IMPROVING PULSAR TIMING PRECISION

A dissertation submitted to the University of Manchester
for the degree of Master of Physics
in the Department of Physics and Astronomy

2025

Jade Salisbury
(10867885)
Performed under the supervision of
Patrick Weltevrede

Department of Physics and Astronomy

Contents

Abstract

There are more than 3000 known pulsars in our galaxy. It is particularly challenging to adequately form accurate timing solutions for all of them for use in study of the pulsar population. A small portion of these pulsars exhibit 'glitches': instantaneous spin-up events which can provide great insight into the structure of a neutron star. The measurement of glitches is necessary to place constraints on physical models of pulsars. The majority of pulsar timing is done via pseudo-periodic observations at roughly constant intervals of time. Issues arise when glitches exhibit transient behaviours which can only be observed if the time between observations is short. We present 3 alternative scheduling strategies to compare retrieved glitch parameters to those found in periodically scheduled observations, each of which decreases the density of observations to a minimum cadence before restarting at a high density. Fitted values are found via a downhill simplex algorithm acting on simulated arrival times from PSR B1800–21, a Vela-like glitcher. Varying both the time when the glitch occurs in each strategy and the characteristic timescale of the shortest transient property of the glitch, it was found that periodic and arithmetic scheduling yielded the best results consistently. Due to the existence of high and low observation density regions in non-periodic scheduling, the chances for the fastest transient components to be observed are increased.

Acknowledgments

Importantly, I would like to acknowledge the usage of GitHub's Copilot coding tool in VSCode for its suggestion and line-completion functionality in some of the code for this project. Like my previous submission, I must stress that no other AI tools, generative or otherwise, were intentionally used during the work, code, and report for this project.

I've really enjoyed writing this project. I personally believe it is some of the best writing I have ever done, and is (hopefully) monumentally better than my submission last semester. I feel as though I have learned more physics in the last 12 months than I have in the entirety of what came before. I really do feel as though this project work (alongside my final-year modules) has really developed my intuition for physics.

Working on pulsars for my final year as an undergraduate has strangely changed my outlook on physics and altered the direction in which I have decided to take my career. At the end of my 3rd year, I had no intention of staying in academia following my graduation. The study of compact objects with the JBCA PET group has helped to nurture my enjoyment of learning and research, and I will be continuing onwards with a PhD at Swinburne, Australia soon after the submission of this document

I would like to thank my lab partner, Lauren Maclean for her work and company. I am glad we collaborate so well despite the fact we had never met before this (academic) year.

Finally, I would also like to ensure my thanks towards all four of my project supervisors, who do not all appear on the front page of this document: Danai Antonopoulou, Avishek Basu, Benjamin Shaw and Patrick Weltevrede.

1 Introduction

Neutron stars are a type of stellar remnant resulting from core collapse (type II) supernovae. They exist in a stability sustained only by degenerate neutron pressure rather than degenerate electron pressure (like a white dwarf) or nuclear reactions (like a main-sequence star). Neutron stars often rotate very quickly and have exceptionally strong magnetic fields. Interactions in the magnetic field often cause neutron stars to emit electromagnetic radiation along their magnetic poles, which can be misaligned with their rotational axis, creating sweeping beams of radiation detectable from Earth as relatively weak periodic pulses. Neutron stars that behave this way are therefore given the name pulsars, characterised by their lighthouse-like emission [1].

Pulsar astronomy as an area of physics is important due to its applications to other fields. The incredibly dense cores and extreme magnetic fields make them good astrophysical laboratories to study poorly-understood regions of otherwise well-understood physics. At such densities, there are theorised new states of matter beyond the assumed neutron superfluids, having large implications on early-universe cosmology [2]. It is important to get the best quality of results to bridge an understanding to this physics.

Pulsars alone provide very little information to an observer through direct observation. The evolution and behaviour of their spin as a function of time can be modelled precisely; these properties then allow us to infer much about the processes occurring within and around these stars. This being said, pulsars which are completely predictable tell us very little about the nature of neutron stars. Predictable pulsars which spin much faster than the average population are known as millisecond pulsars. Their spin periods are of order 1ms and are well suited as tools for other physics, such as detection of the gravitational wave background [3].

On the other hand, unpredictable pulsar behaviour can provide insight into a spinning neutron star. A pulsar is expected to slow by energy conservations, but a pulsar spontaneously and instantaneously increasing the rate of its rotations can provide information on the internal structure of the star. Wandering delayed or premature pulses observed from Earth can provide similarly useful information on the neutron star. Luckily, both behaviours are observed in pulsars and are known as glitches and timing noise respectively. Pulsar glitches can be parametrised through various fitting algorithms, and the quality of the retrieved results can place constraints on theoretical neutron star models [4, 5].

The majority of pulsars are observed quasi-periodically: their observations are separated in time by roughly constant intervals [6]. However, scheduling telescope time for pulsar observations is not strictly algorithmically decided and often involves much subjective decision-making. At Jodrell Bank Observatory*, researchers maintain a pulsar observation priority list, where pulsars are assigned points based on how interesting they are. A pulsar which has just glitched, for instance, might receive a number of points and therefore be more likely to be observed soon. This sort of scheduling gives rise to the 'quasi-periodic' observations seen in most literature, as those which have not been observed for some arbitrary decided time will receive points automatically, and move up in priority. Additionally, the telescope controller will select an order of observations such that the telescope spends its time most efficiently and does not waste time slewing across large portions of the sky between observations.

It is also possible to adjust the scheduling strategy once a glitch has been seen to occur: a warning system can alert an astronomer if a recent observation resulted in an unexpected result indicative of glitch behaviour. Unfortunately, these systems are subject to many external factors and often return false positives. Astronomers, constrained by limited time, will not always have the opportunity to tend to unlikely glitch candidates right away and therefore run the risk of missing the occurrence of a glitch until some weeks following. With longer gaps between observations around the time of a glitch, there is a risk the parameters found through retroactive fitting can deviate significantly from their true values and be of a lower quality overall.

Therefore, the primary motivation for this project is to improve the quality of fitted parameters for pulsar glitches with alternative scheduling. There have only been a few examples in literature for specifically non-periodic observation scheduling. Marshall et al. (2004) [7] used what they referred to as "logarithmic" spacing to ensure phase errors were kept low and Padmanabh et al. (2023) [8] use an initially higher density of observations decaying "pseudo-logarithmically" following the discovery of new pulsars in a MeerKAT† galactic plane survey. Recently, Dunn et al. (2021) [9] have shown that rigidly periodic observations can instead be detrimental to results, further justifying investigation into the viability of non-periodic observations.

With the advent of the Square Kilometre Array (SKA) on the horizon, which is predicted to discover over $10,000$ new pulsars in our own galaxy‡, it will become increasingly difficult to regularly observe and time each

*B. Shaw, private communication.

†MeerKAT is a radio telescope made from 64 13.5m dishes. It is a precursor to the Square Kilometre Array.

‡We currently know of $\sim 3,000$.

of them. As a result, it is important that suitable time is spent in exploring potential other, more time efficient scheduling strategies.

2 Theory

2.1 Pulsars and Radio Astronomy

At the end of the life of a massive main-sequence star, it can undergo a type II supernova. This process is also referred to as a core-collapse supernova after the violent internal process which compresses the internal core of the star to a much reduced size. Progenitor stars below around $8M_\odot$ can be supported by electron degeneracy pressure as a white dwarf; those above around $20M_\odot$ will lose to gravity and collapse further into a black hole. Between these two extremes, the core continues beyond the white dwarf stage until it is instead supported by neutron degeneracy pressure, allowing for a stable extremely dense compact object [1]. The density of neutron stars and therefore degenerate neutron matter is approximately equal to that of nuclei. A neutron star itself harbours the mass of a star ($\sim 1M_\odot$) compressed into an object of diameter ~ 10km [1, 5], providing a density of $\rho \sim 10^{17}\,\mathrm{kgm}^{-3}$.

The angular momentum and magnetic flux of the progenitor star must be conserved in the process of the collapse and therefore the resulting neutron stars often have extreme rotational speeds and high magnetic field strengths. The emission of radiation from the pulsar occurs within the magnetosphere, a strongly magnetic region of space around the pulsar, which co-rotates with the star. The magnetic poles of the neutron star are often misaligned with the rotational axis, so that they are quickly travelling through space. The mechanisms by which a pulsar emits in the radio spectrum are still not well understood and beyond scope to be described in full detail in this report. It is generally assumed to be related to emission processes in the creation and movement of plasma clouds [10, 11]. Importantly, it is known that the emission regions are associated with the magnetic field of the strong rotating dipole. Since, the emission regions are also rotating, this is observed as pulses separated by time equal to the rotation period. If the line of sight from Earth to the pulsar happens to coincide with the sweeping beam, the emission will intermittently radiate towards our telescopes.

2.2 Pulsar Timing

In astronomy there is often a trade-off between high time resolution and high sensitivity [5]. For both, one must use a larger telescope. Many pulsars are weak enough that it is impossible to discern individual pulses through simple analysis of high time resolution observations. In such cases, a Fourier analysis of the observations should be performed in order to reveal the hidden periodic behaviour which was lost in the noise of the observing system [5]. In frequency space, the pulses appear as harmonics separated by the frequency of the pulsar's rotation.

Once an estimate of the pulsar's spin frequency, ν, has been made through the above process, observations thereafter can be folded to create an averaged pulse profile over many subsequent pulses. Provided its pulse profile is consistent*, with a long enough initial observation, the pulse profile developed can be compared against the folded pulses of further observations in a process called 'template matching' [5] to precisely derive the time of arrival (TOA) for just a single pulse in an observation. There are several algorithms for generation of TOAs, the most popular of which is the Fourier phase gradient scheme detailed by Taylor (1992) [12]†.

TOAs themselves are heavily influenced by behaviours within our own solar system. Depending on the Earth's location in its rotation and orbit, the arrival time of pulses from a distant pulsar will vary throughout the day or year: this effect must be adjusted for. The easiest way to account for this locational delay (Roemer delay) is to adjust all TOAs so that they are instead representing the time at which a pulse would arrive at the solar system's barycentre. In addition to the classical effects of light travelling through the solar system there are several relativistic effects which depend on the location of massive bodies in the solar system‡, as the travelling light climbs through gravitational potentials and lenses through curved spacetime.

Pulses from distant pulsars may additionally be significantly smeared in arrival time over the bandwidth of observation. This process, known as dispersion, is a result of free electrons in the interstellar medium (ISM) along the line of sight causing a frequency dependent scattering on the incoming light towards Earth. Light travelling from further will travel through more of the ISM and therefore be more significantly scattered on its journey to earth. To create TOAs which are representative of a pulsar's behaviour, the smearing must be eliminated via dedispersion. Additionally, understanding the amount by which the data is dispersed can help to probe the ISM or derive estimates of distances to pulsars.

*Which, as it turns out, is not always the case.

†A comparison of several others is detailed by Wang et al. (2022) [13].

‡And potentially also in the pulsar system.

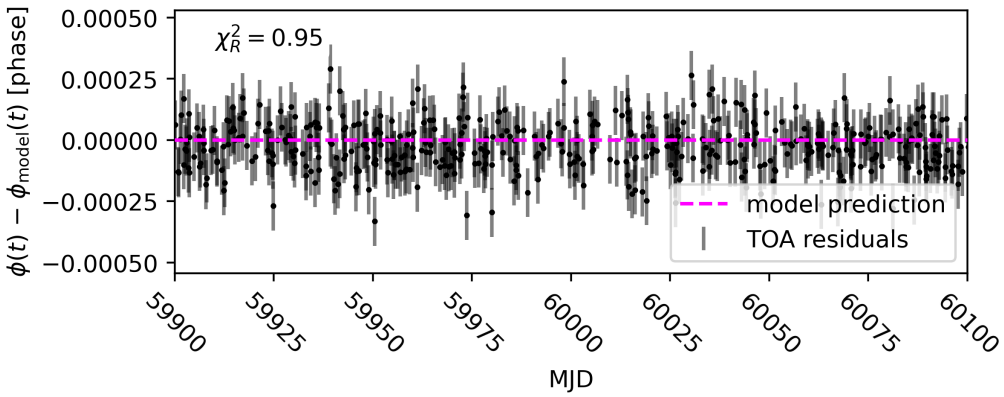

Figure 2.1: 'White' residuals for a pulsar with a well-described model. Example data is generated and not from a real pulsar.

To establish a suitable timing model for a pulsar, which fully describes the spin evolution of the system, one should represent the TOAs as residuals. Residuals describe the arrival of each pulse in terms of rotational phase in reference to an expected arrival time, according to a model. This phase is best described by a Taylor series expansion given by

$$\phi_{\mathrm{model}}(t) = \phi_0 + \nu_0(t - t_0) + \frac{1}{2}\dot{\nu}_0(t - t_0)^2 + \ldots, \tag{2.1}$$

where ϕ_{model} is the expected phase of the pulsar for a given time, t. Then, ϕ_0 is some arbitrary phase offset, ν_0 is the rotational frequency of the pulsar at epoch t_0 and $\dot{\nu}_0$ refers to the spin-down rate at the same epoch. Pulses are expected to arrive at integer values of ϕ_{model}, and those that don't are an indication of an incorrect model or spurious behaviour. Subtracting the model phase from the observed phases at each observation will return white/Gaussian residuals, as shown in Figure 2.1, with root mean square comparable to the TOA errors, with a $\chi^2_R \sim 1$.

Recognising the characteristic shapes of polynomials emerging in the structure of residuals plots can help to determine which values are incorrect in the model. The polynomial coefficients found this way are proportional to the Taylor expansion terms in Equation 2.1. It is usually up to a fitting algorithm to decide the best set of constants to describe the behaviour observed. This is typically performed by a pulsar timing package such as `Tempo2` [14] through chi-square minimisation, as was used for some parts of this project. A full explanation is given in Section 3.2.

One way to describe the spin evolution is through measurement of a braking index, n, which is defined by the power-law equation, $\dot{\nu} \propto -\nu^n$, and can be found via differentiation to be equal to $n = \nu\ddot{\nu}/\dot{\nu}^2$. Measuring the braking index for a pulsar can be difficult, particularly for younger pulsars due to the need for $\ddot{\nu}$, which often requires data over a long time-period. There has been recent work into in estimating n for specific times in a pulsar's spin evolution [15], giving insight to the pulsar's slowing behaviour. For instance, a pulsar slowing via magnetic dipole braking only would have a breaking index $n = 3$.

2.3 Mechanisms for Pulsar Irregularity

2.3.1 Glitches and Glitch Physics

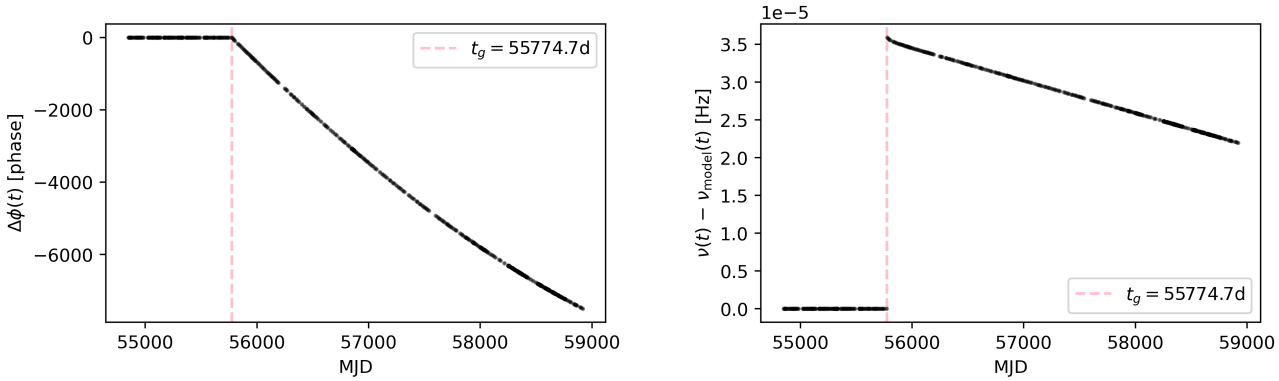

(a) Residuals

(b) Frequency residuals

Figure 2.2: Residuals (2.2a) and Frequency residuals (2.2b) for the glitch in PSR B1800–21 occurring on the 1st August, 2011. A short transient component can be seen in the small 'flick' occurring just after the glitch epoch in spin frequency, caused by a quickly decaying exponential change in ν, $\Delta\nu_d$.

Rarely, and somewhat unpredictably, a pulsar may exhibit an instantaneous spin-up event. The rotational rate and spin evolution of the pulsar are interrupted by some significant jump in several of the constants in Equation 2.1. As much as $\sim 6\%$ of the observed pulsar population are seen to glitch [16]. While only a small portion of the pulsar population is observed to have glitched, it is generally assumed that many more exhibit glitching behaviour. Due to the rarity of such events and an apparent age correlation it can be thought that older pulsars have long-since stopped glitching and that these events are too rare to be observed frequently in all young pulsars [4]. There have even been two instances of observed glitches in old millisecond pulsars [17, 18].

The effects of a glitch on pulsar residuals can be seen in Figure 2.2. The nature of these changes in behaviour can be revealing for the physical structure of neutron stars. For instance, Figure 2.2b shows us not only a change in the spin, $\Delta\nu$ (the discontinuity), but also the spin-down, $\Delta\dot{\nu}$ (the change in gradient), which was explained by Baym et al. (1969) [19] using a multi-component interior containing a neutron superfluid.

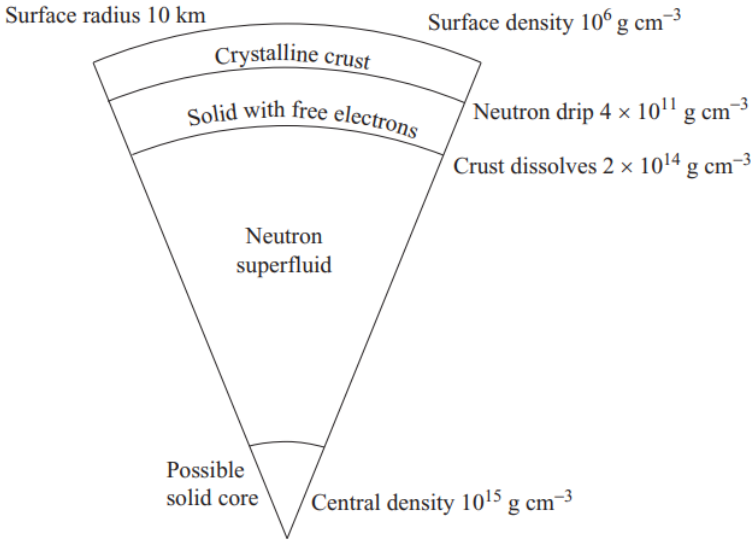

Figure 2.3: The assumed cross-section of a pulsar. Figure from [1].

Now, it is commonly assumed that the interior of a pulsar consists of at least two components [20] as can be seen in Figure 2.3, one of which is the neutron superfluid and the other a rigid crust containing massive nuclei and charged particles [1, 4]. Both components rotate, but the superfluid stores its angular momentum in vortices, rather than by exhibiting bulk motion as a classical fluid would. The total angular momentum of the

superfluid is proportional to the number of these quanta, with those further from the rotational axis contributing less to the total angular momentum.

Since it is known that the magnetosphere and therefore emission mechanisms are coupled to the solid, outer crust component of the star [1], it can be assumed that as the observed pulse period increases, it is the crust that is slowing via magnetic dipole braking and other energy loss mechanisms. On long enough timescales, the internals and externals of the pulsar will be co-rotating due to a weak mutual friction between the crust and the superfluid core [4]. An instantaneous change in the differential speed of both components results in an observable relaxation period on the timescales of minutes to years as the components come to recouple with one another.

The above behaviours have all been observed in real pulsar glitches. There are over 600 [21, 22] glitches which have been measured. Often found in young pulsars, these spontaneous changes in ν are often followed by a significant change in $\dot{\nu}$ and sometimes one or more transient exponential recovery components towards (but not always completely returning to) pre-glitch values on varying timescales, describing the recoupling/relaxation behaviour.

The origins of glitches can be attributed to one of two reasons. The less likely of the two is via "starquake". Ruderman (1968) [23] theorised that the crust is crystallised into tectonic plates, not unlike the Earth's, early in the life of the pulsar. As the star loses energy in its rotational evolution, it eventually becomes energetically favourable to reduce its oblateness in a tectonic event, quickly changing its moment of inertia and therefore causing an observed spin-up. Such events could describe some glitches in Crab pulsar-like glitchers[§], but the frequency of glitches in many pulsars exceeds that predicted by starquakes, and so an exclusively starquake induced glitch theory is unfeasible.

The primarily accepted theory is one of vortex pinning. There exists some border region between the interior superfluid and the crustal crystal lattice. For a superfluid to reduce its angular momentum as the pulsar slows, the contributing vortices must migrate outwards. Upon a vortex reaching the outermost layers, they can become pinned in the border region [24], fixing their contribution to the total angular momentum of the fluid. With time, the crust continues to slow and the superfluid maintains a higher rotational speed: the two components effectively become decoupled [1, 4]. The glitch eventually manifests in a cascade unpinning of the built-up pinned vortices, rapidly transferring angular momentum from the superfluid to the crust. It is currently unknown what drives enough vortices to simultaneously unpin at once, but there are several recent theories and it is thought that in some cases starquakes could be partially responsible [25].

When a glitch occurs, the process of timing as described by Section 2.2 is altered slightly. Equation 2.1 requires additional terms as described by

$$\phi_g(t \geq t_g) = \Delta\phi + \Delta\nu(t - t_g) + \frac{1}{2}\Delta\dot{\nu}(t - t_g)^2 + \sum_i \Delta\nu_{d,i}\tau_{d,i}(1 - \exp(-(t - t_g)/\tau_{d,i})), \tag{2.2}$$

which are added to the original Equation 2.1 for the times beyond the glitch occurrence epoch ($t > t_g$). Here, $\Delta\phi$ refers to a jump in rotational phase; $\Delta\nu$ is a change in ν; $\Delta\dot{\nu}$ is a change in $\dot{\nu}$. The final summation term contains the behaviour of the superfluid recovery as described by a linear combination of decaying exponential terms with $\Delta\nu_{d,i}$ being the changes in ν and $\tau_{d,i}$ being the timescale over which the response behaviour tends towards its pre-glitch values and is of particular interest to this project. Again, this equation is usually utilised in the context of fitting for the best values of all the corresponding constants to describe the behaviour seen in TOAs following a glitch.

Most ordinary glitching pulsars are usually categorised into one of two different behaviours: Vela-like or Crab-like, with namesake belonging to the two most observed pulsars, which appear to exhibit different glitching behaviours [21]. Vela-like glitches are typically more predictable in size and time of occurrence whereas Crab-like glitches occur much more sporadically and have vastly varying sizes. The physical difference is believed to be in relation to the internal interactions of the two-component pulsar model: Crab-like glitching is expected to at least partially originate from starquakes. On the other hand, Vela-like glitches are expected to occur wholly via superfluid vortex pinning and unpinning and so occur much more predictably. In this project, we have specifically looked to apply findings to PSR B1800–21, a typical Vela-like glitcher. Residuals for real glitch occurring in PSR B1800–21 an be seen in Figure 2.2b

2.3.2 Timing Noise

Timing noise is a stochastic process by which pulses may wander, appearing earlier or later on a quasi-periodic basis. Lyne et al. (2010) [26] have noted that there appears to be a correlation between timing noise and many other, beyond-scope, abruptly-changing phenomena such as mode-switching, nulling and pulse-shape variability. They state that timing noise likely arises from rapid switching between two states with differing spin-down rates on quasi-periodic timescales. Alternative theoretical explanations include Tkachenko modes [27], superfluid

[§]'Glitcher' refers to a glitching pulsar.

turbulence [28] and more [25]. In reality, it is likely that timing noise is the emergent behaviour from many such physical interactions occurring simultaneously. Interestingly, it has been pointed out by Hobbs et al. (2010) [29] that in younger pulsars, timing noise appears to be dominated by the relaxations of glitches detailed in Section 2.3.1: there is an obvious connection between the two phenomena and in better understanding either we can hope to place improved constraints on the other.

This quasi-periodicity in the behaviour of timing noise has led to its modelling via a power-law relating the strengths of periodic signals to their periodicity. This can be described in the Fourier domain (from sinusoids) as a power spectral density described by

$$P(f) = \frac{A_0^2}{12\pi^2}\left(\frac{f}{1\,\mathrm{yr}}\right)^{-\beta}, \tag{2.3}$$

where $P(f)$ is the strength of a periodic variation at frequency f and β is a spectral index [30, 31]. A_0 is some random-walk (red) noise amplitude in $\mathrm{yr}^{3/2}$. The full timing noise model will be made from of any number of these sinusoids[¶] with reducing contributions at higher frequencies. Similarly to pulsar timing and glitch modelling, the amplitude and spectral index is left to a fitting algorithm to find. This constraint of timing noise to abide by a power law ensures that a fitting algorithm will not just 'absorb' glitch parameters into the sinusoids when present, as it is possible to model any function with sufficient unconstrained Fourier terms.

3 Method

3.1 Scheduling Strategies and Data Generation

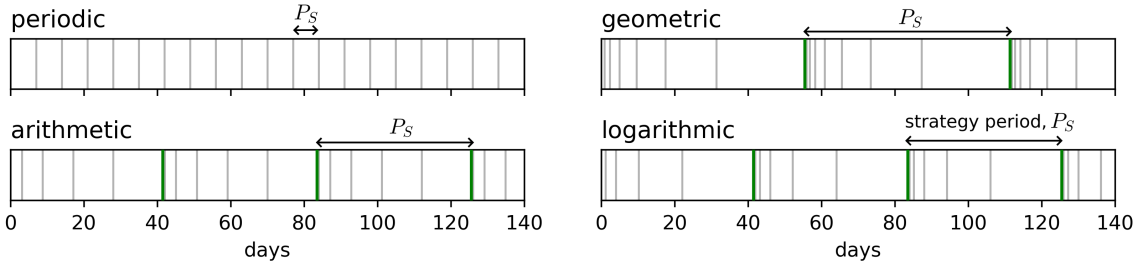

Figure 3.1: A plot demonstrating the observation behaviours of each scheduling strategy at the parametrisations given in Table 3.1. The grey lines indicate when an observation occurs and the green lines indicate a point where the next observation would be too long after the previous, and so the observation instead occurs at ΔT_{min}. Distance between green lines is equal to the periodicity of each strategy, P_S. This image is a modification of a similar image from [32].

Suppose at the time of an observation or measure of a TOA, the next observation will take place at time $\Delta T'$ later, and the previous observation is known to have occurred at a time ΔT before. Hence, we can write equations governing the new observation gap, $\Delta T'$, as a function of the previous observation gap, ΔT. In mathematics, these are known as constant recursive sequences

Rigidly periodic scheduling of observations can be described by the equation

$$\Delta T' = \Delta T = k_p, \tag{3.1}$$

which tells us that, following an observation, the time until the next observation, $\Delta T'$, is equal to the time since the previous observation, ΔT. The time between consecutive observations in the periodic scheduling regime is always the same and can be parametrised by a single constant, k_p. In a similar manner we present the equations for three more scheduling strategies: arithmetic, geometric and logarithmic.

Arithmetic observations can be described by the equation

$$\Delta T' = \Delta T + k_a, \tag{3.2}$$

where each subsequent observation occurs a time after the previous observation equal to the separation of the previous two observations, plus some constant k_a. Each observation gap increases by a constant amount of time, k_a, as shown in Figure 3.1.

[¶]Typically 60-100

Similarly, a geometric scheduling strategy can be described by the equation

$$\Delta T' = k_g \Delta T, \tag{3.3}$$

where k_g describes the constant by which every subsequent observation gap is multiplied by. A set of observations governed by a geometric strategy with $k_g = 2$ would double the wait time following every measure (e.g. 1 day, 2d, 4d, 8d, etc.).

Finally, a logarithmic scheduling strategy is described by

$$\Delta T' = k_l \ln\left(\frac{\Delta T}{10} + 1\right). \tag{3.4}$$

The initial intended equation for this strategy was $\Delta T' = k_l \ln \Delta T$, which created some difficulties in fine-tuning the constant, k_l. The alterations were made, resulting in Equation 3.4, to ensure it was easy to find values of k_l which only increase the observation gap without converging and to keep k_l closer in magnitude to the values of k for other strategies. For small values of ΔT, the logarithmic strategy is the same as the geometric strategy, meaning if $k_g \approx k_l/10$, both strategies will observe initially similarly but diverge with time.

However, for these types of sequences there must also be some starting number, ΔT_{start}, which in all cases in this experiment is 0.5d, decided as to match the highest observation cadence of an interesting pulsar such as the Crab or Vela. Similarly, each equation defines a sequence which is always increasing, so to ensure that the observation strategies remain useful beyond some arbitrary time limit, we have enforced a maximum observation gap, ΔT_{max}. If an observation defined by the strategy equations would return a value where $\Delta T' > \Delta T_{\text{max}}$, then the sequence instead enforces $\Delta T' = \Delta T_{\text{min}}$. This change makes each of the strategies periodic in time, as seen in Figure 3.1. Once the observation cadence is too infrequent, the scheduling will return to a high density of observations, where it will decrease their frequency until exceeding the maximum observation gap once more. This emergent periodic behaviour gives rise to each strategy having a length of time over which it repeats, called its strategy period, P_s.

It is generally well-established that an increased number of observations can better constrain parameters in fitting [6]. Therefore, comparing strategies with vastly differing observation cadences will yield biased results towards whichever strategy observes more often. As a result, the values chosen for ΔT_{max}, ΔT_{min} and the constants $k_{a,g,l,p}$ were selected as to enforce an average observation cadence of a single observation every 7 days, typical for current periodic observations of the PSR B1800–21*. They can be seen in Table 3.1.

Using a bespoke software package†, we generated a set of synthetic TOAs for a glitch in PSR B1800–21 occurring at time t = 60,000 MJD spanning both 1000 days pre and post-glitch‡, with a TOA placed periodically every 0.25 days. Naturally, this is an unrealistic sampling which results in a total of 8000 TOAs over the entire data span. The TOAs were then fed through a sampling algorithm which selected $\sim$ 300 TOAs per strategy in accordance to the equations detailed in this section and the constants in Table 3.1. By sampling TOAs in this way there are only 4 allowed observation times per day, and any scheduling strategy which would observe at another time instead selects the closest TOA, causing a natural deviation in scheduled observation time of up to 0.25 days. This should not be confused with the errors in arrival times/TOAs themselves, which are much smaller. The deviation is more akin to scheduling difficulties which may arise in real observatories where telescope time is more contested.

The parameters chosen in the creation of the TOAs are shown in Table 3.2. In magnitude and morphology, they follow typical behaviour of PSR B1800–21 as described in the work of Espinoza et al. (2011) [21]. As such, we modelled PSR B1800–21 as having a large glitch and two recovery components: short and long. In the creation of TOAs via this method, we introduced timing noise following an appropriate power law, also as observed in PSR B1800–21§ to act as a simulation of the behaviour seen exhibited by timing noise as described in Section 2.3.2.

Despite each strategy having an equal average observation distance, they all have differing periods, i.e. the time taken before they are return to the highest frequency region of observations. As each scheduling strategy

*B. Shaw, private communication.

†Developed by P. Weltevrede

‡MJD 59,000-61,000.

§M. J. Keith's timing solution, provided by B. Shaw, private communication.

Strategy	$k_{a,g,l,p}$	ΔT_{max}	ΔT_{start}	P_s
arithmetic	2.6019	15d	0.5d	42.0d
geometric	1.7391	25d	0.5d	55.9d
logarithmic	24.7678	20d	0.5d	42.0d
periodic	7.0000	7.0d	7.0d	7.0d

Table 3.1: The values chosen to parametrise each scheduling strategy used.

Parameter	Tempo2 Name	Value	Units
ν	F0	7.48	Hz
$\dot{\nu}$	F1	-7.52×10^{-12}	N/A
$-\beta$	TNRedGam	-4*	N/A
A_0	-	6.998×10^{-10}	$\mathrm{yr}^{3/2}$
$\log_{10}(A_0)$	TNRedAmp	-9.155	N/A
$\Delta\nu$	GLF0_1	3.0×10^{-6}	Hz
$\Delta\dot{\nu}$	GLF1_1	-1.5×10^{-14}	N/A
$\Delta\nu_{d,\mathrm{long}}$	GLF0D_1	3.0×10^{-7}	Hz
$\tau_{d,\mathrm{long}}$	GLTD_1	100	days
$\Delta\nu_{d,\mathrm{short}}$	GLF0D_2	1.5×10^{-7}	Hz
$\tau_{d,\mathrm{short}}$	GLTD_2	5, 7	days

Table 3.2: A table showing the true values for each of the fitted glitches. The only varied parameter was short timescale, τ_{short}. Values were decided based on true values found in [21]. Appropriate values for A_0 and β are from a timing model by M. J. Keith, provided by B. Shaw, private communication.
*-4 is not the actual spectral index seen in the red-noise power spectrum for this pulsar, the reasoning for this is discussed in Section 3.3.

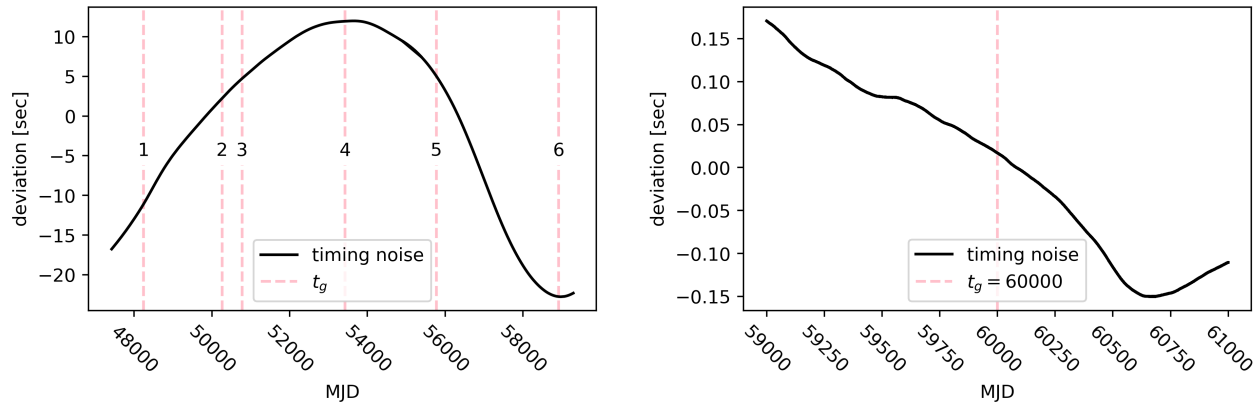

(a) Timing noise for PSR B1800–21

(b) Simulated timing noise

Figure 3.2: The timing noise exhibited by PSR B1800–21 and simulated red noise from Equation 2.3 and the values in Table 3.1. These plots were created by subtracting all spin-evolution behaviour and glitch behaviour from a plot of residuals. 3.2a shows the locations of all observed glitches in B1800–21. It is important to note that Figure 3.2b takes place over a much shorter time than Figure 3.2a. As a result, the deviations appear much larger in Figure 3.2a than those in Figure 3.2b by y-scale comparison.

exhibits this cyclic behaviour, the location in time where a glitch occurs in relation to the strategy's period (which dictates the locations in time of high and low densities of observation) is important. We can assume this is the only factor by which post-fit parameters would vary within a single strategy. As such, every dataset contained the TOA sets for between 8 to 17 different identical glitches, occurring at different points in each strategy. The specific locations of these were picked to enforce approximately 3 glitches per TOA gap¶, with the exception of periodic due to having only a single TOA gap per cycle. Therefore, in this case it was decided to increase this number to 8 distributed equally between periodically scheduled observations.

3.2 Data Analysis Techniques

The results of this project were derived through the main functionality of the aforementioned software package, which utilises a downhill simplex algorithm, a multi-dimensional numerical method for minimising functions [33]. Errors are derived through estimating the covariance matrix via calculation of the Hessian matrix [33, 34]. Additionally, while the software assumes some parameters are correlated, which is shown to be true, errors are only quoted as individual numbers on parameters. They therefore cannot represent the full matrix. A notable limitation of this software package is its inability to model timing noise as characterised in Section 2.3.2.

Simply fitting for all desired parameters at once will return spurious results. The values are correlated and therefore minimisation in one direction for a single parameter (or dimension in parameter space) might be

¶Occurring at roughly the beginning, middle and end of each TOA gap.

compensated for by similar changes in other parameters. To counteract this effect, pulsars are instead 'manually' fitted, referring to the process of an astronomer systematically deciding to introduce particular parameters to the fitting between fits, ensuring that results continue to make sense over the whole process. For instance, it is typical to include larger effects over a small number of TOAs in the fit first, such as changes in parameters $\Delta\nu$ and $\Delta\dot{\nu}$. Once the largest effects have been estimated through the fitting, one can justify the inclusion of more subtle phenomena if their effect is obvious enough in the new post-fit residual structure. The fitting region will also be increased to encompass more observations, alongside the complexity increases. Sometimes, in regions of low observation density, it was impossible to discern the existence of the shortest relaxation component with characteristic timescale, τ_{short}. In such scenarios we instead opted to not try and include a second exponential in the final fit. Once all desired parameters are included in the final fit, the model should return white/Gaussian residuals similar to those shown in Figure 2.1. A case where this is true can be referred to as 'valid solution' or that the pulsar is 'solved'. The aim of the process is to account and describe with a single model every pulse between any two observations [5].

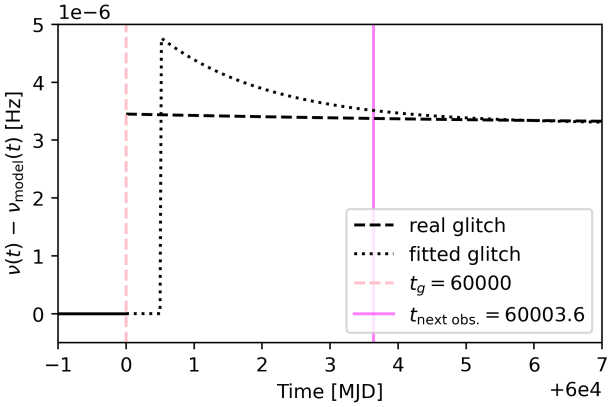

(a) Comparison between a fitted solution and the true value solution (see Table 3.2 for a double-exponential glitch).

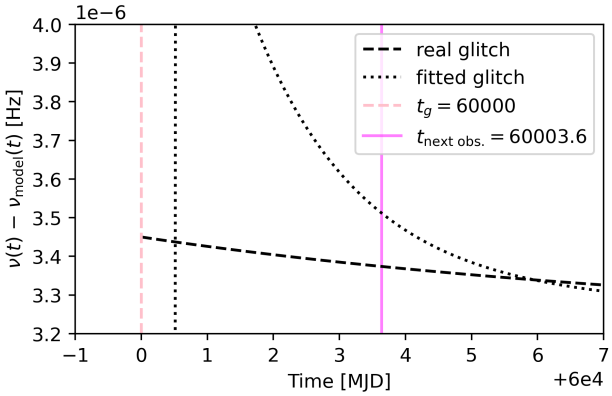

(b) Same as Figure 3.3a but with zoomed y-axis to better show the morphology of the true solution, which otherwise appears flat.

Figure 3.3: Comparing a poor fitted solution to a glitch with the true values for a glitch containing two exponential recovery components. While both the true values and fitted values give similar χ^2_R meaning they both suitably describe the observed data, a correlation is demonstrated. In fitting, the algorithm has preferentially chosen a very large decaying component $\Delta\nu_{\text{short}}$ with a very short timescale τ_{short}. For $\Delta\nu_{\text{short}}$ to be large and still describe the data, it must fully decay before the next observation (the magenta vertical line), and therefore there is a negative correlation between $\Delta\nu_{\text{short}}$ and τ_{short}.

Due to the correlations of parameters, multiple 'valid solutions' exist. There are a large amount of parameter sets which each accurately describe the true data. This effect is best demonstrated in Figure 3.3, where two different solutions well describe the observed data: the further away the next observation is from the glitch epoch, the less constrained the glitch model is in that TOA gap. This makes the model more prone to fitting undesirable and inaccurate results which can be explained by parameter correlation. It is therefore useful to create a corner plot: a set of 2 dimensional histograms which well-describe the distribution of parameters in a many-dimensional system. This allows simple investigation into correlations and a tool to justify observed under or overestimates in particular parameters. To create a corner plot for our glitch parameters we utilised Bayesian inference techniques [35]. We utilised a Markov chain Monte Carlo (MCMC) algorithm implementation called `emcee` [36] using pulsar package, `ENTERPRISE` [37] via `run_enterprise` [38]. MCMC algorithms can generate distributions in multi-dimensional parameter spaces via Bayes' theorem, which was then used to generate the corner plot.

In total, our total data used in this project contains many full sets of fitted glitch parameters for every used value of τ_{short}, fitted manually for each strategy with the glitch occurring at 7 to 18 differing time locations in the strategy's full cycle.

3.3 Assumptions and Caveats

Due to the nature of generating data, considerable care must be taken in the assumptions made during fitting and analysis. A good understanding of the assumptions can ensure knowledge of the regimes under which any findings are valid.

Firstly, the simulated PSR B1800–21 is assumed to have a previously known pre-glitch value for ν and $\dot{\nu}$. This is typically the case for glitching pulsars as they are actively observed and timed on a somewhat regular basis, though it may not hold valid for a newly discovered pulsar. Therefore, in fitting, the true values of ν and $\dot{\nu}$ are known perfectly to the fitting algorithm, but are still allowed to freely move (i.e. are not constrained to the true values). In most cases, where timing noise is not modelled and subtracted, glitch parameter fitting can be improved by deviating the values of ν and $\dot{\nu}$ as to 'absorb' some timing noise behaviour by approximating it as a small polynomial deviation in spin evolution around the glitch epoch. This is done via a re-fit in parameters ν and $\dot{\nu}$ over a small segment of pre-glitch data[‖], resulting in more accurate glitch parameters at the detriment of the quality of overall fitted spin evolution parameters.

We are making the assumption that over the 2000 days of generated data, only one glitch occurs and has a meaningful effect on timing. Realistically, for many Vela-like glitching pulsars there is good evidence pointing towards previous glitches having a large effect on the time evolution of a pulsar [15]. Similarly, Espinoza et al. (2017) [15] show that PSR B1800–21 has a non-insignificant inter-glitch rate of change of spin-down (high breaking-index, n), $\ddot{\nu}$, which likely arises from superfluid recoupling between and that we have opted not to model. Since we are only modelling a single glitch, an exclusively inter-glitch parameter would make little sense and be difficult to model. Additionally, due to the generated TOAs used in this project, introducing an overall $\ddot{\nu}$ behaviour everywhere would serve only to be removed by the fitting algorithm and have no overall effect on results.

In the generation of TOAs for use in timing, we selected to have an arrival time error of 100µs on every TOA. This is a larger-than-typical TOA error for the Fourier phase gradient scheme [13]. The number was chosen as 100µs as this is roughly the average error in arrival times for Crab pulsar observations by the Lovell telescope. In real observations, variations on TOA errors are typical for multitudes of reasons from differences in observation time to disturbances to the temperature of the telescope. There are many effects which contribute to the noise of a receiver system, and it is unrealistic to assume a constant error. Luckily, there is very little difference to fitted results: fitting algorithms will weight the contribution of data-points based on the size of their errors, which the only effect lost by making a constant error assumption. Instead, every TOA will contribute equally to the fit. Modelling varying TOAs would prove difficult and move the problem to an investigation on how to best simulate TOA errors.

Another assumption made is one of a 'coherent' solution. This can also be described as having perfectly numbered pulses. It takes significant work to ensure that the measured pulses are correctly numbered [5]. Each pulse is assigned a number which allows pulsar timing packages to more easily run fitting algorithms without bogus results. There will always be an integer number of pulses between any two moments in time, arriving at the Earth at roughly regular intervals, and therefore each one can theoretically be counted based on the time since the last measured pulse and what is known about the spin evolution. Measurements taken long intervals after previous ones may be mistaken as a different pulse, in which case in fitting it will be treated as arriving some integer number of spin periods early or late.

The occurrence of a glitch creates additional difficulty in maintaining accurate pulse numbering due to their instantaneous and unpredictable nature. Particularly large glitches may cause a loss of coherence, meaning the post-glitch rotation count of the pulsar is unclear. Coherence can often be re-established by an experienced astronomer and so, in the interest of project scope, we have opted to ensure that pulse numbers are known perfectly. Losing coherence is a more significant problem for pulsars which go without observation for large amounts of time, but PSR B1800–21 is regularly observed at least once per week[**], and so does not deviate significantly between subsequent observations.

Glitches in PSR B1800–21 are relatively consistent [39, 21]. The parameters $\Delta\nu$ and $\Delta\dot{\nu}$ are similar in each glitch. Similarly, the morphology of the glitch, which is to say the specific behaviours it exhibits as represented by Equation 2.2, is similar each time[††]. It should be noted that Liu et al. (2024) [39] present the glitches they work with as containing either a single or double exponential term[‡‡], but in the cases of two exponential terms, they cannot place good constraints on its response timescale, τ_{short}. Therefore, for the sake of this investigation, we make the assumption that the glitch morphology of PSR B1800–21 will always contain two exponential recovery terms, and that the timescale of the shorter recovery may be shorter than can be consistently observed via purely periodic scheduling. In fitting, where the effects of the shorter recovery are invisible, the second exponential is omitted from the fit.

We also make the assumption that the time when the glitch occurred is known to the correct TOA gap, i.e. the assumption that the last pre-glitch TOA and the first post-glitch TOA are always known. In realistic pulsar timing, small glitches or glitches with dominant effects in parameters other than $\Delta\nu$ can often be difficult to pinpoint in time. This effect is amplified if TOA errors are large or if TOAs near the glitch epoch happen to stochastically vary significantly from their expected arrival time. However, as the glitch morphology for

[‖] $\lesssim$ 10 TOAs

[**] B. Shaw, private communication.

[††] With the exception of the glitch occurring on 30th June 1996 (MJD 50264) [21], which has an uncharacteristically small $\Delta\nu$.

[‡‡] I.e. $i = 1$, 2 in the final summation term in Equation 2.2.

B1800–21 is relatively consistent, this assumption is not unrealistic; all glitches in this pulsar are likely to be large enough to be informed of the exact two TOAs between which the glitch falls. It is possible that currently undetected small glitches have occurred in PSR B1800–21, but this is unlikely as it is bright [40], frequently observed, and exhibits only subtle timing noise on small timescales.

The found power spectral index, β, for the wandering timing noise as described in section 2.3.2 for PSR B1800–21 is -7.185[§§], but this is relatively steep and therefore only exhibits strong timing noise on large timescales. Since we wanted to investigate the effects of timing noise in non-periodic scheduling for a 2000 day dataset, we opted to change the spectral index to $\beta = -4$, as to contain sufficient deviations over the timespan of generated TOAs. $\beta = -4$ is typical of timing noise in younger pulsars [30].

Finally, there are a handful of human-caused biases that come into play through manual pulsar fitting as described in Section 3.2. As we generated the values we have always known, to the exact number, which parameters make up a perfect solution. As a result, we found ourselves unintentionally informing our decision-making by the closeness of post-fit results to their true values, and would restart the process if any deviated too significantly from an expected model. Investigation into this effect is difficult, but we processed a small subset of our TOA groups using Markov chain Monte Carlo fitting with uninformed (uniform distribution) priors, removing the biases. This was also done as to retrieve information on potential parameter correlations, which can help us to understand the effects of biases.

4 Results

4.1 Parameter and Strategy Comparison

Presented in Figure 4.1 are the fitted results for all parameters of a simulated glitch according to Table 3.2. The average cadence of observations was been set at 7d for all strategies. As expected, across the majority of parameters (rows in Figure 4.1), retrieved fits are closer to their true values towards to the beginning of the strategy cycle, showing that a higher density of observations are generally good for pulsar timing. Other conclusions can be made from this plot and others to better understand the dominant effects of non-periodic scheduling and try to derive a new, improved method observing glitching pulsars.

There is a very evident relation between strategy position and retrieved glitch morphology. In nearly every case, the scenarios where the second, shorter exponential relaxation component cannot be measured is when the glitch has occurred immediately after a TOA and the time before the subsequent observation is large. The most extreme cases of this can be seen in the logarithmic scheduling case, where the penultimate TOA occurs around 23d after the strategy begins, initiating a TOA gap of ~ 20d before the final observation. The only exception to this is for the glitch occurring at 0 days in periodic scheduling. The strategy can apparently register the effects of a 5d recovery when the next observation is in 7d, despite the fact that data for a glitch occurring at 6d prior to the observation cannot. This is likely due to the systematic error on observation time (indicated by the horizontal error bars in the periodic plots), and the glitch actually occurred immediately before the observation, allowing good constraints on the timescale, τ_{short}, and magnitude, $\Delta\nu_{\text{short}}$.

Similarly, on all strategies, the best retrieved values are found when the glitch occurs immediately before an observation. This relation appears suggests that a very important factor in glitch parameter retrieval quality is the time until the first post-glitch observation. Even in low density regions of a strategy, such as around 30d in geometric (column G), the retrieved parameters for a glitch occurring immediately before the next TOA are of equal quality to those occurring in the high density regions (e.g. 0-5 days), for all parameters.

A particular fact to note about all manually fitted points is that they generally appear to have underestimated errors. Statistically, roughly 67% of the points should have their error bars intersect with the true values (this indicates a $\chi^2_R \approx 1$), which is not the case. In many of the plots, particularly those belonging to the non-periodic strategies have negligible derived errors on their y-axis. This is likely because the fitting model is not wholly representative of the glitching pulsar, as it does not model vital information such as the timing noise, and therefore its effects alter the derivations of errors.

4.2 Short Recovery Timescale

In parameter τ_{short} (row 7), the periodic scheduling strategy frequently yielded significantly worse (further deviation from the true values) results, but there are measurements of the short timescale for proportionally more of the strategy than geometric and logarithmic. This suggests that manual fitting periodically scheduled observations is more likely to enable an astronomer to deduce the existence of a short timescale exponential component to a glitch. This is expected, as each of the other strategies contain gaps which are much larger

[§§]From a timing model by M. J. Keith, provided by B. Shaw, private communication.

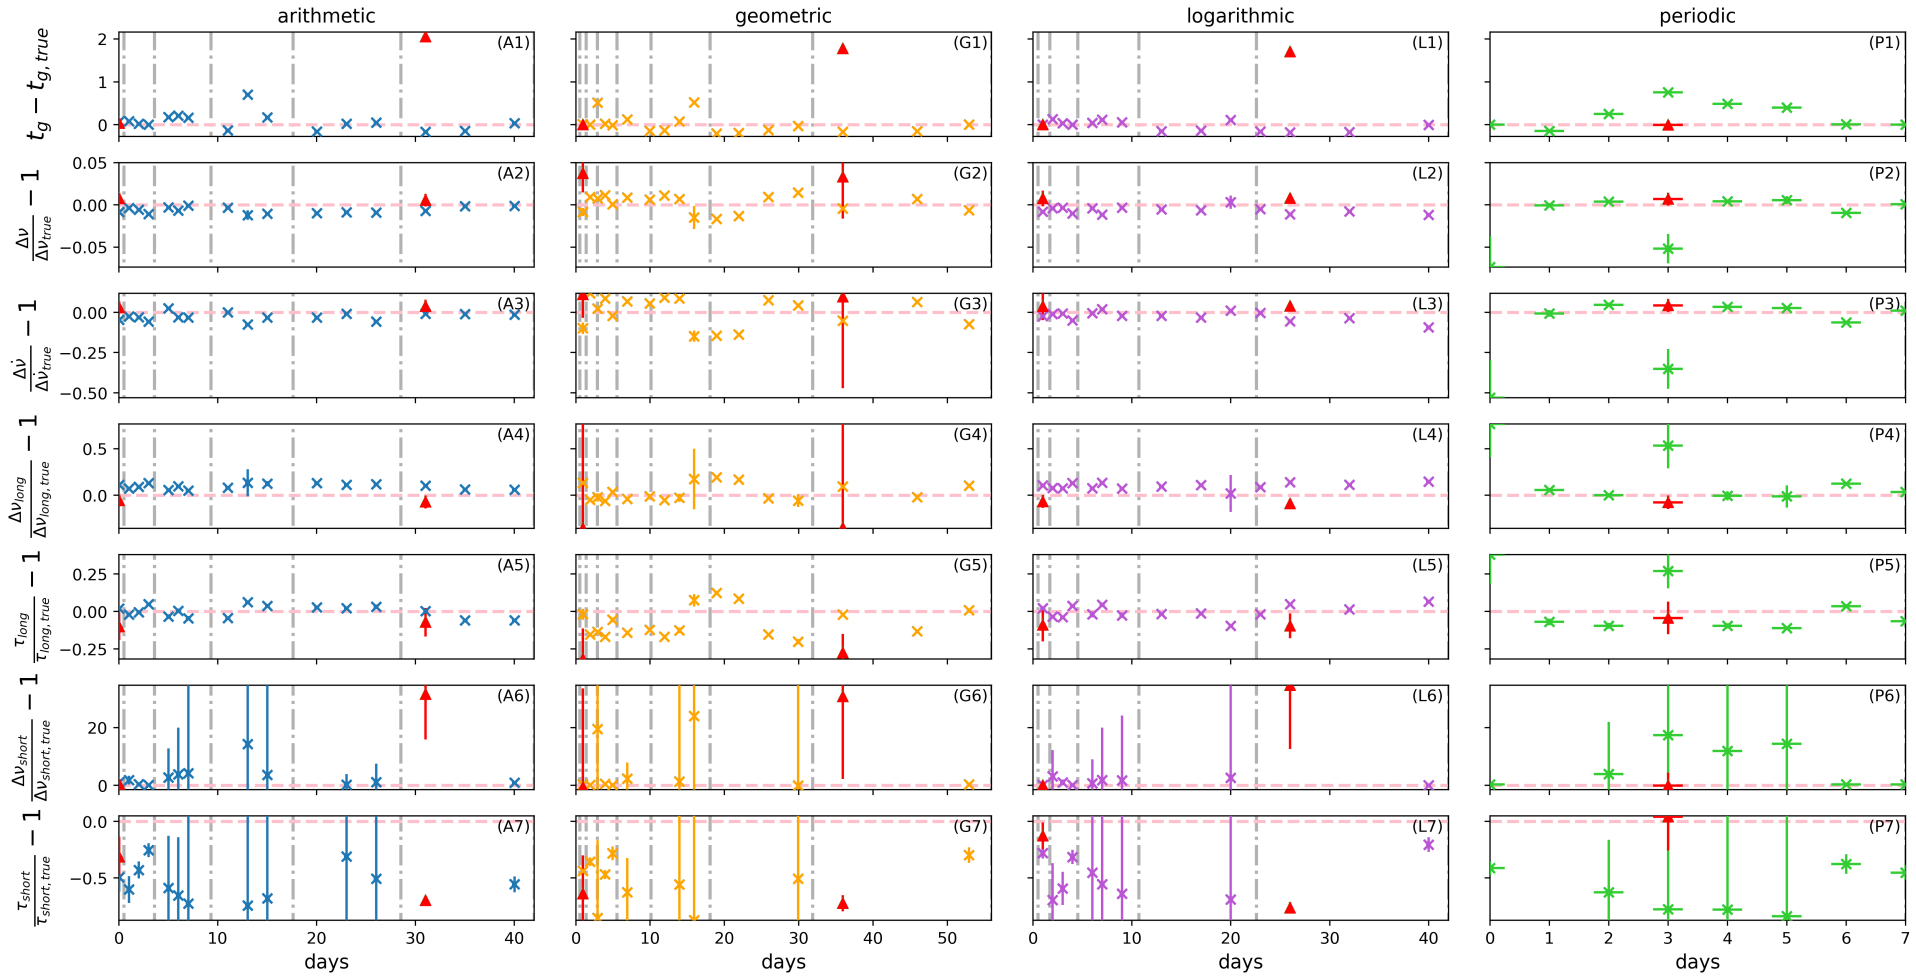

Figure 4.1: Comparison of fitted results for a glitch as described by Table 3.2. The x and y axis are shared between plots for easy comparison between scheduling strategies and parameters. The horizontal pink dashed line on every plot marks the true parameter value, and deviations from this are fractional deviations above or below the true value, except for the case of glitch epoch, t_g, where the y-axis is in days. The vertical grey dash-dot lines are representative of the observation times for each scheduling strategy, and the points' horizontal position on the plot shows where the glitch occurred in relation to the prior and subsequent observations. As expected, the results are best (closest to the magenta line) when the observation cadence is highest for all strategies. Red triangular points describe fits performed by an MCMC algorithm for comparison.

than the consistent 7d observation gap belonging to the periodic strategy. A glitch occurring in an infrequent observation region of an alternative strategy will reduce the number of TOAs coinciding with a quickly relaxing component. This implies that it may be better to employ a non-periodic strategy in cases where high accuracy parameters are a high priority priority, but not when glitch morphology is more important.

The arithmetic strategy's ability to retrieve short transient morphology (plots (A6) and (A7)) appear much more closer in ratio to periodic. Across 42 days, arithmetic scheduling allowed us to deduce a second exponential component in 12 of 15 positions along the strategy. Periodic scheduling allowed the retrieval of the short exponential in 7 of 8 sampled positions across 7 days. It should be noted that since the chosen sampling points are not equally distributed, these ratios are not indicative of the probability for a strategy to allow retrieval of both recovery components. This could be explained by the arithmetic scheduling strategy having a much lower $\Delta T_{\max}$ than the other strategies, which is a relation that has been suggested previously [32]. Arithmetic observations as described by Eq. 3.2 are much slower varying than geometric or logarithmic, and so the distance between observations at the start versus at the end of the strategy period, P_s is much smaller than for the other two alternative strategies. Fitting via arithmetic scheduling to deduce glitch morphology over a much larger portion of the strategy than geometric and logarithmic appears to exhibiting a much improved overall quality of fitted parameters. This ability makes it a good candidate for a scheduling strategy to improve glitch fitting.

For all parameters other than τ_{short} and $\Delta\nu_{\text{short}}$, for all strategies there appears to be little effect on the quality of retrieved results in scenarios where the parameters associated with the short recovery aren't fitted versus scenarios where they are. This is surprising and suggests little to no correlation between the short recovery parameters and all others. This being said, some correlation is most clearly visible in the effects which fitting the second exponential has on the glitch epoch, t_g. A larger glitch would require more time to decay to describe a set of TOAs, therefore pushing the fitted solution to an earlier t_g. All strategies have over-estimated the glitch epoch by up to 0.5d in the cases where the short timescale response was fitted, and underestimated in the cases where it was not. This effect can be best seen in the logarithmic and arithmetic strategies (plots (A1) and (L1)), where the data-points associated with unfit secondary exponentials are below (early) the true t_g at the pink line. This relation can similarly be seen in the corner plot, Figure 4.2, where there is a clear negative correlation between τ_{short} and t_g.

The biggest deviations from the true values are seen in parameter $\Delta\nu_{\text{short}}$, reaching overestimates by over 20 times the true value in a few cases. Every situation where $\Delta\nu_{\text{short}}$ is significantly overestimated, τ_{short} is underestimated by up to $\sim 50\%$. We can understand this as a correlation in Figure 4.2, which shows a negative correlation between these two numbers. The true short timescale from the fitted data in Figure 4.1 is 5d, which is of order unity and lower than the average observation cadence of the strategies, this means that when a glitch occurs, there is a high likelihood that most of the transient component will have decayed before the first observation occurs post-glitch, poorly constraining the parameters. Overestimates by such a significant degree are often indication that the glitch morphology cannot be reasonably ascertained from the data (i.e. it is unclear if a second exponential component is even present on the glitch). We tried to ensure only fitting for an additional component if there was convincing enough visual evidence on the residuals, but naturally this is subject to the biases discussed in 3.3. Further investigation could be done via a Bayesian model comparison [39].

Some fitting, indicated by the red triangles on Figure 4.1, was performed using an MCMC method. It is generally accepted that fitting pulsar properties with Bayesian statistics is superior to other methods [35, 39]. However, in some parameters, some of the worst fit values were found via this method, indicating that biases in manually fitted results had a much larger effect than was expected, causing comparatively better fits. In manually fitting, more information was known about the pulsar than would be known by an astronomer for a real pulsar and therefore our results were artificially improved over realistic model solutions. An additional reason for these poor values is that our MCMC fits were performed with an enforced strict two recovery exponentials, without allowing for fitting only one in scenarios where the second one was not visible, leading to the poorly constrained numbers, which again could be improved by Bayesian model comparison [39].

4.3 Varying the Short Recovery Timescale

Upon repeating the manual fitting for $\tau_{\text{short}} = 10$d, we present only the timescale plots alongside those for $\tau_{\text{short}} = 5$d in Figure 4.3. For the sake of space, the other parameters were omitted for the report and can be found in Appendix A. It is important to note that for a given strategy, the observation and scheduling was completely identical. It should also be noted that the random manifestation of timing noise was simulated to also be identical in both cases (Figure 3.2b).

As was pointed out previously, seemingly no strategy excels in retrieving the shorter of two timescales for the $\tau_{\text{short}} = 5$d case. However, once the shorter exponential was doubled to 10d, the periodic scheduling strategy was able to much more consistently retrieve accurate values for τ_{short}. This similarly has a knock-on effect to some of the other parameters via their correlations. This suggests that there are some limiting scenarios or particular values of τ_{short} where some strategies may be better or worse in retrieval.

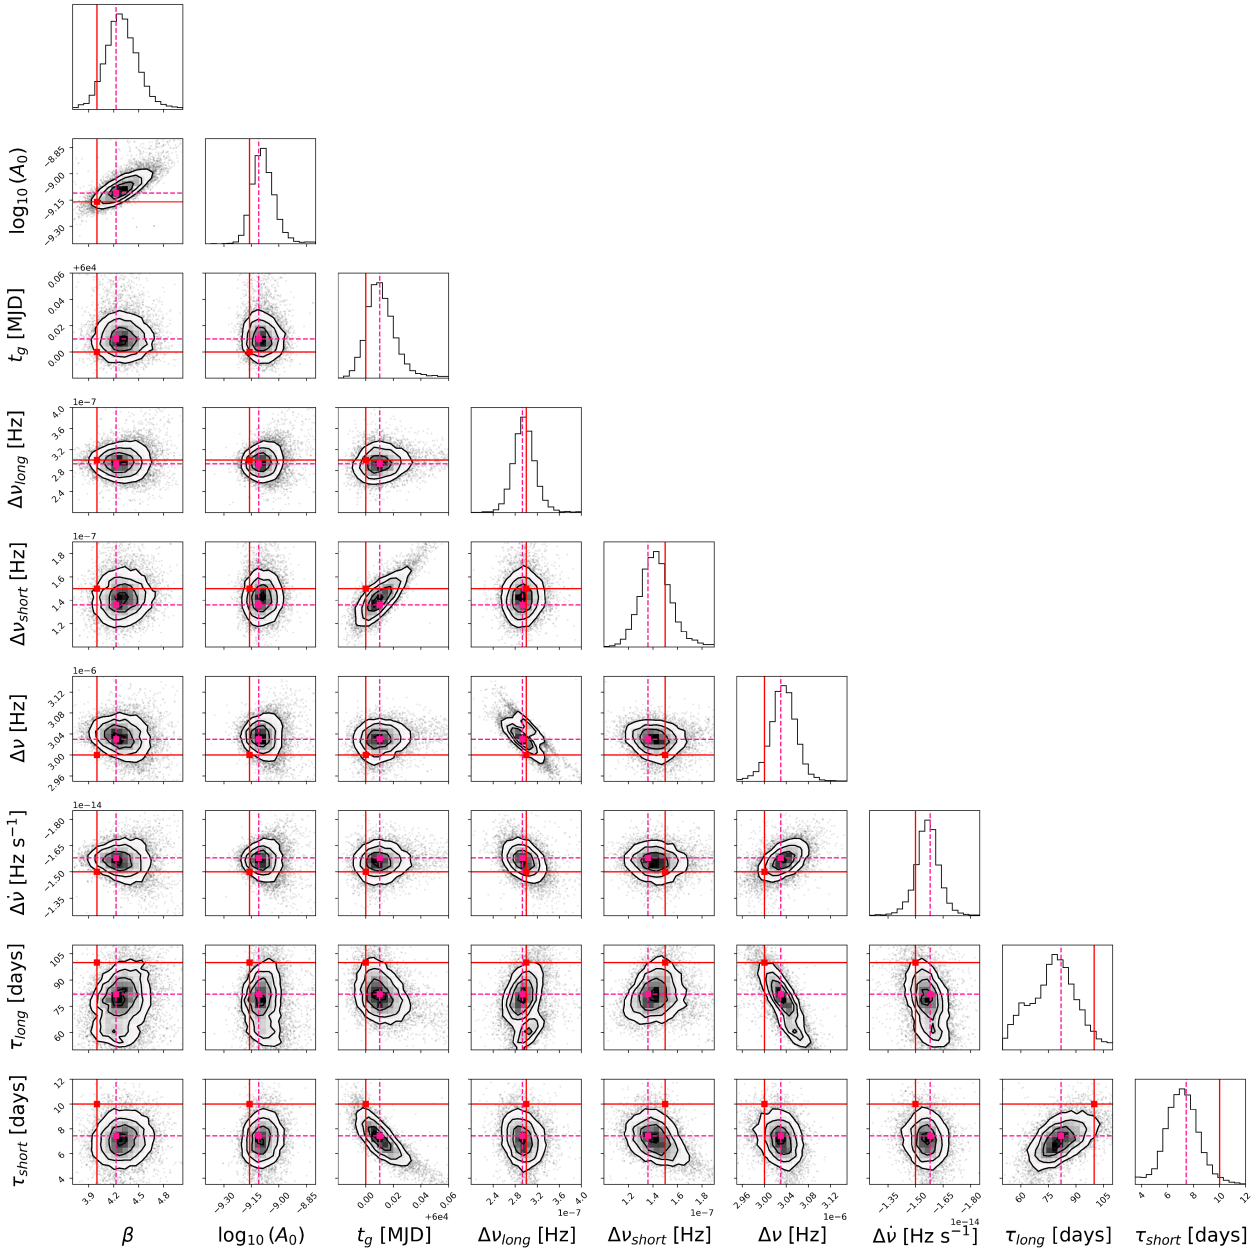

Figure 4.2: Corner plot demonstrating correlations between glitch parameters found via MCMC fitting. As shown in Figure 3.3, there is a negative correlation between $\tau_{d,\text{short}}$ and $\Delta\nu_{d,\text{short}}$. The red lines are representative of the true values and the pink dashed lines are the highest likelihood values according to the histograms. This plot was generated in the case where $\tau_{\text{short}} = 10$d with the glitch occurring at an average cadence location in the arithmetic strategy. Plot produced via [41].

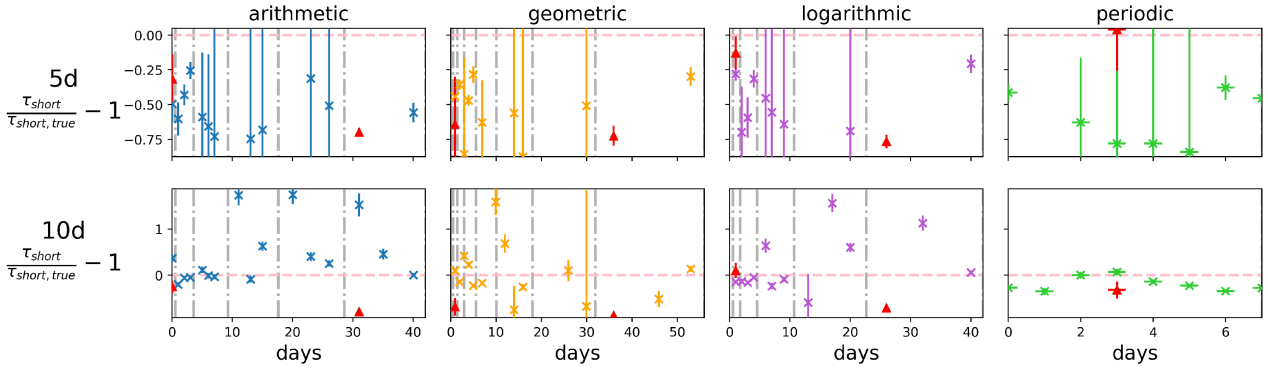

Figure 4.3: Comparison of fitted results for two different glitches with timescales $\tau_{\text{short}} = 5\text{d}$ and $\tau_{\text{short}} = 10\text{d}$ as described by Table 3.2. The x and y axis are shared between plots for easy comparison between scheduling strategies and timescales. The horizontal pink dashed line on every plot marks the true τ_{short}, and deviations from this are fractional deviations above or below the true value. The vertical grey dash-dot lines are representative of the observation times for each scheduling strategy, and the points' horizontal position on the plot shows where the glitch occurred in relation to the prior and subsequent observations. As expected, the results are best (closest to the magenta line) when the observation cadence is highest for all strategies. Red triangular points describe fits performed by an MCMC algorithm for comparison.

The relation to retrieved morphology and time until next TOA still persists in the case of $\tau_{\text{short}} = 10\text{d}$, but it is clear and expected that all strategies are able to retrieve the full morphology in more cases than when $\tau_{\text{short}} = 5\text{d}$, suggesting that there might be a particular relation between TOA gap size and retrieved response timescale. For instance, below a particular threshold value proportional to τ_{short}, where an observation occurring in less time than this threshold post-glitch will allow it to be seen. If such a relation exists, it is also likely to be influenced by the current cadence and the time to subsequent observations, as to avoid morphology being informed by a singular TOA. Though it should be noted that it it takes $\sim 4 \times \tau_{\text{short}}$ before one can consider the effects of the response to have fully dissipated, which even in the case of $\tau_{\text{short}} = 5$, is less than or equal to the maximum observation wait time, ΔT_{max}, for arithmetic, periodic and geometric. It is clear that these strategies do not have enough information to discern the short response in locations where more than one TOA would be effected by its transient effect, suggesting that the size of the effect, or magnitude of $\Delta\nu_{\text{short}}$, is also important to morphology retrieval.

It is first worth noting that it appears the short timescale was best retrieved in the 10d case. It can be seen that there are a few cases, particularly when the glitch occurred at the beginning of a TOA gap, where τ_{short} was significantly overestimated, but all other points are much more in accordance to the true value of 10d. This is expected as there is a much higher likelihood of the exponential recovery affecting a handful of TOAs, and much better constraining its shape and therefore timescale.

An structure apparent in the comparison of the 10d and 5d case is that for incorrect values, the longer timescale is more often overestimate and the shorter timescale is more often underestimated. It is not immediately clear why this is the case. The 5d τ_{short} can be attributed to too few TOAs significantly influenced by the relaxation component of the glitch, leading to an underestimate based on the few TOAs affected, but it is not expected for the inverse (overestimates) to hold true for a longer timescale. A likely explanation for this phenomenon could be as a result of the timing noise, which is described in Figure 3.2b. This particular manifestation is always decreasing and could easily be mistaken as part of the recovery component. This also explains why when the timing noise was also fitted for, as shown by the red Bayesian-fitted points, these are underestimated similarly to the 5d exponential case, as expected.

5 Discussion

This experiment has been performed with restrictive assumptions previously and the results could not rule out that alternative scheduling strategies could provide improved or equal quality glitch parameter estimates [32]. In this project, the generated data and fitting processes are much more representative of real data, and therefore conclusions made can be applied to a wider range of scenarios.

The obvious limiting factor to this research is from the usage of a single glitch size. While this is justified in study of PSR B1800–21 and many other pulsars which exhibit similarly large glitches, it does somewhat limit the applications to the field of pulsar astronomy. A similar effect originates from testing only a single manifestation of timing noise at a single β and A_0. There has been plenty of research into the interplay of

glitches and timing noise and one can very easily be mistaken for the other in scenarios of infrequent observing, and so if in future research the glitch sizes of $\Delta\nu/\nu$, $\Delta\dot{\nu}/\dot{\nu}$, etc. were to vary, care must be taken to appropriately manage the timing noise. Inversely, if stronger timing noise is investigated, care must be taken in understanding the constraints on visible glitch sizes [42].

Generated timing noise can exhibit many different shapes and structures in the residuals for the same power spectrum, making it possible to vary the measured data without changing any of the parameters used in generation. It is intuitive to believe that some manifestations may appear as additional structures on glitches due to the random wandering (red) nature of the noise. While not presented in this project, the simulations required for the data analysis were performed on two different manifestations of the timing noise presented. The difference in the results for the two were only marginally different, but it was decided that investigation into the effects of differing manifestations would require many more than two for appropriate discussion, and so it was omitted.

The four strategies presented in this project were chosen somewhat arbitrarily. Our goal was to ensure that there was some variety in structure (as shown in Figure 3.1), whilst maintaining a degree of simplicity to assist with algorithmic implementation. There exist countless other strategies which could be created; any time varying function could be converted to an observation strategy dictating cadence at any particular time. Every scheduling method here is always increasing in observation gap, when the reverse may be preferable, or even a strategy which increases then decreases. Scheduling done more randomly may act as a better comparative baseline to periodic observations, but random observation gaps about some mean value reduces to quasi-periodicity with small enough variances. In a more realistic scenario where many pulsars are being timed simultaneously and telescopes switch between hundreds of sources within the span of days, it may be better to focus on developing algorithms which can more fairly allocate telescope time to many sources, rather than focusing on the best way to observe a single pulsar. With SKA and the many galactic pulsars it is predicted to discover, there will be a shortage of telescopes to observe them and so the already unfeasible quasi-periodic scheduling will be increasingly difficult to maintain.

If different parameters are better retrieved via different observation methods then an idealistic scheduling system would allocate both telescope time and observation strategy informed by research intent. Our results have shown that periodic may be the best strategy for simple pulsar timing but in research on neutron star interiors it may be integral to get good constraints on a glitch's morphology, by which a varying strategy would have some > 0 chance to well define relaxation behaviour, whereas infrequent periodic observations at the same average cadence may only every be able to have a maximum of one TOA affected by any short-term transient properties.

Due to their regularity, one can somewhat anticipate glitches in Vela-like glitchers, which could be utilised in good scheduling, although it can never be known the exact moment a glitch will occur and the degree of predictability is minimal [4]. To utilise this coarse foreknowledge of when a glitch might occur, it may be beneficial to slowly decrease the distance between observations, potentially arithmetically or geometrically with $k_g \sim 1$ on timescales of months to years. This way, if the glitch occurs earlier than expected, the observations are already more frequent than they were previously, yielding better results. Similarly, if the pulsar glitches late, this could be seen as lucky if the observing frequency had continued to increase in anticipation. What is already an interesting prospect would be amplified if there are new pulsars discovered which glitch more predictably or if methods become available to predict when a glitch may occur.

A relatively easy but long-term improvement to research in this area would exist as the creation of a bespoke system which fully automates the processes of TOA generation and fitting of the values according to a number of settings. The main limiting factor for the amount of data that was collected was required human-hours in the manual fitting process, which could be eliminated, alongside the associated biases, if Bayesian fits with equal priors were performed for every glitch location on every different dataset. A good portion of the code utilised for the project so far is already modular and could be extended to integrate with `run_enterprise` automatically, starting a Bayesian fit immediately after the generation of the dataset. This was deemed unfeasible for the project in the scope of the software required, but its existence would vastly speed up data collection, allowing for simple access to variations in glitch sizes and other parameters. The current method detailed in Section 3 allowed us time to vary a single parameter, τ_{short}, in the time allocated, but automation would also allow for an increase in the number of glitch locations fitted, in addition to the possibility of multiple timing noise manifestations at a point, which can be averaged for improved insight for a glitch at a particular location within a strategy. Only at the cost of additional computation time.

With improved methods for data collection, the changes investigated can be much higher resolution, such as exploring small incremental changes to τ_{short} rather than only $\tau_{\text{short}} = 5, 10\text{d}$. This project has evidenced that there may be limiting values of time before next observation and short response timescale which dictate when particular structures may be appropriately fitted by MCMC or through manual fitting. Exploring these relations further may inform some empirical relation between observation distribution and ability to retrieve glitch morphology and parameter estimation, which would then further inform future research towards the most ideal observation strategy in particular scenarios, which is ultimate goal of understanding non-periodic

scheduling. It is also likely that it becomes known that high frequency periodic scheduling is the best for pulsar observation, but it should not be assumed until understood.

6 Conclusion

We have presented four different scheduling strategies: arithmetic, geometric, logarithmic and periodic. Of these, periodic acted as the baseline for comparison akin to somewhat realistic rigidly periodic pulsar observations. Of the others, each is structurally different from one other but maintain consistent properties of always increasing and cyclically repeating once observation gaps become too large. Each strategy is parametrised by a strategy constant, k, and a start and maximum size of observations, ΔT_{start} and ΔT_{max}.

Timing data for a simulated version of PSR B1800–21 was generated in accordance to each observation strategy over a common amount of time, resulting in a roughly similar number of total observations to maintain fairness in comparison between the strategies. On average, each scheduling strategy observed the fake pulsar once every 7 days. Generated datasets were created for multiple shifts to every TOA up to each strategy's full cyclic period as to simulate a glitch occurring at regions of differing observation density. The data was then algorithmically fitted with manual assistance to retrieve parameters and their associated errors.

Simulating timing noise in the generated data introduced some issues with manual fitting methods, justifying the usage of more complex fitting methods, which shed light on unintentional fitting biases which arose in pulsar simulation. We concluded that we cannot comment on the full effects of timing noise on results without simulating many more manifestations at a particular parametrisation.

It was found that for an average observation cadence of 7 days, an arithmetically varying scheduling strategy appears to retrieve glitch morphology better than or equal to a periodic scheduling method. Variations in observation distance leading to regions of low and high cadence result in "good" and "bad" regions for a glitch to occur in, allowing for a chance to retrieve otherwise invisible behaviours. It was expected and shown that a higher cadence of observations is advantageous for parameter fitting, but also that there is an apparent high level of dependence on the time until the first post-glitch observation in retrieval of glitch morphology. Arithmetic scheduling is a slower varying strategy than geometric or logarithmic and therefore does not have very large TOA gaps, which would obscure the morphology.

Comparing a set of manually fitted pulsar results revealed information about parameter correlation which was presented in a corner plot. It appears as though, differing from previous work [32], the short transient components are not significantly correlated to the other glitch parameters, though they are correlated highly to each other. This behaviour emerges through other parameters being insignificantly affected by the model fitting for more than a single transient recovery.

To conclude we suggest that an alternative scheduling strategy, particularly the arithmetic strategy described by Equation 3.2, would be beneficial for the observation of Vela-like glitching pulsars such as B1800–21, particularly if the observation of fast transient components is desired. Periodic scheduling will always have a small number of TOAs affected by any short relaxation, but with an arithmetic observation cadence, glitches have a chance to occur at a time when observation frequency is higher. This allows for more accurate parameter fitting without changing the overall frequency of observations.

Bibliography

[1] A. G. Lyne and F. Graham-Smith. *Pulsar astronomy*. eng. 4th edition. Cambridge astrophysics ; 48. Cambridge ; Cambridge University Press, 2012. ISBN: 9781107010147.

[2] F. Weber et al. "Pulsars as astrophysical laboratories for nuclear and particle physics". eng. In: *Progress in particle and nuclear physics* 59.1 (2007), pp. 94–113. ISSN: 0146-6410.

[3] G. Agazie et al. "The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background". In: ApJ 951.1, L8 (July 2023), p. L8. DOI: 10.3847/2041-8213/acdac6. arXiv: 2306.16213 [astro-ph.HE].

[4] D. Antonopoulou et al. "Pulsar glitches: observations and physical interpretation". In: *Reports on Progress in Physics* 85.12 (Dec. 2022), p. 126901. DOI: 10.1088/1361-6633/ac9ced. URL: https://dx.doi.org/10.1088/1361-6633/ac9ced.

[5] D. R. Lorimer and M. Kramer. *Handbook of pulsar astronomy*. eng. Cambridge observing handbooks for research astronomers ; 4. Cambridge: Cambridge University Press, 2005. ISBN: 0521828236.

[6] A. Basu et al. "Observed glitches in eight young pulsars". In: *Monthly Notices of the Royal Astronomical Society* 491.3 (Nov. 2019), pp. 3182–3191. ISSN: 0035-8711. DOI: 10.1093/mnras/stz3230. eprint: https://academic.oup.com/mnras/article-pdf/491/3/3182/31490452/stz3230.pdf. URL: https://doi.org/10.1093/mnras/stz3230.

[7] F. E. Marshall et al. "The Big Glitcher: The Rotation History of PSR J0537-6910". In: *The Astrophysical Journal* 603.2 (Mar. 2004), pp. 682–689. DOI: 10.1086/381567.

[8] P. V. Padmanabh et al. "The MPIfR–MeerKAT Galactic Plane Survey – I. System set-up and early results". In: *Monthly Notices of the Royal Astronomical Society* 524.1 (June 2023), pp. 1291–1315. ISSN: 1365-2966. DOI: 10.1093/mnras/stad1900. URL: http://dx.doi.org/10.1093/mnras/stad1900.

[9] L. Dunn et al. "Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches". In: *Monthly Notices of the Royal Astronomical Society* 504.3 (Apr. 2021), pp. 3399–3411. ISSN: 1365-2966. DOI: 10.1093/mnras/stab1097. URL: http://dx.doi.org/10.1093/mnras/stab1097.

[10] M. Ruderman and P. Sutherland. *Theory of pulsars: polar gaps, sparks, and coherent microwave radiation*. en. Tech. rep. COO–2271-31, 4258471. Jan. 1974, COO–2271–31, 4258471. DOI: 10.2172/4258471. URL: http://www.osti.gov/servlets/purl/4258471/.

[11] A. Philippov et al. "Origin of Pulsar Radio Emission". en. In: *Physical Review Letters* 124.24 (June 2020), p. 245101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.124.245101. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.245101.

[12] J. H. Taylor. "Pulsar Timing and Relativistic Gravity". In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660 (Oct. 1992), pp. 117–134. DOI: 10.1098/rsta.1992.0088.

[13] J. Wang et al. "A comparative analysis of pulse time-of-arrival creation methods". In: *Astronomy & Astrophysics* 658 (Feb. 2022), A181. ISSN: 1432-0746. DOI: 10.1051/0004-6361/202141121. URL: http://dx.doi.org/10.1051/0004-6361/202141121.

[14] G. B. Hobbs et al. "tempo2, a new pulsar-timing package - I. An overview: tempo2, a new pulsar-timing package - I. Overview". In: *Monthly Notices of the Royal Astronomical Society* 369.2 (May 2006), pp. 655–672. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2006.10302.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2006.10302.x.

[15] C. M. Espinoza et al. "New long-term braking index measurements for glitching pulsars using a glitch-template method". In: *Monthly Notices of the Royal Astronomical Society* 466.1 (Apr. 1, 2017), pp. 147–162. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stw3081. URL: https://academic.oup.com/mnras/article-lookup/doi/10.1093/mnras/stw3081.

[16] S. Zhou et al. *Pulsar Glitches: A Review*. 2022. arXiv: 2211.13885 [astro-ph.HE]. URL: https://arxiv.org/abs/2211.13885.

[17] I. Cognard and D. C. Backer. "A Microglitch in the Millisecond Pulsar PSR B1821-24 in M28". In: ApJ 612.2 (Sept. 2004), pp. L125–L127. DOI: 10.1086/424692. arXiv: astro-ph/0407546 [astro-ph].

[18] J. W. McKee et al. "A glitch in the millisecond pulsar J0613-0200". In: *Monthly Notices of the Royal Astronomical Society* 461.3 (June 2016), pp. 2809–2817. ISSN: 0035-8711. DOI: 10.1093/mnras/stw1442. eprint: https://academic.oup.com/mnras/article-pdf/461/3/2809/8107796/stw1442.pdf. URL: https://doi.org/10.1093/mnras/stw1442.

[19] G. Baym et al. "Superfluidity in Neutron Stars". en. In: *Nature* 224.5220 (Nov. 1969). Publisher: Nature Publishing Group, pp. 673–674. ISSN: 1476-4687. DOI: 10.1038/224673a0. URL: https://www.nature.com/articles/224673a0.

[20] G. Baym et al. "Spin Up in Neutron Stars : The Future of the Vela Pulsar". en. In: *Nature* 224.5222 (Nov. 1969). Publisher: Nature Publishing Group, pp. 872–874. ISSN: 1476-4687. DOI: 10.1038/224872a0. URL: https://www.nature.com/articles/224872a0.

[21] C. M. Espinoza et al. "A study of 315 glitches in the rotation of 102 pulsars: A study of 315 glitches in 102 pulsars". In: *Monthly Notices of the Royal Astronomical Society* 414.2 (Apr. 2011), pp. 1679–1704. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2011.18503.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2011.18503.x.

[22] A. Basu et al. "The Jodrell bank glitch catalogue: 106 new rotational glitches in 70 pulsars". In: MNRAS 510.3 (Mar. 2022), pp. 4049–4062. DOI: 10.1093/mnras/stab3336. arXiv: 2111.06835 [astro-ph.HE].

[23] M. A. Ruderman. "Crystallization and Torsional Oscillations of Superdense Stars". In: Nature 218.5147 (June 1968), pp. 1128–1129. DOI: 10.1038/2181128a0.

[24] P. W. Anderson and N. Itoh. "Pulsar glitches and restlessness as a hard superfluidity phenomenon". In: Nature 256.5512 (July 1975), pp. 25–27. DOI: 10.1038/256025a0.

[25] B. Haskell and A. Melatos. "Models of pulsar glitches". In: *International Journal of Modern Physics D* 24.03 (Feb. 2015), p. 1530008. ISSN: 1793-6594. DOI: 10.1142/s0218271815300086. URL: http://dx.doi.org/10.1142/S0218271815300086.

[26] A. Lyne et al. "Switched Magnetospheric Regulation of Pulsar Spin-Down". In: *Science* 329 (July 2010). ADS Bibcode: 2010Sci...329..408L, p. 408. ISSN: 0036-8075. DOI: 10.1126/science.1186683. URL: https://ui.adsabs.harvard.edu/abs/2010Sci...329..408L.

[27] B. Haskell. "Tkachenko modes in rotating neutron stars: The effect of compressibility and implications for pulsar timing noise". In: *Physical Review D* 83.4 (Feb. 2011). Publisher: American Physical Society, p. 043006. DOI: 10.1103/PhysRevD.83.043006. URL: https://link.aps.org/doi/10.1103/PhysRevD.83.043006.

[28] A. Melatos and B. Link. "Pulsar timing noise from superfluid turbulence". In: *Monthly Notices of the Royal Astronomical Society* 437.1 (Nov. 2013), pp. 21–31. ISSN: 1365-2966. DOI: 10.1093/mnras/stt1828. URL: http://dx.doi.org/10.1093/mnras/stt1828.

[29] G. Hobbs et al. "An analysis of the timing irregularities for 366 pulsars". In: *Monthly Notices of the Royal Astronomical Society* 402.2 (Feb. 2010), pp. 1027–1048. ISSN: 1365-2966. DOI: 10.1111/j.1365-2966.2009.15938.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2009.15938.x.

[30] A. Parthasarathy et al. "Timing of young radio pulsars – I. Timing noise, periodic modulation, and proper motion". In: *Monthly Notices of the Royal Astronomical Society* 489.3 (Aug. 2019), pp. 3810–3826. ISSN: 1365-2966. DOI: 10.1093/mnras/stz2383. URL: http://dx.doi.org/10.1093/mnras/stz2383.

[31] M. E. Lower et al. "The UTMOST pulsar timing programme - II. Timing noise across the pulsar population". In: *Monthly Notices of the Royal Astronomical Society* 494 (May 2020). Publisher: OUP ADS Bibcode: 2020MNRAS.494..228L, pp. 228–245. ISSN: 0035-8711. DOI: 10.1093/mnras/staa615. URL: https://ui.adsabs.harvard.edu/abs/2020MNRAS.494..228L.

[32] L. Maclean and J. Salisbury. "Improving Pulsar Timing Precision". MPhys Project Research, Unpublished, J. Salisbury's Writeup. University of Manchester, 2025. URL: https://nyft-i.github.io/assets/pdf/jsalis_improving_pulsar_timing_precision_i.pdf.

[33] P. Weltevrede et al. "The glitch-induced identity changes of PSR J1119-6127: The glitch-induced identity changes of PSR J1119-6127". In: *Monthly Notices of the Royal Astronomical Society* 411.3 (Mar. 1, 2011), pp. 1917–1934. ISSN: 00358711. DOI: 10.1111/j.1365-2966.2010.17821.x. URL: https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2010.17821.x.

[34] W. H. Press. *Numerical recipes in C the art of scientific computing / William H. Press ... [et al.]*. eng. Cambridge, 1995.

[35] L. Lentati et al. "temponest: a Bayesian approach to pulsar timing analysis". In: *Monthly Notices of the Royal Astronomical Society* 437.3 (Dec. 2013), pp. 3004–3023. ISSN: 1365-2966. DOI: 10.1093/mnras/stt2122. URL: http://dx.doi.org/10.1093/mnras/stt2122.

[36] D. Foreman-Mackey et al. “emcee: The MCMC Hammer”. In: PASP 125.925 (Mar. 2013), p. 306. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].

[37] J. A. Ellis et al. *ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust PulsaR Inference SuitE*. Zenodo. Sept. 2020. DOI: 10.5281/zenodo.4059815. URL: https://doi.org/10.5281/zenodo.4059815.

[38] M. J. Keith et al. *run_enterprise*. 2022.

[39] Y. Liu et al. “Measuring glitch recoveries and braking indices with Bayesian model selection”. In: *Monthly Notices of the Royal Astronomical Society* 532.1 (July 2024), pp. 859–882. ISSN: 0035-8711. DOI: 10.1093/mnras/stae1499. URL: https://doi.org/10.1093/mnras/stae1499.

[40] T. R. Clifton and A. G. Lyne. “High-radio-frequency survey for young and millisecond pulsars”. en. In: *Nature* 320.6057 (Mar. 1986), pp. 43–45. ISSN: 1476-4687. DOI: 10.1038/320043a0. URL: https://www.nature.com/articles/320043a0.

[41] D. Foreman-Mackey. “corner.py: Scatterplot matrices in Python”. In: *The Journal of Open Source Software* 1.2 (2016), p. 24. DOI: 10.21105/joss.00024. URL: https://doi.org/10.21105/joss.00024.

[42] C. M. Espinoza et al. “Neutron star glitches have a substantial minimum size”. In: *Monthly Notices of the Royal Astronomical Society* 440.3 (Apr. 2014), pp. 2755–2762. ISSN: 1365-2966. DOI: 10.1093/mnras/stu395. URL: http://dx.doi.org/10.1093/mnras/stu395.

A Additional results: $\tau_{\mathbf{short}} = 10\mathbf{d}$

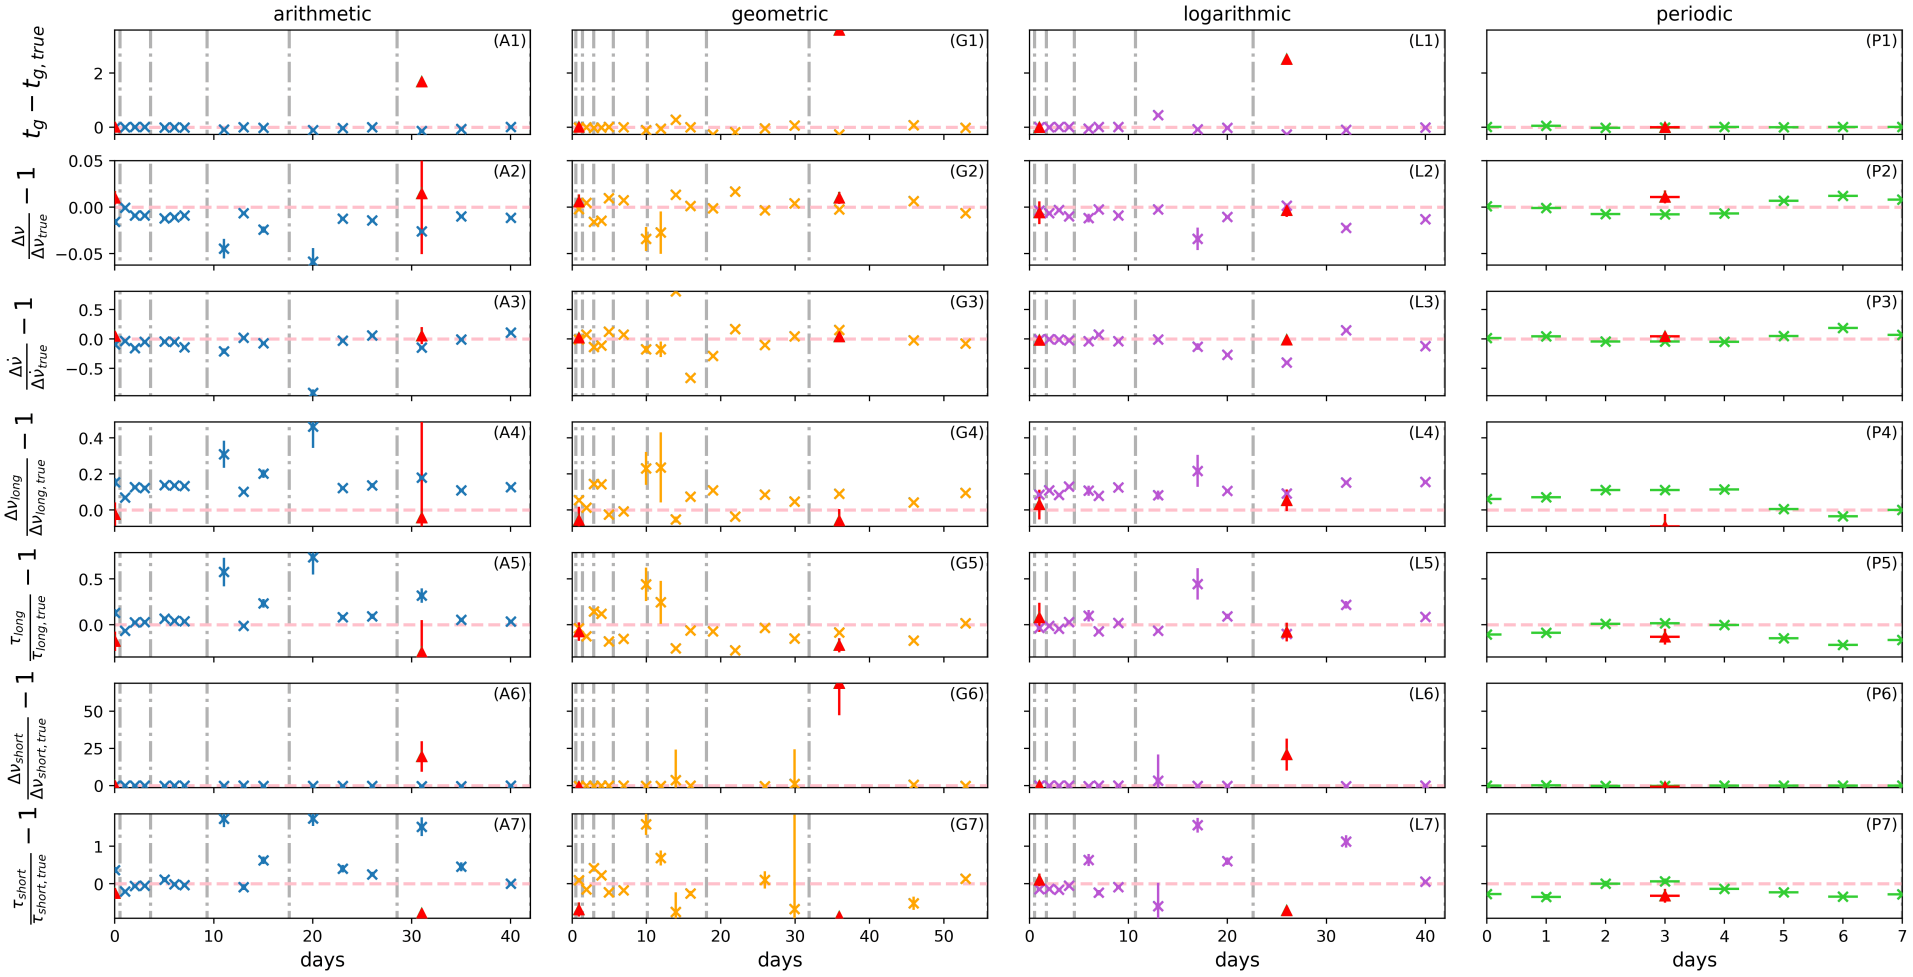

Figure A.1: Comparison of fitted results for a glitch as described by Table 3.2. The y and x axis are scaled for easy comparison between scheduling strategies and parameters. The horizontal pink line on every plot marks the true value, and deviations from this are fractional deviations above or below the true value, except for the case of glitch epoch, t_g, where the y-axis is in days. The vertical grey dash-dot lines are representative of the observation times for each scheduling strategy, and the points' horizontal position on the plot shows where the glitch occurred in relation to the prior and subsequent observations. As expected, the results are best (closest to the magenta line) when the observation cadence is highest for all strategies.