CAMBRIDGE
UNIVERSITY PRESS

Research Article

Increasing the detectability of long-period and nulling pulsars in next-generation pulsar surveys

Garvit Grover ⓘD, Ramesh Bhat ⓘD, and Samuel McSweeney ⓘD

International Centre for Radio Astronomy Research, Curtin University, Bentley, WA, Australia

Abstract

Recent discoveries of multiple long-period pulsars (periods ∼10 s or larger) are starting to challenge the conventional notion that coherent radio emission cannot be produced by objects that are below the many theorised death lines. Many of the past pulsar surveys and software have been prone to selection effects that restricted their sensitivities towards long-period and sporadically emitting objects. Pulsar surveys using new-generation low-frequency facilities are starting to employ longer dwell times, which makes them significantly more sensitive in detecting long-period or nulling pulsars. There have also been software advancements to aid more sensitive searches towards long-period objects. Furthermore, recent discoveries suggest that nulling may be a key aspect of the long-period pulsar population. We simulate both long-period and nulling pulsar signals, using the Southern-sky MWA Rapid Two-meter (SMART) survey data as reference and explore the detection efficacy of popular search methods such as the fast Fourier transform (FFT), fast-folding algorithm (FFA) and single pulse search (SPS). For FFT-based search and SPS, we make use of the PRESTO implementation, and for FFA we use RIPTIDE. We find RIPTIDE's FFA to be more sensitive; however, it is also the slowest algorithm. PRESTO's FFT, although faster than others, also shows some unexpected inaccuracies in detection properties. SPS is highly sensitive to long-period and nulling signals, but only for pulses with high intrinsic signal-to-noise ratios. We use these findings to inform current and future pulsar surveys that aim to uncover a large population of long-period or nulling objects and comment on how to make optimal use of these methods in unison.

1. Introduction

Coherent radio emission from pulsars is thought to rely critically on pair production within their magnetospheres (Ruderman & Sutherland 1975). At longer rotation periods (P), the potential difference in the magnetosphere is suspected to be insufficient for pair production to occur, and hence, no emission can be produced (Ruderman & Sutherland 1975; Arons & Scharlemann 1979). This naturally places a constraint on the rotation period P and the characteristic magnetic field $B \propto \sqrt{P\dot{P}}$ of the pulsar, where $\dot{P}$ is the period derivative. The related constraints can be identified as 'death lines' on a P-$\dot{P}$ diagram. Many death lines have been proposed, each with different assumptions such as the magnetic field orientation, assumptions about the pair production, etc (e.g. Ruderman & Sutherland 1975; Chen & Ruderman 1993; Zhang, Harding, & Muslimov 2000; Rea et al. 2023). However, the one that most completely encompassed the pulsar population was by Zhang et al. (2000).

In recent years, there has been a flurry of discoveries of pulsars with very long periods that challenge some of the most conservative death line models, as shown in Fig. 1. These include a 12.1-s pulsar J2251-3711 discovered by Morello et al. (2020b), a 23.5-s pulsar J0250+5854 by Tan et al. (2018) and PSR J0901-4046 with a rotation period of 76 s (Caleb et al. 2022). These are currently

among the slowest known radio pulsars. Along with their relatively small $\dot{P}$, their positions on the P-$\dot{P}$ diagram are also near the edge of the furthest (least stringent) death line. This death line is shown as Zhang et al. (2000) III' in Fig. 1. The discoveries of such pulsars also suggest the possible existence of a larger population of long-period pulsars that could be uncovered using more optimal search strategies.

Remarkably, all three objects show some degree of 'nulling,' which is the cessation of detectable radio emission for a certain duration or number of rotations. For example, PSRs J2251-3711 and J0250+5854 have nulling fractions of 67% and 27%, respectively (Tan et al. 2018; Morello et al. 2020b), and PSR J0901-4046 shows nulling occurring on time scales shorter than a single pulse (Caleb et al. 2022). Assuming these objects are representative of the population, nulling may be a common property among these objects. It has been suggested that nulling is indicative of the radio emission becoming unsustainable for a pulsar and the nulling fraction increases progressively until there is no detectable emission from the pulsar (Ritchings 1976; Zhang, Gil, & Dyks 2007). However, no convincing correlation has been found between nulling and physical parameters such as P or $\dot{P}$ (Anumarlapudi et al. 2023; Sheikh & Macdonald 2021), though this may be also due to observational selection effects.

In addition to these three pulsars, a couple of pulsar-like sources with extremely long periods have also been detected in recent years. A noteworthy source is GLEAMX J162759.5-523504.3, with a rotation period of ∼18.18 min, discovered by Hurley-Walker et al. (2022). Another similar source, with an even
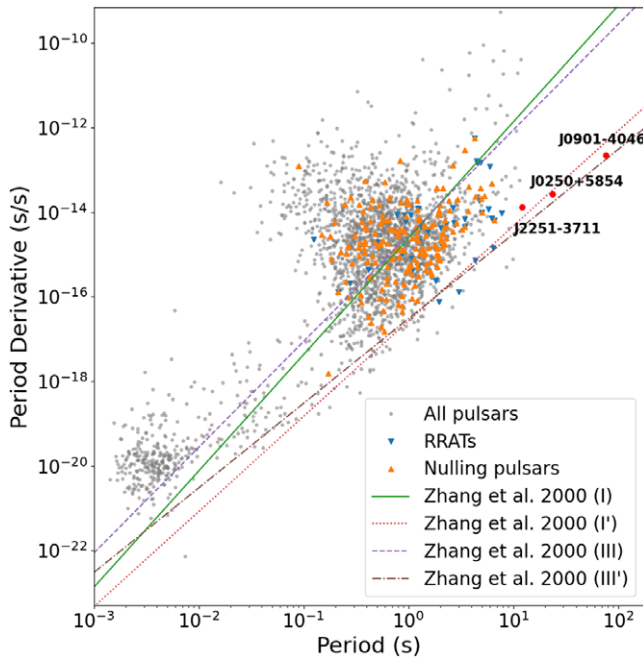
Figure 1. A $P\dot{P}$ diagram showing all pulsars (grey), nulling pulsars (orange) and RRATs (blue). Also highlighted are the three recent long-period pulsar discoveries (Tan et al. 2018; Morello et al. 2020b; Caleb et al. 2022). The various death lines, taken from Zhang et al. (2000), are also shown, as per the legend. These death lines are improvements based on the works of Ruderman & Sutherland (1975) and Chen & Ruderman (1993).

longer period of ∼21 min, has just been discovered (Hurley-Walker, Rea, & McSweeney 2023). The origins of this new class of sources remain currently unknown. The first source does seem to go into a quiescent state, which could be interpreted as an extended nulling state, but interestingly the second source has been active for over three decades. These discoveries have motivated Rea et al. (2023) to extend the original models for death lines.

Yet another recent addition is a ∼10 s pulsar discovered by Su et al. (2023) using the Five-hundred-metre Aperture Spherical Telescope (FAST). Intriguingly, over a 2.29-yr observational time span, there is no evidence of nulling in this object. Regardless, it sits well beyond the current death line models on the P-$\dot{P}$ diagram, thereby further strengthening the evidence of a possible long-period pulsar population. Uncovering a larger population of sources (both pulsars or pulsar-like objects) can help significantly advance our understanding of coherent radio emission in pulsars and current models for the death lines.

Death line models are typically modified to constrain new long-period pulsars as they are discovered; however, long-period pulsar discoveries are made difficult due to observational selection effects. Most large pulsar surveys that scan large swaths of the skies typically have dwell times between 2 and 5 min and therefore are less sensitive to long-period pulsar signals (Stovall et al. 2014; Han et al. 2021; Keith et al. 2010). Longer dwell times naturally increase the sensitivity to these objects, particularly those with long durations of nulling or rotation periods larger than tens of seconds. Specifically, longer observing duration increases the chance of detecting a larger number of pulses, especially when $P \sim 10$–100 s or longer. The emergence of wide-field instruments at low frequencies such as the Murchison Widefield Array (MWA; Tingay et al. 2013) and Low Frequency Array (LOFAR; van Haarlem et al.

2013) and their adaptation to large pulsar surveys (e.g. Bhat et al. 2023a; Sanidas et al. 2019), means longer dwell times are now affordable, at least at low frequencies. This increases the achievable sensitivity in general, particularly for those objects with rotation periods $\gtrsim 10$ s, or with long null durations (∼tens of minutes). The LOFAR Tied-Array All-Sky Survey (LOTAAS; Sanidas et al. 2019) has a 60 min dwell time, and the Southern-sky MWA Rapid Two-meter (SMART) survey uses an 80 min dwell time. The SMART survey uses the MWA to observe the whole southern sky between 140–170 MHz. The data has 100 μs time resolution and 10 kHz frequency resolution, and are stored as raw voltages, and can be beamformed offline to tessellate the full field of view using thousands of sensitive tied-array (phased array) beams. For more details on the SMART survey see Bhat et al. (2023a). In detecting a long-period or highly intermittent population of pulsars, these surveys will have a competitive advantage.

1.1 Overview of search methods

In addition to the use of more optimal survey strategies, it is also important to adopt search algorithms that are sensitive to sporadic, intermittent and or long-period objects. In principle, long-duration data from surveys like the SMART can also be exploited to search for pulsars in binary systems; however, our focus will primarily be on isolated or single pulsars. There are three popular pulsar searching methods, including the fast Fourier transform (FFT) based search, the fast folding algorithm (FFA; first described by Staelin 1969) and the single pulse search (SPS). Fourier methods have been the primary ones used in pulsar searching and have been very successful, having discovered the vast majority of the pulsars currently known. However, in searching for long-period pulsars, Fourier-based methods have been shown to be inherently less sensitive towards long-period and narrow duty-cycle signals, which is generally the case for most long-period pulsars (van Heerden, Karastergiou, & Roberts 2017; Cameron et al. 2017; Singh et al. 2022; Kondratiev et al. 2009; Morello et al. 2020a; Parent et al. 2018). A major obstacle is the presence of low-frequency 'red' noise that arises from slowly varying power levels in the system (e.g. due to gain variations), or radio frequency interference (RFI) that is persistent in nature. This tends to mask long-period signals in Fourier space. Moreover, this method relies on the signal to be periodic to make a significant detection and is therefore inherently less sensitive towards the sporadic emission from pulsars that tend to null for long durations. PRESTO (Ransom 2001), includes algorithms for suppressing red noise, and has not been robustly tested for its efficacy in detecting such objects that emit sporadic or intermittent signals.

FFA is a time-domain periodicity search. It searches for pulsars by efficiently folding signals at multiple trial periods. There have been a few different implementations of this algorithm: ffaGo (Parent et al. 2018), ffancy (Cameron et al. 2017) and RIPTIDE (Morello et al. 2020a), all highlighting the advantages of using FFA over FFT-based searches, especially for long-period pulsars. The algorithm has been demonstrated in its ability to find new pulsars, having recently detected 6 new pulsars in the GMRT High-Resolution Southern-sky (GHRSS) survey (Singh et al. 2023). FFA has also been shown to be effective at finding long-period pulsars, and in fact, has been able to successfully make independent detections of PSRs J0250+5854 and J2251-3711, at a higher detection significance than FFT methods (Morello et al. 2020b; Tan et al. 2018).

Comparisons between FFA and FFT have also been explored in Singh et al. 2022 (hereafter S22). They made use of both simulations and examples of real data and investigated the efficacies of the FFT and FFA implementations in order to compare the quality of detections. They varied the period, duty cycle and morphology of pulses and reported FFA to be able to detect pulsar signals at a consistently higher S/N than FFT.

Similar to FFT, FFA is less effective at detecting sporadic signals. However, a robust analysis comparing the efficacies of FFA's implementations in detecting nulling signals has not been carried out yet. Some limited studies have been published on this front; e.g. the recent work by Singh et al. (2023) who compare FFA and FFT in detecting nulling signals. However, this was done for only one particular source.

Finally, single pulse search (SPS) is quite effective in detecting pulsed, transient-like emission, and is routinely used in transient, pulsar and fast radio burst (FRB) surveys. It searches for emission by convolving the time series with box car filters with various trial widths (e.g. in logarithmic steps of 2^n, where $n = 0, 1, 2, ..., n$). This method is demonstrably effective at finding long-period pulsars, having successfully made the initial detections of both PSRs J2251-3711 and (Tan et al. 2018) J0901-4046 (Caleb et al. 2022). SPS has also been effective in detecting sporadically emitting pulsars, such as Rotating Radio Transients (RRATs), which are, by definition, better detected by SPS (Keane & McLaughlin 2011). Although the efficacy of SPS has been demonstrated in detecting both sporadic and long-period sources, the efficacy of SPS towards sporadic and long-period objects has not been robustly tested and compared with other algorithms.

Ongoing large surveys, such as the SMART survey, intend to incorporate FFA and SPS into their search pipelines, alongside FFT-based methods, in order to increase sensitivity towards long-period objects (Bhat et al. 2023b). Given that recent long-period pulsar discoveries tend to show a significant amount of nulling, it is valuable to test the efficacies of various methods in detecting such long-period nulling pulsars. In order to test these methods, we chose PRESTO's implementations of FFT and SPS, and RIPTIDE's implementation of FFA (see justifications in Section 2.1). We perform a simulation analysis to explore the large search parameter space, covering periods, nulling parameters and S/Ns for a robust analysis. Section 2 describes the process and the justifications for our analysis. For reference, we try to duplicate typical observations from the SMART survey. The results are presented in Section 3. As a sanity check, we also explore the effectiveness of the search methods in detecting pulsars in real data (Section 4). Finally, we discuss our results and make some recommendations to consider in future search processing to increase the detectability of long-period pulsars (Section 5). Our conclusions are summarised in Section 6.

2. Analysis

In order to test how well FFT, FFA and SPS would perform in detecting a variety of long-period and nulling pulsars, we generated a large number of simulated pulsar signals. This allows us to sample a large parameter space, uninfluenced by selection biases. A high-level overview of our simulation and searching procedure is shown in Fig. 2.

S22 performed a similar study where they tested the detection significance of RIPTIDE's FFA and PRESTO's FFT implementations using simulated data. We chose to keep parts of our analysis
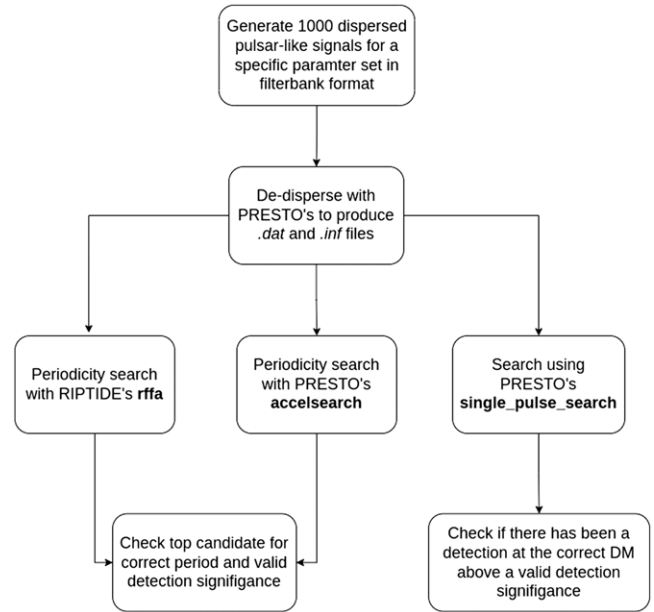

Figure 2. A flow diagram for the simulation of the synthetic pulsar signal and their processing through FFT, FFA, and SPS methods.

similar to S22 to allow for meaningful comparison. S22 varied periodicity, subpulse morphology and duty cycles. As we are interested in nulling properties of long-period pulsars, we chose to vary the pulse period, nulling fraction, nulling duration and signal strength, searching with RIPTIDE's FFA, PRESTO's FFT and PRESTO's SPS implementations.

2.1 Simulating pulsar signals

The simulated signals consisted of trains of Gaussian pulses (mimicking pulsar signals) injected into Gaussian noise (receiver noise). We based our simulations on the SMART survey, as it has the longest dwell time of any current survey (4 800 s), and therefore, in principle, offers better detection prospects (within the sensitivity limits) for detecting long-period signals.

Our main goal is to explore the long-period and nulling parameter spaces. We, therefore, simulate signals with varying periods (P), nulling properties (nulling fraction, n_f; nulling duration, N_d) and average signal-to-noise ratio per pulse (S/N$_{PP}$). Within each simulated signal, we allowed individual pulses to vary in brightness according to a modulation index that was kept constant for all simulations. We also kept the pulse duty cycle constant at 4%.

The target parameters were varied logarithmically with 6–7 increments in order to encapsulate a large range of values. The periods ranged from 0.5–50 s. The choice of a large value for the maximum period was driven by our intent to simulate very long-period pulsars ($P \gtrsim 10$ s).

To emulate nulling, we varied n_f from 1%–90% and N_d from 5 pulses – 1 000 pulses. This accounts for the majority of the nulling pulsar population as well as extremely nulling pulsars such as RRATs. For each null sequence within a simulation, its duration was drawn from a Poissonian distribution with an expected value N_d. Burst sequence durations were drawn from a Poissonian distribution with expected value:

$$B_d = \frac{N_d}{n_f} - N_d. \tag{1}$$

These choices guarantee that the long-term behaviour of the simulated pulse train matches the expected behaviour for a pulsar with the selected average N_d and n_f, whilst still reflecting the randomness of nulling behaviours observed in real pulsars.

Finally, as pulsars are known to vary in brightness, we chose S/N per pulse (S/N$_{PP}$) to range from 0.1–10. A large upper limit was chosen primarily because at low frequencies ($\lesssim$300 MHz), pulsars tend to be generally brighter and some RRATs tend to have bright pulses (Burke-Spolaor & Bailes 2010). Given that the number of pulses within a simulated signal was not predetermined due to the randomness of the nulls, the S/N was chosen per pulse rather than the signal as a whole. Varying S/N$_{PP}$ also allows us to more accurately test the threshold of SPS capabilities than we would have with an integrated S/N. The distribution of S/N$_{PP}$ was assumed to be log-normal, in agreement with what is commonly observed Burke-Spolaor et al. 2012:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(\frac{(\ln x - \mu)^2}{2\sigma^2}\right). \tag{2}$$

Here μ is the mean S/N$_{PP}$ (where S/N$_{PP}$ ranges from 0.1–10) and σ^2 is the variance, which is dependent on the modulation index m of the signal via $\sigma = m\mu$. The modulation index was chosen arbitrarily as 0.1 as the value can range from 0–2 for typical and sub-pulse drifting pulsars (Weltevrede, Edwards, & Stappers 2006).

Due to computational constraints, all signals were generated with a single dispersion measure (DM) of 150 cm^{-3} pc. As we chose not include any scattering effects, the choice of DM was arbitrary as it has little to no effect on the search algorithms. If scattering was implemented to the simulated data, a large DM such 150 cm^{-3} pc would imply considerable scatter broadening, which would lower the peak amplitude of the pulse. The effects of scattering appear differently in the time and frequency domain and hence would affect FFT, FFA and SPS differently. As our primary focus is to assess the effects of nulling on various search methods, we have not considered effects such as pulse broadening (scattering) that alter the pulse shape (and thus impact the detectability). As there is no scattering and the data are dedispersed before being searched, the DM is only relevant as metadata. As long as the algorithm is allowed to search signals of the given DM, it has no effect on the search process. In practice, however, the exact DM of the pulsar may not be searched, and there would be some loss of S/N, as described by Cordes & McLaughlin (2003), in the case of searching sporadic emission. This would directly affect the detection significance of the search methods but since we are using simulated data, we take the 'ideal case' approach.

2.2 Searching methods on simulated data

Once generated, these signals were processed and searched with RIPTIDE's FFA, PRESTO's FFT, and PRESTO's SPS. We chose RIPTIDE as it is the latest, progressive implementation of FFA which has parallelised computation and improved detectability with matched filtering (see Morello et al. 2020a for details). This is also the implementation used in S22.

For the FFT implementation, we chose PRESTO's accelsearch as it is the most commonly used version and is already in use in the first pass (i.e. shallow survey) of the SMART survey (Bhat et al. 2023a). Although accelsearch was designed to do an additional acceleration search in order to target binary pulsars, we do not use this feature as our simulated signals do not include any

Doppler modulation of the period due to binary orbital motion. By using the PRESTO implementation for comparison with other searches in this work, we could directly relate our analysis with S22.

The SPS method was originally described in Cordes & McLaughlin (2003). More recent implementations (e.g. Michilli et al. 2018) primarily improve additional aspects such as the machine learning classification of candidates. We chose PRESTO's single_pulse_search as it is essentially an implementation of the original algorithm described in Cordes & McLaughlin (2003).

It is important to acknowledge that FFA, FFT and SPS are inherently different searching methods that analyse different aspects of the signal and report their results using different metrics. For example, PRESTO's accelsearch searches in Fourier space and reports a σ_{FFT} (the detection significance) and, through additional code, a spectral S/N (S/N$_{FFT}$). RIPTIDE's FFA also reports a detection period but it is accompanied by a S/N$_{FFA}$ measured in the time domain. PRESTO's single_pulse_search on the other hand, is a transient search method, that reports a σ_{SPS} for its transient candidate. The σ and S/N for these methods are thus not equivalent though they all signify detection significance. Hence, a more meaningful comparison of these different methods involves testing whether or not the method can detect a signal based on their relevant criteria.

2.2.1 FFA & FFT

For both FFA and FFT searches, detections would be successful if a candidate was produced with the correct period, within some tolerance, and with a detection significance $\geq 6\sigma$. As FFA does not report a σ value, we used the reported S/N$_{FFA}$. A S/N$_{FFA}$ cut-off of $\sim$8 was recently chosen by Wongphechauxsorn et al. (2024) and the σ_{FFT} cut-off can be $2\sigma_{FFT} - 6\sigma_{FFT}$ (Sanidas et al. 2019; Stovall et al. 2014) for surveys, depending on the instrument. It was also acceptable if the detection was a fractional harmonic of the fundamental period (up to the 32nd harmonic) while still being within the tolerance. The 32nd harmonic limit was chosen as the FFT algorithm was allowed to sum up to 32 harmonics to search for periodic signals. We also lowered the default low-frequency cut-off from one to 0.5 Hz to accommodate long-period signals. The period tolerance was defined as

$$\text{Tol} = \frac{P^2 \times D}{T} \times 1.2 \tag{3}$$

where D is the duty cycle of the signal (which was set to $\sim$ 4%) and T is the total signal time (4 800 s). This was formulated such that the difference in the detected period and original period should not be larger than the full-width half max of the pulse. This ensured that signals folded at the detected period retained a high S/N and resembled the original signal. The additional factor of 1.2 was enforced as the tolerance was too strict and did not reflect a realistic tolerance for surveys, which are more flexible for the detection period.

RIPTIDE was our chosen implementation for FFA. It comes with a search pipeline, rffa, which uses all the tools within the package to search data for periodic signals. This pipeline must be accompanied by a configuration file that specifies the search parameters, such as DM, S/N$_{FFA}$ cut-off, running median width, period range, etc.

2.2.2 SPS

SPS is not a periodicity search; for a given candidate, single_pulse_search (PRESTO) primarily returns a DM and detection significance, σ_{SPS} and S/N$_{SPS}$. The σ_{SPS} is calculated as S/$N_{raw}/\sqrt{bin\_width}$ and is reported in a text file, and S/N$_{SPS}$ is reported in a diagnostic plot; however, the two values are identical, implying that the plot may have intended to display σ_{SPS} but it was incorrectly labelled, or vice versa. Hence in this paper, we will only report σ_{SPS}. Note that the bins mentioned here refer to time bins of the convolved time series generated by single_pulse_search. As mentioned earlier, all our signals were generated with an arbitrary DM of 150 cm^{-3} pc. If SPS produced a candidate above a significance of $6\sigma_{SPS}$, that would be considered a detection. This cutoff is similar to the cut-off of $5\sigma_{SPS}$ used by (Sanidas et al. 2019). For FFA and FFT, the S/N cut-off of 6 was based on other surveys, however, for SPS, the cut-off was chosen based on trials with false candidates. We intentionally de-dispersed the signals at various incorrect DMs and found all false candidates at $\sigma_{SPS} < 6$, motivating the choice of $\sigma_{SPS} = 6$ as our threshold for real detections.

Realistically, pulsar data are trialled at multiple DMs as a real pulsar signal retains its signal strength over multiple DMs, as described by Cordes & McLaughlin (2003). Detections of a transient at multiple DMs help confirm the validity of a candidate. However, the aforementioned approach would only assist the detection quality, if the quality was assessed by a machine learning algorithm or by visual inspection. In this work, we are testing the algorithm's ability to find simulated signals, hence the quality of detection is irrelevant.

3. Results from simulated data

We simulated pulsar signals with varying periods, nulling fractions, nulling durations and S/N per pulse. Each combination of parameters had 1 000 signals generated. These were searched using PRESTO's accelsearch (FFT), RIPTIDE's rffa (FFA) and PRESTO's single_pulse_search (SPS).

The results from our simulations are summarised in Figs. 3 and 4. They are shown as heatmaps with the colour representing the percentage of simulated signals detected out of 1 000, as indicated by the colour bar. Each figure has 7 heatmap panels with a progressively increasing S/N$_{PP}$. Each of the 7 plots shows the period of the signal vs. either the nulling fraction (Fig. 3) or nulling duration (Fig. 4). In either case, the nulling parameter that is not plotted has been averaged in order to make the data presentable. For example, Fig. 3 plots the nulling fraction and rotation period, hence the information relating to the nulling duration has been averaged.

In general, all search methods share some common characteristics. The number of detections increases for all methods as the S/N$_{PP}$ increases. There are also more detections made for shorter periods. This is reconcilable as signals with shorter periods and brighter pulses tend to be present more frequently, even at high values of nulling parameters (i.e. n_f and N_d).

When comparing nulling duration and period, a slight decrease in the number of detections is seen between FFA and FFT. This may be due to our approach of drawing the nulling duration values from a Poisson distribution. At low mean values ($\sim$5), the Poisson distribution has a low chance of outputting 0. Based on Equation (1), if the null duration N_d is 0 pulses, the burst duration B_d would also be 0 pulses. Effectively, this would result in

an additional null sequence, deviating the mean null duration to a larger value. This is a possible reason why the number of detections is lower for $N_d = 5$ pulses than for consecutive higher nulling durations.

3.1 Benchmarking

PRESTO's accelsearch was the fastest process and took typically $\lesssim$0.5 s to search a signal summing up to 8 harmonics; however, when summing up to 32 harmonics searching took up to 2–3 s on some occasions. The accelsearch routine was used with mostly default settings, searching for signals with 0 acceleration, summing up to 32 harmonics and a lower frequency minimum of 0.5 Hz to accommodate for the long-period signals. PRESTO's single_pulse_search was the second fastest, computing for $\lesssim$1 s for low S/N$_{PP}$ ($<$0.5) signals and 1–2 s for high S/N$_{PP}$ (0.5–1) signals. This was also used with the default option, meaning the width of boxcar filters used were $\leq$30 bins; if the whole range of widths were used (1–300 bins), the processing time would likely be longer. Finally, RIPTIDE's rffa was the slowest implementation, which could range 3–20 s depending on the period ranges searched and the definition of other parameter ranges in the configuration file. Note these process times were from a local computer and not a supercomputer, hence computing times may change with devices and different algorithm implementations.[a]

3.2 FFT

Overall accelsearch performed well in detecting long-period and nulling signals, contrary to common perception. The performance was comparable to that of FFA's; however, FFT had noticeably fewer detections at low S/N$_{PP}$ (0.1–0.5) and at high nulling fractions ($n_f = 0.5$–0.9). FFT's better performance is likely due to the very long observation length (4 800 s), which can greatly increase the sensitivity to long-period and nulling signals. We did not consider effects such as RFI and red noise, which can heavily limit the detection of long-period signals in the Fourier space (Parent et al. 2018; Cameron et al. 2017; van Heerden et al. 2017). The duty-cycle was also kept constant at 4%, rather than decreasing with increasing period; a narrow duty-cycle has been shown to be difficult to detect by accelsearch (Singh et al. 2022; Morello et al. 2020a).

Another subtlety we noted is more detections for weaker signals (S/N$_{PP} < 0.3$) at higher periods at $n_f = 0.9$ (top four panels of Fig. 3a). At higher S/N$_{PP}$, this reverses to more detections at lower periods, which is the expected trend.

Initially, we processed the signals summing up to 8 harmonics (by default) for all period ranges; however, this introduced some inaccuracies and artefacts in the simulated results. As the candidate period is dependent on the highest harmonic summed, the candidate periods were not precise (e.g. the number of decimal places), which led to a decrement in detections (for $P = 1.5$ s and 3 s), and some periods could not be detected ($P = 10$ and 50 s) at low S/N$_{PP}$. When reprocessing the FFT search, summing up to 32 harmonics, we found a significant increase in the number of long-period detections and an increase in the precision of the detection period.

---

[a]Local machine has an AMD Ryzen 6 core/12 thread CPU and 2x32GB RAM.
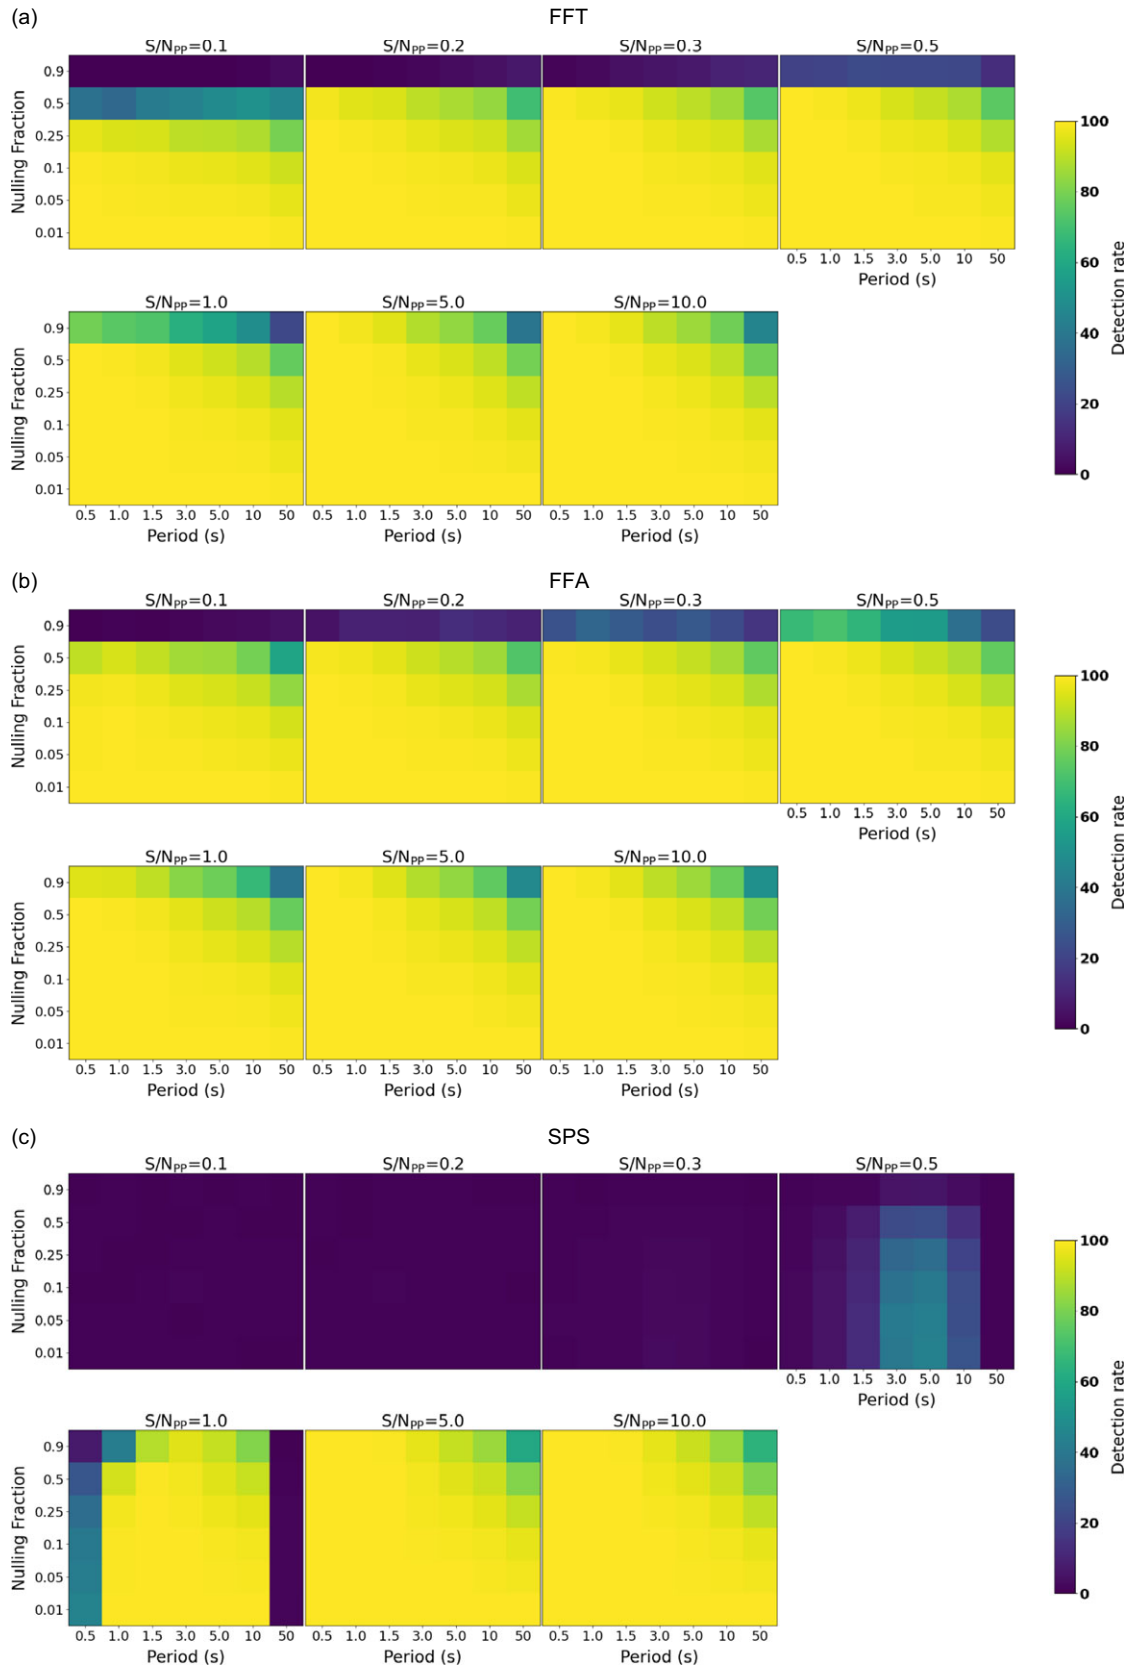
Figure 3. This figure shows the result of searching the simulated data with FFT (a), FFA (B) and SPS (c). The *x*-axis shows the period (s) and the *y*-axis shows the nulling fraction. The 7 plots inside the subfigures show heatmaps with the colour of each pixel indicating the percentage of signals, with the indicated period and nulling fraction, detected with the corresponding search method. This is accompanied by the colour bar on the right-hand side. The 7 plots represent different S/N$_{PP}$ and are labelled as such.
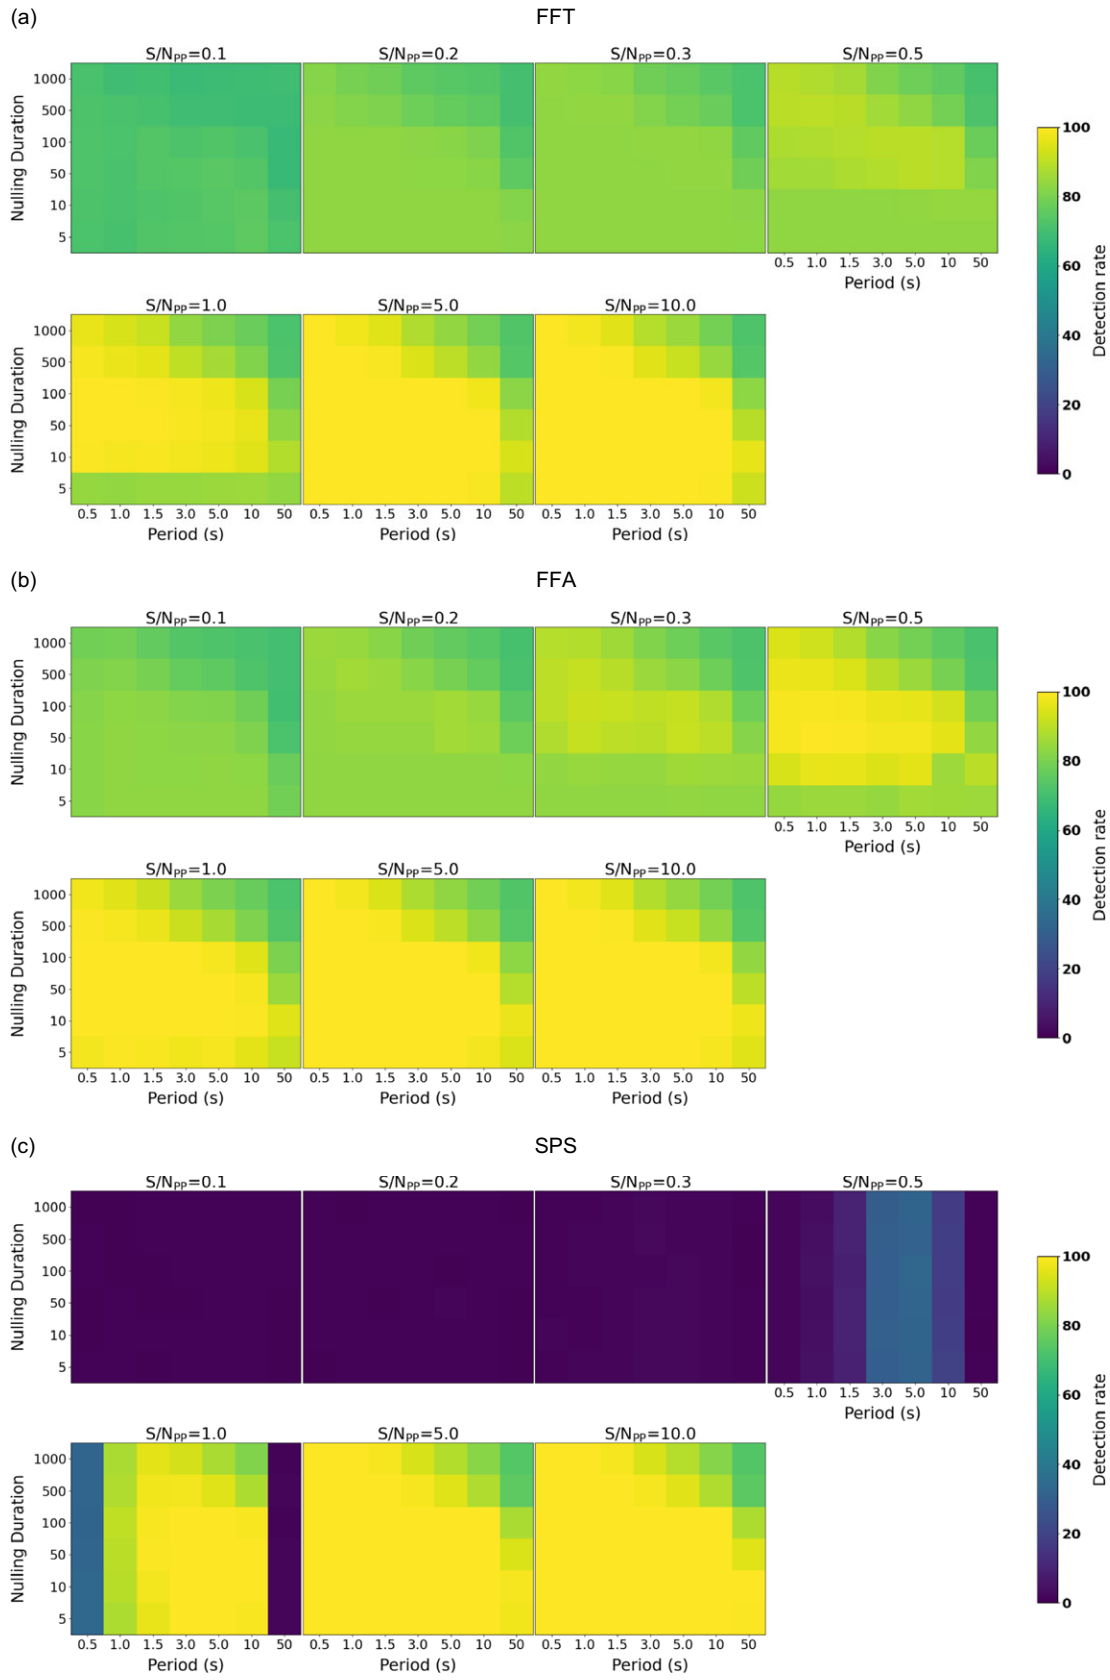
Figure 4. This figure is very similar to Figure 6. It shows the result of searching the simulated data with FFT (a), FFA (B) and SPS (c). The x-axis shows the period (s) and the y-axis shows the nulling duration. The 7 plots inside the subfigures show heatmaps with the colour of each pixel indicating the percentage of signals, with the indicated period and nulling duration, detected with the corresponding search method. This is accompanied by the colour bar on the right-hand side. The 7 plots represent different S/N_{PP} and are labelled as such.

Table 1. The details of the observations and pulsars therein were used for the sanity check analysis. For more on these pulsars (excluding J0452-3418) see Bhat et al. (2023b). Analysis on J0452-3418 (Grover et al. 2024).

Pulsar	Period (s)	DM cm^{-3} pc	S/N$_{\text{FFA}}$	S/N$_{\text{FFT}}$	σ_{FFT}	σ_{SPS}	Observation length (s)	Observation MJD
J0036-1033	0.90	23.1(2)	15.0	22.2	9.7	5.3	600	58774
J0036-1033	0.90	23.1(2)	17.5	23.7	10.6	6.0	2 000	58792
J0026-1955	1.30	23.81(1)	14.5	21.5	9.4	9.2	4 800	58792
J0026-1955	1.30	23.81(1)	62.0	54.6	26.4	14.0	4 800	58799
J1002-2036	1.67	42.28	14.6	18.9	8.0	5.7	600	58890
J1357-2530	0.91	16.04	11.5	19.0	8.2	5.3	900	59309
J0452-3418	1.66	19.77(2)	100.9	118.2	58.1	10.1	4 200	58757

3.3 FFA

RIPTIDE's rffa was very sensitive to all periods, signal strengths and nulling properties. It had the same if not more detections than accelsearch for each parameter set and was more consistent than single_pulse_search in detecting signals at multiple periods and of multiple strengths. Aside from the expected decrease in the number of detections for very high nulling properties at high periods, FFA was efficacious for a large range of parameters in comparison to FFT and SPS.

3.4 SPS

PRESTO's single_pulse_search was primarily only effective at high S/N$_{\text{PP}}$, as one would expect. Periodicity searches fold individual pulses to gain an integrated S/N, however, single_pulse_search does not benefit from that. At best, the algorithm can downsample the time series in order to increase the combined σ_{SPS}. This is why single_pulse_search has very few detections until S/N$_{\text{PP}} = 0.5$.

The algorithm shows a clear preference for signals with periods between 1.5–10 s as shown in the top right plots of Figs. 3c and 4c. The single_pulse_search routine evaluates σ_{SPS} of the pulse based on the matched filters of different widths and their convolution with the pulse. All pulses have the same duty cycle, meaning the pulse width scales with the period. Hence a matched filter would favour the larger period, which would have wider pulses. This is likely the reason why there are fewer detections for smaller periods at small S/N$_{\text{PP}}$ for single_pulse_search.

However, there are a discrete number of matched filters with different widths, and as mentioned previously, the maximum widths of the filters by default range up to 30 bins. As the increment step (dt) of the time series is 3.2×10^{-3} s, the width of the pulses for the periods under consideration go as 6.25, 12.5, 18.75, 37.5, 62.5, 125 and 625 bins. For $P = 5$ s where the pulse is $\sim$60 bins wide, which is double the maximum width allowed for the box car filter, the filtering returns large σ_{SPS} which can allow for better detections; however, for larger periods, the pulse is too large and the maximum filter width of 30 is not able to encapsulate a significant enough portion of the pulse. This limits the σ_{SPS} and hence the detectability. Hence the feature we see is due to the choices of the filters and the shape of the pulses.

At large S/N$_{\text{PP}}$, single_pulse_search performs effectively for both nulling parameters, and at times even performs better than rffa. Although it may be less meaningful to compare SPS with FFA or FFT as their methods have different criteria for detections, it is important to compare the efficacy of these methods

to have the best likelihood of finding long-period and nulling pulsars.

4. Results from real data

4.1 Observational data

We have compared the performance of various search methods on real survey data; we made use of example data from the ongoing SMART survey and measured the detection significance to assess the search methods. Specifically, we made use of data from the discovery and follow-up observations of new pulsars from the SMART survey. In all cases, the original detections were made using PRESTO's FFT implementation; the main purpose is to assess whether the inclusion of other search methods would have led to quicker or more significant detections of these pulsars. Observations of real pulsars used in this study are listed in Table 1.

The observation lengths varied from 600 to 4 800 s, depending on the data that were readily available for analysis. Although all five pulsars have rotation periods of $\sim$1-2 s, two pulsars (J0026-1955 & J0452-3418) also show significant amounts of nulling ($\sim$50% and 33%, respectively for PSRs J0026-1955 & J0452-3418; cf. McSweeney et al. 2022; Grover et al. 2024, making them good targets for testing the efficacy of the methods explored in this work.

4.2 Quantitative performance

S22 compared the efficacies of FFA and FFT using S/N as their detection significance. They report RIPTIDE's S/N$_{\text{FFA}}$ to be much larger than PRESTO's S/N$_{\text{FFT}}$ for both simulated and real data. PRESTO's accelsearch reports a detection probability in the form of a different parameter, σ_{FFT}, which may mistakenly be compared with rffa's S/N$_{\text{FFA}}$. Here, we have presented both values for a clear comparison (see Table 1).

We find S/N$_{\text{FFA}}$ to be consistently larger than σ_{FFT}, which is not meaningful since σ_{FFT} is essentially a probability statistic. However, we also find that S/N$_{\text{FFT}}$ is larger in most cases than the S/N$_{\text{FFA}}$, which is opposite to the findings of S22.

SPS, on the other hand, provides a σ_{SPS}, which is a multiple of the S/N$_{\text{raw}}$. This value is consistently lower than σ_{FFT} value; however, the σ_{FFT} is based on many other parameters as well as the integrated profile, whereas the σ_{SPS} is based on individual pulses and the widths of the filters. Interestingly, we note that SPS returns a very high σ_{SPS} value for both nulling pulsars, with $\sigma_{\text{SPS}} \sim 9$

and 10 for J0026-1955 and J0452-3418, respectively (as shown in Table 1).

5. Discussion

On the basis of the analysis, we now comment on the efficacies of various search methods in detecting long-period and/or nulling pulsars. This is especially relevant for searches using wide-field instruments that also employ longer dwell times (e.g. the SMART and LOTAAS surveys).

5.1 Nulling duration vs Nulling fraction

As described earlier (Section 2), nulling was characterised using two basic parameters: nulling fraction and nulling duration. Our results are shown in Figs. 6 and 10, where the detections (for signals with varying n_f and N_d) are a function of the rotation period. At first glance, the number of detections for both parameters is similar for all search methods, particularly as they approach larger values of n_f and N_d at large S/N_{PP}. It is difficult to compare n_f and N_d directly as the definition of 'large' is different for the two parameters. n_f has a clearly defined minimum and maximum value, whereas the range of N_d values are based entirely on observations.

Based on our analysis, the detections tend to vary more with an increase in n_f than that in N_d, for the periodicity searches. At low S/N_{PP}, the number of detections is roughly constant for all values of N_d. For larger S/N_{PP}, the detections tend to plateau near two to three values, where the number of detections remains consistent for some consecutive values of N_d. This is only visible for FFT at higher S/N_{PP}, whereas for FFA, this is also observed at low S/N_{PP}. The plateauing is likely due to coarse increment sizes for N_d.

The increments in N_d could be increased to better sample the range covered, allowing for a clear gradual change in the number of detections. In any case, the range we cover likely encapsulates the majority of nulling pulsars, as only a small percentage are known to null for longer durations, and as such, for larger values, we may expect to see only a small number of pulses from them. This is discussed further in Section 5.3.

5.2 Comparison of the search methods on simulated data

To analyse the efficacy of FFA, FFT and SPS implementations in detecting long-period and nulling pulsars, we used analysis of simulated data. As there were 1.764×10^6 individual time series streams to process, the computation time for the three search methods was essential to understand. This is a key factor that will influence the consideration of these methods for future surveys. Of all the three methods, the FFT implementation is the fastest and performs up to an order of magnitude faster compared to the other two. FFA implementation is the slowest. Although the results in Figs. 3b and 4b show FFA to be more sensitive than FFT, FFT is computationally efficient. For long-duration surveys (e.g. the SMART survey, which can produce tens of millions of candidates), the processing efficiency may outweigh the benefits of extra sensitivity.

The results presented in Section 3 suggest that FFA and FFT are surprisingly effective at finding pulsars with high nulling properties at low to moderate rotation periods. At larger periods and high nulling properties, both methods lose effectiveness as the signals become less periodic and less frequent.

To further elaborate, rffa accurately and consistently detects pulsars throughout all S/N_{PP}. PRESTO's **accelsearch** is also very effective, however, it has fewer detections at low signal strengths (S/N_{PP}=0.1–0.5) and high nulling fractions ($n_f = 0.5$–0.9). For longer periods ($P \geq 10$ s), it heavily relies on the harmonics sums to be ≥ 16 for the signal to be detected accurately at the fundamental frequency.

Our analysis also suggests that SPS is most effective in detecting long-period signals at high nulling fractions and durations, albeit this applies primarily at high S/N_{PP}. Although SPS may miss a large majority of generic and low-luminosity pulsars that FFT and FFA may detect, it is effective in detecting bright pulses, especially when emission is sporadic. This is vividly demonstrated by the way the extremely long-period pulsar, J0901-4046, was discovered, where the initial detection was made by searching for individual dispersed pulses (Caleb et al. 2022).

5.3 Critique of the simulation

Our choice of a detection significance above 6 was based on generic cut-offs for candidate selection typically employed in search pipelines. For SPS, a threshold of 6σ was based on experiments with incorrectly de-dispersed data as described in Section 2.2.2. If the threshold was to be increased from 6 to 7 for example, it would primarily affect SPS. As FFA and FFT integrate the signal and base their significance values on the sum, their values are generally higher than 6, especially for simulated data in ideal noise. SPS's detection significance is based on the strength of individual pulses. Hence increasing the threshold from 6 to 7 would likely further decrease the detections made by SPS at low S/N_{PP}. However, by definition, this cut-off would also lead to more probable detections.

The periods analysed in this study were spaced logarithmically to cover a large range of values. However, this led to large jumps in periods, specifically between 10–50 s, which reflect as drastic changes in the detectability of the search methods in Figs. 3 and 4. Conducting a simulation for these intermediate periods would reveal more precisely how sensitive the methods are to detecting pulsars like J2251-3711 (Tan et al. 2018) and J0901-4046 (Caleb et al. 2022).

The limitation in terms of time and computational resources prompted us to consider simpler characteristics of pulsar signals; for example, we considered pulse shapes that are composed only of repeating Gaussian-shaped pulses with an additive Gaussian noise component to mimic receiver noise. This was sufficient for our purpose. Furthermore, as mentioned earlier in Section 2, the data were de-dispersed at a single DM of 150 cm^{-3} pc. Exploring the impact of incorrect dedispersion (i.e. true DM different from an assumed DM) will involve more computation and resources, but can help to assess loss in detection significance resulting from an error in DM.

Our analysis did not consider the effect of *red noise* that may arise from variations in receiver noise or gain or varying levels of RFI. As such, such effects tend to be specific to the characteristics of the instrument and therefore a generalisation may not be meaningful. In any case, we note that the inclusion of red noise would have affected the periodicity searches, making detections more difficult (van Heerden et al. 2017). We also note that S22 performed some useful comparison in this regard, between RIPTIDE and PRESTO, highlighting a significant reduction in the detection S/N in PRESTO searches than in RIPTIDE searches.

Furthermore, in our simulations, we allowed nulling to occur randomly (see Section 2 for details). The choice of our considerations also meant that for some data sets, no pulses were present, due to either a high nulling fraction or a large nulling duration. In some ways, this reflects the realistic expectations in searching for objects with high degrees of intermittency. As a result, detections for $S/N_{PP} = 5$ and 10 look very similar for all three searches, and highlights the difficulty in finding such signals, regardless of their brightness.

The modulation was fixed to a value of 0.1, and we did not test how changing it might affect the different search methods. We can anticipate, however, that it would not have a strong effect on periodicity searches, which primarily rely on the integrated signal power and only to a lesser extent on the 'equivalent' red noise introduced by the modulation. However, a change in modulation would affect the SPS; e.g. a larger modulation index will allow certain pulses in the signal to be more readily detectable. That is, intrinsically low S/N_{PP} signals would still have a higher probability of detection by SPS, especially when certain pulses are temporally boosted in brightness as a result of modulation. This can be relevant for objects like giant pulse emitters and RRATs that tend to modulate heavily and emit bright pulses.

The duty cycle of the pulse was also kept constant in our simulations. We note that previous studies (e.g. Morello et al. 2020a, S22) have demonstrated the loss of sensitivity in FFT searches for narrow duty cycles. We therefore chose to keep this constant. Long-period pulsars tend to have very small duty cycles, $\lesssim 1\%$; however, a recently discovered pulsar-like source with a period of ~ 20 min showed a duty cycle of up to 25%. For such small duty cycles, FFT's detection ability will be somewhat hindered, specifically for lower harmonics (Morello et al. 2020a).

5.4 Real data

In order to verify the performance of various search methods, we tested rffa, accelsearch and single_pulse_search on examples of real data, chosen from SMART data sets that are currently being processed for a first-pass shallow survey (cf. Bhat et al. 2023b).

As summarised in Table 1, both FFT and FFA comfortably found all five pulsars in all observations that we used. Based on the criteria employed for the simulated analysis, all sources would be classified as successful detections. We note that SPS finds only two pulsars above a significance of 6; both nulling and bright enough to conduct single pulse analysis. This demonstrates that SPS is primarily sensitive to detecting bright pulsars.

In general, we find S/N_{FFT} to be higher than S/N_{FFA}. The difference in S/N may arise due to the differences in the measurement of this value. For instance, FFA integrates the pulses in the time domain and measures the S/N from the combined profile, therefore the S/N_{FFA} is a function of the trial period and the trial width of the matched filters and is based on a defined Z-statistic(Morello et al. 2020a). FFT instead sums the square root of the amplitude of the harmonics in the Fourier space to calculate S/N_{FFT} (Ransom 2001). Theoretically, for an S/N measured in the Fourier domain to equal one in the time domain, the power in all harmonic bins (up to infinity) should be summed. This implies either RIPTIDE is underestimating the S/N or PRESTO is overestimating the S/N.

However, S22 report that S/N_{FFA} is up to orders of magnitude larger than S/N_{FFT} for a majority of their sources. A similar difference is seen in our case only when we compare S/N_{FFA} to

Table 2. A summary of the recommendations for the three search algorithms based on this work.

Search method	Software	Period range	σ or S/N threshold
FFT	PRESTO	$\lesssim 10$ s	~ 6
FFA	RIPTIDE	$\gtrsim 1$ s	~ 6
SPS	PRESTO	–	> 6

σ_{FFT}, the latter being the commonly adopted metric in accelsearch candidates.

5.5 Boolean vs quantitative comparison

As mentioned earlier, comparing the different search methods by merely using commonly adopted detection significance may be less meaningful than using a standard applicable to all methods. Unravelling how the different values are connected is a complex task as there are many parameters and calculations involved. This can lead to unfair comparisons between search methods as many generally assume a larger detection significance means a stronger detection; however, if the detection significance parameters are defined differently by two methods, then the larger value may not always mean a better detection.

Further, this can also lead to confusion about which numbers to compare when conducting a thorough analysis. SPS lists σ_{SPS} as a detection significance of individual pulses; however, FFT or FFA perform periodicity searches where they consider the sum of the pulses. Hence it is not meaningful to directly compare their S/N's and σ values. Similarly, the σ_{FFT} should not be compared to S/N_{FFA} as they also measure different properties. This is made further difficult as the PRESTO implementation of FFT does not initially report a S/N_{FFT}; it is obtained once candidates are pruned with quick_prune_cands or ACCEL_sift.

Such inherent differences prompted us to consider a Boolean approach in our simulations where we ascertain the efficacies of the search methods based on their ability to make a detection. Although there may still be some bias due to the definition of the detection criteria, it is a more meaningful comparison for survey needs.

5.6 Recommendations for future surveys

Based on the performance of FFT, FFA, and SPS on our simulated data, we outline some considerations to guide their effective use in long-duration surveys such as the SMART survey (see Table 2).

Being a computationally efficient and inexpensive algorithm, FFT has been very effective for pulsar searching. Our analysis suggests that it is still the most effective method for searching pulsars with $P \lesssim 10$ s. Summing up to 32 harmonics, we found many long-period and nulling candidates, however, there was no red noise injected and the duty cycle of signals was kept constant, which can affect the effectiveness of FFT as noted by Singh et al. (2022), Morello et al. (2020a). FFT also detected fewer signals than FFA at low S/N_{PP} (0.1–0.5) and at high nulling fractions (n_f =0.5–0.9). Overall, the maximum period range that could be searched with FFT should be $P \lesssim 10$ s. This accounts for the vast majority of the pulsar population. This upper limit also means that the number of harmonics summed can be smaller (16 instead of 32), hence the processing can be faster.

FFA, in general, shows sensitivity throughout the parameter space; however, it is the slowest of the three search methods (by

1-2 order of magnitudes on local machine). As FFT had fewer detections when summing weaker and highly nulling signals than FFA has been shown to be more sensitive to faint pulsars, which was also shown by Morello et al. (2020a), FFA should have some overlap with the FFT searching periods e.g. $P \gtrsim 1$ s, to ensure a more thorough search. This is also computationally a less expensive option as rffa can be more than 6 times faster for $P \gtrsim 1$ s (Morello et al. 2020a).

SPS is also a computationally inexpensive algorithm that is independent of the exact period and can be quite effective in detecting sporadically emitting sources such as RRATs and FRBs. Most ongoing pulsar surveys routinely use this as part of their processing pipelines (e.g. Sanidas et al. 2019; Bhattacharyya et al. 2016), and as such it is among the future processing plans envisioned for SMART Bhat et al. (2023a). However, we note it is primarily effective at detecting bright single pulses. This can be useful in detecting giant pulse emitters or bright RRATs. Considering the sporadicity aspect, SPS will offer increased sensitivity for detecting new objects, for longer observations. It appears that a higher detection threshold ($>6\ \sigma_{SPS}$) can be a more efficient way of finding bright sporadically emitting objects.

Incorporation of the above considerations will help long-duration survey processing attain an increased detection sensitivity to a wider spectrum of pulsars and search a large part of the parameter space. This can be beneficial for future next-generation surveys with MeerKAT or SKA-Low, which have the potential to uncover the comprehensive population of radio-emitting neutron stars in our Galaxy. These can also be applicable to non-traditional search strategies such as those considered by CHIME where they take daily observations of the same position in the sky and sum their power spectra (similar strategies to Cadelano et al. 2018). This allows for an essentially longer time series which can be useful in detecting long-period and nulling objects. Consideration of such extended search strategies will also enable ongoing and future searches to explore possible populations of pulsars that have been missed earlier due to selection effects.

6. Conclusion

To develop a more comprehensive picture of pulsar emission properties and various aspects relating to their populations, it is important to extend current search approaches to include algorithms and techniques that will allow a more extensive exploration of the long-period parts of the parameter space. This can potentially lead to an improved understanding of pulsar death lines and the related physical modes, alongside yielding new insights into linking nulling to long periods, as hinted by some of the recent discoveries. With the advent of wide-field radio-astronomy facilities like the MWA and LOFAR, and their adoption for large (untargeted) pulsar surveys, longer dwell times per pointing become more affordable, and this can be well exploited for increasing the detection sensitivity to such long-period and nulling pulsars. Popular search algorithms and software tend to vary both in their efficacy and sensitivity to the related classes of objects. Motivated by this, we explored the efficacies of different search methods in detecting long-period and nulling pulsars, using both simulations and examples of real data. The software used includes PRESTO's FFT implementation (accelsearch), RIPTIDE's FFA implementation (rffa) and PRESTO's SPS implementation (single_pulse_search). The simulation analysis was approached from the perspective of exploring a wider range in parameter space.

Our analysis finds FFA to be most sensitive to pulsar signals over a much wider range in period ($\sim$ tens of seconds) even when the signal strengths are low; however, it is also slower than FFT. While FFT is computationally more efficient and generally sensitive towards long periods, it loses sensitivity for very weak and highly nulling signals. SPS on the other hand is primarily sensitive to sporadic or intermittent emission, but largely sensitive to high signal strengths. Therefore, a more effective (if not ideal) search strategy could include the use of FFT searches for $P \lesssim 10$ s, FFA searches for $P \gtrsim 1$ s, alongside SPS with a detection threshold of $> 6\sigma$, and preferably with long observations (>10 min), for increased sensitivity to intermittent or sporadic emitters. While such a strategy was derived from simulations and examples of real data selected from SMART survey data sets, they are also applicable for many of the current and planned large pulsar surveys, including those envisioned with SKA-Low. In general, more extensive explorations of the long-period parts of the search parameter space can help overcome some of the inherent observational biases in detecting long-period or nulling pulsars, and thus uncover a larger population (and variety) of objects that will help us to develop a better understanding of pulsar population and emission physics.

References

Anumarlapudi, A., Swiggum, J. K., Kaplan, D. L., & Fichtenbauer, T. D. J. 2023, ApJ, 948, 32

Arons, J., & Scharlemann, E. T. 1979, ApJ, 231, 854

Bhat, N. D. R., et al. 2023a, PASA, 1

Bhat, N. D. R., et al. 2023b, Publications of the Astronomical Society of Australia, 1

Bhattacharyya, B., et al. 2016, ApJ, 817, 130

Burke-Spolaor, S., et al. 2012, MNRAS, 423, 1351

Burke-Spolaor, S., & Bailes, M. 2010, MNRAS, 402, 855

Cadelano, M., et al. 2018, ApJ, 855, 125

Caleb, M., et al. 2022, NatAs, doi: 10.1038/s41550-022-01688-x

Cameron, A. D., et al. 2017, MNRAS, 468, 1994

Chen, K., & Ruderman, M. 1993, ApJ, 402, 264
Cordes, J. M., & McLaughlin, M. A. 2003, ApJ, 596, 1142
Grover, G., et al. 2024, arXiv e-prints, arXiv:2405.16725
Han, J. L., et al. 2021, RAA, 21, 107
Hurley-Walker, N., et al. 2022, Natur, 601, 526
Hurley-Walker, N., Rea, N., & McSweeney, S. 2023, Natur, 619, doi: 10.1038/s41586-023-06202-5
Keane, E. F., & McLaughlin, M. A. 2011, BASI, 39, 333
Keith, M. J., et al. 2010, MNRAS, 409, 619
Kondratiev, V. I., et al. 2009, ApJ, 702, 692
McSweeney, S. J., et al. 2022, ApJ, 933, 210
Michilli, D., et al. 2018, MNRAS, 480, 3457
Morello, V., et al. 2020b, MNRAS, 493, 1165
Morello, V., Barr, E. D., Stappers, B. W., Keane, E. F., & Lyne, A. G. 2020a, MNRAS, 497, 4654
Parent, E., et al. 2018, ApJ, 861, 44
Ransom, S. M. 2001, PhD thesis, Harvard University, Massachusetts
Rea, N., et al. 2023, arXiv e-prints, arXiv:2307.10351
Ritchings, R. T. 1976, MNRAS, 176, 249
Ruderman, M. A., & Sutherland, P. G. 1975, ApJ, 196, 51
Sanidas, S., et al. 2019, A&A, 626, A104
Sheikh, S. Z., & Macdonald, M. G. 2021, MNRAS, 502, 4669
Singh, S., et al. 2022, ApJ, 934, 138
Singh, S., et al. 2023, ApJ, 944, 54
Singh, S., Roy, J., Sunder, S., Bhattacharyya, B., & Kudale, S. 2023, ApJ, 947, 88
Staelin, D. 1969, Proc. IEEE, 57, 724
Stovall, K., et al. 2014, ApJ, 791, 67
Su, W. Q., et al. 2023, arXiv e-prints, arXiv:2305.16754
Tan, C. M., et al. 2018, ApJ, 866, 54
Tingay, S. J., et al. 2013, PASA, 30, e007
van Haarlem, M. P., et al. 2013, A&A, 556, A2
van Heerden, E., Karastergiou, A., & Roberts, S. J. 2017, MNRAS, 467, 1661
Weltevrede, P., Edwards, R. T., & Stappers, B. W. 2006, A&A, 445, 243
Wongphechauxsorn, J., et al. 2024, MNRAS, 527, 3208
Zhang, B., Gil, J., & Dyks, J. 2007, MNRAS, 374, 1103
Zhang, B., Harding, A. K., & Muslimov, A. G. 2000, ApJ, 531, L135