



The Implementation of a Fast-folding Pipeline for Long-period Pulsar Searching in the PALFA Survey

E. Parent¹, V. M. Kaspi¹, S. M. Ransom², M. Krasteva³, C. Patel¹, P. Scholz⁴, A. Brazier^{5,6}, M. A. McLaughlin^{7,8}, M. Boyce¹, W. W. Zhu^{9,10}, Z. Pleunis¹, B. Allen^{11,12,13}, S. Bogdanov¹⁴, K. Caballero¹⁵, F. Camilo¹⁶, R. Camuccio¹⁵, S. Chatterjee¹⁶, J. M. Cordes⁵, F. Crawford¹⁷, J. S. Deneva¹⁸, R. Ferdman¹⁹, P. C. C. Freire¹⁰, J. W. T. Hessels^{20,21}, F. A. Jenet²², B. Knispel^{11,12}, P. Lazarus¹⁰, J. van Leeuwen^{20,21}, A. G. Lyne²³, R. Lynch², A. Seymour¹⁰, X. Siemens¹³, I. H. Stairs²⁴, K. Stovall²⁵, and J. Swiggum¹³

¹ Department of Physics and McGill Space Institute, McGill University, Montreal, QC H3A 2T8, Canada; parente@physics.mcgill.ca

² National Radio Astronomy Observatory, Charlottesville, VA 22903, USA

³ Department of Physics, Concordia University, Montreal, QC H4B 1R6, Canada

⁴ National Research Council of Canada, Herzberg Astronomy and Astrophysics, Dominion Radio Astrophysical Observatory, P.O. Box 248, Penticton, BC V2A 6J9, Canada

⁵ Department of Astronomy, Cornell University, Ithaca, NY 14853, USA

⁶ Center for Advanced Computing, Cornell University, Ithaca, NY 14853, USA

⁷ Department of Physics and Astronomy, West Virginia University, Morgantown, WV 26506, USA

⁸ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

⁹ National Astronomical Observatories, Chinese Academy of Science, 20A Datun Road, Chaoyang District, Beijing 100012, People's Republic of China

¹⁰ Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

¹¹ Leibniz Universität at Hannover, D-30167 Hannover, Germany

¹² Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany

¹³ Physics Department, University of Wisconsin–Milwaukee, 3135 N. Maryland Avenue, Milwaukee, WI 53211, USA

¹⁴ Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

¹⁵ Center for Advanced Radio Astronomy, University of Texas Rio Grande Valley, Brownsville, TX 78520, USA

¹⁶ SKA South Africa, Pinelands, 7405, South Africa

¹⁷ Department of Physics and Astronomy, Franklin and Marshall College, Lancaster, PA 17604-3003, USA

¹⁸ George Mason University, resident at the Naval Research Laboratory, Washington, DC 20375, USA

¹⁹ Faculty of Science, University of East Anglia, Norwich Research Park, Norwich NR4 7TJ, UK

²⁰ ASTRON, Netherlands Institute for Radio Astronomy, Postbus 2, 7990 AA, Dwingeloo, The Netherlands

²¹ Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

²² Center for Gravitational Wave Astronomy, University Texas Rio Grande Valley—Brownsville, TX 78520, M13 9PL, UK

²³ Jodrell Bank Centre for Astrophysics, School of Physics and Astrophysics, University of Manchester, Manchester, M13 9PL, UK

²⁴ Department of Physics and Astronomy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada

²⁵ NRAO, PO Box 0, Socorro, NM 87801, USA

Received 2017 September 4; revised 2018 May 14; accepted 2018 May 15; published 2018 June 29

Abstract

The Pulsar Arecibo L-Band Feed Array (PALFA) survey, the most sensitive blind search for radio pulsars yet conducted, is ongoing at the Arecibo Observatory in Puerto Rico. The vast majority of the 180 pulsars discovered by PALFA have spin periods shorter than 2 s. Pulsar surveys may miss long-period radio pulsars owing to the summing of a finite number of harmonic components in conventional Fourier analyses (typically ~ 16), or as a result of the strong effect of red noise at low modulation frequencies. We address this reduction in sensitivity by using a time-domain search technique: the fast-folding algorithm (FFA). We designed a program that implements an FFA-based search in the PALFA processing pipeline and tested the efficiency of the algorithm by performing tests under both ideal, white-noise conditions, as well as with real PALFA observational data. In the two scenarios, we show that the time-domain algorithm has the ability to outperform the FFT-based periodicity search implemented in the survey. We perform simulations to compare the previously reported PALFA sensitivity with that obtained using our new FFA implementation. These simulations show that for a pulsar having a pulse duty cycle of roughly 3%, the performance of our FFA pipeline exceeds that of our FFT pipeline for pulses with dispersion measure $\lesssim 40 \text{ pc cm}^{-3}$ and for periods as short as $\sim 500 \text{ ms}$, and that the survey sensitivity is improved by at least a factor of two for periods $\gtrsim 6 \text{ s}$. Early results from the implementation of the algorithm in PALFA, including discoveries, are also presented in this paper.

Key words: methods: data analysis – pulsars: general

1. Introduction

One characteristic of the population of known radio pulsars is that 93% of them have spin periods (P) shorter than 2 s²⁶ (<http://www.atnf.csiro.au/people/pulsar/psrcat/>). The notable lack of long-period pulsars could be an intrinsic property of the population. For instance, the observed population of slowly

rotating pulsars (defined here as having $P > 2 \text{ s}$) have radio beam widths smaller than typical pulsars. Indeed, the median pulse duty cycle, δ , defined as the ratio of the FWHM of the pulse to the pulsar period, for this class of pulsars is 1.6%, while it is 3.1% for pulsars with spin periods shorter than 2 s. The beaming of the radiation would therefore play a role in the detectability of slow pulsars. The lower spin-down luminosity of long-period pulsars is another factor that could explain why these pulsars are particularly difficult to detect.

²⁶ Based on the ATNF Pulsar Database, version 1.56.

In addition to effects that are intrinsic to the pulsars themselves, the lack of long-period pulsars in the known population may also be due to selection bias in pulsar surveys. One of the reasons why surveys are likely to miss slowly rotating pulsars is that pulsar search radio data are often badly affected by red noise, or excess noise at low modulation frequencies. This non-Gaussian noise is the result of the combined effects of various factors such as receiver gain fluctuations and radio frequency interference (RFI). The broad features introduced in the time series by red noise increase the number of false positives in the low modulation frequency regime (defined in this paper as $f < 0.5$ Hz), where red noise is strongest, causing a considerable reduction in the sensitivity of pulsar surveys at this end of the spectrum. For the Pulsar Arecibo L-Band Feed Array (PALFA) survey, the fact that the integration time of the observations is only 268 s is another limiting factor of the detectability of the survey to long-period pulsars.

While Fourier-based search techniques have been commonly used in blind searches for pulsars, their performance is highly compromised by red noise. By recovering synthetic pulsar signals injected in real observational data with PRESTO's (Ransom 2001) fast Fourier transform (FFT) search program (*accelsearch*), Lazarus et al. (2015) demonstrated that there are major discrepancies between the true sensitivity of the PALFA survey and the sensitivity predicted by the radiometer equation (Dewey et al. 1985). For a hypothetical pulsar having a spin period of 10 s and a dispersion measure (DM) of 10 pc cm^{-3} , the minimum mean flux density the FFT can detect is 20 times larger than the value predicted by the radiometer equation. The degradation in the true sensitivity is noticeable at pulsar spin periods as short as few hundreds of milliseconds.

One way to partially address this reduction in sensitivity is by the use of a fast-folding algorithm (FFA; Staelin 1969), a time-domain search technique especially well suited for finding long-period signals. The FFA folds a dedispersed time series at multiple trial periods and avoids redundant summation of bins by storing in memory the resulting sum of each folding step and later reusing these stored quantities when needed. The main advantage of the FFA over a frequency-domain search is that by producing a phase-coherent result, it retains all harmonic structure, as opposed to FFT-based searches, where only a limited number of harmonics²⁷ are incoherently summed (i.e., without using the phase information in the harmonics). Having a search technique that is efficient at finding narrow-pulsed signals in the long-period regime is therefore desirable.

Recovering the loss in sensitivity reported in Lazarus et al. (2015) is important, as it has the potential for scientific advancements in pulsar astronomy. Our understanding of the Galactic pulsar population is heavily biased by various selection effects. These include the propagation effects in the interstellar medium, the nonuniform radio sky background, the distances and proper motions of pulsars, and the sizes of the emission beams. In addition to the observational effects mentioned above, red noise also likely affects the observed period distribution of the pulsar population. Finding more slowly rotating pulsars would help us constrain the radio emission mechanism: one of the longest-period radio pulsars known (Young et al. 1999), PSR J2144–3933 ($P = 8.5$ s), challenges existing models, as this object is located beyond the

theoretical death line (Chen & Ruderman 1993; Zhang et al. 2000; Hibschan & Arons 2001) in the $P-\dot{P}$ diagram. The very recent discovery of a 23.5 s pulsar in the LOFAR Tied-Array All-Sky Survey,²⁸ PSR J0249+58 (C. M. Tan et al. 2018, in preparation), further motivates the search for long-period pulsars. Furthermore, optimizing our detection capabilities at low modulation frequencies increases the chances of discovering the first neutron star—black hole binary system. Since the black hole will presumably have resulted from the supernova explosion of the initially more massive star in the binary, a pulsar companion will not have been recycled and so would generally have similar periods to the nonrecycled pulsar population (Lipunov et al. 2005; Pfahl et al. 2005; Eatough 2007). Such a discovery could provide valuable insights into stellar evolution and serve as a test bed for theories of gravity. Increased sensitivity to low modulation frequencies also makes pulsar surveys more likely to find radio-loud magnetars: the four known radio-loud magnetars have rotation periods between 2 and 6 s (see, e.g., Kaspi & Beloborodov 2017).

The use of the FFA has been fairly limited over the past decades. Lovelace et al. (1969) implemented the algorithm when working at the Arecibo Observatory, resulting in the discovery of PSR B2016+28 ($P = 0.56$ s; Craft et al. 1968). The Parkes Multibeam Pulsar Survey used the FFA to search for periodic signals in the data collected by the survey, which led to the discovery of the 7.7 s pulsar J1001–5939 (Faulkner et al. 2004; Lorimer et al. 2006). It was also used in a search for radio pulsations in observations of the 6.85 s X-ray pulsar XTE J0103–728, but it resulted in no significant detections (Crawford et al. 2009). Kondratiev et al. (2009) used the FFA to perform a search for periodicity on radio observations of six X-ray-dim isolated neutron stars (XDINSs) and then compared the sensitivity of the time-domain algorithm to that of a typical Fourier-based technique. This work demonstrated the ability of the FFA to exceed the performance of the FFT in the white-noise regime, especially when searching for pulsars having high harmonic content. Cameron et al. (2017) recently obtained results similar to those presented by Kondratiev et al. (2009), where an in-depth study of the behavior of the time-domain algorithm was conducted both in a Gaussian noise regime and in real observational data collected by the High Time Resolution Universe (HTRU) pulsar survey. This analysis showed an enhancement in the detectability of long-period pulsars when using the FFA in the two regimes. The use of the algorithm also extends exoplanet hunting, which is similar to pulsar searches, except that dips are observed in the time series rather than pulses. It was used to search for transits by Earth-size planets around G- and K-type dwarf stars in *Kepler* data (Petigura et al. 2013), and it led to the discovery of a number of exoplanet candidates.

Deploying an FFA-based search in a large-scale pulsar survey is computationally expensive, and this is the main reason why the use of this alternative technique has been limited in the past. Nevertheless, the increasing power of modern supercomputers allows us to use the FFA in a large-scale pulsar survey.

In this paper, we present the results from the implementation of an FFA-based search, *ffaGo*,²⁹ in the PALFA survey. We

²⁷ A maximum of 32 summed harmonics are used in the case of the PALFA PRESTO-based pipeline; see Lazarus et al. (2015).

²⁸ <http://www.astron.nl/lotaas/>

²⁹ Available at <https://github.com/emilieparent/ffaGo>.

compare the efficiency of `ffaGo` with that of an FFT pulsar searching program both in the ideal, white-noise regime and in real PALFA survey data. The expected sensitivity of the FFA in the large-scale PALFA survey is evaluated by reproducing an analysis similar to that presented in Lazarus et al. (2015), where various pulsar signals are injected in a selection of PALFA observation files free of astrophysical signals and then recovered using `ffaGo` to determine the minimum mean flux density our FFA-based pipeline can detect in the PALFA survey.

The organization of this paper is as follows: Section 2 offers a brief mathematical description of the FFA. Details regarding the implementation of the algorithm and the testing of significance metrics used to evaluate FFA-generated profiles in the PALFA survey are discussed in Section 3. In Section 4, we compare the performance of the FFA with that of the FFT using both simulated and real data collected at Arecibo containing long-period pulsars. We then report on the sensitivity analysis conducted with the FFA in Section 5, where we recover synthetic pulsar signals injected in real PALFA data. Section 6 presents the results from the implementation of the time-domain algorithm in PALFA, along with new discoveries made by the survey. Finally, we summarize the main results of this paper in Section 7.

2. The Fast-folding Algorithm

The FFA was originally developed by Staelin (1969) for searching for periodic signals in the presence of noise in the time domain, in contrast to the FFT search technique, which operates in the frequency domain. By avoiding redundant summations, the FFA is much faster than standard folding at all possible trial periods: it performs summations through $N \log_2(N/p - 1)$ steps rather than $N(N/p - 1)$, where N and p are the number of samples in the time series and the trial folding period in units of samples, respectively. Large computational power is still required when applying the FFA over a very wide range of trial periods, and this is why the use of the FFA in large-scale pulsar searches has been limited in the past.

The FFA folds each dedispersed time series with sampling interval Δt at multiple periods (p , in units of sample time), and our implementation of the algorithm then looks for statistically significant features in the generated profiles. The algorithm performs partial summations, while avoiding redundancy, into a series of $\log_2 p$ stages and then combines those sums in different ways so that the data are folded with a trial period between p and $p + 1$. A time series containing N time samples folded in an FFA execution at the folding period p (corresponding to a period in time units of $P = p \times \Delta t$) will result in $M = N/p$ different pulse profiles with slightly different periods ranging from p_i to $p_i + 1$:

$$p_i = p_0 + \left(\frac{i}{M - 1} \right), \quad (1)$$

where p_0 is the effective folding period and $0 \leq i \leq M - 1$.

While the folding procedure is a core component of an FFA-based search, the statistical evaluation of the resulting profiles is another crucial component of the search. This is discussed in Section 3.3. Figure 1 shows an example periodogram one obtains from applying the FFA on a 268 s PALFA observation

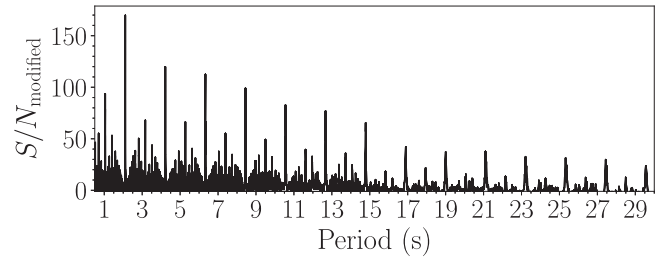


Figure 1. Periodogram of PSR J2004+3137 generated by `ffaGo`. One can clearly identify the fundamental period of the pulsar ($P = 2.11$ s), as well as many harmonics.

of the bright, long-period pulsar J2004+3137, when looking for periodicity between 500 ms and 30 s. The pulse profile of this source is shown in Figure 2. The peak in signal-to-noise ratio (S/N) is at the pulsar’s fundamental period, 2.11 s, and the secondary peaks are the harmonics and subharmonics of the spin period. The FFA requires $\log_2(N/p)$ to be an integer, or equivalently, M to be a power of 2. If this condition is not satisfied, our implementation of the algorithm will pad the time series by its median value. A more complete description of the FFA algorithm can be found in Staelin (1969), Lovelace et al. (1969), and Lorimer & Kramer (2004).

The main advantages of the FFA over the FFT are that the FFA offers greater frequency resolution (especially important in the low-frequency end of the spectrum) and, most importantly, that it coherently sums all harmonics of a signal (i.e., it folds the data in phase). Indeed, the incoherent harmonic summing that is used in Fourier-domain searches inevitably misses power in higher harmonics, since one must choose a finite number of harmonics to be summed when using this technique. Hence, the FFA is more sensitive to narrow pulses.

3. An FFA-based Pipeline in the PALFA Survey

PALFA has two independent search pipelines performing a full-resolution analysis: the primary PRESTO-based pipeline (Lazarus et al. 2015) and the Einstein@Home-based pipeline (Allen et al. 2013). A reduced-resolution analysis is also performed on site at the Arecibo Observatory: this “Quicklook” pipeline (Stovall et al. 2013) allows rapid discovery and confirmation of bright pulsars. The work presented here will focus only on the PRESTO-based pipeline, which has been modified to additionally perform the FFA-based search for long-period pulsars.

This pipeline runs on the Guillimin Supercomputer, part of McGill University’s High Performance Computing Centre operated by Compute Canada and Calcul Québec. PALFA data are transferred from the Arecibo Observatory to the Cornell University Center for Advanced Computing (CAC), from where they are downloaded to Guillimin. The results from the data processing pipeline are uploaded upon completion to the PALFA database, also located at the CAC, for future human inspection.

In the PRESTO-based pipeline, the 4-bit data files (PSRFITS format) are first subject to RFI mitigation routines. The data are then dedispersed at a wide range of trial DMs. A Fourier-based periodicity search is subsequently performed on the dedispersed time series using PRESTO’s `accelsearch` software. The pipeline also has a single pulse search component (Patel 2016) that searches for single, dispersed

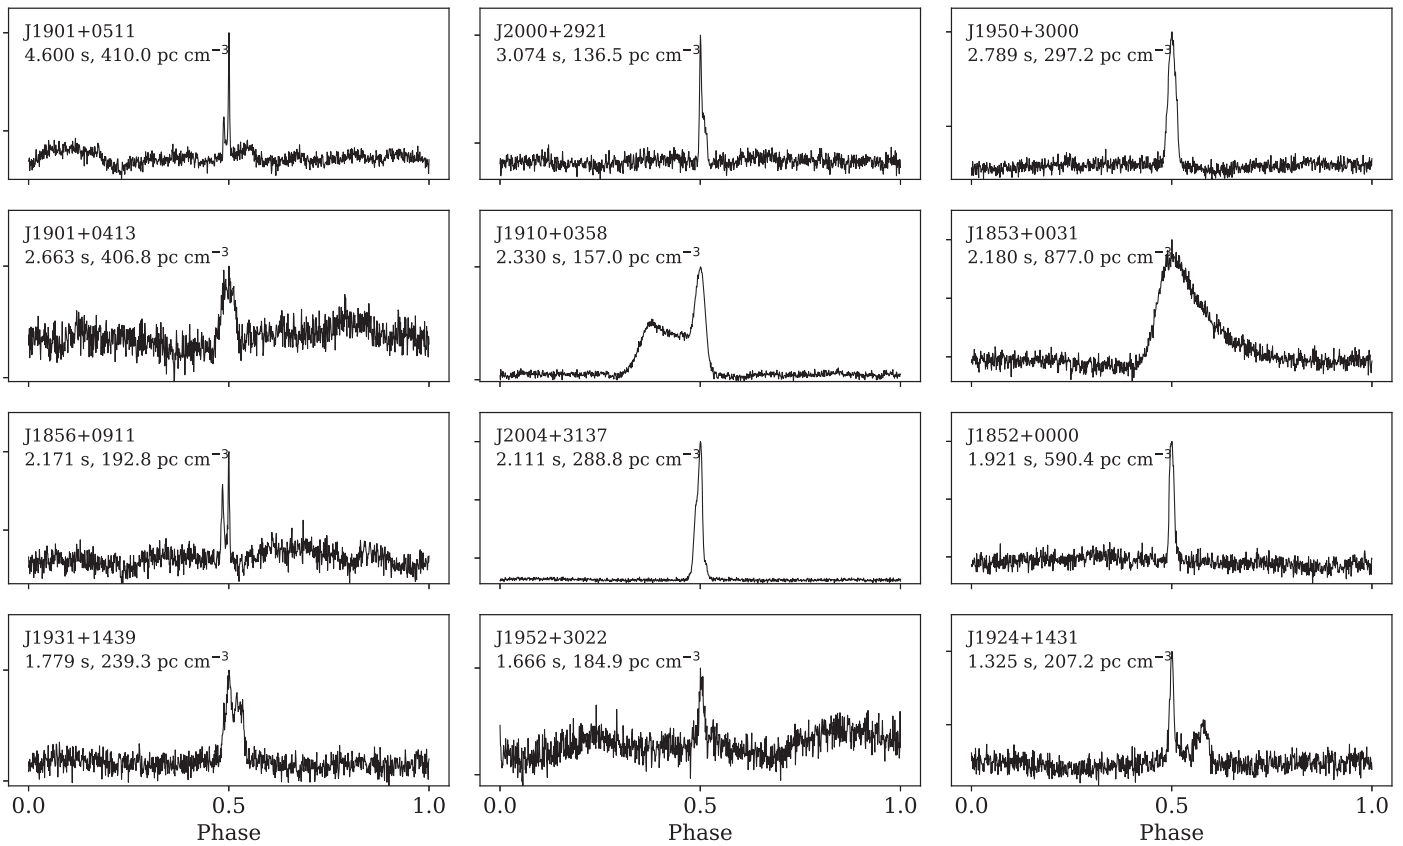


Figure 2. Pulse profiles of 12 long-period pulsars discovered with the PRESTO-based PALFA pipeline in observations with the Mock spectrometer at 1.4 GHz. The profiles were folded using PRESTO’s `prepfold` program. The name, period, and DM of the pulsars are specified above each profile. One can see the broad features in the baseline introduced by red noise and interference in the data, especially prominent for PSR J1901+0413, PSR J1856+0911, and PSR J1952+3022.

pulses up to DM values of $10,000 \text{ pc cm}^{-3}$ (C. Patel et al. 2018, in preparation). The PALFA Consortium then uses the online collaborative tool on the CyberSKA platform³⁰ (Kiddle et al. 2011) to classify generated pulsar and transient candidates. For more details about PALFA’s data processing, see Lazarus et al. (2015).

3.1. Implementation of the FFA in the PALFA Pipeline

We have designed a Python program, `ffaGo`,³¹ that implements the FFA-based periodicity search into the PALFA analysis software. `ffaGo` reads any 32-bit float time series produced by PRESTO, and it includes a de-reddening procedure aimed at reducing the effect of red noise on the input time series. This de-reddening is done by applying a dynamic median filter, where the size of the filtering window is, by default, set to twice the largest trial period searched. To shorten FFA executions, downsampling of the data is performed initially such that the sampling interval is approximately 2 ms. The data are then normalized by dividing by the maximum value before calculating the standard deviation, σ , of the time series for future profile evaluations (see Section 3.1.2). Subsequent dynamical rebinning routines are carried out to search for multiple pulse widths.

Parts of our FFA code are taken from an open-source FFA package,³² written as a Python and C program developed for

transit searches in *Kepler* data (Petigura et al. 2013). More specifically, the parts of our code that wrap the time series, pad it, and perform the folding and the summations were taken from Petigura et al. (2013).

S/N calculations, candidate selection, and sifting are also incorporated in this program. Periodograms similar to Figure 1 can also be generated by `ffaGo`. We note that the primary focus while designing the CPU-based `ffaGo` was not to minimize the computation time. Large-scale, real-time analyses should consider parallelized versions of FFA-based searches.

3.2. Search Parameters

The pulsar parameter space that we consider in the implementation of `ffaGo` in the pipeline consists of the following:

- A. *The Period:* We search for periods ranging from a minimum of 500 ms to a maximum of 30 s. Even though the FFA is designed to be fast, it is still computationally expensive to apply to higher modulation frequencies, since they produce a large number of profiles that need to be statistically evaluated. This in turn results in an important increase in the computational burden: searching down to 100 ms nearly doubles the time required to process one time series with `ffaGo`. This is one of the reasons why the blind search is restricted to periods longer than 500 ms. Moreover, Lazarus et al. (2015) demonstrated that 500 ms is approximately the period at which one notices a decrease in the sensitivity of PALFA at low DMs. It is not worth looking for periods larger

³⁰ <https://ca.cyberska.org/>

³¹ Available at <https://github.com/emilieparent/ffaGo>.

³² Available at <https://github.com/petigura/FFA>.

than 30 s with `ffaGo` since the integration times of PALFA observations are 268 s and 180 s for the inner ($32^\circ \lesssim l \lesssim 77^\circ$) and outer ($168^\circ \lesssim l \lesssim 214^\circ$) Galactic regions, respectively: it is unlikely that folding fewer than ~ 10 pulses will result in significant detections, especially in the presence of red noise. We rely on the single pulse search conducted in the pipeline to identify pulses from very slow ($P > 30$ s) pulsars (C. Patel et al. 2018, in preparation).

- B. *The Pulse Width:* To explore the pulse width parameter while optimizing the S/N and minimizing the computation time, we perform rebinning by a factor of X at multiple stages during the search such that the sampling interval ranges from ~ 2 ms up to a few seconds, depending on the trial period and the pulse duty cycle δ we are searching for at each step of the process. The PALFA time series, which initially have a sampling interval of $65 \mu\text{s}$, are first decimated so that each bin has a width of approximately 2 ms. Afterward, we divide the 500 ms–30 s full range of trial periods into six subranges, processed separately, such that the fixed sampling interval is no smaller than $1/1000$ of the shortest trial period in the subrange and no larger than $1/100$ of that period. In other words, the minimum δ we search is kept between 0.1% and 1%, when assuming a pulse fully enclosed within one bin. We impose this lower limit on the searched range of pulse widths to reduce the execution time. Additional rebinning is applied to the time series before entering FFA executions in each subrange to ensure that the ratio of the sampling interval to the shortest trial period is greater than $1/1000$. Further downsampling of the time series is performed within each period subrange in order to efficiently search δ values ranging from approximately 0.2%–0.5% up to 10%–13%. The downsampling factors we use are 2^k and 3^k , where $1 \leq k \leq 3$. To ensure optimal sensitivity, this last downsampling stage is performed at different phases (i.e., adjacent bins are summed in different ways).
- C. *The DM:* Since the values of DM of the pulsars to be discovered are unknown, a large number of DM trials must be used in the search. We search with the FFA from $\text{DM} = 0$ to 3265 pc cm^{-3} in steps of 5 pc cm^{-3} , resulting in 653 dedispersed time series to be processed through `ffaGo`. Using finer DM steps is unnecessary, as we are searching in the long-period phase space, where the pulse widths are typically from a few to hundreds of milliseconds. The only scenario where our sensitivity could be affected by this coarse DM spacing is one where a pulsar had a value of DM that sits exactly between two trial DMs, which corresponds to a dispersive smearing of 2.6 ms, and if that particular pulsar had a short spin period and a narrow pulse width (for example, shorter than 500 ms and a pulse duty cycle smaller than 0.5%). We are searching up to DMs higher than the maximum Galactic value predicted by NE2001 (Cordes & Lazio 2002), which is about 2000 pc cm^{-3} in the region surveyed by PALFA, to account for any possible dense, local regions that could not be included in the model. The DM step size was chosen such that the amount of processing is minimized while avoiding sensitivity loss from channel smearing due to dispersion. We are not searching above $\text{DM} = 3265 \text{ pc cm}^{-3}$ since the probability of finding normal

pulsars with mean flux densities of a few mJy outside our Galaxy observation is quite low considering the relatively short integration time of PALFA observations (see Section 3).

3.3. Profile Evaluation

The significance metric that we use to evaluate profiles generated by the algorithm assumes that the profile has one single-peaked pulse, that this pulse is constant in phase, and that it is captured within a single bin (i.e., the detection is optimal when the bin size is equal to the width of the profile). The mathematical description of the metric (Metric A) is as follows:

$$\text{S/N} = \frac{I_{\text{max}} - I_{\text{med}}}{(\sigma\sqrt{X})\sqrt{M - z}}, \quad (2)$$

where I_{max} and I_{med} are the maximum and the median intensities of the folded profile, and σ is the standard deviation of the time series, calculated after the initial downsampling, detrending, and normalization of the time series. Subsequent rebinning is accounted for by multiplying the standard deviation by the square root of the downsampling factor, X . Finally, z is the fraction of a profile that requires padding such that the necessity of the number of profiles M being a power of two is respected.

We also explore other metrics for evaluating profiles, such as one in which the median and the standard deviation would be calculated only over the off-pulse portion of the profile, so that the on-pulse component is not included when statistically characterizing the baseline noise in each profile. Kondratiev et al. (2009) and Cameron et al. (2017) used such a metric to evaluate profiles generated by an FFA program.³³ Specifically, we tested Metric B, where we exclude a 20% window centered on the peak of the profile when calculating the median, $I_{\text{med,off}}$, and the standard deviation, σ_{off} , of the profile. As opposed to Metric A, in which the denominator of the expression for the S/N is constant for a given FFA execution, the standard deviation σ_{off} in Metric B is calculated directly on the off-pulse portion of individual profiles produced within an FFA execution. While using this algorithm to evaluate FFA-generated profiles, we explore the pulse width phase space by applying the downsampling procedure described previously, rather than a boxcar matched-filtering approach (Cordes & McLaughlin 2003), as was done in Kondratiev et al. (2009) and in Cameron et al. (2017). The S/N of the peak in each FFA-generated profile is then calculated as follows:

$$\text{S/N} = \frac{I_{\text{max}} - I_{\text{med,off}}}{\sigma_{\text{off}}}. \quad (3)$$

To compare the efficiency of Metric A and B, we performed a search using both metrics on a data set of simulated pulsar signals constructed with SIGPROC's³⁴ `fake` program, which injects periodic top-hat pulses in Gaussian noise. The synthetic pulsars have spin periods, P , ranging from 2 to 20 s (in increments of 2 s), with pulse duty cycles δ of 0.5%, 1%, and from 2% to 20% with a step size of 2%, resulting in 120 different trial combinations of period/pulse width. Each of

³³ The respective programs can be downloaded from <https://github.com/vkond/ffasearch> and <https://github.com/adcameron/ffancy>.

³⁴ <https://github.com/SixByNine/sigproc>

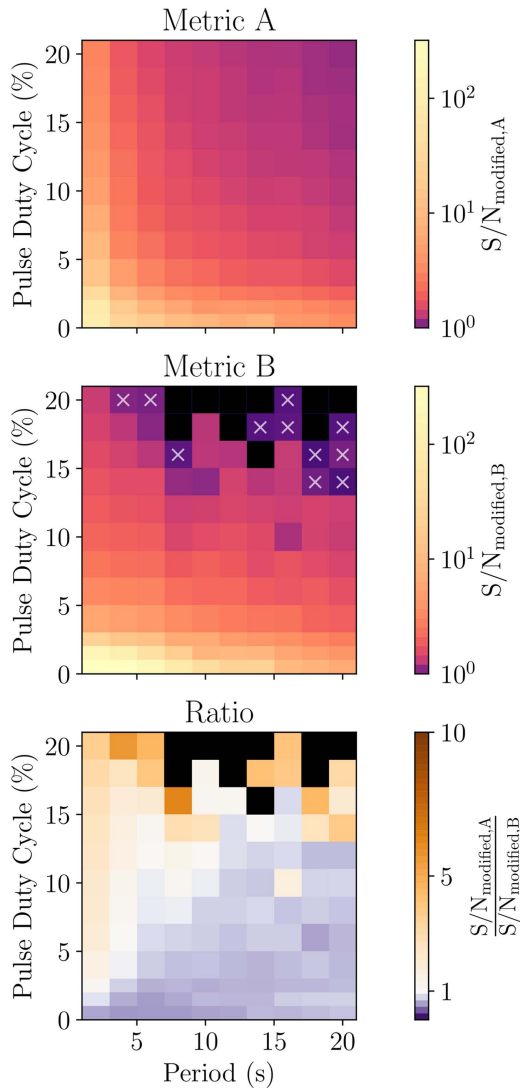


Figure 3. Response patterns of the two FFA significance metrics, Metric A (top panel) and Metric B (middle panel), investigated in the white-noise simulation described in Section 3.3. The ratio of the two values of S/N_{modified} is shown in the bottom panel. The values reported are the average S/N_{modified} from the five simulations. Black pixels represent trials that were not detected in all five data sets, while pixels with white crosses represent those having an average S/N_{modified} below 6 (i.e., trials that were classified as nondetections).

these trials was constructed and tested five times to ensure that no statistical anomalies were introduced in our data set when using the `fake` program. In total, 600 data files were searched with both metrics. The amplitude of individual pulses, S , was chosen such that the total pulse energy, $E = PS\delta$, was kept fixed for each trial. Broader pulses therefore have lower peak fluxes compared to narrow pulses. The sampling interval of the fake observations was set to $65 \mu\text{s}$ with a 268 s integration time at a central observing frequency of 1375 MHz and a bandwidth of 322 MHz to match real PALFA data when observing inner Galaxy regions. The DM value at which all signals were injected was arbitrarily chosen to be 150 pc cm^{-3} .

The simulated observation files were dedispersed at the appropriate DM prior to searching periodicities between 500 ms and 30 s with both metrics. Once the search was completed, the lists of candidates were inspected by eye to identify the highest S/N_{modified} values (Section 3.4 describes how S/N_{modified} differs

from S/N) at which the artificial pulsars were detected. The results of this simulation are shown in Figure 3.

The response pattern from Metric A shows that it provides the best detections for narrow-pulsed, short-period signals. The optimal detection occurs at the shortest trial period of 2 s and $\delta = 0.5\%$. The S/N_{modified} values then gradually fall off. This is expected because, for longer periods/wider profiles, the amplitude of the pulse is reduced since we require the total pulse energy to remain constant.

For Metric B, the response pattern suggests that the determining factor when it comes to the metric responsiveness is the pulse width: this metric responds strongly to narrow profiles, and its sensitivity decreases only slightly with increasing period. Moreover, this metric reaches higher S/N_{modified} values for the trials with narrow pulse widths compared to Metric A. Metric B remains significantly responsive up to $\delta = 8\%–12\%$, above which it practically vanishes. This behavior is also shown in the bottom panel of Figure 3, where we see that, for all periods, Metric B is outperformed by Metric A at large values of pulse duty cycle δ . Figure 3 also suggests that Metric A is better at detecting signals with short periods ($P \lesssim 4$ s) and δ larger than $\sim 2\%–5\%$. However, we also see that Metric B yields larger S/N_{modified} values than Metric A for narrow-pulsed signals having long periods.

One clear distinction between the two metrics is that Metric A detected all artificial pulsars, while 10 trials having broad profiles were missed by Metric B in all five simulations (black pixels in Figure 3). Furthermore, 11 trials were detected by Metric B with an average S/N_{modified} below the threshold for candidate folding set in the pipeline, meaning that we consider those trials as being not successfully detected by Metric B. Therefore, 21 out of 120 fake pulsars were not detected by Metric B.

Cameron et al. (2017) also investigated a significance metric similar to Metric B when evaluating FFA-generated pulse profiles and concluded that even if such a metric possesses the ability to outperform the FFT in the long-period regime, it suffers from sensitivity deterioration when it comes to broad pulses. This characteristic can, however, help in reducing the number of false positives generated by red noise in the data. The analysis presented here is consistent with the results presented in Cameron et al. (2017) and demonstrates that it is likely that the survey would miss pulsars having broad profiles if this metric were used in the FFA search. For an interpretation of the difference in the performance of the two metrics, see Appendix A.

We also designed an alternative, Metric C, which, similarly to Metric B, excludes a 20% window centered on the peak to calculate the median intensity of the profile, $I_{\text{med,off}}$. The standard deviation of Metric C is similar to that used in Metric A, only we include an extra factor of $\sqrt{0.8}$ in the profile's standard deviation to account for the on-pulse exclusion (see Equation (5) in Appendix B). The same set of synthetic pulsars injected in white noise described above was searched with Metric C. Results from this analysis suggest that Metric C has a response pattern very similar to Metric A and that there is no significant difference between the two metrics. Unlike Metric B, Metric C suffers negligible loss in sensitivity for large δ values. Therefore, we conclude that Metric A and Metric C are equivalent. More details on profile evaluation with Metric C,

including the response pattern obtained from the white-noise simulation, can be found in Appendix B.

Due to the nondetection of wider pulses by Metric B, we opted for implementing Metric A to evaluate FFA-generated profiles in the PALFA processing pipeline, which successfully detected all trials and showed a response pattern that suggests overall broader sensitivity. We note that Metric C would also have been a reasonable option. When downloading `ffaGo`, the user can select any of the three metrics described in this work.

3.4. Candidate Selection

For each dedispersed time series processed through the pipeline, all FFA-generated profiles are statistically evaluated (see Section 3.3) to identify periodic signals. A set of S/N values (i.e., a periodogram) is produced each time we downsample the initial time series at a specific phase (i.e., at each possible way of summing adjacent bins) by a factor of 2^k or 3^k , as described in Section 3.2. These sets have different statistical distributions, because the number of profiles generated for a specific period will vary as the number of samples in the rebinned time series changes. To avoid being biased in the candidate selection process, we make the S/N sets uniform by subtracting the mode of that S/N value's distribution and then dividing by its median absolute deviation (MAD):

$$S/N_{\text{modified},i} = \frac{S/N_i - \text{mode}_i}{\text{MAD}_i}, \quad (4)$$

where i represents a specific set of S/Ns (i.e., the periodogram obtained for a specific rebinned time series). All candidates are therefore characterized by a modified S/N value, S/N_{modified} , which estimates the significance of the S/N calculated by the selected metric. The mode and the MAD were chosen for their robustness when evaluating statistics of largely skewed distributions, as is the case when pulsar signals are present in the data.

All candidate periods detected with an $S/N_{\text{modified}} \geq 5$ are recorded to a list along with the S/N_{modified} , the sampling interval, and the value of DM at which the candidate was detected. This is done for all 653 dedispersed time series, and the full FFA search uses approximately 10% of the PALFA pipeline total processing time, which corresponds to a few hours. The set of candidate lists is subsequently sifted using a modified version of PRESTO's sifting routine, also included in the open-source `ffaGo` package. This sifting removes weaker, harmonically related periods and RFI-like signals and groups candidates according to their DM. More details regarding the general candidate sifting procedure can be found in Lazarus et al. (2015).

Once the time series have been searched and the FFA candidates have been sifted, only candidates having $S/N_{\text{modified}} \geq 6$ are selected for folding. This limit is also applied to the candidates produced by `accelsearch` in the PALFA pipeline to reduce the number of false positives that have to be inspected. The raw data are folded with PRESTO's `prepfold` routine at each candidate period. Similarly to FFT-generated candidates, we do not allow `prepfold` to search in period and DM space if the candidate has a period greater than 500 ms to avoid converging to nearby RFI. The resulting plots, along with ratings calculations (Lazarus et al. 2015) and one rating from a candidate-ranking artificial intelligence (AI)

system (Zhu et al. 2014), are then uploaded to a PALFA online Candidate Viewer application for final human inspection and classification. FFA-generated candidates generally represent approximately 10%–25% of the total number of folded periodicity candidates, which varies between 150 and 250 total candidates per beam.

4. Comparing the FFA to the FFT

4.1. Comparison Using Simulated Data

To compare the performance of the `ffaGo` program to that of a typical Fourier-based search, PRESTO's `accelsearch` program was applied to the five data sets of 120 artificial pulsar signals that were used in the analysis presented in Section 3.3. The Fourier-based search summed up to 32 harmonics incoherently, and the significance of the FFT candidates was characterized by a σ_{fft} value, the quantity used in the PALFA survey to evaluate the strength of an FFT candidate. The value of σ_{fft} is determined by calculating the equivalent Gaussian significance of the candidate based on the probability that the same amount of incoherently summed power is noise. In the PALFA pipeline, candidates with σ_{fft} values greater than 2 are recorded to a list of candidates that are later sifted, but only candidates with σ_{fft} above 6 are folded and uploaded to the online Candidate Viewer for human inspection. Therefore, we consider here only signals having $\sigma_{\text{fft}} \geq 6$ as successfully detected by the program. The S/N_{modified} from the FFA search (Metric A) and the σ_{fft} from `accelsearch` at which the simulated pulsars were detected were recorded for the two periodicity searches, and the strengths of the detections are illustrated in Figure 4. It is important to note that the types of statistics used to characterize the detections made by the algorithms are fundamentally different. Therefore, numerical scores from the two searches should not be directly compared.

Both FFA and FFT searches show similar response patterns with similar regions of maximum sensitivity: even under ideal white-noise conditions, the detected S/N_{modified} values decrease with increasing period and increasing pulse width (i.e., decreasing peak amplitude). This is expected since we require the per-pulse energy to be constant and there are fewer pulses in the 268 s time series when injecting longer periods. The response from the frequency-domain algorithm, however, falls off more sharply with period as compared to the time-domain search.

A major difference that arises between the two techniques is that, while the FFA successfully recovered all trials, `accelsearch` detected 10 trials (pixels with white crosses in Figure 4) showing broad profiles with an average σ_{fft} value below 6 (some of these trials were totally missed by `accelsearch`). These are not considered as successful detections since such candidates would have been excluded from the final list of potential candidates generated by the processing pipeline. While we expect the FFT to be particularly sensitive to signals having low harmonic content, the lowest modulation frequencies are effectively searched via their highest harmonics, and, in the PALFA processing pipeline, `accelsearch` searches down to a minimum of 1 Hz. The program is therefore intrinsically less sensitive to very long period pulsars having low harmonic content. This restriction on the lowest frequencies searched is set in order to reduce the number of false-positive candidates produced by red noise in the data. This explains why the algorithm is outperformed by

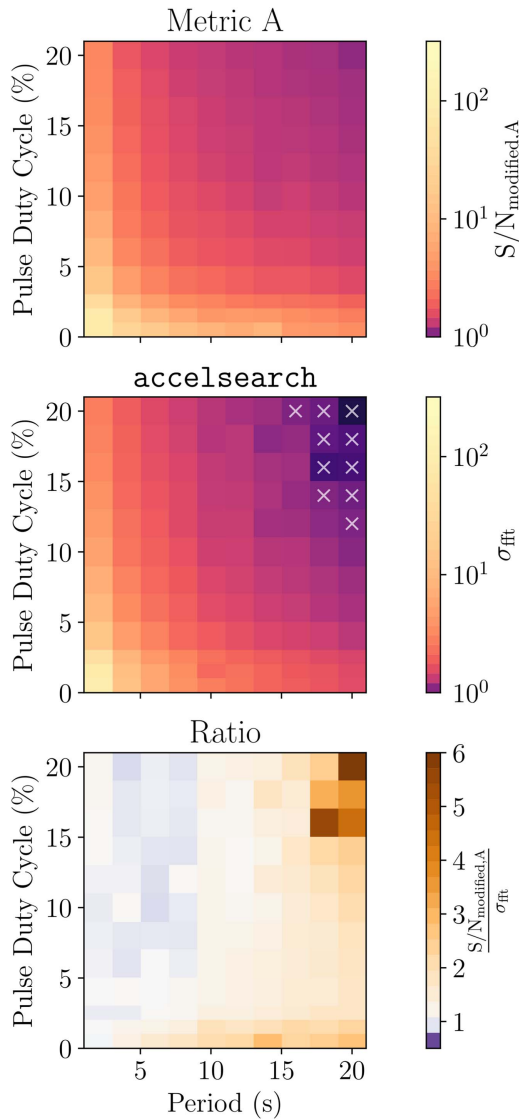


Figure 4. Response patterns of the FFA when using Metric A (top panel) and `accelsearch` (middle panel), investigated in the white-noise simulation described in Section 3.3. Pixels with white crosses represent those having an average σ_{fit} below 6. The bottom panel represents the ratio of the S/N_{modified} over the σ_{fit} for each trial. Although the numerical values of the ratios do not reflect directly the sensitivity gain achieved by the FFA, they allow us to visualize where the improvement is maximal. The values reported are the average S/N_{modified} from the five simulations. Note that the scales for the top and the middle panels are logarithmic, while the bottom panel is displayed on a linear scale.

the FFA in the broad pulse regime and why some trials were missed by the frequency-domain search.

The bottom panel in Figure 4 shows the ratio of the S/N_{modified} to the σ_{fit} values. The resulting pattern can be used to illustrate the phase space where the use of the FFA is the most advantageous. Although the two numerical scores cannot be compared directly owing to the fundamental difference in their nature, the displayed pattern suggests that there are two particular regions where the FFA is more responsive. First, we see that the coherent summing of all harmonics makes the time-domain algorithm more efficient at finding the pulsar signals having the smallest pulse widths, and this advantage grows with increasing period. The second region is where trials have

the broadest pulses and the lowest spin frequencies. We emphasize once more the arbitrary nature of the values of ratio shown in Figure 4, especially considering the fact that the two quantities compared do not scale equivalently to increasingly bright signals.

In summary, this analysis demonstrated the ability of the FFA to outperform the frequency-domain search in the long-period regime in the presence of white noise. Similar studies were carried out by Kondratiev et al. (2009) and Cameron et al. (2017) and also demonstrated that, even if every trial were detected by an FFT, the performance of an FFA exceeds that of an FFT. We also showed that an FFT can fail to detect broad signals with $P > 18$ s even in ideal conditions for a 268 s integration time. This shows that even in the absence of red noise, the coherent summing of all harmonics is necessary to detect some long-period pulsars.

A similar simulation is presented in Section 6, where artificial pulsars have been injected in real observational data rather than in white noise to quantify the efficiency of the FFA when searching for pulsars in a large-scale survey under real RFI conditions.

4.2. Comparison Using Real Pulsar Data

To evaluate the response of `ffaGo` to pulsars in the presence of RFI and red noise and compare it to an FFT-based search, we applied the program to a data set of 12 PALFA observations collected at the Arecibo Observatory containing a variety of long-period pulsars discovered by the survey (Swiggum et al. 2014; Lazarus et al. 2015; Lyne et al. 2017). We then compared the significance of the detections from the FFA search with that obtained by `accelsearch`. We also processed the data set through the FFA using Metrics B and C to evaluate their responses in the presence of red noise.

The selected observations contained pulsar signals covering a period range from 1.32 to 4.6 s and having values ranging from less than 1% up to $\sim 10\%$. The pulse duty cycle (δ) values reported in Table 1 were measured by calculating the fraction of bins with intensity larger than half the maximum value in the integrated pulse profiles. While most of the sources display single-peaked profiles, some pulsars from our data set have two-component profiles (see profiles in Figure 2). For example, PSR J1901+0511 and PSR J1856+0911 both exhibit two narrow, closely spaced pulse components, while PSR J1924+1431 has a broad component and a narrow component that are separated in phase. We were also interested in quantifying the detectability of pulsars having broad profiles, such as PSR J1852+003 and PSR J1910+035, in the red-noise regime. When considering the width of the entire pulse (i.e., the portion of the profile around the peak that is above the baseline intensity), the on-pulse fractions for these two sources are 30.5% and 21.7%, respectively (but they have δ of 9.7% and 3.3%, respectively, when they are calculated via their pulse FWHM).

Prior to dedispersing of the `PSRFITS` observation files at the appropriate DMs of the pulsars, the data were cleaned of interference by applying `PRESTO`'s `rfifind` routine, which identifies narrowband RFI and produces a mask for bad time and frequency intervals. To optimize detections, we produced time series dedispersed at multiple DM values around the true DM of the pulsars.

Table 1
Results from the Analysis Performed on 12 Long-period Pulsars Discovered by the PALFA Survey

PSR Name	Period (s)	Pulse Duty Cycle (% of Phase)	FFA S/N_{modified}			FFT σ_{fft}
			Metric A	Metric B	Metric C	
J1901+0511	4.600	0.4	33.3	35.8	24.8	12.0 ^a
J2000+2921	3.074	0.8	55.0	63.4	47.1	18.5 ^a
J1950+3000	2.789	2.2	78.9	92.2	72.7	12.1
J1901+0413	2.663	3.1	8.5 ^a	9.7 ^a	7.2 ^a	9.4
J1910+0358	2.330	3.3	47.3	12.8	14.5 ^a	31.3
J1853+0031	2.180	9.7	84.1	19.4	76.3	68.1
J1856+0911	2.171	0.5	21.7	24.3	17.7	8.3 ^a
J2004+3137	2.111	1.6	169.5	163.4	157.9	104.4
J1852+0000	1.921	1.3	59.5	55.0	54.4	34.7
J1931+1439	1.779	1.6	65.2	72.8	52.27	22.1 ^a
J1952+3022	1.666	1.0	15.1	...	13.2	7.8 ^a
J1926+0431	1.325	1.1	53.1	51.2	47.9	27.0

Note.

^a Detection via a harmonic of the fundamental frequency of the pulsar.

We processed each masked and dedispersed time series through `ffaGo` (searching for periodicities ranging from 500 ms to 30 s), as well as through `PRESTO`'s `accelsearch`, to search in the Fourier domain while incoherently summing up to 32 harmonics. The candidate periodicities from both searches were then separately sifted, and the lists of final candidates were then inspected by eye to identify the strongest candidates harmonically related to the pulsar.

Metrics A and C, as well as the FFT search, were successful at detecting the full selection of pulsars (see results in Table 1). Except for PSR J1901+0413, for which the detection was marginal ($S/N_{\text{modified}} \lesssim 10$), the S/N_{modified} of the FFA detections with Metric A were all well above the threshold $S/N_{\text{modified},i} \geq 6$ that we consider for candidate folding, meaning that the regular processing pipeline would have folded the pulsars for final human classification. All pulsars were detected at their fundamental frequency when using Metric A in the FFA search, while there were five instances where the FFT search detected pulsars via their harmonics. Moreover, four of the detections made with `accelsearch` were marginal detections ($\sigma_{\text{fft}} < 10$). The important conclusion we draw from this analysis is that, even if the numerical scores of the detections made by both FFA and FFT searches cannot be compared directly, the FFA successfully recovered the true period of a variety of pulsars with different pulse profiles (some having multiple components), at S/N_{modified} values significantly larger than the detectability threshold set in the PALFA pipeline. The FFT search detected a number of these sources at harmonics of their spin frequencies and in some cases only marginally.

When using Metric B, there were two cases where the source was missed. Interestingly, it was neither the longest-period nor largest δ pulsars that were missed. From their pulse profiles we can see that the nondetected pulsars are the ones that have baselines showing broad features introduced by red noise. This further motivates the choice of Metric A over Metric B for FFA-generated profile evaluation in the pipeline implementation of the algorithm. In general, the significance of the detections with Metric C is marginally lower than Metric A's, and two pulsars were detected at a harmonic of the fundamental frequency. This suggests that Metric A is slightly more efficient in the presence of red noise.

5. Sensitivity of the PALFA Survey

To assess the true sensitivity of the PALFA survey, artificial pulsar signals were constructed and injected in real survey data using `PRESTO`'s `injectpsr` (described in Lazarus et al. 2015). This program generates smeared, scattered, and scaled pulse profiles that are added to real data at regular time intervals corresponding to the chosen spin period. To scale the profiles properly, observations of the radio galaxy 3C 138 (for which measurements of the flux density are available in the literature) were carried out in 2013 December. During these observations, a calibrating noise diode was turned on so that the flux density of the diode could be compared to that of the galaxy. Per-channel scaling factors between flux density and the observation data units were then calculated (see Lazarus et al. 2015, for more details on the calibration procedure). These scaling factors are used to obtain the targeted phase-averaged flux density (S_{mean}) of the artificial pulsar signals.

Signals constructed with `injectpsr` have single von Mises (von Mises 1918) component pulse profiles with FWHM specified by the user. Dispersive smearing and scattering are then applied to the profile, where the amount of broadening caused by scattering (in ms) is determined by the specified value of DM and the observing frequency, ν (in GHz), according to the following (Bhat et al. 2004):

$$\log \tau_{\text{scat}} = -6.46 + 0.154 \log \text{DM} + 1.07 (\log \text{DM})^2 - 3.86 \log \nu. \quad (5)$$

The modified data are then recorded into a SIGPROC “filterbank” file format. `injectpsr` can then be used to create a data set of synthetic pulsars for which we can adjust S_{mean} and then characterize the sensitivity of a survey at each point in (period, DM, pulse FWHM) phase space. More details regarding the construction of synthetic pulsars with `injectpsr` can be found in Lazarus et al. (2015).

5.1. Sensitivity of the PALFA Survey using a Fourier-based Search Technique

A realistic sensitivity analysis of the `PRESTO`-based PALFA pipeline was conducted in Lazarus et al. (2015) to evaluate the true performance of `accelsearch` at finding diverse types of pulsars in PALFA data.

Table 2
Parameters of the Synthetic Pulsar Signals Used
in the FFA Sensitivity Analysis

Parameter	Possible Values
Period (ms)	533.3, 1657.5, 3927.0, 5581.9, 10965.5, 14965.5 ^a , 21427.7 ^a
DM (pc cm ⁻³)	10, 40, 150, 600
FWHM (% Phase)	0.5 ^a , 1.5, 2.6, 5.9, 11.9

Note.

^a New trial elements that were not included in the sensitivity analysis conducted in Lazarus et al. (2015).

An important result from Lazarus et al. (2015) is that there is a clear mismatch between the PALFA survey sensitivity curves measured and the ideal case predicted by the radiometer equation (Dewey et al. 1985) when searching in the long-period regime. At low DMs, a reduction in the survey sensitivity is noticeable for spin periods as short as ~ 100 ms. At a spin period of ~ 11 s and with a pulse FWHM of 2.6%, the measured S_{\min} values are 10 and 20 times larger than the predicted value for DMs of 10 and 600 pc cm⁻³, respectively.

Lazarus et al. (2015) also injected synthetic pulsar signals in Gaussian noise. The measured curves in this scenario also rise at longer periods, although the degradation in sensitivity is not as pronounced as in the real data injection case. At a period of ~ 11 s and for DMs of 10 pc cm⁻³ and 600 pc cm⁻³, while the minimum detectable average flux densities measured from real data injections are approximately 10–20 times larger than the values predicted by the radiometer equation, the minimum detectable average flux densities measured from white-noise data injections are still 3–5 times larger than the predictions. This indicates that RFI and red noise alone cannot explain the discrepancy between the measurements and the predictions and that the periodicity-searching component of the pipeline is subject to potential improvement.

5.2. Sensitivity of the PALFA Survey Using the FFA Search

We reproduced the analysis described above to assess the sensitivity of the survey in the long-period phase space when using our implementation of the FFA (Metric A) to search for pulsar signals in PALFA data. Synthetic pulsars with periods longer than 500 ms were injected into the same real survey data that were used in Lazarus et al. (2015), and four of the DM trial values used in the previous analysis were selected. To avoid confusion with RFI, the trial periods were chosen to be noninteger values. We extended the period parameter space up to periods of ~ 15 and ~ 21 s to evaluate the responsiveness of the FFA in the long-period regime. Synthetic pulsars with FWHM pulses of 0.5%, 1.5%, 2.6%, 5.9%, and 11.9% were injected into the PALFA data sets, and only the signals with a pulse FWHM of 2.6% were injected in all 12 data files. The complete list of pulsar parameters used in the work presented here can be found in Table 2.

The minimum detectable mean flux density required for the FFA to detect the injected signals having an FWHM of 2.6% is shown in Figure 5. Similarly to the FFT-based search, the FFA sensitivity curves do not flatten out at longer periods, as opposed to what the radiometer equation predicts (Dewey et al. 1985). However, the degradation in sensitivity is not as pronounced as the FFT curves reported in Lazarus et al. (2015),

implying a gain in sensitivity: the FFA outperforms the FFT search for periods as short as 550 ms at $DM = 10$ pc cm⁻³. The performance of the FFA search does not seem to vary as strongly with DM at the longest periods compared to the FFT, which has a stronger response to pulsars having large values of DM. Figure 6 illustrates the factors of improvement in the sensitivity resulting from using the FFA- versus the FFT-based search. As expected, the gain in sensitivity is greater in the longer periods/low DM phase space.

The second part of the analysis consisted of injecting signals with different pulse FWHM in one of the observational data files. One can see from the results shown in Figures 7 and 8 that, for periods greater than ~ 8 s, the FFA is more efficient than the FFT for all values of DM and all pulse FWHMs. The gain in sensitivity is much greater when searching for pulsars having broad profiles. The advantage of a time-domain search over a frequency-domain search is that the coherent summing of all harmonics makes the search especially responsive to signals having narrow pulse profiles, while a Fourier transform should be more sensitive to profiles with low harmonic content. Recovering low modulation frequencies is, however, more difficult if red noise obscures the fundamental frequency, as well as low harmonic frequencies, which may contain a considerable fraction of the total power in the case of pulsars with wide profiles. In addition to the deterioration due to red noise, `accelsearch` searches very low modulation frequencies via the highest harmonics of a particular frequency, as mentioned in Section 4.1. This explains partially the important gain in sensitivity observed for broad pulses shown in Figures 7 and 8 and confirms the results from the white-noise simulation presented in Section 4.1.

Following these results, we inspected the `accelsearch` program to determine whether the loss in sensitivity for broad pulses is solely due to the effect of red noise. This investigation was further motivated by the inconsistency between the radiometer predictions and the calculations from the injections under ideal, white-noise conditions reported in Lazarus et al. (2015). An issue in `accelsearch` was identified and corrected.

Despite the noticeable improvement in the PALFA sensitivity when using the FFA to look for long-period pulsar signals in survey data, both the Fourier-based and FFA-based search types are not able to recover weak signals that should be detectable according to the radiometer equation. Nevertheless, this analysis demonstrates the ability of a time-domain search technique to outperform a Fourier domain technique when applied in large-scale pulsar surveys and suggests that the PALFA survey should discover long-period pulsars via the new implementation of `ffaGo` in the data processing pipeline.

6. Preliminary Results from the PALFA Survey

Along with the addition of the FFA search in the PALFA PRESTO-based pipeline, we have modified some of the search parameters of the earlier version of the Fourier-based periodicity search.

The number of harmonics incoherently summed in the *zero-acceleration* search, which is optimized to identify isolated pulsars, was initially set to 16. The new version of the search now sums up to 32 harmonics to increase our sensitivity to signals having narrow pulse profiles. This required changes to `accelsearch`. Doubling the number of harmonics summed in the *zero-acceleration* search approximately triples the

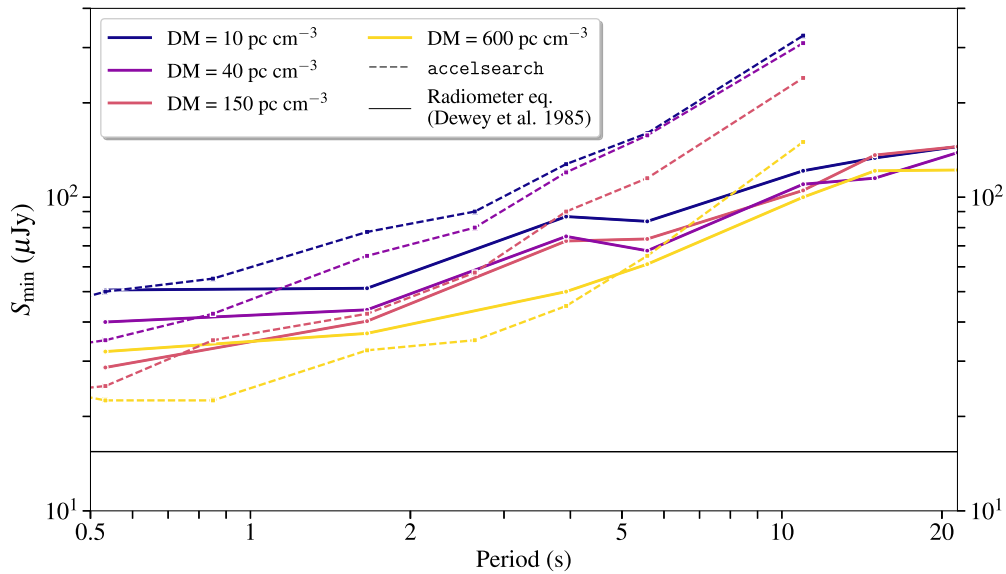


Figure 5. Minimum detectable mean flux density for the PALFA survey as measured when searching synthetic pulsar signals injected in real data with `ffaGo`. The signals have a fixed pulse FWHM of 2.6%. S_{\min} measured by the FFA are illustrated with solid lines, while the dashed lines represent values of S_{\min} obtained from the frequency-domain search reported in Lazarus et al. (2015). Note the greater parameter space covered at long periods in the FFA analysis.

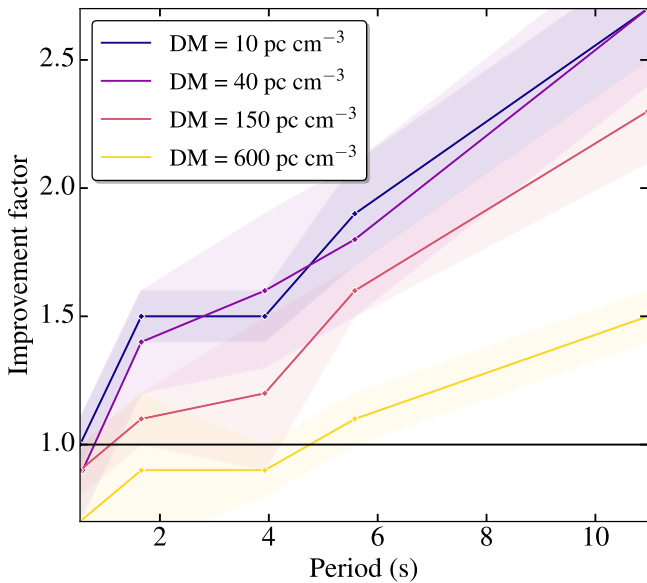


Figure 6. Improvement factor in the survey sensitivity as a function of period for the four DM trial values used, where the pulse FWHM = 2.6% is kept fixed. The factors were determined by dividing the median value of S_{\min} from Lazarus et al. (2015) for a given trial by the S_{\min} obtained from the `ffaGo`-based search for that trial. The shaded regions represent the uncertainty on the improvement factors, which were derived from the difference between the minimum detectable flux densities and the largest value of mean flux densities for which the trials were missed by the searches.

computation time for this specific Fourier-domain search. This remains a small fraction (less than 5%) of the overall processing time.

Also, we lowered the limit on the lowest frequency of the highest harmonic to search. This parameter was previously set to 1 and 2 Hz for accelerated and nonaccelerated searches, respectively, and is now reduced to 0.5 and 1 Hz, respectively. While this modification increases our sensitivity to pulsars having low modulation frequencies, it potentially increases the resulting number of false-positive candidates.

Following the recent implementation of the FFA and the above changes, the PALFA PRESTO-based pipeline has discovered five new sources³⁵ with periods longer than a few hundreds of milliseconds: PSR J1843+01 ($P = 1.267$ s), PSR J1911+13 ($P = 0.300$ s), PSR J1913+05 ($P = 0.662$ s), PSR J1914+08 ($P = 0.456$ s), and PSR J1924+19 (1.278 s). The FFA detected all five pulsars. Additionally, the Quicklook pipeline discovered J1901+11, a 0.409 s pulsar that was later redetected with both the FFA and the FFT searches of the full-resolution PALFA pipeline. Finally, `accelsearch` has found three pulsars with periods shorter than 100 ms, too short to be detected by `ffaGo`.

One pulsar, PSR J1913+05 ($P = 0.662$ s), was uniquely detected by the FFA search. It is interesting to note that it is not the slowest pulsar, but the weakest among the new discoveries. We estimate its flux density to be $11 \mu\text{Jy}$ with a pulse duty cycle δ of 2% when using the FWHM as the pulse width. This demonstrates that in addition to being more sensitive to long-period pulsars, `ffaGo` can outperform an FFT for weak pulsars with significant harmonic structures. This is consistent with our results from the simulations described in Section 5.2.

PSR J1911+13 has a rotation period of 0.300 s, was found at a DM of 322.3 pc cm^{-3} in inner Galaxy data, and appears to be a nulling pulsar. Both frequency-domain and time-domain analyses identified the source. Results from future timing observations of this nulling pulsar will be provided in a separate paper.

Pulsars with spin periods shorter than 0.5 s were detected by `ffaGo` via their second subharmonic, while it detected the ones with longer periods at their fundamental frequency. Timing solutions are not yet available for these new discoveries; precise parameters will be included in a future paper.

Thus far, our FFA pipeline detected more than 50 known sources having periods in the FFA range. There are a few

³⁵ Discovery plots of those sources are available at <http://www.naic.edu/~palfa/newpulsars/>.

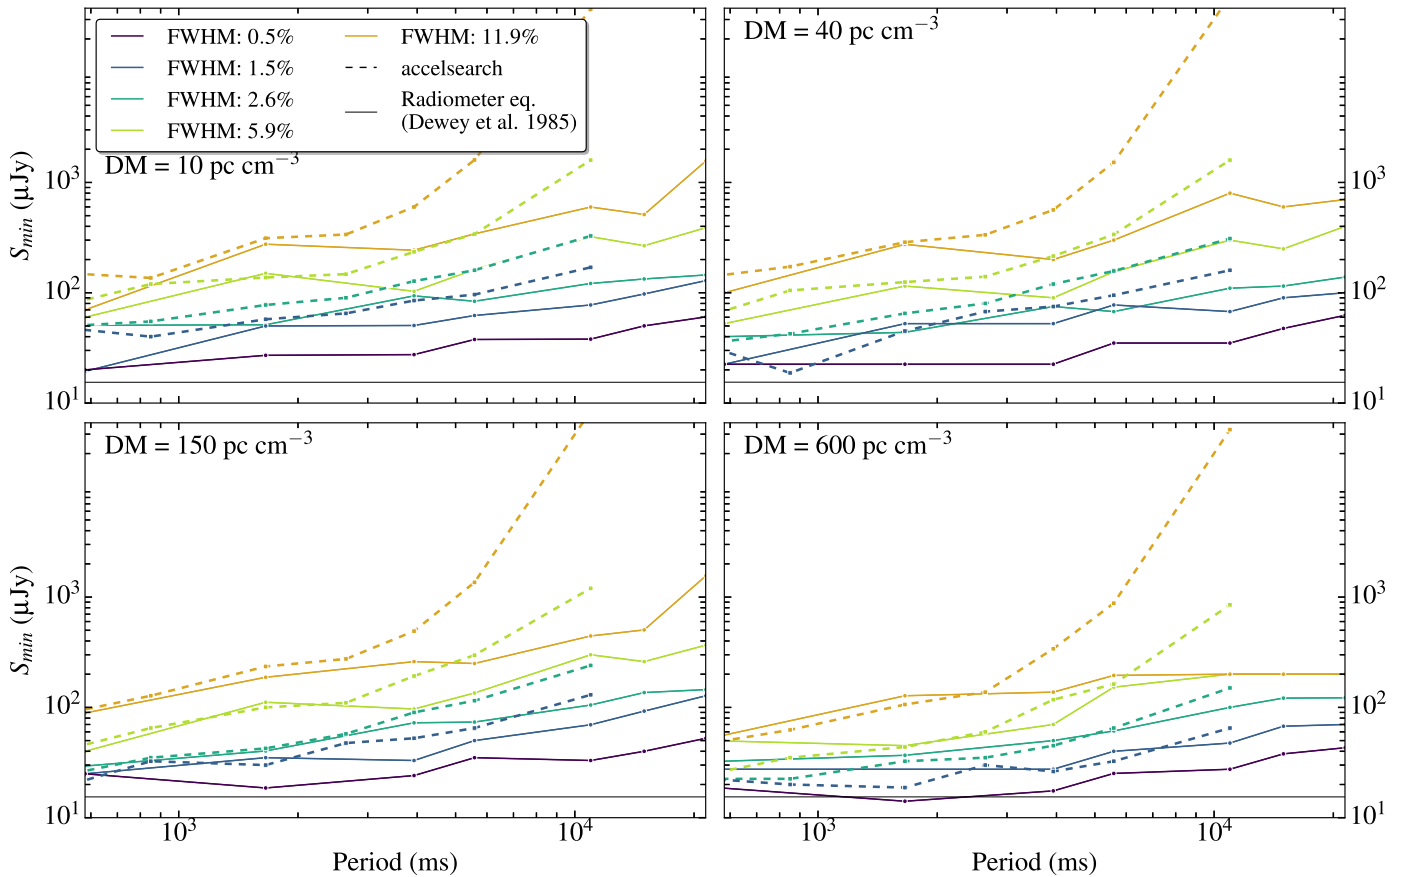


Figure 7. Minimum detectable mean flux density for the PALFA survey, as measured when searching synthetic pulsar signals injected in real data with `ffaGo`, for various pulse widths. Each panel corresponds to a different value of DM. Results from the FFA analysis presented in this work are illustrated with solid lines, while dashed lines represent the results reported in Lazarus et al. (2015). Note that the injections for the FFA analysis included narrow pulse profiles having an FWHM of 0.5% (purple lines) that were not included in Lazarus et al. (2015).

instances where the FFA detected the second subharmonics of pulsars. All of these have periods shorter than 500 ms. Moreover, two known sources were redetected by the FFA in beams a few arcminutes away from the true pulsar positions, but they were not detected by the FFT search. This is therefore promising and demonstrates once more the ability of `ffaGo` to detect pulsar signals in the survey.

7. Conclusion

In this paper, we have discussed the use of an FFA-based search, `ffaGo`, in the PALFA pulsar survey. In the PALFA implementation, `ffaGo` searches for periodic signals with $500 \text{ ms} \leq P \leq 30 \text{ s}$ in time series dedispersed at DM values under 3265 pc cm^{-3} .

We compared the FFA program with `PRESTO`'s frequency-domain periodicity search, `accelsearch`, using a constructed data set of synthetic pulsars having periods between 2 and 20 s, with pulse duty cycles δ ranging from 0.5% to 20%. Results showed that the FFA exceeds the performance of the FFT in the white-noise regime in the case of long-period pulsars, especially when the signals have low harmonic content. We then selected a variety of long-period pulsar observations with periods between 1.32 and 4.6 s discovered with PALFA and compared the response of both algorithms when searching in the presence of red noise and interference. The time-domain algorithm successfully detected all sources at

their fundamental frequencies, at S/N_{modified} values significantly larger than the detection threshold set in the pipeline.

The sensitivity of the PALFA survey was then assessed by conducting an analysis where we used `ffaGo` to recover a variety of synthetic pulsars injected into real PALFA survey observations. Comparing our results with those obtained by Lazarus et al. (2015), we showed that for a pulse width of 2.6%, the FFA outperforms the FFT for $DM \lesssim 40 \text{ pc cm}^{-3}$ and for periods as short as $\sim 500 \text{ ms}$, and the survey sensitivity is improved by at least a factor of two for periods $\gtrsim 6 \text{ s}$. For the same width, the FFA exceeds the performance of the FFT for all trial DMs for periods longer than 5 s. Moreover, for these periods, the sensitivity of the survey increases by at least a factor of three for pulsars having width $\gtrsim 11.9\%$ for all trial DMs. For periods greater than $\sim 8 \text{ s}$, the FFA performs better than the FFT for all tested values of DM and all pulse FWHMs. This simulation demonstrated that the coherent summing of all harmonics greatly enhances the sensitivity of pulsars survey in the red-noise regime.

As of now, our FFA search has uniquely discovered one pulsar, and four others were discovered by both the FFA and the FFT searches of the PALFA survey. It has also redetected more than 50 known pulsars that were present in the data. We are optimistic that our implementation of an FFA pipeline in the PALFA survey will lead to the discovery of new long-period pulsars in the future.

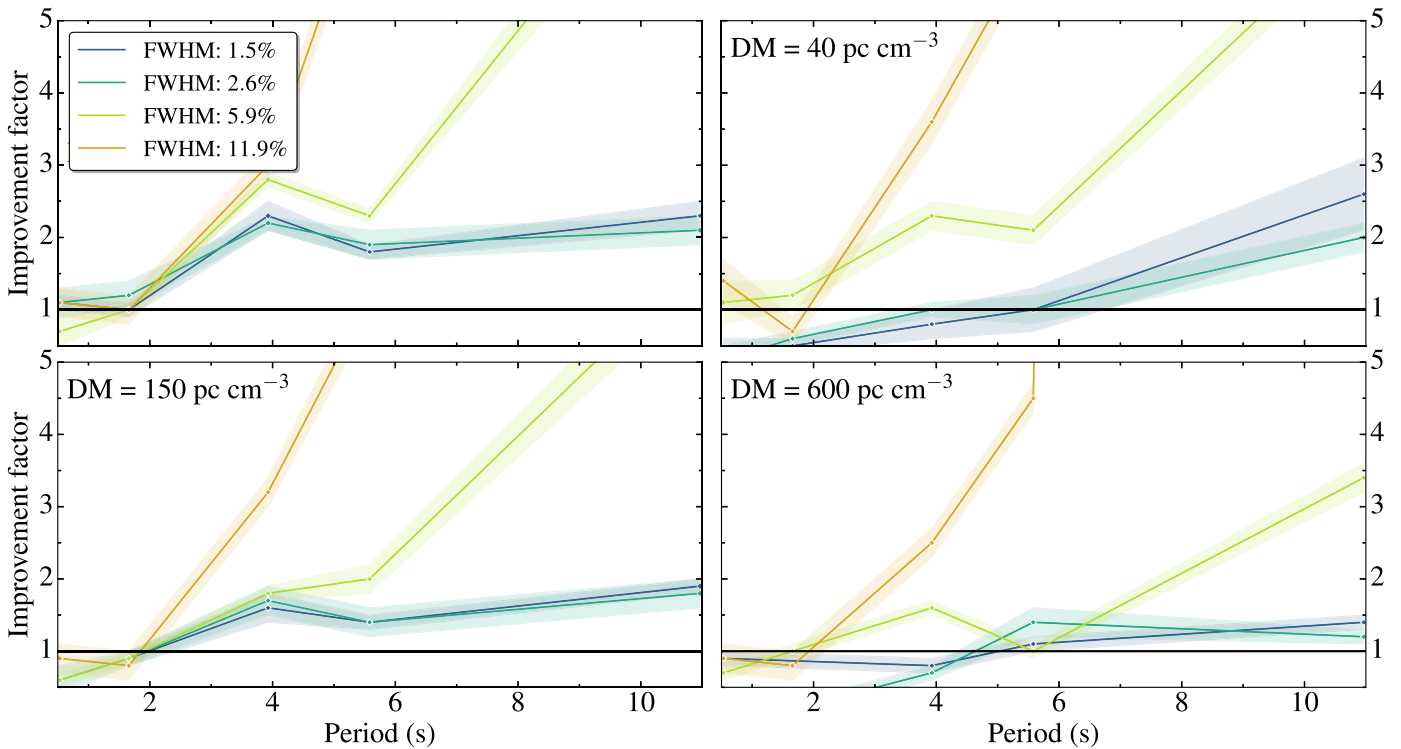


Figure 8. Improvement factor in the survey sensitivity as a function of period for the four DM trial values used and for various pulse FWHMs. The factors were determined by dividing the value of S_{\min} from Lazarus et al. (2015) for a given trial by the S_{\min} obtained from the $f\bar{f}aGo$ -based search for that trial. The shaded regions represent the uncertainty on the improvement factors, which were derived from the difference between the minimum detectable flux densities and the maximum value of mean flux densities for which the trials were missed by the searches.

We gratefully thank Vlad Kondratiev for useful discussions and his valuable assistance in the comparison of the different metrics used to evaluate FFA-generated profiles. E.P. acknowledges the support of NSERC (CGS M) and FQRNT B1. V.M.K. acknowledges support from an NSERC Discovery grant and Herzberg Award, the Canada Research Chairs program, the Canadian Institute for Advanced Research, and FRQ-NT. S.R. is a Senior Fellow in the Canadian Institute for Advanced Research. P.S. is a Covington Fellow at DRAO. M.A.M. is supported by NSF award no. 1458952. W.W.Z. is supported by the CAS Pioneer Hundred Talents Program. J.W.T.H. acknowledges funding from an NWO Vidi fellowship and from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013)/ERC Starting Grant agreement no. 337062 (“DRAGNET”). J.S.D. was supported by the NASA Fermi program. K.S., S.R., and F.C. are supported by the NANOGrav NSF Physics Frontiers Center award no. 1430284. Pulsar research at UBC is supported by an NSERC Discovery Grant and by the Canadian Institute for Advanced Research.

The Arecibo Observatory is operated by SRI International under a cooperative agreement with the National Science Foundation (AST-1100968), and in alliance with Ana G. Méndez-Universidad Metropolitana and the Universities Space Research Association. The CyberSKA project was funded by a CANARIE NEP-2 grant. Computations were made on the supercomputer Guillimin at McGill University, managed by Calcul Québec and Compute Canada. The operation of this supercomputer is funded by the Canada Foundation for Innovation (CFI), NanoQuébec, RMGA, and the Fonds de recherche du Québec-Nature et technologies (FRQ-NT).

Appendix A Comparison between Metric A and Metric B

The most important distinction between Metric A and Metric B is the calculation of a profile’s standard deviation. For Metric A, this value is constant for a given trial period (i.e., the M folded profiles for a given trial period will have the same value of standard deviation), while it is profile dependent for S/N calculations with Metric B. Analytically, regardless of the data distribution, a consequence of calculating the standard deviation directly on the profiles is that there is a greater spread in the standard deviation, especially significant in the case where fewer padded samples are folded. Excluding a 20% window around the profile peak in the calculations further decreases the number of bins within a profile, resulting in even larger standard deviation values compared to Metric A. However, when using real, non-Gaussian-distributed data, removing some of the high-valued bins will reduce the sum of the squared deviations, therefore decreasing the standard deviation. It is therefore nontrivial to predict the expected reduction in S/N (or increase in standard deviation) from Metric B as M , the number of folded profiles, and z , the number of padded profiles, change. The simulations presented in Section 3.3 show that the S/N values from Metric B are the lowest for broad, very long period pulsars. This suggests that the effective increase in the standard deviation in Metric B, which is more significant at longer periods (i.e., for smaller values of z), affects more the S/N calculations than the

decrease in standard deviation resulting from the exclusion of the profile peak.

Appendix B Profile Evaluation with Metric C

To evaluate FFA-generated profiles, we designed a third profile significance metric in `ffaGo`, Metric C. Similarly to Metric B, Metric C excludes a 20% window centered on the peak of the profile when calculating the median intensity of the off-pulse, $I_{\text{med,off}}$. This median value is therefore the same as for Metric B. However, the standard deviation of the off-pulse profile is a scaled-down version of the standard deviation computed with Metric A. The mathematical expression for the S/N calculated with this third metric is as follows:

$$S/N = \frac{I_{\text{max}} - I_{\text{med,off}}}{(\sigma\sqrt{X})\sqrt{0.8}(M - z)}, \quad (6)$$

where the factor $\sqrt{0.8}$ accounts for the exclusion of the on-pulse portion of the profile.

We investigated the performance of Metric C when searching for pulsar signals in white noise and compared it to Metric A. The same five data sets of 120 synthetic pulsars described in Section 3.3 were processed through `ffaGo` with Metric C, and we illustrate the results in Figure 9. This metric's response pattern is very similar to that of Metric A. Slightly larger values of S/N_{modified} are obtained with Metric C when recovering narrow-pulsed pulsars, while S/N_{modified} values are slightly lower when recovering broad, long-period signals. The mean value of the ratio matrix shown in the bottom panel of Figure 9 is 1.01 ± 0.10 , which implies that there is no significant difference in the performance of the two metrics when considering the phase space that our search is covering.

It is interesting to note that Metric C does not suffer the same sensitivity degradation in the long-period/large- δ phase space as Metric B. One can then conclude that it is not the exclusion of the on-pulse fraction that causes the most significant differences between Metric A and Metric B, but rather the calculation of the standard deviation directly on profiles rather than on the time series directly.

The performance of Metric C was also tested when searching for pulsars in the presence of red noise by performing two distinct analyses. First, we processed the selection of 12 real PALFA long-period pulsar observations (see Table 1) through `ffaGo` while using Metric C. Metric A produced marginally larger values of S/N_{modified} than Metric C for all pulsars. While Metric A identified the fundamental spin frequencies of all pulsars, there are two instances where Metric C identified the pulsar via its harmonics.

The second analysis consisted of the partial reproduction of the sensitivity analysis described in Section 5. Synthetic pulsars for a subset of the parameters listed in Table 2 (only the trial $\text{DM} = 150 \text{ pc cm}^{-3}$ was selected) were injected in one of the PALFA observation files used in the PALFA sensitivity analysis. The minimum mean flux density detectable by Metric C was then established. The results suggest that the performances of the two metrics are similar, regardless of the spin period of the pulsar.

We conclude that in the presence of RFI and red noise, Metrics A and C behave similarly.

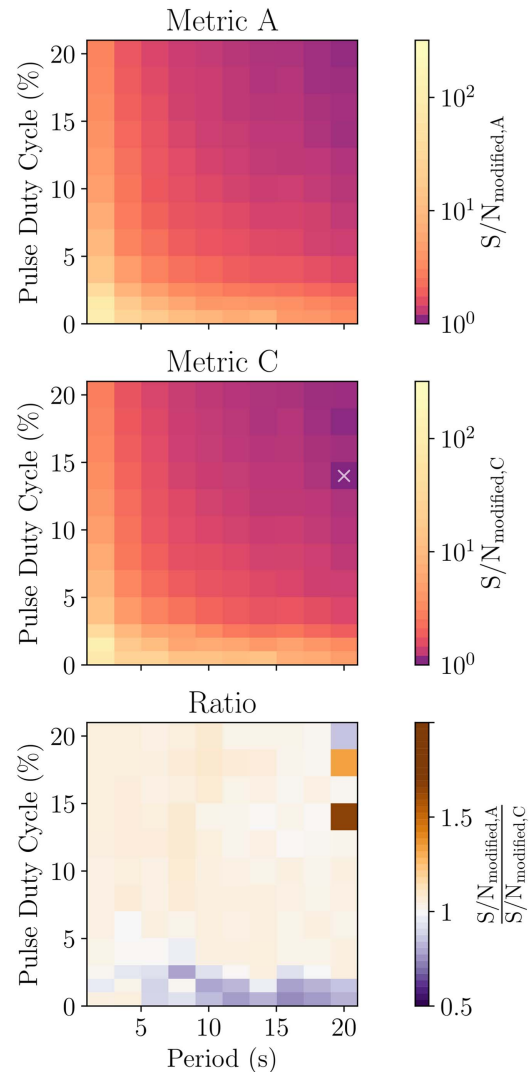


Figure 9. Response patterns of the two FFA significance metrics, Metric A (top panel) and Metric C (middle panel), obtained from the white-noise simulation described in Section 3.3. The ratio of the two values of S/N_{modified} is shown in the bottom panel, which illustrates the relative responses of Metrics A and C. The values reported are the average S/N_{modified} from the five simulations. The pixel with a white cross represents the trial having an average S/N_{modified} below 6 (i.e., classified as a nondetection). Note that the scales for the top and the middle panels are logarithmic, while the bottom panel is displayed on a linear scale.

ORCID iDs

- E. Parent <https://orcid.org/0000-0002-0430-6504>
V. M. Kaspi <https://orcid.org/0000-0001-9345-0307>
S. M. Ransom <https://orcid.org/0000-0001-5799-9714>
P. Scholz <https://orcid.org/0000-0002-7374-7119>
M. A. McLaughlin <https://orcid.org/0000-0001-7697-7422>
W. W. Zhu <https://orcid.org/0000-0001-5105-4058>
Z. Pleunis <https://orcid.org/0000-0002-4795-697X>
B. Allen <https://orcid.org/0000-0003-4285-6256>
S. Bogdanov <https://orcid.org/0000-0002-9870-2742>
F. Camilo <https://orcid.org/0000-0002-1873-3718>
S. Chatterjee <https://orcid.org/0000-0002-2878-1502>
J. M. Cordes <https://orcid.org/0000-0002-4049-1882>
F. Crawford <https://orcid.org/0000-0002-2578-0360>
J. S. Deneva <https://orcid.org/0000-0003-1226-0793>
R. Ferdman <https://orcid.org/0000-0002-2223-1235>

P. C. C. Freire  <https://orcid.org/0000-0003-1307-9435>
 J. W. T. Hessels  <https://orcid.org/0000-0003-2317-1446>
 R. Lynch  <https://orcid.org/0000-0001-5229-7430>
 I. H. Stairs  <https://orcid.org/0000-0001-9784-8670>
 K. Stovall  <https://orcid.org/0000-0002-7261-594X>
 J. Swiggum  <https://orcid.org/0000-0002-1075-3837>

References

- Allen, B., Knispel, B., Cordes, J. M., et al. 2013, *ApJ*, **773**, 91
 Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, *ApJ*, **605**, 759
 Cameron, A. D., Barr, E. D., Champion, D. J., Kramer, M., & Zhu, W. W. 2017, *MNRAS*, **468**, 1994
 Chen, K., & Ruderman, M. 1993, *ApJ*, **402**, 264
 Cordes, J. M., & Lazio, T. J. W. 2002, arXiv:astro-ph/0207156
 Cordes, J. M., & McLaughlin, M. A. 2003, *ApJ*, **596**, 1142
 Craft, H. D., Sutton, J. M., & Comella, J. M. 1968, *Natur*, **219**, 1237
 Crawford, F., Lorimer, D. R., Devour, B. M., Takacs, B. P., & Kondratiev, V. I. 2009, *ApJ*, **696**, 574
 Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, *ApJL*, **294**, L25
 Eatough, R. 2007, in *From Planets to Dark Energy: The Modern Radio Universe* (Trieste: SISSA), 92
 Faulkner, A. J., Stairs, I. H., Kramer, M., et al. 2004, *MNRAS*, **355**, 147
 Hibschan, J. A., & Arons, J. 2001, *ApJ*, **554**, 624
 Kaspi, V. M., & Beloborodov, A. M. 2017, *ARA&A*, **55**, 261
 Kiddle, C., Andrecut, M., Brazier, A., et al. 2011, in *ASP Conf. Ser. 442, Astronomical Data Analysis Software and Systems XX*, ed. I. N. Evans et al. (San Francisco, CA: ASP), 669
 Kondratiev, V. I., McLaughlin, M. A., Lorimer, D. R., et al. 2009, *ApJ*, **702**, 692
 Lazarus, P., Brazier, A., Hessels, J. W. T., et al. 2015, *ApJ*, **812**, 81
 Lipunov, V. M., Bogomazov, A. I., & Abubekkerov, M. K. 2005, *MNRAS*, **359**, 1517
 Lorimer, D. R., Faulkner, A. J., Lyne, A. G., et al. 2006, *MNRAS*, **372**, 777
 Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy* (Cambridge: Cambridge Univ. Press)
 Lovelace, R. V. E., Sutton, J. M., & Salpeter, E. E. 1969, *Natur*, **222**, 231
 Lyne, A. G., Stappers, B. W., Bogdanov, S., et al. 2017, *ApJ*, **834**, 137
 Patel, C. 2016, Master's thesis, McGill Univ. Montreal
 Petigura, E. A., Marcy, G. W., & Howard, A. W. 2013, *ApJ*, **770**, 69
 Pfahl, E., Podsiadlowski, P., & Rappaport, S. 2005, *ApJ*, **628**, 343
 Ransom, S. M. 2001, PhD thesis, Harvard Univ.
 Staelin, D. H. 1969, *IEEEP*, **57**, 724
 Stovall, K., Jenet, F., Siemens, X., et al. 2013, AAS Meeting, **221**, 154.05
 Swiggum, J. K., Lorimer, D. R., McLaughlin, M. A., et al. 2014, *ApJ*, **787**, 137
 von Mises, R. 1918, *PhyZ*, **19**, 490
 Young, M. D., Manchester, R. N., & Johnston, S. 1999, *Natur*, **400**, 848
 Zhang, B., Harding, A. K., & Muslimov, A. G. 2000, AAS Meeting, **32**, 880
 Zhu, W. W., Berndsen, A., Madsen, E. C., et al. 2014, *ApJ*, **781**, 117