

A Blind-search Survey for Gamma-ray Pulsars

Von der Fakultät für Mathematik und Physik
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades

DOKTOR DER NATURWISSENSCHAFTEN
Dr. rer. nat.

genehmigte Dissertation
von

Colin James Clark M.Sci.

geboren am 16.12.1990 in Paisley, UK

2017

Referent:

Dr. rer. nat. Holger J. Pletsch
Albert-Einstein-Institut
Leibniz Universität Hannover

Korreferenten:

Prof. Dr. Bruce Allen
Albert-Einstein-Institut
University of Wisconsin-Milwaukee
Leibniz Universität Hannover

Dr. David A. Smith
Centre d'Études Nucléaires de Bordeaux Gradignan, IN2P3/CNRS
Université Bordeaux

Tag der Promotion:

05.08.2016

Abstract

It has long been known that young pulsars can emit strongly in gamma rays. However, it is only recently, using observations by the Large Area Telescope (LAT) on-board the *Fermi Gamma-ray Space Telescope* which launched in 2008, that it has been possible to study a large number of pulsars, including millisecond pulsars, through their gamma-ray pulsations. Many of the pulsars discovered by the LAT have not been detected in subsequent radio searches. Blind searches, which directly search for pulsations at unknown frequencies in gamma-ray data, are necessary to find such “radio-quiet” pulsars, and to provide a more complete view of the local neutron star population.

This thesis is a collection of works describing the development of and results from the latest blind gamma-ray pulsar searches.

Firstly, techniques were developed to enhance the sensitivity of existing search methods. These include a revised parameter-space metric approximation which enables more efficient search grids to be constructed; interpolation schemes to mitigate potential signal losses; and improved automatic follow-up and refinement stages. By simulating and searching for thousands of weak pulsar signals, these improved methods were shown to result in searches that were almost 50% more sensitive than those performed with previous methods, without requiring additional computational power.

A blind-search survey was performed on the distributed volunteer computing system *Einstein@Home* using these new techniques to search for pulsations in unidentified LAT sources. This resulted in the discovery of nineteen new isolated gamma-ray pulsars; almost one-third of all pulsars that have been discovered through their gamma-ray pulsations. Using the improved understanding of the survey’s sensitivity to pulsations, upper limits were estimated for the fraction of pulsed gamma-ray flux from sources which remain unidentified. The survey was estimated to be sensitive to pulsations from unidentified LAT sources with a point-source significance above 10σ . The later detection of a faint pulsar, PSR J1817–1742, with a similar point-source significance lends credence to this estimated threshold.

One of the first detections made by this survey, of the young, glitching PSR J1906+0722, explained the nature of a highly significant unidentified gamma-ray source, and revealed additional gamma-ray emission, likely from a nearby unrelated supernova remnant. Another pulsar detected by this survey, PSR J1208–6238, is the youngest pulsar to be found through its gamma-ray pulsations, with an estimated age of just 2,700 years. Its youth, and high magnetic field strength allowed its braking index to be measured; a first for a radio-quiet gamma-ray pulsar.

This survey also discovered two isolated gamma-ray millisecond pulsars, both of which have remained undetected in radio observations despite extensive searches. These pulsars could possibly be the first two members of a long-awaited pulsar class: radio-quiet millisecond pulsars.

Keywords: pulsars, gamma rays, astronomy

Kurzfassung

Es ist seit langem bekannt, dass junge Pulsare starke Gamma-Strahler sein können. Aber erst dank der Messdaten vom Large Area Telescope (LAT) an Bord des *Fermi Gamma-ray Space Telescope*, welches im Jahr 2008 gestartet wurde, ist es möglich geworden eine große Anzahl solcher Pulsare, einschließlich Millisekunden-Pulsare, anhand ihrer Gammastrahlung zu studieren. Viele der vom LAT entdeckten Pulsare konnten in darauf folgenden Suchen im Radiospektrum nicht gefunden werden. „Blindsuchen“, bei denen nach Pulsationen mit unbekannter Frequenz direkt in den Gammastrahlungsdaten gesucht wird, sind notwendig zum Auffinden solcher „Radio-stillen“ Pulsare, um damit ein vollständigeres Bild der lokalen Neutronenstern-Population zu ergeben.

Diese Dissertation ist eine Sammlung von Arbeiten, die die Entwicklung und die Ergebnisse der neuesten Blindsuchen nach Gamma-Pulsaren beschreiben.

Zuerst wurden Verfahren entwickelt, um die Sensitivität bestehender Suchmethoden zu verbessern: eine überarbeitete Approximation der Metrik über den Parameterraum zur Konstruktion effizienterer Suchgitter; Interpolationsmethoden zur Verringerung möglicher Signalverluste; verbesserte automatische Verfeinerungs- und Nachfolgeuntersuchungen. Durch die Simulation und Suche von tausenden schwachen Pulsarsignalen konnte gezeigt werden, dass diese Verbesserungen, verglichen mit den bisherigen Methoden, die Empfindlichkeit um etwa 50% erhöhen, ohne dabei mehr Rechenleistung zu benötigen.

Eine Blindsuche mit diesen neuen Verfahren wurde auf dem verteilten Rechenprojekt *Einstein@Home* durchgeführt, bei der nach Pulsationen in bisher unidentifizierten LAT-Quellen gesucht wurde. Dies führte zur Entdeckung von 19 neuen isolierten Gamma-Pulsaren und entspricht fast einem Drittel aller bisher durch Gammastrahlen entdeckten Pulsare. Für weiterhin unidentifizierte Quellen war es dank der verbesserten Charakterisierung der Empfindlichkeit dieser Suche möglich, Obergrenzen für die pulsierenden Anteile der Gammastrahlung zu setzen. Es wurde abgeschätzt, dass diese Suche sensitiv für Pulsationen von unidentifizierten LAT-Quellen ist, die eine Punktquellen-Signifikanz über 10σ besitzen. Die spätere Entdeckung des schwachen Pulsars, PSR J1817–1742, mit einer ähnlichen Punktquellen-Signifikanz verleiht dieser geschätzten Grenze Glaubwürdigkeit.

Eine der ersten Entdeckungen dieser Suche, der junge und glitchende Pulsar PSR J1906+0722, war eine sehr signifikante aber unidentifizierte Gammaquelle und deutet auf zusätzliche Gammastrahlung von einem vermutlichen nahe gelegenen Überrest einer Supernova hin. Ein weiterer in dieser Suche entdeckter Pulsar, PSR J1208-6238, ist der jüngste durch Gammastrahlung entdeckte Pulsar mit einem geschätzten Alter von nur 2 700 Jahren. Durch sein geringes Alter und sein starkes Magnetfeld war es möglich, zum ersten Mal den Brems-Index für einen Radio-stillen Gamma-Pulsar zu bestimmen.

Diese Suche fand weiterhin zwei isolierte Millisekunden-Pulsare. Beide konnten trotz intensiver Suchen im Radio-Spektrum nicht nachgewiesen werden, was sie zu den ersten möglichen Mitgliedern einer lang erwarteten Pulsar-Klasse macht: Radio-stille Millisekunden-Pulsare.

Schlagworte: Pulsare, Gammastrahlung, Astronomie

Contents

1	Introduction	1
1.1	Pulsars	2
1.2	Gamma-ray Pulsars	3
1.2.1	Pulsars Detected with Previous Gamma-ray Telescopes	3
1.2.2	Pulsar Observations with Current Gamma-ray Telescopes	4
1.3	Blind Searches for Gamma-ray Pulsars	7
1.4	Timing Gamma-ray Pulsars	11
1.5	Chapter Descriptions and Authorship Clarifications	13
1.5.1	Description of Chapter 2	13
1.5.2	Description of Chapter 3	14
1.5.3	Description of Chapter 4	14
1.5.4	Description of Chapter 5	15
1.5.5	Description of Chapter 6	15
1.5.6	Description of Chapter 7	16
2	Optimized Blind Gamma-ray Pulsar Searches at Fixed Computing Budget	17
2.1	Introduction	18
2.2	Statistical Detection of Pulsations	18
2.3	Coherent Test Statistics	20
2.3.1	Statistical Properties	21
2.3.2	Effects of Pulse Profile on Sensitivity	23
2.3.3	Grid-point Counting for Coherent Search	27
2.3.4	Coherent Search Sensitivity at Fixed Computing Cost	30
2.3.5	Efficiency of Harmonic Summing at Fixed Computing Cost	31
2.4	Semicoherent Test Statistics	33
2.4.1	Statistical Properties	35
2.4.2	Grid-point Counting for Semicoherent Search	38
2.4.3	Semicoherent Search Sensitivity at Fixed Computing Cost	39
2.5	Efficient Implementation of a Multistage Search Scheme	41
2.5.1	Efficient Computation of Semicoherent Test Statistic	43
2.5.2	Frequency Domain Interpolation	45
2.5.3	Complex Heterodyning	48
2.5.4	Lag Domain Interpolation	48
2.6	Performance Demonstration	51
2.7	Conclusions	52
2.8	Appendix to Chapter 2	54
2.8.1	Derivation of statistical properties of coherent test statistic	54

2.8.2	Coherent Metric	55
2.8.3	Coherent Metric with Incoherent Harmonic Summing	58
2.8.4	Approximate Harmonic-summing Computing Cost	60
2.8.5	Optimal Mismatch in Coherent Search	61
2.8.6	Derivation of Statistical Properties of Semicoherent Test Statistic	62
2.8.7	Semicoherent Metric	64
2.8.8	Optimal Mismatch in Semicoherent Search	66
2.8.9	Sky-grid Construction	66
3	PSR J1906+0722: An Elusive Gamma-Ray Pulsar	69
3.1	Introduction	70
3.2	Discovery	70
3.2.1	Data Preparation	70
3.2.2	Blind Search Method	71
3.3	Follow-up Analysis	71
3.3.1	Glitch Identification	73
3.3.2	Timing Analysis	73
3.3.3	Off-pulse Analysis	76
3.4	Analysis in Other Wavelengths	78
3.4.1	Radio and X-ray Observations	78
3.4.2	Possible SNR Associations	78
3.5	Discussion	80
4	The Einstein@Home Gamma-ray Pulsar Survey: Paper I	83
4.1	Introduction	84
4.2	Search Scheme	84
4.2.1	Data	84
4.2.2	Parameter Space	85
4.2.3	Detection Statistics	86
4.2.4	Initial Search Stage	89
4.2.5	Follow-up and Refinement Stages	91
4.3	Sensitivity	91
4.3.1	Pulsed Fraction Upper Limit for W49B	98
4.3.2	Pulsed Fraction Upper Limit for Cassiopeia A	98
4.3.3	Pulsed Fraction Upper Limit for the Galactic Center	99
4.4	Timing Analysis	99
4.4.1	Timing Methods	100
4.4.2	Spin-down vs. Timing Noise	101
4.4.3	Timing Large Glitches	104
4.4.4	Pulse Profile Variability	106
4.5	Discussion	109

4.5.1	Sensitivity	109
4.5.2	New Pulsars	114
4.6	Conclusions	114
4.7	Appendix to Chapter 4	116
4.7.1	Candidate Ranking	116
4.7.2	Distribution of S_1 with a Signal	118
4.7.3	Efficient Sampling of Glitch Parameters	119
5	The Braking Index of a Radio-quiet Gamma-ray Pulsar	121
5.1	Introduction	122
5.2	Observations	123
5.2.1	LAT Data	123
5.2.2	Discovery in a Blind Search	123
5.2.3	Off-pulse Analysis	124
5.2.4	Radio Observations	124
5.3	Timing Analysis	126
5.4	Discussion	127
5.5	Additional Material: Alternative Spin-down Models	131
6	The Discovery of Two Possibly Radio-Quiet Gamma-ray Millisecond Pulsars	135
6.1	Introduction	136
6.2	MSP Discoveries by the <i>Einstein@Home</i> Search	137
6.3	Pulsar Properties	137
6.4	Radio Observations	140
6.5	Pulse Profile Modeling	142
6.6	Off-pulse and Spectral Analysis	143
6.7	X-ray Observations	145
6.8	Implications	146
7	Latest Results from the <i>Einstein@Home</i> Survey	149
7.1	PSR J1641–5317	149
7.2	PSR J1817–1742	152
7.3	Future Gamma-ray Pulsar Surveys	154
	Bibliography	157
	Acknowledgments	171
	Curriculum Vitae	173
	Publication List	175

Chapter 1

Introduction

This thesis concerns the detection and analysis of pulsed signals in the gamma-ray emission from pulsars, the ultra-dense rapidly-spinning compact remnants that result from the deaths of massive stars. The vast majority of the large population of known pulsars have been discovered through their radio emission, however a number of gamma-ray pulsars remain undetected in radio searches. The ability to detect and study pulsars via their gamma-ray pulsations alone is therefore necessary to fully explore the local pulsar population.

This chapter will briefly introduce pulsars (Section 1.1); observations of gamma-ray pulsars in particular (Section 1.2); the history of “blind” searches for gamma-ray pulsars (Section 1.3); and the methods used to time gamma-ray pulsars (Section 1.4).

The next five chapters following this introduction consist of individual papers, either written or co-written by the author of this thesis, which have either been published or are soon to be submitted to scientific journals. These chapters are adaptations of those papers, with minor changes to correct typographical errors found since their publication, for formatting reasons, or to provide references between them. Chapters 4 and 5 are earlier versions of papers from the time this thesis was initially submitted; the published versions contain some minor updates. Chapter 6 describes preliminary analysis from the time this thesis was initially submitted; a full article including updated analyses and additional observations taken since then will be submitted for publication soon. The introductions of some papers have been shortened to avoid overlap with this introductory chapter and with each other. The last part of the introductory chapter (Section 1.5) describes these papers in more detail, and outlines the author’s specific contributions to each.

Chapter 2 investigates the methods used to perform “blind” searches for gamma-ray pulsations, and describes the developments made to these methods during the doctoral studies presented here. These methods were implemented and used to perform a large-scale survey for pulsations in gamma-ray photons detected by the *Fermi* Large Area Telescope from unidentified gamma-ray sources. The survey ran on the distributed volunteer computing system, *Einstein@Home*. The four succeeding chapters (Chapters 3, 4, 5 and 6) describe the results of this survey, and investigations of the new pulsars discovered during it. Finally, in Chapter 7, the most recent discoveries from the *Einstein@Home* survey are presented for the first time.

1.1 Pulsars

Neutron stars, first proposed by Baade & Zwicky [1] to be the ultra-dense compact remnants of the collapsed cores of massive stars, are born rapidly spinning, with periods of just a few milliseconds [2, 3], and with extremely high magnetic fields [4–8]. In their essential textbook, Lorimer & Kramer [9, pp. 63, 80] give the following simplified picture of how these properties result in neutron stars being visible as “Pulsating Sources of Radio” or “pulsars”. A neutron star’s intense magnetic field and rapid rotation result in immense Lorentz forces which can strip charged particles from their surface and produce a dense plasma which fills the “magnetosphere” [10]. The plasma is confined by the dipolar magnetic field and forced to co-rotate with the neutron star out to the “light-cylinder” radius ($R_{LC} = c/2\pi f$, at which the plasma would have to travel faster than the speed of light to corotate with a neutron star spinning at a frequency f). Particles accelerated along curved field lines emit gamma rays, which in turn produce electron–positron pairs as they interact with the intense magnetic field near the magnetic poles. These pairs then emit further gamma rays, which themselves produce more electron–positron pairs, causing a “pair cascade” [11]. This results in bunches of charged particles that emit coherent radio emission in sharp, conical beams around each magnetic pole. These beams corotate with the neutron star and, if they sweep over the line-of-sight to the neutron star, we can observe them as regular pulses of radio emission at the neutron star’s rotation frequency.

Since the first serendipitous discovery of pulsed radio emission in 1967 by Bell & Hewish [12], and the rapid identification of its neutron star origin [13], more than 2,500 pulsars [14] have been detected, revealing a wealth of variety in the properties of neutron stars. Distinct pulsar populations have emerged: “young” pulsars [15], which spin down rapidly and exhibit complex instabilities in their rotation rates [16]; “magnetars” [17], whose X-ray emission is powered not by their rotation, but by their intense but decaying magnetic fields; and older, rapidly spinning “recycled” or “millisecond” pulsars [MSPs; 18], the eventual result of the “spinning-up” of accretion-powered binary X-ray pulsars as they gain angular momentum from their companion stars [19]. Each of these groups are extensively studied, but many of their properties are not yet fully understood [20].

Measurements of pulsars’ rotations (or “timing”) spanning several years, have revealed many important scientific results [21]. Timing measurements of pulsars in binary systems, the first detected by Hulse & Taylor in 1975 [22], provide the best constraints on neutron star masses [e.g. 23, 24], and allow for precise tests of relativistic effects [e.g. 25–27]. The largest known pulsar masses place constraints on the equation of state of matter above nuclear densities [28], which is an important unsolved problem for nuclear physics [e.g. 29] and determines the properties of supernovae explosions [30] and the gravitational emission from black hole–neutron star/neutron star–neutron star coalescences [e.g. 31].

Measurements of “glitches” [e.g. 32–34], sudden changes in a young pulsar’s rotation rate, indicate a layered neutron star structure, in which a superfluid core rotates semi-independently from a solid crust, but with occasional angular momentum transfers between the two [e.g. 35]; or alternatively “starquakes”, sudden reconfigurations of the neutron star’s crust as it loses oblateness due to its spin-down [e.g. 36, 37].

In contrast, the stability of the rotation of some MSPs can rival that of atomic clocks [38], a property that makes them valuable astrophysical tools. One of the major current goals of pulsar astronomy is to use a “timing array” of MSPs to detect the tiny ripples in space-time of passing gravitational waves [e.g. 39–43]. The recent detection of a MSP in a triple system [44] allows for extremely precise tests of the strong equivalence principle [45]. Another example of the scientific potential of MSP timing was the first detection of extra-solar planets, which were found in 1992 by Wolszczan & Frail [46] by measuring periodic variations in the arrival times of an MSP’s pulsations.

The ability to detect and study new pulsars is therefore of importance for understanding the evolutionary paths of the Galactic stellar population; studying the nature of matter under the most extreme conditions; and for testing fundamental physical laws.

1.2 Gamma-ray Pulsars

Charged particles accelerated in “gaps” in pulsars’ magnetospheres can reach GeV energies [9, pp. 79–80], emitting intense co-rotating beams of gamma-ray radiation. Much like the observed radio pulses for which pulsars are named, these beams are observable as pulses of gamma-ray flux. The ability to detect and study pulsars through their gamma-ray emission offers a complementary window with which to observe the Galactic neutron star population. Much of the following two sections, which briefly outline the history of observations of gamma-ray pulsars and recent progress in the field, follows the fascinating account given by Caraveo [47].

1.2.1 Pulsars Detected with Previous Gamma-ray Telescopes

The first tentative detections of pulsed gamma-rays were made in the 1970s by balloon-borne gas-Cerenkov gamma-ray telescopes [e.g. 48], and were confirmed shortly after by the SAS-2 and COS-B satellites [49, 50]. These space-based telescopes featured spark chambers surrounded by “anti-coincidence” guards to discern between gamma-rays and charged cosmic rays [51]. They were sensitive to gamma rays from ~ 20 MeV to a few GeV, and could localize point sources to within a few degrees [52]. Their detections of gamma-ray pulsations from the Crab and Vela pulsars [53, 54] showed that pulsars’ spin-down-power-to-gamma-ray efficiencies dwarf those in radio or X-ray wavelengths [55]. These observations also revealed a striking difference between these two pulsars: while the Crab’s radio, optical and gamma-ray pulses appeared almost identical, the Vela pulsar’s doubly peaked optical and gamma-ray pulsations were not phase-aligned with each other, nor with its single radio peak [54, 56].

The EGRET [57] (the next generation of spark-chamber based gamma-ray telescope, with an effective collecting area of 1500 cm^2 , many times larger than that of SAS-2 or COS-B) and COMPTEL [58] (sensitive to lower-energy photons between 1 and 30 MeV) instruments on board the CGRO satellite [59] detected significant gamma-ray pulsations from a further five pulsars [60] (with inconclusive detections of three more [61–63]). The differences in (rotational) phase and shape between gamma-ray and radio pulses revealed by CGRO favored the “slot gap” [64–66] and “outer gap” [67, 68] models for gamma-ray emission, in which particle acceleration occurs in

narrow vacuum regions far from the neutron star’s surface, over models in which gamma rays were produced near the polar cap [69].

1.2.2 Pulsar Observations with Current Gamma-ray Telescopes

As their gas ages, the performance of spark chambers deteriorates over time, limiting the life-span of a gamma-ray telescope, and requiring gas refills which add to the weight and cost of the mission. Solid-state silicon detectors do not suffer these problems, and offer key advantages including: reduced “dead time” after photon detection allowing for better measurements of short duration transients, for example gamma-ray bursts which often last for less than a second; improved track resolution allowing for better direction reconstruction; and improved efficiency at high energies [70].

The present generation of space-based gamma-ray telescopes based around solid-state silicon trackers began with the launch of the AGILE satellite in 2007 [71]. By folding AGILE data using radio ephemerides, weak gamma-ray pulsations were detected from a handful of radio pulsars [72, 73].

The Fermi Large Area Telescope

The launch of the *Fermi Gamma-ray Space Telescope* in 2008 brought with it a new era for pulsar astronomy, in which large numbers of pulsars have become detectable through their gamma-ray emission.

Fermi’s primary on-board instrument is the Large Area Telescope [LAT; 74, 75]. The LAT is composed of a precision converter–tracker (interleaved tungsten converter layers and silicon-strip detectors), a caesium-iodide calorimeter and a segmented anti-coincidence detector. The LAT’s effective collecting area, field-of-view and energy range are all greatly improved over EGRET’s, making it the most sensitive gamma-ray telescope ever to have flown [47]. *Fermi* normally operates in a “surveying” mode, pointing alternately up and down on consecutive orbits to cover the entire sky every 190 minutes or so. In combination with the LAT’s impressive angular resolution (enabling the localization of gamma-ray sources to within a few arcminutes), this has provided a detailed and continuous view of the entire sky, and revealed thousands of new sources of gamma-ray emission (as shown in Figure 1.1). Furthermore, as the mission continues, the understanding of the instrument improves [76]; the recent “Pass 8” data reprocessing [77] has greatly increased the LAT’s energy range, effective collecting area and resolving power.

Crucially, for the field of pulsar astronomy at least, the LAT’s sensitivity and resolving power is optimized for GeV energies, where many pulsars’ spectra peak, and where different models for pulsars’ gamma-ray emission predict significant differences [78]. Additionally, the photon arrival times measured by the LAT are accurate to less than $1 \mu\text{s}$ [75], a resolution that is essential for timing pulses from MSPs.

The LAT’s unprecedented sensitivity has enabled the detection of pulsed gamma-ray emission

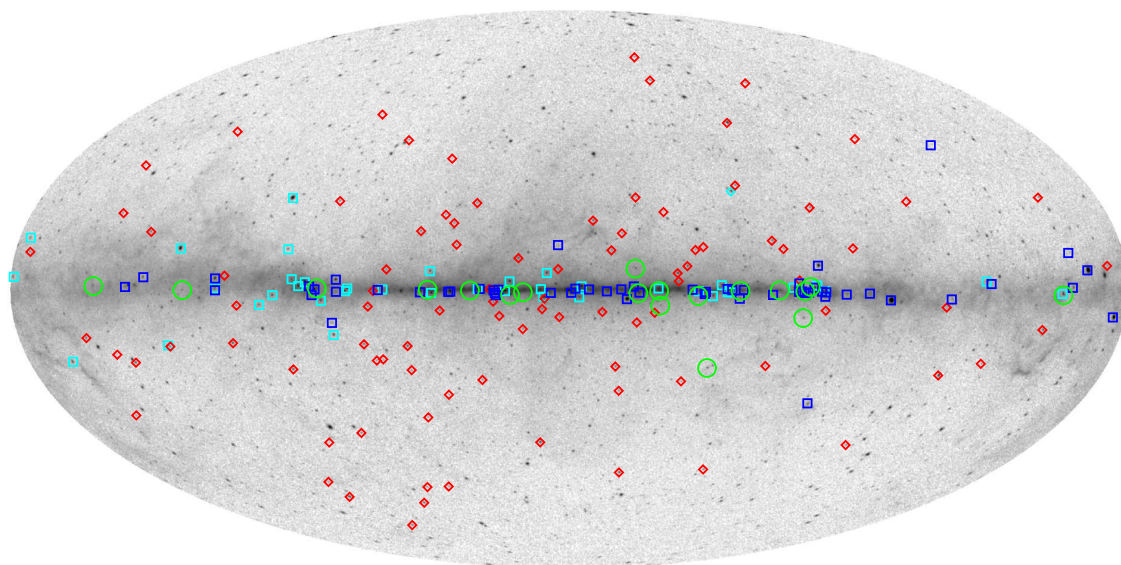


Figure 1.1: 7-year *Fermi*-LAT all-sky image, showing young (blue squares), millisecond (red diamonds) and blind-search gamma-ray pulsars (cyan squares or diamonds for young and millisecond pulsars accordingly). The new pulsars discovered during the survey described in this thesis are shown by green circles. The grayscale shows the logarithmically-scaled gamma-ray intensity above 1 GeV. This image is aligned along the Galactic plane, and centered on the Galactic Center. The bright Galactic diffuse emission from is clearly visible along the plane, where the majority of young pulsars are located, in contrast with the more isotropically distributed MSPs. LAT data courtesy of the *Fermi*-LAT Collaboration, pulsar locations compiled by E. Ferrara.

from more than 200 pulsars¹ [79, 80]. In doing so, it has provided a more representative population of gamma-ray pulsars, no longer limited to the very nearest or most energetic (and hence most luminous) pulsars, as EGRET was [60, 81]. The improved photon statistics given by the LAT’s larger collecting area provide extremely high signal-to-noise ratio pulse profiles for the brightest pulsars (as visible in Figure 1.2) [e.g. 82–84] and a large number of varied pulse profiles with which to perform large-scale tests of magnetosphere models [e.g. 85, 86].

The LAT has made many contributions to our understanding of pulsar’s high energy emissions (see [87] for a recent review). Among many other discoveries, it has: ruled out a low-altitude polar cap emission region for gamma-ray beams [79]; revealed that most young and energetic pulsars emit strongly in gamma rays [79]; probed the gamma-ray “death line”, the minimum spin-down power that a pulsar must have to emit gamma-ray pulsations [88, 89]; and shown that magnetars are not strong gamma-ray emitters [90].

Perhaps most surprisingly, the LAT has detected gamma-ray pulsations from a large (and growing) number of MSPs [e.g. 91–93]. In fact, MSPs now make up almost half of the population

¹An up-to-date list of detected gamma-ray pulsars is available at <http://tinyurl.com/fermipulsars>




Figure 1.2: Gamma-ray and radio pulse profiles from the three brightest gamma-ray pulsars, Crab (PSR J0534+2200), Vela (PSR J0835–4510) and the first radio-quiet pulsar, Geminga (PSR J0633+1746). The phase offset between the Vela pulsar’s gamma-ray and radio pulse is clearly evident, in contrast to the almost aligned peaks of the Crab pulsar. The radio profile has arbitrary intensity units. Gamma-ray (LAT) and radio (Nançay and Parkes) data from Abdo et al. [80].

of gamma-ray pulsars. Many of the MSPs’ gamma-ray pulse profiles are very similar to those of young pulsars, despite the prediction that these might be “pair-starved” by their low magnetic field strengths [87, 94].

The LAT has also become an invaluable tool in the hunt for new MSPs. Most gamma-ray pulsars show very low variability, and have exponentially cut off spectra, characteristics that set them apart from other unidentified gamma-ray sources [95]. Radio searches pointed towards promising pulsar-like LAT sources [e.g. 96–103] have discovered 54 new MSPs [104], many of which are “black widow” or “redback” pulsars in tight interacting binary systems [e.g. 105]. These can provide important tests for understanding whether isolated MSPs could form from the total evaporation of their companion stars [106]. The number of evaporating systems has increased by more than a factor of 6 since *Fermi*’s launch [107]. LAT observations of “missing link” pulsars J1023+0038 and J1227–4853 caught them in the act of switching between being rotationally powered pulsars with visible radio and gamma-ray pulsations and being in low-mass X-ray binary states, where accretion from the companion star suppresses radio pulsations and produces strong X-ray and gamma-ray emission [107, 108]. These transitional pulsars, only three of which are known, are important systems for understanding the recycling process by which MSPs are formed [109].

As the *Fermi* mission continues and more data is collected, fainter pulsars become detectable to the LAT [110], an excellent example of this being the recent detection of PSR J0540–6919 situated in the Large Magellanic Cloud: the first extragalactic gamma-ray pulsar [111]. The LAT’s long observations of pulsars have also revealed surprising new behavior, with the detection of gamma-ray flares from the Crab nebula [e.g. 112], and correlated flux, spin-down and pulse profile variations in the radio-quiet pulsar PSR J2021+4026 [113] challenging the picture of pulsars being non-variable gamma-ray emitters.

The majority of gamma-ray pulsars were first discovered through their radio or X-ray emission,

before gamma-ray pulsations were detected by folding the LAT data using their known rotational ephemerides [e.g. 92, 114] many of which were obtained through a dedicated campaign by the Pulsar Timing Consortium [115]. As a result, this population still suffers from selection biases towards gamma-ray pulsars which also emit radio pulsations [47, 80, 116].

Additionally, there remain a large number of unidentified gamma-ray sources, many of which appear to have pulsar-like spectra but are undetected in targeted radio searches. The typical gamma-ray pulsar spectrum is very similar to the spectrum predicted from the annihilation of dark matter [117]. The ability to confidently discriminate between these source classes could therefore be of importance for our understanding of fundamental physics, and for explaining the observed GeV excess towards the Galactic center [e.g. 118–120].

To attempt to address these issues, it is necessary to be able to discover pulsars through their gamma-ray pulsations alone; we must perform “blind” searches for pulsations.

1.3 Blind Searches for Gamma-ray Pulsars

As the population of *Fermi* pulsars has confirmed, gamma-ray emission is produced in the outer regions of the magnetosphere [79], with the result that gamma-ray pulses are often of very different shape (typically much wider) and not aligned in (rotational) phase with radio pulses. The wider gamma-ray beams are visible from a larger solid angle around the pulsar than radio beams [121]; many pulsars are *only* observable through their pulsed gamma-ray emission. A large sample of “radio-quiet” gamma-ray pulsars can therefore even offer a more complete view of the local neutron star population than is visible to radio telescopes [122].

The first radio-quiet gamma-ray pulsar to be detected, “Geminga”, was long known as the brightest unidentified source of gamma-ray emission [123] until the eventual detection of its pulsed X-ray emission by ROSAT [124]. Using its known rotational period, gamma-ray pulsations were soon found in EGRET data [125], and later in archival COS-B and SAS-2 data, providing a rotational ephemeris reaching back to 1973 [126].

Discovering radio-quiet (or simply previously unknown) pulsars in gamma-ray data *without* prior knowledge of their signal parameters (i.e. rotational frequency, spin-down rate or sky position) is an enormous computational challenge. Gamma-ray pulsars typically have very low photon fluxes (less than one pulsed photon is observed per day from some LAT-discovered pulsars), and therefore long integration times of months or years are required to detect pulsations with high significance. In turn, these long integration times mean that each unknown parameter must be searched with extremely fine resolution; a small offset in any of these can result in the signal quickly being lost.

The first blind searches for gamma-ray pulsars were performed in 2001 by Chandler et al. [127] using EGRET data. While these searches were able to detect previously known gamma-ray pulsars with high significance, no new pulsars were detected. Even at this stage, with searches covering month-long segments of data, computational costs proved problematic. Chandler et al. [127] employed a “stacking” technique, in which shorter segments of data were analyzed coherently before being summed together incoherently [128]. This reduced the required number of trials




Figure 1.3: Blind-search pulsar pulse profiles in 3 years of LAT data. From left to right: PSR J0007+7303, the first pulsar to be found in a blind search of gamma-ray data (data from Abdo et al. [80]); PSR J1311–3430, the first and only binary MSP to be discovered through its gamma-ray pulsations (data produced using a script written by J. Wu, timing solution from Pletsch et al. [135]); and PSR J1932+1916, the first gamma-ray pulsar to be found by the volunteer computing project *Einstein@Home* (data and timing solution from Pletsch et al. [136]). The dashed blue lines indicate the estimated background level, as calculated in Abdo et al. [80]. The weak pulse profile of PSR J1932–1916 highlights the impressive sensitivity of the first *Einstein@Home* survey.

somewhat, allowing for a larger parameter space to be searched at the same cost (in this case, a large range of spin-down values and second frequency derivatives were searched with the goal of detecting gamma-ray pulsations from magnetars).

In anticipation of the launch of *Fermi* (then known as the *Gamma-ray Large Area Space Telescope*, or *GLAST*) Atwood et al. [129] developed a new method to efficiently search for gamma-ray pulsars in data sets spanning many years. In this method, instead of taking the Fourier transform of a series of delta functions centered on the photon arrival times, the test statistic is the Fourier transform of the auto-correlation function (i.e. the set of *differences* between the photon arrival times) up to a maximum lag. This method, which can reduce the number of searches required by orders-of-magnitude, was applied to the archival EGRET data [130], but again could only detect the three brightest gamma-ray pulsars (Vela, Geminga and Crab).

The power of this “time-differencing” technique was quickly demonstrated after the launch of *Fermi* in 2008. In just one month of data, taken during the LAT’s commissioning phase, Abdo et al. [131] were able to discover a previously unknown young gamma-ray pulsar, PSR J0007+7303, in the supernova remnant CTA 1. Shortly afterwards, using four months of data, a further fifteen pulsars were discovered with these methods [132], with another eight and two pulsars being found in the first year and two years of LAT data respectively [133, 134].

However, even with this efficient method, the detection of weak pulsar signals was inhibited by the LAT’s angular resolution. Pulsed signals gain a periodic (annual) Doppler shift, caused by *Fermi*’s motion around the Solar System [e.g. 137]. This must be removed by “barycentering” photon arrival times, i.e. shifting the photon arrival times to when they would arrive at the Solar

System Barycenter, which is an approximately inertial frame of reference. These corrections depend on the true sky position of the pulsar; a small positional offset can result in annual phase modulations proportional to the pulsar’s spin frequency, and the signal can be lost completely.

While the angular resolution of the LAT is greatly improved over previous instruments, point source localizations still do not provide sky positions which are accurate enough to detect pulsations: in the most recent *Fermi*-LAT Third Source Catalog [3FGL; 138], the average 95% uncertainty radius for a point source was around 6 arcmin, whereas the angular resolution required to detect a 10 Hz pulsar is around 1.4 arcmin, and as low as a few arcseconds for MSPs. The first blind searches only searched at the center of unidentified LAT sources (or used the more accurate positions of X-ray counterpart candidates [132]). They could therefore only detect bright pulsars, whose signals could still be detected despite the loss of power due to an offset sky position; pulsars with bright X-ray emission; or slowly spinning pulsars. In particular, this issue prevented the blind-search detection of millisecond pulsars.

The problem of searching for periodic astrophysical signals in very long time series is a familiar one for the gravitational wave community. Searches for continuous gravitational waves (CWs) from spinning neutron stars with slight non-axisymmetric distortions face many of the same challenges as gamma-ray pulsar searching: low signal strengths requiring long observation times to be detectable; enormous parameter spaces in which to search; and limited computing power [128, 139–142]. So-called “semicoherent” methods, in which segments of data are analyzed coherently, and the signal power from each summed incoherently, have been used extensively to perform efficient and sensitive CW searches [e.g. 143–147]. Notably, CW searches have utilized the concept of a “parameter space metric” [e.g. 148, 149], which describes the distance from the parameters of a signal to another point in the parameter space in terms of the loss in measured signal power. This provides an analytical method to determine an optimum search-grid spacing which ensures that the parameter space is adequately covered without wasting time searching at too many locations.

Rather than summing the coherent power from non-overlapping segments, as the stacking method employed by Chandler et al. [127] does, the time-differencing technique of Atwood et al. [129] is equivalent to a “sliding coherence window” approach, which gains extra sensitivity from overlapping segments, as described in Pletsch [150]. These sliding-window techniques developed to search for CWs were adapted for the purposes of searching for gamma-ray pulsars by Pletsch et al. [151]. Importantly, this survey extended the search over frequency and spin-down to a 4-dimensional search including the sky position, using the metric formulation to ensure optimal coverage. The introduction of a photon weighting scheme [152] avoided the need to apply strict energy or angular cuts when producing photon datasets and significantly improved the sensitivity to weak pulsations. This survey, running on the ATLAS cluster computer [153], discovered ten more young gamma-ray pulsars [151, 154], and a further extension to search over binary parameters resulted in the spectacular blind-search discovery of PSR J1311–3430 [135], the only binary MSP to be found through its gamma-ray pulsations so far.

Even with these new methods, as the time span covered by the *Fermi*-LAT data grew, the computational cost of running a search increased, and surveys like this became prohibitively expensive to run, even on a dedicated computing cluster. To meet these costs, Pletsch et al. [136]

turned to the distributed volunteer computing system, *Einstein@Home* [155]. In this scheme, the full parameter space which must be searched is split into several thousand smaller chunks, called “work units”, each of which can be searched within a few hours on a typical private desktop computer. Using the Berkeley Open Infrastructure for Network Computing [BOINC; 156, 157] as its base, *Einstein@Home* sends out work units to volunteered computers which then perform the necessary calculations when they are not otherwise being used. This “citizen science” scheme has proven to be highly popular, and powerful. Volunteers have donated computing cycles from more than 1.5 million computers to the project since it began, and with more than 60,000 currently active computers, *Einstein@Home* has a sustained computing power of > 2 PFLOP/s, making it comparable to the world’s 50 fastest supercomputers [158]. Originally designed to search for gravitational waves [144–146], *Einstein@Home* has also been used to perform pulsar searches, leading to the discovery of more than 50 new radio pulsars [159–161] including the recent detection of a rare highly eccentric binary MSP [162].

The first *Einstein@Home* gamma-ray pulsar survey discovered four new radio-quiet pulsars [136], two of which suffered large glitches during the observation interval, making their discovery all the more impressive. Of these results Caraveo [47] commented that

“the additional computer power provided by Einstein@Home combined with an improved search technique is a good omen for more findings, hopefully also of the long-sought radio-quiet MSP.”

At this stage, 41 pulsars had been detected in blind searches of LAT data, only four of which were subsequently detected in radio searches. However, a large number of bright, pulsar-like LAT sources remain unidentified [138], and the long-predicted class of radio-quiet gamma-ray MSPs [163, 164] remained undiscovered.

The main aims of this author’s doctoral studies were to further develop blind search methods, and to apply these methods to carry out more sensitive surveys for gamma-ray pulsars on *Einstein@Home* using the latest *Fermi*-LAT data. This has led to the discovery of nineteen new pulsars, just under one third of all pulsars to be found through their gamma-ray pulsations. The discovery of the first of these new pulsars, PSR J1906+0722, explained the nature of the previously most significant unassociated gamma-ray source [Chapter 3]. The new pulsars discovered by this survey (shown in their place on the frequency–spin-down diagram in Figure 1.4) span, and even extend, the previously known population of radio-quiet gamma-ray pulsars, including the two most slowly rotating gamma-ray pulsars [Chapter 4], and the youngest known radio-quiet gamma-ray pulsar [Chapter 5]. Perhaps most excitingly, these discoveries include two new isolated MSPs, neither of which were detected in extensive radio searches [Chapter 6].




Figure 1.4: Frequency (f) vs. spin-down (\dot{f}) diagram, showing the new pulsars discovered during the author’s doctoral studies. This is an updated version of Figure 4.1, which reports a subset of the newly discovered pulsars. Non-gamma-ray pulsars from the ATNF Pulsar Catalogue [14] are shown as black crosses, gamma-ray pulsars detected by *Fermi*-LAT as blue circles, and new pulsars discovered by the latest *Einstein@Home* survey as orange squares. Lines of constant characteristic age (dot-dashed), surface magnetic field strength (dotted) and spin-down power (dashed) are also shown. These properties are calculated by assuming canonical values for pulsars’ moments of inertia ($I = 10^{45}$ g cm²), radii (10 km), and a simple rotating magnetic dipole model, as described in [e.g. 9, 80].

1.4 Timing Gamma-ray Pulsars

The low photon flux from gamma-ray pulsars not only makes them harder to detect than radio pulsars, but also complicates the procedures used to “time” their rotations.

The traditional method of timing a pulsar, described by Lorimer & Kramer [9, p. 200] and briefly summarized here, is designed for use with typical radio observations. These are normally short (~ 1 hr) observations separated by longer intervals (the most well-studied pulsars are observed

daily [e.g. 3]). Individual pulses cannot be detected from most radio pulsars, however “folding” these short observations (wrapping the de-dispersed data around at the pulsar’s rotation period, and summing), allows for a significant integrated pulse profile to be built up. The time at which this averaged pulse arrives is known as a “time of arrival”, or TOA.

By carrying out multiple observations, each yielding a single TOA, an ephemeris (or “timing solution”) describing the rotation of the pulsar over the entire observing interval can be produced by finding the set of rotational parameters which best fits the observed set of TOAs.

While the TOA fitting procedure has been successfully applied to *Fermi*-LAT data, for example by Kerr et al. [165], the requirement for this procedure to work well (that is, the signal strength is high enough that the data can be split into short segments, each with a high signal-to-noise ratio) does not always apply to gamma-ray pulsars. Indeed, for dimmer pulsars the full span of the *Fermi* mission is required to accumulate enough photons to detect statistically significant pulsations [110]. The low flux is particularly problematic for measuring effects which occur over timescales shorter than the required integration time, for example binary motion or glitches [166].

During these studies, the methods of Ray et al. [167] were built upon to develop an entirely unbinned (in both phase and time) likelihood maximization procedure to time weak gamma-ray pulsars. This procedure can be used to make precise estimates of pulsars’ rotational and positional properties, and for robust estimates of the statistical uncertainties on each parameter and correlations between them. This method was used to time the redback PSR J2339–0533 [166], and extrapolate its orbital parameters back from its radio discovery to the beginning of the *Fermi* mission. This revealed modulation in its orbital period, possibly due to variations in its companion star’s gravitational quadrupole moment.

This method, described in most detail in Section 4.4.1, was used to study the new pulsars discovered by the *Einstein@Home* survey. Notably, this resulted in measurements of large glitches experienced by two newly discovered pulsars [Chapters 3 & 4]; a new gamma-ray pulsar’s braking index [Chapter 5]; and the proper motion of two new MSPs [Chapter 6].

1.5 Chapter Descriptions and Authorship Clarifications

1.5.1 Description of Chapter 2

Chapter 2 contains a detailed investigation of the blind search methodology, with the goal of maximizing the sensitivity of a blind search to weak pulsations under the constraint of a fixed computational cost. This chapter builds on the earlier work of Pletsch et al. [151] in adapting semicoherent methods, previously used to search for continuous gravitational waves, to search for gamma-ray pulsars. A comparison using simulations of weak pulsar signals showed that the improved methods result in searches which are almost 50% more sensitive than those using the methods developed in Atwood et al. [129], for the same computational cost. This chapter was originally published as Pletsch & Clark [168].

This chapter was a close collaborative effort between the author and supervisor, H. J. Pletsch, who provided the ideas for the study. The majority of the text in this chapter was written by H. J. Pletsch, particularly in Sections 2.1 to 2.4, although both authors contributed to the text, calculations and derivations in all sections.

The author was primarily responsible for:

- the derivation of the coherent (Sections 2.3.3, 2.8.2 and 2.8.3) and semicoherent metrics (Sections 2.4.2 and 2.8.7), resulting in more accurate approximations than those given in Pletsch et al. [151];
- the modification of frequency- and lag-domain interpolation methods (Sections 2.5.2 and 2.5.4) for use with the DFT formulation of the semicoherent detection statistic;
- the simulations used to demonstrate the improved sensitivity of the new techniques (Section 2.6);
- the calculation of optimal mismatches for coherent (Section 2.8.5) and semicoherent (Section 2.8.8) searches;
- the derivation of the statistical properties of the semicoherent detection statistic (Section 2.8.6);
- and the procedure for constructing an efficient grid of sky locations to search (Section 2.8.9).

The software used by Pletsch et al. [136] was updated with the improvements developed in Chapter 2, including an additional semicoherent refinement stage between the initial semicoherent search and the coherent follow-up step. These improvements were mostly implemented by the author although the implementation of the “complex-to-real” FFT algorithm was carried out by H.-B. Eggenstein. This software was then used to perform a new blind search survey running on *Einstein@Home* to search for pulsations from unidentified sources from the 3FGL [138]. The author was responsible for designing and carrying out the survey within the framework of *Einstein@Home*, maintained by C. Beer, O. Bock, H.-B. Eggenstein, B. Machenschalk and staff at the University of

Wisconsin–Milwaukee. The source selection and ranking, and preparation of LAT datasets were performed by J. Wu under the guidance of L. Guillemot. An associated paper describing these is in preparation by J. Wu.

A follow-up pipeline was developed by A. Cuéllar and the author. This ran daily on the ATLAS cluster to automatically identify and refine the parameters of significant pulsar candidates. Many additional candidates were followed-up “manually” by the author, which notably resulted in the discovery of the glitching pulsar PSR J1844–0346 (see Chapter 4).

The remaining chapters describe the new pulsars discovered by this survey.

1.5.2 Description of Chapter 3

Shortly after its commencement, the *Einstein@Home* survey detected pulsations from a young pulsar, PSR J1906+0722, in a highly significant but previously unidentified LAT source. Several factors complicated its discovery, including a large positional offset from its associated LAT source, and the fact that the pulsar suffered a large glitch in the early portion of the data. Chapter 3, originally published as Clark et al. [169], describes the discovery and subsequent analysis of this pulsar.

In order to time this new pulsar, an unbinned likelihood timing analysis method was developed by the author, H. J. Pletsch and L. Nieder, with help from H. Fehrmann in parallelizing the computations to run on new Intel XeonPhi hardware. Template pulse profiles were fit to the observed photon phases using software written by M. Kerr. The timing procedure is first described in Section 3.3.2 and again in later chapters, with most detail provided in Section 4.4.1. Additionally, this procedure was used in Pletsch & Clark [166] to time a binary millisecond pulsar, revealing quasi-periodic modulations in its binary period, suggestive of quadrupole moment variations in its companion star.

The majority of the text in this chapter was written by the author, under the guidance of H. J. Pletsch. The timing analyses (Sections 3.3.1 and 3.3.2) were performed by the author. The LAT data preparation (Section 3.2.1) and off-pulse spectral analysis (Section 3.3.3) was performed by J. Wu and L. Guillemot, who also contributed the text in these sections (although Figure 3.4 was produced by the author using their results). The radio follow-up searches were performed by J. Wu, P. Torne and D. J. Champion. The calculation of the upper limit on the X-ray flux from this pulsar was performed by F. Krauss and M. Marelli. This chapter was reviewed within the *Fermi*-LAT collaboration by P. M. Saz Parkinson and H. Laffon, and additional comments were incorporated from L. Baldini, P. A. Caraveo, S. Digel, X. Hou, M. Razzano, D. J. Thompson and M. Wood.

1.5.3 Description of Chapter 4

Chapter 4 describes in detail the majority of the results of the *Einstein@Home* survey, including thirteen of the newly discovered pulsars, and estimates of the upper limits on the fraction of pulsed gamma-ray emission from the remaining sources in which no pulsations were detected. Following the initial submission of this thesis, this chapter was published as Clark et al. [170]. An accompanying paper is being prepared by J. Wu to describe the target source selection and ranking procedures, data preparation, and the spectral analyses and radio follow-up searches for these new

pulsars.

The text in this chapter was written by the author, with suggested modifications and corrections provided by H. J. Pletsch. Again, the LAT datasets were prepared by J. Wu and L. Guillemot. All of the remaining analyses in this chapter were performed by the author. This chapter was reviewed within the *Fermi*-LAT collaboration by D. A. Smith who, in addition to other comments, suggested important additions to the text regarding the details of the data preparation and their effect on the sensitivity estimates made in this chapter. Additional comments were provided by B. Allen, S. Digel, L. Guillemot, L. Nieder, J. Perkins and D. J. Thompson.

1.5.4 Description of Chapter 5

Chapter 5 describes the discovery and timing analysis of PSR J1208–6238, the youngest known radio-quiet gamma-ray pulsar, and one of a group of only ten pulsars with measurable braking indices. Following the initial submission of this thesis, this chapter was published as Clark et al. [171].

This chapter was written by the author under the guidance of H. J. Pletsch, and the timing analyses (Section 5.3) and calculations in the discussion (Section 5.4) were performed by the author. LAT datasets were prepared by J. Wu and L. Guillemot, with J. Wu additionally performing the off-pulse analysis (Section 5.2.3). The gamma-ray pulse profile modeling was performed by T. J. Johnson, who also suggested some additions to the text describing this analysis. Radio follow-up searches with the Parkes radio telescope were performed by F. Camilo and M. Kerr, the latter of whom provided the text for this (Section 5.2.4) in addition to valuable comments on the manuscript. This chapter was also reviewed within the *Fermi*-LAT collaboration by M. Razzano. Further comments were incorporated from B. Allen, F. Camilo, E. Cavazzuti, L. Guillemot, T. J. Johnson, M. Kramer, F. de Palma, J. Perkins and D. A. Smith.

1.5.5 Description of Chapter 6

While searching in two high-Galactic-latitude unassociated 3FGL sources, the *Einstein@Home* survey discovered pulsations from two new isolated gamma-ray MSPs. These remain undetected in radio observations despite several searches, making these potentially the first radio-quiet MSPs, a long-predicted but never before seen class of pulsar. Their discovery and subsequent analysis is described in Chapter 6.

This is a preliminary version of a report that we intend to adapt for submission to *Science*. Following the initial submission and defense of this thesis, we obtained additional dedicated radio observations of these pulsars. These will be described in the future published article. The abstract and introduction were primarily written by H. J. Pletsch, with additions by the author. LAT datasets were prepared by J. Wu, who additionally performed the spectral and off-pulse analyses (Section 6.6). The timing analysis (Section 6.3) was performed by the author, who wrote the text in this section. The distance estimates from the proper motion measurements were calculated by H. J. Pletsch; the author expanded that analysis to include estimates of the corrected spin-down powers of each pulsar.

The gamma-ray pulse profile modeling was performed by T. J. Johnson, who contributed to the text for this (Section 6.4), and provided the data for Figure 6.3 and Table 6.2. Radio searches were performed with the Parkes radio telescope by F. Camilo, who provided text on these (Section 6.4), and M. Kerr. These observations took place during a search campaign for radio MSPs targeting unidentified pulsar-like LAT sources [102]. Analyses of the X-ray counterparts detected by *XMM-Newton* were performed and written about by D. Salvetti (Section 6.7). The final section discussing the implications of these discoveries was written by the author, with additions from H. J. Pletsch.

1.5.6 Description of Chapter 7

Chapter 7 describes the two latest discoveries made by the *Einstein@Home* survey and gives some short speculation about future blind search surveys. One of the new pulsars, PSR J1817–1742, has the lowest point-source significance of any blind search pulsar, validating the sensitivity estimate made in Section 4.5.

The text and analysis in this chapter was written and performed by the author. The data used in this chapter were produced by scripts written by J. Wu and L. Guillemot, which additionally output the spectral parameters of the new pulsars, given in Table 7.1.

Chapter 2

Optimized Blind Gamma-ray Pulsar Searches at Fixed Computing Budget

*Originally published as Pletsch & Clark, 2014, The Astrophysical Journal, 795, 75.
© 2014. The American Astronomical Society. All rights reserved.*

Abstract

The sensitivity of blind gamma-ray pulsar searches in multiple years worth of photon data, as from the *Fermi* LAT, is primarily limited by the finite computational resources available. Addressing this “needle in a haystack” problem, we here present methods for optimizing blind searches to achieve the highest sensitivity at fixed computing cost. For both coherent and semicoherent methods, we consider their statistical properties and study their search sensitivity under computational constraints. The results validate a multistage strategy, where the first stage scans the entire parameter space using an efficient semicoherent method and promising candidates are then refined through a fully coherent analysis. We also find that for the first stage of a blind search incoherent harmonic summing of powers is not worthwhile at fixed computing cost for typical gamma-ray pulsars. Further enhancing sensitivity, we present efficiency-improved interpolation techniques for the semicoherent search stage. Via realistic simulations we demonstrate that overall these optimizations can significantly lower the minimum detectable pulsed fraction by almost 50% at the same computational expense.

2.1 Introduction

We here give a more detailed description of the strategies and methods exploited in the blind searches performed by Pletsch et al. [136, 151, 154], and how they relate to the seminal “time differencing technique” by Atwood et al. [129, hereafter A06]. We consider related questions one might be faced with when setting up a blind search: Could a fully coherent blind search using a subset of data perhaps be more sensitive than a semicoherent search using all of the data? Is harmonic summing worthwhile under computational constraints? What is the optimal search-grid point density to balance sensitivity versus computing effort? In addressing such questions, we present the technical framework to optimize the sensitivity of blind pulsar searches in gamma-ray data at fixed computing cost. Moreover, we present further important methodological advances to improve the overall blind-search efficiency.

The paper is organized as follows. In Section 2.2, we describe the statistical detection of pulsations in general. In Section 2.3, we discuss the statistical properties of coherent blind searches and study their computational cost scalings using the parameter-space metric. We also investigate the efficiency of harmonic summing for different pulse profiles. In Section 2.4, we describe the statistical properties of a semicoherent blind-search method and compare the respective computing demand using the semicoherent metric. Section 2.5 presents a collection of technical improvements for the implementation of the semicoherent search stage, including efficient interpolation methods and automated candidate follow-up procedures. We demonstrate the superiority from combining these advances through realistic simulations in Section 2.6. Finally, conclusions follow in Section 2.7.

2.2 Statistical Detection of Pulsations

In blind pulsar searches the pulse profile (the periodic light curve) and the exact parameters describing the rotational evolution of the neutron star are *unknown* in advance. As Bickel et al. [172] have pointed out, unless the pulse profile shape is precisely known, there is no universally optimal statistical test, because any most powerful test for one template profile will not be most powerful against another. Any test can only be most sensitive to a finite-dimensional class of targets. Thus, for computational feasibility of a blind search an efficient (potentially suboptimal) template pulse profile to test against should attain only modest reduction in detection sensitivity compared to an optimal template. The construction of such a test can be guided by the profiles of known gamma-ray pulsars, which we will consider below.

For isolated pulsars the search parameters describing the rotational phase of the neutron star is at least four-dimensional, consisting of frequency f , spindown rate \dot{f} , and sky position with right ascension α and declination δ . To the LAT-registered arrival times t_{LAT} sky-position (α, δ) dependent corrections (“barycentric corrections”) are applied in order to obtain the photon arrival times t at the solar system barycenter (SSB). Then the rotational phase $\Phi(t)$ is described by

$$\Phi(t) = \phi_0 + 2\pi f(t - t_0) + 2\pi \dot{f} \frac{(t - t_0)^2}{2}, \quad (2.1)$$

where f and \dot{f} are defined at reference time t_0 , when the phase equals the constant ϕ_0 .

Apart from the arrival time, for each of N detected gamma-ray photons, indexed by j , the LAT also records the photon's reconstructed energy and direction. From these a weight, w_j , can be computed measuring the probability that it has originated from the target source [152, 172]. Using these probability weights efficiently avoids testing different hard selection cuts on energy and direction (implying binary weights), providing near optimal pulsation detection sensitivity [151, 152].

The observed gamma-ray pulse profile $F(\Phi)$, the flux as a function of Φ , can be written as

$$F(\Phi) \propto \frac{1-p}{2\pi} + p F_s(\Phi), \quad (2.2)$$

where p is the *pulsed fraction* that is estimated by the number of pulsed gamma-ray photons divided by the total number of photons. $F_s(\Phi)$ represents the pulse profile (undisturbed by background) and is a probability density function on $[0, 2\pi]$, which can be expressed as a Fourier series

$$F_s(\Phi) = \frac{1}{2\pi} \left(1 + \sum_{n \neq 0} \alpha_n e^{in\Phi} \right), \quad (2.3)$$

with the complex Fourier coefficients α_n , defined at harmonic order n as

$$\alpha_n = \int_0^{2\pi} F_s(\Phi) e^{-in\Phi} d\Phi. \quad (2.4)$$

Hence the total flux $F(\Phi)$ can be rewritten as

$$F(\Phi) \propto 1 + p \sum_{n \neq 0} \alpha_n e^{in\Phi}. \quad (2.5)$$

If $F_s(\Phi)$ is an exact sinusoidal pulse profile, then from Equation (2.4) it follows that $|\alpha_1| = 1/2$ and all other coefficients vanish, $|\alpha_{n>1}| = 0$. As another example, if the pulse profile $F_s(\Phi)$ is a Dirac delta function, i.e. the narrowest possible profile, then all coefficients are equal, $|\alpha_n| = 1$, implying equal Fourier power at all harmonic orders.

In general, the null hypothesis is given by $p = 0$, meaning that all phases are uniformly distributed (i.e. no pulsations). From the likelihood for photon arrival times Bickel et al. [172] derived a score test statistic Q_M for $p > 0$,

$$Q_M = \frac{1}{K^2} \sum_{n=1}^M |\alpha_n|^2 |A_n|^2, \quad (2.6)$$

where we defined the normalization constant K [different from 172] as

$$K^2 = \frac{1}{2M} \sum_{n=1}^M |\alpha_n|^2, \quad (2.7)$$

and A_n is given by

$$A_n = \frac{1}{\kappa} \sum_{j=1}^N w_j e^{-in\phi(t_j)}, \quad (2.8)$$

with the time-dependent part of the phase $\phi(t) = \Phi(t) - \phi_0$ and the normalization constant κ defined as

$$\kappa^2 = \frac{1}{2} \sum_{j=1}^N w_j^2. \quad (2.9)$$

Thus, we denote by \mathcal{P}_n the coherent Fourier power at the n th harmonic,

$$\mathcal{P}_n = |A_n|^2 = \frac{1}{\kappa^2} \left| \sum_{j=1}^N w_j e^{-in\phi(t_j)} \right|^2. \quad (2.10)$$

Appealing to the Central Limit Theorem (since $N \gg 1$ in all practical cases) the normalization choice of Equation (2.9) has the convenient property that the coefficients $\Re(A_n)$ and $\Im(A_n)$ become independent Gaussian random variables with zero mean and unit variance under the null hypothesis. Therefore, to good approximation each \mathcal{P}_n is χ^2 -distributed with 2 degrees of freedom, as will be discussed below. Thus, Q_M is the weighted sum of coherent Fourier powers,

$$Q_M = \sum_{n=1}^M \frac{|\alpha_n|^2}{K^2} \mathcal{P}_n. \quad (2.11)$$

Therefore, as noted by Bickel et al. [172], the test statistic Q_M is invariant under phase shifts (i.e. independent of reference phase ϕ_0) and only depends on the amplitudes of the Fourier coefficients α_n , but not on their phases. Moreover, Beran [173] showed earlier that if the pulse profile is known a priori, a test statistic following from Q_M for binary weights is locally most powerful for testing uniformity of a circular distribution, assuming unknown and weak (small p) signal strength.

2.3 Coherent Test Statistics

In what follows, we examine the sensitivity of coherent blind searches at fixed computational cost, taking into account the statistical properties and sensitivity scalings in terms of relevant quantities. For simplicity, during the remainder of this section we here assume hard photon selection cuts, i.e., binary weights only, $w_j \in \{0, 1\}$, such that \mathcal{P}_n reduces to

$$\mathcal{P}_n = \frac{2}{N} \left| \sum_{j=1}^N e^{-in\phi(t_j)} \right|^2. \quad (2.12)$$

However, the main conclusions obtained will also have applicability when arbitrary (i.e., non-binary) weights are used.

2.3.1 Statistical Properties

Under the null hypothesis $p = 0$ and assuming $N \gg 1$, the coherent power \mathcal{P}_n as of Equation (2.12) follows a central χ^2 -distribution with 2 degrees of freedom (see Appendix 2.8.1), whose the first two moments are,

$$E_0 [\mathcal{P}_n] = 2, \quad Var_0 [\mathcal{P}_n] = 4. \quad (2.13)$$

Suppose the photon data contains a pulsed signal, $p > 0$, whose pulse profile can be expressed in terms of complex Fourier coefficients, γ_n as in Equation (2.4). In this case, we show in Appendix 2.8.1 that for moderately strong pulsed signals the distribution of \mathcal{P}_n can be well approximated by a noncentral χ^2 -distribution [174, 175] with 2 degrees of freedom. Thus, in the perfect-match case (the pulsar parameters f , \dot{f} , and sky position are precisely known), the first two moments are approximately given by¹

$$E_p [\mathcal{P}_n] \approx 2 + 2p^2 N |\gamma_n|^2, \quad (2.14a)$$

$$Var_p [\mathcal{P}_n] \approx 4 + 8p^2 N |\gamma_n|^2, \quad (2.14b)$$

where pN photons are assumed to be ‘‘pulsed’’ and accordingly $(1 - p)N$ photons are ‘‘non-pulsed’’ (i.e., background). Thus, the second summand in Equation (2.14a) represents the noncentrality parameter.² We can also identify the amplitude signal-to-noise ratio (S/N) at the n th harmonic, $\theta_{\mathcal{P}_n}$, as

$$\theta_{\mathcal{P}_n}^2 = \frac{E_p [\mathcal{P}_n] - E_0 [\mathcal{P}_n]}{\sqrt{Var_0 [\mathcal{P}_n]}} \approx p^2 N |\gamma_n|^2. \quad (2.15)$$

Therefore, by comparison to Equation (2.14a) the noncentrality parameter is just $2\theta_{\mathcal{P}_n}^2$.

A similar calculation for Q_M , based on the above relations shows that if $p = 0$,

$$E_0 [Q_M] = 4M, \quad Var_0 [Q_M] = \frac{4}{K^4} \sum_{n=1}^M |\alpha_n|^4, \quad (2.16)$$

and for $p > 0$, one obtains

$$E_p [Q_M] \approx 4M + \frac{2p^2 N}{K^2} \sum_{n=1}^M |\alpha_n|^2 |\gamma_n|^2. \quad (2.17)$$

Thus, the amplitude S/N θ_{Q_M} for the test statistic Q_M can be expressed as

$$\theta_{Q_M}^2 \approx \frac{p^2 N \sum_{n=1}^M |\alpha_n|^2 |\gamma_n|^2}{\sqrt{\sum_{n=1}^M |\alpha_n|^4}}. \quad (2.18)$$

¹In fact, these values scale with $N - 1$ rather than N , since clearly we need more than one photon to detect pulsations. In practice N must be large to detect pulsations, so the approximations are justified. With thanks to L. Nieder for noticing this.

²A random variable X following a non-central χ^2 -distribution with 2 degrees of freedom and noncentrality parameter λ , has expectation value $2 + \lambda$.

A similar expression has been derived by Bickel et al. [172] who used this parameter as an approximate measure of the sensitivity of the test statistic Q_M , since the larger the S/N θ_{Q_M} the higher the probability of detection. However, it is only an approximate sensitivity measure, because any meaningful sensitivity comparison must be done at fixed probability of false alarm as will be described below. Equation (2.18) also shows that the S/N is maximized if $|\alpha_n|^2 \propto |\gamma_n|^2$, i.e., when the template pulse profile α_n perfectly matches the γ_n , representing the signal pulse profile. However, as Bickel et al. [172] correctly note, practical blind searches can only test for a finite-dimensional class of template pulse profiles.

A particularly simple template profile for a given value of M is

$$|\alpha_n| = \begin{cases} 1, & n \leq M \\ 0, & n > M \end{cases}. \quad (2.19)$$

With this choice, Q_M measures the coherent Fourier power summed over the first M harmonics, which we therefore refer to as *incoherent harmonic summing*. The resulting statistic is also known as Z_M^2 [55],

$$Z_M^2 = \sum_{n=1}^M \mathcal{P}_n. \quad (2.20)$$

Maximizing Z_M^2 over different values of M as $H = \max_{1 \leq M \leq 20} (Z_M^2 - 4M + 4)$ also recovers the widely used H -test by de Jager et al. [176].

The template of Equation (2.19) has the additional benefit that the statistical distribution of Z_M^2 is known analytically. Therefore, we use this to obtain realistic sensitivity scalings for such coherent test statistics. Since \mathcal{P}_n is χ_2^2 -distributed³, it follows that Z_M^2 is distributed as χ_{2M}^2 . Thus, one obtains

$$E_0 [Z_M^2] = 2M, \quad \text{Var}_0 [Z_M^2] = 4M, \quad (2.21)$$

and

$$E_p [Z_M^2] \approx 2M + 2\theta_M^2 \sqrt{M}. \quad (2.22)$$

Correspondingly, the S/N θ_M is written as

$$\theta_M^2 = \frac{1}{\sqrt{M}} \sum_{n=1}^M \theta_{\mathcal{P}_n}^2 = \frac{p^2 N}{\sqrt{M}} \sum_{n=1}^M |\gamma_n|^2. \quad (2.23)$$

In the Neyman–Pearson sense, we define *search sensitivity* from the *lowest threshold pulsed fraction* required to achieve a certain detection probability P_{DET}^* for a given number of photons N and at given false alarm probability P_{FA}^* . For Z_M^2 the false alarm probability is computed as

$$P_{\text{FA}}(Z_{M,\text{th}}^2) = \int_{Z_{M,\text{th}}^2}^{\infty} \chi_{2M}^2(Z_M^2; 0) dZ_M^2, \quad (2.24)$$

³We use the notation χ_k^2 to indicate a χ^2 -distribution with k degrees of freedom.

where $\chi_k^2(X; \lambda)$ denotes the probability density function for the χ_k^2 -distributed variable X with noncentrality parameter λ . The probability of detection for a noncentrality parameter of $2\theta_M^2\sqrt{M}$ is

$$P_{\text{DET}}(Z_{M,\text{th}}^2, 2\theta_M^2\sqrt{M}) = \int_{Z_{M,\text{th}}^2}^{\infty} \chi_{2M}^2(Z_M^2; 2\theta_M^2\sqrt{M}) dZ_M^2. \quad (2.25)$$

The minimum detectable pulsed-fraction threshold for summing coherent power from M harmonics, $p_{\text{coh},M}$, is obtained by first inverting Equation (2.24) to get the threshold test-statistic value $Z_{M,\text{th}}^2(P_{\text{FA}}^*)$, which in a second step is substituted in Equation (2.25) to numerically find the required threshold S/N:

$$\theta_M^* = \theta_M(P_{\text{FA}}^*, P_{\text{DET}}^*). \quad (2.26)$$

Finally, Equation (2.23) can be used to convert the threshold S/N θ_M^* into $p_{\text{coh},M}$, which defines the coherent search sensitivity as

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{N}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (2.27)$$

Assuming the overall photon count rate, $\mu = N/T_{\text{coh},1}$, is constant throughout the entire coherent integration time, $T_{\text{coh},1}$ then the search sensitivity increases with the well-known square-root scaling of $T_{\text{coh},1}$,

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{\mu T_{\text{coh},1}}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (2.28)$$

Thus, we have obtained an expression for the search sensitivity, separating the two effects of photon count rate (or integration time) and pulse profile shape. Regarding the latter effect, Equation (2.28) reveals that the sensitivity only improves with including higher harmonics (i.e. increasing M) if the pulse profile shape is such that $(\sum_{n=1}^M |\gamma_n|^2)^{1/2}$ increases more quickly than the ‘‘statistical penalty’’ factor $M^{1/4} \theta_M^*$. While this is true for the narrowest possible pulse profile (a Dirac delta function), we show below that the same does not hold in general for typical gamma-ray pulsar profiles.

2.3.2 Effects of Pulse Profile on Sensitivity

From Equation (2.28) in the previous section, we have seen how the sensitivity for pulsation detection depends on the shape of the pulse profile, represented by the Fourier coefficients γ_n . Therefore, it is instructive to examine the change in sensitivity as a function of the number of harmonics M for some exemplary profiles. Thus, we consider the following ratio,

$$\frac{p_{\text{coh},M}^{-1}}{p_{\text{coh},1}^{-1}} = \frac{\theta_1^*}{M^{1/4} \theta_M^*} \frac{1}{|\gamma_1|} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}, \quad (2.29)$$

which compares in the statistical sense the search sensitivity of including M harmonics, compared to using the fundamental only (in absence of any computational constraints).

In the ideal case, where all harmonics have equal power $|\gamma_n|^2 = 1$, the pulse profile is a Dirac delta function as described above. In this case, $(\sum_{n=1}^M |\gamma_n|^2)^{1/2} = M^{1/2}$, and the sensitivity is a monotonically increasing function of M at fixed detection probability, P_{DET}^* , and fixed false alarm probability, P_{FA}^* . To illustrate this, consider the following example, assuming that $P_{\text{FA}}^* = 1\%$ and $P_{\text{DET}}^* = 90\%$. Then, to good approximation, the corresponding S/N threshold θ_M^* can be described by

$$\theta_M^* \approx \left(3.715 + \frac{4.987}{\sqrt{M}} \right)^{1/2}. \quad (2.30)$$

Hence, with increasing M , the threshold S/N θ_M^* decreases and becomes constant in the limit of large M , in which case the statistical penalty factor $(M^{1/4} \theta_M^*)$ becomes $\propto M^{1/4}$. Since this scaling is slower than the pulse profile factor $(\sum_{n=1}^M |\gamma_n|^2)^{1/2} = M^{1/2}$ in this case, the sensitivity is monotonically increasing with M . This is also shown in Figure 2.1, using the exact values for θ_M^* that we calculated numerically.

To obtain a more realistic signal pulse-profile model, we considered those of the known gamma-ray pulsars. We carried out a harmonic analysis of the pulse profile shapes of the 117 known gamma-ray pulsars listed in the second *Fermi* LAT pulsar catalog [80] and computed their Fourier coefficients, γ_n . These are shown in Figure 2.2 (top panel) and illustrate that for most of the known gamma-ray pulsars the largest fraction of Fourier power is typically in a single harmonic that is either the first (mostly single-peaked profiles) or the second (mostly two-peaked profiles). Therefore, before computing an average profile (by averaging the $|\gamma_n|$), it makes sense to divide the pulsars into these two groups (based on whether or not $|\gamma_1| > |\gamma_2|$). These results, separately for each group, are displayed in the two bottom panels of Figure 2.2.

We use the resulting two sets of coefficients γ_n to calculate the sensitivity scaling with M from Equation (2.28) as also shown in Figure 2.1. Notice that for the typical pulse profiles, in contrast to the Dirac delta pulse-profile, when summing more than a certain number of harmonics, the sensitivity starts to decrease (at fixed P_{DET}^* and P_{FA}^*). This is because the Fourier powers $|\gamma_n|^2$ at the higher harmonics become vanishingly small and thus effectively only contribute “noise” when summed (i.e. the statistical penalty factor cannot be overcome anymore).

These results also illustrate the success of the H -test for targeted pulsation searches in gamma-ray data with known pulsar ephemerides, because this test maximizes the Fourier power sums over the first 20 harmonics. Maximizing only over fewer harmonics could likely already be sufficient (or even be more sensitive due to the reduced trials factor) in most cases, as suggested by Figure 2.1. Besides, further improvements over the H -test could also be achieved by employing one or more template profiles α_n that are more representative of the typical gamma-ray profile (than the delta function) to compute the Q_M test statistic. Using the average profile from the known pulsars from above for this seems the simplest first step. While also conducting a principal component analysis appears worthwhile, we defer a detailed study of this to future work.

So far, we have not considered the computational costs involved, which is only justifiable

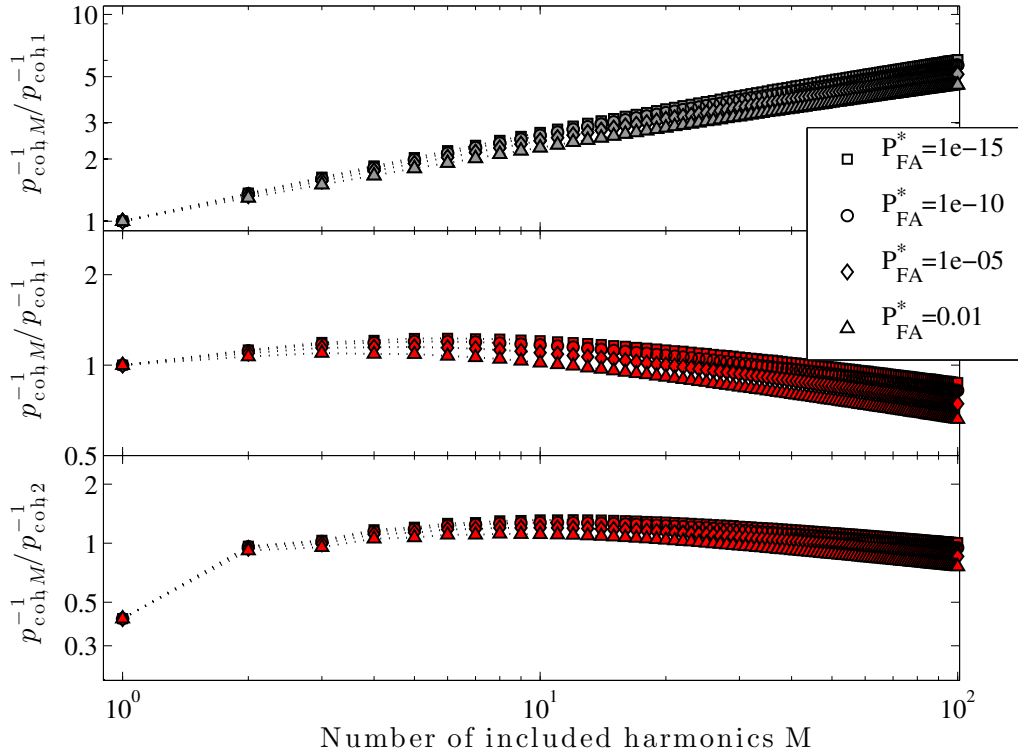


Figure 2.1: Sensitivity as a function of the number of harmonics M included in absence of computational constraints and for three different pulse-profile models. In each panel, we fixed the detection probability $P_{\text{DET}}^* = 90\%$ and the four curves correspond to different values of false alarm probability P_{FA}^* as shown by the legend. The upper panel is for a Dirac delta function pulse profile (implying equal Fourier power at all harmonics). The middle panel is for a typical pulse profile, obtained from the known gamma-ray pulsars by averaging those profiles that are mostly single-peaked (i.e. the γ_n values shown in the bottom left panel in Figure 2.2). The bottom panel is also for a realistic pulse profile, obtained from the known gamma-ray pulsars by averaging those profiles that are mostly two-peaked (i.e. the γ_n values shown in the bottom right panel in Figure 2.2). Since for these profiles the Fourier power $|\gamma_2|^2$ is highest at the second harmonic ($n = 2$), in this plot the vertical axis shows the sensitivity compared to a blind search which would report the highest detection significance at the second harmonic (i.e. “misidentify” the fundamental).




Figure 2.2: Harmonic analysis of pulse profiles of the 117 gamma-ray pulsars in the second *Fermi* LAT pulsar catalog [80]. *Top panel:* Fourier power $|\gamma_n|^2$ (color-coded) at the n th harmonic (vertical axis) for each of the 117 pulsars (horizontal axis). *Bottom left panel:* Fourier power $|\gamma_n|^2$ at the n th harmonic averaged over the 76 out of the 117 pulsars, whose power at the fundamental is highest (mostly single-peaked profiles). *Bottom right panel:* Fourier power $|\gamma_n|^2$ at the n th harmonic averaged over 41 out of the 117 pulsars, whose power at the second harmonic is highest (mostly two-peaked profiles).

for computationally inexpensive targeted searches. In contrast, blind searches are limited by computational power. Therefore, in the following section, we will revisit the efficiency of harmonic summing under the constraint of a fixed computational cost.

2.3.3 Grid-point Counting for Coherent Search

In blind searches, the pulsar's rotational and positional parameters are unknown a priori. Therefore, one has to construct a grid in the multidimensional search parameter space that is explicitly searched, i.e., the test statistic is to be computed at each grid point. Therefore the question arises: What is the most efficient scheme for constructing the search grid? If grid points are placed too far apart potential pulsar signals might be missed. On the other hand, it is highly inefficient to place grid points too closely together, because of redundancy resulting from strongly correlated nearby grid points. The problem of constructing efficient search grids has been intensively studied in the context of gravitational-wave searches [see e.g. 128, 148, 177–179] and we employ some of these concepts here.

The key element is a distance *metric* on the search space [180, 181]. The metric provides an analytic geometric tool measuring the expected fractional loss in squared S/N for any given pulsar-signal location at a nearby grid point.

Let the vector \mathbf{u}_{sig} collect the actual pulsar signal parameters. In a blind search for isolated pulsars, this vector is at least four-dimensional, $\mathbf{u}_{\text{sig}} = (f_{\text{sig}}, \dot{f}_{\text{sig}}, \alpha_{\text{sig}}, \delta_{\text{sig}})$. For simplicity, we begin by considering the metric at the fundamental harmonic ($n = 1$). As will be shown below, it is subsequently straightforward to generalize the results to higher harmonic orders. Following Equation (2.15), let $\theta_{\mathcal{P}_1}(\mathbf{u}_{\text{sig}})$ denote the S/N for the perfect-match case, i.e., at the signal parameter-space location. In a blind search the signal parameters generally will not coincide with a grid point \mathbf{u} , but will typically have some offset,

$$\Delta\mathbf{u} = \mathbf{u} - \mathbf{u}_{\text{sig}}. \quad (2.31)$$

These offsets lead to a (time-dependent) residual phase $\phi(t; \mathbf{u}) - \phi(t; \mathbf{u}_{\text{sig}})$ and therefore a fractional loss in squared S/N results, which is commonly referred to as *mismatch*,

$$m(\Delta\mathbf{u}) = 1 - \frac{\theta_{\mathcal{P}_1}^2(\mathbf{u})}{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}})} = 1 - \frac{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u})}{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}})}. \quad (2.32)$$

The metric is obtained from a Taylor expansion of the mismatch to second order in the offsets $\Delta\mathbf{u}$ at the signal location \mathbf{u}_{sig} ,

$$m(\Delta\mathbf{u}) \approx \sum_{k,\ell} G_{k\ell} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta u^3), \quad (2.33)$$

This equation defines a positive definite metric tensor G with components $G_{k\ell}$, where k and ℓ label the tensor indices. In Appendix 2.8.2, we derive explicit expressions for the coherent metric for a simplified phase model that is appropriate for the purpose of grid construction. We also find that the

resulting metric tensor G is diagonal, which greatly simplifies the grid construction. The results of this derivation will therefore be used in what follows.

As noted by Prix & Shaltev [147], the probability distribution of signal mismatches in a given search grid constructed with a certain maximal mismatch m depends on the structure and dimensionality of the search parameter space. The corresponding average mismatch in each dimension, ξm , will generally be smaller by a characteristic geometric factor $\xi \in (0, 1)$, depending on the actual search-grid construction. For example, for hyper-cubical lattices, ξ is known to be $\xi = 1/3$. In order to construct a hyper-cubical grid in which the maximum mismatch due to an offset in each parameter is m , then the grid point spacing in each parameter should be,

$$\Delta u^k = 2\sqrt{\frac{m}{G_{kk}}}. \quad (2.34)$$

Denote by \mathcal{U} the four-dimensional parameter space, spanned by \mathbf{u} , which is to be searched. Thus, when searching for pulsars with spin frequencies in the range $[0, f_{\max}]$, with spin-down rates in the range $[\dot{f}_{\max}, 0]$, and whose sky location is confined by the LAT to a region of area A_{sky} , the coordinate volume U can be written as

$$\mathcal{U} = f_{\max} \left| \dot{f}_{\max} \right| A_{\text{sky}}. \quad (2.35)$$

In principle, the metric coefficients (and hence also the grid point spacings) can vary throughout the parameter space. Indeed, for the metrics considered in this work, the grid point spacing in the sky dimensions depends on the spin frequency of the pulsar. In order to avoid having to construct a separate sky grid for each search frequency value, we adopt the conservative approach of using the highest frequency searched f_{\max} for the sky grid construction. The metric (and hence also the grid point spacing) becomes uniform throughout \mathcal{U} . The total number of search-grid points $\mathcal{N}_{\text{coh},1}$ for a coherent blind search over \mathcal{U} is therefore simply the product of the number of grid points in each dimension.

$$\mathcal{N}_{\text{coh},1} = \mathcal{U} \prod_k \frac{1}{\Delta u^k} = \frac{1}{16} \mathcal{U} m^{-2} \sqrt{\det G}, \quad (2.36)$$

as G is found to be diagonal. In Appendix 2.8.2 we derive that

$$\sqrt{\det G} = \frac{\pi^4}{\sqrt{135}} T_{\text{coh},1}^3 f^2 r_E^2 \Psi(T_{\text{coh},1}), \quad (2.37)$$

where we defined,⁴

$$\begin{aligned} \Psi^2(T_{\text{coh},1}) = & \left[1 + \text{sinc}(\Omega_E T_{\text{coh},1}/\pi) - 2 \text{sinc}^2(\Omega_E T_{\text{coh},1}/2\pi) \right] \\ & \times \left[1 - \text{sinc}(\Omega_E T_{\text{coh},1}/\pi) \right], \end{aligned} \quad (2.38)$$

and where we have denoted the Earth's orbital angular frequency as $\Omega_E = 2\pi/1\text{yr}$, and the light travel-time from the Earth to the SSB as $r_E = 1\text{AU}/c \sim 500\text{s}$.

⁴We use the definition $\text{sinc}(x) = \sin(\pi x)/(\pi x)$ throughout this manuscript.




Figure 2.3: Scaling of the determinant of the coherent metric G as function of the coherent integration time $T_{\text{coh},1}$ (black solid curve). The red dot-dashed curve shows the model of the coherent metric determinant from the approximation of Equation (2.39) used to estimate the computing cost scaling.

To analytically study the scaling of $\mathcal{N}_{\text{coh},1}$ as a function of $T_{\text{coh},1}$, the function $\Psi(T_{\text{coh},1})$ can be well approximated by

$$\Psi(T_{\text{coh},1}) \approx \begin{cases} \frac{\Omega_E^3 T_{\text{coh},1}^3}{12\sqrt{15}}, & T_{\text{coh},1} < 0.572\text{yr} \\ 1, & T_{\text{coh},1} \geq 0.572\text{yr} \end{cases}. \quad (2.39)$$

The validity of this approximation is illustrated in Figure 2.3. Hence, the total number of grid points required in a coherent search is

$$\mathcal{N}_{\text{coh},1} = \frac{\pi^4}{48\sqrt{15}} \left(\frac{\Omega_E^3}{12\sqrt{15}} \right)^{(a-3)/3} r_E^2 m^{-2} f_{\text{max}}^2 T_{\text{coh},1}^a \mathcal{U}, \quad (2.40)$$

where

$$a \approx \begin{cases} 6, & T_{\text{coh},1} < 0.572\text{yr} \\ 3, & T_{\text{coh},1} \geq 0.572\text{yr} \end{cases}. \quad (2.41)$$

Equation (2.40) tells us that for coherent integration times much shorter than half a year the sky metric components also still scale with $T_{\text{coh},1}$, such that $\mathcal{N}_{\text{coh},1}$ increases approximately as $T_{\text{coh},1}^6$. After half a year of coherent integration the sky metric components quickly approach the resolution saturation as the maximum baseline (1 AU) is reached, and thereafter become approximately independent of $T_{\text{coh},1}$. Therefore $\mathcal{N}_{\text{coh},1}$ scales only as $T_{\text{coh},1}^3$ in this regime.

2.3.4 Coherent Search Sensitivity at Fixed Computing Cost

For computational efficiency, we use the fast Fourier transform (FFT) algorithm [182] to scan the f -dimension. There are two steps involved in calculating an FFT, each with an associated computational cost. Firstly, it is necessary to construct a discrete time series by interpolating (e.g. by binning) the photon arrival times into equidistant samples. The cost of this step is proportional to the number of photon arrival times which must be interpolated. Secondly, the discrete time series must be transformed into a discretely sampled frequency spectrum, using the FFT algorithm. For a maximum frequency of f_{\max} , and a coherent integration time of $T_{\text{coh},1}$ there are $f_{\max}T_{\text{coh},1}$ frequency samples, and the computational cost of calculating the FFT is proportional to $f_{\max}T_{\text{coh},1} \log_2(f_{\max}T_{\text{coh},1})$. We assume that the cost of calculating the FFT is much larger than the cost of creating the discrete time series. Compared to the cost of computing \mathcal{P}_1 explicitly for N photon times at $f_{\max}T_{\text{coh},1}$ frequencies, which is proportional to $Nf_{\max}T_{\text{coh},1}$, it is clear that the FFT method offers more efficiency provided $N \gg \log_2(f_{\max}T_{\text{coh},1})$.

The spacing of frequency samples output by the FFT is $1/T_{\text{coh},1}$. According to the metric [see Equation (2.121a)] this implies a worst-case mismatch due to frequency offsets of $m = G_{ff}/(4T_{\text{coh},1}^2) = \pi^2/12 = 0.82$, which obviously also leads to a high average mismatch. However, as we will discuss in Section 2.5.2, it is possible to reduce this mismatch at almost no extra computational cost by interpolating the frequency spectrum. In the following derivations, we therefore separate the total mismatch m_{tot} into two components: a constant mismatch due to the frequency spacing, m_f determined by the interpolation method used, which has a negligible effect on the overall computing cost; and the mismatch due to offsets in the remaining parameters, m , which can be freely varied to construct an optimal grid.

For every grid point in $\{f, \alpha, \delta\}$ an FFT must be computed, and hence the overall computation time for the search is simply the cost of calculating one FFT multiplied by the number of FFTs that must be computed. The total cost, $C_{\text{coh},1}$ (measured in units of time), is

$$C_{\text{coh},1} = K_{\text{FFT}} f_{\max} T_{\text{coh},1} \log_2(f_{\max} T_{\text{coh},1}) \frac{\mathcal{N}_{\text{coh},1}}{\mathcal{N}_f}, \quad (2.42)$$

where K_{FFT} is an implementation and computing hardware dependent constant, and where \mathcal{N}_f is the number of frequency samples that would be calculated using a grid with an arbitrary maximum mismatch per dimension of m ,

$$\mathcal{N}_f = \frac{f_{\max}}{2} \sqrt{\frac{G_{ff}}{m}} = \frac{\pi}{2\sqrt{3}m} f_{\max} T_{\text{coh},1}. \quad (2.43)$$

The total computational cost is therefore

$$C_{\text{coh},1} = K_{\text{coh},a} m^{-3/2} T_{\text{coh},1}^a \log_2(T_{\text{coh},1} f_{\max}), \quad (2.44)$$

where the constant $K_{\text{coh},a}$ depends on a ,

$$K_{\text{coh},a} = K_{\text{FFT}} \frac{\pi^3 r_E^2 f_{\max}^2 \mathcal{U}}{24\sqrt{5}} \left(\frac{\Omega_E^3}{12\sqrt{15}} \right)^{(a-3)/3}. \quad (2.45)$$

For a search grid constructed with maximum mismatch $m_{\text{tot}} = m_f + 3m$, the search sensitivity will scale with the average mismatch $\langle m_{\text{tot}} \rangle = \langle m_f \rangle + 3\xi m$ as $\sqrt{1 - \langle m_{\text{tot}} \rangle}$ [147]. Thus, from Equation (2.28) it follows that the search sensitivity without harmonic summing scales as

$$p_{\text{coh},1}^{-1} = \frac{\sqrt{(1 - \langle m_{\text{tot}} \rangle) \mu T_{\text{coh},1}}}{\theta_1^*} |\gamma_1|. \quad (2.46)$$

For a computing cost $C_{\text{coh},1}$, Equation (2.44) can be used to obtain (numerically) the maximum $T_{\text{coh},1}$. Substituting this value of $T_{\text{coh},1}$ in Equation (2.46) finally yields the search sensitivity at the given computational cost.

2.3.5 Efficiency of Harmonic Summing at Fixed Computing Cost

Based on the results of the previous sections, we now investigate the efficiency of incoherent harmonic summing under computational cost constraints. More precisely, we address the question of whether it is more efficient in blind searches to sum M harmonics, or to instead use a longer coherent integration time without harmonic summing at the same computing cost.

Thus, we consider the test statistic Z_M^2 , which incoherently sums Fourier powers \mathcal{P}_n from M higher harmonics. In Appendix 2.8.3 we derive the parameter space metric for the Z_M^2 statistic, denoted by \tilde{G} , and find that $\sqrt{\det \tilde{G}} = r^4 \sqrt{\det G}$, where r represents a refinement factor due to harmonic summing, and G is the metric tensor for \mathcal{P}_1 of Equation (2.37). Therefore, to ensure equal sensitivity throughout the original parameter space⁵ the required number of grid points increases by the factor of r^4 compared to using \mathcal{P}_1 only. The value of $r \geq 1$ depends on the pulse profile γ_n . For a sinusoidal pulse profile ($|\gamma_1| = 1/2$ and $|\gamma_{n>1}| = 0$), obviously $r = 1$ (i.e. no refinement), and for a Dirac delta function ($|\gamma_n| = 1$), one finds $r \sim M$, as derived in Equation (2.128). In principle, one could construct a grid with $r^4 \mathcal{N}_{\text{coh},1}$ points, and calculate and sum M values of \mathcal{P}_n at each point, leading to the cost of a harmonic summing search being simply Mr^4 times greater than that of a coherent search at the fundamental frequency with the same coherent integration time.

In practice, to utilize the efficiency of the FFT, it would be necessary to construct a sub-optimal grid in which the range in f and \dot{f} is extended by a factor of M , and the coherent powers summed appropriately over harmonics. The sky-grid in this case may still be constructed using the refinement factor r , leading to the computing cost being $M^2 r^2$ times $C_{\text{coh},1}$ at the same coherent integration time. While this method may quickly become infeasible due to the amount of memory required, we use this only as a theoretically efficient method to compare to an equally costly search using only the fundamental harmonic power.

We here assume that the small extra cost of actually summing the \mathcal{P}_n is negligible.⁶ The computational expense for incoherent harmonic summing, $C_{\text{coh},M}$, using the Z_M^2 statistic for a

⁵This constraint is imposed to eliminate any detection bias in favor of pulsars with low frequencies and frequency derivatives, allowing for estimates of the true astrophysical pulsar populations.

⁶Note that this makes the computing cost estimate generous in favor of the harmonic summing approach in this comparison.

coherent integration time $T_{\text{coh},M}$ becomes

$$C_{\text{coh},M} = K_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2 \log_2(T_{\text{coh},M} f_{\text{max}} M). \quad (2.47)$$

From Equation (2.27) above, we found that the search sensitivity of incoherent harmonic summing is given by

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{(1 - \langle m_{\text{tot}} \rangle) \mu T_{\text{coh},M}}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (2.48)$$

Hence, to compare the search sensitivities $p_{\text{coh},1}^{-1}$ and $p_{\text{coh},M}^{-1}$ at fixed computing cost, in principle the following steps are required. First, for a given computing cost $C_{\text{coh},1}$, Equations (2.44) and (2.46) provide the corresponding coherence time $T_{\text{coh},1}$ and sensitivity $p_{\text{coh},1}^{-1}$, respectively. Second, by equating $C_{\text{coh},1} = C_{\text{coh},M}$, Equation (2.47) then can be solved (numerically) for $T_{\text{coh},M}$, which finally is used to obtain the sensitivity $p_{\text{coh},M}^{-1}$ from Equation (2.48). It should be noted that in comparing $p_{\text{coh},1}^{-1}$ and $p_{\text{coh},M}^{-1}$ the same values of P_{FA}^* and P_{DET}^* must be assumed. We here also assume the same mismatch m in either case, because as shown in Appendix 2.8.5, the optimal mismatch at fixed computing cost is independent of coherent integration time, number of harmonics summed, and computing power available. Notably, a similar result has been found previously by Prix & Shaltev [147] in the context of gravitational-wave pulsar searches.

In the following, we describe an analytical approximation to the numerical approach above which we show to be sufficiently accurate for typical search setups. This approximation is based on ignoring the slowly varying \log_2 factors in Equations (2.44) and (2.47), such that

$$C_{\text{coh},M} \sim K_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2. \quad (2.49)$$

Then from $C_{\text{coh},1} = C_{\text{coh},M}$, it immediately follows that $T_{\text{coh},M}$ must be shorter by the factor $(M^2 r^2)^{(1/a)}$,

$$T_{\text{coh},M} = T_{\text{coh},1} (M^2 r^2)^{-1/a}. \quad (2.50)$$

We show in Appendix 2.8.4 that the $T_{\text{coh},M}$ obtained from this approximation slightly *overestimates* the sensitivity $p_{\text{coh},M}^{-1}$, while being accurate to within less than about 1% for typical search setups. Using Equation (2.50) to substitute $T_{\text{coh},M}$ in Equation (2.48) one obtains for the ratio of search sensitivities,

$$\frac{p_{\text{coh},1}^{-1}}{p_{\text{coh},M}^{-1}} = \frac{M^{1/4+1/a} r^{1/a} \theta_M^*}{\theta_1^*} |\gamma_1| \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{-1/2}, \quad (2.51)$$

which remarkably is independent of $T_{\text{coh},1}$ and $T_{\text{coh},M}$. This sensitivity ratio $p_{\text{coh},1}^{-1}/p_{\text{coh},M}^{-1}$ of Equation (2.51) is shown in Figure 2.4 and is found to be greater than unity for typical gamma-ray pulsars. Only for unrealistically narrow pulse profiles (i.e. a Dirac delta function), the sensitivity ratio can remain close to or slightly below unity. It also should be pointed out that we obtained these results despite the generous assumptions in favor of the harmonic summing approach. First, we ignored the extra costs of summing the M power values. Second, we neglected the possible




Figure 2.4: Ratio of search sensitivities $p_{\text{coh},1}^{-1}$ (without harmonic summing) and $p_{\text{coh},M}^{-1}$ (summing power from M harmonics) at fixed computational cost. The Z_M^2 test statistic used here, assumes a delta function pulse profile, so optimizing the search grid for this profile leads to the curves shown by the filled squares and circles. The red crosses and pluses are for the same statistic and grid, but where the signal pulse profile is a more typical one (derived from averaging those of the known gamma-ray pulsars of Figure 2.1). The open squares and circles are for the same statistic, but using the same typical signal pulse profile and a grid that is also optimized for that same pulse profile. For each case, the results of two different scalings of the computing cost with $T_{\text{coh},M}^a$ are shown, corresponding to $a = 6$ and $a = 3$ (see text for details). While all points shown are for $P_{\text{DET}}^* = 90\%$ and $P_{\text{FA}}^* = 1\%$, they remain qualitatively similar for lower P_{FA}^* values, too.

extra trials when one would maximize the test statistics over different M . Third, the analytical approximation of Equation (2.50) overestimates the true $T_{\text{coh},M}$ (and hence the sensitivity $p_{\text{coh},M}^{-1}$) as we show by numerical evaluation in Figure 2.11.

Hence the basic moral is clear: For blind searches for isolated gamma-ray pulsars, whose sensitivity is limited by computing power rather than the amount of available data, a more sensitive search strategy is to employ a longer coherence time instead of using incoherent harmonic summing at the same computational cost.

2.4 Semicoherent Test Statistics

The key property of the semicoherent test statistics is that only pairs of photon arrival times (t_j, t_k) whose separation $\tau_{jk} = t_j - t_k$, also called *lag*, is at most T (which is much shorter than T_{obs}) are

combined coherently, otherwise incoherently. Hence, we refer to T as the *coherence window* size and denote by R the ratio of total observational data time span T_{obs} of the semicoherent search and T ,

$$R = T_{\text{obs}}/T. \quad (2.52)$$

Compared to fully coherent methods, this semicoherent approach drastically reduces the computing cost since fewer search grid points are required (due to the lower parameter-space resolution as will be described in Section 2.4.2) at the expense of reduced search sensitivity. In Section 2.4.3 we argue that this tradeoff is a profitable one, because at fixed given computing cost the overall search sensitivity of the semicoherent searches outperform fully coherent searches restricted to data spans shorter than T_{obs} by the computational constraints.

To derive a semicoherent test statistic, notice the (unnormalized) coherent Fourier power from Equation (2.10) for the fundamental frequency (first harmonic) can also be written in the following form,

$$\mathcal{P}_1 \propto \left| \sum_{j=1}^N w_j e^{-i\phi(t_j)} \right|^2 = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j)-\phi(t_k)]}. \quad (2.53)$$

Thus, the semicoherent statistic S_1 is formed by multiplying the terms in the above double sum with a real *lag window* $\hat{W}_T(\tau_{jk})$, such that

$$S_1 = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j)-\phi(t_k)]} \hat{W}_T(\tau_{jk}), \quad (2.54)$$

where the lag window has an effective size T ,

$$\int_{-\infty}^{\infty} \hat{W}_T(\tau) d\tau = T, \quad (2.55)$$

and thus must fall off rapidly outside the interval $[-T/2, T/2]$. Blackman & Tukey [183] were the first to consider power spectral estimators of the form of S_1 , which can be seen as the Fourier transform of the lag-windowed covariance sequence [184]. The semicoherent statistic S_1 is just a more general version of the classic Blackman-Tukey method [183] in spectral analysis, e.g. if the phase model was simply $\phi(t_j) = 2\pi f t_j$ only. Hence, S_1 can also be seen as a local spectral average of \mathcal{P}_1 values over neighboring frequencies weighted according to the frequency response of \hat{W}_T [184].

As outlined in [151], for special forms of the lag window, S_1 can also be obtained by summing time-windowed coherent power from overlapping subsets of data. This implies a lag window that must be always positive semidefinite, because it is formed by the convolution of the time window with itself in this case [184], whereas the more general form as of Equation (2.54) in principle can have arbitrary lag windows.

In general, the choice of lag-window function $\hat{W}_T(\tau)$ has an impact on the sensitivity of the statistic S_1 . In tests with simulated LAT data, for the purpose of pulsation detection we found that

the best sensitivity is provided by the simple rectangular lag window,

$$\hat{W}_T^{\text{rect}}(\tau) = \begin{cases} 1, & |\tau| \leq T/2 \\ 0, & \text{otherwise.} \end{cases} \quad (2.56)$$

which also allows for an efficient implementation as will be described in more detail in Section 2.5. The usage of the rectangular lag window could also be motivated from the following viewpoint. Considering the significant sparseness of the LAT data, typically all pairs of photon times fall at different lags (for any practical sampling time, see Section 2.5.1). Therefore, one could argue that optimally (for minimum variance) all lags (i.e., all photon pairs) should be weighted equally when forming S_1 , which is exactly what $\hat{W}_T^{\text{rect}}(\tau)$ implements. Thus, in the remainder of this manuscript we will keep using the rectangular lag window $\hat{W}_T^{\text{rect}}(\tau)$ to calculate S_1 .

2.4.1 Statistical Properties

To examine the statistical properties of the semicoherent statistic, S_1 , it is useful to rewrite Equation (2.54) as

$$S_1 = \sum_{j=1}^N w_j^2 + 2 \sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \cos[\phi(t_j) - \phi(t_k)] \hat{W}_T^{\text{rect}}(\tau_{jk}). \quad (2.57)$$

Under the null hypothesis, $p = 0$ and assuming $N \gg 1$, we show in Appendix 2.8.6 that S_1 follows a normal distribution, whose first two moments of the noise distribution of S_1 are:

$$E_0[S_1] = \sum_{j=1}^N w_j^2, \quad (2.58)$$

$$Var_0[S_1] = 2 \sum_{j=1}^N \sum_{k=j+1}^N w_j^2 w_k^2 \left[\hat{W}_T^{\text{rect}}(\tau_{jk}) \right]^2, \quad (2.59)$$

Now consider that the photon data contains a pulsed signal (i.e. $p > 0$) with a pulse profile defined by Fourier coefficients γ_n . Then the expectation value of S_1 is obtained as

$$E_p[S_1] \approx E_0[S_1] + 2 E_p \left[\sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \cos(\phi(t_j) - \phi(t_k)) \hat{W}_T^{\text{rect}}(\tau_{jk}) \right]. \quad (2.60)$$

Thus, for S_1 we can identify the amplitude S/N θ_{S_1} as

$$\begin{aligned} \theta_{S_1}^2 &= \frac{E_p[S_1] - E_0[S_1]}{\sqrt{Var_0[S_1]}} \\ &= \frac{\sqrt{2}E_p \left[\sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \cos(\phi(t_j) - \phi(t_k)) \hat{W}_T^{\text{rect}}(\tau_{jk}) \right]}{\sqrt{\sum_{j=1}^N \sum_{k=j+1}^N w_j^2 w_k^2 \left[\hat{W}_T^{\text{rect}}(\tau_{jk}) \right]^2}}. \end{aligned} \quad (2.61)$$

To extract the scalings of the semicoherent S/N θ_{S_1} in terms of the relevant search parameters, we assume hard photon-selection cuts, i.e., binary photon weights, for the remainder of this section. Then Equation (2.57) reduces to

$$S_1 = N + 2 \sum_{j=1}^N \sum_{k=j+1}^N \cos[\phi(t_j) - \phi(t_k)] \hat{W}_T^{\text{rect}}(\tau_{jk}). \quad (2.62)$$

In this case, as derived in Appendix 2.8.6, the first two moments of the noise distribution are

$$E_0[S_1] = N, \quad Var_0[S_1] \approx N^2 R^{-1}. \quad (2.63)$$

We show in Appendix 2.8.6, that for moderately strong signals the first two moments of the distribution of S_1 are approximately given by

$$E_p[S_1] \approx N + p^2 N^2 |\gamma_1|^2 R^{-1}, \quad (2.64a)$$

$$Var_p[S_1] \approx N^2 R^{-1} \left(1 + 2p^2 N |\gamma_1|^2 R^{-1} \right), \quad (2.64b)$$

and the squared S/N of Equation (2.61) becomes

$$\theta_{S_1}^2 \approx p^2 N R^{-1/2} |\gamma_1|^2. \quad (2.65)$$

As shown in Appendix 2.8.6, the probability density function of S_1 can be approximated by a normal distribution with the above expectation values and variances. The sensitivity of a semicoherent search is the lowest threshold pulsed fraction p for a given number of photons N and at given false alarm probability P_{FA}^* to achieve a certain detection probability P_{DET}^* . For a threshold $S_{1,\text{th}}$ the false alarm probability is computed as

$$\begin{aligned} P_{\text{FA}}(S_{1,\text{th}}) &\approx \int_{S_{1,\text{th}}}^{\infty} \mathcal{N}\{S_1; E_0[S_1], Var_0[S_1]\} dS_1 \\ &\approx \frac{1}{2} \text{erfc} \left(\frac{S_{1,\text{th}} - E_0[S_1]}{\sqrt{2 Var_0[S_1]}} \right). \end{aligned} \quad (2.66)$$

Where, in this context, $\mathcal{N}\{X; \mu, \sigma^2\}$ denotes a normal distribution with mean μ and variance σ^2 , and should not be confused with the number of grid-points, $\mathcal{N}_{\text{coh},1}$. We compute the probability of detection using $\text{Var}_p[S_1] \approx \text{Var}_0[S_1](1 + 2p^2 N |\gamma_1|^2 R^{-1})$ as

$$\begin{aligned} P_{\text{DET}}(S_{1,\text{th}}, \theta_{S_1}^2) &\approx \int_{S_{1,\text{th}}}^{\infty} \mathcal{N}\{S_1; E_p[S_1], \text{Var}_p[S_1]\} dS_1 \\ &\approx \frac{1}{2} \text{erfc} \left\{ \left(\frac{S_{1,\text{th}} - E_0[S_1]}{\sqrt{2 \text{Var}_0[S_1]}} - \theta_{S_1}^2 \right) \frac{1}{\sqrt{2 + 4p^2 N |\gamma_1|^2 R^{-1}}} \right\}. \end{aligned} \quad (2.67)$$

The minimum detectable pulsed fraction is obtained by first inverting Equation (2.66) to get $S_{1,\text{th}}(P_{\text{FA}}^*)$, which in a second step is substituted in Equation (2.67) to obtain the threshold S/N $\theta_{S_1}^*$ as

$$\begin{aligned} \theta_{S_1}^* &= \theta_{S_1}(P_{\text{FA}}^*, P_{\text{DET}}^*) \\ &\approx \left[\sqrt{2} \text{erfc}^{-1}(2P_{\text{FA}}^*) - \sqrt{2 + 4p^2 N |\gamma_1|^2 R^{-1}} \text{erfc}^{-1}(2P_{\text{DET}}^*) \right]^{1/2}. \end{aligned} \quad (2.68)$$

Finally, using Equation (2.68) one can convert Equation (2.65) into the threshold pulsed fraction $p_{\text{scoh},1}^{-1}$, determining the semicoherent sensitivity as

$$p_{\text{scoh},1}^{-1} = \frac{\sqrt{N} R^{-1/4}}{\theta_{S_1}^*} |\gamma_1| = \frac{\sqrt{\mu T} R^{1/4}}{\theta_{S_1}^*} |\gamma_1|, \quad (2.69)$$

where we used $N = \mu T R$. This reveals the square-root scaling with the coherence window size T and the expected fourth-root scaling with R of the semicoherent sensitivity. Furthermore using $R = T_{\text{obs}}/T$, we can rewrite the previous equation as

$$p_{\text{scoh},1}^{-1} = \frac{\sqrt{\mu} (T T_{\text{obs}})^{1/4}}{\theta_{S_1}^*} |\gamma_1|. \quad (2.70)$$

As a comparison, recall that the coherent sensitivity as of Equation (2.46), $p_{\text{coh},1}^{-1} \propto \sqrt{T_{\text{coh},1}}$, increases with the square root of the coherent integration time $T_{\text{coh},1}$. Here, Equation (2.70) shows that the semicoherent sensitivity, $p_{\text{scoh},1}^{-1} \propto \sqrt{(T T_{\text{obs}})^{1/2}}$, increases with the square root of the geometric mean of the coherence window size T and the total observation time T_{obs} .

It should be noted that while the semicoherent method allows for the use of short lag-windows, in order to detect pulsations there is the additional requirement that there is at least one pair of pulsed photons which arrive within T of each other. This sets a fundamental lower limit on T . But for typical pulsed fractions and photon arrival rates considered in this work, this lower limit is on the order of only a few hours.

2.4.2 Grid-point Counting for Semicoherent Search

To optimally construct the search grid for the semicoherent statistic S_1 , it is necessary to re-evaluate the appropriate metric on parameter space. Analog to Equation (2.32), we define the mismatch for S_1 as the fractional loss in semicoherent S/N squared,

$$\bar{m} = 1 - \frac{\theta_{S_1}^2(\mathbf{u})}{\theta_{S_1}^2(\mathbf{u}_{\text{sig}})} = 1 - \frac{\theta_{S_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u})}{\theta_{S_1}^2(\mathbf{u}_{\text{sig}})}. \quad (2.71)$$

Expanding the mismatch \bar{m} to second order in the offsets $\Delta\mathbf{u}$ as in Equation (2.33) yields the semicoherent metric tensor \bar{G} ,

$$\bar{m} = \sum_{k,\ell} \bar{G}_{k\ell} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta u^3). \quad (2.72)$$

We derive the components $\bar{G}_{k\ell}$ from the phase model in Appendix 2.8.7 analog to the methods described in [148]. Following the same steps as in Section 2.3.3, we find that \bar{G} is also diagonal and the total number of grid points for a semicoherent step can thus be written as

$$\mathcal{N}_{\text{scoh}} = \frac{1}{16} \mathcal{U} \bar{m}^{-2} \sqrt{\det \bar{G}} \quad (2.73)$$

where \bar{m} here represents the maximum mismatch per dimension used for grid construction. As derived in Appendix 2.8.7, the determinant of the semicoherent metric is

$$\sqrt{\det \bar{G}} = \frac{\pi^4}{12\sqrt{3}} T^3 f^2 r_E^2 R \left[1 - \text{sinc}^2 \left(\frac{\Omega_E T}{2\pi} \right) \right]. \quad (2.74)$$

As in Section 2.3.3, for practical purposes we construct the grid for the highest frequency searched f_{max} in a given frequency band. Thus, we can rewrite Equation (2.73) as

$$\mathcal{N}_{\text{scoh}} = \bar{m}^{-2} \frac{\pi^4}{192\sqrt{3}} T^3 f_{\text{max}}^2 r_E^2 R \left[1 - \text{sinc}^2 \left(\frac{\Omega_E T}{2\pi} \right) \right] \mathcal{U}, \quad (2.75)$$

where the proper search volume \mathcal{U} has been defined previously in Equation (2.35).

To extract the scaling of $\mathcal{N}_{\text{scoh}}$ with T , we use the following approximation⁷,

$$\left[1 - \text{sinc}^2 \left(\frac{\Omega_E T}{2\pi} \right) \right] \approx \begin{cases} \frac{\Omega_E^2 T^2}{12}, & T < 0.551\text{yr} \\ 1, & T \geq 0.551\text{yr}. \end{cases} \quad (2.76)$$

which is illustrated in Figure 2.5. Hence, using $R = T_{\text{obs}}/T$ one finds that the total number of grid

⁷Note that when deriving the semicoherent metric, we assume that $T \ll 1$ yr, which obviously does not hold in the second case here. In fact, if the coherence window is longer than one year twice as many search locations are required. Nevertheless, with current computational restraints, only short lag-windows far shorter than one year are feasible. Again, thanks to L. Nieder for this important note.




Figure 2.5: Scaling of the determinant of the semicoherent metric \bar{G} as function of the coherent window size T (black solid curve). The red dot-dashed curve shows the model for the semicoherent metric determinant from the approximation of Equation (2.76) used to estimate the computing cost scaling.

points in the semicoherent search scales as

$$\mathcal{N}_{\text{scoh}} \propto \bar{m}^{-2} \left(\frac{\Omega_E^2}{12} \right)^{\frac{(s-3)}{2}} T^{s-1} T_{\text{obs}} f_{\text{max}}^2, \quad (2.77)$$

where the exponent s is given by

$$s \approx \begin{cases} 5, & T < 0.551 \text{yr} \\ 3, & T \geq 0.551 \text{yr} \end{cases}. \quad (2.78)$$

2.4.3 Semicoherent Search Sensitivity at Fixed Computing Cost

In analogy to Section 2.3.4, we here adopt a similar model for the computational cost of a semicoherent search, which is proportional to the number of search-grid points $\mathcal{N}_{\text{scoh}}$ needed. We again assume that the FFT algorithm is used to compute S_1 over $f_{\text{max}}T$ frequency bins, and again split the total mismatch \bar{m}_{tot} into the mismatch due to a frequency offset \bar{m}_f , and the mismatch due to offsets in the other parameters \bar{m} . Hence, using Equations (2.77) and (2.52) the semicoherent computing cost model $C_{\text{scoh},1}$ is obtained as

$$C_{\text{scoh},1} = K_{\text{scoh},s} \bar{m}^{-3/2} T^{s-1} T_{\text{obs}} \log_2(T f_{\text{max}}), \quad (2.79)$$

where $K_{\text{scoh},s}$ denotes a constant of proportionality that depends on s ,

$$K_{\text{scoh},s} = K_{\text{FFT}} \frac{\sqrt{2} \pi^3 r_E^2 f_{\text{max}}^2 \mathcal{U}}{96} \left(\frac{\Omega_E^2}{12} \right)^{(s-3)/2}, \quad (2.80)$$

as well as on the implementation and computing-hardware dependent constant K_{FFT} as in Equation (2.45). Analog to Equation (2.44), we here also assume that the FFT algorithm is used, hence the \log_2 factor in Equation (2.79). In Section 2.4.1 we found the sensitivity of the semicoherent search as of Equation (2.70) can be approximately described by

$$p_{\text{scoh},1}^{-1} = \sqrt{(1 - \langle \bar{m}_{\text{tot}} \rangle)} \mu \frac{|\gamma_1|}{\theta_{S_1}^*} T^{1/4} T_{\text{obs}}^{1/4}, \quad (2.81)$$

where $\langle \bar{m}_{\text{tot}} \rangle = \langle \bar{m}_f \rangle + 3\xi \bar{m}$ denotes again the total average mismatch of the search grid.

With the sensitivity and computing-cost model at hand, we can now illustrate the increased efficiency that a semicoherent search offers over a fully coherent search. We compare the sensitivity $p_{\text{scoh},1}^{-1}$ of a semicoherent search with coherence window size T over a data set which in total spans the observational time interval T_{obs} to the sensitivity $p_{\text{coh},1}^{-1}$ of a fully coherent search with coherent integration time $T_{\text{coh},1}$, at the same computational cost: $C_{\text{scoh},1} = C_{\text{coh},1}$. For a given computing cost $C_{\text{scoh},1}$, and observational data set spanning T_{obs} , Equation (2.79) determines T . This value of T can then be used to obtain the sensitivity $p_{\text{scoh},1}^{-1}$ via Equation (2.81). Similarly, as described in Section 2.3.5, the given value of $C_{\text{coh},1}$ determines $T_{\text{coh},1}$ and thus provides the corresponding $p_{\text{coh},1}^{-1}$.

The so-obtained ratio of sensitivities $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ is studied numerically in Figure 2.6 for realistic computational power available, such as *Einstein@Home*. In both cases the optimal mismatch parameters are assumed, which are independent of computing cost (see Appendices 2.8.5 and 2.8.8). As can be seen in the figure, this sensitivity ratio is always greater than unity and increases as T decreases, which is representative of the fact that the sensitivity of a semicoherent search decreases more slowly than that of a coherent search as the available computing power decreases. Whilst this ratio decreases as T (and, therefore, the computing cost) increases, the absolute search sensitivity always increases with T , and so it is still beneficial to use the largest achievable lag-window size T at the available computational power.

Using a simplified approximation for the semicoherent computing cost model of Equation (2.79) allows us to obtain some analytical insight into the ratio $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ at fixed computing cost, similar to what has been done in Section 2.3.5. Ignoring the slowly varying \log_2 term gives the approximate semicoherent computing cost model as

$$C_{\text{scoh},1} \sim K_{\text{scoh},s} \bar{m}^{-3/2} T^{s-1} T_{\text{obs}}. \quad (2.82)$$

With this simplified model, $C_{\text{scoh},1} = C_{\text{coh},1}$ can be rewritten using the approximation of Equation (2.49) as

$$\frac{K_{\text{scoh},s} T_{\text{obs}} T^{s-1}}{\bar{m}^{3/2}} = \frac{K_{\text{coh},a} T_{\text{coh},1}^a}{m^{3/2}}. \quad (2.83)$$

Furthermore, using Equations (2.45) and (2.80) to replace $K_{\text{coh},a}$ and $K_{\text{scoh},s}$, we can rewrite Equation (2.83) as

$$T = \left(\frac{4 \Omega_E \bar{m}^{3/2} T_{\text{coh},1}^6}{5\sqrt{6} m^{3/2} T_{\text{obs}}} \right)^{1/4}. \quad (2.84)$$

where we assume $a = 6$ and $s = 5$, since coherent integration times $T_{\text{coh},1}$ less than half a year will be practically feasible in the near future. This relation can then be used to substitute T in the ratio $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ using Equations (2.81) and (2.46), yielding

$$\frac{p_{\text{scoh},1}^{-1}}{p_{\text{coh},1}^{-1}} \approx 2 \frac{\theta_1^*}{\theta_{S_1}^*} \left(\frac{T_{\text{obs}}}{1\text{yr}} \right)^{1/16} \left(\frac{T_{\text{obs}}}{T_{\text{coh},1}} \right)^{1/8}, \quad (2.85)$$

where we again assumed the optimal mismatch choices for m and \bar{m} (see Appendices 2.8.5 and 2.8.8) that are independent of computational cost. For $a = 6$ and $s = 5$ these are $m_{\text{opt}} = 0.172$ and $\bar{m}_{\text{opt}} = 0.146$. Hence, as *Fermi*-LAT data spans several years (implying $T_{\text{obs}} \gtrsim 1\text{yr}$) and typically $\theta_1^* \gtrsim \theta_{S_1}^*$, the sensitivity ratio of Equation (2.85) exceeds unity in all practically relevant cases. This clearly indicates that at fixed computational cost, a semicoherent blind search is always more sensitive than a fully coherent search over the same parameter space.

2.5 Efficient Implementation of a Multistage Search Scheme

In Section 2.3, we argued that under computational cost constraints, blind fully coherent searches without harmonic summing are more efficient, i.e. can typically achieve higher search sensitivity. In Section 2.4, we showed that at fixed computing cost semicoherent searches are more efficient than fully coherent searches to scan wide parameter space.

These considerations motivate a multistage search strategy, in which the first and by far most computationally expensive stage uses the most efficient method (i.e. a semicoherent search) to explore the entire parameter space. In subsequent stages, the most promising candidates are automatically “followed up” in further, more sensitive steps, ultimately using fully coherent search methods. Since the parameter space relevant for these candidates has been previously narrowed down by the first-stage search, the computing cost constraints are relaxed (i.e. the computing cost of the follow-ups is negligible compared to the overall cost of the first stage of the blind search). Hence then the usage of fully coherent methods offering the highest sensitivity is made possible.

In this multistage search scheme, before statistically significant candidates from the first-stage semicoherent search are followed-up with fully coherent methods, it is advisable to refine the location of each semicoherent candidate by searching, again semicoherently, using a refined grid with a smaller mismatch. We then “zoom in” on each significant candidate by performing a fully coherent search of the local parameter space around the refined location of the semicoherent candidate, using the full observational data time span, T_{obs} . The search-grid construction of each stage is guided by the metric, as described in Appendices 2.8.2, 2.8.7 and 2.8.9.

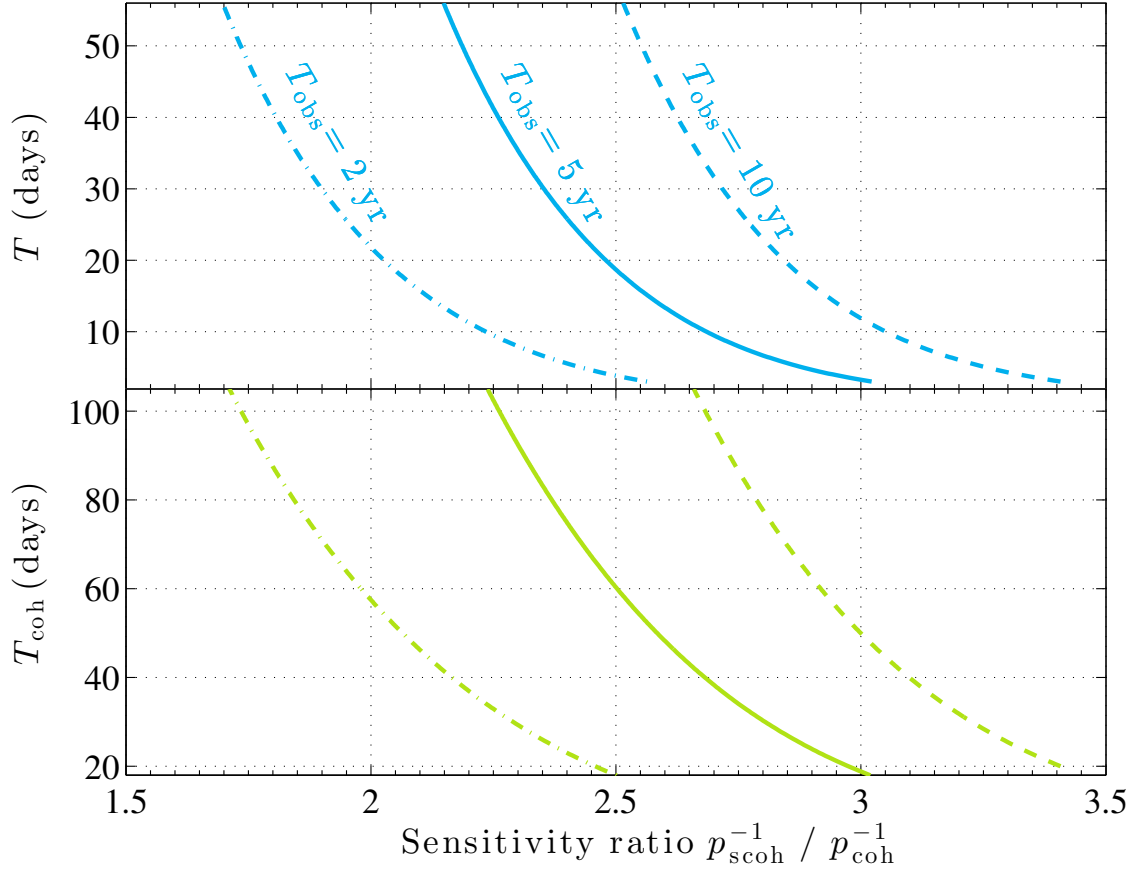


Figure 2.6: Comparison of a semicoherent and fully coherent search sensitivity at equal computing cost and given observational data time spans of $T_{\text{obs}} = 2 \text{ yr}$ (dotted-dashed curves), $T_{\text{obs}} = 5 \text{ yr}$ (solid curves), $T_{\text{obs}} = 10 \text{ yr}$ (dashed curves). The top panel shows the coherence window size T of the corresponding semicoherent search as a function of the sensitivity ratio. The bottom panel shows, for a coherent search, the integration time $T_{\text{coh},1}$, i.e. the subset of T_{obs} that could be fully coherently analyzed with the same computing cost as the semicoherent search with the corresponding T shown in the upper panel. The sensitivity is for $P_{\text{FA}}^* = 10^{-3}$ and $P_{\text{DET}}^* = 0.9$ in each case and a typical pulsed signal with $p = 0.1$ and $|\gamma_1|^2 = 0.35$ (cf. Figure 2.2). Since the sensitivity ratio is in all practically relevant cases much greater than unity, the semicoherent search approach is more efficient.

When searching for weak signals in the presence of noise, this can cause the refined semicoherent candidate to occur at a small but unknown offset from the true signal parameters. This offset depends on the candidate S/N ; candidates with higher S/N have a smaller uncertainty region. In order not to miss weak signals, the coherent follow-up has to cover a conservative region in each dimension around the semicoherent candidate location. Since the parameter space which must be searched coherently has been greatly reduced, this step represents a very small fraction of the overall cost of the search. If the ratio of the coherence window size T used in the first stage and T_{obs} is very large, it is more efficient to insert another intermediate zooming stage that does another semicoherent search with a coherence window size between T and T_{obs} . This would further reduce the parameter space to be searched in the fully coherent step, ensuring that the follow-up remains a negligible fraction compared to overall search. Finally, candidates from this coherent follow-up step are then ranked for further investigation (e.g. by taking into account higher harmonics, or a more complex phase model) according to their false alarm probability.

Since this multistage scheme is designed such that the largest computational burden is associated with the first stage, it is important to optimize this method of calculating the semicoherent test statistic S_1 as much as possible. In the following, we describe various complementary methods which improve the efficiency and sensitivity of a computationally limited semicoherent search.

2.5.1 Efficient Computation of Semicoherent Test Statistic

For each sky-position grid point of the search region the barycentric corrections are applied directly to the LAT-registered arrival times t_{LAT} , to obtain the corresponding photon arrival times t at the SSB. The semicoherent detection statistic S_1 as of Equation (2.54) is then computed over the f - and \dot{f} -ranges. However, directly computing S_1 from Equation (2.54) is computationally inefficient. Therefore, we here discuss more efficient ways of how to do this.

Making the dependence of S_1 on the search parameters f and \dot{f} explicit for clarity, we rewrite Equation (2.57) as

$$S_1(f, \dot{f}) = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j; f, \dot{f}) - \phi(t_k; f, \dot{f})]} \hat{W}_T^{\text{rect}}(\tau_{jk}), \quad (2.86)$$

where the phase differences in terms of f and \dot{f} are given by

$$\begin{aligned} \phi(t_j; f, \dot{f}) - \phi(t_k; f, \dot{f}) &= 2\pi f \tau_{jk} + \pi \dot{f} [(t_j - t_0)^2 - (t_k - t_0)^2] \\ &= 2\pi f \tau_{jk} + \pi \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]. \end{aligned} \quad (2.87)$$

Thus, S_1 of Equation (2.86) takes the following form,

$$S_1(f, \dot{f}) = \sum_{j,k=1}^N w_j w_k e^{-\pi i \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]} \hat{W}_T^{\text{rect}}(\tau_{jk}) e^{-2\pi i f \tau_{jk}}, \quad (2.88)$$

which allows us to utilize the efficiency of the FFT to scan along the f -direction. In the following we describe how to achieve this. First, we construct an equidistant lag series whose separation is the sampling interval $\delta_\tau = 1/(2f_{\max})$, where f_{\max} is equal to the Nyquist frequency f_{Ny} . Then for each pair of times (t_j, t_k) having a lag τ_{jk} smaller than the lag window (i.e. for which $\hat{W}_T^{\text{rect}}(\tau_{jk}) = 1$), we determine the corresponding bin index b of the equidistant lag series via interpolation. While we study the efficiency of different lag-domain-interpolation schemes below, let us assume here nearest-neighbor interpolation for simplicity. Thus, we just round to the nearest lag-bin index b ,

$$b = \text{round} [\tau_{jk}/\delta_\tau] . \quad (2.89)$$

The FFT performance is generally best for input sizes that are a power of 2 (radix-2 FFTs). Therefore, we choose T and f_{\max} , such that the total number of lag bins $B_T = T/\delta_\tau = 2Tf_{\max}$ is a power of 2. We denote the lag-interpolated version of S_1 from Equation (2.88) by \hat{S} , which can be written using the lag-bin index b as

$$\hat{S}(f, \dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y_b(\dot{f}) e^{-2\pi i f \delta_\tau b} , \quad (2.90)$$

where terms depending on \dot{f} and the photon weights have been absorbed into the complex numbers $Y_b(\dot{f})$. More precisely, each $Y_b(\dot{f})$ is the sum of pairwise weight and \dot{f} phase factors, falling into the same lag bin b ,

$$Y_b(\dot{f}) = \sum_{j=1}^N y_j(b; \dot{f}) , \quad (2.91)$$

where

$$y_j(b, \dot{f}) = \sum_{k=j+1}^N \begin{cases} w_j w_k e^{-\pi i \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]} , & \text{round} [\tau_{jk}/\delta_\tau] = b , \\ 0 , & \text{else .} \end{cases} \quad (2.92)$$

Note that the so-constructed lag series Y_b has Hermitian symmetry, i.e. $Y_b = Y_{-b}^*$, and therefore \hat{S} remains entirely real-valued. The above expression for \hat{S} in Equation (2.90) can be seen as a Fourier transform of the complex lag series Y_b , and so \hat{S} can be computed efficiently at many discrete frequencies by exploiting the FFT algorithm, i.e. by calculating

$$\hat{S}_g(\dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y_b(\dot{f}) e^{-2\pi i g b / B_T} . \quad (2.93)$$

where the frequency at the g th bin is $f = g/T$. There exist efficient FFT algorithms [182] which can be used to evaluate this complex-to-real (c2r) transform of Equation (2.93), and which only require the positive lag portion of Y_b to be calculated as an input.

The above formulation of the semicoherent detection statistic, \hat{S}_g , is very similar to the D_ℓ statistic, described in A06 as the DFT of the discrete autocorrelation function of the (binned) photon

arrival times. However, there are some key differences. While further differences are discussed in the following subsections as we encounter them, we here note a first difference between the methods related to the correction of the frequency derivative \dot{f} . When calculating D_ℓ , the frequency derivative is corrected by constructing a new time series in which the photon arrival times are stretched out according to $t_j = \tilde{t}_j + \frac{1}{2} \frac{\dot{f}}{f} \tilde{t}_j^2$. In order to search the $\{f, \dot{f}\}$ parameter space, the ratio \dot{f}/f is increased by small increments. According to this scheme, the search points in the $\{f, \dot{f}\}$ plane lie along straight lines with increasing gradient, intersecting at the origin. As a result, the search grid point density is highly non-uniform in the $\{f, \dot{f}\}$ plane, decreasing from low to high search frequencies. The result is that the search parameter space is highly oversampled in the \dot{f} dimension at low frequencies. This sub-optimal grid-point density implies that far more grid points are needed to cover the parameter space. Decreasing the lag-window size to account for this extra computational cost causes a reduction in sensitivity which more than accounts for the decrease in the average mismatch⁸. Calculating \hat{S}_g in the manner described above, where the effect of the frequency derivative is accounted for by the complex lag-series, $Y_b(\dot{f})$, allows us to uniformly sample the $\{f, \dot{f}\}$ plane with the optimal average mismatch.

2.5.2 Frequency Domain Interpolation

When performing a semicoherent search using \hat{S}_g , computed via the FFT as in Equation (2.93), for a pulsar signal frequency that does not lie exactly at a Fourier frequency (i.e. not at an integer multiple of $1/T$) a loss in signal power (mismatch) will result. To evaluate the response of \hat{S}_g to signals at a non-Fourier frequency, we consider the case when the lag-series contains a pure sinusoid [as in 185], with amplitude \hat{S}_0 , at a frequency h/T . Including an appropriate normalization factor of $1/B_T$ for the Fourier transform, so that

$$Y_b(0) = \frac{\hat{S}_0}{B_T} e^{2\pi i h b / B_T}. \quad (2.94)$$

This represents the (unlikely) case of a strong signal, in the absence of noise, where the frequency derivative and sky location have been perfectly matched. Using Equation (2.93) the response at the g th frequency bin is therefore:

$$\begin{aligned} \hat{S}_g &= \sum_{b=-B_T/2}^{B_T/2} Y_b(0) e^{-2\pi i g b / B_T} \\ &= \frac{\hat{S}_0}{B_T} \sum_{b=-B_T/2}^{B_T/2} e^{-2\pi i b (g-h) / B_T}. \end{aligned} \quad (2.95)$$

⁸This is because despite the reduced mismatch in the \dot{f} dimension, the contributions of the other three dimensions still remain and dominate the total mismatch that is relevant for the search sensitivity.

The above summation over b can be explicitly calculated and is also called the *Dirichlet* kernel, which is given by

$$\mathcal{D}_N(x) = \sum_{b=-N}^N e^{-ibx} = \frac{\sin((N + 1/2)x)}{\sin(x/2)}. \quad (2.96)$$

Using this identity gives rise to rewrite Equation (2.95),

$$\begin{aligned} \hat{S}_g &= \frac{\hat{S}_0}{B_T} \mathcal{D}_{B_T/2}(2\pi(g-h)/B_T) \\ &= \frac{\hat{S}_0}{B_T} \frac{\sin(\pi(g-h)(1-1/B_T))}{\sin(\pi(g-h)/B_T)} \\ &\approx \frac{\hat{S}_0}{B_T} \frac{\sin(\pi(g-h))}{\sin(\pi(g-h)/B_T)} \\ &\approx \hat{S}_0 \operatorname{sinc}(g-h), \end{aligned} \quad (2.97)$$

where in the approximation made in the third step we assumed that $1/2 \gg 1/B_T$, and in the fourth step we used in addition the following approximation $\sin(\pi(g-h)/B_T) \approx \pi(g-h)/B_T$, since typically for nearby frequency bins $B_T \gg (g-h)$. Therefore, the match is well described by a sinc function for signals at non-Fourier frequencies and is smallest (i.e. greatest mismatch) if the signal lies exactly halfway between two Fourier frequencies. This is shown in Figure 2.7, which displays the approximated response of Equation (2.97).

This loss can be reduced by interpolating the Fourier response halfway between two Fourier frequency bins. One method of interpolating the Fourier transform output, known as zero-padding, is to extend the original lag series (or time series) to twice its original length by adding zeros onto the end. However, this requires calculating a Fourier transform which is twice as long, and therefore more than twice as costly. To avoid increasing the computational cost, we use a more efficient interpolation technique in the frequency domain, also known as ‘‘interbinning’’ [185, 186]. Note that [185] gives a formulation for calculating interbin amplitudes for real- or complex-to-complex Fourier transforms. However, in our case, where \hat{S}_g is entirely real-valued, it is sufficient to calculate interbins by summing the amplitude of neighboring frequency bins,

$$\hat{S}_{g+1/2} = \frac{1}{\sqrt{2}} (\hat{S}_g + \hat{S}_{g+1}). \quad (2.98)$$

It is also important to emphasize that our chosen normalization differs from that used by [185, 186], where the interbins are normalized to ensure that all of the signal power is recovered in an interbin if the signal lies exactly halfway between two Fourier bins. Instead, we here use a normalization factor of $1/\sqrt{2}$ ensuring that interbins have the same noise variance as the standard Fourier bins [as was first done by 141]. Whilst the method used in Equation (2.98) results in a mismatch even for signals at the center of an interbin, ensuring that the noise variance is consistent between bins and interbins facilitates semicoherent candidate ranking for follow-up procedures.




Figure 2.7: Illustration of frequency domain interpolation. The dashed blue curve shows the relative response (match) of \hat{S}_g at neighboring Fourier bins as a function of the signal frequency offset. The solid black curve represents the overall DFT response. The overlaid dotted-dashed red curve is the overall DFT response match obtained via the frequency domain interpolation as described in the text.

The overall response for signals at non-Fourier frequencies before and after interbinning is shown in Figure 2.7. Using the interbinning method, the average mismatch due to a frequency offset is reduced from ~ 0.13 to ~ 0.075 , whilst the maximum mismatch is reduced from ~ 0.36 to ~ 0.14 . Thanks to their simplicity, interbins can be calculated very quickly, and so this performance gain comes at negligible extra computing cost (when compared to the dominant FFT computing cost).

2.5.3 Complex Heterodyning

Searching a wide range of frequencies (i.e., large f_{\max}) using the test statistic \hat{S} would require computing a single FFT of large size, B_T . The length of an FFT which can be computed is limited by the amount of memory accessible. In particular, extending the frequency search band to the millisecond pulsar regime (i.e. near 1kHz frequencies) would require a large increase in the sampling rate, and would potentially require decreasing the lag-window size (and hence the sensitivity of the search) to make the FFT short enough to fit into memory.

To address this problem, we divide the total frequency range into smaller bands of size Δf (that can be efficiently searched in parallel) using complex heterodyning, without sacrificing sensitivity. Using this method, the center frequency, f_H , of a given subband is shifted to DC, which in the lag domain corresponds to multiplying each lag bin by $e^{-2\pi i f_H \delta_\tau b}$. The heterodyned lag series is therefore defined as

$$Y'_b(\dot{f}, f_H) = Y_b(\dot{f}) e^{-2\pi i f_H \delta_\tau b}, \quad (2.99)$$

and the frequency at the g th bin becomes

$$f = g/T + f_H. \quad (2.100)$$

One can therefore compute $\hat{S}_g(\dot{f})$ over the subband $[f_H - \Delta f/2; f_H + \Delta f/2]$ via

$$\hat{S}_g(\dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y'_b(\dot{f}, f_H) e^{-2\pi i g b/B_T}, \quad (2.101)$$

in the same way as described in Equation (2.93), but using a sampling interval of only $\delta_\tau = 1/(\Delta f)$. Hence, we can search subbands in the millisecond-pulsar regime, while the FFT size remains at $B_T = T\Delta f$.

2.5.4 Lag Domain Interpolation

As outlined above, before the FFT can be performed the lags τ_{jk} have to be binned into an equidistant lag series. Because the lags τ_{jk} will in general not coincide with the lag-bin centers, the nearest-neighbor interpolation of Equation (2.89) introduces an additional, frequency-dependent loss (mismatch) of signal power across the frequency band analyzed [e.g., 185, 186].

The process of binning in lag can be thought of as convolving the lag series with a binning function. By the convolution theorem, the resulting response across the frequency band is the Fourier transform of this convolving function. For \tilde{S}_g as derived above, the binning function (for nearest-neighbor interpolation) is a simple rectangular function of width δ_τ , leading to the sinc response in the frequency domain. As a consequence, this results in an average loss (mismatch) in signal power of $\sim 13\%$ across the entire search band, illustrated in Figure 2.8.

Improved lag domain interpolation can reduce these losses. A given frequency response can be achieved by weighting the lag series bins around each τ_{jk} with an appropriate interpolation function. Ideal (i.e. lossless) interpolation would lead to a frequency response that is a rectangular function: unity within the search band to remove all bias in the spectrum, and zero outwith to prevent any noise from being aliased into the band. Therefore, this ideal case of a rectangular frequency response requires a lag interpolation function that is the sinc function. However, this interpolation function has infinite extent in the lag domain and is therefore impossible to realize in practice.

A practical solution is to truncate the sinc function in the lag domain around each τ_{jk} , such that the computational cost of this interpolation remains a negligible fraction of the overall computation time. In fact, one can show that using lag domain interpolation with the sinc function truncated to only the d nearest lag bins for each τ_{jk} is the best d th order approximation in the least squares sense to the ideal (rectangular) response function [e.g., 187]. As a result, the average loss (mismatch) across the frequency search band is drastically reduced. In the example shown in Figure 2.8, with a truncated sinc kernel using the $d = 15$ nearest lag bins on either side reduces this average mismatch to only $\sim 1\%$, as compared to the nearest-neighbor interpolation. Generally, it is often practical to use even more neighboring bins without significantly affecting the computational cost, but reducing the average mismatch even further.

However, as can also be seen in Figure 2.8, an inconvenient property of the truncated sinc kernel is the Gibbs oscillation throughout the frequency band. These oscillations mean that the false alarm probabilities of candidates can vary significantly across the frequency band, making it difficult to rank candidate pulsars for follow-up. This problem can be mitigated by multiplying the sinc kernel by another windowing function [188, p. 176]. This windowing function is required to be simple (and therefore efficient) to compute, and must still have a reasonably sharp fall-off in frequency near the edges of the bands. We find that the Welch window [an inverted parabola, 189]) provides a useful compromise between these requirements. The interpolated lag series, \tilde{Y}_b , is constructed by spreading the original lag-series Y'_b amongst the first d bins on either side of the nearest bin to a single photon pair with lag τ_{jk} ,

$$\tilde{Y}_{b+l}(\dot{f}, f_H) = Y'_b(\dot{f}, f_H) \operatorname{sinc}\left(b+l - \frac{\tau_{jk}}{\delta_\tau}\right) \left[1 - \left(b+l - \frac{\tau_{jk}}{\delta_\tau}\right)^2 \frac{1}{d^2}\right], \quad (2.102)$$

for $l = 0, \pm 1, \dots, \pm d$. The frequency response of the Welch-windowed sinc kernel is displayed Figure 2.8. Whilst the average mismatch with the Welch-windowed sinc kernel is comparable to the truncated sinc kernel, the reduced Gibbs oscillation means that the false alarm probabilities of candidates are much more consistent across the frequency band, allowing candidate pulsars to




Figure 2.8: Upper panel: comparison of different lag domain interpolation functions, with the interpolating bin weights calculated over the range of the nearest 15 lag bins on either side of the center. For clarity the inset shows a zoom of the central region. Lower panel: overall frequency response of each interpolation function from the upper panel, that \hat{S} is multiplied with in the frequency domain. The ideal response would be unity within the search band, i.e. for $-f_{Ny} < (f - f_H) < f_{Ny}$, and zero otherwise. For this specific choice of using the nearest 15 bins for the interpolation, the average mismatch (loss in signal power) across the search band from is $\sim 23\%$ for the rectangular binning function in time, $\sim 13\%$ for the rectangular binning function in lag domain, but only $\sim 1\%$ for the sinc kernel and also for the Welch-windowed sinc kernel that shows reduced Gibbs oscillations.

be more easily ranked, albeit with almost no increase in the cost of interpolating the lag-series. Fortunately, the interpolation functions can be efficiently computed using trigonometric look-up tables and recurrence relations. When this efficiency is combined with the typical sparseness of the lag-series, the interpolation step remains a negligible fraction of the overall computation time.

Within this framework of lag domain interpolation, another key difference to the A06 method is worth pointing out. In A06, the SSB photon arrival times t_j are binned directly *prior* to calculating the lags τ_{jk} and the DFT (the D_ℓ in their notation). This implies a rectangular window function in time, which then is convolved with itself leading to a triangular window shape in the lag domain. Hence, the resulting frequency response is effectively the sinc function squared (also shown in Figure 2.8). This causes significant loss in signal power, especially at the edges of the frequency band, and amounts to a loss of $\sim 23\%$ averaged across the entire frequency band. For comparison, by using the lag domain interpolation technique with the Welch-windowed sinc kernel as presented above, this average loss can be reduced by more than an order of magnitude, from $\sim 23\%$ to $\sim 1\%$, at about the same computational expense.

2.6 Performance Demonstration

In order to validate the expected sensitivity gain from the improved methods presented in this paper, we perform extensive Monte-Carlo simulations. The false alarm probabilities are obtained using simulated data sets with different realizations of 8000 photon arrival times (with unit weights), spanning a realistic observation time of $T_{\text{obs}} = 5$ yr. To find the detection probabilities (for a given false alarm probability) simulated pulsar signals are added, which have the same pulse profile of Gaussian shape whose Fourier coefficient at the fundamental frequency is $|\gamma_1| = 0.82$, and varying pulsed fractions p .

While for computational reasons, the actual parameter space searched in each simulation was chosen smaller than in a real search, the main conclusions from these results are unaffected by this. In each simulation, the search covered a frequency bandwidth of 1 Hz and a frequency derivative range of 10^{-13} Hz s $^{-1}$. Each simulation searched the nearest nine sky positions around the signal location, at a uniformly random location on the sky. In the semicoherent search stage we used a coherence window size of $T = 2^{20}$ s ≈ 12 d.

For further comparison, we also apply the A06 method to the simulated data sets. However, here we obtain a generous sensitivity estimation. This is because the non-uniform sampling of the $\{f, \dot{f}\}$ parameter space (discussed in more detail in Section 2.5.1) was not accounted for. While this is justifiable for a search for isolated millisecond pulsars, at lower frequencies and larger frequency derivatives (i.e. where most young pulsars are found) this non-optimal sampling requires reducing the lag-window size (and therefore reducing the sensitivity) to achieve the same computational cost.

The results from all simulations are summarized in Figure 2.9, which shows the detection probability as a function of pulsed fraction for each of the search methods discussed in this paper. From best-fit curves (of typical sigmoid shape) shown in Figure 2.9, we compare the pulsed fraction required to give a detection probability of 95% at a false alarm probability of 0.1%. We find that this

pulsed fraction is around 48% lower for the full multistage method presented here than for the A06 method with approximately the same computational cost. This sensitivity increase is due to several improvements described in previous sections, in particular: use of the parameter space metric to allow optimally spaced grid-points; lag- and frequency-domain interpolation to reduce mismatch; and an automated coherent follow-up step to increase sensitivity to weak gamma-ray pulsar signals.

2.7 Conclusions

We have presented optimized strategies to improve the efficiency of blind searches for isolated gamma-ray pulsars, whose search sensitivity is computationally limited. Under these conditions, our results confirm that fully coherent searches are generally less efficient than semicoherent searches, as well as that harmonic summing is typically less efficient than searching only for the strongest individual harmonic. We also derived the parameters for most efficient search grids. As motivated by these results, we presented and studied the implementation of a multistage search strategy. We have also presented efficient computation and interpolation techniques for the semicoherent test statistic, offering further important sensitivity gains. Finally, we have conducted realistic simulations which demonstrate the improved performance from our combined advances, providing in a substantial increase in sensitivity (i.e. lowering the minimum detectable pulsed fraction by almost 50%) over previous methods at the same computational cost.

The methods presented here are being implemented with the *Einstein@Home* volunteer computing project to increase the chances of detecting new gamma-ray pulsars among the unidentified LAT sources. While here we have focused on searches for isolated pulsars, the methods also apply to searches for pulsars in binaries, where partial knowledge of the orbit is available from observations at other wavelengths [135].

Furthermore, the framework derived in this work in order to obtain an improved understanding of the pulsation search sensitivities underlying the different methods should also be useful for population studies. Specifically, these estimates can facilitate identifying the selection biases in the known gamma-ray pulsar sample, for example due to the difference in pulse profile shape. In future work, we shall also explore using this framework to improve the efficiency of harmonic summing employing one or more realistic pulse profile templates built from the existing population of known gamma-ray pulsars.

Acknowledgments

This work was supported by the Max-Planck-Gesellschaft (MPG), as well as by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether research grant PL 710/1-1 (PI: Holger J. Pletsch). We also thank the anonymous referee for suggestions that helped improve the manuscript.

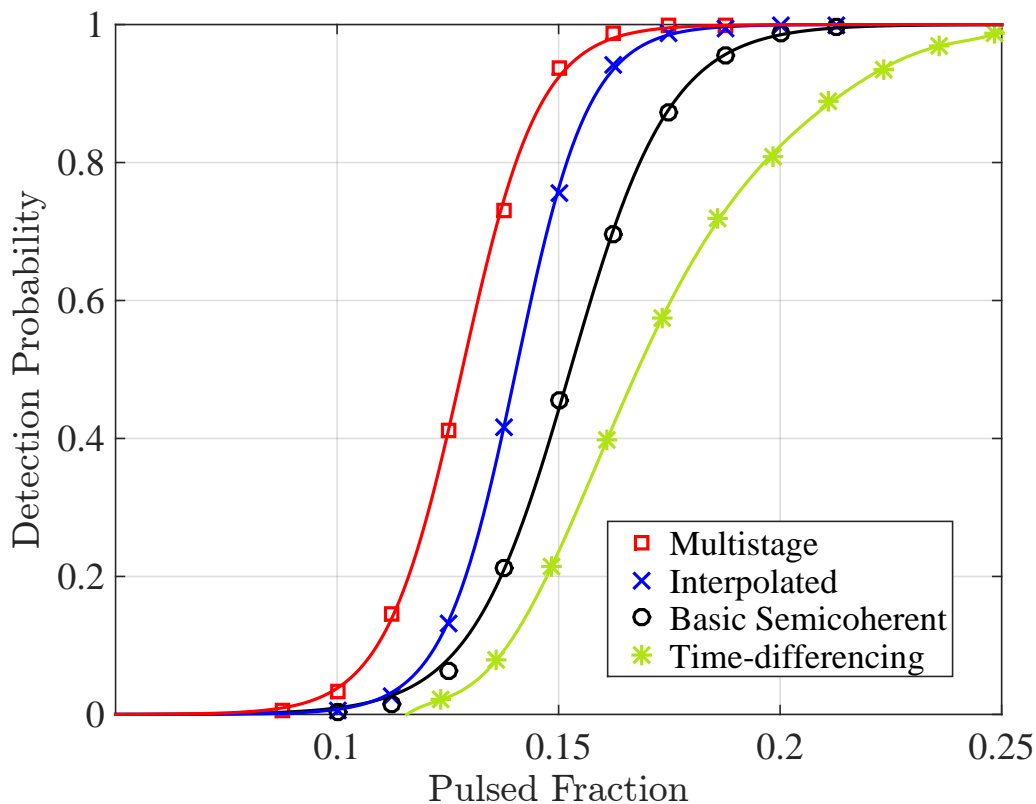


Figure 2.9: Comparison of search efficiency of different search methods at fixed computational cost. Shown is the detection probability P_{DET} at increasing pulsed fractions p for the simulated gamma-ray pulsar signals at $P_{\text{FA}} = 10^{-3}$. The solid curves represent fits to each set of data points. Green stars: estimated sensitivity using the A06 method for the same computing cost. Black circles: semicoherent search method, using only nearest-neighbor lag-domain interpolation. Blue crosses: semicoherent search method, using interbinning frequency-domain interpolation, and lag-domain interpolation with a Welch-windowed sinc kernel. Red squares: multistage search method (Semicoherent search method using lag-domain and interbinning frequency-domain interpolation, plus fully coherent follow-up). In all cases, the number of simulations was chosen large enough so that the uncertainties of the data points become smaller than the size of the data markers.

2.8 Appendix to Chapter 2

2.8.1 Derivation of statistical properties of coherent test statistic

From Equation (2.12) in Section 2.3.1 the coherent power \mathcal{P}_n can be rewritten as $\mathcal{P}_n = c_n^2 + s_n^2$, where

$$c_n = \sqrt{\frac{2}{N}} \sum_{j=1}^N \cos[n\phi(t_j)] \quad (2.103)$$

$$s_n = \sqrt{\frac{2}{N}} \sum_{j=1}^N \sin[n\phi(t_j)]. \quad (2.104)$$

Under the null hypothesis $p = 0$, the phases $\phi(t_j)$ are uniformly distributed on $[0, 2\pi]$ and it is straightforward to show that

$$E_0[\cos(n\phi(t_j))] = E_0[\sin(n\phi(t_j))] = 0, \quad (2.105a)$$

$$Var_0[\cos(n\phi(t_j))] = Var_0[\sin(n\phi(t_j))] = 1/2. \quad (2.105b)$$

Since we have typically $N \gg 1$, by appealing to the Central Limit Theorem, the random variables c_n and s_n are normally distributed with zero mean and unit variance,

$$E_0[c_n] = E_0[s_n] = 0, \quad (2.106a)$$

$$Var_0[c_n] = Var_0[s_n] = 1. \quad (2.106b)$$

Hence, \mathcal{P}_n follows a central χ^2 -distribution with 2 degrees of freedom [e.g., 183]. Therefore the first two moments are $E_0[\mathcal{P}_n] = 2$ and $Var_0[\mathcal{P}_n] = 4$, as given in Equation (2.13).

Suppose a pulsed signal is present, $p > 0$, with a pulse profile having the complex Fourier coefficients γ_n as defined by Equation (2.4). While in this case for the $(1-p)N$ “non-pulsed” photons (i.e. background) Equations (2.105) still hold, however for the pN “pulsed” photons (i.e. not background) one obtains

$$E_p[\cos(n\phi(t_j))] = \Re(\gamma_n), \quad (2.107a)$$

$$E_p[\sin(n\phi(t_j))] = -\Im(\gamma_n), \quad (2.107b)$$

$$Var_p[\cos(n\phi(t_j))] = \frac{1}{2} + \frac{\Re(\gamma_{2n})}{2} - \Re(\gamma_n)^2, \quad (2.107c)$$

$$Var_p[\sin(n\phi(t_j))] = \frac{1}{2} - \frac{\Re(\gamma_{2n})}{2} - \Im(\gamma_n)^2. \quad (2.107d)$$

Therefore, the random variables c_n and s_n are normally distributed (since $N \gg 1$) with the following

mean values and variances,

$$E_p[c_n] = p \sqrt{2N} \Re(\gamma_n), \quad (2.108a)$$

$$E_p[s_n] = -p \sqrt{2N} \Im(\gamma_n), \quad (2.108b)$$

$$\text{Var}_p[c_n] = 1 + p \Re(\gamma_n^2) - 2p \Re(\gamma_n)^2, \quad (2.108c)$$

$$\text{Var}_p[s_n] = 1 - p \Re(\gamma_n^2) - 2p \Im(\gamma_n)^2. \quad (2.108d)$$

For weak signals (i.e. small pulsed fractions) and typical gamma-ray pulse profiles (see Figure 2.2), we can approximate these variances as

$$\text{Var}_p[c_n] \approx \text{Var}_p[s_n] \approx 1. \quad (2.109)$$

With this approximation, the distribution of \mathcal{P}_n follows a noncentral χ^2 -distribution [174, 175] with 2 degrees of freedom, whose the first two moments are

$$E_p[\mathcal{P}_n] \approx 2 + 2p^2 N |\gamma_n|^2, \quad (2.110a)$$

$$\text{Var}_p[\mathcal{P}_n] \approx 4 + 8p^2 N |\gamma_n|^2, \quad (2.110b)$$

recovering Equations (2.14a) and (2.14b). The noncentrality parameter of that distribution is the second summand in Equation (2.110a), $2p^2 N |\gamma_n|^2$.

2.8.2 Coherent Metric

For the purpose of efficient search-grid construction we exploit a simplified phase model which captures the most dominant effects. It is to be emphasized that we do *not* use this phase model in the actual search when computing the phases at the photon arrival times. Thus, we here assume that the LAT data set spans at least one year, such that the Doppler modulation is dominated by the Earth motion around the SSB.

For very short coherent integration times, the orbital motion of the *Fermi* satellite around the Earth could also introduce further Doppler effects. Comparing this effect to the much larger effect of the Earth's orbital motion around the sun, which is responsible for the behavior of the metric visible in, e.g., Figure 2.3, it is clear that this effect would saturate after a small number of orbits. Hence for coherent integration times of more than a few hours, here it is safe to neglect the rapidly oscillating components of the motion of the *Fermi* satellite around the Earth. Doing so yields the following phase model,

$$\begin{aligned} \phi(t, \mathbf{u}) &= 2\pi f(t - t_0) + \pi \dot{f}(t - t_0)^2 + 2\pi f \frac{\vec{n} \cdot \vec{r}_E(t)}{c} \\ &= 2\pi f(t - t_0) + \pi \dot{f}(t - t_0)^2 \\ &\quad + 2\pi f r_E [n_x \cos(\Omega_E t) + n_y \sin(\Omega_E t)], \end{aligned} \quad (2.111)$$




Figure 2.10: Comparison of mismatch in \mathcal{P}_1 (dashed curves) with coherent metric prediction (solid curves). In each panel the horizontal axis shows the offset from the signal parameters in f (left), \dot{f} (middle), and sky position (right). The sky-location offset is $\sqrt{\Delta n_x^2 + \Delta n_y^2}$, which measures the offset in coordinates (n_x, n_y) in the ecliptic plane. The underlying pulsar signal has been simulated with spin parameters $f = 32$ Hz, $\dot{f} = -10^{-12}$ Hz s $^{-1}$ for a total coherent observation time of $T_{\text{coh}} = 3.4$ yr.

where n_x and n_y are the components of \vec{n} , the unit vector pointing from the SSB to the sky location (α, δ) , projected into the ecliptic plane (using the obliquity of the ecliptic, ϵ),

$$n_x = \cos(\alpha) \cos(\delta), \quad (2.112)$$

$$n_y = \cos(\epsilon) \sin(\alpha) \cos(\delta) + \sin(\epsilon) \sin(\delta), \quad (2.113)$$

and $\Omega_E = 2\pi/1\text{yr}$, and $r_E = 1\text{AU}/c \sim 500\text{s}$.

In the presence of a small offset $\Delta\mathbf{u}$ from a signal's location in parameter space \mathbf{u}_{sig} , we can write the mismatch, $m^{[t_j]}$, in the coherent power in a window of length T , centered on the j th photon as

$$m^{[t_j]} = 1 - \frac{(\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u}))^{[t_j]}}{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}})} \quad (2.114)$$

$$= 1 - \left| \langle e^{-i\phi(t, \Delta\mathbf{u})} \rangle^{[t_j]} \right|^2, \quad (2.115)$$

where we have replaced the discrete sum of Equation (2.10) for simplicity by a continuous integral over the coherent integration time T , i.e.,

$$\langle x \rangle^{[t_j]} \equiv \frac{1}{T} \int_{t_j - T/2}^{t_j + T/2} x(t) dt. \quad (2.116)$$

Following the derivation in [148], the mismatch can be Taylor expanded up to second order in terms of the parameter offsets, $\Delta\mathbf{u}^k$ to give

$$m^{[t_j]} = \sum_{k, \ell} G_{k\ell}^{[t_j]} \Delta\mathbf{u}^k \Delta\mathbf{u}^\ell + \mathcal{O}(\Delta\mathbf{u}^3). \quad (2.117)$$

The coherent metric components are defined as

$$G_{k\ell}^{[t_j]} = \langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]} - \langle \partial_k \phi \rangle^{[t_j]} \langle \partial_\ell \phi \rangle^{[t_j]}, \quad (2.118)$$

where $\partial_k \phi$ is the partial derivative of the phase at the signal location with respect to the k th component of the parameter offset:

$$\partial_k \phi \equiv \left. \frac{\partial \phi(t; \mathbf{u}_{\text{sig}} + \Delta\mathbf{u})}{\partial(\Delta\mathbf{u}^k)} \right|_{\Delta\mathbf{u}=\mathbf{0}}. \quad (2.119)$$

Using the simplified phase model of Equation (2.111), the metric components for a coherent window,

centered on t_j are given by

$$G_{ff}^{[t_j]} = \frac{\pi^2 T^2}{3}, \quad (2.120a)$$

$$G_{fj}^{[t_j]} = \frac{\pi^2 T^4}{180} + \frac{\pi^2 (t_j - t_0)^2 T^2}{3}, \quad (2.120b)$$

$$G_{n_x n_x}^{[t_j]} = 2\pi^2 f^2 r_E^2 \left[1 + \text{sinc}(\Omega_E T / \pi) \cos(2\Omega_E t_j) - 2 \text{sinc}^2(\Omega_E T / 2\pi) \cos^2(\Omega_E t_j) \right], \quad (2.120c)$$

$$G_{n_y n_y}^{[t_j]} = 2\pi^2 f^2 r_E^2 \left[1 - \text{sinc}(\Omega_E T / \pi) \cos(2\Omega_E t_j) - 2 \text{sinc}^2(\Omega_E T / 2\pi) \sin^2(\Omega_E t_j) \right]. \quad (2.120d)$$

For the specific case of the general expressions above, where $t_j = t_0 = 0$, the metric components for the coherent detection statistic simplify to the following form,

$$G_{ff} = \frac{\pi^2 T^2}{3}, \quad (2.121a)$$

$$G_{fj} = \frac{\pi^2 T^4}{180}, \quad (2.121b)$$

$$G_{n_x n_x} = 2\pi^2 f^2 r_E^2 \left[1 + \text{sinc}(\Omega_E T / \pi) - 2 \text{sinc}^2(\Omega_E T / 2\pi) \right], \quad (2.121c)$$

$$G_{n_y n_y} = 2\pi^2 f^2 r_E^2 \left[1 - \text{sinc}(\Omega_E T / \pi) \right]. \quad (2.121d)$$

The mismatches predicted by these derived metric components are compared to the measured mismatches in \mathcal{P}_1 for a simulated pulsar signal in Figure 2.10.

Therefore, the determinant of the coherent metric is found as

$$\begin{aligned} \sqrt{\det G} &= \frac{\pi^4}{\sqrt{135}} T^3 f^2 r_E^2 \\ &\times \left[1 + \text{sinc}(\Omega_E T_{\text{coh},1} / \pi) - 2 \text{sinc}^2(\Omega_E T_{\text{coh},1} / 2\pi) \right] \\ &\times \left[1 - \text{sinc}(\Omega_E T_{\text{coh},1} / \pi) \right]. \end{aligned} \quad (2.122)$$

2.8.3 Coherent Metric with Incoherent Harmonic Summing

If a search is performed using the Z_M^2 statistic, i.e., incoherently summing the coherent power \mathcal{P}_n in the first M harmonics, the mismatch, \tilde{m} , becomes

$$\begin{aligned} \tilde{m} &= 1 - \frac{\sum_{n=1}^M \theta_{\mathcal{P}_n}^2(\mathbf{u}_{\text{sig}} + \Delta \mathbf{u})}{\sum_{n=1}^M \theta_{\mathcal{P}_n}^2(\mathbf{u}_{\text{sig}})} \\ &= 1 - \frac{\sum_{n=1}^M |\gamma_n|^2 \left| \langle e^{-in\phi(t, \Delta \mathbf{u})} \rangle^{[t_j]} \right|^2}{\sum_{n=1}^M |\gamma_n|^2}. \end{aligned} \quad (2.123)$$

Taylor expanding this mismatch to second order gives the metric components,

$$\tilde{m} = \sum_{k,\ell} \tilde{G}_{k\ell}^{[t_j]} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta \mathbf{u}^3), \quad (2.124)$$

which can be expressed using Equation (2.118) as

$$\tilde{G}_{k\ell}^{[t_j]} = r^2 G_{k\ell}^{[t_j]}, \quad (2.125)$$

where we defined the harmonic refinement factor r from

$$r^2 = \frac{\sum_{n=1}^M |\gamma_n|^2 n^2}{\sum_{n=1}^M |\gamma_n|^2}. \quad (2.126)$$

Thus, Equation (2.125) indicates that the parameter space must be sampled r times more finely in each dimension when summing the power from M harmonics,

$$\sqrt{\det \tilde{G}} = r^4 \sqrt{\det G}. \quad (2.127)$$

The value of this refinement factor r also depends on the signal pulse profile γ_n , which of course is unknown in advance. However, we can consider the two limiting cases. First, for the narrowest possible pulse profile, a Delta function, all coefficients are equal, $|\gamma_n| = 1$, such that

$$r^2 = \frac{1}{M} \sum_{n=1}^M n^2 = \frac{M^2}{3} + \frac{M}{2} + \frac{1}{6}. \quad (2.128)$$

Therefore, for $M > 1$ the parameter space must be sampled more finely in each dimension by a factor of approximately $M^2/3$ (to leading order). On the other limiting case, for a sinusoidal pulse profile, where $|\gamma_{n>1}| = 0$, $r = 1$ and thus $\tilde{G}_{k\ell}^{[t_j]} = G_{k\ell}^{[t_j]}$, requiring no refinement. Therefore, the range of the harmonic-summing refinement factor is approximately limited to $r \in [1, M]$.

Finally, we would like to point out a further generalization. Suppose a search is performed using the Q_M statistic and a template pulse profile α_n , which is not equal to the Dirac delta function (in this case Q_M would reduce again to Z_M^2). Then by a straightforward repetition of arguments from the beginning of this section one obtains the resulting metric tensor $\hat{G}_{k\ell}^{[t_j]}$ for the Q_M test statistic as

$$\hat{G}_{k\ell}^{[t_j]} = \hat{r}^2 G_{k\ell}^{[t_j]}, \quad (2.129)$$

where the harmonic refinement factor \hat{r} in this case would be different from Equation (2.126), namely

$$\hat{r}^2 = \frac{\sum_{n=1}^M |\gamma_n|^4 n^2}{\sum_{n=1}^M |\gamma_n|^4}. \quad (2.130)$$




Figure 2.11: Comparison of the analytical approximation for the harmonic-summing computing cost model (leading to $T_{\text{coh},M}^{\text{approx}}$) to the results obtained from fully numerical evaluation (leading to $T_{\text{coh},M}^{\text{exact}}$), as a function of $T_{\text{coh},1}$ corresponding to the same computing cost $C_{\text{coh},1} = C_{\text{coh},M}$. Since we are interested in the impact on search sensitivity $p_{\text{coh},M}^{-1} \propto \sqrt{T_{\text{coh},M}}$, the vertical axis shows the square root of the ratio. As indicated by the legend the different curves are for different values of scaling exponent a of Equation (2.41) and number of harmonics summed M .

2.8.4 Approximate Harmonic-summing Computing Cost

In Section 2.3.5, we describe an analytical approximation for the computing cost model of incoherent harmonic summing. This approximation is based on ignoring the slowly varying \log_2 factors in Equations (2.44) and (2.47). If then one equates $C_{\text{coh},1} = C_{\text{coh},M}$, it follows that $T_{\text{coh},M}$ must be shorter by the factor $(M^2 r^2)^{(1/a)}$, as given in Equation (2.50). Here, we study the accuracy of the analytical approximation in terms of the search sensitivity $p_{\text{coh},M}^{-1} \propto \sqrt{T_{\text{coh},M}}$, by comparison to the exact value for $T_{\text{coh},M}$ obtained from numerical evaluation. For a given value of $T_{\text{coh},1}$, we find numerically the exact value of $T_{\text{coh},M}$ such that $C_{\text{coh},1} = C_{\text{coh},M}$. We here assume a wide search frequency range, $f_{\text{max}} = 1000$ Hz. The results are displayed in Figure 2.11, showing that the approximation is accurate to within less than 1% for typical search setups. As can also be seen, for the realistic case of $a = 6$ the approximation is generous in favor of the harmonic summing approach, because $T_{\text{coh},M}^{\text{approx}} \gtrsim T_{\text{coh},M}^{\text{exact}}$, the approximation overestimates the true search sensitivity.

2.8.5 Optimal Mismatch in Coherent Search

In this section, we use the method of Lagrange multipliers as in Prix & Shaltev [147] to obtain the optimal average mismatch for a fully coherent search. We use the scalings of the sensitivity $p_{\text{coh},M}^{-1}$ and computing cost $C_{\text{coh},M}$, ignoring the \log_2 FFT scaling factor, from Equations (2.28) and (2.49), respectively. In order to find the optimal mismatch at a fixed computing cost C_0 , we search for stationary points of the Lagrange function,

$$\begin{aligned} L(T_{\text{coh},M}, m, M, \lambda) &= p_{\text{coh},M}^{-1} - \lambda(C_{\text{coh},M} - C_0) \\ &= (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh},M}^{1/2} h^*(M) \\ &\quad + \lambda \left(K'_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2(M) - C_0 \right), \end{aligned} \quad (2.131)$$

where λ is a Lagrange multiplier, and we defined $K'_{\text{coh},a} = K_{\text{coh},a} f_{\text{max}}^2$, as well as the function $h^*(M)$ as,

$$h^*(M) = \frac{1}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}, \quad (2.132)$$

using * to indicate the implicit dependence on P_{FA}^* and P_{DET}^* through θ_M^* . Taking partial derivatives with respect to $T_{\text{coh},M}$, m and M respectively yields:

$$\frac{\partial L}{\partial T_{\text{coh},M}} = \frac{1}{2} (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh}}^{-1/2} h^*(M) + \frac{a\lambda C_{\text{coh},M}}{T_{\text{coh},M}} = 0, \quad (2.133)$$

$$\frac{\partial L}{\partial m} = \frac{1}{2} (1 - \langle m_{\text{tot}} \rangle)^{-1/2} 3\xi T_{\text{coh}}^{1/2} h^*(M) + \frac{3\lambda C_{\text{coh},M}}{2m} = 0, \quad (2.134)$$

$$\begin{aligned} \frac{\partial L}{\partial M} &= (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh},M}^{1/2} \frac{\partial h^*(M)}{\partial M} \\ &\quad + \lambda C_{\text{coh},M} \left(\frac{2}{M} + \frac{2}{r(M)} \frac{\partial r}{\partial M} \right) = 0. \end{aligned} \quad (2.135)$$

Equating these and rearranging for ξm , we find that the optimal average mismatch for a fully coherent search is

$$3\xi m_{\text{opt}} = \frac{1 - \langle m_f \rangle}{\frac{2a}{3} + 1}. \quad (2.136)$$

As we argue in Section 2.3.4, practical fully coherent searches are computationally limited to integration times $T_{\text{coh},M}$ less than half a year, implying $a = 6$. If the frequency dimension is interpolated using interbinning, $\langle m_f \rangle \approx 0.14$, giving $m_{\text{opt}} = 0.172$ for a total average mismatch of $\langle m_{\text{tot}} \rangle = 0.312$. It is noteworthy that this result is independent of the computational cost, the coherent integration time, and the number of harmonics summed.

In principle, one can also rearrange for M to find the optimal number of harmonics, which then requires solving a complicated differential equation. However, the derivatives of the functions $h^*(M)$ defined in Equation (2.132), and $r(M)$ defined in Equation (2.126) are difficult to obtain

for most pulse profiles. Therefore, we followed the approach presented in Section 2.3.5 to find the optimal M at fixed computing cost, which does not require calculating these derivatives.

2.8.6 Derivation of Statistical Properties of Semicohherent Test Statistic

From Equation (2.54), the expectation value of S_1 can be written as

$$E_0 [S_1] = E_0 \left[\sum_{j,k}^N w_j w_k e^{-i(\phi(t_j) - \phi(t_k))} \hat{W}_T^{\text{rect}}(\tau_{jk}) \right]. \quad (2.137)$$

In order to evaluate this expectation value, we must take into account terms in the double sum where the photon indexes (j, k) are equal, giving

$$E_0 [S_1] = \sum_{j=1}^N w_j^2 \hat{W}_T(0) + \sum_{j \neq k}^N w_j w_k E_0 \left[e^{-i(\phi(t_j) - \phi(t_k))} \right] \hat{W}_T(\tau_{jk}), \quad (2.138)$$

where $\sum_{j \neq k}^N$ denotes a double sum over all photons, excluding terms where $j = k$. Under the null hypothesis, $p = 0$, it holds

$$E_0 \left[e^{-i\phi(t_j)} \right] = E_0 \left[e^{i\phi(t_k)} \right] = 0, \quad (2.139)$$

and hence we find that the expectation value of S_1 is simply

$$E_0 [S_1] = \sum_{j=1}^N w_j^2 \hat{W}_T(0). \quad (2.140)$$

To find the variance of S_1 , we must evaluate

$$E_0 [S_1^2] = E_0 \left[\sum_{j,k,l,m}^N w_j w_k w_l w_m e^{-i(\phi(t_j) - \phi(t_k) + \phi(t_l) - \phi(t_m))} \hat{W}_T(\tau_{jk}) \hat{W}_T(\tau_{lm}) \right]. \quad (2.141)$$

Again, taking into account terms where photon indexes are equal, and using Equation (2.139), we find that

$$E_0 [S_1^2] = \sum_{j=1}^N w_j^4 \hat{W}_T(0)^2 + \sum_{j \neq k}^N w_j^2 w_k^2 \hat{W}_T(0)^2 + \sum_{j \neq k}^N w_j^2 w_k^2 \hat{W}_T(\tau_{jk})^2, \quad (2.142)$$

and hence the variance of S_1 under the null hypothesis is

$$\begin{aligned} \text{Var}_0 [S_1] &= E_0 [S_1^2] - E_0 [S_1]^2 \\ &= \sum_{j \neq k}^N w_j^2 w_k^2 \hat{W}_T(\tau_{jk})^2. \end{aligned} \quad (2.143)$$




Figure 2.12: Comparison of empirical and analytically predicted probability density function (PDF) of the semicoherent test statistic S_1 . The blue curves (left) refer to the noise-only case, where S_1 has been calculated from many simulated data sets of $N = 10^4$ unit-weight photons with $R = 60$ to obtain the empirical PDF (solid curve) which is compared to the analytical PDF (dashed). The green curves represent the PDF of S_1 for simulated data sets containing signals with a pulsed fraction of $p = 0.1$ and a profile with $|\gamma_1|^2 = 0.668$, where again the empirical PDF (solid curve) is compared to the analytical PDF (dashed).

From now on in this section, we will use the rectangular lag-window $\hat{W}_T^{\text{rect}}(\tau_{jk})$ of Equation (2.56). In addition, we assume binary photon weights for simplicity. In this case one obtains

$$E_0[S_1] = N, \quad \text{Var}_0[S_1] \approx N^2 R^{-1}. \quad (2.144)$$

To derive the moments of the distribution of S_1 in the presence of a perfectly-matched signal, we need to distinguish times t_j of non-pulsed photons (i.e., background) from pulsed photons by denoting the latter times as t'_j . We then use the definitions of the Fourier coefficients of the pulse profile to evaluate the expectation values

$$E_p \left[e^{-in\phi(t'_j)} \right] = \gamma_n, \quad E_p \left[e^{in\phi(t'_j)} \right] = \gamma_n^*. \quad (2.145)$$

Evaluating Equations (2.137) and (2.141), using the expectation values from Equations (2.139) and (2.145), with a pulsed fraction, $p \sim \mathcal{O}(10^{-1})$ and a typical pulse profile γ_n (cf. Figure 2.2), gives

the first two moments of the distribution of S_1 in the presence of a weak signal as

$$E_p [S_1] \approx N + p^2 N^2 |\gamma_1|^2 R^{-1}, \quad (2.146)$$

$$Var_p [S_1] \approx \frac{N^2}{R} \left(1 + 2p^2 N |\gamma_1|^2 R^{-1} \right), \quad (2.147)$$

where we have assumed a large number of photons $N \gg 1$, and that R is large enough such that edge effects (e.g., effectively shorter windows near the end of the observational data time span) become negligible.

Again, appealing to the central limit theorem (i.e., assuming that there are many photon pairs within the double sums of Equation (2.137)), we can approximate the distribution of S_1 by a normal distribution with the same mean and variance. By comparison with numerical simulations Figure 2.12 validates this approximation for the purpose of the sensitivity estimation as presented in Section 2.4.1.⁹

2.8.7 Semicoherent Metric

To derive the semicoherent metric, we investigate the mismatch in the semicoherent detection statistic in the presence of a strong signal. Starting from Equation (2.54), using binary photon weights and the rectangular lag window,

$$\begin{aligned} S_1 &= \sum_{j=1}^N \sum_{k=1}^N e^{-i[\phi(t_j) - \phi(t_k)]} \hat{W}_T^{\text{rect}}(\tau_{jk}) \\ &= \sum_{j=1}^N e^{-i\phi(t_j)} \sum_{k=1}^N e^{i\phi(t_k)} \hat{W}_T^{\text{rect}}(\tau_{jk}). \end{aligned} \quad (2.148)$$

Again, replacing the sum over k with a continuous integral allows us to write the mismatch as:

$$\bar{m} = 1 - \frac{\sum_{j=1}^N e^{-i\phi(t_j, \mathbf{u}_{\text{sig}} + \Delta \mathbf{u})} \langle e^{i\phi(t, \mathbf{u}_{\text{sig}} + \Delta \mathbf{u})} \rangle^{[t_j]}}{\sum_{j=1}^N e^{-i\phi(t_j, \mathbf{u}_{\text{sig}})} \langle e^{i\phi(t, \mathbf{u}_{\text{sig}})} \rangle^{[t_j]}}. \quad (2.149)$$

Assuming that each coherent window contains the same power (and hence has the same S/N at \mathbf{u}_{sig}), this can be simplified to:

$$\bar{m} = 1 - \frac{1}{N} \sum_{j=1}^N e^{-i\phi(t_j, \Delta \mathbf{u})} \langle e^{i\phi(t, \Delta \mathbf{u})} \rangle^{[t_j]}. \quad (2.150)$$

⁹We revisit the distribution of S_1 in the presence of a signal, and generalised to non-binary photon weights, in Section 4.7.2.

Taylor expanding this mismatch around $\Delta \mathbf{u} = 0$ to second order in $\Delta \mathbf{u}$ gives:

$$\begin{aligned}
\bar{m} &= \frac{i}{N} \sum_{j=1}^N \left(\partial_k \phi|_{t=t_j} - \langle \partial_k \phi \rangle^{[t_j]} \right) \Delta u^k \\
&+ \frac{1}{2N} \sum_{j=1}^N \left(\partial_k \phi|_{t=t_j} \partial_\ell \phi|_{t=t_j} + \langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]} \right) \Delta u^k \Delta u^\ell \\
&- \frac{1}{N} \sum_{j=1}^N \left(\partial_k \phi|_{t=t_j} \langle \partial_\ell \phi \rangle^{[t_j]} \right) \Delta u^k \Delta u^\ell \\
&+ \frac{i}{2N} \sum_{j=1}^N \left(\partial_k \partial_\ell \phi|_{t=t_j} - \langle \partial_k \partial_\ell \phi \rangle^{[t_j]} \right) \Delta u^k \Delta u^\ell \\
&+ \mathcal{O}(\Delta \mathbf{u}^3), \tag{2.151}
\end{aligned}$$

where there are implicit sums over repeated indices. Evaluating the partial derivatives at t_j , under the assumption that $T \ll T_{\text{obs}}$ and $T \ll 1$ yr gives¹⁰ requirements :

$$\partial_k \phi|_{t=t_j} \approx \langle \partial_k \phi \rangle^{[t_j]} \tag{2.152}$$

$$\partial_k \partial_\ell \phi|_{t=t_j} \approx \langle \partial_k \partial_\ell \phi \rangle^{[t_j]} \tag{2.153}$$

Thus, the mismatch of Equation (2.151) becomes,

$$\begin{aligned}
\bar{m} &\approx \frac{1}{2N} \sum_{j=1}^N \left(\langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]} - \langle \partial_k \phi \rangle^{[t_j]} \langle \partial_\ell \phi \rangle^{[t_j]} \right) \Delta u^k \Delta u^\ell \\
&= \frac{1}{2N} \sum_{j=1}^N G_{k\ell}^{[t_j]} \Delta u^k \Delta u^\ell. \tag{2.154}
\end{aligned}$$

Hence, the semicoherent metric components can be found by taking half the average of the coherent metric components of Equations (2.120) over all photons in the observation time. Using the approximations given in [148], which are valid under the assumption that the data set spans many years, we find

$$\bar{G}_{ff} = \frac{\pi^2 T^2}{6}, \tag{2.155a}$$

$$\bar{G}_{f\dot{f}} = \frac{\pi^2 T^4}{360} \gamma^2, \tag{2.155b}$$

$$\bar{G}_{n_x n_x} = \bar{G}_{n_y n_y} = \pi^2 f^2 r_E^2 [1 - \text{sinc}^2(\Omega_E T / 2\pi)], \tag{2.155c}$$

¹⁰Note that these assumptions only hold for isolated pulsar searches. If searching for a pulsar in a binary system, since the coherence window is typically longer than the orbital period, these approximations do not hold. For the relevant parameters, the semicoherent metric components are approximately twice as large in this case. Credit to L. Nieder for noting this.

where γ is the semicoherent refinement factor [148, 179] defined as

$$\gamma^2 = 1 + \frac{60}{N} \sum_{j=1}^N \frac{(t_j - t_0)^2}{T^2}. \quad (2.156)$$

The mismatches predicted by these derived metric components are compared to the measured mismatches in S_1 for a simulated pulsar signal in Figure 2.13.

For the purpose of the analytic study of the computing cost scaling in this paper, we employ the approximation $\gamma \approx \sqrt{5}T_{\text{obs}}/T = \sqrt{5}R$. Hence the determinant of the semicoherent metric is obtained as,

$$\sqrt{\det \bar{G}} \approx \frac{\pi^4}{4\sqrt{27}} T^3 f^2 r_E^2 R \left[1 - \text{sinc}^2 \left(\frac{\Omega_E T}{2\pi} \right) \right]. \quad (2.157)$$

2.8.8 Optimal Mismatch in Semicoherent Search

Following the same steps as in Appendix 2.8.5, we can find the optimal average mismatch for a semicoherent search with sensitivity $p_{\text{scoh},1}^{-1}$ at a fixed computing cost C_0 by consideration of the following Lagrange function:

$$\begin{aligned} L(T, \bar{m}, \lambda) &= p_{\text{scoh},1}^{-1} + \lambda(C_{\text{scoh}} - C_0) \\ &= (1 - \langle \bar{m}_{\text{tot}} \rangle)^{1/2} T^{1/4} + \lambda(K'_{\text{scoh}} \bar{m}^{-3/2} T^{(s-1)} - C_0). \end{aligned} \quad (2.158)$$

Applying the method of Lagrange multipliers as above, we find that

$$3\xi \bar{m}_{\text{opt}} = \frac{1 - \xi \bar{m}_f}{\frac{4(s-1)}{3} + 1}. \quad (2.159)$$

As argued in Section 2.4.3, an efficient strategy uses coherence window sizes T much less than half a year. In this regime of interest, $s = 5$. Using interbinning to interpolate the frequency spectrum gives $\langle m_f \rangle \approx 0.075$, giving the optimal maximum mismatch in the remaining three parameters as $\bar{m}_{\text{opt}} = 0.146$.

2.8.9 Sky-grid Construction

From the metrics derived above, in Appendices 2.8.2 and 2.8.7, we know when searching over a grid of sky locations that these grid points should be defined by a uniform grid in the ecliptic plane.

To construct the sky search grid for a source within an angular radius of θ from (α_0, δ_0) , this central point is rotated from equatorial to ecliptic coordinates according to the Earth's axial tilt (using the obliquity of the ecliptic, ϵ) and projected into the ecliptic plane, with Cartesian coordinates (x_0, y_0) ,

$$x_0 = \cos(\alpha_0) \cos(\delta_0), \quad (2.160)$$

$$y_0 = \cos(\epsilon) \sin(\alpha_0) \cos(\delta_0) + \sin(\epsilon) \sin(\delta_0). \quad (2.161)$$




Figure 2.13: Comparison of mismatch in S_1 (dashed curves) with semicoherent metric prediction (solid curves). In each panel the horizontal axis shows the offset from the signal parameters in f (left), \dot{f} (middle), and sky position (right). The sky-location offset is $\sqrt{\Delta n_x^2 + \Delta n_y^2}$, which measures the offset in coordinates (n_x, n_y) in the ecliptic plane. The underlying pulsar signal has been simulated with parameters $f = 32$ Hz, $\dot{f} = -10^{-12}$ Hz s $^{-1}$ for a total observational data time span of $T_{\text{obs}} = 3.4$ yr and a coherent window size of $T = 524288$ s.

A square of side length θ on the unit circle is calculated around this point, and sampled (using the semicoherent or coherent metric components as appropriate) with spacings

$$\Delta n_x = \Delta n_y = 2\sqrt{m/G_{n_x n_x}} . \quad (2.162)$$

These grid points are then projected back onto the unit sphere, and rotated into equatorial coordinates for barycentering.

Since a square region is sampled in the ecliptic plane, many of the resulting sky-points lie outwith the radius defining the search region on the sky. These points are simply discarded, resulting in the original circular search region on the sky in equatorial coordinates, sampled by a uniform grid defined in the ecliptic plane.

A possible problem arises when the search region crosses the ecliptic equator, since when the square is constructed in the ecliptic plane, some points lie outwith the unit circle, and therefore cannot be projected onto a unit sphere. This can be overcome by reflecting points, (x, y) , which lie outside the unit circle back into the sphere around the ecliptic longitude, l , of the center of the search region:

$$l = \tan^{-1}(y_0/x_0) , \quad (2.163a)$$

$$x' = \cos(l) - [x - \cos(l)] , \quad (2.163b)$$

$$y' = \sin(l) - [y - \sin(l)] . \quad (2.163c)$$

The new points (x', y') are then projected into the *opposite* hemisphere from the central point of the search region, resulting in a grid which covers an area of the sky which wraps around the ecliptic equator.

Chapter 3

PSR J1906+0722: An Elusive Gamma-Ray Pulsar

*Originally published as Clark, C. J., Pletsch, H. J., Wu, J., et al. 2015,
The Astrophysical Journal Letters, 809, L2.
© 2015. The American Astronomical Society. All rights reserved.*

Abstract

We report the discovery of PSR J1906+0722, a gamma-ray pulsar detected as part of a blind survey of unidentified *Fermi* Large Area Telescope (LAT) sources being carried out on the volunteer distributed computing system, *Einstein@Home*. This newly discovered pulsar previously appeared as the most significant remaining unidentified gamma-ray source without a known association in the second *Fermi*-LAT source catalog (2FGL) and was among the top ten most significant unassociated sources in the recent third catalog (3FGL). PSR J1906+0722 is a young, energetic, isolated pulsar, with a spin frequency of 8.9 Hz, a characteristic age of 49 kyr, and spin-down power 1.0×10^{36} erg s⁻¹. In 2009 August it suffered one of the largest glitches detected from a gamma-ray pulsar ($\Delta f/f \approx 4.5 \times 10^{-6}$). Remaining undetected in dedicated radio follow-up observations, the pulsar is likely radio-quiet. An off-pulse analysis of the gamma-ray flux from the location of PSR J1906+0722 revealed the presence of an additional nearby source, which may be emission from the interaction between a neighboring supernova remnant and a molecular cloud. We discuss possible effects which may have hindered the detection of PSR J1906+0722 in previous searches and describe the methods by which these effects were mitigated in this survey. We also demonstrate the use of advanced timing methods for estimating the positional, spin and glitch parameters of difficult-to-time pulsars such as this.

3.1 Introduction

In Chapter 2 we presented newly advanced methods designed to increase the sensitivity of blind searches without increasing the computational cost. These improvements were incorporated into a new blind survey of unidentified, pulsar-like *Fermi*-LAT sources which was conducted on the distributed volunteer computing system, *Einstein@Home*.¹ Previous surveys have been extremely successful in detecting new gamma-ray pulsars [136, 151, 154], and the newly improved search methods, in combination with the latest *Fermi*-LAT data, offer a significant increase in sensitivity. As part of this survey, we carried out a blind search for pulsed emission from a point source in the third *Fermi*-LAT source catalog [3FGL, 138], 3FGL J1906.6+0720. This source, previously known as 2FGL J1906.5+0720 [190], is highly significant and stands out as the most significant unassociated 2FGL source. Moreover, it was included in the “bright” pulsar-like source list described by Romani [191]. An investigation of the spectral properties of 2FGL sources found that, after the source associated with the Galactic Center, 2FGL J1906.5+0720 was the unidentified source most likely to contain a pulsar [95]. As such, over recent years, this source has been searched for pulsations, both in gamma rays [e.g., 151, 192] and in radio observations [e.g., 101]. However, despite these attempts, pulsed emission from this source remained undetected until now.

Here, we present the discovery and follow-up study of PSR J1906+0722, a young isolated gamma-ray pulsar detected by the *Einstein@Home* survey.

3.2 Discovery

3.2.1 Data Preparation

In the blind search we analyzed *Fermi*-LAT data recorded between 2008 August 4 and 2014 April 6. The *Fermi* Science Tools² were used to extract Pass 8 source class photons, which were analyzed using the P8_SOURCE_V3 instrument response functions (IRFs).³ We used `gtselect` to select photons with reconstructed directions within an 8° region of interest (ROI) around 3FGL J1906.6+0720, photon energies > 100 MeV and zenith angles $< 100^\circ$. We only included photons detected when the LAT was working in normal science mode, and with rocking angle $< 52^\circ$.

To assign photon weights representing the probability of each photon having been emitted by the target source [152], we performed a likelihood spectral analysis using the `pointlike` package. We built a source model by including all 3FGL catalog sources located within 13° of 3FGL J1906.6+0720, while allowing the spectral parameters of point sources within 5° to vary. We modeled the gamma-ray spectrum of this source with an exponentially cutoff power law, typical of gamma-ray pulsar spectra [190]. We used the `template_4years_P8_V2_scaled.fits` map cube and `isotropic_source_4years_P8V3` template to model the Galactic diffuse emission

¹<http://www.einsteinathome.org>

²<http://fermi.gsfc.nasa.gov/ssc/data/analysis/software>

³The Science Tools, IRFs and diffuse models used here are internal pre-release versions of the Pass 8 data analysis. Our results did not change substantially with the final release versions.

and the isotropic diffuse background respectively.⁴ The normalization parameters of both diffuse components were left free. Finally, the photon weights were computed using `gtsrcprob`, based on the best-fit source model resulting from the likelihood analysis.

3.2.2 Blind Search Method

For the blind search, we assumed a canonical isolated pulsar model, making it necessary to search in four parameters: spin frequency, f , spin-down rate, \dot{f} , R.A., α and decl., δ .

The basis for most blind searches for gamma-ray pulsars is the well-known multistage scheme based around an initial semicoherent search [e.g., 129, 151]. For this survey, we implemented the form of the multistage search scheme described in Chapter 2, where the initial semicoherent stage uses a lag-window of duration $2^{21}\text{s} \approx 24$ days.

Notably, this survey incorporates an intermediate semicoherent refinement step, with a longer (more sensitive) lag-window of $2^{22}\text{s} \approx 48$ days, reducing the parameter space around each first-stage candidate to be searched in the final fully-coherent follow-up step. This improves the efficiency of the follow-up stage, and allows the search to “walk” away (in all 4 search parameters) from the original location of the candidate if necessary.

Figure 3.1 illustrates the importance of these new techniques. In the blind survey, we searched a conservatively large circular region around the 3FGL sky location with a radius 50% larger than the 3FGL 95% confidence region. As evident from Figure 3.1, the pulsar lies far outside the original source’s confidence region, and also outside our search region. We therefore owe its detection to the large resolution of the semicoherent step, and the flexibility of the follow-up steps, which allow for signals to be detected despite a large offset between the signal parameters and the search location.

The most significant pulsar candidates from the blind search were automatically refined using the H -test statistic [176]. This revealed an interesting candidate; however the measured signal-to-noise ratio (S/N) was slightly below the detection threshold for a blind search involving a very high number of trials. Upon manual inspection, clear pulsations were observed in the photon data after April 2010; however the phase of these pulsations was not constant, and exhibited “wraps” in which the pulsation phase quickly jumped by one full rotation. These features indicated that the canonical isolated pulsar model used for the blind survey was insufficient, and hence follow-up studies were required to describe the pulsar’s rotation over the entire dataset.

3.3 Follow-up Analysis

Before carrying out follow-up analyses, we extended the dataset to include photons observed until 2014 October 1 and increased the ROI to 10° . To speed up the timing procedure computations, we discarded photons with a probability weight below 5%.

⁴<http://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>

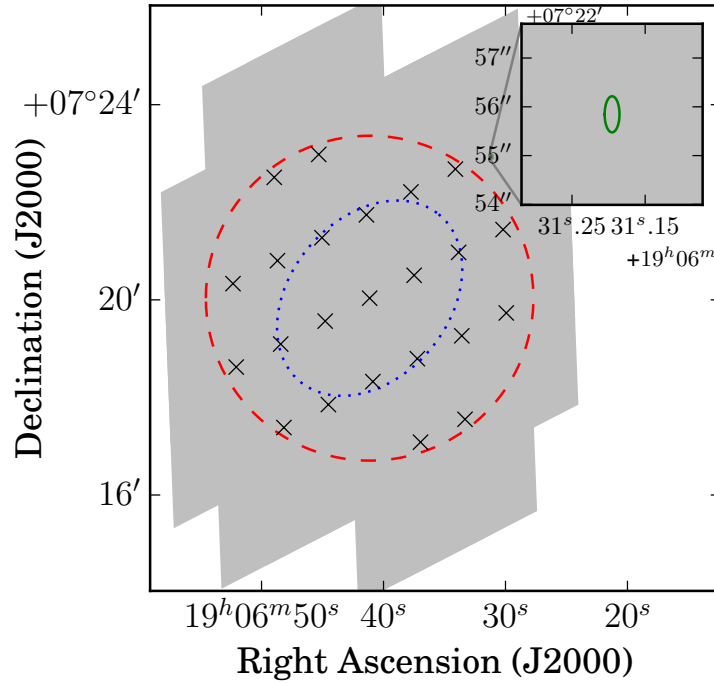


Figure 3.1: Sky location of PSR J1906+0722 and positional offset from the catalog location. The dotted ellipse shows the 3FGL 95% confidence region. The dashed ellipse shows the region in which the search grid (crosses) was constructed for the initial semicoherent search stage. The filled area shows the region that can be reached by the follow-up stage as it moves away from the initial candidate location. The 1σ region from the timing solution, shown by the solid ellipse, is highlighted in the inset.

3.3.1 Glitch Identification

Since pulsations were initially only detected during the final 4 years of data, the first step was to identify any glitches within the observation time. To achieve this, we searched the local $\{f, \dot{f}\}$ space around the observed signal, in 150-day segments with approximately 90% overlap, using the Q_M -test [172, and Chapter 2],

$$Q_M = 2M \frac{\sum_{n=1}^M |\alpha_n|^2 \mathcal{P}_n}{\sum_{n=1}^M |\alpha_n|^2}, \quad (3.1)$$

where α_n and \mathcal{P}_n are the Fourier coefficients of the measured pulse profile and the coherent power at the n th harmonic respectively. Using the Q_M test method to weight the contributions from each harmonic, as opposed to the commonly used H -test, offers a significant sensitivity improvement [173], making it particularly useful when analyzing weak pulsar signals. For this step, we included the first 10 Fourier coefficients with appreciable power from a segment of the data in which the signal was reasonably stable. Using the results of this scan, shown in Figure 3.2, an initial ephemeris was produced for the timing procedure described in the following section.

3.3.2 Timing Analysis

To accurately estimate the pulsar's rotational, glitch and sky location parameters we used a variation of the timing method used by [167], based on unbinned likelihood maximization. For all N photons in the dataset, with weights $\{w_j\}$, we assigned a rotational phase $\phi \equiv \phi(t_j, \mathbf{u})$, determined by the photon's arrival time, t_j and the set of model parameters, denoted by the vector \mathbf{u} . For a template pulse profile, $F(\phi)$, the likelihood is

$$\mathcal{L}(\mathbf{u}) = \prod_{j=1}^N [w_j F(\phi(t_j, \mathbf{u})) + (1 - w_j)]. \quad (3.2)$$

We first constructed a template pulse profile from the (background subtracted, see Figure 3.3) photons within a sub-section of the data set in which the initial ephemeris was believed to be accurate. When timing PSR J1906+0722 we used a template pulse profile consisting of 3 wrapped Gaussian functions [80], which were fit by maximizing the likelihood within the segment.

With a template profile at hand, we then estimated the pulsar's parameters (given in Table 3.1) by varying them around their initial estimate to maximize the likelihood over the *entire* dataset. The result is a likelihood maximization which is unbinned in both phase (via the template profile) and time. This avoids the need to construct a set of data subsegments for pulse times of arrival (TOA) determination. This is especially beneficial for faint pulsars, which require longer subsegments (and hence fewer TOAs) to ensure the S/N is large enough in each for accurate TOA measurement. Subsequently, using the most likely parameters, the template profile was updated and the process was iterated to maximize the overall likelihood.




Figure 3.2: Evolution of the PSR J1906+0722 signal including the glitch at MJD 55067. Left: Phase–time diagram where each point represents one photon, with the intensity representing the photon weight. Center and right: The Q_{10} -test (shown by the color-bar) calculated over small ranges in $\{f, \dot{f}\}$, centered on the pre-glitch parameters, in overlapping 150-day segments, and maximized over \dot{f} and f respectively. The dashed line indicates the maximum likelihood timing solution.

Table 3.1. Parameters for PSR J1906+0722

Parameter	Value
Range of Photon Data (MJD)	54682–56931
Reference epoch (MJD)	55716
Timing Parameters	
R.A., α (J2000.0)	19 ^h 06 ^m 31 ^s .20(1)
Decl., δ (J2000.0)	+07° 22′ 55″.8(4)
Frequency, f (Hz)	8.9666688432(1)
1st frequency derivative, \dot{f} , (Hz s ⁻¹)	$-2.884709(2) \times 10^{-12}$
2nd frequency derivative, \ddot{f} , (Hz s ⁻²)	$3.18(1) \times 10^{-23}$
Glitch epoch ^a (MJD)	55067 ⁺² ₋₉
Permanent f glitch increment ^a , Δf (Hz)	$4.033(1) \times 10^{-5}$
Perm. \dot{f} glitch increment ^a , $\Delta \dot{f}$ (Hz s ⁻¹)	$-2.56(3) \times 10^{-14}$
Decaying f glitch increment ^a , Δf_d (Hz)	$3.64(9) \times 10^{-7}$
Glitch decay time constant ^a , τ_d (days)	221(12)
Spectral Properties	
Spectral index, Γ	1.9 ± 0.1
Cutoff energy, E_c (GeV)	5.5 ± 1.2
Photon flux, ^b F_{100} (photons cm ⁻² s ⁻¹)	$(1.1 \pm 0.3) \times 10^{-7}$
Energy flux, ^b G_{100} (erg cm ⁻² s ⁻¹)	$(7.3 \pm 1.3) \times 10^{-11}$
Derived Properties	
Period, P (ms)	111.524136498(1)
1st period derivative, \dot{P} (s s ⁻¹)	$3.587895(2) \times 10^{-14}$
Weighted H -test	731.2
Characteristic age ^c , τ_c (kyr)	49.2
Spin-down power ^c , \dot{E} (erg s ⁻¹)	1.02×10^{36}
Surface B -field strength ^c , B_s (G)	2.02×10^{12}
Light-cylinder B -field ^c , B_{LC} (G)	1.34×10^4
Heuristic distance ^c , d_h (kpc)	1.91

Note. — Values for timing parameters are the mean values of the marginalized posterior distributions from the timing analysis, with 1σ uncertainties in the final digits quoted in parentheses.

^aGlitch model parameters are defined in Edwards et al. [137], with the correction noted by Yu et al. [34].

^bFluxes above 100 MeV, F_{100} and G_{100} , were calculated by extrapolation from the $E > 200$ MeV spectrum.

^cDerived pulsar properties are defined in Abdo et al. [80]. The heuristic distance, $d_h = (L_\gamma^h/4\pi G_{100})^{1/2}$, is calculated from the heuristic luminosity, L_γ^h , described therein.

To explore the multi-dimensional parameter space we used the `MultiNest` nested sampling algorithm [193], which offers high sampling efficiency, and allows posterior distributions to be calculated as a by-product.

The timing procedure was carried out in two stages: firstly, all timing parameters were allowed to vary. Due to the shortness of the pre-glitch segment, the uncertainties in the glitch parameters dominated those of the remaining timing parameters. We therefore fixed the glitch parameters at their maximum likelihood values, and fit again for the remaining timing parameters.

When timing radio pulsar glitches, Yu et al. [34] noted that unique solutions for glitch epochs could not be found for large glitches occurring during an interval between two radio observations. We observe a similar effect here, although our limiting factor is the photon flux. When phase folding, a full rotation can be lost/gained if the offset between the model glitch epoch and the true glitch epoch is more than $1/\Delta f \approx 0.3$ days; however an average of only 1.4 weighted photons are observed from the pulsar within this time, making this phase wrap simply undetectable. We assumed that no phase increment occurred at the glitch, and found that the posterior distribution for the glitch epoch features several bands, separated by $1/\Delta f$. Due to the multi-modal shape of the posterior distribution, in Table 3.1 we report the glitch epoch that results in the maximum likelihood and the 95% credible interval.

The inclusion of an additional nearby source in the source model and raising the energy threshold to 200 MeV when calculating the photon weights for PSR J1906+0722 increased the S/N (see Section 3.3.3). Therefore the timing analysis was repeated with the updated photon weights, and the results are given in Table 3.1. The time versus rotational phase diagram based on this timing solution is shown in Figure 3.2 and the integrated pulse profile is displayed in Figure 3.3. Through these refinement and timing procedures, the initial candidate's Q_{10} -test S/N (Equation 2.18) was increased from $\theta_{10} = 6.86$ to the highly significant value of $\theta_{10} = 16.55$ given by the final timing solution.

3.3.3 Off-pulse Analysis

Fitting an exponential cutoff model to the spectrum of PSR J1906+0722 revealed a relatively high cutoff energy compared to typical gamma-ray pulsars ($E_c = 6.5 \pm 0.9$ GeV), suggesting that the spectrum could be contaminated by the presence of a nearby source as was also noted by Xing & Wang [192].

To investigate this possibility, we analyzed the off-pulse part of the data using photons with energies between 200 MeV and 300 GeV. A residual test statistic (TS) map for the off-pulse data (see Figure 3.3) revealed an excess $(0.28 \pm 0.02)^\circ$ away from PSR J1906+0722, at $(\alpha, \delta) = (286.84^\circ, 7.15^\circ)$, with a TS value of 288.

Modeling this secondary source with a power-law spectrum, we added it to the spectral model for the region, keeping its location fixed from the off-pulse analysis, but leaving its normalization and spectral index free, and analyzed again the full phase interval data. As a result, we found that the log-likelihood value increased slightly, and the new photon weights increased the S/N of the pulsations from $\theta_{10} = 16.38$ to $\theta_{10} = 16.55$.




Figure 3.3: Top panel: weighted pulse profile of PSR J1906+0722 given by the timing solution. Lower panels: weighted pulse profiles in increasing energy bands. In each panel, the estimated background level, calculated from the photon weights [194], is shown by the dashed line. The dot-dashed lines mark the off-pulse phase interval used in Section 3.3.3. The error bars show 1σ statistical uncertainties [151]. The pulsed fraction, p , and pulsed S/N, θ_{10} , in each energy band is also shown.

The low energy threshold of 200 MeV was chosen to provide improved angular resolution in order to better separate the pulsar emission from that of the new source. When lower energy (100–200 MeV) photons were included in the spectral analysis, the pulsation S/N calculated with the resulting photon weights decreased, suggesting that source confusion at low energies leads to a less reliable source model.

Figure 3.4 shows TS maps and spectral energy distributions for PSR J1906+0722 and the new source found in this off-pulse analysis. The integrated energy flux of the secondary source above 100 MeV is $4.34_{-0.67}^{+0.91} \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ with a spectral index of 2.17 ± 0.07 .

The best-fitting location of the secondary source is very close to the western edge of the supernova remnant (SNR), G41.1-0.3 [3C 397, 195]. Jiang et al. [196] observed a molecular cloud interacting with the SNR at this location; it is possible that we are observing gamma-ray emission resulting from this interaction.

3.4 Analysis in Other Wavelengths

3.4.1 Radio and X-ray Observations

In probing for radio emission from PSR J1906+0722, we carried out a 120-minute follow-up observation with the L-band (1.4 GHz) single-pixel receiver mounted on the 100 m Effelsberg Radio Telescope in Germany. The gamma-ray-timing ephemeris allowed us to search the data over dispersion measure (DM) only. No evidence for radio pulsations was found. Assuming a 10% pulse width, bandwidth $\Delta F = 150 \text{ MHz}$, telescope gain $G = 1.55$, $n_p = 2$ polarization channels, system temperature $T_{\text{sys}} = 24 \text{ K}$, digitization factor $\beta = 1.2$ and a signal-to-noise threshold of 5, by the radiometer equation [Equation (A1.22), 9, p265], we computed a flux density limit of $\approx 21 \mu\text{Jy}$. While this is below the conventional radio-quiet level of $30 \mu\text{Jy}$ [80], we note that the nearby LAT-discovered pulsar PSR J1907+0602 has been observed in radio observations with a flux density of just $3.4 \mu\text{Jy}$ [197], and would therefore not have been detected in this radio search.

To check for a possible X-ray counterpart, we analyzed a 10 ks observation with *Swift*'s X-ray Telescope [198]. No counterpart source was detected, with an unabsorbed-flux (0.5–10 keV) upper limit of $2 \times 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ at the pulsar position. This limit yields a gamma-ray-to-X-ray flux ratio of > 365 , or an efficiency $L_X/\dot{E} \lesssim 8.7 \times 10^{-5}$ at distance d_h , similar to other gamma-ray pulsars [151, 199].

3.4.2 Possible SNR Associations

There are 4 known SNRs lying within 1° from the timing position of PSR J1906+0722 [200]. There is strong evidence that the closest of these, G41.1–0.3, is a Type Ia SNR from a Chandrasekhar mass progenitor [201], making it unlikely to be the birthplace of a pulsar. Each of the remaining nearby SNRs lies closer to other young pulsars than to PSR J1906+0722 (G41.5+0.4 and G42.0–0.1 to PSR J1906+0746; G40.5–0.5 to PSR J1907+0602), making a physical association between any

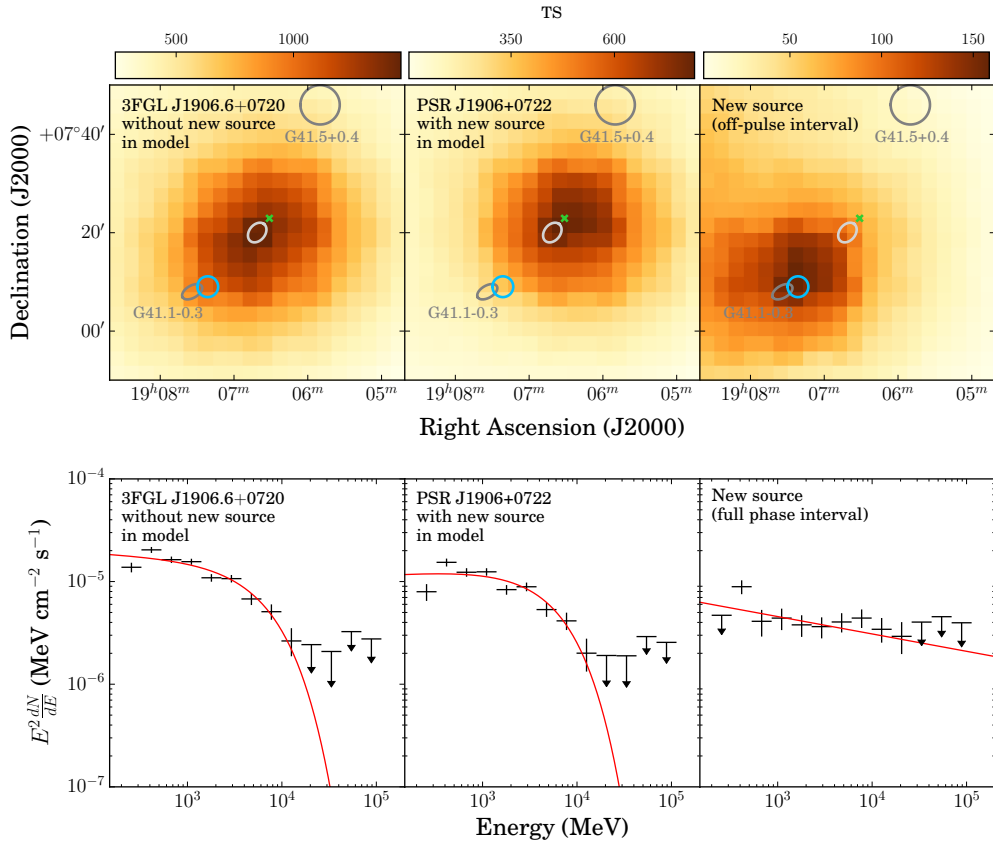


Figure 3.4: Top panels: test statistic (TS) sky-maps of the PSR J1906+0722 region above 200 MeV. Each pixel shows the TS value for a point source located at the pixel position. The cross represents the timing position of PSR J1906+0722. The light ellipse shows the 95% confidence region of the 3FGL source, the blue ellipse shows the 95% confidence region of the new secondary source, and the darker ellipses show the approximate extents of nearby SNRs. Bottom panels: Spectral energy distributions for the full-pulse interval. The solid curves present the results of the likelihood analyses of Section 3.3.3.

of these difficult to verify. Kick-velocity requirements based on the pulsar’s characteristic age and heuristic distance do not rule out any of these SNRs as the birthplace of the pulsar.

3.5 Discussion

Despite several years of attempts, the identification of 2FGL J1906.5+0720 remained elusive. Now that this source has been identified as PSR J1906+0722, we here investigate potential reasons for the failure of previous searches to detect it.

Perhaps the most significant source of difficulty in the detection of PSR J1906+0722 was the large positional offset between its 3FGL catalog position and its true position. This offset, which could only be accommodated by the new follow-up method outlined in Section 3.2.2, is most likely due to the presence of the secondary source described in Section 3.3.3.

The close proximity of PSR J1906+0722 to the Galactic plane ($b = 0.03^\circ$) likely also hindered its detection, as the large majority of the weighted photons can be attributed to the background. From the pulse profile shown in Figure 3.3, we estimate that the pulsed fraction of the total weighted photon flux (as defined in Equation 2.2) is as low as 6%⁵. This low pulsed fraction leads to a low observable S/N, making detection more challenging.

A further complication for detecting PSR J1906+0722 was the presence of the glitch about one year into the *Fermi* mission. This glitch is among the largest detected from a gamma-ray pulsar in terms of relative magnitude ($\Delta f/f \approx 4.5 \times 10^{-6}$) [154]. In previous searches using a shorter total observation time, the data segment after the glitch represented a much shorter fraction of the total observation time. As the time interval covered by *Fermi*’s observations since 2008 August 1 continues to increase, the existence of a long timespan in which a pulsar’s signal is stable becomes ever more likely. The increase in the weighted photon flux offered by the Pass 8 analysis [77] further increases the observable S/N throughout the observation time, and results in searches that are not only more sensitive overall [202], but also more robust against glitching or noisy pulsars.

The ability to detect young gamma-ray pulsars in blind searches can be of significant importance to the overall study of energetic pulsars. For example, [203] use the observed population of radio-quiet pulsars to investigate the dependence of properties of pulsar emission geometries on the spin-down energy, \dot{E} . Since pulsars with a high \dot{E} tend to exhibit timing noise and glitches (which do not typically affect radio searches), they are hard to find in gamma-ray data, where long integration times are required. Advanced search methods that can detect complicated signals such as that from PSR J1906+0722 are therefore crucial for reducing a potential bias against young, energetic and glitching pulsars in the radio-quiet population. As noted by Abdo et al. [80] and Caraveo [47], such pulsars are indeed lacking in the *Fermi* pulsar sample.

⁵Here we define the pulsed fraction as the fraction of the *total* weighted photon flux. In Chapter 4 we introduce an alternative quantity, based on the pulsed fraction of the flux *attributable to the pulsar*. Nevertheless, the high background level here did make detecting pulsations from PSR J1906+0722 more difficult.

Acknowledgments

This work was supported by the Max-Planck-Gesellschaft (MPG), as well as by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether research grant PL 710/1-1 (PI: Holger J. Pletsch). We are very grateful to all *Einstein@Home* volunteers who have donated their spare computing time, especially Connor Barry of Lafayette, Colorado, USA and Rich Johnson of Hayward, California, USA on whose computers PSR J1906+0722 was first detected.

The *Fermi*-LAT Collaboration acknowledges support for LAT development, operation and data analysis from NASA and DOE (United States), CEA/Irfu and IN2P3/CNRS (France), ASI and INFN (Italy), MEXT, KEK, and JAXA (Japan), and the K.A. Wallenberg Foundation, the Swedish Research Council and the National Space Board (Sweden). Science analysis support in the operations phase from INAF (Italy) and CNES (France) is also gratefully acknowledged.

Chapter 4

The Einstein@Home Gamma-ray Pulsar Survey. I.

Search Methods, Sensitivity and Discovery of New Young Gamma-ray Pulsars

*Following the initial submission of this thesis, this chapter was published as
Clark, C. J., Wu, J., Pletsch, H. J., et al., 2017, The Astrophysical Journal, 834, 106.
© 2017. The American Astronomical Society. All rights reserved.*

Abstract

We report on the results of a recent blind search survey for gamma-ray pulsars in *Fermi* Large Area Telescope (LAT) data being carried out on the distributed volunteer computing system, *Einstein@Home*. The survey has searched for pulsations in 118 unidentified pulsar-like sources, requiring about 10,000 years of CPU core time. In total, this survey has resulted in the discovery of 17 new gamma-ray pulsars, of which 13 are newly reported in this work, and an accompanying paper (in preparation). These pulsars are all young, isolated pulsars with characteristic ages between 12 kyr and 2 Myr, and spin-down powers between 10^{34} and 4×10^{36} erg s⁻¹. Two of these are the slowest spinning gamma-ray pulsars yet known. One pulsar experienced a very large glitch $\Delta f/f \approx 3.5 \times 10^{-6}$ during the *Fermi* mission. In this, the first of two associated papers, we describe the search scheme used in this survey, and estimate the sensitivity of our search to pulsations in unidentified *Fermi*-LAT sources. One such estimate results in an upper limit of 57% for the fraction of pulsed emission from the gamma-ray source associated with the Cas A supernova remnant, constraining the pulsed gamma-ray photon flux that can be produced by the neutron star at its center. We also present the results of precise timing analyses for each of the newly detected pulsars.

4.1 Introduction

In this chapter we present 13 new pulsar discoveries from the full *Einstein@Home* survey of 118 unidentified, but pulsar-like *Fermi*-LAT sources. These are the result of around 10,000 years of CPU time generously donated by volunteers. An accompanying paper is being prepared by J. Wu et al. (hereafter Paper II), which will describe in more detail the source selection and data preparation procedures; spectral and off-pulse analyses of the newly discovered pulsars; and the results of radio follow-up searches for these new pulsars.

The chapter is organized as follows: In Section 4.2 the search methods are described; an investigation of the sensitivity of the search follows in Section 4.3; details of the newly discovered pulsars and their timing solutions are given in Section 4.4; Section 4.5 contains a discussion of the sensitivity of blind searches to unidentified gamma-ray pulsars; and finally we summarize our conclusions in Section 4.6.

4.2 Search Scheme

The following contains some overlap with the two previous chapters. This has been left in to ensure that this chapter contains a complete description of the Einstein@Home search.

4.2.1 Data

The data searched during the survey consisted of gamma-ray photons detected by the LAT between 2008 August 4 and 2014 April 6 (2014 October 1 for some sources searched later in the survey) with energies above 100 MeV. Photons were included if they arrived within 8° of a target source, with a zenith angle $< 100^\circ$ and when the LAT's rocking angle was $< 52^\circ$. The photons were selected and analyzed using the `P8_SOURCE_V3` instrument response functions (IRFs).

For each target source, we performed a likelihood spectral analysis using the `pointlike` package [204]. Our source model included all 3FGL catalog sources within 13° of the target source and used the `template_4years_P8_V2_scaled.fits` map cube and `isotropic_source_4years_P8V3` template to model the Galactic diffuse emission [205] and isotropic background respectively.

Target sources were modeled with an exponentially cutoff power law typical of gamma-ray pulsars. During the likelihood fitting, we allowed the normalization of the diffuse models, and the spectral parameters of the target source and all 3FGL sources within 5° to vary. Sources searched near the beginning of the survey had their sky positions fixed at the 3FGL location. Later sources were relocalized during the likelihood fitting to exploit the improved angular resolution offered by the Pass 8 data. Spectral energy distribution (SED) plots and Test Statistic (TS) maps were visually compared to the corresponding 3FGL sources to diagnose any problems with the fitting. With the best-fitting source model, we used `gtsrcprob`¹ to compute weights representing the probability

¹`gtsrcprob` is part of the *Fermi* Science Tools, available at <http://fermi.gsfc.nasa.gov/ssc/data/>

of each photon having come from our target source based on their reconstructed energy and arrival direction. Full details of the data preparation methods can be found in Paper II.

The IRFs and diffuse templates used here were internal pre-release versions of the Pass 8 analysis tools, as the final release versions were not yet available when the survey began. When investigating the gamma-ray emission from the region surrounding PSR J1906+0722 (Chapter 3), we found that these preliminary IRFs and templates resulted in spectral parameters consistent with those found using the final Pass 8 release. However, photon weights calculated with the most recent Pass 8 data usually result in slightly higher pulsation significance within the same time interval; the sensitivity estimates in Section 4.3 are likely to be more conservative as a result.

4.2.2 Parameter Space

To search for gamma-ray pulsations in LAT data, it is necessary to assume a certain “phase model” (i.e. a rotation ephemeris) relating the arrival time of every photon to a certain rotational phase, and test all possible combinations of the model parameters for pulsations, indicated by large values of a detection statistic (described in Section 4.2.3). In the case where a signal is present, the distribution of rotational phases will deviate significantly from uniformity. For isolated pulsars, the phase model² is typically described by a Taylor series expansion in time around a chosen reference epoch t_{ref} , for photon arrival time t at the Solar System barycenter (SSB),

$$\Phi(t) = \Phi_0 + 2\pi \sum_{m=1} \frac{f^{(m-1)}}{m!} (t - t_{\text{ref}})^m, \quad (4.1)$$

where $f^{(m)}$ denotes the m -th time derivative of the pulsar’s rotational frequency, f . While the higher derivative terms are often measurable for young pulsars, it is usually sufficient (and often only feasible) to include only the first two terms in the blind search, resulting in a simplified phase model in which the spin frequency decreases by a constant spin-down rate, $\dot{f} \equiv f^{(1)}$.

Aside from correcting for this constant spin-down, it is also necessary to account for the apparent Doppler modulation of pulsations that results from the Earth’s orbit around the SSB. This can be achieved by applying position-dependent corrections to the measured photon arrival times, to retrieve the set of arrival times at the SSB, hereafter denoted as $\{t_j\}$. The angular resolution at which sky positions must be searched increases linearly with the pulsar’s spin frequency (see Equation 4.11). For all but the slowest of pulsars, the required resolution is finer than the gamma-ray source localization, determined by the LAT’s point spread function. For a blind survey of unidentified gamma-ray point sources, it is therefore necessary to search in two sky positional parameters (R.A. α and Decl. δ), making the overall search parameter space four-dimensional. For sources for which we used the original 3FGL locations, we searched a circular region around the source with an angular radius that was 50% larger than the semi-major axis of the 95% confidence region. For

analysis/software/

²While we define the phase in radians, in all plots we show phase in rotations for clarity, and re-normalize the pulse profiles accordingly.

relocalized sources, we searched a conservatively large region with a radius three times larger than the semi-major axis of the 68% confidence region.

Some pulsars (e.g. PSRs J2017+3625 and J1350–6225) were found near the edge, or even slightly outside of their search regions, indicating that the confidence regions may be underestimated, and pulsars may have been missed by our survey as a result. This could be due to nearby, unmodeled gamma-ray sources “pulling” the apparent position of the source away from its true position, as was seen with PSR J1906+0722 (Section 3.3.3). To mitigate this effect in future surveys it may be necessary to search over larger regions, especially for sources at low Galactic latitude where source confusion is more likely. However increasing the solid angle over which we search increases the computational cost of the search by the same factor.

We split the search parameter space into two main regions: the young pulsar region, with spin frequencies below 80 Hz; and the millisecond pulsar (MSP) region at higher spin frequencies. This parameter space is shown by the shaded area in Figure 4.1, and covers all currently known young pulsars, MSPs and magnetars. In the low frequency region we extend the \dot{f} range from 0 down to -10^{-9} Hz s $^{-1}$ to be sensitive to the youngest and most energetic pulsars. Older, recycled MSPs have much lower spin-down rates, and we therefore only search from 0 down to -10^{-13} Hz s $^{-1}$ in this region. Since more sky locations must be searched at higher frequencies, the majority of the computational cost of the search is spent in the high frequency and high spin-down regions. Pulsars whose pulse profile features two similarly sized peaks separated by half a rotation have most power in the second harmonic of their spin frequency. For this reason, we search up to 1520 Hz, more than twice the frequency of the fastest known MSP, 716 Hz [206]. Only one known pulsar, PSR J0537–6910, has its second spin harmonic outside our parameter space [207].

For each of the 118 unidentified LAT sources in which we searched, this parameter space is split into $\sim 10^3$ – 10^6 smaller “work units”, each of which can be searched in a few hours on a typical home computer. These work units are then distributed amongst *Einstein@Home* volunteers’ computers.

4.2.3 Detection Statistics

In all stages of a gamma-ray pulsar search, statistical tests are used to measure the strength of pulsations for given rotational parameters. The detection statistics used in this survey are described in detail in Chapter 2, and briefly defined here.

Kerr [152] demonstrated the advantages of applying a weight to each photon indicating its probability of having come from the target source. The photon probability weights mentioned in Section 5.2.1, denoted by $\{w_j\}$, were therefore used to weight the contributions of each photon to a detection statistic. Weighting photons improves the sensitivity of a blind search by avoiding the need to apply specific photon energy and angular offset cuts, and by increasing the apparent fraction of flux that is pulsed.

To mitigate the computational cost of a blind search, semicoherent methods can be used, in which only photons arriving within a certain time difference from one another are combined coherently.




Figure 4.1: Frequency–spin-down diagram, showing the locations of non-gamma-ray pulsars in the ATNF Pulsar Catalogue [black crosses, 14], gamma-ray pulsars detected by *Fermi*-LAT (blue circles), and the newly detected *Einstein@Home* pulsars reported in this work (orange squares). The parameter space covered by the *Einstein@Home* survey is shown by the gray shaded area. Lines of constant characteristic age (dot-dashed), surface magnetic field strength (dotted) and spin-down power (dashed) are also shown. These properties are calculated as described in Abdo et al. [80].

The (real-valued) semicoherent detection statistic used in this search is defined as³

$$S_1 = \frac{1}{\kappa_{S_1}} \sum_{j=1}^N \sum_{k \neq j}^N w_j w_k e^{-i[\Phi(t_j) - \Phi(t_k)]} \hat{W}_T^{\text{rect}}(\tau_{jk}), \quad (4.2)$$

where τ_{jk} is the time difference, or lag, between the arrivals of the j -th and k -th photons, and \hat{W}_T^{rect} is a rectangular window, of length T ,

$$\hat{W}_T^{\text{rect}}(\tau) = \begin{cases} 1, & |\tau| \leq T/2 \\ 0, & \text{otherwise.} \end{cases} \quad (4.3)$$

The κ_{S_1} term of Equation (4.2) is a normalizing factor,

$$\kappa_{S_1} = \sqrt{\sum_{j=1}^N \sum_{k \neq j}^N w_j^2 w_k^2 \hat{W}_T^{\text{rect}}(\tau_{jk})}, \quad (4.4)$$

making the noise distribution of S_1 well approximated by a normal distribution with zero mean and unit variance.

The lag window length, T , is an important tuneable parameter for a search based on semicoherent methods. A longer lag window offers more sensitivity, but requires a finer grid in all four search parameters, and therefore results in a more costly search.

In the case where the lag window covers the entire observation span, then all photons are combined fully coherently, and the test statistic reduces to the well known Rayleigh test (modulo a constant term, and normalization) at the fundamental harmonic, $n = 1$,

$$\mathcal{P}_n = \frac{1}{\kappa^2} \left| \sum_{j=1}^N w_j e^{-in\Phi(t_j)} \right|^2, \quad (4.5)$$

with the normalization constant,

$$\kappa^2 = \frac{1}{2} \sum_{j=1}^N w_j^2. \quad (4.6)$$

This is hereafter referred to as the coherent Fourier power at the n -th harmonic.

To gain further sensitivity to weak signals, one can also combine the coherent Fourier power from several harmonics of the fundamental spin frequency. The well-known H -test developed by de Jager et al. [176] offers a heuristic method for combining these harmonics in the typical case where the pulsar's pulse profile (and hence the distribution of Fourier power amongst the different harmonics) is unknown in advance by maximizing over the number of included harmonics, M , via

$$H = \max_{1 \leq M \leq 20} \left(\sum_{n=1}^M \mathcal{P}_n - 4M + 4 \right). \quad (4.7)$$

³The subscript 1 here denotes that the detection statistic only sums power in the fundamental harmonic.

Combining Fourier power from higher harmonics requires finer resolution in all phase model parameters. It therefore only becomes feasible in later search stages, in which the parameter space within which a candidate signal could lie is constrained to be very narrow.

As discussed in Chapter 2, a multistage search scheme can be used to combine the efficiency of a semicoherent search with the superior sensitivity of fully coherent methods. In this scheme, the majority of the search is spent scanning the entire parameter space with the most efficient method available, before “following-up” the most interesting candidates in more sensitive stages.

4.2.4 Initial Search Stage

In this survey, the first stage used the semicoherent detection statistic, S_1 , with a lag window of length $T = 2^{21}\text{s} \approx 24$ days. This lag window is a factor of 2 longer than in previous *Einstein@Home* searches [136].

As described in Section 2.5, the semicoherent detection statistic, S_1 , defined in Equation (4.2), can be approximated more efficiently as a discrete Fourier transform (DFT), by utilising the FFTW fast Fourier transform (FFT) algorithms [182]. We hereby refer to the DFT form of the semicoherent detection statistic as \hat{S}_1 .

Each FFT searched over a frequency bandwidth of $\Delta f_{\text{BW}} = 32$ Hz. We applied the technique of complex heterodyning, i.e. multiplying the FFT input vector by an additional sine wave at the heterodyning frequency, f_{H} , to shift the search band to higher frequencies, $[f_{\text{H}} - 16, f_{\text{H}} + 16)$ Hz, without increasing the FFT memory size, $\Delta f_{\text{BW}}T = 256$ MiB. This allows us to search for high frequency signals, such as those from millisecond pulsars, using typical computing hardware. Furthermore, since the required resolution in the sky position becomes finer at higher frequencies, we can construct individual sky grids for each frequency band to avoid oversampling sky positions at low frequencies. The first frequency band was centered at 0 Hz, and all frequencies below 5 mHz were ignored to prevent harmonics of *Fermi*’s orbital frequency (~ 0.175 mHz) from “drowning out” any astrophysical signal.

To ensure approximately equal sensitivity throughout the frequency band, we performed lag-domain interpolation (see Section 2.5.4), whereby each photon pair is interpolated into the 15 nearest bins on either side in the lag-series using a Welch-windowed sinc kernel [188, 189, p. 176]. Since this technique introduces an additional computational cost per pair of photons, we performed a photon weight cutoff to include at most the 30,000 highest-weight photons, ensuring that the FFT computation time remained the dominant factor. Identifying the photon weights as the probability of each photon being from a pulsar, the maximum (coherent) signal-to-noise ratio (S/N) is proportional to $\sum_{j=1}^N w_j^2$. For sources where fewer than 30,000 photons were required to reach 95% of this maximum (typically sources far from the Galactic plane, where the diffuse background is lower) we increased the number of interpolation bins, up to a maximum of 30.

A signal whose parameters, denoted by the vector \mathbf{u} , lie within the search space will in general not lie exactly at one of our search-grid points, and some of the S/N is lost as a result of this offset,

$\Delta \mathbf{u}$. We call this (fractional) loss in S/N “mismatch”,

$$m = 1 - \frac{\hat{S}_1(\mathbf{u} + \Delta \mathbf{u})}{\hat{S}_1(\mathbf{u})}. \quad (4.8)$$

We can predict the expected mismatch as a function of the distance to the nearest search-grid point using an analytical “metric” approximation, as described in Section 2.8.7. This prediction can then be used to construct an efficiently spaced grid of points in the parameter space at which to test for pulsations.

The spacing of frequency trials is fixed by the DFT formulation of \hat{S}_1 to be

$$\Delta f = \frac{1}{T}. \quad (4.9)$$

While this spacing would result in a large average mismatch, we can improve upon this by performing simple “interbinning” [141, 186] to partially recover the lost S/N experienced by signals lying between our frequency bins. This technique does not recover the full S/N for such signals, but is far more efficient than the alternative of “zero-padding” the FFT to double length.

In the remaining parameters, we construct a cubic lattice with spacings chosen to provide the optimal maximum mismatch in each parameter of $m = 0.15$ according to the metric approximation. In \dot{f} the spacing depends on the lag window T but also requires a refinement based on the full data set [179],

$$\Delta \dot{f} = \frac{12\sqrt{10m}}{\pi T^2} \left[1 + \frac{60}{N} \sum_{j=1}^N \frac{(t_j - t_{\text{ref}})^2}{T^2} \right]^{-1/2}. \quad (4.10)$$

The grid of sky locations is first defined within a circle (with unit radius) in the ecliptic plane as a square grid with spacings in each direction of

$$\Delta n_x = \Delta n_y = \frac{2\sqrt{m}}{\pi f_{\text{max}} r_E} [1 - \text{sinc}^2(\Omega_E T/2)]^{-1/2}, \quad (4.11)$$

where f_{max} is the maximum frequency in the searched frequency band, r_E and Ω_E are the Earth’s orbital semi-major axis (in light seconds) and orbital angular frequency respectively, and $\text{sinc}(z) = \sin(z)/z$. These locations are then projected back into the celestial sphere to cover the LAT source localization region. At each location, barycentering corrections are applied to each photon’s arrival time according to the JPL DE405 solar system ephemeris.

Each work-unit performs an FFT at every location in this cubic lattice within its assigned portion of the parameter space. The five highest values of \hat{S}_1 (including interbinned samples) are stored in a running short list that is updated after each FFT. At the end of the semicoherent stage, this short list is saved, and each short-listed candidate is automatically “followed up” in additional, more sensitive search stages.

4.2.5 Follow-up and Refinement Stages

After the semicoherent stage, we are left with a small number of candidates in each work unit that have been localized to a small region of the parameter space. However, due to the large number of work units for each *Fermi*-LAT source, weak signals in these short lists can be of low overall significance. To separate weak signals from noise candidates, we can carry out more sensitive follow-up stages to act as a veto for the large number of candidate signals reported back by the semicoherent stage.

In the *Einstein@Home* survey, we implemented an intermediate refinement stage, in which candidates from the first stage are refined using a double-length lag window ($T = 2^{22}$ s \approx 48 days). This step is computationally cheap, and narrows down the volume in which the candidate signal lies by a factor of ~ 16 .

Following the semicoherent refinement stage, we now have a parameter space volume around each candidate that is small enough for a fully coherent search to be feasible with just a small associated computing cost. For this stage, we search only in the fundamental harmonic using the \mathcal{P}_1 test, with grid spacings according to the coherent metric approximation derived in Section 2.8.2.

All search stages up to this point are carried out on the *Einstein@Home* volunteers' computers, after which the short-listed candidates from the initial semicoherent stage (each of which were followed-up), and the top candidates from the coherent follow-up stage are sent back to our servers.

As results are sent back, we update the top 20 most significant coherent candidates (see Appendix 4.7.1 for a description of the ranking procedure) overall from each source, and perform further follow-up and refinement procedures on them. Firstly, we refine the location of the candidate using the \mathcal{P}_1 statistic, but calculated over a grid with a smaller mismatch ($m = 0.05$) than that used in the third stage. We then perform a fully coherent search using the H -test to incoherently sum the Fourier power in the first 5 harmonics.

After this refinement step, diagnostic plots for each candidate are produced that illustrate the candidates' signals and their evolution throughout the *Fermi*-LAT observation time. This allows us to identify pulsars with timing noise, whose pulsations may be visible in these plots despite having a low apparent coherent power due to variations in their signal phase.

4.3 Sensitivity

In Chapter 2 the sensitivity of a blind search for gamma-ray pulsations was defined as the minimum pulsed fraction of the observed photon flux that can be detected with a fixed probability, P_{det}^* , and with a fixed false-alarm probability, P_{FA}^* . We now apply this definition to investigate the sensitivity of our search to each source in the survey.

The quantity of interest is the fraction of the background-subtracted weighted photon flux that is pulsed, denoted p_s . Given a set of photon weights, the fractions of the weighted photon counts that




Figure 4.2: Gamma-ray pulse profile of PSR J2017+3625 illustrating the definition of the pulsed fraction. The blue dashed line indicates the background level, b . The pulsed fraction is defined by the area under the template pulse profile above its lowest level (the orange shaded area), divided by the source fraction, s (that is the sum of blue and orange areas).

can be attributed to the source, s , and background fraction b , are estimated as [80, 194]

$$s = \frac{\sum_{j=1}^N w_j^2}{\sum_{j=1}^N w_j}, \quad b = 1 - s. \quad (4.12)$$

The probability of the j -th photon being pulsed is $w_j p_s$, and the overall weighted pulse profile takes the form

$$g(\Phi) = \frac{b}{2\pi} + s g_s(\Phi), \quad (4.13)$$

where $g_s(\Phi)$ is the background-subtracted pulse profile,

$$g_s(\Phi) = \frac{1 - p_s}{2\pi} + p_s g_p(\Phi), \quad (4.14)$$

where $g_p(\Phi)$ is the pulse profile after subtracting all unpulsed emission (background or otherwise). These quantities are illustrated in Figure 4.2.

Note that this definition of the pulsed fraction is equal to the area under the pulse in the normalised pulse profile, as opposed to the ‘‘rms pulsed flux’’ used by e.g. Dib et al. [208], which

is additionally dependent on the shape of the pulse profile. While the rms pulsed flux provides a measure of the power of pulsations, this does not provide a physically meaningful measure of the proportion of pulsed flux emitted by the pulsar [209].

For a pulsar to be detected by this survey, its signal must be strong enough to enter the short list of semicoherent candidates within the work unit that covers the region of the parameter space in which the signal lies. That is, the measured value of \hat{S}_1 at the grid point nearest the signal's location in the parameter space must be greater than the lowest value in the short list of candidates, \hat{S}_1^* . The probability that a signal with a pulse profile described by the complex Fourier coefficients $\{\gamma_n\}$ will be detected by the survey, as a function of the pulsed fraction is

$$P_{\text{det}}(p_s | \{\gamma_n\}) = \int_{-\infty}^{\infty} P(\hat{S}_1^* < \hat{S}_1) p(\hat{S}_1 | p_s, \{\gamma_n\}) d\hat{S}_1, \quad (4.15)$$

where $P(\hat{S}_1^* < \hat{S}_1)$ is the (empirically measured) probability that \hat{S}_1^* is less than \hat{S}_1 , and $p(\hat{S}_1 | p_s, \{\gamma_n\})$ is the probability density function of the measured value of \hat{S}_1 for a signal at a random location within the searched parameter space (see Appendix 4.7.2 for the derivation of this distribution). Each of these quantities depends additionally on the set of photon weights for each source; we have omitted these dependencies from Equation (4.15) for readability. This definition of the detection probability is illustrated in Figure 4.3. This equation can be solved numerically to recover the minimum pulsed fraction, p_s^* , that can be detected at a given probability.

The purpose of the coherent follow-up stage of the survey is to greatly improve the significance of any true signal that is detected by the first stage, and we apply our final false-alarm probability threshold to the candidates from this stage. For typical values of p_s given by solving Equation (4.15), the expected coherent power corresponds to an extremely significant signal ($\mathcal{P}_1 \approx 270$, $P_{\text{FA}} \sim 10^{-59}$) and hence even a conservative false-alarm threshold has no real effect on the overall sensitivity estimate. In practice, effects such as glitches or timing noise that are not included in our simplified isolated pulsar phase model can severely reduce the observed coherent power, resulting in true signals with large p_s appearing with low significance. We attempt to mitigate these effects somewhat by monitoring the twenty most significant candidates from both the semicoherent stage and the follow-up stages rather than applying the false-alarm threshold rigorously.

In Table 4.1 we estimate the minimum pulsed fraction, p_s^* , that can be detected with $P_{\text{det}}^* = 0.95$ for each source in our survey, averaged over the pulse profiles from the 30 most significant pulsars in the 2PC, and assuming constant signal parameters (i.e. no glitches or significant timing noise). For sources in which a new gamma-ray pulsar was discovered we also report the measured pulsed fraction for illustration. Note that while some sources have $p_s^* > 1$, this does not necessarily preclude the detection of pulsed emission from this source, since the assumptions on which these estimates are based might not always apply. For example, the true pulse profile could be narrower than average or the photon weights might not accurately represent the probability of each photon coming from the target source. The results of this sensitivity estimation displayed in Figure 4.4 therefore also illustrate the range of thresholds covered by the various pulse profiles observed.




Figure 4.3: Illustration of the definition of the detection probability. The orange line indicates the pdf of the semicoherent detection statistic (including mismatch) in the presence of a signal with fixed pulse profile and pulsed fraction. The blue dashed line shows the empirical probability that a true signal resulting in a detection statistic \hat{S}_1 will be followed-up (and hence detected). The detection probability is therefore the area under the product of these functions, shown by the grey shaded area.

Table 4.1. Pulsed fraction upper limits

3FGL Name	Pulsar Name	Max. ^a p_g^*	p_s	$N_{\text{FFT}}^b/10^6$	3FGL Name	Pulsar Name	Max. ^a p_g^*	p_s	$N_{\text{FFT}}^b/10^6$
J0002.6+6218	J0002+6216	0.82	0.96	38.78	J1726.6-3530c	...	> 1	...	292.96
J0212.1+5320	...	0.68	...	73.31	J1736.0-2701	...	> 1	...	160.02
J0223.6+6204	...	0.65	...	304.66	J1740.5-2642	...	> 1	...	56.71
J0225.8+6159	...	0.96	...	400.20	J1740.5-2726	...	> 1	...	177.86
J0359.5+5413	J0359+5414	0.78	0.87	258.10	J1740.5-2843	...	0.87	...	117.70
J0426.7+5437	...	0.61	...	45.44	J1742.6-3321	...	> 1	...	231.76
J0541.1+3553	...	> 1	...	776.03	J1744.1-7619	J1744-7619^c	0.66	0.72	275.49
J0631.6+0644	J0631+0646	0.91	0.90	10.11	J1745.1-3011	...	0.94	...	165.41
J0634.1+0424	...	0.56	...	1140.90	J1745.3-2903c	...	0.40	...	16.65
J0744.1-2523	...	0.87	...	26.92	J1746.3-2851c	...	0.48	...	17.16
J0854.8-4503	...	0.87	...	579.82	J1747.0-2828	...	0.60	...	45.69
J0855.4-4818	...	> 1	...	4899.56	J1747.7-2904	...	0.98	...	43.13
J0901.6-4700	...	> 1	...	415.84	J1748.3-2815c	...	0.95	...	23.44
J0907.0-4802	...	> 1	...	860.15	J1749.2-2911	...	> 1	...	193.65
J0933.9-6232	...	0.90	...	133.98	J1754.0-2538	...	0.99	...	11.29
J1026.2-5730	...	0.97	...	713.27	J1754.0-2930	...	0.94	...	79.32
J1035.7-6720	J1035-6720^c	0.69	0.93	353.57	J1758.8-2346	...	> 1	...	21.08
J1037.9-5843	...	> 1	...	748.28	J1800.8-2402	...	0.92	...	12.44
J1039.1-5809	...	> 1	...	7498.83	J1814.0-1757c	...	> 1	...	33.23
J1047.3-6005	...	> 1	...	1607.52	J1814.1-1734c	...	0.98	...	43.69
J1048.2-5928	...	> 1	...	651.27	J1823.2-1339	...	0.71	...	41.60
J1056.7-5853	J1057-5851	0.85	0.68	1793.05	J1827.3-1446	J1827-1446	0.95	1.00	169.81
J1101.9-6053	...	0.94	...	1676.99	J1831.7-0230	...	> 1	...	624.00
J1104.9-6036	J1105-6037	0.80	0.71	499.71	J1833.9-0711	...	> 1	...	27.47
J1111.9-6038	...	0.42	...	113.63	J1834.5-0841	...	> 1	...	208.28
J1112.0-6135	...	> 1	...	2341.02	J1839.3-0552	...	0.78	...	104.52
J1119.9-2204	...	0.62	...	13.72	J1840.1-0412	...	> 1	...	145.45
J1139.0-6244	...	> 1	...	42.51	J1843.7-0322	...	0.84	...	747.26
J1208.4-6239	J1208-6238^d	0.77	0.53	92.51	J1844.3-0344	J1844-0346	0.95	0.88	318.79
J1212.2-6251	...	> 1	...	61.67	J1848.4-0141	...	0.71	...	727.88
J1214.0-6236	...	0.81	...	146.94	J1849.4-0057	...	0.86	...	176.94
J1306.4-6043	...	0.75	...	161.88	J1850.5-0024	...	> 1	...	458.08
J1317.6-6315	...	> 1	...	314.90	J1852.8+0158	...	0.95	...	379.46
J1329.8-6109	...	> 1	...	42.10	J1855.4+0454	...	> 1	...	95.62
J1345.1-6224	...	> 1	...	182.65	J1857.2+0059	...	> 1	...	220.81
J1350.4-6224	J1350-6225	> 1	1.00	100.80	J1857.8+0129c	...	> 1	...	109.05
J1358.5-6025	...	0.87	...	199.81	J1857.9+0210	...	0.89	...	459.00

Table 4.1 (cont'd)

3FGL Name	Pulsar Name	Max. ^a p_s^*	p_s	$N_{\text{FFT}}^b/10^6$	3FGL Name	Pulsar Name	Max. ^a p_s^*	p_s	$N_{\text{FFT}}^b/10^6$
J1405.4–6119	...	0.58	...	198.20	J1857.9+0355	...	> 1	...	270.39
J1503.5–5801	...	> 1	...	459.50	J1859.6+0102	...	> 1	...	120.74
J1528.3–5836	J1528–5838	> 1	0.98	21.22	J1900.8+0337	...	> 1	...	354.12
J1539.2–3324	...	> 1	...	7.42	J1901.1+0728	...	> 1	...	278.80
J1549.1–5347c	...	0.84	...	663.95	J1906.6+0720	J1906+0722^e	0.58	0.77	206.02
J1552.8–5330	...	> 1	...	930.32	J1910.9+0906	...	0.38	...	47.32
J1620.0–5101	...	> 1	...	384.50	J1919.9+1407	...	> 1	...	383.20
J1622.9–5004	J1623–5005	0.76	0.74	78.62	J1925.4+1727	...	> 1	...	1704.54
J1624.2–4041	J1624–4041	0.74	0.85	192.42	J1928.9+1739	...	> 1	...	2680.58
J1625.1–0021	...	0.69	...	154.54	J2004.4+3338	...	0.86	...	84.87
J1626.2–2428c	...	> 1	...	10.64	J2017.9+3627	J2017+3625	0.53	0.64	178.21
J1636.2–4709c	...	> 1	...	405.59	J2023.5+4126	...	> 1	...	776.94
J1636.2–4734	...	0.66	...	321.30	J2032.5+3921	...	> 1	...	541.70
J1638.6–4654	...	0.89	...	102.88	J2034.6+4302	...	> 1	...	1601.98
J1639.4–5146	...	0.77	...	9.24	J2035.0+3634	...	> 1	...	26.52
J1641.1–4619c	...	> 1	...	14.19	J2038.4+4212	...	> 1	...	845.17
J1650.0–4438c	...	> 1	...	161.85	J2039.4+4111	...	> 1	...	525.98
J1650.3–4600	J1650–4601	0.74	0.78	139.09	J2039.6–5618	...	0.74	...	91.84
J1652.8–4351	...	> 1	...	741.14	J2041.1+4736	...	> 1	...	170.62
J1702.8–5656	...	0.52	...	96.18	J2042.4+4209	...	> 1	...	5359.83
J1714.5–3832	...	0.48	...	67.93	J2112.5–3044	...	0.69	...	22.79
J1718.0–3726	...	0.96	...	2.56	J2323.4+5849	...	0.57	...	63.63

Note. — Sources in which pulsars were discovered by the *Einstein@Home* survey are shown in bold. For some of these pulsars, the measured pulsed fraction is well below our estimated upper limit. This can be due to the pulsar having a narrower-than-average pulse profile, a very low spin frequency (at which the sky grid, constructed for the highest frequency in the search band greatly overcovers the search region), or simple “luck” in that the signal lay close to one of our search points and had a lower-than-average mismatch. The 95% detection probability requirement therefore results in conservative limits.

^aEstimated values for the pulsed fraction above which we expect to detect a signal from each source with 95% probability.

^bNumber of FFTs required to search the entire parameter space for each source.

^cThe discoveries and analyses of PSRs J1035–6720 and J1744–7619 are presented in Chapter 6.

^dThe discovery and analysis of PSR J1208–6238 is presented in Chapter 5.

^eThe discovery and analysis of PSR J1906+0722 is presented in Chapter 3.




Figure 4.4: Cumulative fraction of the total computing cost of the survey, as a function of the estimated pulsed fraction threshold. The solid line shows the 95% detection probability pulsed fraction threshold averaged over the pulse profiles of the 30 pulsars in the 2PC with the highest significance. The shaded area shows the range of pulsed fraction thresholds over these pulse profiles. The dashed vertical line denotes the maximum expected pulsed fraction of $p_s^* = 1$.

While the semicoherent search stage is rather robust to the presence of timing noise, any large enough deviation from our simplified constant spin-down model will significantly affect the sensitivity of our search. Glitching pulsars, for example, are particularly difficult to detect, as their large jumps in spin frequency prevent the S/N from accumulating steadily throughout the observations. Pulsars in binary systems are all but impossible to find in a search for isolated pulsars. Our pulsed fraction thresholds therefore only represent our sensitivity to well-behaved isolated pulsars. In particular, our sensitivity estimates are likely to be most reliable for sources far from the Galactic plane, where we would expect to find old, stable millisecond pulsars. Our sensitivity estimates rely on the assumption that a pulsar lies within the sky region in which we search. As mentioned in Section 4.2.2, in some cases the extension of this region may have been underestimated.

We have also not considered the fact that our sensitivity is not exactly constant throughout the parameter space. Rather, the grid of sky locations slightly overcovers the region at low frequencies and, due to the lower number of sky points required at low frequencies, work units searching the lowest frequency band often search at a smaller number of trials. The result is that our survey is sometimes slightly more sensitive at low spin frequencies. The results of Table 4.1 have been

averaged over all spin frequencies, and assume that the entire parameter space is equally well-covered by search points.

In the following sections we highlight and discuss the implications of our measured pulsed fraction upper limits for three sources with well-known counterparts from observations at other wavelengths, from which no pulsations were detected by our survey.

4.3.1 Pulsed Fraction Upper Limit for W49B

Chandra observations of the supernova remnant (SNR) W49B (3FGL J1910.9+0906), believed to be the remnant of a jet-driven, core-collapse supernova, place strong upper limits on the presence of a neutron star [210].

Gamma-rays from W49B have been detected at energies far higher than observed from a typical gamma-ray pulsar [211, 212], indicating that the majority of the high energy flux from W49B comes from the SNR itself. Any pulsed emission from a gamma-ray pulsar would therefore likely represent only a small fraction of the photon flux. Our results place a 95% pulsed fraction limit of 38% of the weighted photon flux from this source.

Placing W49B at the distance of 10 kpc obtained by Zhu et al. [213] suggests a gamma-ray luminosity of $\sim 2 \times 10^{36}$ erg s⁻¹. In order to provide a significant fraction of this emission, any gamma-ray pulsar would have to have a very large spin-down power, since the efficiency of converting spin-down power into gamma-ray luminosity tends to be lower for energetic pulsars [80]. The estimated age of the SNR is in the range 1000–4000 yr. Together, these observations suggest that any potential gamma-ray pulsar would be very young and extremely energetic, and would likely exhibit large timing noise and glitches as a result. This would seriously reduce the sensitivity of our blind search of this target, making our upper limit estimate for the pulsed flux unreliable for this source.

To check for signals with large timing noise, we manually followed up semicoherent candidates from this source using refinement steps with increasing lag-window lengths, but none revealed a significant pulsed signal.

4.3.2 Pulsed Fraction Upper Limit for Cassiopeia A

The SNR Cassiopeia A (Cas A) contains a point-like, central X-ray source, most likely a neutron star [214], from which no pulsations have been detected in X-rays, gamma rays or radio observations. The gamma-ray spectrum for this source is also unlike that of any pulsar, again suggesting that any pulsed emission would likely only contribute a fraction of the total observed flux. The position of this central compact object (CCO) is within our search region for 3FGL J2323.4+5849, a source for which our survey sets a pulsed fraction upper limit of $p_s^* = 57\%$. The photon flux above 100 MeV from this source was $3.1(2) \times 10^{-8}$ photons cm⁻² s⁻¹, making our 95% upper limit more than an order of magnitude lower than the 5σ limit for pulsed flux reported in Abdo et al. [215]. A dedicated search for pulsations at the known position of the Cas A neutron star, excluding photons above typical pulsar emission energies, could likely bring this limit down further.

However, since Cas A is known to be a young SNR (the supernova occurred around A.D. 1680), if the NS is indeed a pulsar, it will be very energetic and likely have a highly unstable spin, making detection in a blind search extremely challenging even if the pulsed fraction is far higher than our stated upper limit. Indeed, the SNR is young enough that the pulsar’s spin-down could even be outside our search range [216]. Again, we followed up semicoherent candidates from this source, without success.

4.3.3 Pulsed Fraction Upper Limit for the Galactic Center

As a result of intense and difficult-to-model interstellar emission, the area around the Galactic Center (GC) is one of the most complicated, and hence poorly understood regions of gamma-ray emission. Both the 3FGL and the recent First *Fermi*-LAT Inner Galaxy point source Catalog [217] identify several bright point sources within a few degrees of the GC, although some of these could be due to misattributed interstellar emission. Nevertheless, a substantial contribution to the flux from the GC region is expected to come from other astrophysical sources, such as young or millisecond pulsars [e.g. 120, 218, submitted, and references therein], or possibly even annihilating dark matter particles [118, and references therein]. The detection of a gamma-ray pulsar near the GC would have important implications for these two competing interpretations of the GC GeV flux.

The bright fore-/background from the interstellar medium makes blindly searching for pulsars near the GC particularly difficult. In order for one single pulsar to be detectable above this background, it must be extremely bright, especially if it lies at a similar distance as the GC, ~ 8 kpc. As an example, the bright source 3FGL J1745.3–2903 searched during this survey has $p_s^* = 40\%$, corresponding to a pulsed photon flux above 1GeV of $\sim 10^{-8}$ photons $\text{cm}^{-2} \text{s}^{-1}$. This flux is similar to the photon flux that the Crab pulsar would produce if it was at the same distance as the GC. Again, such highly luminous pulsars also exhibit the most timing noise and glitches, further adding to the difficulty of detecting their pulsations above the bright background flux.

4.4 Timing Analysis

In this section, we describe the methods used to precisely measure the spin and positional parameters of the 13 new pulsars discovered by this survey.

For the purpose of these follow-up timing analyses, we produced extended LAT data sets until 2015 September 9 for each of the pulsars newly reported in this work. These updated data sets were produced using the `P8R2_SOURCE_V6` IRFs, `gll_iem_v06.fits` Galactic diffuse emission template [205], and `iso_P8R2_SOURCE_V6_v06.txt` isotropic diffuse background template⁴. The extended data sets had a lower zenith angle cutoff of 90° . The pulsar’s position was fixed at its initial timing position. Photons from within a larger 15° radius were included in the likelihood fitting, which was performed using `gtlike`. Photon weights were then calculated for all photons

⁴<http://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>

from within 5° of the pulsar using `gtsrcprob`. Further details of the preparation of these data will be given in Paper II, including the spectral properties of each newly detected pulsar.

To reduce the number of photons included in the timing analysis for computational efficiency, we applied a photon weight cutoff with the minimum photon weight chosen such that no more than 1% of the maximum coherent Fourier power was lost (again assuming that the maximum coherent S/N is proportional to $\sum_{j=1}^N w_j^2$).

4.4.1 Timing Methods

We performed detailed timing analyses for each new pulsar to precisely determine their sky positions and rotational parameters, again denoted by the vector \mathbf{u} . The analysis follows the procedure described in Section 4.4.1, as an extension of the method described by Ray et al. [167].

Starting from the spin and positional parameters of the pulsar reported by the refinement stage, we phase-folded the photon data to obtain a weighted pulse profile. We also phase-folded at half, and one third of the measured frequency to ensure that the original signal was not a higher harmonic of the fundamental spin frequency. In two pulsars, J1350–6225 and J1624–4041, this revealed sharply double-peaked profiles at half of the original candidate frequency, and greatly increased their measured H -test values, indicative of having identified the true spin frequency.

From the phase-folded data, we constructed a template pulse profile, $\hat{g}_s(\Phi)$, consisting of a combination of symmetrical wrapped Gaussian peaks (as defined in Abdo et al. [80]), which were fit by maximizing the likelihood,

$$\mathcal{L}(\hat{g}_s, \mathbf{u}) = \prod_{j=1}^N [w_j \hat{g}_s(\Phi(t_j, \mathbf{u})) + (1 - w_j)] . \quad (4.16)$$

The number of peaks in the template profile was chosen by the template that minimizes the Bayesian Information Criterion [BIC, 219],

$$BIC = -2 \log(\mathcal{L}(\hat{g}_s, \mathbf{u})) + k \log \left(\sum_{j=1}^N w_j \right) , \quad (4.17)$$

where k is the number of free parameters in the model. Because each Gaussian peak consists of three parameters (central phase, width and amplitude), when fitting the template pulse profile, $k = 3N_{\text{peaks}}$. Due to the presence of the second term in Equation (4.17), a new component was only added to the template profile if its presence significantly increased the likelihood. It therefore acts as a penalty factor, discriminating against a template profile featuring many “spiky” components, unless this is warranted by the data. The parameters of the template pulse profiles used to time each pulsar are given in Table 4.2, and the profiles themselves are shown in Figure 4.5.

After obtaining the template pulse profile, we varied the positional and spin parameters and explored the resulting multi-dimensional likelihood surface to find the most likely parameter values. To explore the likelihood surface, we used the Affine Invariant Monte Carlo (AIMC) method

described by Goodman & Weare [220], in which many Monte Carlo chains are run in parallel, with proposal jumps for each chain depending on the locations of the other chains⁵. We used the scheme described by Foreman-Mackey et al. [221] to efficiently parallelize the likelihood computations amongst several CPU cores.

With the new parameter values we re-folded the photon data to obtain a new template pulse profile as above. Additional parameters could then be added to the timing model, and the procedure was repeated. For each pulsar, we started from the simplified timing model (i.e. up to \dot{f}), added higher frequency derivative terms, found the most likely parameters, and updated the template pulse profile until the BIC of Equation 4.17 (with k now as the number of parameters in the timing model) stopped decreasing. With the number of timing parameters selected in this manner, we performed a final longer Monte Carlo run, using a large number of chains, to obtain precise estimates of the mean value and uncertainty of each parameter.

With over 100,000 photons whose individual barycentric arrival times must be computed each time, each likelihood evaluation is relatively computationally expensive. Hence, efficient convergence of the Monte Carlo step is crucial to allow us to perform the timing analysis in a reasonable amount of time. To avoid the possibility of chains getting stuck in low-likelihood regions, we start all of the chains in a tight ball near our current most-likely point, as advocated by Foreman-Mackey et al. [221]. Using a lengthy burn-in period, we allow these chains to spread out throughout the most likely regions of our parameter space. While this initialization can in principle lead to the Monte Carlo sampling only reaching a local likelihood maximum rather than exploring the full parameter space to find a global maximum, visual inspection of the phase-folded photon data can typically reveal any significant residuals in the timing solution requiring further fitting.

The results of these analyses are summarized in Table 4.3. The physical properties of each of the new pulsars, as derived from their spin frequency and spin-down rate are given in Table 4.4.

These timing solutions allow for sensitive follow-up searches, the identification of candidate multiwavelength counterparts, and phase-resolved spectroscopy of the on- and off-pulse photons. Dedicated radio observations of the newly discovered pulsars were also performed, which used the gamma-ray timing solution to fold the data. These analyses are described in Paper II.

4.4.2 Spin-down vs. Timing Noise

The long-term spin-down behavior of a pulsar can be characterized by the braking index [e.g., 3], n , where,

$$\dot{f} \propto -f^n, \quad (4.18)$$

$$n = \frac{f \ddot{f}}{\dot{f}^2}. \quad (4.19)$$

⁵In Chapter 3 this step was performed using the nested sampling algorithm `MultiNest` [193]. We changed to the AIMC method due to the extra degree of parallelization that it offers, which greatly speeds up the sampling for pulsars near the Galactic plane with many low-weight photons.

Table 4.2. Template Pulse Profile Parameters

Pulsar	a_1	σ_1	$\Delta\mu_2$	a_2	σ_2	$\Delta\mu_3$	a_3	σ_3	$\Delta\mu_4$	a_4	σ_4
J0002+6216	0.32	0.088	-2.476	0.25	0.186	-0.319	0.26	0.217	-1.531	0.68	0.697
J0359+5414	1.36	0.511
J0631+0646	0.29	0.145	-1.821	0.62	0.374	-0.663	0.50	0.408
J1057-5851	0.43	0.301	-1.143	0.65	0.907
J1105-6037	0.19	0.116	-2.026	0.46	0.294	-0.632	0.47	0.866
J1350-6225	0.57	0.077	3.053	0.31	0.079	0.278	0.62	0.453
J1528-5838	0.13	0.053	0.972	1.42	0.762
J1623-5005	0.30	0.151	-2.215	0.17	0.141	-1.438	0.70	0.855
J1624-4041	0.31	0.080	-2.715	0.11	0.075	-2.328	0.50	0.383	-0.072	0.42	0.604
J1650-4601	0.38	0.210	2.079	0.33	0.211	1.227	0.52	0.571
J1827-1446	1.44	0.311	-1.581	0.13	0.182
J1844-0346	1.38	0.467
J2017+3625	0.37	0.200	2.356	0.25	0.160	1.843	0.39	0.494

Note. — Columns 2–12 give the amplitude (a_i), offset in radians from the tallest component ($\Delta\mu_i$), and width parameter (σ_i) for each wrapped Gaussian component in the template pulse profile.

Table 4.3. Pulsar timing parameters

Pulsar	t_{ref} (MJD)	R.A.	Decl.	f (Hz)	\dot{f} (10^{-12} Hz s $^{-1}$)
J0002+6216	55806	00 ^h 02 ^m 58 ^s .17(2)	+62° 16' 09''4(1)	8.6682478274(1)	-0.448354(5)
J0359+5414	55716	03 ^h 59 ^m 26 ^s .01(2)	+54° 14' 55''7(3)	12.5901403227(2)	-2.65247(1)
J0631+0646	55806	06 ^h 31 ^m 52 ^s .38(2)	+06° 46' 14''2(7)	9.01071834910(6)	-0.293694(2)
J1057-5851	55716	10 ^h 57 ^m 09 ^s .5(2)	-58° 51' 07(2)''	1.6119541713(3)	-0.26135(1)
J1105-6037	55716	11 ^h 05 ^m 00 ^s .48(4)	-60° 37' 16''3(3)	5.12982912390(8)	-0.574649(2)
J1350-6225	55806	13 ^h 50 ^m 44 ^s .45(1)	-62° 25' 43''8(1)	7.23810134280(6)	-0.465408(2)
J1528-5838	55806	15 ^h 28 ^m 24 ^s .3(1)	-58° 38' 01(1)''	2.81146362521(6)	-0.195700(1)
J1623-5005	55716	16 ^h 23 ^m 04 ^s .11(1)	-50° 05' 15''1(2)	11.7547287226(1)	-0.574965(3)
J1624-4041	55716	16 ^h 24 ^m 09 ^s .927(9)	-40° 41' 29''7(3)	5.95730476591(3)	-0.1676839(9)
J1650-4601	55716	16 ^h 50 ^m 18 ^s .62(2)	-46° 01' 18''6(4)	7.8664037135(1)	-0.937157(3)
J1827-1446	55716	18 ^h 27 ^m 24 ^s .60(5)	-14° 46' 28(4)''	2.0032588600(1)	-0.181932(3)
J1844-0346	55716	18 ^h 44 ^m 32 ^s .89(2)	-03° 46' 30''6(9)	8.8609552273(8)	-12.14675(5)
J2017+3625	55716	20 ^h 17 ^m 55 ^s .84(1)	+36° 25' 07''9(2)	5.99703102436(3)	-0.0489063(8)

Note. — Reported values of timing parameters are the mean values obtained from the Monte-Carlo analysis described in Section 4.4.1 at the reference epoch, t_{ref} , with 1σ uncertainties in the final digits given in brackets. Two different observation spans were used during this survey and the reference epochs were chosen to lie at the middle of the observation, hence the two distinct values shown in column 2.




Figure 4.5: Weighted pulse profiles of all pulsars reported in this chapter. The template pulse profiles used for timing analyses are shown by orange curves. The background fraction is illustrated by the dashed blue line in each plot.

Table 4.4. Derived pulsar properties

Pulsar	l ($^{\circ}$)	b ($^{\circ}$)	P (ms)	\dot{P} (10^{-15} s s $^{-1}$)	τ_c (kyr)	\dot{E} (10^{33} erg s $^{-1}$)	B_S (10^{12} G)
J0002+6216	117.33	-0.07	115.363568268(2)	5.96703(7)	306	153	0.8
J0359+5414	148.23	+0.88	79.427232292(1)	16.73359(7)	75	1318	1.2
J0631+0646	204.68	-1.24	110.9789432160(7)	3.61723(2)	486	104	0.6
J1057-5851	288.61	+0.80	620.3650313(1)	100.583(5)	98	17	8.0
J1105-6037	290.24	-0.40	194.938267113(3)	21.83720(6)	141	116	2.1
J1350-6225	309.73	-0.34	138.157778213(1)	8.88352(4)	246	133	1.1
J1528-5838	322.17	-1.75	355.686622097(8)	24.7586(2)	228	22	3.0
J1623-5005	333.72	-0.31	85.0721461635(8)	4.16118(2)	324	267	0.6
J1624-4041	340.56	+6.15	167.861145148(1)	4.72489(2)	563	39	0.9
J1650-4601	339.78	-0.95	127.122893310(2)	15.14468(6)	133	291	1.4
J1827-1446	17.08	-1.50	499.18661037(3)	45.3351(9)	174	14	4.8
J1844-0346	28.79	-0.19	112.85464991(1)	154.7031(6)	12	4249	4.2
J2017+3625	74.51	+0.39	166.7491790419(8)	1.35985(2)	1943	12	0.5

Note. — Columns 2 and 3 give the pulsars’ Galactic longitudes (l) and latitudes (b) respectively. Columns 4 and 5 give the derived spin period ($P = 1/f$) and period derivative ($\dot{P} = -\dot{f}/f^2$). Characteristic ages, τ_c , spin-down luminosities, \dot{E} , and surface magnetic field strengths, B_S are calculated as described in Abdo et al. [80].

The exact value of the braking index depends on the physical mechanism causing the pulsar to spin down; a pulsar whose braking is entirely due to magnetic dipole radiation will have $n = 3$, whereas one whose spin-down power is entirely due to the radiation of gravitational waves will have $n = 5$ or $n = 7$ [216].

The vast majority of pulsars, however, also exhibit red-spectrum “timing noise”, manifesting as low-frequency quasi-periodic variations in the arrival times of pulses [e.g., 16, 165]. The amplitude of this timing noise appears to correlate with the spin-down energy, \dot{E} , which is typically higher for gamma-ray pulsars than the rest of the pulsar population. For all but the youngest pulsars or those with the highest magnetic fields, this timing noise obscures the true long time-scale braking behavior.

In all pulsars measured here, n deviates significantly from any of these integer values (except for PSR J1650-4601, where the index is low, but poorly constrained), suggesting that the measured values of \ddot{f} are due to short time-scale timing noise. For pulsars with measurable frequency derivative terms beyond the first derivative, the evolution of the spin frequency and spin-down rate is shown in Figure 4.6.

4.4.3 Timing Large Glitches

In addition to slowly varying timing noise behavior, young pulsars also occasionally exhibit large, sudden changes in their spin frequency, known as “glitches” [e.g. 34]. Pulsars with large glitches are particularly difficult to detect in blind searches, which require long intervals containing a stable




Figure 4.6: Evolution of each pulsar’s rotational frequency during the *Fermi*-LAT observation time. For each pulsar, the top and bottom panels show the deviations from a constant spin-down model of the frequency, $\Delta f(t)$, and first frequency derivative, $\Delta \dot{f}(t)$, respectively. The shaded areas show 1σ uncertainty regions. These deviations are most likely due to the pulsars’ intrinsic timing noise, as suggested by their unphysical braking indices (n , described in Section 4.4.2 and shown in the lower panel for each pulsar).

signal to accumulate sufficient S/N.

Not only are pulsars with large glitches harder to detect in a blind search, they are also significantly more complicated to time. A glitch occurring at time t_g , with increments in the spin frequency (Δf_g), the first two frequency derivatives ($\Delta \dot{f}_g$, $\Delta \ddot{f}_g$) and an exponentially decaying frequency increment ($\Delta f_{D,g}$, decay timescale $\tau_{D,g}$) causes a phase offset at time $t > t_g$ of

$$\Delta \Phi_g(t) = 2\pi \left[\Delta f_g(t - t_g) + \frac{\Delta \dot{f}_g}{2}(t - t_g)^2 + \frac{\Delta \ddot{f}_g}{6}(t - t_g)^3 + \Delta f_{D,g} \tau_{D,g} \left(1 - e^{-(t-t_g)/\tau_{D,g}} \right) \right]. \quad (4.20)$$

If the parameter increments associated with the glitch are large enough then $\Delta \Phi_g$ can quickly exceed a small integer number of rotations. If the photon data are sparse, as is often the case, this can happen before we have even detected any pulsed photons after the glitch [34, and Chapter 3].

A result is that the likelihood distribution in the glitch epoch, t_g , resembles a comb of possible epochs, with each maximum occurring at a time where $\Delta \Phi_g(t)$ equals an integer number of rotations. Such a highly multi-modal likelihood surface causes problems for our Monte Carlo sampling method, as chains can easily get stuck in low-likelihood regions between maxima, greatly reducing the efficiency of the sampling procedure.

To avoid this, we can include an arbitrary phase increment at the time of the glitch in our phase model, and allow it to vary as part of the Monte Carlo sampling. This phase increment removes the multi-modal nature of the likelihood surface by accounting for the phase offset for glitch models that do not occur at one of the maxima described above. While the phase increment is highly correlated with the glitch epoch, we can predict and account for the dominant contribution to the time-dependent part of the glitch increment to remove this correlation, ensuring efficient sampling (see Appendix 4.7.3 for further details).

Apart from PSR J1906+0722, (described in Chapter 3) one other pulsar detected by the *Einstein@Home* survey, PSR J1844–0346, experienced a glitch during the *Fermi* mission. Occurring in July–August 2012, with $\Delta f/f \approx 3.5 \times 10^{-6}$, it was comparable to some of the largest glitches detected from gamma-ray pulsars [154]. With a characteristic age of $\tau_c = 11.6$ kyr, and spin-down energy $\dot{E} = 4.2 \times 10^{36}$ erg s⁻¹, this pulsar is by far the most energetic pulsar found by our survey, and also displays a correspondingly large degree of timing noise. The evolution of the spin frequency and spin-down rate of PSR J1844+0346, including the glitch, are shown in Figure 4.7. The glitch parameters obtained from the timing analysis are given in Table 4.5.

4.4.4 Pulse Profile Variability

We checked for variations in the pulse profiles of all pulsars detected in this survey by visually inspecting their phase–time diagrams, and by measuring their Fourier coefficients in a small number of time intervals. In one pulsar, PSR J1350–6225, small changes in the first and second Fourier




Figure 4.7: Evolution of the rotational frequency of PSR J1844–0346 during the *Fermi*-LAT observation time, as in Figure 4.6.

coefficients were observed. This was also observed in the phase–time diagram for this pulsar (shown in Figure 4.8) where the trailing peak seems to appear less prominently in the latter portion of the *Fermi* mission than in the earlier data.

To investigate the significance of this effect, we compared the distributions of the observed photon phases before and after MJD 55750 (chosen to maximize the change in the pulse profile’s first Fourier coefficient) by calculating the Weighted Kuiper Test statistic [222],

$$V_{w_1, w_2} = \max \{G_1(\Phi) - G_2(\Phi)\} + \max \{G_2(\Phi) - G_1(\Phi)\}, \quad (4.21)$$

where $G_1(\Phi)$ and $G_2(\Phi)$ are the empirical weighted cumulative distributions of the photon phases before and after MJD 55750 respectively.

The distribution of V_{w_1, w_2} under the null hypothesis is unknown, and calculating significances based on the properties of the unweighted statistic always *under-estimates* the false-alarm rate. To estimate the significance, we therefore performed a Monte Carlo analysis. Using the observed sets of photon weights (before and after MJD 55750), we randomly generated two sets of photon phases,




Figure 4.8: Pulse profile of PSR J1350–6225 showing a possible change around MJD 55750. Left panel: phase–time diagram of the detected photons from PSR J1350–6225, with photon probability weight indicated by the intensity. MJD 55750 is marked by the dashed red line. Right panels, bottom and top: normalised phase histograms showing the total pulse profile before ($g_1(\Phi)$) and after MJD 55750 ($g_2(\Phi)$) respectively. The estimated background levels, b , are shown by the dashed blue line.

Table 4.5. PSR J1844–0346 Glitch Parameters

Parameter	Value
Glitch epoch, t_g (MJD)	56135(7)
Frequency increment, Δf_g (Hz) . . .	$3.06(1) \times 10^{-5}$
\dot{f} increment, $\Delta \dot{f}_g$ (Hz s $^{-1}$)	$-9.4(3) \times 10^{-14}$
\ddot{f} increment, $\Delta \ddot{f}_g$ (Hz s $^{-2}$)	$-7.0(9) \times 10^{-22}$
Decaying f increment, $\Delta f_{D,g}$ (Hz)	$4.5(7) \times 10^{-7}$
Decay time constant, $\tau_{D,g}$ (d)	117(22)

Note. — Reported values of glitch parameters are the mean values obtained from the Monte-Carlo analysis described in Section 4.4.1 with 1σ uncertainties in the final digits given in brackets.

with the j -th photon’s phase drawn from a common pulse profile⁶ with the probability $p = w_j$, otherwise distributed uniformly. This process was repeated many times to estimate the distribution of V_{w_1, w_2} under the null hypothesis.

From this analysis, we find that the observed value of V_{w_1, w_2} corresponds to a p -value of 0.038. Given that all 13 pulsars were checked for pulsations, and that a small number of trials were performed when choosing the date defining the boundary between the two intervals, we conclude that this is not a significant variation.

Long-term monitoring of the *Fermi*-LAT data from this pulsar would be required to detect the presence of any pulse profile mode changes, either by observing another such mode change, or by reducing the uncertainty on the new template pulse profile. A change in pulse profile has only been detected in one gamma-ray pulsar to date, PSR J2021+4026 [113]. This variation was accompanied by abrupt changes in the pulsar’s gamma-ray flux and spin-down rate, neither of which are observed from PSR J1350–6225.

4.5 Discussion

4.5.1 Sensitivity

Recent publications have argued, both by modeling the emission mechanisms of known radio and gamma-ray pulsars [223] and by constructing an unbiased sampling of radio-loud and radio-quiet gamma-ray pulsars [224], that *Fermi*-LAT should detect significantly more non-recycled gamma-ray

⁶The distribution of V_{w_1, w_2} under the null hypothesis should be independent of the chosen pulse profile, since it only tests the possibility that the observed phases are drawn from the *same* distribution, regardless of the true underlying distribution.

pulsars that are radio quiet than are radio loud, by a factor of ~ 2 . The 13 new pulsars reported here, only two of which appear to be radio loud (see Paper II) bring the total number of radio-quiet and radio-loud non-recycled gamma-ray pulsars to 51 and 61 respectively⁷. This would suggest that there are still large numbers of unidentified radio-quiet gamma-ray pulsars requiring blind gamma-ray pulsation searches to be detected. In this section, we compare the newly discovered pulsars to the earlier population of blind-search discovered gamma-ray pulsars to identify and discuss the overall trends in blind search sensitivity.

One question that we may like to address is, how bright does an unidentified gamma-ray pulsar need to be in order to be detectable in a blind search? In particular, we would like to know the lowest point source significance within which pulsations can be detected, and how this threshold changes as more data are accumulated.

In Figure 4.9, we have plotted the point source significance vs H -test value for each pulsar in the 2PC, as well as the blind search pulsars detected after this catalog was produced [136], and those discovered by this survey, including PSR J1906+0722 from Chapter 3. We can see that the H -test value for a pulsar can be well approximated by its point source Test Statistic (TS) value (as shown by the dashed line in Figure 4.9). The values plotted here for pulsars detected in previous blind searches have been scaled back to represent their value at the time of their discovery.

We begin by looking at how the coherent detection statistic, \mathcal{P}_1 , varies with the observation duration, T_{obs} . Since the H -test is a maximized sum over \mathcal{P}_n values, the relevant scalings with respect to the observation time will be unchanged. As shown in Section 2.3.1, the expected values of \mathcal{P}_1 for a signal with average photon arrival rate, μ , pulse profile Fourier coefficients, $\{\gamma_n\}$, and pulsed fraction p_s over an observation lasting T_{obs} is approximately

$$E_p[\mathcal{P}_1] \approx 2p_s^2 \mu |\gamma_1|^2 T_{\text{obs}} + 2. \quad (4.22)$$

This shows the well-known result that \mathcal{P}_1 (and hence H) increases linearly with time. This is relevant for detecting gamma-ray pulsations using known radio or X-ray ephemerides; since only a small number of trials are required, fully coherent searches are perfectly feasible and signals only need to overcome a low threshold to be detected. As T_{obs} increases, so too does the point source significance, and pulsars whose gamma-ray pulsations are not initially above the detection threshold will eventually be detectable. Indeed, pulsations have been detected in this way from sources all the way down to the point source detection threshold [110].

However, the limiting factor in our blind searches is the sensitivity of the initial semicoherent stage. The expected semicoherent (power) S/N, given a lag-window length, T , is

$$\theta_{S_1}^2 = E_p[S_1] \approx p_s^2 \mu |\gamma_1|^2 T^{1/2} T_{\text{obs}}^{1/2}.$$

The semicoherent S/N accumulates much more slowly, only with the square root of T_{obs} . Substituting this into Equation (4.22), we can identify the effective coherent threshold, \mathcal{P}_1^* , in terms of the

⁷We use the definition from the 2PC that a radio-loud pulsar has a flux density $S_{1400} > 30 \mu\text{Jy}$. Two gamma-ray pulsars have radio detections with lower fluxes, we count them here as radio-quiet.




Figure 4.9: Point source significance vs. H -test value for all pulsars in 2PC, and those detected by the *Einstein@Home* survey. The dashed line denotes where $H = TS$. The H -test values and significances for previously detected blind-search pulsars have been scaled back from their 2PC values to estimate their value in the data used in their original searches.

semicoherent S/N threshold, S_1^* , as

$$\mathcal{P}_1^* \approx 2 S_1^* \left(\frac{T_{\text{obs}}}{T} \right)^{\frac{1}{2}} + 2. \quad (4.23)$$

This reveals an unintuitive result: with a fixed semicoherent threshold and lag-window size, as the observation time increases our sensitivity threshold in terms of the coherent signal power (and hence point source significance) actually *increases*.

Equation (4.23) also reveals the solution to this problem: if we are to maintain the same source significance threshold in searches using longer observation times, we must also increase the lag-window duration by the same factor. However, as was derived in Section 2.4.2, the computational cost associated with a blind semicoherent search scales with $T^4 T_{\text{obs}}$.

In Figure 4.10, we have estimated the semicoherent S/N at the time of discovery for all gamma-ray pulsars detected in previous blind searches by calculating their \mathcal{P}_1 values from the data provided by the 2PC, and scaling these down to the T_{obs} used in each search. We have also estimated the computational cost that would be required to perform each search (only covering the young pulsar parameter space) using the original lag-window size ($T = 2^{20}$ s) and observation length, but with otherwise the same search scheme described in Section 4.2. Searches prior to Pletsch et al. [151] only searched for pulsations from the LAT point source location, rather than searching over many




Figure 4.10: Computing cost vs. Semicohherent signal-to-noise ratio comparison for all LAT-discovered gamma-ray pulsars. Those newly discovered in this work are shown by orange squares, previously detected pulsars are shown as blue circles. The locations on this plot of the newly detected pulsars represent their expected location if our survey had used the 50% shorter lag-window used in previous searches. The lowest semicoherent S/N for all blind-search pulsars is given by the blue dashed line, we consider signals below this line to be undetectable. The orange diagonal lines show how the computational cost and S/N increased for the newly discovered pulsars due to the longer lag window used in this survey.

possible sky locations. While this significantly decreases the required computational cost, this requires pulsars either to be close to the LAT source’s estimated position, or to have a strong pulsed signal, such that they can be detected despite a large positional offset. This also rules out the detection of isolated millisecond pulsars, whose high frequency requires us to search sky locations with much finer resolution. We therefore exclude the cost of searching for millisecond pulsars from our estimated computing costs.

In this figure, we have also estimated the computational cost and S/N that each of the new pulsars would have had if we had performed these searches with the lag-window duration used by previous searches. At least 4, possibly as many as 7 pulsars would have been below the lowest semicoherent S/N from a previously detected pulsar, and therefore most likely would not have been detected had we not performed the more expensive search with the longer lag window.

Figure 4.10 also shows how the computing cost and S/N increased for these pulsars as a result of using a lag-window with double the length. We see that only one pulsar detected here required less computing time than any previously detected gamma-ray pulsar, with the most expensive detection being more than 100 times more costly. Based on this argument, the current survey should have

taken many years to complete, even taking into account a generous estimate of a doubling of the computing power between the two surveys computed on *Einstein@Home*. In actuality, the first 118 sources were searched in just over one year, similar in duration to the first gamma-ray pulsar survey. We consider this a testament to the improved efficiency of the search methods developed in Chapter 2.

The signal detected here with the lowest coherent power was from PSR J1105–6037 with $\mathcal{P}_1 = 165$. From the spread about the diagonal of Figure 4.9, we estimate that a pulsar with a particularly narrow pulse profile should be detectable with a TS of around 100 (a significance of approximately 10σ)⁸. The lowest-significance point source in which a pulsar was detected here (PSR J1350–6225) had a TS of 356.6 (approximate significance 19σ). This could be because the ranking method for our target sources depends on the sources’ significance; brighter *Fermi*-LAT sources were searched earlier than dimmer sources with similar spectral properties. It is therefore possible that as our survey continues we may yet detect a pulsar with a lower point-source significance.

In any case, the following trend is clear: as the *Fermi* mission continues, dimmer pulsars will gradually become detectable to blind searches, but the computing cost required to detect equally low H -test signals will rise quickly as we search longer data sets. This survey received a large boost to its sensitivity in the form of the “free” (in terms of computational cost) increase in the observed photon counts offered by the Pass 8 data reprocessing [77]. It is unlikely that such a welcome jump in sensitivity will occur again in the near future, and therefore it may become necessary to make some sacrifices to maintain the current search sensitivity without the computing cost requirements becoming unrealistic. For example, reducing the spin-down range to which we are sensitive by a factor of ~ 10 will reduce the computational cost by the same factor, while only losing sensitivity to the most energetic of pulsars. Indeed, one could argue that we already have greatly reduced sensitivity to these pulsars since they typically exhibit the most timing noise and large glitches.

Another option would be to focus our available computing resources to perform more sensitive searches of only the most promising target sources, rather than performing wide-ranging surveys of many unidentified sources as was done here. Indeed, if we look at how the computing cost was distributed amongst sources, as shown in Figure 4.4, we can see that, even with a favorable pulse profile, some sources would require a pulsed fraction far higher than unity for us to detect pulsations. The procedures described in Section 4.3 allow us to reliably predict our search sensitivity to new sources for the first time. Now that we are able to predict this in advance, we can remove such sources from our search list and focus our efforts on improving our sensitivity to the more promising sources. If we relax the requirement of 95% detection probability used in producing Figure 4.4 slightly to reflect a more moderate chance of detecting pulsations from each source, we find that $\sim 25\%$ of the computing cost of this survey could perhaps have been better spent by removing unpromising sources.

⁸Note that one of the latest pulsars to be found by *Einstein@Home*, described in Chapter 7, has a TS of 108.9, making it the faintest pulsar to be found in a blind search, and likely close to the limit of the survey’s sensitivity; this fact was only discovered after this chapter was written.

4.5.2 New Pulsars

The group of pulsars detected here appears to be very consistent with the overall population of young gamma-ray pulsars (see Figure 4.1). Two of the new pulsars reported here (PSR J1057–5851 and PSR J1827–1446) are the most slowly spinning gamma-ray pulsars yet discovered, at 1.6 Hz and 2.0 Hz respectively. While the existence of these pulsars is not in tension with estimates of the gamma-ray pulsar “death line” [88], their discovery does extend the known population of gamma-ray pulsars down to lower spin frequencies. Sensitive blind searches that can detect these pulsars are important to ensure that the entire pulsar population is being fully explored, and to reduce biases inherent in the radio-quiet gamma-ray pulsar population.

Also of interest are the pulsars detected here despite exhibiting significant timing noise, PSR J1844–0346, and PSR J0359+5414. The large contributions of higher frequency derivatives mean that the original phase model used in the blind search could not maintain phase coherence over the full duration of the data. These effects make it very difficult for the coherent follow-up stages to pick up these signals. On the other hand, the semicoherent stage is largely unaffected by timing noise, although large glitches are also detrimental to the semicoherent sensitivity. Noisy pulsars will therefore only appear as semicoherent candidates, and may easily escape detection from our pipeline, which focuses on the results from the final coherent follow-up. Further investigation of the vast number of semicoherent candidates reported by *Einstein@Home* may yet reveal more noisy pulsars lurking in our results.

4.6 Conclusions

We have presented the discovery of 13 new gamma-ray pulsars found by the ongoing *Einstein@Home* survey of unidentified *Fermi*-LAT sources. Amongst these pulsars are two new energetic pulsars with $\dot{E} > 10^{36}$ erg s⁻¹, one of which experienced a large glitch; and the two slowest spinning gamma-ray pulsars yet detected.

As the *Fermi* mission continues and the LAT gathers more data, the sensitivity to weak pulsar signals will increase, and many currently undetectable pulsars could rise above the detection threshold in the near future, although future searches with more data will also require even more computing power to be sensitive to similarly weak signals.

We also placed realistic upper limits on the pulsed flux from point sources from which no pulsations were detected. The framework for this allows us to also predict our sensitivity to other sources, enabling us to identify promising targets for searching, and also to veto sources from which pulsations would be almost impossible to detect. This will allow us to focus our computing power on increasing our sensitivity to the most promising sources in future surveys.

The *Einstein@Home* survey continues to run, and while it is now searching sources that could be considered somewhat less promising than the 118 sources searched in this work, we are still hopeful that more interesting detections can be made by this survey. A further exciting new advancement is the development of the first *Einstein@Home* survey for gamma-ray pulsars in candidate binary systems with well constrained orbital parameters, similar to the search that

discovered PSR J1311–3430 [135]. The additional computing power of *Einstein@Home* will enable more complicated searches, allowing for searches from sources with larger uncertainties in their orbital parameters, or even with slight eccentricities.

Acknowledgments

We are extremely grateful to all volunteers who have donated their CPU time to the *Einstein@Home* project, without whom this survey could not have been performed. We are especially grateful to those users whose computers discovered the 13 new pulsars reported here. They are⁹:

- PSR J0002+6216: James Drews of UW-Madison, WI, USA and Ralph Elwell of Richland, WA, USA;
- PSR J0359+5414: Whelton A. Miller III, Lincoln University of Pennsylvania & University of Pennsylvania, USA; the ATLAS Cluster, AEI, Hannover, Germany and Philip “Dely” Horney of the GPU Users Group, Fort Wright, KY, USA;
- PSR J0631+0646: Katagiri, Atsushi of Kawasaki, Japan and Nicholas Huwar of Houston, TX, USA;
- PSR J1057–5851: Syracuse University HTC Campus Grid¹⁰, NY, USA and Igor Yakushin of Chicago, IL, USA and the LIGO Laboratory, USA;
- PSR J1105–6037: The ATLAS Cluster, AEI, Hannover, Germany and Syracuse University HTC Campus Grid¹⁰, NY, USA;
- PSR J1350–6225: Petr Ruzicka of Brno, Czech Republic and Bryden Kanngiesser of Calgary, Canada;
- PSR J1528–5838: “fred c” and Gabriel Vasquez of Miami, FL, USA;
- PSR J1623–5005: Lars Bollwinkel, of Kiel, Germany and Greg Dorais of Martinez, CA, USA;
- PSR J1624–4041: Xio of NYC and Hung Tran of Chandler, AZ, USA;
- PSR J1650–4601: Syracuse University HTC Campus Grid¹⁰, NY, USA and Eric Schwartz of Vashon Island, WA, USA;
- PSR J1827–1446: The ATLAS Cluster, AEI, Hannover, Germany and Igor Yakushin of Chicago, IL, USA and the LIGO Laboratory, USA;

⁹Where the volunteer’s name is unknown or private, we give the *Einstein@Home* username in quotation marks.

¹⁰Supported by NSF awards ACI-1341006 and ACI-1541396

- PSR J1844–0346: Aurélien FAUCHEUX of Antibes, France and Roger Gulbranson, Ph.D. of Wickliffe, OH, USA;
- PSR J2017+3625: Kurt Gramoll, Ph.D., University of Oklahoma, OK, USA and Michael Brandau, of Kassel, Germany.

This work was supported by the Max-Planck-Gesellschaft (MPG), by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether research grant PL 710/1-1 (PI: Holger J. Pletsch), and by NSF award 1104902.

The *Fermi* LAT Collaboration acknowledges generous ongoing support from a number of agencies and institutes that have supported both the development and the operation of the LAT as well as scientific data analysis. These include the National Aeronautics and Space Administration and the Department of Energy in the United States, the Commissariat à l’Energie Atomique and the Centre National de la Recherche Scientifique / Institut National de Physique Nucléaire et de Physique des Particules in France, the Agenzia Spaziale Italiana and the Istituto Nazionale di Fisica Nucleare in Italy, the Ministry of Education, Culture, Sports, Science and Technology (MEXT), High Energy Accelerator Research Organization (KEK) and Japan Aerospace Exploration Agency (JAXA) in Japan, and the K. A. Wallenberg Foundation, the Swedish Research Council and the Swedish National Space Board in Sweden.

Additional support for science analysis during the operations phase is gratefully acknowledged from the Istituto Nazionale di Astrofisica in Italy and the Centre National d’Études Spatiales in France.

4.7 Appendix to Chapter 4

4.7.1 Candidate Ranking

In the final follow-up stages, performed offline, we would like to only search the most significant candidates. To rank candidates by their significance, we need to account for the effective number of trials from which each candidate has resulted. While the number of (semicoherent) independent trials is approximately the same in each work unit within each source, there are more work units in higher frequency bands due to the density of sky locations increasing with frequency. Additionally, since the grid of sky locations searched in the first stage is constructed at the highest frequency in the band, whereas the “zoomed in” grids are defined by the candidates’ spin frequencies, the number of trials in the refinement step varies from candidate to candidate.

The overall result of these effects is that candidates with high detection statistic values are more likely to occur by chance in higher frequency bands than in lower frequency bands, and at the higher end of the frequency band. We construct a consistent ranking statistic by comparing candidates’ false alarm probabilities whilst taking the differing number of trials into account.

We start from the result [225] that the cumulative distribution function (cdf), $G(X_{\max})$, of the

maximum value of N samples of the random variable X , is related to $F(X)$, the cdf of X , by

$$G(X_{\max}) = [F(X_{\max})]^N . \quad (4.24)$$

The false alarm probability of X_{\max} after N samples is therefore

$$\begin{aligned} P_{\text{FA},N}(X_{\max}) &= 1 - [1 - P_{\text{FA},1}(X_{\max})]^N \\ &\approx N P_{\text{FA},1}(X_{\max}) , \end{aligned} \quad (4.25)$$

where we have assumed that $P_{\text{FA},1} \ll 1/N$. In our case, the single-trial false alarm for a candidate with coherent power \mathcal{P}_1 is

$$P_{\text{FA},1}(\mathcal{P}_1) = \int_{\mathcal{P}_1}^{\infty} \chi_2^2(\mathcal{P}'_1) d\mathcal{P}'_1 \quad (4.26)$$

$$= e^{-\mathcal{P}_1/2} , \quad (4.27)$$

where $\chi_2^2(X)$ is the central chi-squared distribution with two degrees of freedom.

It is considerably more difficult to estimate the effective number of independent trials, since each candidate is the result of a large number of trials in previous search stages using different detection statistics. However, since at this stage we are only interested in ranking candidates within each source, and the number of independent trials in the semicoherent step is approximately the same for each candidate, we only need to consider the varying number of trials in the follow-up stages, and the total number of work units in each frequency band.

The overall false-alarm probability is therefore a function of the frequency of the candidate, f , and the coherent power:

$$P_{\text{FA}}(\mathcal{P}_1, f) = K N_W(f) N_F(f) e^{-\mathcal{P}_1/2} , \quad (4.28)$$

where the constant of proportionality, K , is the (unknown) number of independent trials per work unit, $N_W(f)$ is the number of work units within the appropriate frequency band and $N_F(f)$ is the number of trials in the coherent follow-up stage for a candidate at frequency f .

We define the ranking statistic, \hat{R} for follow-up analyses according to the logarithm of the inverse of the false alarm probability,

$$\begin{aligned} \hat{R}(\mathcal{P}_1, f) &\equiv -\log \left[\frac{P_{\text{FA}}(\mathcal{P}_1, f)}{K} \right] \\ &= \frac{\mathcal{P}_1}{2} - \log [N_W(f)] - \log [N_F(f)] , \end{aligned} \quad (4.29)$$

where we have removed the constant term corresponding to K . We note that the above formulation of \hat{R} can *only* be used to rank pulsar candidates from the same source, as the effective number of independent trials per work unit (K) varies from source to source.

4.7.2 Distribution of S_1 with a Signal

To estimate the sensitivity of the search, it is necessary to know the expected distribution of \hat{S}_1 in the presence of a signal for each source. To derive this, we first expand the double sum of Equation (4.2) and separate it into terms in which photon indices are never equal. Identifying the photon weights as the probability that each photon originated from the source in question, and therefore assuming that each photon has a probability $w_j p_s$ of having been pulsed, we find the following expressions,

$$E_p \left[w_j^m e^{-in\Phi(t_j)} \right] = w_j^{m+1} p_s \gamma_n, \quad (4.30a)$$

$$E_p \left[w_j^m e^{in\Phi(t_j)} \right] = w_j^{m+1} p_s \gamma_n^*. \quad (4.30b)$$

The expectation value and variance of S_1 in the presence of a signal are therefore found, after some relabeling and rearranging, to be

$$E_p(p_s, \{\gamma_n\})[S_1] = \frac{p_s^2 |\gamma_1|^2}{\kappa_{S_1}} \sum_{j=1}^N \sum_{k \neq j}^N w_j^2 w_k^2 \hat{W}_T^{\text{rect}}(\tau_{jk}), \quad (4.31)$$

$$\begin{aligned} \sigma_p^2(p_s, \{\gamma_n\})[S_1] &= \frac{1}{\kappa_{S_1}^2} \sum_{j=1}^N \sum_{k \neq j}^N \left[w_j^2 w_k^2 \hat{W}_T^{\text{rect}}(\tau_{jk}) \left(1 + p_s^2 |\gamma_2|^2 w_j w_k - 2p_s^4 |\gamma_1|^4 w_j^2 w_k^2 \right. \right. \\ &\quad \left. \left. + \sum_{l \neq j \neq k}^N w_l^2 \hat{W}_T^{\text{rect}}(\tau_{jl}) \left[2p_s^2 |\gamma_1|^2 + 2p_s^3 \Re \left(\gamma_2 (\gamma_1^*)^2 \right) w_j - 4p_s^4 |\gamma_1|^4 w_j^2 \right] \right) \right]. \quad (4.32) \end{aligned}$$

The expected semicoherent S/N for a signal with pulsed fraction p_s and a pulse profile with Fourier coefficients $\{\gamma_n\}$ is therefore

$$\begin{aligned} \theta_{S_1}^2(p_s, \{\gamma_n\}) &= \frac{E_p[S_1] - E_0[S_1]}{\sqrt{\sigma_0^2[S_1]}} \\ &= p_s^2 |\gamma_1|^2 \sqrt{\sum_{j=1}^N \sum_{k \neq j}^N w_j^2 w_k^2 \hat{W}_T^{\text{rect}}(\tau_{jk})}. \quad (4.33) \end{aligned}$$

In addition to the statistical variance of S_1 , a signal at a random location in the parameter space will be detected at the nearest grid point, and some signal power will be lost as a result of this offset. Denoting this mismatch by m , the pdf of \hat{S}_1 is therefore the pdf of the product of S_1 and $(1 - m)$, which we approximate as a Gaussian with the same mean and variance,

$$E_p \left[\hat{S}_1 \right] = E_p \left[S_1 (1 - m) \right] = E_p \left[S_1 \right] (1 - E[m]), \quad (4.34a)$$

$$\sigma_p^2 \left[\hat{S}_1 \right] = \sigma_p^2 \left[S_1 \right] \sigma^2[m] + \sigma_p^2 \left[S_1 \right] (1 - E[m])^2 + \sigma^2[m] E_p \left[S_1 \right]^2. \quad (4.34b)$$




Figure 4.11: Glitch epochs and phase increments from Monte Carlo sampling during timing of PSR J1844–0346. In the left panel, the phase increment was completely free to vary and exhibits high correlation with the glitch epoch. In the right panel, the glitch epoch-dependent part of the phase increment was accounted for, and a small random increment allowed to vary on top of this, removing the correlation.

The values of $E[m]$ and $\sigma^2[m]$ depend on the geometry of the search grid. For the grid used in this survey, which had a maximum mismatch per parameter of 0.15 and an interbinned frequency spectrum, $E[m] \approx 0.22$ and $\sigma^2[m] \approx 8 \times 10^{-3}$. This pdf is illustrated in Figure 4.3.

4.7.3 Efficient Sampling of Glitch Parameters

As mentioned in Section 4.4.3, when timing large glitches, a phase increment occurring at the time of the glitch can be included in the phase model to ensure that the likelihood surface in the glitch epoch is continuous and easy to sample.

However, this increment is strongly correlated with the glitch epoch, as can be seen in the first panel of Figure 4.11, which can lead to inefficient Monte Carlo sampling. Once an initial combination of valid glitch parameters has been found (including a suitable combination of glitch epoch and phase increment that removes any phase discontinuity at the glitch) we can remove this correlation by separating the total phase increment in our phase model into separate terms: a glitch epoch-dependent term accounting for the known difference in phase between the glitch model parameters being sampled and the initial “reference” glitch model, plus a random offset that is allowed to vary as part of the Monte Carlo sampling. Denoting the reference glitch model with the subscript g_0 , and the sampled glitch model parameters by the subscript g_1 , the total phase increment is

$$\Phi_{g_1} = \delta\Phi_{g_1} + \Phi_{g_0} + \begin{cases} -\Delta\Phi_{g_1}(t_{g_0}), & t_{g_1} < t_{g_0} \\ \Delta\Phi_{g_0}(t_{g_1}), & t_{g_1} > t_{g_0} \end{cases}, \quad (4.35)$$

where, for $i = 0, 1$,

$$\Delta\Phi_{g_i}(t) = 2\pi \left[\Delta f_{g_i}(t - t_{g_i}) + \frac{\Delta \dot{f}_{g_i}}{2}(t - t_{g_i})^2 + \frac{\Delta \ddot{f}_{g_i}}{6}(t - t_{g_i})^3 + \Delta f_{D,g_i} \tau_{D,g_i} \left(1 - e^{-(t-t_{g_i})/\tau_{D,g}} \right) \right]. \quad (4.36)$$

The first term of Equation (4.35) ensures that the phase increment is free to vary over a small range to find the glitch parameters with the highest likelihood. The other terms ensure that the reference glitch model's desirable property of causing no large phase discontinuity between the pulse before and after the glitch also applies to the sampled glitch parameters. This ensures that the sampling rate (and hence the efficiency of the timing procedure) is not unnecessarily burdened by having to phase-fold the data with glitch models resulting in large phase discontinuities, which will obviously have a low likelihood and be rejected.

Removing the correlation between the glitch epoch and the part of the phase increment that is being sampled ensures that the Monte Carlo chains explore the parameter space efficiently. This is especially helpful at the beginning of the Monte Carlo run, as the starting locations of the chain are spread uniformly throughout the parameter space, and could otherwise easily get stuck in low-likelihood regions, as they struggle to jump to the very narrow, highly correlated region of high likelihood.

Chapter 5

The Braking Index of a Radio-quiet Gamma-ray Pulsar

Following the initial submission of this thesis, this chapter was published as Clark, C. J., Pletsch, H. J., Wu, J., et al., 2016, The Astrophysical Journal Letters, 832, L15. © 2016. The American Astronomical Society. All rights reserved.

Abstract

We report the discovery and timing measurements of PSR J1208–6238, a young and highly magnetized gamma-ray pulsar, with a spin period of 440 ms. The pulsar was discovered in gamma-ray photon data from the *Fermi* Large Area Telescope (LAT) during a blind-search survey of unidentified LAT sources, running on the distributed volunteer computing system *Einstein@Home*. No radio pulsations were detected in dedicated follow-up searches with the Parkes radio telescope, with a flux density upper limit at 1369 MHz of $30 \mu\text{Jy}$. By timing this pulsar, we measure its braking index during five years of LAT observations to be $n = 2.598 \pm 0.001 \pm 0.1$, where the first uncertainty is statistical and the second estimates the bias due to timing noise. Assuming its braking index has been similar since birth, the pulsar has an estimated age of around 2,700 yr, making it the youngest pulsar to be found in a blind search of gamma-ray data and the youngest known radio-quiet gamma-ray pulsar. Despite its young age the pulsar is not associated with any known supernova remnant or pulsar wind nebula. The pulsar’s inferred surface magnetic field strength is 3.8×10^{13} G, almost 90% of the quantum-critical level. We investigate some potential physical causes of the braking index deviating from the simple dipole model but find that LAT data covering a longer time interval will be necessary to distinguish between these.

5.1 Introduction

The physical mechanisms by which pulsars radiate rotational energy are as yet unclear. The dominant process can be inferred by measuring a pulsar’s *braking index*, n , the index of a power law relating the pulsar’s spin frequency, f , to its spin-down rate, \dot{f} , via

$$\dot{f} \propto -f^n. \quad (5.1)$$

For example, the simple model of a pulsar as a spinning magnetic dipole predicts $n = 3$ [226].

A pulsar’s braking index can be calculated from measurements of the first two time-derivatives of its spin frequency,

$$n = \frac{f\ddot{f}}{\dot{f}^2}. \quad (5.2)$$

However, most young pulsars exhibit unpredictable fluctuations in their spin frequency on top of their long-term spin-down behavior [16] known as “timing noise”. Observations spanning a long time interval are required to discriminate the overall braking behavior in these fluctuations.

Braking indices can therefore only be measured for the youngest or most highly magnetized pulsars, whose long-term braking is still large enough to dominate their spin-down variation. Reliable measurements of braking indices have been possible for just nine pulsars [3, 227–233]. All of these braking indices deviate significantly from $n = 3$, with only PSR J1640–4631 having $n > 3$ [233].

One pulsar with a measurable braking index, PSR J1119–6127 [234], is particularly unusual. Its emission properties (e.g. radio pulsations, exponentially cut-off power-law gamma-ray spectrum) were, until recently, typical for a “normal” rotationally powered pulsar [235], despite its almost magnetar-level magnetic field (4.1×10^{13} G). The recent magnetar-like outburst from PSR J1119–6127 [236, 237], and similar events from PSR J1846–0258 [238], therefore offer insights into the connection between magnetars and rotationally-powered pulsars.

The Large Area Telescope [LAT; 74] on board the *Fermi Gamma-ray Space Telescope* has proven to be a valuable instrument for the study of young pulsars [87]. The LAT’s 8 years of almost continuous coverage of the entire gamma-ray sky has led to the detection of more than 200 gamma-ray pulsars¹. This long observation span enables precise timing analyses of gamma-ray pulsars [e.g. 165–167] immediately after their detection.

Blind searches in LAT data have led to the discovery of a sizeable population of young, radio-quiet gamma-ray pulsars [132–134, 136, 151, and Chapters 3 and 4]. However, until now, only pulsars discovered through their radio or X-ray pulsations have had measurable braking indices.

In this chapter, we describe the discovery and investigation of PSR J1208–6238, a very young, highly magnetized gamma-ray pulsar with very similar properties to those of PSR J1119–6127, including a measurable braking index.

¹<http://tinyurl.com/fermipulsars>

5.2 Observations

5.2.1 LAT Data

In the *Einstein@Home* survey, we searched for pulsations in gamma-ray photons from unidentified sources in the *Fermi*-LAT Third Source Catalog [3FGL; 138]. One such source, lying close to the Galactic plane, was 3FGL J1208.4–6239.

Our data consist of the LAT arrival times of SOURCE-class² photons above 100 MeV, and weights representing the probability of each photon having come from the target source [152]. The initial data in which the pulsar was originally discovered spanned 2008 August 4 to 2014 October 1 and were produced using internal preliminary versions of the Pass 8 [77] instrument response functions (IRFs) and background models. We included photons from within a 8° region of interest (ROI) around the 3FGL position, with a maximum zenith angle cut-off of 100°, and a maximum cut-off on the LAT’s rocking angle of 52°.

After discovering the pulsar, we produced an extended dataset covering observations until 2016 February 16 for use in follow-up spectral (Section 5.2.3) and timing analyses (Section 5.3). This used the P8R2_SOURCE_V6 IRFs, a lower zenith angle cut-off of 90°, and a larger 15° ROI.

We calculated probability weights for photons within 5° of the target source with `gtsrcprob` using the results of a binned likelihood spectral analysis performed with `pointlike`. Our source model included all 3FGL sources within a radius 5° larger than the ROI. The target source’s spectrum was modeled with an exponentially cut-off power law. Its sky position and spectral parameters were free to vary, as were the spectra of 3FGL sources within 5°, and the normalizations of the background models.

When producing the extended dataset the spectral analysis was performed with `gtlike`, the Galactic diffuse emission and isotropic diffuse background were modeled with the `glliem_v06.fits` [205] and `iso_P8R2_SOURCE_V6_v06.txt` templates³ respectively, and the pulsar’s position was fixed at its preliminary timing position.

5.2.2 Discovery in a Blind Search

The first stage of this survey employs a “semicoherent” search, in which only photons arriving within 2²¹ s (≈ 24 days) of one another are combined coherently (Section 4.2). Candidates from this stage are then “followed up” in more sensitive stages (including a fully coherent stage) to increase their significance. While our search model assumes a constant spin-down rate, the semicoherent stage is sensitive to signals with varying spin-down⁴, provided $\left| \ddot{f} \right| \lesssim 3 \times 10^{-22} \text{ Hz s}^{-2}$. Nevertheless,

²The photon class refers to the selection of cuts that each potential photon event must pass in order to be selected for analysis. SOURCE-class photons are suitable for long-duration analyses of point sources. In comparison, when studying short transient events one can accept higher background rates; or conversely if studying the diffuse gamma-ray background, one requires more stringent cuts to further reduce cosmic ray contamination. See Ackermann et al. [76] for details.

³<http://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>

⁴This can be seen by calculating the $\dot{G}_{\ddot{f}\ddot{f}}$ component of the semicoherent metric of Section 2.8.7

the overall search is less sensitive to pulsars with $|\ddot{f}| \gtrsim 5 \times 10^{-24} \text{ Hz s}^{-2}$, since these will not be detected in the coherent follow-up stage⁵, and must have strong enough signals to appear as highly significant semicoherent candidates.

Visual inspection of a candidate signal from 3FGL J1208.4–6239 revealed clear pulsations, but with large variations in the pulse phase. Including a second frequency derivative of $3.3 \times 10^{-22} \text{ Hz s}^{-2}$ in the phase model removed the majority of these phase variations.

5.2.3 Off-pulse Analysis

After refining the pulsar’s spin parameters, the integrated pulse profile showed significant emission above the background at all phases, indicative either of flux from the pulsar or from a nearby unmodeled source.

To investigate this unpulsed flux, we assigned rotational phases to all photons using our initial timing solution, and performed a spectral analysis of the off-pulse emission. This revealed additional nearby gamma-ray sources, the closest and most significant of which lies $20'$ from the pulsar’s timing position, with a spectral index of $\Gamma = 2.56 \pm 0.09$ and $(40 \pm 13)\%$ of the energy flux of PSR J1208–6238.

The nature of these sources is unclear; without likely counterparts at other wavelengths it is difficult to tell whether these are point sources or residuals from the Galactic diffuse emission template. Their steep spectra are consistent with LAT-detected supernova remnants [SNRs; 239], but kick velocity requirements make an association with the pulsar unlikely (see Section 3.5).

After including these additional sources in our model, re-fitting the ROI and re-calculating the photon weights for PSR J1208–6238, the pulsation significance increased. These weights were subsequently used in the timing analyses described in Section 5.3, resulting in the pulse profile shown in Figure 5.1, with a final H -test value 1454.2. The spectral properties of the pulsar are given in Table 5.1.

5.2.4 Radio Observations

On 2016 March 28 and 2016 April 14 we observed the pulsar timing position with the 64-m Parkes radio telescope for 2.5 hr and 4.3 hr respectively using the H-OH receiver at a center frequency of 1369 MHz. With PDFB4 we recorded 256 MHz of bandwidth filtered into 512-channel spectra with $256 \mu\text{s}$ sampling. Using the gamma-ray ephemeris, we folded the data over 690 dispersion measure (DM) trials from $0\text{--}1500 \text{ pc cm}^{-3}$ and visually inspected the summed profiles over a range of phase and time binnings. We found no plausible pulsar candidates. Recent flux-calibrated observations using the same configuration show a root-mean-square noise level of $130 \mu\text{Jy/hr}$ for a 32-bin pulse profile. With a detection threshold of one or more phase bins with $S/N > 8$, then optimally our longer observation would have detected a source with a mean flux density $> 17 \mu\text{Jy}$. Accounting for scalloping losses due to binning in time and DM, and for the unknown pulse duty cycle, we estimate an upper limit of $30 \mu\text{Jy}$, equal to the radio-quiet threshold defined by Abdo et al. [80].

⁵Similarly to footnote 4, this can be seen from the relevant component of the coherent metric

Table 5.1. Properties of PSR J1208–6238

Parameter	Value
Range of photon data (MJD)	54682–57434
Reference epoch, t_{ref} (MJD)	56040
Timing parameters ^a	
R.A. (J2000.0)	12 ^h 08 ^m 13 ^s .96(6)
Decl. (J2000.0)	–62° 38′ 02″.3(4)
Spin frequency, f (Hz)	2.26968010518(7)
Spin-down rate, \dot{f} (10^{-12} Hz s ⁻¹)	–16.842733(5)
Braking index, n	2.598(1)
\dot{f} -increment epoch (MJD)	55548(23)
\dot{f} -increment, $\Delta\dot{f}$ (10^{-15} Hz s ⁻¹)	0.59(9)
n -increment, Δn	–0.10(2)
Derived properties ^b	
Galactic longitude, l (°)	297.99
Galactic latitude, b (°)	–0.18
Spin period, P (ms)	440.59072365(1)
Period derivative, \dot{P} (10^{-12} s s ⁻¹)	3.2695145(9)
Surface B-field strength, B_{S} (10^{12} G)	38.4
Estimated age ^c , τ (yr)	2672
Spin-down power, \dot{E} (10^{36} erg s ⁻¹)	1.5
Maximum distance, $d_{100\%}$ (kpc)	18.9
Heuristic distance, d_{h} (kpc)	3.0
Spectral parameters above 100 MeV ^d	
Spectral index, Γ	$1.73 \pm 0.08 \pm 0.04$
Cutoff energy, E_{c} (GeV)	$4.86 \pm 0.59 \pm 0.70$
Photon flux, F_{100} (cm ⁻² s ⁻¹)	$(4.41 \pm 0.86 \pm 0.37) \times 10^{-8}$
Energy flux, G_{100} (erg cm ⁻² s ⁻¹)	$(3.49 \pm 0.44 \pm 0.29) \times 10^{-11}$

Note. — The reported values for f and \dot{f} at the reference time include the effect of the earlier \dot{f} increment.

^aFor timing parameters, we report mean values and 1σ uncertainties on the final digits in brackets from the results of the timing analysis described in Section 5.3 using the phase model in Equation 5.4.

^bDerived properties are calculated as described in Abdo et al. [80]. Maximum and heuristic distances are calculated assuming isotropic emission and gamma-ray luminosities of \dot{E} and $\sqrt{10^{33}\dot{E}}$ respectively.

^cThe estimated age was calculated using the measured braking index.

^dThe first uncertainty is statistical, the second estimates systematic uncertainties in the LAT’s effective area, estimated by performing the same spectral analysis with rescaled effective areas (see http://fermi.gsfc.nasa.gov/ssc/data/analysis/scitools/Aeff_Systematics.html for details), and in the Galactic diffuse emission model, estimated by performing the spectral analysis with the normalization of the Galactic diffuse emission rescaled by $\pm 6\%$.




Figure 5.1: Gamma-ray pulse profile of PSR J1208–6238, using the photon weights described in Section 5.2.3 and the Taylor series timing solution from Section 5.3. The dashed blue line shows the background level, estimated from the photon probability weights as described in Abdo et al. [80]. The solid orange curve shows the template pulse profile used in the timing analysis, described in Section 5.3.

5.3 Timing Analysis

To investigate the braking properties of PSR J1208–6238, we performed a dedicated timing analysis, following the procedure described in Pletsch & Clark [166] and Chapters 3 & 4, a modification of the methods developed by Ray et al. [167]. For efficiency, we removed around 750,000 photons with a weight below 0.7%, which contribute negligibly to the signal.

Our phase model consisted of a Taylor series,

$$\Phi(t) = \Phi_0 + \sum_{m=0} \frac{f^{(m)}}{(m+1)!} (t - t_{\text{ref}})^{m+1}, \quad (5.3)$$

where $f^{(m)}$ denotes the m -th time-derivative of the pulsar’s rotational frequency at the reference epoch, t_{ref} .

Starting with our initial solution, consisting of the pulsar’s sky position, frequency and first two frequency derivatives ($\dot{f} \equiv f^{(1)}$ and $\ddot{f} \equiv f^{(2)}$), we phase-folded the photon data and fit a template pulse profile (shown in Figure 3.3) by maximizing the likelihood. We used the Bayesian Information Criterion [BIC; 219] to estimate the appropriate number of components to include in the template, finding that two wrapped Gaussian peaks were sufficient.

We then sampled the pulsar’s spin and positional parameters with an Affine Invariant Monte Carlo algorithm [220, 221], using the template pulse profile to evaluate the likelihood at each point. The process continued iteratively; after each sampling stage the photons were re-folded with the most likely parameters and the template pulse profile was updated. Higher frequency derivatives were added to the phase model after each stage, again using the BIC to find the appropriate number of terms.

The solution that optimized the BIC contained nine frequency derivatives. However, the phase residuals from this model contain some remaining “red” noise, indicating that the timing noise is not entirely captured in this solution. We were unable to remove this effect by including higher frequency derivatives, up to the 12th derivative.

We can calculate n with Equation (5.2) using the measured values of \dot{f} and \ddot{f} , finding $n = 2.578 \pm 0.007$. However, as explained by Antonopoulou et al. [240], truncating the Taylor series at the second derivative means that the measured value of n depends on the arbitrarily chosen reference epoch. Making the simple assumption of a constant braking index, a self-consistent physically motivated phase model can be derived, which avoids this problem [240],

$$\Phi(t) = \Phi_0 + \frac{f^2}{\dot{f}(2-n)} \left(\left[1 + \frac{\dot{f}}{f}(1-n)(t - t_{\text{ref}}) \right]^{\frac{2-n}{1-n}} - 1 \right). \quad (5.4)$$

In this model, one only fits for f , \dot{f} and n . Faster timing noise variations on top of this long term spin-down are left unmodelled.

Carrying out this fit for PSR J1208–6238, we find that a constant braking index fits well in the last five years of LAT data, but cannot account for the pulsar’s spin-down over the entire span of the LAT data, as is clear from Figure 5.2. Instead, we find that increments in the pulsar’s braking index and spin-down rate are required. The results of this fit, including these increments, are given in Table 5.1. We cannot tell whether these were sudden or gradual changes as it takes many months for an offset in either of these parameters to cause a detectable phase offset. We note that Archibald et al. [241] required a similar change in \dot{f} and n in their timing model for PSR J1846–0258 covering six years of observations.

5.4 Discussion

PSR J1208–6238 is the first radio-quiet gamma-ray pulsar with a measured braking index⁶, and only the tenth pulsar of any kind with such a measurement. Its estimated age, assuming its braking index has been constant since birth, is around 2,700 yr, making it the youngest known radio-quiet gamma-ray pulsar.

PSR J1208–6238 shares many similarities (e.g. age, spin-down power) with PSR J1119–6127, an otherwise unique pulsar. They are both highly magnetized, with inferred surface magnetic field

⁶Although Kuiper & Dekker [242] recently detected faint pulsations from PSR J1846–0258 in LAT data below 100 MeV.




Figure 5.2: Measured phase residuals from our timing models. The blue lines and gray shaded regions represent the best-fitting Taylor series phase model and 1σ uncertainties. The “TOAs” shown here were not used to perform the fit, which was an entirely unbinned likelihood maximization, but are included here to illustrate the validity of the phase models. Upper panel: residuals between the Taylor series and a pure dipole-braking model with $n = 3$. The decreasing cubic curve is characteristic of a significantly over-estimated braking index. Middle panel: residuals between the Taylor series model and a best-fitting constant braking index model. A significant deviation is evident in the early mission. Lower panel: as above, but with increments in the braking index and spin-down rate occurring at the time marked by the dashed vertical line included in the phase model.

strengths, $B_S = 3.2 \times 10^{19} (-\dot{f} f^{-3})^{1/2}$, of 3.8×10^{13} G (PSR J1208–6238) and 4.1×10^{13} G (PSR J1119–6127), close to the quantum-critical limit $B_{\text{cr}} = 4.4 \times 10^{13}$ G [243]. Such a high magnetic field is expected to affect a pulsar’s high-energy emissions, however both PSR J1208–6238 and PSR J1119–6127 have gamma-ray spectra that are typical of young pulsars [80]. Given their similar magnetic fields, one may speculate that PSR J1208–6238 could in the future exhibit similar magnetar-like activity as was recently detected from PSR J1119–6127 [236]. However, no X-ray emission has yet been detected from PSR J1208–6238 [198].

Radio emission is thought to be suppressed by high magnetic fields, as pair production in the magnetosphere is inhibited by photon splitting [234, 243]. To predict whether to otherwise expect radio emission from PSR J1208–6238, we modeled the gamma-ray emission geometry using the methods of Johnson et al. [85], with outer gap (OG) and two-pole caustic (TPC) models described therein. Our simulations used $P = 100$ ms and $\dot{P} = 10^{-15}$ s s⁻¹. While these are different from those in Table 5.1 these parameters only determine the size of the magnetic polar cap and radio emission altitude and the gamma-ray profile dependence on these is weak at these values. We used steps in gap widths of 1% the polar cap opening angle. Additionally, while Johnson et al. [85] fit the unweighted counts using Poisson likelihood, we used weighted counts and the χ^2 statistic. The TPC model is slightly, although not significantly, favored over the OG model by the likelihood, with the angle between the line-of-sight and magnetic angle, $\beta = 2_{-1}^{+1}$ °, predicting visible radio pulsations. The OG fit gives $\beta = 28_{-6}^{+8}$ °, with no visible radio emission. We caution, however, that performing these fits without constraints on the radio emission can lead to large systematic errors [86], and note that neither model perfectly reproduces the observed peak separation.

Despite the pulsar’s very young inferred age, no positionally associated SNR or pulsar wind nebula has been detected by radio imaging [244], X-ray observations [198, 245], or in TeV emission [246]. There is, however, a luminous H II region [IRAS 12073–6233; 247] less than 0.5° away from PSR J1208–6238, whose radio emission could mask a faint SNR near the pulsar. A dedicated search for an associated SNR may be required to clarify this. The two nearest known SNRs (G298.6–0.0 and G298.5–0.3) are more than half a degree away from the pulsar [200]. Due to the pulsar’s young age, an association with either of these SNRs implies an unrealistically high kick velocity, approximately $3000 (d/1 \text{ kpc}) \text{ km s}^{-1}$ for pulsar distance $3 \text{ kpc} \lesssim d < 19 \text{ kpc}$ (see Table 5.1). A bias on the pulsar’s timing position due to timing noise [165] is insufficient to account for the positional offsets between the pulsar and the nearest SNRs.

The braking index measurement is also sensitive to the pulsar’s intrinsic timing noise, and is likely to contain some bias due to this effect. The uncertainty on n of $\pm 10^{-3}$ quoted in Table 5.1 is statistical only; low frequency timing noise will lead to an additional unknown bias on this value. Assuming that the observed change in n around MJD 55600 was due to timing noise, we take this increment, $\Delta n = 0.1$, as an estimate of this systematic uncertainty.

The larger braking index in the earlier data could alternatively be caused by the pulsar relaxing from a glitch occurring before the start of the *Fermi* mission, of the kind observed from PSR J1119–6127 [231, 240, 248]. However, a timing model featuring an exponentially decaying frequency term in the early mission is disfavored by the BIC.

Nevertheless, our timing measurements constrain the braking index to be below the $n = 3$

predicted by a simple dipole braking model, as evident from the large cubic residuals shown in Figure 5.2. A reduced braking index can be explained by the pulsar’s physical properties varying over time. Differentiation of the dipole braking model gives [3],

$$n = 3 + 2 \frac{\dot{f}}{f} \left[\frac{\dot{\mu}}{\mu} + \frac{\dot{\alpha}}{\tan \alpha} - \frac{\dot{I}}{2I} \right], \quad (5.5)$$

where μ is the magnetic dipole moment, inclined at an angle α to the spin axis, and I is the pulsar’s moment of inertia. The measured braking index can be explained by fractional changes in these parameters of $\sim 5 \times 10^{-5} \text{ yr}^{-1}$.

A low braking index can be explained by a growing magnetic field [249]. One possible cause of this could be that the pulsar’s initial magnetic field was “buried” by matter accretion shortly after birth, and is gradually growing back to its original strength [250]. This implies a braking index that evolves back to 3 over a timescale of $\sim 10^5 \text{ yr}$. This evolution of n is undetectable with current observation lengths even in the absence of timing noise.

A varying α naturally leads to evolution of the pulsar’s observed pulse profile. This has been observed from the Crab pulsar, where a measured increase in the (angular) peak separation of $\approx 6 \times 10^{-3} \text{ }^\circ \text{ yr}^{-1}$ implies a similar magnitude for $\dot{\alpha}$ [251]. This effect, which could be caused by free precession [252], is sufficient to explain its low braking index.

Fits to TPC and OG models described above estimate $\alpha = 81_{-1}^{+1} \text{ }^\circ$ and $\alpha = 57_{-5}^{+8} \text{ }^\circ$ respectively for PSR J1208–6238. For the variation in α to account for the reduced braking index requires $\dot{\alpha} \approx 2 \times 10^{-2} \text{ }^\circ \text{ yr}^{-1}$ (TPC) or $\dot{\alpha} \approx 4 \times 10^{-3} \text{ }^\circ \text{ yr}^{-1}$ (OG). For either model, the expected evolution of the gamma-ray pulse profile caused by the required $\dot{\alpha}$ is too small to be measured with the current data.

Lyne et al. [3] note that a decreasing moment of inertia could be caused by decoupling between the pulsar’s crust and interior, although they conclude that such a mechanism should not last long enough to explain the even lower braking index of the older Vela pulsar.

Low braking indices can also be explained by a portion of the total spin-down torque being applied by a process with a different braking index, e.g. angular momentum being lost to an outflowing particle wind [$n = 1$; 253, 254], or propeller torque from an in-falling disk [$n = -1$; 255]. The fraction, ϵ , of the total spin-down torque that a process with a braking index of n_2 must account for to explain an observed index n is⁷

$$\epsilon \approx \frac{3 - n}{3 - n_2}. \quad (5.6)$$

Under these models, as the pulsar spins down ϵ will increase (provided $n_2 < 3$), leading to a time-varying braking index, where

$$\dot{n} \approx \frac{\dot{f}}{f} \epsilon (n_2 - 3) (n_2 - n). \quad (5.7)$$

⁷ Lyne et al. [3] define ϵ as the ratio of wind-induced torque to dipole-induced torque, rather than wind-induced torque to total spin-down torque. This definition may have been used in error by Archibald et al. [233, 241], although their conclusions are unchanged.

For the above wind (disk) model, we find $\epsilon \approx 20\%$ (10%) and $\dot{n} \approx -1.5 \times 10^{-4} \text{ yr}^{-1}$ ($-3.4 \times 10^{-4} \text{ yr}^{-1}$). The braking index variations are not currently measurable, but may become so with future LAT data.

Acknowledgments

We gratefully acknowledge the support of all volunteers who have donated spare CPU cycles to *Einstein@Home*. We are especially grateful to James Drews of UW-Madison, WI, USA and University of Houston, IT High Performance Computing, TX, USA, whose computers discovered PSR J1208–6238.

This work was supported by the Max-Planck-Gesellschaft (MPG), by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether research grant PL 710/1-1 (PI: Holger J. Pletsch), and by NSF award 1104902.

The Parkes radio telescope is part of the Australia Telescope, which is funded by the Commonwealth Government for operation as a National Facility managed by CSIRO.

The *Fermi*-LAT Collaboration acknowledges support for LAT development, operation and data analysis from NASA and DOE (United States), CEA/Irfu and IN2P3/CNRS (France), ASI and INFN (Italy), MEXT, KEK, and JAXA (Japan), and the K.A. Wallenberg Foundation, the Swedish Research Council and the National Space Board (Sweden). Science analysis support in the operations phase from INAF (Italy) and CNES (France) is also gratefully acknowledged.

5.5 Additional Material: Alternative Spin-down Models

This chapter was written to a strict word limit for submission to *The Astrophysical Journal Letters*. The following section will not appear in the final published version, but is included here for additional information.

In Section 5.4, it was stated that a pulsar whose spin-down torque is partially provided by a process other than dipole braking will have a lower braking index, and that its braking index will change over time. In this case, the assumption behind the phase model of Equation (5.4), that the braking index is constant, clearly does not hold. We therefore derive self-consistent, exact expressions for a pulsar’s spin frequency as a function of time, assuming the wind ($n_2 = 1$) and disk ($n_2 = -1$) models discussed above.

We start with the simple spin-down models, and defining ϵ as the fraction of the spin-down torque due to an alternative braking process (denoting the wind and disk models by the subscript w and d respectively),

$$\dot{f}_{\text{wind}} = -Af^3 - Bf, \quad A = -(1 - \epsilon_{0,w}) \frac{\dot{f}_0}{f_0^3}, \quad B = -\epsilon_{0,w} \frac{\dot{f}_0}{f_0}, \quad (5.8)$$

$$\dot{f}_{\text{disk}} = -Cf^3 - \frac{D}{f}, \quad C = -(1 - \epsilon_{0,d}) \frac{\dot{f}_0}{f_0^3}, \quad D = -\epsilon_{0,d} f_0 \dot{f}_0, \quad (5.9)$$

where the values of parameters with a subscript 0 are taken at the reference epoch. Integrating these models from $t' = t_{\text{ref}}$ to $t' = t$ gives the resulting expressions for the pulsar’s spin frequency as a function of time.

$$f_{\text{wind}}(t) = f_0 \left[1 - \frac{1}{\epsilon_{0,w}} \left(e^{-2\epsilon_{0,w} \frac{\dot{f}_0}{f_0} (t-t_{\text{ref}})} - 1 \right) \right]^{-\frac{1}{2}} \quad (5.10)$$

$$f_{\text{disk}}(t) = f_0 \left(\frac{\epsilon_{0,d}}{1 - \epsilon_{0,d}} \right)^{\frac{1}{4}} \left(\tan \left[-2\sqrt{\epsilon_{0,d}(1 - \epsilon_{0,d})} \frac{\dot{f}_0}{f_0} (t - t_{\text{ref}}) + \tan^{-1} \left(\sqrt{\frac{1 - \epsilon_{0,d}}{\epsilon_{0,d}}} \right) \right] \right)^{\frac{1}{2}}. \quad (5.11)$$

Equation (5.11) corresponds to Equation (8) of Menou et al. [255]. To the author’s knowledge, Equation (5.10) has not been published previously, although Harding et al. [254] derive the corresponding “age” of such a pulsar via a similar integration.

In Figure 5.3 these models are compared to the constant braking index model underlying Equation (5.4), whose frequency model (also presented in [37]) is,

$$f(t) = f_0 \left[1 - \frac{\dot{f}_0}{f_0} (n - 1)(t - t_{\text{ref}}) \right]^{\frac{1}{1-n}}. \quad (5.12)$$

The lower panel of Figure 5.3 shows the corresponding phase residuals between these models.

The dominant term in the frequency residuals is cubic, hence the value of the third frequency derivative could alternatively be used to discern between these models. This would be equivalent to measuring the “second deceleration parameter”, $p = \ddot{f} f^2 / \dot{f}^3$ [e.g. 3, 249]. However, this parameter suffers from the same problem as the braking index does when calculated via Equation (5.2), namely that if there is timing noise its value changes depending on the arbitrary reference time. Furthermore, to measure this parameter, one needs to fit both n and p , rather than only fitting for ϵ as for either of the above models.

Unfortunately, the differences between these models will be even more biased by timing noise than measurements of braking indices, likely rendering them indistinguishable for most young pulsars. For the youngest pulsars, where these models are of most interest, one would need to jointly fit their parameters with the parameters of a timing noise model [e.g. 165, and references therein]. Also, while these models have pleasingly simple underlying assumptions, dipole braking and wind/disk torques are far from the only processes which can brake pulsars; for example alignment torque [e.g. 256], precession [e.g. 252, 257] or even quadrupole radiation [e.g. 226, 233] can introduce further terms (which are not necessarily frequency dependent) into the initial spin-down equation.




Figure 5.3: Differences between the constant braking index model, and models where the pulsar’s spin-down torque contains contributions from a wind or a disk. The top panel shows the differences ($\Delta f(t)$) in the spin frequency of PSR J1208–6238 between each model and the constant braking index model (dashed line at $\Delta f(t) = 0$) over the observation interval. The lower panel shows the resulting phase residuals between these models obtained by numerically integrating the above frequency offsets. Note that the largest phase differences at the start and the end of the observation time are comparable to the absolute uncertainty on the pulsar’s phase, shown in Figure 5.2. However the presence of the apparent increments in the spin-down rate and braking index of PSR J1208–6238, which we speculate in Section 5.4 could be due to timing noise or potentially a glitch recovery, reduce the time-span over which we could apply one of these models consistently, making these residuals unmeasurable in this case.

Chapter 6

The Discovery of Two Possibly Radio-Quiet Gamma-ray Millisecond Pulsars

Following the initial submission and defense of this thesis we obtained dedicated follow-up radio observations of both MSPs described in this chapter. The results of these observations will be described in an upcoming publication.

Abstract

Millisecond pulsars (MSPs) are old neutron stars, with typical ages of at least a billion years, that spin rapidly at hundreds of times per second. As yet, all spin-powered MSPs have shown radio pulsations, interpreted as co-rotating beams of emission intersecting our view from Earth with each turn. While many radio-detected MSPs are also bright gamma-ray sources, until now the existence of a radio-quiet gamma-ray MSP population has remained hidden. Exploiting the aggregated computing power of the *Einstein@Home* distributed computing project, we conducted a blind search for pulsations in gamma-ray data from the *Fermi* Large Area Telescope. This survey discovered two isolated MSPs, both of which remain undetected in extensive radio follow-up searches. One explanation could be that their radio beams miss Earth, putting existing emission models to test.

6.1 Introduction

Over the last few years, the Large Area Telescope (LAT) on the Fermi Gamma-ray Space Telescope [74] has revealed entirely new populations of gamma-ray emitting objects, radiating in the 0.1–100 GeV energy band. Among these are millisecond pulsars (MSPs): old neutron stars thought to be spun-up by accreting matter from a companion star [19, 258], reaching fast rotation rates of hundreds of revolutions per second [206]. Recent observational evidence underpins this model [259]. The origin of the 20% fraction of Galactic MSPs that are isolated [14] is still less clear. One explanation for their formation could be that the companion was fully evaporated by the high-energy radiation of the MSP [260, 261]. However, in several cases the time-scales inferred from observations for complete evaporation of the companion seem too long [106].

Shortly after launch, by folding gamma-ray photon arrival times using rotational ephemerides from concurrent radio telescope observations [115], the LAT discovered that many radio-detected MSPs also emit gamma-rays [92]. In addition, dedicated radio searches in unidentified LAT sources have been conducted with great success [96–103].

The LAT is also sensitive enough for direct pulsar detection through blind searches for periodicity in gamma-ray photon arrival times. Pulsars' gamma-ray beams are typically wider than their radio beams and therefore are visible from a larger range of inclination angles [121]. They are also not subject to dispersion, scattering or scintillation as they propagate through the interstellar medium. Blind gamma-ray pulsar searches are therefore sensitive to a population of pulsars that is complementary to that which can be found in radio searches.

So far, blind gamma-ray pulsar searches have found several tens of solitary younger, slower spinning pulsars [132–134, 136, 151, Chapters 3, 4 & 5]. However, the blind-search detection of gamma-ray MSPs has proven challenging. In just one case a direct search of LAT data revealed an MSP in a binary system [135] by exploiting partial knowledge of the orbit from prior identification of the heated companion at optical and X-ray wavelengths [191]. Subsequently, this MSP was also detected in the radio band [262], unlike the majority of young blind-search gamma-ray pulsars which appear to be radio-quiet [80].

Thus, until now, no gamma-ray MSP without detectable radio emission has been found, despite population syntheses [163, 164] predicting that several such objects reside among the unidentified LAT sources.

The apparent absence of this source class is most likely due to the inherent difficulty of detecting the signal of an MSP with an unknown sky position in a blind gamma-ray pulsation search. Such signals gain additional Doppler modulations due to *Fermi's* motion through the Solar System, which must be corrected by applying position-dependent corrections to the LAT-recorded photon arrival times to retrieve their arrival times at the Solar System Barycenter. The precision with which these must be applied depends on the pulsation frequency. The localization of unidentified gamma-ray sources, limited by the LAT's angular resolution to a few arcmin [138], is insufficiently precise to detect pulsations from MSPs, where a precision of a few arcseconds is often required. Tens to hundreds of thousands of sky locations covering the source localization region must therefore be searched, incurring a large computational cost. Additionally, without the complex heterodyning

technique (see Section 2.5.3) first used for blind gamma-ray pulsar searching in Pletsch et al. [151], very long FFTs are required to search at high spin frequencies, in which case memory constraints can become limiting [e.g. 127].

6.2 MSP Discoveries by the *Einstein@Home* Search

In Chapter 2 we made improvements to the semicoherent search method which, in addition to increasing their sensitivity to weak pulsations, mean that fewer search locations are required to cover a given solid angle, reducing the cost of searching for MSPs.

To avoid repetition with previous chapters, we here refer the reader to the relevant sections. The LAT data were prepared using the methods described in Section 4.2.1. The blind search methods are described in Section 4.2.

Among the sources searched by the *Einstein@Home* survey described in this thesis were two of the three most significant unassociated 3FGL sources, 3FGL J1035.7–6720 and 3FGL J1744.1–7619 [138] (the remaining source likely being a binary MSP [263]). Both of these have significantly curved spectra and low variability, marking them as promising pulsar search targets. Their locations off the Galactic plane ($b = -7.8^\circ$ and $b = -22.5^\circ$ respectively) also made young pulsar explanations unlikely. While searching in 5.5 years of data from these sources, the *Einstein@Home* survey discovered strong pulsations at 348 Hz and 213 Hz respectively, with the final timing solutions giving weighted H -test [152, 176] values of 459 and 722 corresponding to single-trial false alarm probabilities $< 10^{-79}$ and $< 10^{-125}$ respectively. These gamma-ray sources are therefore identified as isolated MSPs, now known as PSR J1035–6720 and PSR J1744–7619.

6.3 Pulsar Properties

Following their detections we extended our data for each pulsar to cover observations until 2016 February 16¹. We refined the parameters of each signal through an unbinned timing analysis². The resulting estimates and uncertainties on the astrometric and rotational parameters for both pulsars are given in Table 6.1. The photon phases and integrated pulse profiles given by the final timing solution are shown in Figure 6.1.

With the pulsars’ locations fixed at their timing positions, we performed binned likelihood spectral analyses using `gtlike`, following the procedure described in Section 4.4.1. The resulting spectral parameters for each pulsar are given in Table 6.1. The spectral properties (Γ & E_{cut}) of both pulsars are similar to those of the MSP population seen in the Second *Fermi*-LAT Catalog of Gamma-ray Pulsars [2PC; 80].

The inclusion of proper motion in our timing models led to a slight decrease in the BIC (see Equation (4.17)), $\Delta\text{BIC} = -5.2$, for PSR J1035–6720 and a significant decrease, $\Delta\text{BIC} = -29.0$, for PSR J1744–7619. Assuming this is due to each pulsars’ motion, and not an effect of timing

¹The extended datasets were produced using the procedure described in Section 4.4.

²The unbinned likelihood timing procedure is described in Section 4.4.

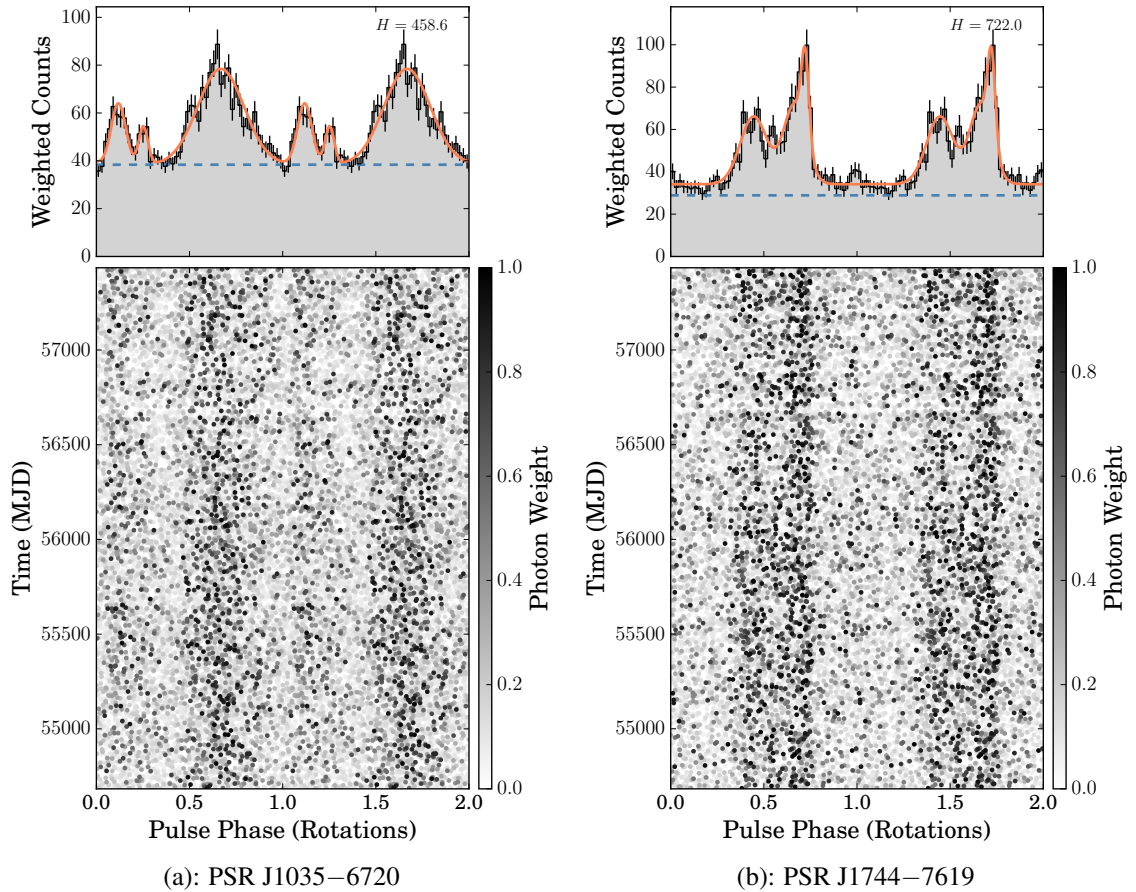


Figure 6.1: Pulsed signals from the newly discovered MSPs. Lower panels: Phase-time diagram showing the arrival times and phases of gamma-ray photons from the pulsar. The grayscale intensity represents the weight of each photon, or the probability that each photon is from the pulsar. Upper panels: The integrated gamma-ray pulse profile. The dashed horizontal blue lines show the estimated backgrounds, based on the distributions of photon weights [via Equation (4.12)]. The orange curves show the template pulse profile used to time each pulsar.

Table 6.1. Measured and derived MSP parameters

Parameter	PSR J1035–6720	PSR J1744–7619
Range of Photon Data (MJD)	54682–57434	
Reference epoch (MJD)	55716	
Timing parameters ^a		
R. A. (J2000.0), α	10 ^h 35 ^m 27 ^s .477(1)	17 ^h 44 ^m 00 ^s .491(2)
Decl. (J2000.0), δ	–67° 20′ 12″.693(6)	–76° 19′ 14″.709(1)
Proper motion in R. A., $\mu_\alpha \cos \delta$ (mas yr ^{–1})	–12(3)	–22(3)
Proper motion in Decl., μ_δ (mas yr ^{–1})	–1(3)	–8(4)
Spin frequency, f (Hz)	348.18864014054(8)	213.33223675349(5)
1st frequency derivative, \dot{f} (10 ^{–15} Hz s ^{–1})	–5.634(2)	–0.439(1)
2nd frequency derivative, $ \ddot{f} $ (10 ^{–25} Hz s ^{–2})	< 1.6	< 1.2
Derived quantities ^b		
Galactic longitude, l (deg)	290.37	317.11
Galactic latitude, b (deg)	–7.84	–22.46
Spin period, P (ms)	2.8720063916972(7)	4.687524094896(1)
Period-derivative, \dot{P} (10 ^{–20} s s ^{–1})	4.647(1)	0.967(3)
Spin-down luminosity, \dot{E} (10 ³⁴ erg s ^{–1})	7.7	0.4
Characteristic age, τ_c (10 ⁹ yr)	1.0	7.6
Surface magnetic field, B_S (10 ⁸ G)	3.7	2.2
Light-cylinder magnetic field, B_{LC} (10 ⁵ G)	1.4	0.2
Phase-averaged gamma-ray spectral parameters ^c		
Photon index, Γ	1.37 ± 0.07 ± 0.23	1.06 ± 0.11 ± 0.01
Cutoff energy, E_c (GeV)	2.33 ± 0.18 ± 0.51	1.87 ± 0.21 ± 0.04
Photon flux above 0.1 GeV, F (10 ^{–9} photons cm ^{–2} s ^{–1})	22.8 ± 3.5 ± 2.9	19.2 ± 4.6 ± 3.9
Energy flux above 0.1 GeV, G (10 ^{–12} erg cm ^{–2} s ^{–1})	20.6 ± 2.3 ± 1.1	21.2 ± 3.8 ± 3.7
Off-pulse spectral parameters ^c		
Off-pulse phase range (rotations)	0.297–0.437	0.812–1.272
Test statistic (TS)	4.2	127.0
TS of exponential cutoff, TS _{cut}	...	42.3
Photon index, Γ	...	0.01 ± 0.03 ± 0.01
Cutoff energy, E_c (GeV)	...	0.53 ± 0.05 ± 0.01
Photon flux above 0.1 GeV, F (10 ^{–9} photons cm ^{–2} s ^{–1})	...	2.6 ± 1.2 ± 0.3
Energy flux above 0.1 GeV, G (10 ^{–12} erg cm ^{–2} s ^{–1})	...	2.6 ± 1.4 ± 0.3

Note. — The reference time scale is Barycentric Dynamical Time (TDB).

^aQuoted values for the timing parameters are the values giving the highest likelihood in the unbinned timing analysis, with 1 σ uncertainties on the final digits quoted in brackets.

^bDerived quantities are calculated as in [e.g. 80]. The reported values for \dot{E} have not been corrected for the pulsars’ proper motions (see Section 6.3 and Figure 6.2).

^cThe first uncertainty is statistical, the second estimates systematic uncertainties in the LAT’s effective area, estimated by performing the same spectral analysis with rescaled effective areas (see http://fermi.gsfc.nasa.gov/ssc/data/analysis/scitools/Aeff_Systematics.html for details), and in the Galactic diffuse emission model, estimated by performing the spectral analysis with the normalization of the Galactic diffuse emission rescaled by $\pm 6\%$.

noise (which can bias MSPs’ timing positions and proper motions [165]), a portion of the observed spin-down rate (\dot{f}_{OBS}) for each must be caused by the Doppler shift introduced by their velocity gaining an increasing radial component, also known as the Shklovskii effect [80, 264]. Correcting the pulsar’s intrinsic spin-down rate (\dot{f}_{INT}) for the Doppler-induced apparent spin-down (\dot{f}_{SHK}) requires knowledge of the pulsars’ distances, which are uncertain without dispersion or parallax measurements. However, we can use the pulsars’ observed gamma-ray fluxes (G_{100}) and proper motions (μ), and assume certain gamma-ray efficiencies (η), to retrieve constraints on the pulsars’ distances (d) by solving,

$$\dot{f}_{\text{OBS}} = \dot{f}_{\text{INT}} + \dot{f}_{\text{SHK}} \quad (6.1)$$

$$= -\frac{G_{100}f_{\Omega}d^2}{\eta\pi If} - \frac{\mu^2 df}{c}. \quad (6.2)$$

Figure 6.2 shows the results of this, assuming a canonical moment of inertia $I = 10^{45} \text{g cm}^2$, a gamma-ray beaming factor of $f_{\Omega} = 1$ and a realistic range of gamma-ray efficiencies $0.01 < \eta < 1$ [80]. From the inferred distances we can retrieve Shklovskii-corrected spin-down powers, (\dot{E}), via $\dot{E} = G_{100}4\pi f_{\Omega}d^2/\eta$. These ranges are shown in the right-hand panels of Figure 6.2.

We find that this Doppler correction can account for up to 20% and 65% of the spin-down rates of PSR J1035–6720 and PSR J1744–7619 respectively at an efficiency of $\eta = 1$. If the efficiency of PSR J1744–7619 is high its spin-down power could be as low as $1.3 \times 10^{33} \text{ erg s}^{-1}$, which would make this one of the least energetic gamma-ray MSPs and close to the empirical gamma-ray “death line”, i.e. the minimum spin-down power an MSP must have to emit in gamma-rays [89].

6.4 Radio Observations

In a search with the Parkes radio telescope of 56 unidentified *Fermi*-LAT sources, Camilo et al. [102], detected 11 MSPs, 10 of them discoveries. The LAT sources associated with PSRs J1035–6720 and J1744–7619 were among the targets of this survey. Here we summarize these searches, which yielded no radio counterparts to the new gamma-ray MSPs.

We observed the LAT sources now known to be associated with the two MSPs multiple times between 2009 and 2012, at a center frequency of 1390 MHz using an analog filterbank system. The total power from the central beam of the Parkes multibeam receiver (FWHM = $14'$), filtered through 512 frequency channels spanning a bandwidth of 256 MHz, was sampled every $125 \mu\text{s}$ and recorded for off-line analysis.

Typical integration times were around 1 hr, although they ranged between 41 and 136 minutes for these two sources (see Table 1 of Camilo et al. [102]). The LAT source associated with PSR J1035–6720 was observed a total of nine times, at a typical offset from the actual pulsar position of $4'.3$. This reduced the sensitivity of those observations to 80% of the maximum (on-axis) sensitivity. The LAT source associated with PSR J1744–7619 was observed ten times at a typical offset of $1'$, with no material impact on the sensitivity.




Figure 6.2: Left panels: constraints on MSP distances, as in Figure 11 of Abdo et al. [80]. Assuming certain gamma-ray efficiencies (η), constraints on the pulsar distances are inferred from their gamma-ray fluxes and spin-down powers (\dot{E}), using the measured proper motion (maximum likelihood value shown by the dashed line with the 1σ range given by the grey shaded region) to correct for the Shklovskii effect. The physically realistic region is to the left of the 100% efficiency line. The $\dot{f}_{\text{INT}} = 0$ line shows the required distance if the observed spin-down is purely due to the Shklovskii effect. Lines of constant tangential velocity are also shown. Right panels: Spin-down powers after correcting \dot{f} . For each assumed efficiency the range of valid distances (i.e. where the curves in the left panels cross the shaded proper motion ranges), and corresponding \dot{E} range, is shown by a shaded region. The dashed line shows the uncorrected \dot{E} , i.e. if there was no proper motion.

All the data were analyzed using standard pulsar search techniques implemented in PRESTO [265]. Before the new gamma-ray pulsars were known, the data were searched as described in Camilo et al. [102], considering possible dispersion measures (DMs) up to approximately twice the maximum DM predicted by the NE2001 model [266] for the corresponding line of sight (about 500 pc cm^{-3} for PSR J1035–6720 and 200 pc cm^{-3} for PSR J1744–7619; Table 1 of Camilo et al. [102]). Subsequently we re-analyzed the data sets using the timing ephemerides obtained from gamma-ray data, folding the radio data while searching only in DM. None of the observations yielded significant radio pulsar candidates.

Some of the MSPs discovered by Camilo et al. [102] were not detected in all their search observations. This can be caused by a combination of orbital acceleration, eclipses, or interstellar scintillation. The first two effects are not relevant for the two new isolated MSPs. Depending on the unknown DM (or equivalently distance) of the new MSPs, especially if they are nearby objects, interstellar scintillation could very much modulate any putative radio flux density recorded by our Parkes observing system. However, we know of no established radio-loud MSP in the Galactic disk that is undetectable as much as 90% of the time.

It therefore seems likely that PSRs J1035–6720 and J1744–7619 either have radio flux densities beamed towards the Earth that are below the nominal threshold of these searches, or if they emit radio waves their beams do not intersect the Earth. The nominal threshold for these searches depends on the putative radio duty cycle of these pulsars and their DM. For an assumed duty cycle of 25% and $\text{DM} \lesssim 50 \text{ pc cm}^{-3}$, an indicative 1.4 GHz flux density threshold for PSR J1035–6720 is about 0.25 mJy normalized to a 1 hr integration. For PSR J1744–7619 the equivalent figure is approximately 0.15 mJy.

How restrictive these S_{1400} flux density limits are depends on the assumed distance to the pulsars, which are as yet uncertain. For PSR J1035–6720, with its higher spin-down rate and therefore less well-constrained distance, we find a “pseudo-luminosity” upper limit of $L_{1400} = S_{1400}d^2 \approx 3.75 \text{ mJy kpc}^2$, which is greater than the average L_{1400} of MSPs in the ATNF Catalog [14] with flux density and distance measurements. However, for PSR J1744–7619, our estimated distance upper limit, assuming 100% gamma-ray efficiency, implies a maximum $L_{1400} \approx 0.15 \text{ mJy kpc}^2$. Just 3% of the MSPs in the ATNF Catalog have a lower pseudo-luminosity, and this number drops to zero if we take a more typical 10% efficiency. PSR J1744–7619 therefore has (at least) an unusually low radio luminosity. Deeper searches would be required to determine just how radio-quiet these pulsars are.

6.5 Pulse Profile Modeling

To investigate the possibility that the MSPs’ radio beams do not cross our line-of-sight, we modeled the pulsars’ gamma-ray emission geometry by fitting simulated pulse profiles to the observed photon phases, using the fitting technique described in Johnson et al. [85]. We considered three emission models: an outer gap (OG) model, a two-pole caustic (TPC) model and a pair-starved polar cap (PSPC) model. Johnson et al. [85] also give descriptions of these models, which we briefly

summarize here.

In the OG and TPC models, particle acceleration takes place in narrow “gaps” in the magnetosphere, where the plasma charge density deviates from the force-free configuration. In both models these gaps border the last-closed magnetic field lines. In the OG model [e.g. 67, 68], the lower bound of the gap is defined by the “null-charge surface”, where the plasma charge density changes sign. In the TPC model [267], the gap begins at the pulsar surface and extends to the light cylinder.

The PSPC model [268] is valid for low- \dot{E} pulsars, where pair creation may be insufficient to reach the force-free configuration [94], allowing for particle acceleration throughout the regions of open-field-lines. For illustrations of the acceleration regions in these models, we direct the reader to Figure 2 of Caraveo [47].

We used simulated rotation periods of 2.5 ms and 4.5 ms respectively and a simulated period derivative of $1 \times 10^{-20} \text{ s s}^{-1}$ for both. Instead of the Poisson likelihood, we used a χ^2 statistic to fit the weighted-counts light curves, using 60 bins for both MSPs, with background levels and uncertainties calculated as in Abdo et al. [80]. The best-fitting magnetic inclination angles (α), line-of-sight inclinations (ζ) and estimated beaming fractions (f_Ω) for each model are given in Table 6.2, with 95% confidence-level uncertainties estimated as described in Johnson et al. [85]. The best-fitting light curves from all models for each pulsar are shown in Figure 6.3. It should be noted that, as pointed out by Pierbattista et al. [86], fitting only the gamma-ray light curves with these geometric models can lead to systematic biases in the best-fit parameters of $\sim 10^\circ$.

For both MSPs, the models favored by the likelihood and best matching the qualitative features of the light curves predict radio emission, assuming a hollow-cone model [163], that should be beamed in our direction at both 1400 and 300 MHz. These predictions have our line of sight cutting across the emission cone and not merely clipping it. The lack of radio emission is therefore not easily explained by small offsets of a few degrees in the best-fitting parameters of these models. For the models that do not predict radio emission (TPC and OG for PSR J1035–6720 and OG for PSR J1744–7619), the lines-of-sight barely miss the edges of the emission cones at 300 MHz, suggesting that it might be possible to detect these MSPs at even lower frequencies. We note, however, that no model considered here successfully reproduces the broad double-peaked pulse profile of PSR J1035–6720.

6.6 Off-pulse and Spectral Analysis

Gamma-ray emission from pulsars is normally concentrated in one or more narrow peaks. However, depending on the viewing geometry and emission mechanism, a pulsar can have significant magnetospheric emission outside of the peaks of the pulse. The integrated pulse profile for PSR J1744–7619 includes unpulsed emission at all phases above the estimated background level, which is consistent with the best-fitting (TPC) emission model (shown in Figure 6.3). We performed a spectral analysis of this “off-pulse” emission to determine whether this is indeed magnetospheric, or contamination from a nearby source, and also to search for unpulsed emission from PSR J1035–6720.

We defined off-pulse regions for the two pulsars using the Bayesian Block decomposition method

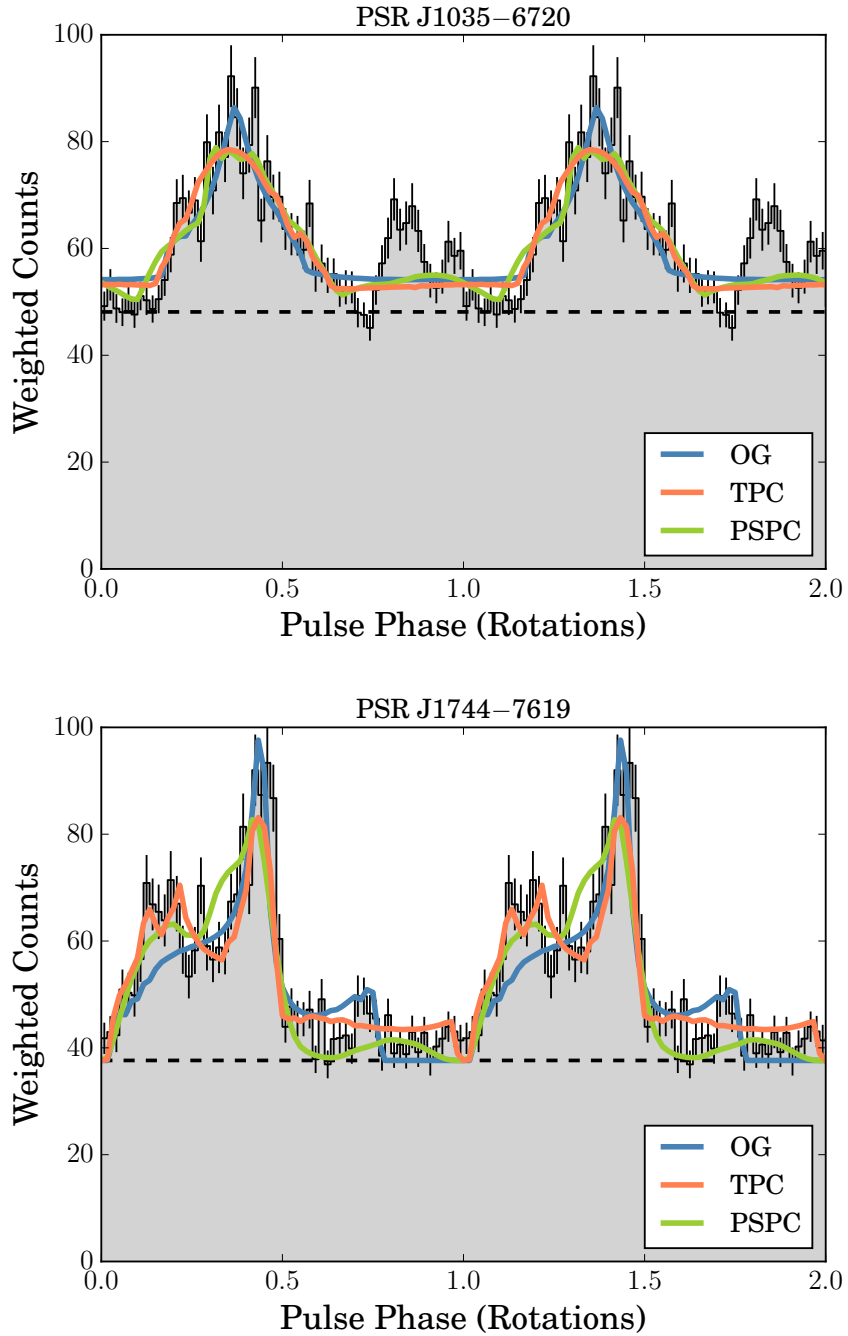


Figure 6.3: Gamma-ray pulse profiles of the newly detected MSPs. The overlaying curves are the best-fitting pulse profiles predicted by fits to outer-gap (OG), two-pole caustic (TPC) and pair-starved polar cap (PSPC) gamma-ray emission models. The dashed black line is the estimated background level, derived from the photon weights as in Abdo et al. [80].

Table 6.2. Best-fit parameters from gamma-ray pulse profile modeling

Parameter	PSR J1035–6720			PSR J1744–7619		
	TPC	OG	PSPC	TPC	OG	PSPC
$-\ln(L)$	75.5464	81.99390	72.8106	55.5374	81.3759	68.4602
α (deg)	7_{-2}^{+3}	9_{-3}^{+3}	51_{-5}^{+4}	62_{-2}^{+2}	66_{-2}^{+19}	62_{-1}^{+8}
ζ (deg)	70_{-1}^{+1}	75_{-1}^{+1}	78_{-4}^{+3}	44_{-4}^{+2}	32_{-6}^{+2}	74_{-3}^{+6}
f_{Ω}	$0.62_{-0.1}^{+0.08}$	$0.22_{-0.03}^{+0.25}$	$0.98_{-0.04}^{+0.16}$	$0.89_{-0.10}^{+0.02}$	$0.91_{-0.07}^{+0.08}$	$1.14_{-0.14}^{+0.15}$

Note. — For PSR J1744–7619, the confidence contours of the parameters of the OG model are not simply connected, see [85] for details.

described by Lande et al. [269]. The definitions of the off-pulse regions can be found in Table 6.1. It is worth noting that the duty cycle for PSR J1035–6720 is close to 100%; in fact only $\sim 14\%$ of the total spin phase was selected for the analysis. We first computed residual test statistic (TS) maps with gamma-ray photons coming solely from the off-peak region to look for putative sources around the pulsar position. No significant off-pulse emission was detected from PSR J1035–6720, although this could be due to the low photon statistics in the defined off-pulse region. Off-pulse emission consistent with the position of PSR J1744–7619 was detected with $TS = 127$. We computed the TS value at the pulsar position with both power-law (PL) and exponentially-cutoff power-law (PLEC) models to test for curvature of the gamma-ray spectrum. The off-pulse emission from PSR J1744–7619 has a significant cutoff, with $TS_{\text{cut}} = 42$, further suggesting a magnetospheric origin [80].

6.7 X-ray Observations

In a quest to find X-ray counterparts for unidentified 3FGL sources [270], the fields of PSRs J1035–6720 and J1744–7619 were observed by *XMM-Newton* for 25 ks. They were selected prior to the pulsars’ discoveries because their gamma-ray properties indicated a probable pulsar nature.

Plausible X-ray counterparts for both pulsars were detected with a significance greater than 10σ at locations consistent with the newly discovered pulsars’ timing positions. Their unabsorbed X-ray flux in the 0.3–10 keV energy range is $3.06_{-0.50}^{+0.96} \times 10^{-14}$ erg cm $^{-2}$ s $^{-1}$ for PSR J1035–6720 and $1.92_{-0.39}^{+0.59} \times 10^{-14}$ erg cm $^{-2}$ s $^{-1}$ for PSR J1744–7619. See Table 11 of Saz Parkinson et al. [270] for additional information. We computed the probability that the association between the X-ray source and the MSP is due to a chance coincidence using $P = 1 - \exp(-\pi\rho r^2)$, where r

is the matching radius (in our case the X-ray source error radius) and ρ is the density of X-ray objects in the *XMM-Newton* EPIC field, regardless of their flux. We estimated $P \sim 1.3 \times 10^{-4}$ for PSR J1035–6720 ($r = 2''$ and $\rho \sim 0.038 \text{ arcmin}^{-2}$) and $P \sim 1.4 \times 10^{-4}$ for PSR J1744–7619 ($r = 1.9''$ and $\rho \sim 0.046 \text{ arcmin}^{-2}$), which make a chance positional coincidence unlikely. In addition, the gamma-ray-to-X-ray flux ratios of the two likely X-ray counterparts (~ 700 and ~ 1100 respectively) are consistent with an MSP nature [199] confirming their associations.

6.8 Implications

The two new pulsars described in this chapter are the only known MSPs to remain undetected in radio observations. Direct gamma-ray pulsation searches such as those described in this thesis are the only way to detect and study such pulsars, and hence to gain a full picture of the local MSP population. The fact that our survey discovered two isolated MSPs also implies that there could be a far larger number of binary MSPs that are potentially only detectable by the LAT, assuming the ratio of binary/isolated radio-quiet MSPs is similar to that for radio-loud MSPs. The detection of binary MSPs requires orbital constraints from studies of their X-ray or optical counterparts, such as those that led to the discovery of PSR J1311–3430 [135, 191], to reduce the otherwise prohibitively large parameter space. Several such candidate systems have already been identified for 3FGL sources [e.g. 263, 271–274].

Mirabal et al. [117] highlighted a number of high Galactic latitude, unidentified LAT sources with curved spectra consistent with dark matter annihilation in ultra-faint dwarf spheroidal galaxies or dark matter subhalos [275, 276]. However, such sources are perhaps more likely to be MSPs (or globular clusters featuring many MSPs), whose spectra are almost indistinguishable. A sensitive survey specifically targeting these sources may be able to identify isolated MSPs amongst the group. While such sources are perhaps most likely to be MSPs in binary systems, blind searches such as this are necessary to rule out at least an isolated, radio-quiet MSP explanation for these dark matter candidates.

It has been suggested that MSPs have wide radio beams [203, 277] that are detectable over a broad range of possible viewing geometries, implying very few radio-quiet MSPs. The discovery of these new, potentially radio-quiet, MSPs could therefore have important implications and help discriminate between existing radio and gamma-ray emission models. Indeed, the lack of detectable radio pulsations from PSR J1744–7619 in particular already implies an unusually low radio luminosity, and is in tension with the best-fitting gamma-ray emission model.

Acknowledgments

We are extremely grateful to all volunteers who have supported the *Einstein@Home* project. Special thanks goes to those volunteers whose computers discovered the new MSPs³: PSR J1035–6720: “WSyS”, “Magpoo” and the ATLAS Cluster, AEI, Hannover, Germany; PSR J1744–7619: Darrell

³Where the volunteer’s name is unknown or private, we give the *Einstein@Home* username in quotation marks.

Hoberer, of Gainesville, TX, USA; the ATLAS Cluster, AEI, Hannover, Germany; Igor Yakushin of State College, PA, USA; and Keith Pickstone of Oldham, UK.

This work was supported by the Max-Planck-Gesellschaft (MPG), by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether Research Grant, No. PL 710/1-1 (PI: H. J. Pletsch), and by NSF award 1104902.

The Parkes Observatory is part of the Australia Telescope, which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO.

The Fermi LAT Collaboration acknowledges support from several agencies and institutes for both development and the operation of the LAT as well as scientific data analysis. These include NASA and Department of Energy (United States), CEA/Irfu and IN2P3/CNRS (France), ASI and INFN (Italy), MEXT, KEK, and JAXA (Japan), and the K. A. Wallenberg Foundation, the Swedish Research Council and the National Space Board (Sweden). Additional support from INAF in Italy and CNES in France for science analysis during the operations phase is also gratefully acknowledged. Fermi LAT data are available from the Fermi Science Support Center (<http://fermi.gsfc.nasa.gov/ssc>).

Chapter 7

Latest Results from the *Einstein@Home* Survey

While Chapter 4, which describes the results from the first 118 sources, was being written, the *Einstein@Home* survey continued to search through a further 34 sources. Because we sorted the target sources by how likely they were to contain pulsars, these latest sources could be considered to be less promising targets, either because their spectra are less typical of gamma-ray pulsars, or because they have a lower photon flux (details of the ranking procedure, which followed the techniques used by Lee et al. [95], will be given in J. Wu, et al., 2017, in preparation). However, despite this, the *Einstein@Home* survey has discovered two further pulsars from these sources.

These pulsars were found in Pass 8 data, covering LAT observations until 2015 July 07. Their initial search datasets were produced using the methods described in Section 4.2.1, but using the `P8R2_SOURCE_V6` IRFs, a zenith angle cut-off of $< 90^\circ$, and a larger 15° ROI. The Galactic and isotropic diffuse emission were modeled with the latest `gll_iem_v06.fits` [205] and `iso_P8R2_SOURCE_V6_v06.txt` templates respectively¹.

Again, after discovering the pulsars, we produced extended datasets including photons up to 2016 February 16, using `gtlike` to perform the binned spectral analysis with the pulsars' locations fixed at their preliminary timing positions. The unbinned timing analysis described in Section 4.4.1 was performed on these pulsars. The resulting photon phases and integrated pulse profiles are illustrated in Figure 7.1 and the timing and spectral properties of each pulsar are given in Table 7.1.

7.1 PSR J1641–5317

PSR J1641–5317 was discovered by the computers of *Einstein@Home* volunteers James Drews of UW-Madison, WI, USA and The ATLAS Cluster, AEI, Hannover, Germany, during a search of 3FGL J1641.5–5319. It is a fairly typical, “middle-aged” gamma-ray pulsar, and displays little

¹<http://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html>




Figure 7.1: Photon phases and integrated pulse profiles from newly discovered pulsars as in Figure 6.1.

Table 7.1. Properties of New Pulsars

Parameter	PSR J1641-5317	PSR J1817-1742
Range of Photon Data (MJD)	54682–57434	
Reference epoch (MJD)	55946	
Timing Model Parameters ^a		
R.A. (J2000.0)	16 ^h 41 ^m 15 ^s .91(6)	18 ^h 17 ^m 09 ^s .034(5)
Decl. (J2000.0)	−53° 17′ 48″.8(9)	−17° 42′ 01″.1(6)
Spin frequency, f (Hz)	5.71081786086(8)	6.67806213863(7)
Spin-down rate, \dot{f} (10^{-12} Hz s ^{−1})	−0.120706(2)	−0.917844(4)
Second frequency derivative, \ddot{f} (10^{-24} Hz s ^{−2})	...	6.0(2)
Derived Parameters ^b		
Galactic longitude, l (°)	333.29	13.34
Galactic latitude, b (°)	−4.56	−0.70
Spin Period, P (ms)	175.106267502(2)	149.744039399(2)
Period derivative, \dot{P} (10^{-15} s s ^{−1})	3.70110(7)	20.58107(9)
Spin-down power, \dot{E} (10^{33} erg s ^{−1})	27	242
Characteristic age, τ_c (kyr)	750	115
Surface magnetic field strength, B_S (10^{12} G)	0.8	1.8
Light-cylinder magnetic field strength, B_{LC} (10^3 G)	1.4	4.9
Spectral Parameters ^c		
Spectral index, Γ	1.95 ± 0.02	0.55 ± 0.11
Cutoff Energy, E_c (GeV)	3.47 ± 0.12	1.94 ± 0.15
Photon flux > 100 MeV, F_{100} (10^{-8} cm ^{−2} s ^{−1})	3.40 ± 0.13	0.67 ± 0.09
Energy flux > 100 MeV, G_{100} (10^{-11} erg cm ^{−2} s ^{−1})	1.93 ± 0.05	1.31 ± 0.24

^aThe quoted values of the timing parameters are the mean values of the posterior distributions from the results of timing analyses following the procedures described in Section 4.4.1 with 1σ uncertainties on the final digits given in brackets.

^bThe derived parameters (\dot{E} , τ_c , B_S , B_{LC}) are calculated as in [e.g. 80, 151].

^cThe uncertainties quoted on the spectral parameters are statistical only, and do not include estimates of the systematic uncertainties from the LAT’s collecting area and Galactic diffuse emission, which are likely to be larger. Additionally, PSR J1817–1742 lies less than 15° from the Galactic center, where the systematic uncertainties from the Galactic diffuse emission are likely to be even higher. The reported spectral parameters should therefore be treated with caution.

timing noise, in that the inclusion of a second frequency derivative in the timing model did not decrease the BIC, giving a 2σ upper limit of $|\ddot{f}| < 3.8 \times 10^{-25}$. We note that, given the pulsar’s singly peaked pulse profile, it is possible that the timing solution given in Table 7.1 corresponds to its second frequency harmonic. Folding at half the spin frequency resulted in an H -test value which neither increased nor decreased significantly. However, the two resulting peaks appeared very similar, as one would expect from folding the data at half of the pulsar’s true spin frequency.

7.2 PSR J1817–1742

PSR J1817–1742 was discovered by the computers of *Einstein@Home* volunteers Dave Billenness of Victoria BC, Canada and Jonathan L. Kerfoot of Bixby, Oklahoma, USA while they searched for pulsations from 3FGL J1817.2–1739. The source in which it was discovered has a LAT point-source test statistic (TS) of just 108.9 (corresponding to a significance of $\approx 10\sigma$), making it the least significant unidentified source in which pulsations have been detected (see Figure 4.9 for the H -test and significances of blind-search gamma-ray pulsars at their time of discovery). This is very close to the threshold value estimated in Chapter 4 as the minimum TS a pulsar can have while still being detectable by our survey. Despite its weak photon flux, PSR J1817–1742 was significantly detected by our survey at its second and fourth spin harmonics, most likely as a result of its extremely narrow pulse profile, visible in Figure 7.1b.

Seven frequency derivative terms were required in the timing solution for PSR J1817–1742 to minimize the BIC. In the early part of the *Fermi* mission, its spin frequency appears to decrease exponentially, as one may expect if the pulsar was relaxing from a glitch prior to the beginning of the LAT data. The evolution of the spin frequency and spin-down rate of PSR J1817–1742 is shown in Figure 7.2. Shannon et al. [278] recently reported a similar detection of a glitch from the Vela pulsar occurring before the start of their observations².

A timing model consisting of five frequency derivatives and an exponentially decaying frequency term in the early mission gives a slightly lower likelihood ($\Delta \log \mathcal{L} \approx -10$) with the same number of free parameters as the pure Taylor series model. This decrease in the log-likelihood could potentially be explained by the fact that we have removed two higher frequency derivatives, which can reproduce faster phase variations throughout the data, and replaced these with the two glitch recovery parameters, which can only model a decreasing frequency in the early data. That is, the glitch recovery could accurately model the pulsar’s phase in the early mission, but the higher frequency derivative terms included in the Taylor series model additionally improve the log-likelihood in the late mission. The best-fitting decaying frequency term corresponds to a recovering glitch increment of $\Delta f_{\text{D}} = (1.4 \pm 0.1) \times 10^{-7} e^{-(t_g - t_0)/\tau_{\text{D}}}$ Hz, with a decay constant $\tau_{\text{D}} = 218 \pm 11$ days, where t_g is the unknown glitch epoch, occurring before the start of the data interval at $t_0 = 54682$ MJD. The magnitudes of these parameters are consistent with those of typical pulsar glitch recoveries [33, 34]. This could therefore be a hint that PSR J1817–1742 may suffer large glitches, which may be observed in future LAT data as the *Fermi* mission continues.

²With thanks to M. Kerr for bringing this to the author’s attention.




Figure 7.2: Evolution of the spin frequency and spin-down rate of PSR J1817–1742 during the LAT observation time, after subtracting the spin-down rate measured at the reference epoch, as in Figure 4.6. The decreasing frequency residuals visible in the first part of the data could suggest that the pulsar is relaxing following a large glitch occurring before the start of the *Fermi* mission.

There is a supernova remnant, G013.3–1.3 [200], just over half a degree away from the timing position of PSR J1817–1742. Given the pulsar’s characteristic age, and the apparent distance to the SNR of ~ 3 kpc [279], then PSR J1817–1742 would require a proper motion of ~ 9 mas yr $^{-1}$ and a kick velocity of around 120 km s $^{-1}$. We added proper motion parameters to our timing model, and measured a significant ($> 3\sigma$) velocity in the α direction of -120 ± 30 mas yr $^{-1}$. This is at least qualitatively in the right direction for an association with the SNR, however at a distance of 3 kpc this implies a tangential velocity of $\sim 1,700$ km s $^{-1}$, which is very high compared to known pulsar velocities.

Pulsar timing positions and proper motions can be significantly biased by timing noise [165]; an unbiased determination of the pulsar’s motion would require joint modeling of the pulsar’s spin and positional parameters with the parameters of an assumed timing noise power spectrum.

7.3 Future Gamma-ray Pulsar Surveys

One lesson learned from our sensitivity investigation in Chapter 4 is that there remain a number of promising pulsar-like candidates amongst the unidentified 3FGL sources. The ability to determine quantitatively in advance which sources our survey is most sensitive to will allow future surveys to focus most of their computational power on promising sources, and waste little on the others. As mentioned in the discussion in Section 4.5, the computational resources required to maintain a low sensitivity threshold will increase as the *Fermi* mission continues, so savings like this will become even more crucial.

One way to substantially reduce the cost of a pulsar search is to target the searched sky position at the location of a potential counterpart [e.g. 132]. There are a number of 3FGL sources for which likely multiwavelength counterparts, with precise positions, have been identified [e.g. 270] at which we could target searches. A notable example would be the Cas A neutron star discussed in Section 4.3.2. Such searches could feasibly be run quickly (and sensitively) “offline” on a cluster computer.

Perhaps one of the most important outstanding questions of high energy astronomy is the origin of the GeV excess observed towards the Galactic Center (GC) [217]. The high dispersion and absorption along the line-of-sight to the GC limit the sensitivity of radio searches; blind gamma-ray searches are potentially the only way to discern between a pulsar explanation and a dark matter interpretation. Searching for pulsations from sources in the inner galaxy will therefore be a key focus. In Section 4.3.3, we estimate that our blind search may have been able to detect a Crab-like pulsar at the Galactic Center, if such a pulsar existed and did not glitch during the observation interval. A dedicated survey focusing only on this region may be able to gain additional sensitivity, and allow more common, less luminous pulsars to be detected.

As mentioned in Chapter 6, blind searches targeting high galactic latitude sources may also be required to rule out a radio-quiet MSP origin for potential dark matter candidates. However, such searches can only detect isolated pulsars, with the detection of binary MSPs requiring additional observations to constrain their orbital parameters.

This brings us nicely on to the current work being performed on *Einstein@Home*. X-ray and optical counterparts with periodic flux modulations consistent with binary pulsar systems have been identified for a number of 3FGL sources [e.g. 263, 271–274]. Using the constraints on the orbital parameters that spectroscopic and photometric observations have provided for these candidates, targeted searches, much like the one in which PSR J1311–3430 was discovered [135], are being prepared by L. Nieder. The first of these is currently running on *Einstein@Home*, and includes many of the sensitivity improvements developed in Chapter 2.

We are hopeful that the methods developed during this work will lead to interesting new pulsar discoveries, and that the pulsars discovered by the *Einstein@Home* survey will contribute to the understanding of the properties and behavior of neutron stars.

Bibliography

- [1] Baade, W., & Zwicky, F. 1934, *Physical Review*, 46, 76
- [2] Kaspi, V. M., & Helfand, D. J. 2002, in *Astronomical Society of the Pacific Conference Series*, Vol. 271, *Neutron Stars in Supernova Remnants*, ed. P. O. Slane & B. M. Gaensler, 271, 3, arXiv:astro-ph/0201183
- [3] Lyne, A. G., Jordan, C. A., Graham-Smith, F., et al. 2015, *MNRAS*, 446, 857, arXiv:1410.0886
- [4] Truemper, J., Pietsch, W., Reppin, C., et al. 1978, *ApJ*, 219, L105
- [5] Wheaton, W. A., Doty, J. P., Primini, F. A., et al. 1979, *Nature*, 282, 240
- [6] Thompson, C., & Duncan, R. C. 1993, *ApJ*, 408, 194
- [7] Reisenegger, A. 2001, in *Astronomical Society of the Pacific Conference Series*, Vol. 248, *Magnetic Fields Across the Hertzsprung-Russell Diagram*, ed. G. Mathys, S. K. Solanki, & D. T. Wickramasinghe, 248, 469, arXiv:astro-ph/0103010
- [8] Bignami, G. F., Caraveo, P. A., De Luca, A., & Mereghetti, S. 2003, *Nature*, 423, 725, arXiv:astro-ph/0306189
- [9] Lorimer, D., & Kramer, M. 2005, *Handbook of Pulsar Astronomy* (Cambridge: Cambridge University Press)
- [10] Goldreich, P., & Julian, W. H. 1969, *ApJ*, 157, 869
- [11] Sturrock, P. A. 1971, *ApJ*, 164, 529
- [12] Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- [13] Gold, T. 1968, *Nature*, 218, 731
- [14] Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993, arXiv:astro-ph/0412641
- [15] Bell, J. F. 1998, in *NATO Advanced Science Institutes (ASI) Series C*, Vol. 515, *NATO Advanced Science Institutes (ASI) Series C*, ed. R. Buccheri, J. van Paradijs, & A. Alpar, 515, 159, arXiv:astro-ph/9702150
- [16] Hobbs, G., Lyne, A. G., & Kramer, M. 2010, *MNRAS*, 402, 1027, arXiv:0912.4537
- [17] Turolla, R., Zane, S., & Watts, A. L. 2015, *Reports on Progress in Physics*, 78, 116901, arXiv:1507.02924

- [18] Lorimer, D. R. 2008, *Living Reviews in Relativity*, 11, 8, arXiv:0811.0762, cited on 21-May-2016
- [19] Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *Nature*, 300, 728
- [20] Tauris, T. M., Kaspi, V. M., Breton, R. P., et al. 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 39, arXiv:1501.00005
- [21] Lyne, A. G. 1994, in *Astrophysics and Space Science Library*, Vol. 187, *Frontiers of Space and Ground-Based Astronomy*, ed. W. Wamsteker, M. S. Longair, & Y. Kondo, 187, 277
- [22] Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, 195, L51
- [23] Nice, D. J., Stairs, I. H., & Kasian, L. E. 2008, in *American Institute of Physics Conference Series*, Vol. 983, *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More*, ed. C. Bassa, Z. Wang, A. Cumming, & V. M. Kaspi, 983, 453
- [24] Kiziltan, B., Kottas, A., De Yoreo, M., & Thorsett, S. E. 2013, *ApJ*, 778, 66, arXiv:1011.4291
- [25] Blandford, R., & Teukolsky, S. A. 1976, *ApJ*, 205, 580
- [26] Lyne, A. G., Burgay, M., Kramer, M., et al. 2004, *Science*, 303, 1153, arXiv:astro-ph/0401086
- [27] Damour, T. 2015, *Classical and Quantum Gravity*, 32, 124009, arXiv:1411.3930
- [28] Ozel, F., & Freire, P. 2016, *ArXiv e-prints*, arXiv:1603.02698
- [29] Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, *Reviews of Modern Physics*, 88, 021001, arXiv:1602.01081
- [30] Janka, H.-T., Langanke, K., Marek, A., Martínez-Pinedo, G., & Müller, B. 2007, *Phys. Rep.*, 442, 38, arXiv:astro-ph/0612072
- [31] Hotokezaka, K., Kyutoku, K., Okawa, H., Shibata, M., & Kiuchi, K. 2011, *Phys. Rev. D*, 83, 124008, arXiv:1105.4370
- [32] Radhakrishnan, V., & Manchester, R. N. 1969, *Nature*, 222, 228
- [33] Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, *MNRAS*, 414, 1679, arXiv:1102.1743
- [34] Yu, M., Manchester, R. N., Hobbs, G., et al. 2013, *MNRAS*, 429, 688, arXiv:1211.2035
- [35] Alpar, M. A., Anderson, P. W., Pines, D., & Shaham, J. 1981, *ApJ*, 249, L29
- [36] Baym, G., Pethick, C., Pines, D., & Ruderman, M. 1969, *Nature*, 224, 872

- [37] Middleditch, J., Marshall, F. E., Wang, Q. D., Gotthelf, E. V., & Zhang, W. 2006, *ApJ*, 652, 1531, arXiv:astro-ph/0605007
- [38] Taylor, Jr., J. H. 1991, *IEEE Proceedings*, 79, 1054
- [39] Foster, R. S., & Backer, D. C. 1990, *ApJ*, 361, 300
- [40] Manchester, R. N., Hobbs, G., Bailes, M., et al. 2013, *PASA*, 30, e017, arXiv:1210.6130
- [41] The NANOGrav Collaboration, Arzoumanian, Z., Brazier, A., et al. 2015, *ApJ*, 813, 65, arXiv:1505.07540
- [42] Verbiest, J. P. W., Lentati, L., Hobbs, G., et al. 2016, *MNRAS*, 458, 1267, arXiv:1602.03640
- [43] Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, *MNRAS*, 458, 3341, arXiv:1602.08511
- [44] Ransom, S. M., Stairs, I. H., Archibald, A. M., et al. 2014, *Nature*, 505, 520, arXiv:1401.0535
- [45] Shao, L. 2016, *Phys. Rev. D*, 93, 084023, arXiv:1602.05725
- [46] Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145
- [47] Caraveo, P. A. 2014, *ARA&A*, 52, 211, arXiv:1312.2913
- [48] McBreen, B., Ball, Jr., S. E., Campbell, M., Greisen, K., & Koch, D. 1973, *ApJ*, 184, 571
- [49] Bignami, G. F., Boella, G., Burger, J. J., et al. 1975, *Space Science Instrumentation*, 1, 245
- [50] Fichtel, C. E., Hartman, R. C., Kniffen, D. A., et al. 1975, *ApJ*, 198, 163
- [51] Derdeyn, S. M., Ehrmann, C. H., Fichtel, C. E., Kniffen, D. A., & Ross, R. W. 1972, *Nuclear Instruments and Methods*, 98, 557
- [52] The Caravane Collaboration. 1977, in *NASA Conference Publication*, Vol. 2, in *The Structure and Content of the Galaxy and Galactic Gamma Rays*, ed. C. E. Fichtel & F. W. Stecker, 2, 29
- [53] Kniffen, D. A., Hartman, R. C., Thompson, D. J., Bignami, G. F., & Fichtel, C. E. 1974, *Nature*, 251, 397
- [54] Thompson, D. J., Fichtel, C. E., Kniffen, D. A., & Ogelman, H. B. 1975, *ApJ*, 200, L79
- [55] Buccheri, R., Bennett, K., Bignami, G. F., et al. 1983, *A&A*, 128, 245
- [56] Buccheri, R., Kanbach, G., Caraveo, P., et al. 1978, *A&A*, 69, 141
- [57] Hartman, R. C., Bertsch, D. L., Fichtel, C. E., et al. 1992, in *The Compton Observatory Science Workshop*, Vol. 3137, *NASA Conference Publication*, ed. C. R. Shrader, N. Gehrels, & B. Dennis, 3137, 116

- [58] Schoenfelder, V., Aarts, H., Bennett, K., et al. 1993, *ApJS*, 86, 657
- [59] Gehrels, N., Chipman, E., & Kniffen, D. A. 1993, *A&AS*, 97, 5
- [60] Thompson, D. J. 2004, *Cosmic Gamma-Ray Sources*, ed. K. S. Cheng & G. E. Romero (Dordrecht: Springer Netherlands), 149–168
- [61] Ramanamurthy, P. V., Fichtel, C. E., Kniffen, D. A., Sreekumar, P., & Thompson, D. J. 1996, *ApJ*, 458, 755
- [62] Kuiper, L., Hermsen, W., Verbunt, F., et al. 2000, *A&A*, 359, 615, arXiv:astro-ph/0005338
- [63] Kaspi, V. M., Lackey, J. R., Mattox, J., et al. 2000, *ApJ*, 528, 445, arXiv:astro-ph/9906373
- [64] Arons, J. 1983, *ApJ*, 266, 215
- [65] Muslimov, A. G., & Harding, A. K. 2003, *ApJ*, 588, 430, arXiv:astro-ph/0301023
- [66] Harding, A. K., Stern, J. V., Dyks, J., & Frackowiak, M. 2008, *ApJ*, 680, 1378, arXiv:0803.0699
- [67] Cheng, K. S., Ho, C., & Ruderman, M. 1986, *ApJ*, 300, 500
- [68] Romani, R. W. 1996, *ApJ*, 470, 469
- [69] Arons, J. 1996, *A&AS*, 120, 49
- [70] Barbiellini, G., Boezio, M., Casolino, M., et al. 1995, *Nuclear Instruments and Methods in Physics Research A*, 354, 547
- [71] Tavani, M., Barbiellini, G., Argan, A., et al. 2009, *A&A*, 502, 995, arXiv:0807.4254
- [72] Halpern, J. P., Camilo, F., Giuliani, A., et al. 2008, *ApJ*, 688, L33, arXiv:0810.0008
- [73] Pellizzoni, A., Pilia, M., Possenti, A., et al. 2009, *ApJ*, 695, L115, arXiv:0903.0087
- [74] Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *ApJ*, 697, 1071, arXiv:0902.1089
- [75] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *Astroparticle Physics*, 32, 193, arXiv:0904.2226
- [76] Ackermann, M., Ajello, M., Albert, A., et al. 2012, *ApJS*, 203, 4, arXiv:1206.1896
- [77] Atwood, W., Albert, A., Baldini, L., et al. 2012, in *Proceedings of the 4th Fermi Symposium*, ed. T. J. Brandt, N. Omodei, & C. Wilson-Hodge, eConf C121028, 8, arXiv:1303.3514
- [78] Aliu, E., Anderhub, H., Antonelli, L. A., et al. 2008, *Science*, 322, 1221, arXiv:0809.2998
- [79] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJS*, 187, 460, arXiv:0910.1608

- [80] Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *ApJS*, 208, 17, arXiv:1305.4385
- [81] Smith, D. A., & Thompson, D. J. 2009, *Neutron Stars and Pulsars*, ed. W. Becker (Berlin, Heidelberg: Springer Berlin Heidelberg), 621–649
- [82] Abdo, A. A., Ackermann, M., Atwood, W. B., et al. 2009, *ApJ*, 696, 1084, arXiv:0812.2960
- [83] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 708, 1254, arXiv:0911.2412
- [84] —. 2010, *ApJ*, 720, 272, arXiv:1007.1142
- [85] Johnson, T. J., Venter, C., Harding, A. K., et al. 2014, *ApJS*, 213, 6, arXiv:1404.2264
- [86] Pierbattista, M., Harding, A. K., Grenier, I. A., et al. 2015, *A&A*, 575, A3, arXiv:1403.3849
- [87] Grenier, I. A., & Harding, A. K. 2015, *Comptes Rendus Physique*, 16, 641, arXiv:1509.08823
- [88] Wang, R.-B., & Hirotani, K. 2011, *ApJ*, 736, 127, arXiv:1105.3030
- [89] Guillemot, L., Smith, D. A., Laffon, H., et al. 2016, *A&A*, 587, A109, arXiv:1601.05987
- [90] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 725, L73, arXiv:1011.0091
- [91] Abdo, A. A., Ackermann, M., Atwood, W. B., et al. 2009, *ApJ*, 699, 1171, arXiv:0904.4377
- [92] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *Science*, 325, 848
- [93] Freire, P. C. C., Abdo, A. A., Ajello, M., et al. 2011, *Science*, 334, 1107, arXiv:1111.3754
- [94] Harding, A. K., Muslimov, A. G., & Zhang, B. 2002, *ApJ*, 576, 366, arXiv:astro-ph/0205077
- [95] Lee, K. J., Guillemot, L., Yue, Y. L., Kramer, M., & Champion, D. J. 2012, *MNRAS*, 424, 2832, arXiv:1205.6221
- [96] Cognard, I., Guillemot, L., Johnson, T. J., et al. 2011, *ApJ*, 732, 47, arXiv:1102.4192
- [97] Keith, M. J., Johnston, S., Ray, P. S., et al. 2011, *MNRAS*, 414, 1292, arXiv:1102.0648
- [98] Hessels, J. W. T., Roberts, M. S. E., McLaughlin, M. A., et al. 2011, in *American Institute of Physics Conference Series*, Vol. 1357, American Institute of Physics Conference Series, ed. M. Burgay, N. D’Amico, P. Esposito, A. Pellizzoni, & A. Possenti, 1357, 40, arXiv:1101.1742
- [99] Ransom, S. M., Ray, P. S., Camilo, F., et al. 2011, *ApJ*, 727, L16, arXiv:1012.2862
- [100] Kerr, M., Camilo, F., Johnson, T. J., et al. 2012, *ApJ*, 748, L2, arXiv:1201.5160
- [101] Barr, E. D., Guillemot, L., Champion, D. J., et al. 2013, *MNRAS*, 429, 1633, arXiv:1301.0359
- [102] Camilo, F., Kerr, M., Ray, P. S., et al. 2015, *ApJ*, 810, 85, arXiv:1507.04451

- [103] Cromartie, H. T., Camilo, F., Kerr, M., et al. 2016, *ApJ*, 819, 34, arXiv:1601.05343
- [104] Ray, P. S. 2016, Fermi's contribution to understanding the Galactic population of millisecond pulsars, JINA-CEE International Symposium on Neutron Stars in the Multi-Messenger Era: Prospects & Challenges Ohio University, Athens, Ohio, USA, http://www.phy.ohiou.edu/~SoNS/speakers/Symposium/Ray,Pauls_FermiGalacticMSPs.pdf, [Online; accessed 2-Jun-2016]
- [105] Deneva, J. S., Ray, P. S., Camilo, F., et al. 2016, *ApJ*, 823, 105, arXiv:1601.03681
- [106] Stappers, B. W., Bailes, M., Manchester, R. N., Sandhu, J. S., & Toscano, M. 1998, *ApJ*, 499, L183, arXiv:astro-ph/9804108
- [107] Johnson, T. J., Ray, P. S., Roy, J., et al. 2015, *ApJ*, 806, 91, arXiv:1502.06862
- [108] Stappers, B. W., Archibald, A. M., Hessels, J. W. T., et al. 2014, *ApJ*, 790, 39, arXiv:1311.7506
- [109] Papitto, A., Ferrigno, C., Bozzo, E., et al. 2013, *Nature*, 501, 517, arXiv:1305.3884
- [110] Hou, X., Smith, D. A., Guillemot, L., et al. 2014, *A&A*, 570, A44, arXiv:1407.6271
- [111] Fermi LAT Collaboration. 2015, *Science*, 350, 801
- [112] Buehler, R., Scargle, J. D., Blandford, R. D., et al. 2012, *ApJ*, 749, 26, arXiv:1112.1979
- [113] Allafort, A., Baldini, L., Ballet, J., et al. 2013, *ApJ*, 777, L2, arXiv:1308.0358
- [114] Weltevrede, P., Abdo, A. A., Ackermann, M., et al. 2010, *ApJ*, 708, 1426, arXiv:0911.3063
- [115] Smith, D. A., Guillemot, L., Camilo, F., et al. 2008, *A&A*, 492, 923, arXiv:0810.1637
- [116] Pierbattista, M., Grenier, I. A., Harding, A. K., & Gonthier, P. L. 2012, *A&A*, 545, A42, arXiv:1206.5634
- [117] Mirabal, N., Charles, E., Ferrara, E. C., et al. 2016, *ApJ*, 825, 69, arXiv:1605.00711
- [118] Hooper, D., Cholis, I., Linden, T., Siegal-Gaskins, J. M., & Slatyer, T. R. 2013, *Phys. Rev. D*, 88, 083009, arXiv:1305.0830
- [119] Lee, S. K., Lisanti, M., Safdi, B. R., Slatyer, T. R., & Xue, W. 2016, *Physical Review Letters*, 116, 051103, arXiv:1506.05124
- [120] O'Leary, R. M., Kistler, M. D., Kerr, M., & Dexter, J. 2016, *ArXiv e-prints*, arXiv:1601.05797
- [121] Watters, K. P., Romani, R. W., Weltevrede, P., & Johnston, S. 2009, *ApJ*, 695, 1289, arXiv:0812.3931

- [122] Watters, K. P., & Romani, R. W. 2011, *ApJ*, 727, 123, arXiv:1009.5305
- [123] Bignami, G. F., Caraveo, P. A., & Lamb, R. C. 1983, *ApJ*, 272, L9
- [124] Halpern, J. P., & Holt, S. S. 1992, *Nature*, 357, 222
- [125] Bertsch, D. L., Brazier, K. T. S., Fichtel, C. E., et al. 1992, *Nature*, 357, 306
- [126] Mattox, J. R., Halpern, J. P., & Caraveo, P. A. 1998, *ApJ*, 493, 891
- [127] Chandler, A. M., Koh, D. T., Lamb, R. C., et al. 2001, *ApJ*, 556, 59
- [128] Brady, P. R., & Creighton, T. 2000, *Phys. Rev. D*, 61, 082001, arXiv:gr-qc/9812014
- [129] Atwood, W. B., Ziegler, M., Johnson, R. P., & Baughman, B. M. 2006, *ApJ*, 652, L49
- [130] Ziegler, M., Baughman, B. M., Johnson, R. P., & Atwood, W. B. 2008, *ApJ*, 680, 620
- [131] Abdo, A. A., Ackermann, M., Atwood, W. B., et al. 2008, *Science*, 322, 1218, arXiv:0810.3562
- [132] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *Science*, 325, 840, arXiv:1009.0748
- [133] Saz Parkinson, P. M., Dormody, M., Ziegler, M., Ray, P. S., et al. 2010, *ApJ*, 725, 571, arXiv:1006.2134
- [134] Saz Parkinson, P. M., Belfiore, A., Dormody, M., Ziegler, M., & Razzano, M. 2011, Blind Searches for γ -ray Pulsars in 2 years of Fermi-LAT data, Fermi Symposium 2011, https://confluence.slac.stanford.edu/download/attachments/102860834/PSP_Fermi_Symposium_2011_v3.pdf, [Online; accessed 16-May-2016]
- [135] Pletsch, H. J., Guillemot, L., Fehrmann, H., et al. 2012, *Science*, 338, 1314, arXiv:1211.1385
- [136] Pletsch, H. J., Guillemot, L., Allen, B., et al. 2013, *ApJ*, 779, L11, arXiv:1311.6427
- [137] Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, 372, 1549, arXiv:astro-ph/0607664
- [138] Acero, F., Ackermann, M., Ajello, M., et al. 2015, *ApJS*, 218, 23, arXiv:1501.02003
- [139] Cutler, C., Gholami, I., & Krishnan, B. 2005, *Phys. Rev. D*, 72, 042004, arXiv:gr-qc/0505082
- [140] Prix, R. 2009, *Neutron Stars and Pulsars*, ed. W. Becker (Berlin, Heidelberg: Springer Berlin Heidelberg), 651–685
- [141] Astone, P., Borkowski, K. M., Jaranowski, P., Pietka, M., & Królak, A. 2010, *Phys. Rev. D*, 82, 022005, arXiv:1003.0844

- [142] Krishnan, B., Sintès, A. M., Papa, M. A., et al. 2004, *Phys. Rev. D*, 70, 082001, arXiv:gr-qc/0407001
- [143] Abbott, B., Abbott, R., Adhikari, R., et al. 2008, *Phys. Rev. D*, 77, 022001, arXiv:0708.3818
- [144] —. 2009, *Phys. Rev. D*, 79, 022001, arXiv:0804.1747
- [145] Abbott, B. P., Abbott, R., Adhikari, R., et al. 2009, *Phys. Rev. D*, 80, 042003, arXiv:0905.1705
- [146] Aasi, J., Abadie, J., Abbott, B. P., et al. 2013, *Phys. Rev. D*, 87, 042001, arXiv:1207.7176
- [147] Prix, R., & Shaltev, M. 2012, *Phys. Rev. D*, 85, 084010, arXiv:1201.4321
- [148] Pletsch, H. J. 2010, *Phys. Rev. D*, 82, 042002, arXiv:1005.0395
- [149] Wette, K. 2015, *Phys. Rev. D*, 92, 082003, arXiv:1508.02372
- [150] Pletsch, H. J. 2011, *Phys. Rev. D*, 83, 122003, arXiv:1101.5396
- [151] Pletsch, H. J., Guillemot, L., Allen, B., et al. 2012, *ApJ*, 744, 105, arXiv:1111.0523
- [152] Kerr, M. 2011, *ApJ*, 732, 38, arXiv:1103.2128
- [153] Aulbert, C., & Fehrmann, H. 2009, *Gravitationswellensuche mit einem der weltweit größten Supercomputer*, <https://www.mpg.de/308429/forschungsSchwerpunkt?c=166022>, [Online; accessed 1-Jun-2016]
- [154] Pletsch, H. J., Guillemot, L., Allen, B., et al. 2012, *ApJ*, 755, L20, arXiv:1207.5333
- [155] Allen, B., Knispel, B., Cordes, J. M., et al. 2013, *ApJ*, 773, 91, arXiv:1303.0028
- [156] Anderson, D. P. 2004, in *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing, GRID '04* (Washington, DC, USA: IEEE Computer Society), 4–10
- [157] Anderson, D. P., Christensen, C., & Allen, B. 2006, in *Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, SC '06* (New York, NY, USA: ACM), doi:10.1145/1188455.1188586
- [158] Strohmaier, E., Dongarra, J., Simon, H., & Meuer, M. 2015, *Top500 List - November 2015*, <http://www.top500.org/list/2015/11/>, [Online; accessed 16-May-2016]
- [159] Knispel, B., Allen, B., Cordes, J. M., et al. 2010, *Science*, 329, 1305, arXiv:1008.2172
- [160] Knispel, B., Lazarus, P., Allen, B., et al. 2011, *ApJ*, 732, L1, arXiv:1102.5340
- [161] Knispel, B., Eatough, R. P., Kim, H., et al. 2013, *ApJ*, 774, 93, arXiv:1302.0467
- [162] Knispel, B., Lyne, A. G., Stappers, B. W., et al. 2015, *ApJ*, 806, 140, arXiv:1504.03684

- [163] Story, S. A., Gonthier, P. L., & Harding, A. K. 2007, *ApJ*, 671, 713, arXiv:0706.3041
- [164] Takata, J., Wang, Y., & Cheng, K. S. 2011, *MNRAS*, 415, 1827, arXiv:1104.0065
- [165] Kerr, M., Ray, P. S., Johnston, S., Shannon, R. M., & Camilo, F. 2015, *ApJ*, 814, 128, arXiv:1510.05099
- [166] Pletsch, H. J., & Clark, C. J. 2015, *ApJ*, 807, 18, arXiv:1504.07466
- [167] Ray, P. S., Kerr, M., Parent, D., et al. 2011, *ApJS*, 194, 17, arXiv:1011.2468
- [168] Pletsch, H. J., & Clark, C. J. 2014, *ApJ*, 795, 75, arXiv:1408.6962
- [169] Clark, C. J., Pletsch, H. J., Wu, J., et al. 2015, *ApJ*, 809, L2, arXiv:1508.00779
- [170] Clark, C. J., Wu, J., Pletsch, H. J., et al. 2017, *ApJ*, 834, 106, arXiv:1611.01015
- [171] Clark, C. J., Pletsch, H. J., Wu, J., et al. 2016, *ApJ*, 832, L15, arXiv:1611.01292
- [172] Bickel, P., Kleijn, B., & Rice, J. 2008, *ApJ*, 685, 384, arXiv:0706.4108
- [173] Beran, R. J. 1969, *Ann. Math. Stat.*, 40, 1196
- [174] Groth, E. J. 1975, *ApJS*, 29, 285
- [175] Guidorzi, C. 2011, *MNRAS*, 415, 3561, arXiv:1104.5308
- [176] de Jager, O. C., Raubenheimer, B. C., & Swanepoel, J. W. H. 1989, *A&A*, 221, 180
- [177] Brady, P. R., Creighton, T., Cutler, C., & Schutz, B. F. 1998, *Phys. Rev. D*, 57, 2101, arXiv:gr-qc/9702050
- [178] Prix, R. 2007, *Classical and Quantum Gravity*, 24, S481, arXiv:0707.0428
- [179] Pletsch, H. J., & Allen, B. 2009, *Physical Review Letters*, 103, 181102, arXiv:0906.0023
- [180] Balasubramanian, R., Sathyaprakash, B. S., & Dhurandhar, S. V. 1996, *Phys. Rev. D*, 53, 3033, arXiv:gr-qc/9508011
- [181] Owen, B. J. 1996, *Phys. Rev. D*, 53, 6749, arXiv:gr-qc/9511032
- [182] Frigo, M., & Johnson, S. G. 2005, *Proc. of the IEEE*, 93, 216, special issue on “Program Generation, Optimization, and Platform Adaptation”
- [183] Blackman, R. B., & Tukey, J. W. 1958, *Bell Syst. Tech. J.*, 37, 485
- [184] Stoica, P., & Moses, R. 2005, *Spectral Analysis of Signals* (Upper Saddle River, New Jersey: Pearson Prentice Hall)

- [185] Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *AJ*, 124, 1788, arXiv:astro-ph/0204349
- [186] van der Klis, M. 1989, in *NATO ASI Series, Vol. 262, Timing Neutron Stars*, ed. H. Oegelman & E. Heuvel (Springer Netherlands), 262, 27
- [187] Percival, D. B., & Walden, A. T. 1993, *Spectral Analysis for Physical Applications* (Cambridge: Cambridge University Press)
- [188] Lyons, R. G. 2004, *Understanding Digital Signal Processing* (2nd ed.; Prentice Hall Professional Technical Reference)
- [189] Welch, P. 1967, *IEEE Transactions on Audio and Electroacoustics*, 15, 70
- [190] Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, *ApJS*, 199, 31, arXiv:1108.1435
- [191] Romani, R. W. 2012, *ApJ*, 754, L25, arXiv:1207.1736
- [192] Xing, Y., & Wang, Z. 2014, *PASJ*, 66, 72, arXiv:1402.6401
- [193] Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, *ArXiv e-prints*, arXiv:1306.2144
- [194] Guillemot, L., Johnson, T. J., Venter, C., et al. 2012, *ApJ*, 744, 33, arXiv:1110.1271
- [195] Safi-Harb, S., Dubner, G., Petre, R., Holt, S. S., & Durouchoux, P. 2005, *ApJ*, 618, 321, arXiv:astro-ph/0407121
- [196] Jiang, B., Chen, Y., Wang, J., et al. 2010, *ApJ*, 712, 1147, arXiv:1001.2204
- [197] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 711, 64, arXiv:1001.0792
- [198] Stroh, M. C., & Falcone, A. D. 2013, *ApJS*, 207, 28, arXiv:1305.4949
- [199] Marelli, M., De Luca, A., & Caraveo, P. A. 2011, *ApJ*, 733, 82, arXiv:1103.0572
- [200] Green, D. A. 2014, *Bulletin of the Astronomical Society of India*, 42, 47, arXiv:1409.0637
- [201] Yamaguchi, H., Badenes, C., Foster, A. R., et al. 2015, *ApJ*, 801, L31, arXiv:1502.04255
- [202] Laffon, H., Smith, D. A., & Guillemot, L. for the Fermi-LAT Collaboration. 2015, *ArXiv e-prints*, arXiv:1502.03251
- [203] Ravi, V., Manchester, R. N., & Hobbs, G. 2010, *ApJ*, 716, L85, arXiv:1005.1966
- [204] Kerr, M. 2010, PhD thesis, University of Washington
- [205] Acero, F., Ackermann, M., Ajello, M., et al. 2016, *ApJS*, 223, 26, arXiv:1602.07246

- [206] Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, *Science*, 311, 1901, arXiv:astro-ph/0601337
- [207] Marshall, F. E., Gotthelf, E. V., Zhang, W., Middleditch, J., & Wang, Q. D. 1998, *ApJ*, 499, L179, arXiv:astro-ph/9803214
- [208] Dib, R., Kaspi, V. M., & Gavriil, F. P. 2009, *ApJ*, 702, 614, arXiv:0811.2659
- [209] Zhu, W., Kaspi, V. M., Dib, R., et al. 2008, *ApJ*, 686, 520
- [210] Lopez, L. A., Ramirez-Ruiz, E., Castro, D., & Pearson, S. 2013, *ApJ*, 764, 50, arXiv:1301.0618
- [211] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 722, 1303, arXiv:1008.4190
- [212] H. E. S. S. Collaboration, Abdalla, H., Abramowski, A., et al. 2016, *ArXiv e-prints*, arXiv:1609.00600
- [213] Zhu, H., Tian, W. W., & Zuo, P. 2014, *ApJ*, 793, 95, arXiv:1407.8260
- [214] Ho, W. C. G., & Heinke, C. O. 2009, *Nature*, 462, 71, arXiv:0911.0672
- [215] Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, 710, L92, arXiv:1001.1419
- [216] Abadie, J., Abbott, B. P., Abbott, R., et al. 2010, *ApJ*, 722, 1504, arXiv:1006.2535
- [217] Ajello, M., Albert, A., Atwood, W. B., et al. 2016, *ApJ*, 819, 44, arXiv:1511.02938
- [218] Bartels, R., Krishnamurthy, S., & Weniger, C. 2016, *Physical Review Letters*, 116, 051102, arXiv:1506.05104
- [219] Schwarz, G. 1978, *Ann. Stat.*, 6, 461
- [220] Goodman, J., & Weare, J. 2010, *Comm. App. Math. Comp. Sci.*, 5, 65
- [221] Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306, arXiv:1202.3665
- [222] Jetsu, L., & Pelt, J. 1996, *A&AS*, 118, 587
- [223] Perera, B. B. P., McLaughlin, M. A., Cordes, J. M., et al. 2013, *ApJ*, 776, 61, arXiv:1309.1982
- [224] Rubtsov, G. I., & Sokolova, E. V. 2015, *Soviet Journal of Experimental and Theoretical Physics Letters*, 100, 689, arXiv:1406.0608
- [225] Kruger, A. T., Lored, T. J., & Wasserman, I. 2002, *ApJ*, 576, 932, arXiv:astro-ph/0112192
- [226] Ostriker, J. P., & Gunn, J. E. 1969, *ApJ*, 157, 1395

- [227] Lyne, A. G., Pritchard, R. S., Graham-Smith, F., & Camilo, F. 1996, *Nature*, 381, 497
- [228] Livingstone, M. A., Kaspi, V. M., Gotthelf, E. V., & Kuiper, L. 2006, *ApJ*, 647, 1286, arXiv:astro-ph/0601530
- [229] Livingstone, M. A., Kaspi, V. M., Gavriil, F. P., et al. 2007, *Ap&SS*, 308, 317, arXiv:astro-ph/0702196
- [230] Espinoza, C. M., Lyne, A. G., Kramer, M., Manchester, R. N., & Kaspi, V. M. 2011, *ApJ*, 741, L13, arXiv:1109.2740
- [231] Weltevrede, P., Johnston, S., & Espinoza, C. M. 2011, *MNRAS*, 411, 1917, arXiv:1010.0857
- [232] Roy, J., Gupta, Y., & Lewandowski, W. 2012, *MNRAS*, 424, 2213, arXiv:1205.6264
- [233] Archibald, R. F., Gotthelf, E. V., Ferdman, R. D., et al. 2016, *ApJ*, 819, L16, arXiv:1603.00305
- [234] Camilo, F., Kaspi, V. M., Lyne, A. G., et al. 2000, *ApJ*, 541, 367, arXiv:astro-ph/0004330
- [235] Parent, D., Kerr, M., den Hartog, P. R., et al. 2011, *ApJ*, 743, 170, arXiv:1109.1590
- [236] Archibald, R. F., Kaspi, V. M., Tendulkar, S. P., & Scholz, P. 2016, *ApJ*, 829, L21, arXiv:1608.01007
- [237] Göğüş, E., Lin, L., Kaneko, Y., et al. 2016, *ApJ*, 829, L25, arXiv:1608.07133
- [238] Gavriil, F. P., Gonzalez, M. E., Gotthelf, E. V., et al. 2008, *Science*, 319, 1802, arXiv:0802.1704
- [239] Acero, F., Ackermann, M., Ajello, M., et al. 2016, *ApJS*, 224, 8, arXiv:1511.06778
- [240] Antonopoulou, D., Weltevrede, P., Espinoza, C. M., et al. 2015, *MNRAS*, 447, 3924, arXiv:1412.5853
- [241] Archibald, R. F., Kaspi, V. M., Beardmore, A. P., Gehrels, N., & Kennea, J. A. 2015, *ApJ*, 810, 67, arXiv:1506.06104
- [242] Kuiper, L., & Dekker, A. 2016, *The Astronomer's Telegram*, 9077
- [243] Baring, M. G., & Harding, A. K. 1998, *ApJ*, 507, L55, arXiv:astro-ph/9809115
- [244] Murphy, T., Mauch, T., Green, A., et al. 2007, *MNRAS*, 382, 382, arXiv:0708.3092
- [245] Hwang, U., & Markert, T. H. 1994, *ApJ*, 431, 819
- [246] Wakely, S. P., & Horan, D. 2008, *International Cosmic Ray Conference*, 3, 1341
- [247] Martín-Hernández, N. L., van der Hulst, J. M., & Tielens, A. G. G. M. 2003, *A&A*, 407, 957

- [248] Akbal, O., Gügercinoğlu, E., Şaşmaz Muş, S., & Alpar, M. A. 2015, *MNRAS*, 449, 933, arXiv:1502.03786
- [249] Blandford, R. D., & Romani, R. W. 1988, *MNRAS*, 234, 57P
- [250] Ho, W. C. G. 2015, *MNRAS*, 452, 845, arXiv:1506.03933
- [251] Lyne, A., Graham-Smith, F., Weltevrede, P., et al. 2013, *Science*, 342, 598, arXiv:1311.0408
- [252] Arzamasskiy, L., Philippov, A., & Tchekhovskoy, A. 2015, *MNRAS*, 453, 3540, arXiv:1504.06626
- [253] Michel, F. C., & Tucker, W. H. 1969, *Nature*, 223, 277
- [254] Harding, A. K., Contopoulos, I., & Kazanas, D. 1999, *ApJ*, 525, L125, arXiv:astro-ph/9908279
- [255] Menou, K., Perna, R., & Hernquist, L. 2001, *ApJ*, 554, L63, arXiv:astro-ph/0103326
- [256] Ekşi, K. Y., Andaç, I. C., Çıkıntoğlu, S., et al. 2016, *ApJ*, 823, 34, arXiv:1603.01487
- [257] Kerr, M., Hobbs, G., Johnston, S., & Shannon, R. M. 2016, *MNRAS*, 455, 1845, arXiv:1510.06078
- [258] Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *Phys. Rep.*, 203, 1
- [259] Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, *Science*, 324, 1411, arXiv:0905.3397
- [260] Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, *Nature*, 333, 237
- [261] van den Heuvel, E. P. J., & van Paradijs, J. 1988, *Nature*, 334, 227
- [262] Ray, P. S., Ransom, S. M., Cheung, C. C., et al. 2013, *ApJ*, 763, L13, arXiv:1210.6676
- [263] Romani, R. W., Filippenko, A. V., & Cenko, S. B. 2014, *ApJ*, 793, L20, arXiv:1408.2886
- [264] Shklovskii, I. S. 1970, *Soviet Astron.*, 13, 562
- [265] Ransom, S. M. 2001, PhD thesis, Harvard University
- [266] Cordes, J. M., & Lazio, T. J. W. 2002, *ArXiv Astrophysics e-prints*, arXiv:astro-ph/0207156
- [267] Dyks, J., & Rudak, B. 2003, *ApJ*, 598, 1201, arXiv:astro-ph/0303006
- [268] Harding, A. K., Usov, V. V., & Muslimov, A. G. 2005, *ApJ*, 622, 531, arXiv:astro-ph/0411805
- [269] Lande, J., Ackermann, M., Allafort, A., et al. 2012, *ApJ*, 756, 5, arXiv:1207.0027

- [270] Saz Parkinson, P. M., Xu, H., Yu, P. L. H., et al. 2016, *ApJ*, 820, 8, arXiv:1602.00385
- [271] Strader, J., Chomiuk, L., Sonbas, E., et al. 2014, *ApJ*, 788, L27, arXiv:1405.5533
- [272] Xing, Y., Wang, Z., & Ng, C.-Y. 2014, *ApJ*, 795, 88, arXiv:1407.1607
- [273] Romani, R. W. 2015, *ApJ*, 812, L24, arXiv:1509.07834
- [274] Salvetti, D., Mignani, R. P., De Luca, A., et al. 2015, *ApJ*, 814, 88, arXiv:1509.07474
- [275] Ackermann, M., Albert, A., Anderson, B., et al. 2015, *Physical Review Letters*, 115, 231301, arXiv:1503.02641
- [276] Charles, E., Sanchez-Conde, M., Anderson, B., et al. 2016, *ArXiv e-prints*, arXiv:1605.02016
- [277] Manchester, R. N. 2005, *Ap&SS*, 297, 101
- [278] Shannon, R. M., Lentati, L. T., Kerr, M., et al. 2016, *MNRAS*, 459, 3104, arXiv:1604.02661
- [279] Seward, F. D., Dame, T. M., Fesen, R. A., & Aschenbach, B. 1995, *ApJ*, 449, 681

Acknowledgments

Firstly, I would like to thank Holger Pletsch for introducing me to this exciting area of research. I could not have asked for a more kind or patient mentor to begin my academic career with. I wish you and your family the best of luck in your next adventure!

I would like to thank Bruce Allen, and all of the Observational Cosmology and Relativity division at the Albert Einstein Institute for welcoming me into the group. I am grateful to Carsten Aulbert, Henning Fehrmann and Alex Post for their help with ATLAS, and Oliver Bock, Heinz-Bernd Eggenstein and Bernd Machenschalk (and later Christian Beer) for working so hard to make our *Einstein@Home* survey a reality. To have not one but two supercomputers to work with was a dream come true!

I would like to thank all those with whom I have worked directly: Andrés, Amar, James, Lukas, Pia, Stanislav and especially Lars whose careful reading has prevented numerous mistakes on my part. It was also a pleasure to have shared an office with those people, and with Avneet, Giulio, Miriam, Miroslav and Nat.

I would also like to thank Reinhard Prix and Lars Nieder for translating the abstract into German.

I am grateful to everyone in the *Fermi*-LAT collaboration for their warm welcome, especially Patrizia Caraveo, Elizabeth Ferrara, Eric Grove, Lucas Guillemot, Alice Harding, Tyrel Johnson, Matthew Kerr, Julie McEnery, Paul Ray, Pablo Saz Parkinson, David Smith and David Thompson; and the Galactic Sources and Analysis Coordinators: Luca Baldini, Philippe Bruel, Elisabetta Cavazzuti, Lucas Guillemot, Jack Hewitt, Francesco de Palma, Jeremy Perkins and Max Razzano, for their hard work and helpful advice. I am extremely grateful to all those who have contributed to or helped in any way with the papers that make up my thesis, and to those who reviewed them within the collaboration.

Jason Wu deserves special thanks for patiently answering my many (many) questions and requests without complaint, and for doing an incredible job of picking pulsar targets!

I would like to thank the many friends who have made my time in Hannover so enjoyable, especially Alex, Kahae, Felix & Aerin; Andy; Avneet; Bridget & Dave; Collin & Kate; Evan; Germán & Karolina; Giulio; Grant; Josh; Karl; Melanie & Stefan; Lars; Miriam; Natalia; Pablo; Sebastian; Sylvia; Tito & Mariela; Tom A; Tom D; and Vaishali.

I am so grateful for the love and support of my family; their constant encouragement has made it feel like I've hardly left home.

Finally I would like to thank Heather, whose kindness and love still amaze me. I am truly grateful to have you by my side.

Curriculum Vitae

Colin James Clark

Date of Birth: 16th December 1990

Place of Birth: Paisley (United Kingdom)

Positions

2016–Present: **Postdoctoral Researcher**
Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Hannover

Education

2013–2016: **Dr. rer. nat.**
Leibniz Universität Hannover, Germany
Supervisor: Dr. Holger J. Pletsch
Research at MPI for Gravitational Physics (Albert Einstein Institute)

2009–2013: **M.Sci. Physics & Astronomy with Honours of the First Class**
University of Glasgow, UK

Awards

2017: *Max Planck Society*
Otto Hahn Medal

2010–2013: *University of Glasgow*
Ede & Ravenscroft Prize — Most distinguished Science graduate
Tannahill Bequest — Best performance in 5th year Astronomy examinations
Archibald McAulay Memorial Prize — Distinction in Physics Honours exams
Mackay Smith Prize — Most distinguished candidate in 4th year Physics exams
Lang Scholarship — Awarded on the recommendation of Heads of Science departments
Tannahill Bequest — For the best performance in 2nd and 3rd year Astronomy labs
Astronomy 2nd, 3rd and 4th year Class Prizes

Research

2013–Present: **Blind Searches for Gamma-ray Pulsars** (this work and postdoctoral research)
Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Hannover
Leibniz Universität Hannover

2012–2013: **Principal Component Analysis of Binary Black Hole Inspirals** (Master's project)
University of Glasgow

Summer 2012: **Laser Cooling of C³⁺ Ions in the Experimental Storage Ring**
Summer Student Program at GSI Helmholtzzentrum für Schwerionenforschung

Summer 2011: **Imaging Spectroscopy of Solar Flares with RHESSI**
University of Glasgow, funded by a Cormack Bequest from the Royal Society of Edinburgh

Publication List

- Clark, C. J., Wu, J., Pletsch, H. J., et al., 2017, *The Einstein@Home Gamma-ray Pulsar Survey. I. Search Methods, Sensitivity and Discovery of New Young Gamma-ray Pulsars*, The Astrophysical Journal, 834, 106.
- Clark, C. J., Pletsch, H. J., Wu, J., et al., 2016, *The Braking Index of a Radio-quiet Gamma-ray Pulsar*, The Astrophysical Journal Letters, 832, L15.
- Clark, C. J., Pletsch, H. J., Wu, J., et al., 2015, *PSR J1906+0722: An Elusive Gamma-Ray Pulsar*, The Astrophysical Journal Letters, 809, L2.
- Pletsch, H. J., & Clark, C. J. 2015, *Gamma-ray Timing of Redback PSR J2339–0533: Hints for Gravitational Quadrupole Moment Changes*, The Astrophysical Journal, 807, 18.
- Pletsch, H. J. & Clark, C. J. 2014, *Optimized Blind Gamma-ray Pulsar Searches at Fixed Computing Budget*, The Astrophysical Journal, 795, 75.