. There are noticeable differences in the shapes of these profiles at the leading edge (~ 0.42 in phase) and the shoulder (~ 0.30 in phase).

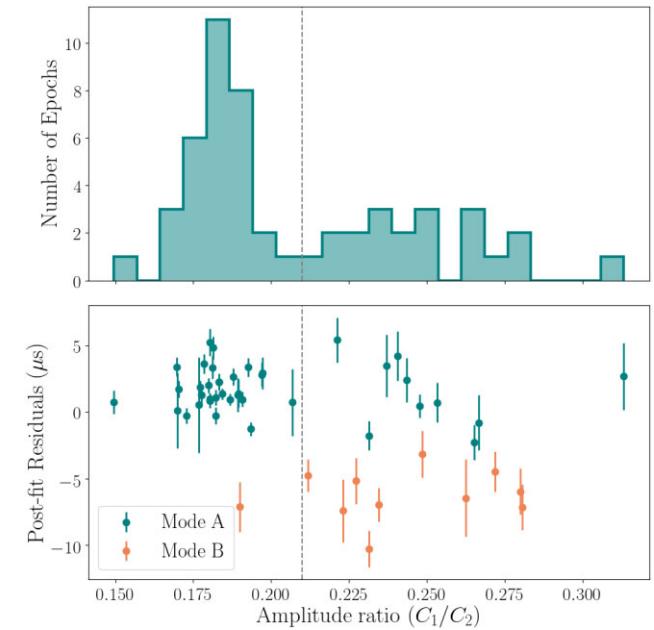


Figure 6. The ratio of Gaussian-peak-height one to Gaussian-peak-height two C_1/C_2 . The top panel shows the bimodality arising from the different pulse shapes of Mode A and Mode B, which we use to design our cut. The bottom panel shows the timing solution residuals versus C_1/C_2 , highlighting the effect of the cut to obtain a more pure sample of Mode A pulses.

⁴We use a single symmetric error estimate as TEMPO2 does not take a full posterior.

sub-populations hinted at in Figs 4 and 5 are strongly correlated with the different pulse profile modes (similar to Lyne et al. 2010). Small values of C_1/C_2 are associated with the mode clustered about a residual of zero (Mode A, shown in teal), while larger values are associated with the mode clustered around residuals of $-7.5\,\mu\text{s}$ (Mode B, shown in orange). The correlation between C_1/C_2 and the residuals confirms that the two modes are related to changes in

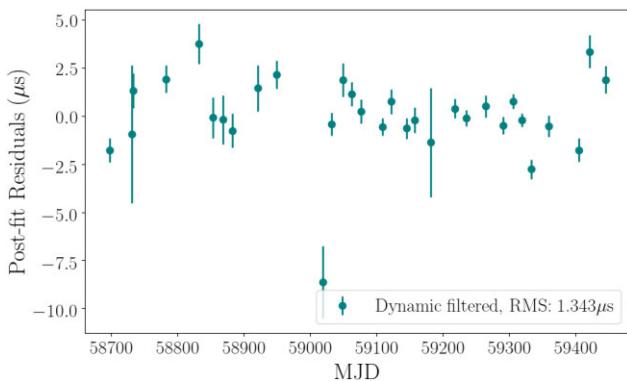


Figure 7. Timing solution produced with times of arrival from a flexible Gaussian fit, where pulses with an amplitude ratio (C_1/C_2) greater than 0.21 are removed. This timing solution has an RMS of 1.343 μ s, which is a factor of 1.78 improvement on the matched-filter template method. There is an outlier with a large negative residual at around 59020 MJD, due to imperfect separation between Mode A and Mode B when using the C_1/C_2 metric.

C_1/C_2 . We design a cut using C_1/C_2 that yields a subset of timing measurements targeting just a single mode.⁵

We cut the data, keeping epochs only if $C_1/C_2 < 0.21$ (the dashed vertical line in Fig. 6). This cut results in the removal of the majority of Mode B while preserving the bulk of the Mode A epochs. The cut is not perfect: there is still some contamination from Mode B epochs, and some of the data that are removed contain Mode A epochs. This implies that there are additional pulse-shape changes connected to the assumed mode changing that this method is unable to capture. However, the cut produces a more pure sample of Mode A. Fig. 7 shows the result of timing after the application of the cut. Our timing solution has an RMS of only 1.343 μ s, which improves on the matched-filter template timing solution when all data are used, by a factor of 1.78.

Next, we determine how our analysis affects the usefulness of these data for searches for gravitational waves. When modelling pulsar data, an uncorrelated white-noise term is included to account for pulse jitter, the random variation in pulse shape.⁶ This white-noise term in the full data set is 3.981 μ s, whereas in the data set with the cut it is 0.933 μ s. This factor 4.3 reduction in white noise suggests that the mode-changing behaviour was misspecified by the model, and was included instead as jitter noise. By accounting for mode changing in the timing and reducing the white noise, our noise modelling can be more sensitive to time-correlated red noise processes, such as the stochastic gravitational-wave background.

The signal-to-noise ratio (SNR) of the gravitational-wave background scales as (Siemens et al. 2013)

$$\langle \rho \rangle \propto (c\sigma^{-2})^{1/(2\beta)}, \quad (2)$$

where $\langle \rho \rangle$ is the (expectation value of the) SNR of the gravitational-wave background, σ is the white-noise RMS, β is the spectral index of

the gravitational-wave background, and c is the observing cadence.⁷ We use this scaling to assess the change in SNR from our analysis, which produces a smaller value for σ , but at the cost of throwing out some data:

$$\frac{\langle \rho_{\text{filtered}} \rangle}{\langle \rho_{\text{full}} \rangle} = \left(\frac{c_{\text{filtered}} \sigma_{\text{filtered}}^{-2}}{c_{\text{full}} \sigma_{\text{full}}^{-2}} \right)^{3/26} = 1.32. \quad (3)$$

Thus, for this particular pulsar, we can achieve a 32 per cent reduction in the noise for a stochastic background analysis.

4 DISCUSSION

We show, using PSR J1103–5403 as a case study, that employing a fitting method that can identify and then time a single mode (using flexible, epoch-to-epoch fitting) can improve timing solution residuals in mode-changing pulsars. This method may be able to identify other pulsars exhibiting mode changing and improve their timing. We hope to apply this method to other pulsars to see whether we are able to reduce the inferred white noise similarly. Additionally, we hope to fit basis functions to pulsars experiencing other kinds of pulse-shape variability to see whether these can be characterized and improved as well. Mode changing was recently found in J1909–3744, which is the most precisely timed pulsar (Miles et al. 2022). Therefore, dynamic pulse fitting may offer new insights to some of our most important and precise pulsars for timing.

While this work primarily focuses on the improvement of timing precision for gravitational-wave detection, pulse profile variability has the potential to provide insights into pulsar physics. In particular, the pulse emission mechanism is currently poorly understood, but the geometry of the pulsar magnetic field and emission heights are thought to influence the intensity of the observed radio flux. By using basis functions to characterize each observing epoch (or in the case of bright pulsars, individual pulses), we can start to piece together a picture of pulse-shape evolution. Large shape-change events such as mode changing, nulling, and sub-pulse drifting have been studied in order to constrain pulsar emission models (e.g. Kramer et al. 2006a; Lyne et al. 2010; Basu et al. 2016; Brook et al. 2016, 2018; Chen et al. 2022; Shang et al. 2022; Shaw et al. 2022), but we believe that a broader data set of subtle, pulse-to-pulse, or longer time span changes, may shed light on the emission mechanism further. Applying our method to many pulsars may therefore provide a large data set of pulse-shape evolution across multiple pulsars.

Mode changing and nulling are related to interesting shape-change phenomena. They are both thought to arise from a change in the structure in the magnetosphere in the neutron star. Janagal et al. (2022) show a relationship between periods of nulling and mode changing in PSR J1822–2256, where nulling is always followed by a particular mode. This correlation suggests that the mechanisms between nulling and mode changing may be strongly related. Mode changing and nulling are not commonly observed features of millisecond pulsars, but with PSR J1103–5403 being the fourth millisecond pulsar now observed to exhibit mode changing, it is clear that these events happen to pulsars with millisecond periods. This demonstrates that mode changing can occur on millisecond time-scales, and this can continue to constrain the physical processes causing these phenomena. The

⁵The ratio C_0/C_2 is also correlated with the post-fit residual, albeit less strongly than the ratio C_1/C_2 .

⁶In pulsar timing, white noise is typically described by the sum of the error factor (EFAC), error-correction factor (ECORR), and quadrature-added error (EQUAD; see Verbiest et al. 2016, for details). When white noise is mentioned in this paper, we refer only to ECORR, the pulse jitter term.

⁷Time-correlated noise terms were not able to be modelled in this data set due to an insufficient time span, so instead of providing a direct upper limit for gravitational waves we provide a projection based on Siemens et al. (2013).

observing epochs for PSR J1103–5403 are separated by days, but Miles et al. (2022) show that the millisecond pulsar PSR J1909–3744 has an upper limit of 2.941 ms on the time-scale for magnetospheric changes causing the mode changes.

There are a number of natural extensions to be made to the method introduced in this work. The observing epoch fits here are frequency- and polarization-averaged. Pulse profiles are known to have a large dependence on frequency in many pulsars, and averaging over large bandwidths may induce larger timing solution residuals (Demorest et al. 2013; Shannon & Cordes 2017). Low-frequency pulse profile changes specifically offer insight about pulsar emission geometry, offering a rare study of the configuration of the pulse emission mechanism (Olszanski et al. 2022). Applying this method to frequency sub-banded data may yield a simple yet impactful extension. Additionally, polarization profiles provide information about the electron density in the interstellar medium and magnetic fields in globular clusters (Abbate et al. 2023). The science to be gleaned from frequency and polarization banded pulse profile fits motivates the application of Bayesian basis-function fitting to them (see van Straten 2006; Lentati et al. 2017). Given that each profile can be analysed in parallel, this extra computation may be practical with a computing cluster, while offering the potential for rich pulsar science.

As discussed in Section 2, the choice for a three-Gaussian model was based on a visual inspection of the average pulse morphology. This model worked well for PSR J1103–5403, but every pulsar will have a unique pulse shape and pulse-shape evolution. For this reason, the extension of this method to other pulsars may require different combinations of basis functions for each pulsar. The options are endless, but there is still the question of how many basis functions are needed to best model data. To assist this, we are developing a transdimensional sampler that is able to fit the number of basis functions. This will allow us to remove the arbitrary choice of the number of basis functions.

Another pulsar of interest is PSR J1713+0747. This pulsar is observed by all constituents of the IPTA due to its apparent stability and large SNR. However, in 2021 April it underwent a dramatic pulse-shape-change event (Lam 2021; Singha et al. 2021; Jennings et al. 2022), which induced large residuals in the timing solution using the matched-filter template approach. This shape-change event seems to mainly effect the leading and trailing edges of the pulse, rather than the tallest component of the pulse. An appropriately modelled basis-function-fitting method would be able to reduce the timing residuals connected to these shape changes. In addition, quantifying the pulse-shape parameters of this event by recording how the parameters of the basis function change from pulse to pulse may offer insight as to how portions of the emission region are changing over the recovery period.

ACKNOWLEDGEMENTS

We acknowledge and pay respects to the Elders and Traditional Owners of the land on which our institutions stand, the Bunurong and Wurundjeri Peoples of the Kulin Nation. The authors wish to thank Marcus Lower for his constructive feedback. We wish to thank the anonymous reviewer for the helpful comments. The authors acknowledge support from the Australian Research Council (ARC) Centre of Excellence CE170100004, Future Fellowship FT190100155, and Discovery Projects DP220101610 and DP230103088. This work was performed on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR programme receives funding in part from the Astronomy National Collaborative Research Infrastructure

Strategy (NCRIS) allocation provided by the Australian Government. The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation.

DATA AVAILABILITY

The data underlying this article are available in the MPTA first data release on the GPU Supercomputer for Theoretical Astrophysics Research (gSTAR) Data Sharing Portal, at <https://dx.doi.org/10.26185/6392814b27073>.

REFERENCES

Abbate F. et al., 2023, *MNRAS*, 518, 1642
 Antoniadis J. et al., 2013, *Science*, 340, 448
 Antoniadis J. et al., 2022, *MNRAS*, 510, 4873
 Archibald A. M. et al., 2018, *Nature*, 559, 73
 Arzoumanian Z. et al., 2020, *ApJ*, 905, L34
 Ashton G., Nathan R. S., 2020, kookaburra, available at <https://github.com/GregoryAshton/kookaburra.git>
 Ashton G. et al., 2019a, Astrophysics Source Code Library, record ascl:1901.011
 Ashton G., Lasky P. D., Gruber V., Palfreyman J., 2019b, *Nat. Astron.*, 3, 1143
 Askew J., Reardon D. J., Shannon R. M., 2023, *MNRAS*, 519, 5086
 Backer D. C., 1970, *Nature*, 228, 1297
 Backer D. C., 1973, *ApJ*, 182, 245
 Bailes M. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e028
 Bartel N., Morris D., Sieber W., Hankins T. H., 1982, *ApJ*, 258, 776
 Basu R., Mitra D., Melikidze G. I., Maciesiak K., Skrzypczak A., Szary A., 2016, *ApJ*, 833, 29
 Bilous A. V., Ransom S. M., Demorest P., 2019, *ApJ*, 877, 125
 Brook P. R., Karastergiou A., Johnston S., Kerr M., Shannon R. M., Roberts S. J., 2016, *MNRAS*, 456, 1374
 Brook P. R. et al., 2018, *ApJ*, 868, 122
 Buchner J. et al., 2014, *A&A*, 564, A125
 Burke-Spolaor S. et al., 2019, *A&AR*, 27, 5
 Cairns I. H., Johnston S., Das P., 2004, *MNRAS*, 353, 270
 Caleb M. et al., 2022, *MNRAS*, 510, 1996
 Cameron A. D. et al., 2022, preprint ([arXiv:2203.15995](https://arxiv.org/abs/2203.15995))
 Chen S. et al., 2021, *MNRAS*, 508, 4970
 Chen J.-L. et al., 2022, *Res. Astron. Astrophys.*, 22, 115014
 Chen J. L. et al., 2023, *MNRAS*, 519, 2709
 Demorest P. B., Pennucci T., Ransom S. M., Roberts M. S. E., Hessels J. W. T., 2010, *Nature*, 467, 1081
 Demorest P. B. et al., 2013, *ApJ*, 762, 94
 Desvignes G. et al., 2019, *Science*, 365, 1013
 Drake F. D., Craft H. D., 1968, *Nature*, 220, 231
 Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
 Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
 Fonseca E., Stairs I. H., Thorsett S. E., 2014, *ApJ*, 787, 82
 Fonseca E. et al., 2021, *ApJ*, 915, L12
 Freire P. C. C. et al., 2012, *MNRAS*, 423, 3328
 Gajjar V., Joshi B. C., Kramer M., 2012, *MNRAS*, 424, 1197
 Geyer M. et al., 2021, *MNRAS*, 505, 4468
 Goncharov B. et al., 2021, *ApJ*, 917, L19
 Grishchuk L. P., 2005, *Phys. Usp.*, 48, 1235
 Hellings R. W., Downs G. S., 1983, *ApJ*, 265, L39
 Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655
 Hotan A. W., van Straten W., Manchester R. N., 2004, *Publ. Astron. Soc. Aust.*, 21, 302
 Janagal P., Chakraborty M., Bhat N. D. R., Bhattacharyya B., McSweeney S. J., 2022, *MNRAS*, 509, 4573
 Jennings R. J. et al., 2022, preprint ([arXiv:2210.12266](https://arxiv.org/abs/2210.12266))

Kocsis B., Sesana A., 2011, *MNRAS*, 411, 1467

Kramer M., 1998, *ApJ*, 509, 856

Kramer M., Champion D. J., 2013, *Class. Quantum Gravity*, 30, 224009

Kramer M., Wielebinski R., Jessner A., Gil J. A., Seiradakis J. H., 1994, *A&AS*, 107, 515

Kramer M., Lyne A. G., O'Brien J. T., Jordan C. A., Lorimer D. R., 2006a, *Science*, 312, 549

Kramer M. et al., 2006b, *Science*, 314, 97

Kramer M. et al., 2021a, *Phys. Rev. X*, 11, 041050

Kramer M. et al., 2021b, *MNRAS*, 504, 2094

Kurkela A., Fraga E. S., Schaffner-Bielich J., Vuorinen A., 2014, *ApJ*, 789, 127

Lam M. T., 2021, *Res. Notes Am. Astron. Soc.*, 5, 167

Lasky P. D. et al., 2016, *Phys. Rev. X*, 6, 011035

Lentati L., Alexander P., Hobson M. P., 2015, *MNRAS*, 447, 2159

Lentati L. et al., 2017, *MNRAS*, 466, 3706

Lin F. X., Main R. A., Verbiest J. P. W., Kramer M., Shaifullah G., 2021, *MNRAS*, 506, 2824

Liu K. et al., 2014, *MNRAS*, 443, 3752

Lyne A., Hobbs G., Kramer M., Stairs I., Stappers B., 2010, *Science*, 329, 408

Mahajan N., van Kerkwijk M. H., Main R., Pen U.-L., 2018, *ApJ*, 867, L2

Manchester R. N. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, e017

McLaughlin M. A., 2013, *Class. Quantum Gravity*, 30, 224008

Miles M. T., Shannon R. M., Bailes M., Reardon D. J., Buchner S., Middleton H., Spiewak R., 2022, *MNRAS*, 510, 5908

Miles M. T. et al., 2023, *MNRAS*, 519, 3976

Miller M. C. et al., 2021, *ApJ*, 918, L28

Noutsos A. et al., 2020, *A&A*, 643, A143

Olszanski T., Rankin J., Venkataraman A., Wahl H., 2022, *MNRAS*, 517, 1189

Oppenheimer J. R., Volkoff G. M., 1939, *Phys. Rev.*, 55, 374

Özel F., Psaltis D., Güver T., Baym G., Heinke C., Guillot S., 2016, *ApJ*, 820, 28

Padmanabhan P. V., Barr E. D., Champion D. J., Karuppusamy R., Kramer M., Jessner A., Lazarus P., 2021, *MNRAS*, 500, 1178

Parthasarathy A. et al., 2021, *MNRAS*, 502, 407

Pennucci T. T., 2019, *ApJ*, 871, 34

Pennucci T. T., Demorest P. B., Ransom S. M., 2014, *ApJ*, 790, 93

Rankin J. M., 1986, *ApJ*, 301, 901

Refregier A., 2003, *MNRAS*, 338, 35

Rickett B. J., 1970, *MNRAS*, 150, 67

Riley T. E. et al., 2021, *ApJ*, 918, L27

Sesana A., Haardt F., Madau P., Volonteri M., 2004, *ApJ*, 611, 623

Shang L.-H., Bai J.-T., Dang S.-J., Zhi Q.-J., 2022, *Res. Astron. Astrophys.*, 22, 025018

Shannon R. M., Cordes J. M., 2017, *MNRAS*, 464, 2075

Shannon R. M. et al., 2014, *MNRAS*, 443, 1463

Shaw B. et al., 2022, *MNRAS*, 513, 5861

Siemens X., Ellis J., Jenet F., Romano J. D., 2013, *Class. Quantum Gravity*, 30, 224015

Singha J. et al., 2021, *MNRAS*, 507, L57

Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V., eds, AIP Conf. Proc. Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. Am. Inst. Phys., New York, p. 395

Skilling J., 2009, in Goggans P. M., Chan C.-Y., eds, AIP Conf. Proc. Vol. 1193, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: The 29th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. Am. Inst. Phys., New York, p. 277

Staelin D. H., Reifenstein E. C. III, 1968, *Science*, 162, 1481

Starobinsky A. A., 1980, *Phys. Lett. B*, 91, 99

Taylor J. H., 1992, *Phil. Trans. R. Soc. A*, 341, 117

Taylor S. R., Simon J., Sampson L., 2017, *Phys. Rev. Lett.*, 118, 181102

van Straten W., 2006, *ApJ*, 642, 1004

Venkatraman Krishnan V., Bailes M., van Straten W., Keane E. F., Kramer M., Bhat N. D. R., Flynn C., Osłowski S., 2019, *ApJ*, 873, L15

Verbiest J. P. W. et al., 2016, *MNRAS*, 458, 1267

Voisin G., Cognard I., Freire P. C. C., Wex N., Guillermot L., Desvignes G., Kramer M., Theureau G., 2020, *A&A*, 638, A24

Wang N., Manchester R. N., Johnston S., 2007, *MNRAS*, 377, 1383

Wang S. Q. et al., 2021, *ApJ*, 913, 67

APPENDIX A: PRIORS ON DYNAMIC MODEL PARAMETERS

In Table A1, we show the prior distributions of the parameter fits to every observing epoch.

Table A1. The priors on the parameter fits to each epoch. There are three Gaussians for each fit, thus $i = \{0, 1, 2\}$. $\mathcal{U}(a, b)$ denotes a uniform prior distribution between a and b .

Parameter	Prior
ϕ_i (mean)	$\mathcal{U}(0, C_{i+1, i \neq 2})$
ϕ_2 (mean)	$\mathcal{U}(0, 1)$
C_i (amplitude)	$\mathcal{U}(0, C_{i+1, i \neq 2})$
C_2 (amplitude)	$\mathcal{U}(0, 1)$
β_i (width)	$\log_{10}[\mathcal{U}(\frac{1}{1024} \times 5, 0.032)]$

APPENDIX B: POSTERIORS OF INDIVIDUAL EPOCH FIT

In Fig. B1, we show the posterior distributions of the parameters of the three-Gaussian fit to the epoch shown in Fig. 2. The ϕ_i represent the centring of the Gaussians in phase, the β_i are the widths in phase, and the C_i are the amplitudes of each Gaussian in normalized flux. We use ϕ_2 as the time of arrival for producing a timing solution. We place an arbitrary maximum on β_i to prevent the Gaussian fitting the shoulder of the pulse (Gaussian 1) from interfering with the main pulse (which we use to determine the time of arrival as described above). We note that our posterior has support at this arbitrary maximum (see Fig. B1). However, we do not use the full posterior and, if we increase this maximum, the post-fit residuals are increased; therefore, we argue that the arbitrary maximum is justified.

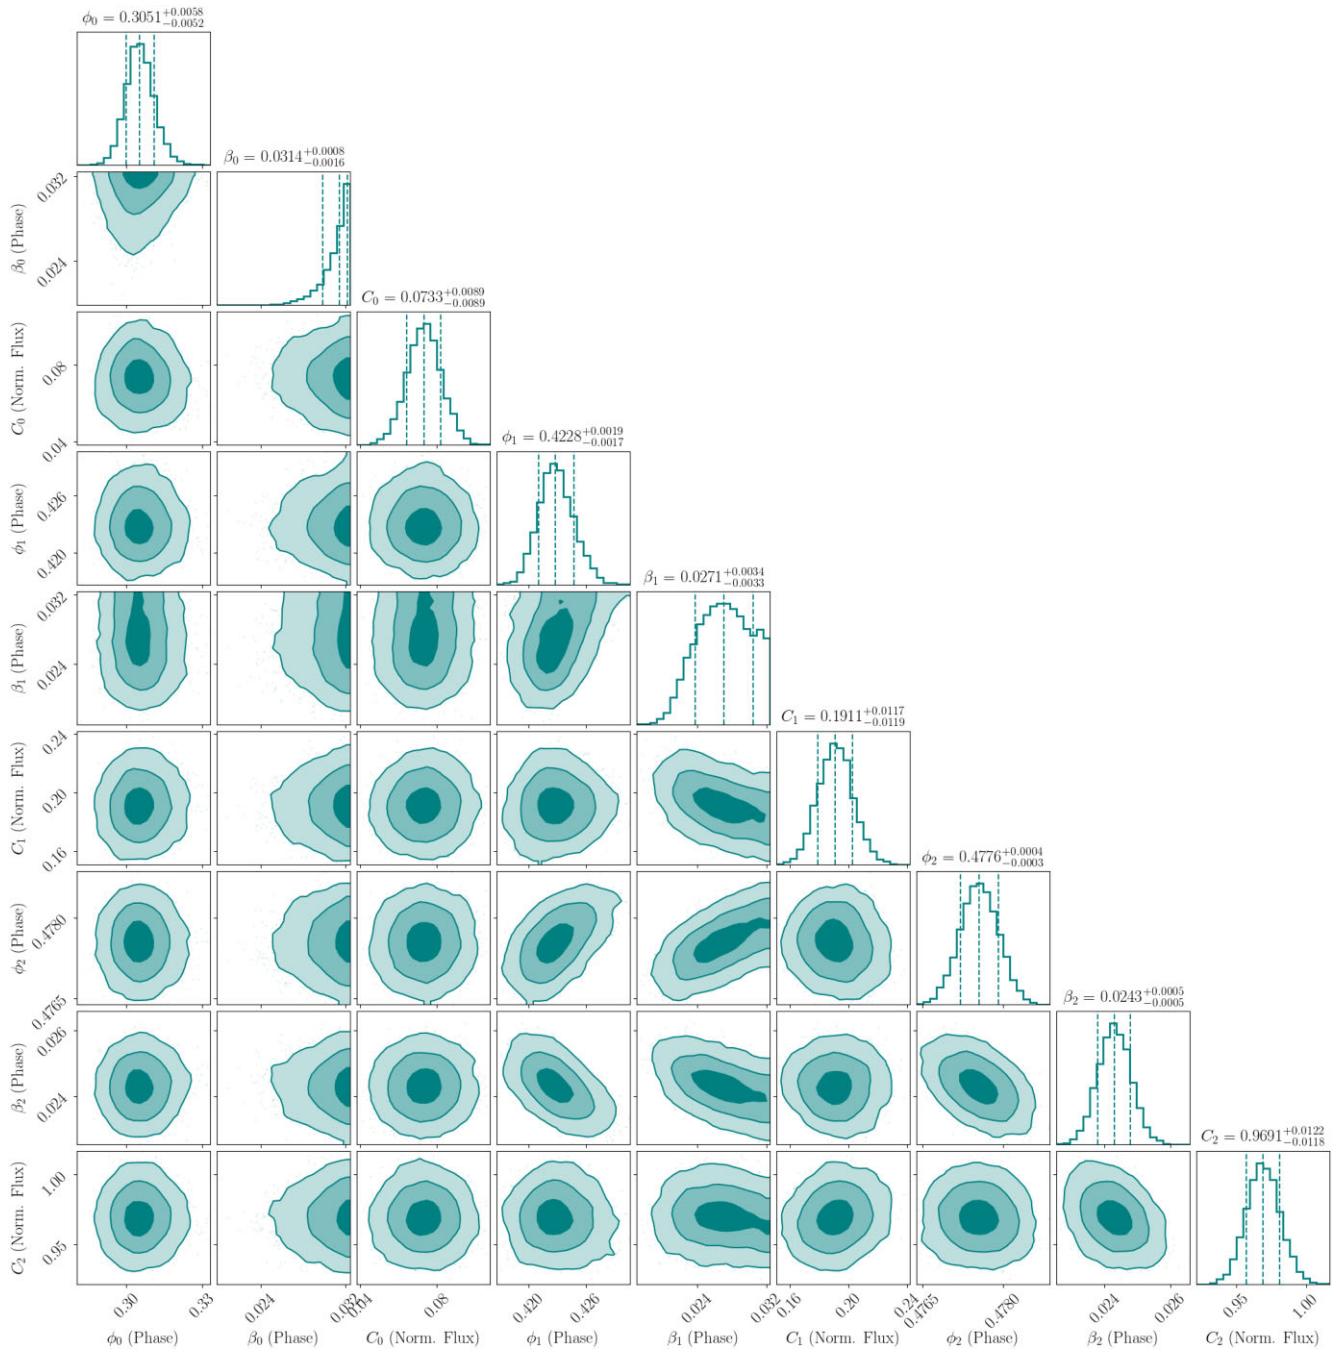


Figure B1. Posterior distribution of the parameter fits to the epoch shown in Fig. 2.

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