

BINARY GAMMA-RAY PULSARS

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

Doktor der Naturwissenschaften
Dr. rer. nat.

genehmigte Dissertation
von

M.Sc. Lars Nieder
geboren am 15.12.1989 in Hannover

2021

Referent:

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

Korreferenten:

Dr. David A. Smith
Laboratoire d'Astrophysique de Bordeaux, Université de Bordeaux
Centre d'Études Nucléaires de Bordeaux Gradignan, Université de Bordeaux

Prof. Dr. Domenico Giulini
Leibniz Universität Hannover
Center for Applied Space Technology and Microgravity (ZARM), Universität Bremen

Tag der Promotion:

04.12.2020

Abstract

Pulsar astronomy was born in 1967 with the first pulsar being discovered in radio waves. In the last decades more than 2,800 pulsars were found, most of them in radio. Only with the launch of the Large Area Telescope (LAT) onboard the *Fermi Gamma-ray Space Telescope* in 2008 did a large number of pulsars become detectable in gamma rays. To date, more than 250 gamma-ray pulsars are known, a quarter of them were first found in gamma rays, and some of them remained undetected in radio despite extensive follow-up searches. Computationally intensive gamma-ray searches with only the rough position of the targeted gamma-ray source known, also called partially informed searches, have been very successful in finding many new isolated pulsars and one binary pulsar.

This thesis concerns the development of sensitive and efficient methods to search for binary gamma-ray pulsars, as well as the methods' applications and the results of their execution.

The novel methods to search for binary pulsars were developed by building upon the search framework for isolated pulsars with the goal of scouring the pulsar parameter space with efficient test statistics using optimized search grids. To construct these grids, a metric was derived that allows one to build cost-efficient grids which ensure no signal is lost.

Within this thesis, the methods have been applied to two different kinds of pulsar searches. Firstly, two of the exotic “spider” pulsars have been discovered in partially informed searches using the volunteer computing project *Einstein@Home*. Secondly, for two recently found radio pulsars, the detection of their gamma-ray pulsations in follow-up searches enabled the precise measurement of their system parameters.

All four pulsars are remarkable and in some ways unusual compared to the rest of the known pulsar population. One pulsar is the fastest spinning pulsar for which its intrinsic spin-down rate is reliably constrained, revealing that it has one of the weakest inferred surface magnetic field strengths. Another pulsar is in a binary with a record-low 75 min orbital period, with the pulsar potentially having a mass $\gtrsim 2 M_{\odot}$. The next pulsar is the first “redback” millisecond pulsar ever to be discovered via its gamma-ray pulsations despite strong orbital-period variability, which may be linked to quadrupole changes in the companion star. The fourth pulsar shows variability in the spin frequency, which is either extraordinarily large, or indicates that the binary system is orbited by another planetary-mass companion.

Keywords: neutron stars, gamma rays, binary pulsars, search methods

Kurzfassung

Die Pulsar-Astronomie nahm 1967 mit der Entdeckung des ersten Pulsars in Radiowellen ihren Anfang. In den letzten Jahrzehnten wurden mehr als 2800 Pulsare gefunden, die meisten im Radiobereich. Erst mit dem Start des Large Area Telescope (LAT) an Bord des *Fermi Gamma-ray Space Telescope* im Jahr 2008 wurde eine große Zahl Pulsare im Gammabereich beobachtbar. Bislang sind über 250 Gammapulsare bekannt, ein Viertel davon wurde zuerst im Gammabereich entdeckt und einige von ihnen sind trotz ausgedehnter Suchen im Radiobereich nicht gefunden worden. Rechenintensive Gamma-Suchen, bei denen nur die ungefähre Position der anvisierten Gammaquelle bekannt ist, waren sehr erfolgreich darin, viele neue isolierte Pulsare und einen Binär-Pulsar zu finden.

Diese Doktorarbeit behandelt die Entwicklung sensitiver und effizienter Methoden, um nach Binär-Gammapulsaren zu suchen. Außerdem werden die Anwendung der Methoden und die Ergebnisse ihrer Umsetzung beschrieben.

Die neuartigen Suchmethoden für Binär-Pulsare wurden auf Basis der Suchstrukturen für isolierte Pulsare mit dem Ziel entwickelt, den Pulsar-Parameterraum mittels effizienter Teststatistiken und optimierter Suchgitter zu durchkämmen. Um solche Gitter zu konstruieren, wurde eine Metrik hergeleitet, die es erlaubt, kosteneffiziente Gitter zu designen, die sicherstellen, dass kein Signal verloren geht.

Innerhalb dieser Doktorarbeit wurden die Methoden bei zwei Arten von Pulsarsuchen angewendet. Zum einen wurden zwei der exotischen “Spinnenpulsare” in teilinformierten Suchen mit dem freiwilligen Rechenprojekt *Einstein@Home* entdeckt. Zum anderen ermöglichte die Folgeentdeckung von Gammapulsen zweier kürzlich gefundener Radiopulsare, dass deren Systemparameter präzise vermessen werden konnten.

Die vier Pulsare zeigen jeweils außergewöhnliche Eigenschaften im Vergleich zu der restlichen bekannten Pulsar-Population. Der eine ist der am schnellsten rotierende Pulsar, für den das intrinsische Abbremsen der Rotationsrate belastbar eingegrenzt ist. Dadurch konnte eine der schwächsten Oberflächenmagnetfeldstärken abgeleitet werden. Bei einem anderen Pulsar ist die Orbitalperiode mit 75 min die kürzeste. Die Masse des Pulsars ist möglicherweise $\gtrsim 2 M_{\odot}$. Der nächste Pulsar ist der erste “Redback” Millisekundenpulsar, der jemals trotz starker Variationen der Orbitalperiode durch seine Gammapulse entdeckt wurde. Diese Variationen könnten mit Quadrupoländerungen des Begleitsterns zusammenhängen. Der vierte Pulsar weist eine Variabilität der Rotationsrate auf, die entweder außergewöhnlich groß ist oder darauf hinweist, dass das System von einem weiteren Begleiter mit der Masse eines Planeten umkreist wird.

Schlagworte: Neutronensterne, Gamma-Strahlen, Binär-Pulsare, Suchmethoden

Contents

Abstract	3
Kurzfassung	5
Table of Contents	7
1 Introduction	11
1.1 Pulsars	11
1.2 Gamma-ray Pulsars	12
1.3 Spider Pulsars	14
1.3.1 Radio Searches Guided by Gamma-ray Sources	14
1.3.2 Partially Informed Gamma-ray Searches	15
1.4 Timing Analysis of Gamma-Ray Pulsars in Spider Systems	16
1.5 Clarification of Contributions to Publications	18
1.5.1 Chapter 2	18
1.5.2 Chapter 3	18
1.5.3 Chapter 4	19
1.5.4 Chapter 5	20
1.5.5 Chapter 6	20
1.5.6 Chapter 7	21
2 Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-ray Pulsars	23
2.1 Introduction	24
2.2 Partially-informed gamma-ray searches for pulsars	25
2.2.1 Pulse profile and photon arrival probability	26
2.2.2 Relationship of detector time t to t_{psr}	27
2.2.3 Searching for pulsations	28
2.2.4 Coherent power test statistic P	28
2.2.5 Semicoherent power test statistic S	33
2.2.6 Multiple harmonic test statistic H	36
2.2.7 Searches for isolated pulsars	36
2.3 Search Method: Circular Binary Orbits	37

2.3.1	Parameter space metrics	38
2.3.2	Search design for circular binary	40
2.4	Search Method: Eccentric Binary Orbits	45
2.4.1	Parameter space metrics	47
2.4.2	Search design for low-eccentricity binary	49
2.5	Comparison with other Methods	53
2.5.1	Acceleration search	53
2.5.2	Stack/slide search	54
2.5.3	Power spectrum search	55
2.5.4	Sideband search	55
2.5.5	Discussion	60
2.6	Conclusions	60
2.7	Appendix to Chapter 2	63
2.7.1	Expectation values of signal statistics	63
2.7.2	Maximal Sensitivity at Fixed Computing Cost	64
2.7.3	High-order phase model for elliptical binaries	69
3	Detection and Timing of Gamma-ray Pulsations	
	from the 707 Hz Pulsar J0952–0607	75
3.1	Introduction	76
3.2	Gamma-ray Pulsation Discovery	78
3.2.1	Data Preparation	78
3.2.2	Search	80
3.2.3	Detection	82
3.3	Gamma-ray Timing	83
3.3.1	Methods	83
3.3.2	Solution	84
3.4	Multi wavelength	88
3.4.1	Updated Radio Timing	88
3.4.2	Optical Photometry	88
3.4.3	Optical Light-Curve Modeling	90
3.4.4	Search for Continuous Gravitational Waves	91
3.5	Discussion	93
3.6	Conclusions	100
3.7	Appendix to Chapter 3	102
3.7.1	Estimating the false-alarm probability for a multi-dimensional H statistic search	102
4	Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated	
	Einstein@Home	105
4.1	Introduction	106
4.2	Gamma-ray pulsations	106
4.2.1	Data preparation	106
4.2.2	Search	107
4.2.3	Timing	110
4.3	Multiwavelength & Multimessenger	110
4.3.1	Optical Light Curve Modeling and System Masses	110
4.3.2	Radio pulsation searches	114
4.3.3	Continuous gravitational waves	114

4.4	Discussion & Conclusions	116
5	Einstein@Home Discovery of the Gamma-ray Millisecond Pulsar PSR J2039–5617 Confirms its Predicted Redback Nature	121
5.1	Introduction	122
5.2	Summary of previous literature	124
5.3	Gamma-ray Observations	125
5.3.1	Gamma-ray Pulsation Search	126
5.3.2	Gamma-ray Timing	128
5.3.3	Gamma-ray Variability	130
5.4	Optical observations and Modelling	137
5.4.1	New optical observations	137
5.4.2	Light curve modelling	138
5.5	Results and Discussion	142
5.5.1	Binary Inclination and Component Masses	142
5.5.2	Distance and Energetics	148
5.5.3	Optical light curve asymmetry and variability	149
5.5.4	Orbital Period Variability	151
5.5.5	Prospects for binary gamma-ray pulsar searches	154
5.6	Conclusions	155
6	Gamma-ray Follow-up Detection and Timing of PSR J1555–2908	159
6.1	Introduction	159
6.2	Gamma-ray Observations	159
6.2.1	<i>Fermi</i> -Large Area Telescope (LAT) Data Preparation	160
6.2.2	Gamma-ray Pulsation Search and Detection	160
6.2.3	Gamma-ray Pulsation Timing	161
6.3	Discussion	163
7	Does the Black Widow Pulsar PSR J1555–2908 have an additional Planetary-mass Companion?	165
7.1	Introduction	166
7.2	Rotational phase model	166
7.3	Gamma-ray timing analysis	167
7.4	Discussion	168
8	Conclusion and Outlook	171
	Bibliography	175
	Acknowledgments	199
	Curriculum Vitae	201
	Publication List	203

CHAPTER 1

Introduction

The goals of this thesis are to develop methods to enable partially informed searches for gamma-ray pulsars in binary systems, to perform such searches within gamma-ray data taken with the *Fermi Gamma-ray Space Telescope* using powerful computing clusters like the volunteer computing project *Einstein@Home*, and hopefully to discover such pulsars using those methods.

In this chapter, the background to this thesis is briefly presented. Pulsars and some of their properties are described in Section 1.1. Gamma-ray pulsars are introduced in Section 1.2. Routes to discover the exotic “spider pulsars” are outlined in Section 1.3. Section 1.4 explains how the properties of these spider pulsars can be measured using gamma-ray data.

The following six chapters are reproductions of papers that are published in, accepted for publication in, or soon to be submitted to scientific journals. Chapter 2 is published as Nieder et al. [1]. Chapter 3 is published as Nieder et al. [2]. Chapter 4 is accepted for publication in ApJ and on the arXiv as Nieder et al. [3]. Chapter 5 is accepted for publication in MNRAS and on the arXiv as Clark et al. [4]. Chapter 6 is the author’s contribution to a paper, P. S. Ray et al., that will be submitted to ApJ soon. Chapter 7 will be submitted to ApJL as a companion paper, L. Nieder et al., to the aforementioned paper. These chapters are nearly identical to the original papers with only minor changes to correct typos or for formatting reasons. To all of these papers the author of this thesis made substantial contributions which are described in more detail in Section 1.5.

Finally, in Chapter 8, a brief summary of the thesis results is presented, and an outlook into current as well as future projects is attempted.

1.1 Pulsars

A neutron star is born when the core of a massive star collapses under its own gravity at the end of its stellar life, and the outer shells are expelled. This supernova birth was already proposed in 1934 by Baade and Zwicky [5].

More than three decades later, in 1967, a pulsating source of radio was discovered by Jocelyn Bell-Burnell and Antony Hewish, which was speculated to be associated

either with a white dwarf or a neutron star [6]. Within a couple of months three more pulsating sources were found, but the source of the radio emission remained unclear [7].

Just before the publication of the pulsating radio source, Pacini [8] independently suggested that a neutron star with a strong magnetic dipole field not aligned with the rotation axis might dissipate energy via particle acceleration and subsequently electromagnetic radiation. In 1968, Gold [9] proposed a very similar picture clearly drawing the connection between neutron stars and pulsars, establishing the idea of pulsars being cosmic beacons, and even correctly predicting the spin-down process due to the energy loss via electromagnetic radiation.

It is observed that most pulsars evolve with time towards longer spin periods. Their loss in kinetic energy is typically called spin-down luminosity, and is thought to mostly power the beamed pulsar emission (“rotation-powered”), especially at high energies [10, 11].

In 1982, the first millisecond pulsar (MSP) PSR B1937+21, a pulsar rotating 642 times per second, was discovered [12]. Even though over a hundred other MSPs were found in the next decades, PSR B1937+21 kept the record for the fastest spinning pulsar until, in 2006, the 716 Hz pulsar PSR J1748–2446ad, located in the globular cluster Terzan 5, was detected [13]. Since then, only two more pulsars have surpassed the rapid rotation rate of the original MSP [14, 15].

The existence of millisecond pulsars is typically explained by old pulsars being recycled or spun up by accreting matter from a companion star [16, 17]. In fact, some neutron stars have been observed to spin up during accretion from a companion star, causing pulsed X-ray emission (“accretion-powered”) [11]. This is explained by the accreted matter following the magnetic field lines creating hot spots on the neutron star surface [18]. The pulsar-recycling theory was further backed up when “transitional MSPs” were discovered, which have been observed to switch back and forth between rotation-powered MSP state and accretion-powered low-mass X-ray binary (LMXB) state [19–22]. These pulsars all belong to the “redback” class (see Section 1.3).

1.2 Gamma-ray Pulsars

Gamma-ray pulsar astronomy has evolved dramatically over the last decades [23]. It started slowly in 1974 with the identification of gamma-ray pulsations from the Crab pulsar [24] taken with the gamma-ray telescope onboard the Small Astronomy Satellite 2 (SAS-2) [25]. Following up the X-ray discovery of Geminga [26] with data from the COS-B satellite [27], the first radio-quiet¹ gamma-ray pulsar was discovered [29]. The following gamma-ray telescopes EGRET [30] and later AGILE [31] were more sensitive than the ones from the previous generation, but the launch of the *Fermi Gamma-ray Space Telescope* [32] has taken gamma-ray pulsar astronomy to the next level.

The *Fermi Gamma-ray Space Telescope* carries two main instruments: the LAT [32] and the Gamma-ray Burst Monitor (GBM) [33]. The LAT is a pair-conversion

¹Here, “radio-quiet” means that the pulsar is not discovered in radio. Abdo et al. [28] proposed a pulsed-radio-emission threshold of $30 \mu\text{Jy}$ at a radio observing frequency of 1.4 GHz below which a pulsar may be called radio-quiet.




Figure 1.1: Schematic cutaway depiction of the Large Area Telescope, showing a simplistic representation of the operating principle – an incoming gamma-ray photon hits one of the high- Z foils, undergoes pair conversion into an electron and a positron, whose paths are observed by the precision tracker to reconstruct the direction from which the photon came, and whose energies are measured in the calorimeter to infer the photon’s energy. Image credit: Atwood et al. [32].

gamma-ray telescope with the basic composition shown in Figure 1.1. It consists of precision trackers, calorimeters, and an anticoincidence detector. A tracker module has a stack of high- Z planes (tungsten) interleaved with position-sensitive detectors. The high- Z planes convert incoming gamma rays into electron-positron pairs, and the position-sensitive detectors track the particles to infer the direction of the incident gamma-ray photon. The calorimeter modules deduce the gamma-ray photon’s energy from the electromagnetic particle shower produced by the electron-positron pair. The shower development profile, also measured by the calorimeter, and the anticoincidence detector are used to suppress background, e.g. from cosmic rays.

The LAT is a remarkable gamma-ray telescope. With its wide field-of-view and its all-sky survey mode, the exposure is almost uniform after two orbits of the *Fermi* satellite around Earth (~ 3 hours) [32]. The detector has a low deadtime and the photon timing accuracy is $< 10 \mu\text{s}$. The energy range spans 20 MeV – 300 GeV.

Since its first bright gamma-ray source list (also called 0FGL catalog) [34], the *Fermi*-LAT collaboration released four more catalogs of gamma-ray sources [35–38]. The newest of these is the *Fermi*-LAT Fourth Source Catalog (4FGL catalog) [38] which lists 5,064 sources with 1,336 of them being unassociated².

Along with the number of detected gamma-ray sources, the number of known gamma-ray pulsars increased rapidly as well. When the *Fermi* satellite was launched in 2008, fewer than ten gamma-ray pulsars were known [23]. The first *Fermi*-LAT catalog of gamma-ray pulsars already listed 46 pulsars [39] only six months after the launch. After three years of operation, the second *Fermi*-LAT catalog of gamma-ray pulsars contained 117 gamma-ray pulsars. In the year 2020, more than 250

²Here, “unassociated” means that the source has no plausible counterpart belonging to a known gamma-ray emitting source class [37].

gamma-ray pulsars are known thanks to the excellent sensitivity of the LAT³.

The large number of gamma-ray pulsars also enables population studies, and several discoveries of extraordinary pulsars extended our knowledge about pulsars in general. For example, there seems to be a minimum spin-down power below which a pulsar’s gamma radiation turns off [40–42]. The first radio-quiet gamma-ray millisecond pulsar has been discovered by the LAT [43]. The binary pulsars with the shortest [44] and the longest orbital period [45] have first been detected in LAT data.

The high-energy emission mechanism of pulsars is not yet fully understood. The first models assumed that pulsars have a simple dipole magnetic field and that the high-energy emission originates from close to the magnetic poles at the pulsar surface [9]. Fully taking into account the complexity of a pulsar magnetosphere, the region where this emission is thought to be produced moved further outward with time, with the latest models even suggesting that the gamma-ray emission originates outside the light cylinder [46–48].

1.3 Spider Pulsars

In 1988, Fruchter, Stinebring, and Taylor [49] discovered a millisecond pulsar in a compact binary system, PSR B1957+20. The radio pulsations, observed at 430 MHz with the Arecibo Observatory, were eclipsed for nearly 10% of the 9.2 hr orbit, which is the result of plasma being ablated from the companion star by the energetic pulsar wind [50]. This pulsar is now categorized as a “black widow” pulsar. A black widow consists of a millisecond pulsar in a close binary (orbital period $\lesssim 1$ day) with a low-mass companion star ($M_c \ll 0.1 M_\odot$) that is being evaporated by the pulsar. The nickname was chosen in analogy to a type of spider where the female kills the male after mating.

With several “redback” pulsars being detected, another subcategory of “spider” pulsars was discovered two decades later [51, 52]. Similar to black widow pulsars redback pulsars are in close binary systems with the companion being irradiated and evaporated by the pulsar wind. However, the companions have higher masses ($M_c \sim 0.15 - 0.7 M_\odot$) [53, 54]. One redback pulsar candidate might even have a companion with mass $M_c \gtrsim 0.8 M_\odot$ [55].

1.3.1 Radio Searches Guided by Gamma-ray Sources

Radio searches targeting the sky positions of *Fermi*-LAT unassociated sources have been very successful in identifying new MSPs and spider pulsars, discovering more than 40 MSPs, 14 of them being spider pulsars, within the first three years of the *Fermi* mission [56].

Most large radio telescopes continued using the strategy of targeting gamma-ray sources and found many new pulsars. Such radio surveys are undertaken by the Arecibo telescope [57], the Effelsberg telescope [58], the Five-hundred-meter Aperture Spherical radio Telescope (FAST) [59], the Giant Metrewave Radio Telescope (GMRT) [60], the Green Bank Telescope (GBT) [61, 62], the Low-Frequency Array (LOFAR) telescope [14, 63], the Nançay telescope [64], and the Parkes telescope

³<https://tinyurl.com/fermipulsars>

[65].

For the radio surveys, there are different strategies to select the targets among the unassociated gamma-ray sources that will be searched. The most promising pulsar candidates typically show curved spectra well described by an exponentially cutoff power law and little variability in brightness over time [66, 67].

One example of a spider pulsar being discovered in a search targeting a *Fermi*-LAT source is the GBT discovery of the redback pulsar PSR J2339–0533 [68]. In this case, it was known that this source likely harbors a spider pulsar due to analysis of optical, X-ray, and gamma-ray data [69, 70]. This pulsar was difficult to detect in the radio even in follow-up observations, due to ablated ionized gas surrounding the system. However, a gamma-ray follow-up search, guided by the parameter constraints from the radio discovery, revealed gamma-ray pulsations [71].

Two radio surveys targeting gamma-ray source sky locations with LOFAR and GBT led to the discoveries of PSR J0952–0607 [14] and PSR J1555–2908. Both pulsars were discovered subsequently in gamma rays by exploiting the preliminary radio ephemerides and are the subjects of Chapters 3 and 6 of this thesis, respectively. They are also shown in red on the gamma-ray sky map in Figure 1.2.

1.3.2 Partially Informed Gamma-ray Searches

Partially informed gamma-ray searches⁴ have been extremely successful. In the first ten years of the *Fermi* mission, 60 pulsars were discovered by directly searching the LAT data, with most of them remaining undetected in subsequent radio searches [43, 44, 72–79].

More than half of these pulsars were discovered with search methods [75, 80] similar to those originally developed to find continuous gravitational waves [81–83]. The first ten pulsars were detected [44, 75] using the ATLAS cluster [84]. To enable more sensitive, and thus even more computing intensive searches, the search codes were ported to run on the distributed volunteer computing system *Einstein@Home* [85]. Searches with *Einstein@Home* resulted in 24 more discoveries [43, 76–79].

Einstein@Home, started in 2005 is a citizen science project⁵, and a powerful supercomputer designed to search for weak astrophysical periodic signals. The computing work for a search is divided into smaller work units, which would take a few hours to analyze on an average computer. These work units are then distributed to volunteers’ computers using the Berkeley Open Infrastructure for Network Computing (BOINC) [86] and are searched when the computer is otherwise idle. More than one million people have participated already, with $\sim 20,000$ active users. The total floating point speed of *Einstein@Home* is > 5.5 PFLOP/s.

The only binary or spider pulsar discovered in a partially informed gamma-ray search, previous to this thesis, is the black widow pulsar PSR J1311–3430 [44]. The pulsations were detected by the ATLAS computing cluster [84] using a preliminary version of the search methods described in Chapter 2, exploiting previously published orbital constraints from optical observations [87]. Optical follow-up analysis suggests a pulsar mass $\gtrsim 2 M_{\odot}$.

Using preliminary optical constraints on the orbital parameters, partially informed gamma-ray searches with *Einstein@Home* (Chapter 2) led to the discovery

⁴In previous studies, these searches have been named “blind” searches.

⁵<https://einsteinathome.org>

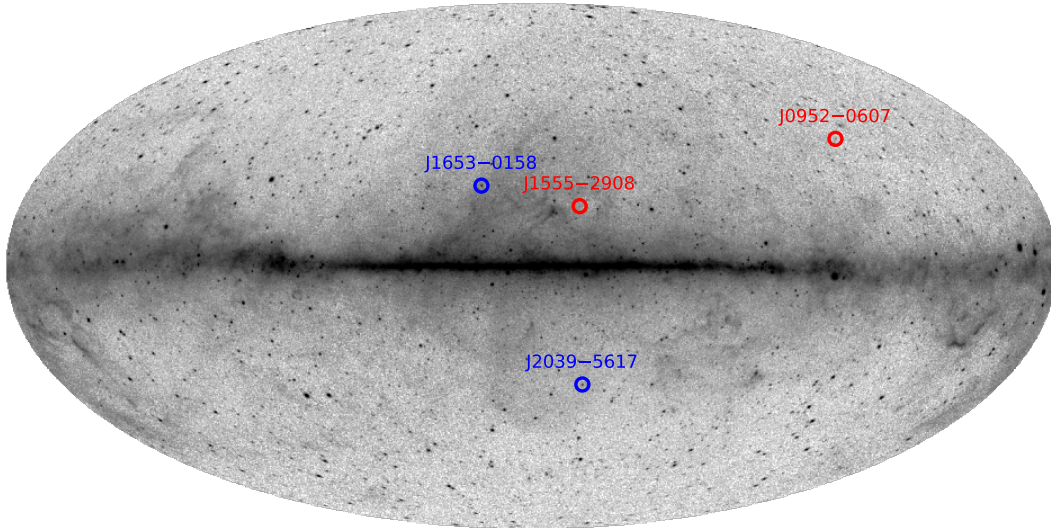


Figure 1.2: The sky locations of the four spider gamma-ray pulsars discovered as part of this thesis, shown on top of a *Fermi*-LAT all-sky image using eleven years of gamma-ray data. The red circles denote the pulsars discovered in follow-up searches of radio pulsars, and the blue circles depict the pulsars discovered in partially informed gamma-ray searches. The greyscale indicates the gamma-ray intensity above 1 GeV in logarithmic units.

of two new spider pulsars. The first one to be discovered was the black widow pulsar PSR J1653–0158 (Chapter 4). The second one was the redback pulsar PSR J2039–5617 (Chapter 5). Both pulsars are shown in blue on the gamma-ray sky map in Figure 1.2.

Several spider pulsar candidates have been identified over the last years as promising targets [88–94] for radio searches (Section 1.3.1) or partially informed gamma-ray searches. However, for the latter searches, tight constraints on the orbital parameters will be required.

1.4 Timing Analysis of Gamma-Ray Pulsars in Spider Systems

The timing analysis is the procedure, coming after a pulsar detection, in which properties of the pulsar, e.g. spin frequency or sky location, are precisely measured. Usually, this requires so-called pulse times of arrival (TOAs) and a rotation phase model. Since only for the brightest radio pulsars single pulses are visible, these are usually obtained by “folding” data taken within a certain period of time. Here, “folding” means cutting the data into segments of the pulsar’s spin-period length, and adding enough of them up to have a pulse distinguishable from the noise [10]. In the case of radio pulsars this can mean folding seconds to hours of data. The phase model assigns a pulse time of arrival with the pulsar’s rotational phase at the time of emission. The goal of the timing analysis is to find the phase model and the associated pulsar parameters which line up all the pulses at the same phase. This analysis can for example be done with one of the most widely used pulsar timing algorithms, TEMPO2 [95, 96].

With gamma-ray data, pulsars may be timed as well. The TOA timing method, known from radio, has also been used successfully on LAT data [52, 97]. Building upon these methods, Clark et al. [79] developed an entirely unbinned timing procedure, i.e. it does not require TOAs, which can be beneficial for faint gamma-ray pulsars. The unbinned timing method has also been used throughout this thesis.

Several spider pulsars, black widows and redbacks, have been observed to show variable orbital periods [19, 71, 98–100], which are proposed to originate from the companion star’s gravitational quadrupole moment changing with its magnetic activity cycles [101, 102]. Recently, a timing model including this has been developed [103], and has been used to measure the variable quadrupole component for the companion of the black widow pulsar PSR J2051–0827 [104].

A timing campaign with radio data of spider pulsars, and particularly of redback pulsars, has two major difficulties, which are both displayed in the timing analysis of the redback pulsar PSR J1048+2339 [100]. Firstly, tracking the orbital period variations seen in many spider systems requires observations of the pulsar frequently without larger gaps in time, while high-cadence observations radio timing campaigns require large amounts of valuable time on large radio telescopes. Secondly, the radio pulsations are often blocked by eclipses of varying degrees around the superior conjunction, sometimes blocking the radio waves for half of the orbit.

Timing spider pulsars with gamma-ray data is in some ways complementary to the radio case. The sparsity of the LAT data requires long integration times for any detection. Typical pulsars rotate millions of times between two gamma-ray photons being detected by the LAT, and might even complete several orbits in their binary system. However, the gamma-ray photons are unaffected by the plasma surrounding the binary system, and the ongoing all-sky survey of the *Fermi*-LAT provides gamma-ray data for any location in the sky since its launch in 2008 [32]. The latter is especially useful because a pulsar, discovered today, may immediately be timed over more than 12 years of data if bright enough in gamma rays.

As described in Section 1.3.1, the redback pulsar PSR J2339–0533 is difficult to see in radio [68]. A radio timing survey of a redback pulsar whose orbital period varies over time would require a sensitive radio telescope observing the pulsar very frequently. The gamma-ray data, unaffected by the ablated plasma within the system, has been successfully used to precisely track the pulsar’s rotation, and to measure the variations of the orbital period [71].

Another example for a redback pulsar being timed using gamma-ray data is PSR J2039–5617, which is described in Chapter 5. In radio the pulsar is not detected in every observation, and in the cases in which it is detected it is eclipsed for about half of the orbit [105]. In gamma-ray data, it was possible to detect pulsations, and measure the orbital period variations over the full *Fermi* mission since 2008.

The long-term timing analysis of black widow pulsar PSR J1555–2908 in *Fermi*-LAT data revealed unexpected changes of the spin frequency over time (Chapter 6), which could also be modelled by the binary system being orbited by a planetary-mass object in a wide, multi-year orbit (Chapter 7). The gamma-ray data taken over the next years of the ongoing *Fermi* mission will reveal whether PSR J1555–2908 is intrinsically an unusual MSP or part of a triple system.

1.5 Clarification of Contributions to Publications

In the following, the paper adaptations presented as chapters of this thesis will be briefly introduced. Furthermore, it is clarified who wrote which parts of the papers and who provided which scientific contributions, emphasizing the work of the author of this thesis.

1.5.1 Chapter 2

Chapter 2 describes methods to efficiently search for binary pulsars within gamma-ray data, exploiting partial orbital information. The chapter was published as Nieder et al. [1] and can be seen as an extension of the work by Pletsch and Clark [80], who developed the search methods for isolated pulsars.

This chapter was written by the author, under the close guidance of the supervisor B. Allen. Parts of this chapter have been derived previously by the author as part of a master thesis with supervisor H. J. Pletsch. For the paper, the methodology was significantly extended, a discussion of search method alternatives was added, and it has been investigated which phase models should be used depending on the expected specifics of the binary system. C. J. Clark suggested multiple additions, modifications and corrections, as well as derived the equations in Section 2.5.2. Comments from B. Allen, C. J. Clark, and H. J. Pletsch were incorporated.

As part of this work, two search pipelines were constructed by the author. One pipeline is set up for partially informed gamma-ray searches exploiting orbital constraints from optical analysis, where the first stages are run on the distributed volunteer computing system *Einstein@Home*, and the final stage on the ATLAS computing cluster. The search and timing codes initially designed for isolated-pulsar searches also had preliminary capabilities for binary pulsars – this was rewritten and extended by the author. This was necessary since some of the assumptions made in the previous versions of the search and timing codes do not hold for the heavier redback pulsars in their wider orbits. The porting of the search code from CPUs to GPUs was mainly done by B. Machenschalk, with help from C. Choquet. The author performed the tests on the newly written code. The other pipeline is designed for rapid follow-up searches with ATLAS, using preliminary information about the pulsar parameters derived after the discovery of radio pulsations.

1.5.2 Chapter 3

Chapter 3 presents the detection of gamma-ray pulsations from the black widow pulsar PSR J0952–0607, first discovered with the LOFAR radio telescope. The pulsar is spinning at the remarkable rate of 707 Hz, and its inferred surface magnetic field strength is among the ten lowest of all pulsars. This chapter describes and discusses the results of a multiwavelength study of this pulsar, and was originally published as Nieder et al. [2].

Most parts of this chapter were written by the author. The preparation of the LAT data set for the search was performed by the author. The preparation for the timing analysis and the final gamma-ray spectral analysis were carried out by J. Wu. Using the initial parameter constraints from the LOFAR discovery, the author was responsible for the design and execution of the search, as well as the timing analysis after detection. The large distance to the pulsar and the consequential weak energy

flux required a different treatment of the pulse-template parameters in the timing analysis. To take the pulse-shape uncertainty into account, the timing code was extended by C. J. Clark and the author. The radio analysis and text were provided by C. G. Bassa, J. Y. Donner, and J. P. W. Verbiest. C. J. Clark and C. Bassa contributed the optical analysis and texts. The gravitational wave search was led and the text was written by A. Singh, M. A. Papa, and their research group. Most of the discussion was written by the author. The text snippets concerning the optical and gravitational-wave search results were contributed by the respective groups mentioned above.

The chapter was reviewed within the *Fermi*-LAT collaboration by P. S. Ray. Further comments were provided by M. Kerr, D. J. Thompson, most co-authors, and the anonymous referee.

1.5.3 Chapter 4

Chapter 4 describes *Einstein@Home*'s first binary pulsar discovery. The pulsar PSR J1653–0158 was detected using the search methods described in Chapter 2, and is remarkable in multiple ways: its inferred surface magnetic field is possibly the weakest known for any pulsar, the optical analysis suggests a pulsar mass $\gtrsim 2 M_{\odot}$, and, with 75 min, its orbital period is the shortest known for any rotation-powered binary pulsar. Additionally, the pulsar remained undetected in radio waves despite several very sensitive observations. The discovery was published as Nieder et al. [3].

The text of this chapter was mainly written by the author, under the guidance of B. Allen. The gamma-ray data were collected and the photons weighted by C. J. Clark. Most of the new optical data were taken by R. P. Breton and V. S. Dhillon. The optical data from August 2017 were taken by C. Bassa and the author with the Wide Field Camera (WFC) on the 2.5m Isaac Newton Telescope (INT). The analysis and subsequently the constraints on the orbital parameters leading to this discovery were provided by C. J. Clark. The search was planned, started, and the results were analyzed by the author. The timing analysis was performed by the author. The estimation of the intrinsic spin-down parameter was computed by the author, taking into account valuable comments by M. Kramer. The optical modelling of the companion light curve after the pulsar discovery was performed and the associated text provided by D. Kandel and R. W. Romani. Many radio observations with large telescopes around the world (Effelsberg, FAST, GBT, GMRT, LOFAR, Lovell, Nançay, Parkes) tried to discover the pulsar in radio as part of the Pulsar Search Consortium (PSC). The associated groups provided the data for Table 4.3, which were used by the author to compute the upper limits on the pulsed radio emission. Text and interpretation of the results were written by the author. The continuous gravitational wave searches were developed, executed, and the texts provided by A. Ashok, M. A. Papa, B. Allen, and their research groups. The discussion of the results presented in this chapter was contributed by the author.

Within the *Fermi*-LAT collaboration, the chapter was reviewed by M. Razzano. Additional comments were given by N. Omodei, J. L. Racusin, and many co-authors.

1.5.4 Chapter 5

Chapter 5 describes the *Einstein@Home* discovery of PSR J2039–5617, the first reback pulsar to be detected in a partially informed gamma-ray search. The ability to find pulsars in reback systems is noteworthy because they typically show a significant variability of the orbital period. For PSR J2039–5617, these variations reduced the detectable signal power by 66% in the *Einstein@Home* search. In this chapter, the available optical and gamma-ray data are used to measure the system parameters, to track the behavior of the binary system, and to understand its evolution.

Analysis of archival data from the Parkes Radio Telescope led to the discovery of strongly eclipsed radio pulsations. The study of PSR J2039–5617 in radio waves is described in a companion paper [105].

This chapter was mostly written by C. J. Clark. The optical constraints on the orbital parameters, and the photon weights were also provided by C. J. Clark. The author planned and executed the *Einstein@Home* gamma-ray search, and contributed the text describing the search and discovery (Section 5.3.1). The initial timing analysis was performed by the author. This used a newly developed method to find a Taylor series approximation to track the orbital period variations – the ansatz used previously for reback pulsars. However, the strong variations (see Fig. 5.1) could not be well described with a Taylor series expansion anymore, which led to the development of another novel timing method by C. J. Clark (see Section 5.3.3). The optical analysis and discussion were mostly written by C. J. Clark, with G. Voisin being strongly involved in the work displayed in Section 5.5.4. The text and analysis describing the prospects for future reback searches (Section 5.5.5) were contributed by the author.

The review within the *Fermi*-LAT collaboration for this chapter was done by D. A. Smith. More comments were provided by D. J. Thompson, S. Digel, M. Kerr, many co-authors, and the anonymous referee.

1.5.5 Chapter 6

Chapter 6 presents the gamma-ray follow-up discovery of the millisecond black-widow pulsar PSR J1555–2908, which was first discovered in radio data taken with the Green Bank Telescope (GBT). PSR J1555–2908 experiences changes of the spin frequency over the course of the *Fermi* mission, unusually large for a millisecond pulsar. These variations require four frequency derivatives to track the pulsar’s rotational phase.

This chapter is a preliminary report describing the gamma-ray discovery and timing analysis, and is intended to be included in a paper led by P. Ray – currently in early draft stage. The paper will also present the radio discovery, and the results of some optical and X-ray observations.

Most of this chapter was written by the author. The text describing the selection and weighting of the gamma-ray photons, as well as the weighted photons used in the follow-up search, were contributed by C. J. Clark. The final photon weights were provided by P. Bruel. The search was planned, executed and described by the author, as well as all further analysis and discussion of the results. Comments from B. Allen and C. J. Clark were incorporated.

1.5.6 Chapter 7

Chapter 7 presents a hypothetical explanation for the spin-frequency variations of PSR J1555–2908, described in Chapter 6. With currently available LAT data, the pulsar’s rotation can be equally accurately described with an additional low-mass companion being in a multi-year around the inner black-widow system.

This chapter was written by the author, and is a preliminary version of a paper intended to be submitted as a companion paper to the one in Chapter 6. The gamma-ray photon data are the same as in Chapter 6. Comments were provided by B. Allen and C. J. Clark.

Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-ray Pulsars

*Published as Nieder, L., Allen, B., Clark, C. J., Pletsch, H. J. 2020, ApJ, 901, 2.
DOI: 10.3847/1538-4357/abaf53
© 2020. The American Astronomical Society.*

Abstract

It is difficult to discover pulsars via their gamma-ray emission because current instruments typically detect fewer than one photon per million rotations. This creates a significant computing challenge for isolated pulsars, where the typical parameter search space spans wide ranges in four dimensions. It is even more demanding when the pulsar is in a binary system, where the orbital motion introduces several additional unknown parameters. Building on earlier work by Pletsch and Clark [80], we present optimal methods for such searches. These can also incorporate external constraints on the parameter space to be searched, for example, from optical observations of a presumed binary companion. The solution has two parts. The first is the construction of optimal search grids in parameter space via a parameter-space metric, for initial semicoherent searches and subsequent fully coherent follow-ups. The second is a method to demodulate and detect the periodic pulsations. These methods have different sensitivity properties than traditional radio searches for binary pulsars and might unveil new populations of pulsars.

2.1 Introduction

The Large Area Telescope [LAT; 32] on the *Fermi* satellite has helped to increase the known Galactic population of gamma-ray pulsars to more than 250 pulsars¹ [for a review see, e.g., 23]. However, in the recent *Fermi*-LAT Fourth Source Catalog [4FGL; 38] 1,336 out of 5,064 gamma-ray sources remain unassociated. Many of those are thought to be pulsars, perhaps in binary systems.

Gamma-ray pulsars may be detected in three ways: (a) A known (radio or X-ray) pulsar position and ephemeris guides a follow-up gamma-ray pulsation search within a nearby LAT source [e.g., 106–108]. (b) A similar gamma-ray pulsation search is done for a known pulsar, but *without* an obvious gamma-ray source being present [109]. (c) A “partially informed” search² hunts for gamma-ray pulsations around a LAT source where no pulsar has yet been identified, and hence several timing parameters, notably the spin period, are unknown in advance.

Partially informed searches are the focus of this paper. Such searches have discovered more than 50 young pulsars (YPs) [e.g., 73–75, 79], and three MSPs [43, 44]. Many of these pulsars could not have been found via radio or X-ray emissions, which were not detected in extensive follow-up searches. Such systems are of particular interest because they constrain models of pulsar emission and beaming. Partially informed searches also have the potential to discover new populations of pulsar/neutron star objects.

So far, most partially informed gamma-ray searches have targeted isolated pulsars. The searches are a substantial computing effort, and have been carried out in campaigns or surveys that last several years. More recent surveys find new systems because the ongoing LAT operations provide additional data, which enables the detection of weaker pulsations [e.g., 79]. However, there is also a downside: the computing power required also increases quickly with longer observation time spans.

Until now, partially informed gamma-ray searches have only found one binary MSP, PSR J1311–3430 [44]. This is tantalizing because three quarters of the known MSPs in the Australia Telescope National Facility (ATNF) Pulsar Catalogue³ [110] are in binaries. So if search sensitivity were not limited by computing power, it might be possible to find many more. But even for isolated pulsars it is expensive to search for high (> 100 Hz) spin frequencies, and adding (at least three) additional orbital parameters makes it even more costly. By improving the techniques, the methods presented here are a first step toward finding more of these systems.

Much of our focus is on binary pulsars in so-called “spider” systems, in which the pulsar companion is being evaporated by an energetic pulsar wind. A typical example is the first “black widow” pulsar to be discovered, PSR B1957+20 [49]. This was found in radio, where pulsations are eclipsed for a large fraction of the orbit, presumably by material ablated from the companion. Spider pulsars are categorized as black widows if the companion mass M_c is very low ($M_c \ll 0.1 M_\odot$) or as “redbacks” (another spider species) for larger companion masses ($M_c \sim 0.15 - 0.7 M_\odot$) [e.g., 53, 54], with one redback candidate likely having an even higher companion mass $\gtrsim 0.8 M_\odot$ [55].

For many of the known MSPs in spider systems, the companions are visible in the

¹<https://tinyurl.com/fermipulsars>

²These searches have been called “blind” searches in previous literature.

³<http://www.atnf.csiro.au/research/pulsar/psrcat>

optical. The light originates from nuclear burning, and/or from pulsar wind heating up the companion. The orbital motion of the companion then leads to a detectable modulation of the orbital brightness. The source of this modulation is not well understood. It might be that the side of the companion facing the pulsar is hotter than the other side and is more visible at the companion's superior conjunction. The companion might also be tidally elongated into an ellipsoid, whose projected cross section onto the line of sight varies over the orbit.

The new search methods presented here are well suited to gamma-ray pulsars in spider systems, with nearly circular orbits (eccentricity $e < 0.05$) and for which optical observations of the pulsar's companion provide information about the orbital motion, and thus constrain the gamma-ray pulsation search space.

For concreteness, we present the search designs for two promising gamma-ray sources: (a) 4FGL J1653.6–0158, a likely MSP in a circular binary [111, 112], and (b) 4FGL J0523.3–2527, a probable MSP in a slightly eccentric binary [55]. These are ranked among the most likely pulsar candidates [67]. We demonstrate the feasibility of a search using the computing resources of the distributed volunteer computing project *Einstein@Home* [85].

The paper is organized as follows. Section 2.2 reviews partially informed search methods for isolated gamma-ray pulsars and introduces the concepts required for such searches. Section 2.3 extends the methods to gamma-ray pulsars in circular orbit binaries, and Section 2.4 further extends these to eccentric orbit binaries. In Section 2.5 our methods are compared with alternatives used in radio and gravitational-wave astronomy. Finally, in Section 2.6 we discuss the feasibility of future partially informed searches for binary gamma-ray pulsars and also consider some specific sources. This is followed by Appendices 2.7.1, 2.7.2, and 2.7.3 containing some technical details.

In this paper, c denotes the speed of light and G denotes Newton's gravitational constant.

2.2 Partially-informed gamma-ray searches for pulsars

Partially informed search methods for isolated gamma-ray pulsars have been studied in detail by Pletsch and Clark [80]. Here we summarize and extend their framework. The following sections generalize the search methods to binary pulsars.

The search for gamma-ray pulsations begins with a list of N photons from a posited source, which we label with the index $j = 1, \dots, N$. The data available for these photons are their detector arrival time t_j , their direction of origin, and their energy, spanning an observation interval T_{obs} .

We are dealing with many sums and products in this paper. Sums and products over j, k, ℓ run from $1, \dots, N$ unless otherwise specified. Furthermore, we adopt the notation

$$\sum_{j \neq k} \equiv \sum_{j=1}^N \sum_{\substack{k=1 \\ j \neq k}}^N \quad (2.1)$$

for simplicity reasons.

Not all photons are equally significant. Photons at low energies are less well localized than those at higher energies and cannot be so readily attributed to a target source. Photons whose energy is more consistent with a distributed background are less likely to come from the pulsar. Photons originating from a nearby point source might contaminate the data set. For such reasons, searches may be improved by modeling the spatial and energy distribution of the sources.

To quantify the significance, we assign a weight $w_j \in [0, 1]$ to each photon. This weight w_j represents the probability that the j th photon originated at the nominal pulsar [113, 114]. The photon weights are determined from an assumed spectral and spatial model of gamma-ray sources in the region around the target pulsar, which is obtained using the standard methods for fitting gamma-ray sky maps⁴.

Each photon’s weight is computed as the predicted fraction that the target pulsar contributes to the total photon flux at the photon’s energy and arrival direction, after convolution with the *Fermi*-LAT’s energy-dependent point-spread function [114, 115]. The weighting process, and hence the resulting w_j , is the only place where the energy and arrival direction of the photons enter our analysis. In practice, the weights are computed using `gtsrcprob` from the *Fermi* Science Tools⁵, using, e.g., the 4FGL catalog [38] and associated Galactic and isotropic diffuse emission templates as the input model. These weights are used for noise suppression and to reduce computing cost by removing the lowest-weighted photons.

In this paper, we assume that these weights have been determined in advance for each photon, so the only information available for the j th photon is its arrival time t_j in the detector and the weight w_j .

The question that we need to answer is, are the arrival times of these photons random, or is there an underlying periodicity? To answer this question (in the statistical sense), we first need a model for the periodicity, which we assume is tied to the physical rotation of the pulsar.

2.2.1 Pulse profile and photon arrival probability

For now, assume that “in isolation” the pulsar would have a linearly changing angular velocity. Using Φ to denote the rotational phase in radians

$$\Phi(t_{\text{psr}}, \boldsymbol{\lambda}) = 2\pi f(t_{\text{psr}} - t_{\text{ref}}) + \pi \dot{f}(t_{\text{psr}} - t_{\text{ref}})^2, \quad (2.2)$$

where t_{psr} is the time that would be measured by a fictitious observer freely falling with the center of mass of the pulsar, and t_{ref} is a reference time. Note that detector time ticks at a different rate than t_{psr} , because the detector is moving around the Earth and the Sun, and because the pulsar might be orbiting a binary companion, or accelerating toward the Galaxy. Also note that without loss of generality we have set the phase at the reference time to zero.

The parameters $\boldsymbol{\lambda}$ describe the pulsar. Here they are the spin frequency f and its first time derivative \dot{f} at reference time t_{ref} . This second-order Taylor approximation holds for many pulsars and most MSPs, but for very young and “glitching” pulsars, additional higher-order terms may be needed.

The flux of photons can be broken into three parts. The first does not come from the pulsar: it is a background that is uncorrelated with pulsar rotation. We call

⁴<https://fermi.gsfc.nasa.gov/ssc/data/analysis/scitools/>

⁵<https://fermi.gsfc.nasa.gov/ssc/data/analysis/software/>

these unpulsed photons “background”. The second part originates from the pulsar itself but is also uncorrelated with pulsar rotation. We call these “unpulsed source” photons. The last part is a periodically time-varying flux from the source, which we call “pulsed”. We use p to denote the ratio of the number of pulsed photons to the total number of source photons (pulsed and unpulsed source).

The pulsed photon flux may be described with a periodic function $F_S(\Phi)$ of the pulsar’s phase around its rotational axis, $\Phi \in [0, 2\pi]$, and is time stable for most pulsars. The normalized probability that a pulsed photon arrives in the phase interval $[\Phi, \Phi + d\Phi]$ is $F_S(\Phi) d\Phi$. The function $F_S(\Phi)$ has minimum value zero and encloses unit area in the interval $[0, 2\pi]$.

We can now give the probability density function for the rotation phase associated with a given photon. This differs from one photon to the next because photons with small weight w_j are more likely to have a phase-independent probability distribution. The probability that the j th photon originates from a rotation phase interval $[\Phi_j, \Phi_j + d\Phi_j]$ is $F_j(\Phi_j) d\Phi_j$, where

$$F_j(\Phi_j) = \frac{1 - w_j}{2\pi} + w_j \left[\frac{1 - p}{2\pi} + pF_S(\Phi_j) \right]. \quad (2.3)$$

The first term (with probability $1 - w_j$) describes the background photons, and the second and third terms (with probability w_j) describe the unpulsed and pulsed source photons, respectively.

The probability distribution of pulsed photons may be expressed as the Fourier series

$$F_S(\Phi) = \frac{1}{2\pi} + \frac{1}{2\pi} \sum_{n=1}^{\infty} (\gamma_n e^{in\Phi} + \gamma_n^* e^{-in\Phi}). \quad (2.4)$$

The complex Fourier coefficients are

$$\gamma_n = \int_0^{2\pi} F_S(\Phi) e^{-in\Phi} d\Phi. \quad (2.5)$$

Note that the Fourier coefficients γ_n are constrained because F_S has minimum value zero. Note also that for known gamma-ray pulsars $|\gamma_n|^2$ decreases quickly with increasing index n [80]. In many cases the first five harmonics are sufficient to describe the pulse profile.

In principle, to detect gamma-ray pulsations, we assume a rotational model f, \dot{f} and then compute the rotational phase associated with each photon. “Binning” these phases (mod 2π) with weights w_j provides an estimate of $F(\Phi) = \sum_j w_j F_j(\Phi) / \sum_j w_j$, from which we can estimate $F_S(\Phi)$ by shifting the minimum value to zero and rescaling to unit area. If that function is compatible with zero (meaning: coefficients γ_n are small), then no pulsations were detected. Conversely, if the γ_n are large for some values of f and \dot{f} , we have found pulsations.

2.2.2 Relationship of detector time t to t_{psr}

The situation is slightly more complicated than described in the previous paragraph because computing t_{psr} for each photon from its time of arrival at the *Fermi* satellite also requires the pulsar’s sky position (right ascension α and declination δ). The sky position allows for “barycentric corrections”, e.g., to account for Doppler shifts

due to the LAT’s movement around the solar system barycenter (SSB). Thus, the photon’s emission time $t_{\text{psr}}(t, \alpha, \delta)$ is a function of its arrival time t at the LAT and the putative pulsar’s sky position. The pulsar’s putative phase is a function of t and the four parameters $\boldsymbol{\lambda} = \{f, \dot{f}, \alpha, \delta\}$.

In partially informed searches the spin parameters are unknown. Although each photon is tagged with an arrival direction α, δ , these are not sufficiently precise to detect pulsations, so those location parameters must also be searched. Hence, the parameter space search volume Λ for isolated pulsars ($\boldsymbol{\lambda} \in \Lambda$) is 4-dimensional. In Sections 2.3 and 2.4, the higher-dimensional search spaces for binary pulsars in circular and elliptical orbits are discussed.

2.2.3 Searching for pulsations

For realistic searches the parameter space Λ is too large to search by the straightforward computational process described above. Instead, Λ is explored with a multistage search based on several different test statistics [e.g., 116]. This gives the greatest sensitivity at fixed computational cost [80]. The approach is hierarchical. In the first stage, a coarse grid covering the parameter space Λ is searched at low sensitivity using inexpensive test statistics. These are relatively insensitive to mismatch between tested parameters and pulsar parameters. In the following stages, smaller regions of Λ around the most promising candidates are searched at higher sensitivity. These use more expensive test statistics on finer, more closely spaced grids. Thus, a search is defined by a test statistic/grid hierarchy.

The spacing of the grids in parameter space is governed by the mismatch described above. For a given test statistic, we calculate a “metric”, which is the fractional loss in the expected signal-to-noise ratio (S/N). The details of this are found later in this section.

The search described in this paper has four stages, which employ detection statistics P_1, S_1 , and H . Here we briefly describe the overall structure. The test statistics are defined and characterized later in this section.

The first three stages search for significant power in the first harmonic $|\gamma_1|^2$. Each discards regions of parameter space that contain no signals; what remains is passed to the following stage. The first stage uses the “semicoherent” test statistic S_1 with a low threshold. The second stage tests S_1 on a finer grid, with a higher threshold. The third stage uses the fully coherent test statistic P_1 . This searches coherently for power $|\gamma_1|^2$ over the full observation span T_{obs} with much greater sensitivity and a finer grid than before.

The fourth stage employs the expensive H statistic, which combines P_1, \dots, P_5 . This coherently integrates over T_{obs} to identify power in the first five harmonics $|\gamma_1|^2, \dots, |\gamma_5|^2$. By searching around the surviving candidate points in parameter space with a still finer grid, this completes the hierarchy.

2.2.4 Coherent power test statistic P

The basis for all of our test statistics is the coherent Fourier power, evaluated over different periods of time. For the n th harmonic, and including all of the photons, this is

$$P_n(\boldsymbol{\lambda}) = \frac{1}{\kappa^2} \left| \sum_j w_j e^{-in\Phi(t_j, \boldsymbol{\lambda})} \right|^2. \quad (2.6)$$

To simplify notation, from here on we use $\Phi(t_j, \boldsymbol{\lambda})$ to denote $\Phi(t_{\text{psr}}(t_j, \alpha, \delta), f, \dot{f})$, where t_j is the photon arrival time measured at the LAT. The normalization constant is

$$\kappa^2 = \frac{1}{2} \sum_j w_j^2. \quad (2.7)$$

How does P_n behave in the absence of pulsations and in the presence of pulsations?

To answer this question, we compute expectation values as shown in Appendix 2.7.1. The power P_n has an expected value (Eq. 2.104) and variance (in the absence of a pulsed signal, $p = 0$)

$$\mathbb{E}_p[P_n] = 2 + \kappa^{-2} p^2 |\gamma_n|^2 \sum_{j \neq k} w_j^2 w_k^2 \quad (2.8)$$

$$\text{Var}_0[P_n] = \kappa^{-4} \sum_{j \neq k} w_j^2 w_k^2. \quad (2.9)$$

The power P_n is a *detection statistic* because it is sensitive to a nonvanishing pulse profile. If γ_n is nonzero, then P_n should be larger than 2. It becomes larger as the fraction p of pulsed to source photons increases (which we cannot control). It also becomes larger as the number of photons (or equivalently, the observation time) grows. But to understand what values of P_n correspond to statistically significant detections, we need to know about its statistical fluctuations, meaning the variance in P_n .

Note that the diagonal-free double sum in these expressions can be reexpressed as $(\sum_j w_j^2)^2 - \sum_j w_j^4$. Thus, the variance can be written as

$$\text{Var}_0[P_n] = 4 - 4 \frac{\sum_j w_j^4}{(\sum_j w_j^2)^2}. \quad (2.10)$$

If there are many photons from the source and the weights are relatively uniformly distributed, then it follows that the numerator in Eq. (2.10) is $\mathcal{O}(N)$ and the denominator is $\mathcal{O}(N^2)$. Hence, the variance $\text{Var}_0[P_n] \rightarrow 4 - \mathcal{O}(1/N)$ approaches 4. In this limit, and with the statistical assumptions of Appendix 2.7.1, P_n has a noncentral χ^2 -distribution with two degrees of freedom [80]. The noncentrality parameter is the second term appearing in Eq. (2.8).

The expected S/N associated with P_n is

$$\begin{aligned} \theta_{P_n}^2 &= \frac{\mathbb{E}_p[P_n] - \mathbb{E}_0[P_n]}{\sqrt{\text{Var}_0[P_n]}} = p^2 |\gamma_n|^2 \sqrt{\sum_{j \neq k} w_j^2 w_k^2} \\ &= p^2 |\gamma_n|^2 \mu T_{\text{obs}}. \end{aligned} \quad (2.11)$$

In the many-photon limit the quantity $\mu \rightarrow \sum_j w_j^2 / T_{\text{obs}}$ is proportional to the mean weighted photon arrival rate.

2.2.4.1 Loss of P from parameter mismatch

In a real search, we compute detection statistics at a grid of discrete values of the signal parameters $\boldsymbol{\lambda}$. If there is a signal present, its actual (true) parameters might be close to one of these discrete values but will not match it exactly. There will

always be *some* offset between the tested parameters and the true parameters. Here we quantify how much S/N is expected to be lost because of this mismatch.

Assume that the tested parameters $\boldsymbol{\lambda}$ are close to the true pulsar parameters $\boldsymbol{\lambda}_{\text{psr}}$, and introduce the notation

$$d\lambda^a = \lambda^a - \lambda_{\text{psr}}^a \quad (2.12)$$

for the small parameter offsets. Here and elsewhere in the paper we index the parameter space dimension with lowercase Latin letters “ a ” and “ b ”. These offsets change the pulsar rotation phase by

$$\Delta\Phi(t) = \Phi(t, \boldsymbol{\lambda}) - \Phi(t, \boldsymbol{\lambda}_{\text{psr}}) \approx \partial_a \Phi d\lambda^a, \quad (2.13)$$

where the notation

$$\partial_a \Phi = \left. \frac{\partial \Phi}{\partial \lambda^a} \right|_{\boldsymbol{\lambda}=\boldsymbol{\lambda}_{\text{psr}}} \quad (2.14)$$

is introduced and we neglect higher powers in $d\lambda$. We also adopt the *Einstein summation convention* that repeated parameter space indices are summed over all the dimensions of the parameter space.

We now compute the fractional loss in expected S/N associated with this parameter mismatch. For the offset parameters the coherent power is

$$P_n(\boldsymbol{\lambda}) = 2 + \kappa^{-2} \sum_{j \neq k} w_j w_k e^{in(\Phi_j - \Phi_k)} e^{in(\Delta\Phi_j - \Delta\Phi_k)}, \quad (2.15)$$

where $\Phi_j = \Phi(t_j, \boldsymbol{\lambda}_{\text{psr}})$ and $\Delta\Phi_j = \Delta\Phi(t_j)$. Following Appendix 2.7.1, the expectation value of this is

$$E_p[P_n(\boldsymbol{\lambda})] = 2 + \kappa^{-2} p^2 |\gamma_n|^2 \sum_{j \neq k} w_j^2 w_k^2 e^{in(\Delta\Phi_j - \Delta\Phi_k)}. \quad (2.16)$$

It follows that for the mismatched signal the expected S/N is

$$\theta_{P_n(\boldsymbol{\lambda})}^2 = p^2 |\gamma_n|^2 \left[\sum_{j \neq k} w_j^2 w_k^2 \right]^{-1/2} \sum_{j \neq k} w_j^2 w_k^2 e^{in(\Delta\Phi_j - \Delta\Phi_k)}. \quad (2.17)$$

The fractional loss in S/N (often called the “mismatch”) is

$$\begin{aligned} m(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) &= \frac{\theta_{P_n}^2(\boldsymbol{\lambda}_{\text{psr}}) - \theta_{P_n}^2(\boldsymbol{\lambda})}{\theta_{P_n}^2(\boldsymbol{\lambda}_{\text{psr}})} \\ &= \sum_{j \neq k} w_j^2 w_k^2 [1 - e^{in(\Delta\Phi_j - \Delta\Phi_k)}] / \sum_{j \neq k} w_j^2 w_k^2 \\ &= \left[\left(\sum_j w_j^2 \right)^2 - \left| \sum_j w_j^2 e^{in\Delta\Phi_j} \right|^2 \right] / \left[\left(\sum_j w_j^2 \right)^2 - \sum_j w_j^4 \right]. \end{aligned} \quad (2.18)$$

We need the mismatch to help set the spacings of the parameter space search grids, but for that purpose, approximations suffice.

Assume that there are many photons and the weights are uniformly distributed in time (or at least slowly varying in a way that is not correlated with the pulsar

rotation phase). The sums over the weights may then be replaced with simple integrals over time, giving

$$m(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) \approx 1 - \left| \left\langle e^{-in\Delta\Phi(t)} \right\rangle_{T_{\text{obs}}}(t_0) \right|^2. \quad (2.19)$$

Here we introduce the ‘‘angle bracket’’ notation for an average over a time interval of length T centered around an arbitrary time t_0 . This takes an input function $Q(t')$ and outputs a new function of time t defined by

$${}_T\langle Q(t') \rangle(t) \equiv \frac{1}{T} \int_{t-T/2}^{t+T/2} Q(t') dt', \quad (2.20)$$

which is the average of Q around the time t .

2.2.4.2 Parameter space metric g_{ab}

Since the sensitivity of these searches is limited by available computing power, we need to construct a grid that covers the relevant parameter space with the smallest number of grid points. This means that the parameters $\boldsymbol{\lambda}_{\text{psr}}$ of any possible pulsar should be close enough to a grid point that we do not lose too much S/N from the mismatch, but the grid should have as few points as possible.

The distance metric on the search space is a useful tool for such constructions [117, 118]. It provides an analytical approximation to the mismatch. For example, the coherent mismatch in Equation (2.19) can be approximated by the ‘‘coherent metric’’ g_{ab}

$$m(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) = n^2 g_{ab}(\boldsymbol{\lambda}) d\lambda^a d\lambda^b + \mathcal{O}(d\lambda^3) \quad (2.21)$$

for small coordinate offsets $d\lambda^a$ from the true pulsar parameters.

Expanding the exponential that appears in Eq. (2.19) to first order, one finds

$$g_{ab} = {}_{T_{\text{obs}}}\langle \partial_a \Phi \partial_b \Phi \rangle(t_0) - {}_{T_{\text{obs}}}\langle \partial_a \Phi \rangle(t_0) {}_{T_{\text{obs}}}\langle \partial_b \Phi \rangle(t_0). \quad (2.22)$$

To evaluate the metrics, we need to account for the way in which the detected pulsar rotation phase depends on the different pulsar parameters.

2.2.4.3 Evaluation of g_{ab} for isolated pulsars

As seen by an observer freely falling at the center of mass of the pulsar, the rotation phase just depends on the intrinsic frequency f and its derivative \dot{f} as given in Eq. (2.2). But as explained in Sec. 2.2.2, these must be converted to detector time.

For computing the metric, we do not need a conversion that is accurate to microseconds, but only one that takes into account the largest shifts between detector and pulsar time, of order ≈ 500 s, arising from the motion of the Earth around the Sun [80]. We denote the orbital angular frequency by $\Omega_E = 2\pi/\text{yr}$, the orbital light-crossing time by $r_E = 1 \text{ AU}/c$, and the obliquity of the ecliptic by $\epsilon = 23.4^\circ$.

If we choose a coordinate axis z along the line of sight to the pulsar, then the projected motion is

$$r_{z,\text{sky}}(t) = r_E [n_x \cos(\Omega_E t + \varphi_{\text{ref}}) + n_y \sin(\Omega_E t + \varphi_{\text{ref}})], \quad (2.23)$$

where

$$n_x = \cos \alpha \cos \delta, \quad (2.24)$$

$$n_y = \cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta, \quad (2.25)$$

and the sky location is given by the right ascension α , and the declination δ . The (arbitrary) choice for the origin of the time coordinate determines the constant φ_{ref} , which is the Earth's orbital phase at that moment.

Note that this simplified version of the Rømer delay does not account for the motion of the *Fermi* satellite around the Earth. It is not accurate enough to use in a search for pulsations and is only used in the metric calculation.

For the purpose of computing the metric we can model the detected pulsar rotation phase as the sum of Eq. (2.2) and the additional phase cycles introduced by the Rømer delay (2.23):

$$\begin{aligned} \Phi(t, \boldsymbol{\lambda}) = & 2\pi f(t - t_{\text{ref}}) + \pi \dot{f}(t - t_{\text{ref}})^2 \\ & + 2\pi f r_E [n_x \cos(\Omega_E t + \varphi_{\text{ref}}) + n_y \sin(\Omega_E t + \varphi_{\text{ref}})]. \end{aligned} \quad (2.26)$$

Here the search parameters are $\boldsymbol{\lambda} = \{f, \dot{f}, n_x, n_y\}$, and the terms correcting the arrival times t have been neglected for the \dot{f} summand.

The metric for the coherent power P_1 follows from Eq. (2.22). The formulae are complicated, but if we keep only the most significant terms, then they simplify. To determine these, consider the relative size of the different quantities:

$$\begin{aligned} T_{\text{obs}} & \approx 10 \text{ yr} \approx 3 \times 10^8 \text{ s}, \\ |t_0 - t_{\text{ref}}| & \lesssim T_{\text{obs}}, \\ \Omega_E & \approx 2\pi/\text{yr} \approx 2 \times 10^{-7} \text{ s}^{-1}, \\ r_E & \approx 5 \times 10^2 \text{ s}, \\ f & \approx (100 - 700) \text{ s}^{-1}, \\ \dot{f} & \approx (10^{-16} - 10^{-14}) \text{ s}^{-2}. \end{aligned} \quad (2.27)$$

Most MSPs have parameters f and \dot{f} in the given range. With these in mind, one finds diagonal metric components

$$\begin{aligned} g_{ff} & = \frac{1}{3} \pi^2 T_{\text{obs}}^2 [1 + \mathcal{O}(r_E/T_{\text{obs}})], \\ g_{\dot{f}\dot{f}} & = \frac{1}{180} \pi^2 T_{\text{obs}}^4 \left[1 + 60 \frac{(t_0 - t_{\text{ref}})^2}{T_{\text{obs}}^2} \right], \\ g_{n_x n_x} & = 2\pi^2 f^2 r_E^2 [1 + \mathcal{O}(1/\Omega_E T_{\text{obs}})], \\ g_{n_y n_y} & = 2\pi^2 f^2 r_E^2 [1 + \mathcal{O}(1/\Omega_E T_{\text{obs}})]. \end{aligned} \quad (2.28)$$

Most of the off-diagonal metric components are negligible.

Determining whether off-diagonal metric components are significant requires some care because they need to be compared to the corresponding diagonal components. This arises here and in several other places in the paper. Here we show in detail how this significance is determined. The same reasoning is used for the other cases that arise later but is not elaborated.

Since the fundamental quantity of interest is the mismatch m , for fixed a and b (no Einstein summation convention), consider $m = g_{aa}(\text{d}\lambda^a)^2 + g_{bb}(\text{d}\lambda^b)^2 + 2g_{ab} \text{d}\lambda^a \text{d}\lambda^b$.

Rescale the coordinates $\{\lambda^a, \lambda^b\}$ to new coordinates $\{\lambda^{a'} = u\lambda^a, \lambda^{b'} = w\lambda^b\}$ such that the two diagonal components of the metric in the new coordinates are both unity. (Here u and w denote the rescaling factors.) This implies that $g_{aa}(\partial\lambda^a/\partial\lambda^{a'})^2 = g_{aa}u^{-2} = 1$ and $g_{bb}(\partial\lambda^b/\partial\lambda^{b'})^2 = g_{bb}w^{-2} = 1$. Then, all off-diagonal metric components are of $\mathcal{O}(1/\Omega_E T_{\text{obs}})$, apart from

$$g_{ff} = \frac{1}{3}\pi^2 T_{\text{obs}}^3 \left[\frac{(t_0 - t_{\text{ref}})}{T_{\text{obs}}} + \mathcal{O}(r_E/T_{\text{obs}}) \right]. \quad (2.29)$$

Note that all the off-diagonal terms may be neglected in the case that the integration time $T_{\text{obs}} \gg 1$ yr and the reference time $t_{\text{ref}} = t_0$.

For this case the diagonal ‘‘coherent metric’’ terms reduce to

$$\begin{aligned} g_{ff} &= \frac{1}{3}\pi^2 T_{\text{obs}}^2, \\ g_{jj} &= \frac{1}{180}\pi^2 T_{\text{obs}}^4, \\ g_{n_x n_x} &= 2\pi^2 f^2 r_E^2, \\ g_{n_y n_y} &= 2\pi^2 f^2 r_E^2. \end{aligned} \quad (2.30)$$

2.2.5 Semicoherent power test statistic S

The coherent power P_n in Eq. (2.6) provides a good statistical basis to find pulsations (meaning γ_n nonzero) but is inefficient to compute. Hence, the first two stages of our searches use the ‘‘semicoherent’’ Fourier power S_n . Its definition is similar to P_n except that photons are only combined if their arrival time difference is smaller than a coherence time, $T_{\text{coh}} \ll T_{\text{obs}}$. This makes it less expensive to compute (but also less sensitive). The coherence time in a typical search in the first stage is $T_{\text{coh}} = 2^{21}$ s ≈ 24 d, in the second stage it is $T_{\text{coh}} = 2^{22}$ s ≈ 48 d, and the observation span T_{obs} (i.e. the operation time of the LAT) is more than 10 years.

For convenience the statistic S_n differs from P_n in one other way: we omit the diagonal $j = k$ terms in the sum. This ensures that in the no-signal ($p = 0$) case the expected value of S_n vanishes, with

$$S_n(\boldsymbol{\lambda}) = \frac{1}{\bar{\kappa}} \sum_{j \neq k} w_j w_k e^{-in[\Phi(t_j, \boldsymbol{\lambda}) - \Phi(t_k, \boldsymbol{\lambda})]} \hat{W}_{T_{\text{coh}}}(\tau_{jk}). \quad (2.31)$$

The rectangular window function restricts the sum to photons in which the arrival time difference $\tau_{jk} = t_j - t_k$ (or ‘‘lag’’) is not larger than T_{coh} :

$$\hat{W}_{T_{\text{coh}}}(\tau) = \begin{cases} 1 & \text{for } |\tau| \leq T_{\text{coh}}/2, \\ 0 & \text{otherwise.} \end{cases} \quad (2.32)$$

The semicoherent normalization constant is chosen to be

$$\bar{\kappa} = \sqrt{\sum_{j \neq k} w_j^2 w_k^2 \hat{W}_{T_{\text{coh}}}^2(\tau_{jk})}, \quad (2.33)$$

which ensures that in the no-signal ($p = 0$) case S_n has unit variance [79].

To characterize this detection statistic, we calculate the expectation value and variance with the calculational framework of Appendix 2.7.1, obtaining

$$\mathbb{E}_p[S_n] = \frac{1}{\bar{\kappa}} p^2 |\gamma_n|^2 \sum_{j \neq k} w_j^2 w_k^2 \hat{W}_{T_{\text{coh}}}(\tau_{jk}), \quad (2.34)$$

$$\text{Var}_0[S_n] = 1. \quad (2.35)$$

The expectation value is the same as the second term of P_n in Eq. (2.8), except that the sum is restricted to the lag window. In fact, the formulae above hold for any choice of window function.

The S/N for the semicoherent Fourier power S_n is simplified by assuming a rectangular window function (which equals its square). This gives

$$\begin{aligned} \theta_{S_n}^2 &= \frac{\mathbb{E}_p[S_n] - \mathbb{E}_0[S_n]}{\sqrt{\text{Var}_0[S_n]}} = p^2 |\gamma_n|^2 \sqrt{\sum_{j \neq k} w_j^2 w_k^2 \hat{W}_{T_{\text{coh}}}(\tau_{jk})} \\ &= p^2 |\gamma_n|^2 \mu \sqrt{T_{\text{coh}} T_{\text{obs}}}. \end{aligned} \quad (2.36)$$

The second line adopts the definition of μ given after Eq. (2.11) and makes the same assumptions of steady photon flux and large photon number.

In practice, how large are these detection statistics? A typical gamma-ray pulsar might have a pulsed flux for which $|\gamma_1|^2 \approx 0.2$ and a 70% fraction of pulsed photons for which $p^2 \approx 0.5$. The weighted flux of source photons detected might be $\sum_j w_j^2 \approx 500$ over $T_{\text{obs}} = 10$ yr, implying a rate $\mu \approx 50 \text{ yr}^{-1}$. With $T_{\text{coh}} = 24$ d, this leads to coherent and incoherent S/Ns of order $\theta_{P_1}^2 \approx 50$ and $\theta_{S_1}^2 \approx 4$, significant at the 50σ and 4σ levels, respectively.

2.2.5.1 Loss of S from parameter mismatch

We now turn to the metric for the semicoherent statistic. To compute the mismatch for the semicoherent detection statistic S_n , with the same assumptions as above, we can replace the sums with integrals, obtaining

$$\begin{aligned} \bar{m}(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) &= 1 - \frac{\theta_{S_n}^2(\boldsymbol{\lambda})}{\theta_{S_n}^2(\boldsymbol{\lambda}_{\text{psr}})} \\ &= 1 - \left\langle \left\langle e^{-in\Delta\Phi(t')} \right\rangle_{T_{\text{coh}}} \left\langle e^{in\Delta\Phi(t'')} \right\rangle_{T_{\text{coh}}} \right\rangle_{T_{\text{obs}}}(t_0). \end{aligned} \quad (2.37)$$

Note that the inner integral in the second line can include times outside the observation span $t'' \in [t_0 - T_{\text{obs}}/2, t_0 + T_{\text{obs}}/2]$, going down to $t'' = t_0 - T_{\text{obs}}/2 - T_{\text{coh}}/2$ or up to $t'' = t_0 + T_{\text{obs}}/2 + T_{\text{coh}}/2$. In such cases the integrand should be set to zero and normalized so that $\langle 1 \rangle = 1$.

2.2.5.2 Parameter space metric \bar{g}_{ab}

We now evaluate these mismatches to lowest order, obtaining a *distance metric* on the parameter space. We evaluate the integrals in Eq. (2.37) naively, without setting the integrands to zero outside of the “valid data range”. This gives rise to terms (complex or linear in $d\boldsymbol{\lambda}^a$) that are not present in the exact expression. We assume

that $T_{\text{coh}} \ll T_{\text{obs}}$ (typically $T_{\text{coh}} = 24$ d and $T_{\text{obs}} > 10$ yr). In that case, these terms are small, and we discard them.

The partial derivatives with respect to $\lambda^a \in \{f, \dot{f}, n_x, n_y\}$, under the assumption that $T_{\text{coh}} \ll 1 \text{ yr} \ll T_{\text{obs}}$, can be approximated as

$$\partial_a \Phi \approx_{T_{\text{coh}}} \langle \partial_a \Phi \rangle(t), \quad (2.38a)$$

$$\partial_a \partial_b \Phi \approx_{T_{\text{coh}}} \langle \partial_a \partial_b \Phi \rangle(t), \quad (2.38b)$$

as [80] did. (Here and in what follows, for readability, the time dependence of phase derivatives such as $\partial_a \Phi$ is not shown explicitly.)

With these assumptions the semicoherent mismatch Eq. (2.37) can be approximated by the semicoherent metric

$$\bar{m}(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) \approx n^2 \bar{g}_{ab} d\lambda^a d\lambda^b + \mathcal{O}(d\lambda^3), \quad (2.39)$$

where $d\lambda^a = \lambda^a - \lambda_{\text{psr}}^a$ as earlier. Note that Eq. (2.39) has the same form as the coherent mismatch in Eq. (2.21).

The metric components are

$$\begin{aligned} \bar{g}_{ab} &= \frac{1}{T_{\text{obs}}} \left\langle T_{\text{coh}} \langle \partial_a \Phi \partial_b \Phi \rangle(t') - T_{\text{coh}} \langle \partial_a \Phi \rangle(t') T_{\text{coh}} \langle \partial_b \Phi \rangle(t') \right\rangle(t_0), \\ &= \frac{1}{2} \frac{1}{T_{\text{obs}}} \left\langle \tilde{g}_{ab}(t') \right\rangle(t_0), \end{aligned} \quad (2.40)$$

where we have introduced

$$\tilde{g}_{ab}(t') = T_{\text{coh}} \langle \partial_a \Phi \partial_b \Phi \rangle(t') - T_{\text{coh}} \langle \partial_a \Phi \rangle(t') T_{\text{coh}} \langle \partial_b \Phi \rangle(t'), \quad (2.41)$$

which is exactly the coherent metric given in Eq. (2.22), but with T_{obs} replaced by $T_{\text{coh}} \ll 1 \text{ yr}$ and t_0 replaced by t' . Thus, terms of $\mathcal{O}(1/\Omega_E T_{\text{coh}})$, similar to those appearing in Eq. (2.28), cannot be neglected.

2.2.5.3 Evaluation of \bar{g}_{ab} for isolated pulsars

The nonvanishing semicoherent metric components are

$$\begin{aligned} \bar{g}_{ff} &= \frac{1}{6} \pi^2 T_{\text{coh}}^2 [1 + \mathcal{O}(r_E/T_{\text{obs}})], \\ \bar{g}_{f\dot{f}} &= \frac{1}{6} \pi^2 T_{\text{coh}}^2 T_{\text{obs}} \left[\frac{(t_0 - t_{\text{ref}})}{T_{\text{obs}}} + \mathcal{O}(r_E/T_{\text{obs}}) \right], \\ \bar{g}_{\dot{f}\dot{f}} &= \frac{1}{72} \pi^2 T_{\text{coh}}^2 T_{\text{obs}}^2 \left(1 + 12 \frac{(t_0 - t_{\text{ref}})^2}{T_{\text{obs}}^2} + \frac{T_{\text{coh}}^2}{5T_{\text{obs}}^2} \right), \\ \bar{g}_{n_x n_x} &= \pi^2 f^2 r_E^2 \left[1 - \frac{4}{\Omega_E^2 T_{\text{coh}}^2} \sin^2 \left(\frac{\Omega_E T_{\text{coh}}}{2} \right) + \mathcal{O}(1/\Omega_E T_{\text{obs}}) \right], \\ \bar{g}_{n_y n_y} &= \pi^2 f^2 r_E^2 \left[1 - \frac{4}{\Omega_E^2 T_{\text{coh}}^2} \sin^2 \left(\frac{\Omega_E T_{\text{coh}}}{2} \right) + \mathcal{O}(1/\Omega_E T_{\text{obs}}) \right]. \end{aligned} \quad (2.42)$$

The semicoherent metric \bar{g} is diagonal for $t_0 = t_{\text{ref}}$, as was the case for the coherent metric g in Eq. (2.30). Note that $g_{n_x n_x}$ and $g_{n_y n_y}$ are not equal because the neglected terms of $\mathcal{O}(1/\Omega_E T_{\text{obs}})$ have opposite signs.

The metric component $\bar{g}_{\dot{j}\dot{j}}$ differs from that given by Pletsch and Clark [80], but our results are identical in the limit of a large number of photons N homogeneously distributed over the observation span. This is the case, since we assumed it in deriving Eqs. (2.19) and (2.37).

Comparison of Eqs. (2.28) and (2.42) illustrates the benefits of the multistage search process described in Section 2.2.3. For grids with the same mismatch $m = \bar{m}$, the ratio between the density of the coherent grid and the semicoherent grid would be

$$\frac{\text{coherent grid density}}{\text{semicoherent grid density}} = \sqrt{\frac{\det g}{\det \bar{g}}} = \frac{48T_{\text{obs}}^2}{\sqrt{5}\Omega_{\text{E}}^2 T_{\text{coh}}^4}. \quad (2.43)$$

For the timescales T_{coh} and T_{obs} given above, the ratio is $\sim 10^6$. This is why the semicoherent search stage is beneficial.

2.2.6 Multiple harmonic test statistic H

In the last and most sensitive stage of the multistage search, we adopt the widely used statistic

$$H(\boldsymbol{\lambda}) = \max_{1 \leq M \leq M_{\text{max}}} \left(4 - 4M + \sum_{n=1}^M P_n(\boldsymbol{\lambda}) \right), \quad (2.44)$$

which incoherently sums the coherent power from up to the first M_{max} harmonics in the pulse profile. The H statistic provides a sensitive test for unknown (generic) pulse profiles. The original simulations by de Jager, Raubenheimer, and Swanepoel [119] recommended $M_{\text{max}} = 20$, and to assess the false-alarm probability, they carried out a numerical study of the distribution of H in pure noise.

Later results by Kerr [114] show that the single-trial probability ρ of exceeding a value $H_{\text{threshold}}$ in pure noise is well modeled by $\rho \approx \exp(-0.398 H_{\text{threshold}})$ if the number of harmonics M_{max} is very large. Obviously, if M_{max} is reduced, then the single-trial probabilities are smaller than this, so $\exp(-0.4 H_{\text{threshold}})$ is a reliable upper bound.

To avoid overfitting, we generally use smaller limits $M_{\text{max}} = 3, 4$, or 5 on the number of harmonics. Typical partially informed search gamma-ray pulsar detections have H values in the hundreds, corresponding to single-trial ρ values that must lie below 10^{-30} .

Normally, the last search stage is not computationally limited. Hence, we use a grid fine enough to secure power in the higher harmonics, while overcovering the search space for power in the lower harmonics. In practice, the grid is built using the coherent metric presented in Section 2.2.4.2 with $n = M_{\text{max}}$.

2.2.7 Searches for isolated pulsars

Partially informed searches for isolated pulsars within gamma-ray data recorded by the LAT have been very successful [see, e.g., 79]. The key ingredients are the utilization of the powerful volunteer-distributed computing system *Einstein@Home* [85] and searches that use these computing resources as efficiently as possible.

Most of the tools for constructing efficient searches have been presented in the earlier sections. To discard unpromising regions in parameter space, the multistage approach is used as described in Section 2.2.3. For the first and computationally

most crucial search stage, efficient grids covering the parameters \dot{f} , n_x , n_y are built based on the distance metric, and f is searched using fast Fourier transform (FFT) algorithms [120]. In later search stages, f is also gridded with the metric, but it is not efficient to use FFTs on the small ranges in f around the few most significant candidates from the semicoherent search stage.

2.3 Search Method: Circular Binary Orbits

The main problem in partially informed gamma-ray searches for pulsars is that the phase model from Eq. (2.2) depends on the (photon emission) time at the pulsar, while a gamma-ray detector records the time of arrival at the telescope. For binary pulsars the largest corrections to shift between these two times arise from the line-of-sight motion of the *Fermi* satellite around Earth and Sun $r_{z,\text{sky}}(t)$ and of the pulsar around its companion $r_{z,\text{cir}}(t_{\text{psr}})$.

The line-of-sight motion of a binary pulsar in a circular orbit can be described via three parameters, which are usually taken to be the orbital frequency Ω_{orb} , the projected semimajor axis x in seconds, and the epoch of ascending node t_{asc} . With these, the two times are related by

$$t_{\text{psr}} + r_{z,\text{cir}}(t_{\text{psr}}) = t + r_{z,\text{sky}}(t), \quad (2.45)$$

where the corrections, also called Rømer delays, are expressed in seconds.

The simplest expression of the pulsar's orbital line-of-sight motion $r_{z,\text{cir}}$ depends on the time measured at the pulsar t_{psr} . In many cases, this time may be replaced with the detector time because

$$r_{z,\text{cir}}(t_{\text{psr}}) = r_{z,\text{cir}}(t) [1 + \mathcal{O}(x\Omega_{\text{orb}})], \quad (2.46)$$

and the quantity $x\Omega_{\text{orb}} \ll 1$. In such cases

$$t_{\text{psr}} \approx t + r_{z,\text{sky}}(t) - r_{z,\text{cir}}(t). \quad (2.47)$$

This holds for most black widow and some redback systems with projected semimajor axes on the order of a few light-seconds [see, e.g., the ATNF Pulsar Catalogue⁶ by 110]. In all cases, it is accurate enough to compute the metric, and in many cases accurate enough for maintaining phase coherence in a search.

The Rømer delay can be expressed in terms of the three orbital parameters as

$$r_{z,\text{cir}}(t) = x \sin[\Omega_{\text{orb}}(t - t_{\text{asc}})]. \quad (2.48)$$

Here the orbital frequency Ω_{orb} is connected to the orbital period P_{orb} via $P_{\text{orb}} = 2\pi/\Omega_{\text{orb}}$.

In gamma-ray searches, in addition to the Rømer delay, we also have to correct for other effects like the Shapiro and Einstein delays. In contrast to radio observations, we do *not* have to account for the frequency-dependent dispersion caused by the interstellar medium (ISM) because gamma rays are well above the plasma frequency of the ISM.

All of these effects are described by Lorimer and Kramer [10], and Edwards, Hobbs, and Manchester [96]. While these corrections must be included in gamma-ray searches, only the largest effects need to be included in the phase model for the derivation of a distance metric approximation.

⁶<http://www.atnf.csiro.au/research/pulsar/psrcat>

2.3.1 Parameter space metrics

In order to compute the metric, a simplified phase model can be used that accounts for the corrections (2.23) and (2.48):

$$\begin{aligned}\Phi(t, \boldsymbol{\lambda}) = & 2\pi f(t - t_{\text{ref}}) + \pi \dot{f}(t - t_{\text{ref}})^2 \\ & + 2\pi f r_E [n_x \cos(\Omega_E t + \varphi_{\text{ref}}) + n_y \sin(\Omega_E t + \varphi_{\text{ref}})] \\ & - 2\pi f x \sin[\Omega_{\text{orb}}(t - t_{\text{asc}})].\end{aligned}\quad (2.49)$$

Here the search parameters are $\boldsymbol{\lambda} = \{f, \dot{f}, n_x, n_y, \Omega_{\text{orb}}, x, t_{\text{asc}}\}$ and the terms correcting the arrival times t have been neglected for the \dot{f} summand. This phase model is not sufficient for searches because it would not maintain phase coherence with a true pulsar signal. However, it is sufficient to describe how varying the signal parameters leads to loss of S/N.

The dominant components of the coherent metric for the orbital parameters are

$$\begin{aligned}g_{xx} &= 2\pi^2 f^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})], \\ g_{\Omega_{\text{orb}} \Omega_{\text{orb}}} &= \frac{1}{6} \pi^2 f^2 x^2 T_{\text{obs}}^2 \left[1 + 12 \frac{(t_o - t_{\text{asc}})^2}{T_{\text{obs}}^2} + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}}) \right], \\ g_{t_{\text{asc}} t_{\text{asc}}} &= 2\pi^2 f^2 x^2 \Omega_{\text{orb}}^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})], \\ g_{\Omega_{\text{orb}} t_{\text{asc}}} &= -2\pi^2 f^2 x^2 \Omega_{\text{orb}} T_{\text{obs}} \left[\frac{(t_o - t_{\text{asc}})}{T_{\text{obs}}} + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}}) \right],\end{aligned}\quad (2.50)$$

where we have assumed that the integration time span T_{obs} is much larger than the orbital period P_{orb} . Compared to the diagonal terms, as done in the text below Eq. (2.28), all other components are of $\mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})$.

The off-diagonal component $g_{\Omega_{\text{orb}} t_{\text{asc}}}$ is vanishingly small if the epoch of the ascending node is close to the middle of the gamma-ray data set, $t_{\text{asc}} \approx t_0$. In principle, t_{asc} can be shifted forward or backward by an integer number \mathcal{N} of orbital periods P_{orb} to achieve this. However, when t_{asc} is constrained, for example, by optical observations, this is undesirable because it introduces uncertainties in the shifted value of t_{asc} that grow linearly with \mathcal{N} .

Even if $t_{\text{asc}} \not\approx t_0$, our current searches ignore the off-diagonal term in the metric. The only negative consequence is that the grids are more closely spaced than needed, which reduces the efficiency of the search.

If we include the additional orbital parameters, the semicoherent mismatch (2.37) can still be written in metric form,

$$\bar{m}(\boldsymbol{\lambda}, \boldsymbol{\lambda}_{\text{psr}}) \approx n^2 \bar{g}_{ab} d\lambda^a d\lambda^b + \mathcal{O}(d\lambda^3). \quad (2.51)$$

However, the assumptions made previously in Eq. (2.38) to calculate this only hold for the ‘‘isolated pulsar’’ parameter space coordinates $\boldsymbol{\lambda}_{\text{iso}} = \{f, \dot{f}, n_x, n_y\}$. They do not hold for the additional orbital parameters $\boldsymbol{\lambda}_{\text{orb}} = \{\Omega_{\text{orb}}, x, t_{\text{asc}}\}$.

If $\lambda^a \in \boldsymbol{\lambda}_{\text{orb}}$ is an orbital parameter and $P_{\text{orb}} \ll T_{\text{coh}}$ (typical coherence time $T_{\text{coh}} \approx 24$ d), the approximations

$$T_{\text{coh}} \langle \partial_a \Phi \rangle(t) \approx 0, \quad (2.52a)$$

$$T_{\text{coh}} \langle \partial_a \partial_b \Phi \rangle(t) \approx 0, \quad (2.52b)$$

$$T_{\text{obs}} \left\langle T_{\text{coh}} \langle \partial_a \Phi \partial_b \Phi \rangle(t') \right\rangle(t_0) \approx T_{\text{obs}} \langle \partial_a \Phi \partial_b \Phi \rangle(t_0) \quad (2.52c)$$

are valid. By this, we mean that the ratio of the resulting metric to the correct metric is $1 + \mathcal{O}(P_{\text{orb}}/T_{\text{coh}})$.

With these assumptions the semicoherent metric \bar{g}_{ab} is composed of three types of components. For the first type, the parameters $\lambda^a, \lambda^b \in \boldsymbol{\lambda}_{\text{orb}}$ are orbital. For these components,

$$\bar{g}_{ab} = T_{\text{obs}} \langle \partial_a \Phi \partial_b \Phi \rangle (t_0) = g_{ab}, \quad (2.53)$$

giving the coherent result from Eq. (2.50).

For the second type, the parameters $\lambda^a, \lambda^b \in \boldsymbol{\lambda}_{\text{iso}}$ are isolated. For these components

$$\begin{aligned} \bar{g}_{ab} &= \frac{1}{2} T_{\text{obs}} \left\langle T_{\text{coh}} \langle \partial_a \Phi \partial_b \Phi \rangle (t') - T_{\text{coh}} \langle \partial_a \Phi \rangle (t') T_{\text{coh}} \langle \partial_b \Phi \rangle (t') \right\rangle (t_0) \\ &= \frac{1}{2} T_{\text{obs}} \langle \tilde{g}_{ab}(t') \rangle (t_0), \end{aligned} \quad (2.54)$$

which is the semicoherent result found in Eq. (2.42).

For the third type, one of a or b is in $\boldsymbol{\lambda}_{\text{orb}}$ and the other is in $\boldsymbol{\lambda}_{\text{iso}}$. One obtains the same equation as for the second type. This vanishes by virtue of Eq. (2.52) and because $T_{\text{coh}} \langle \partial_a \Phi \partial_b \Phi \rangle (t')$ is of order $\mathcal{O}(P_{\text{orb}}/T_{\text{coh}})$.

In short, the nonvanishing semicoherent metric components reduce to earlier results. For the orbital parameters, they are the same as the coherent metric components. For the isolated (spin and celestial) parameters, they are the same as the semicoherent metric components for an isolated pulsar. To reiterate, the nonvanishing semicoherent orbital metric components are

$$\begin{aligned} \bar{g}_{xx} &= 2\pi^2 f^2, \\ \bar{g}_{\Omega_{\text{orb}}\Omega_{\text{orb}}} &= \frac{1}{6} \pi^2 f^2 x^2 T_{\text{obs}}^2 \left(1 + 12 \frac{(t_0 - t_{\text{asc}})^2}{T_{\text{obs}}^2} \right), \\ \bar{g}_{t_{\text{asc}}t_{\text{asc}}} &= 2\pi^2 f^2 x^2 \Omega_{\text{orb}}^2, \\ \bar{g}_{\Omega_{\text{orb}}t_{\text{asc}}} &= -2\pi^2 f^2 x^2 \Omega_{\text{orb}} (t_0 - t_{\text{asc}}). \end{aligned} \quad (2.55)$$

As before, for epoch of ascending node close to the middle of the dataset, i.e. $t_{\text{asc}} \approx t_0$, the semicoherent metric is diagonal.

At the end of Section 2.2.5.3, we discussed the relative densities of the coherent and semicoherent grids for isolated sources. Now we have added three additional (orbital) dimensions to the parameter space. Because the metric factors into a product of a metric on the orbital parameters and a metric on the isolated parameters, the grid may also be constructed as a product of the grids on the corresponding subspaces. For the isolated parameters, the ratio between the density of the coherent grid and the semicoherent grid is the same as for the search for isolated pulsars. For the orbital parameters, the number of grid points needed is the same as in the coherent case. Hence, the ratio of grid densities is the same as in Eq. (2.43).

In Figure 2.1 the mismatch and its coherent metric approximation are compared for small parameter offsets, for a realistic simulated pulsar. The corresponding plot for the semicoherent mismatch looks very similar but has different f - and \dot{f} -scales. The mismatch and its metric approximation agree well for mismatch $m \leq 0.4$. This is a typical value for a search: in Appendix 2.7.2, we show that maximum sensitivity for a given computing resource is obtained for an average mismatch $\hat{m} = 0.383$ (see Table 2.3).




Figure 2.1: Comparison of the coherent metric approximation to the actual mismatch, for parameters of a simulated circular orbit binary pulsar in 4FGL J1653.6–0158. Blue contours show the actual mismatch and red contours the metric approximation, at $m = 0.2$ and 0.4 . As is generally the case [121], the metric contours are conservative and lie inside the actual mismatch contours.

The celestial parameters are not shown in Figure 2.1; for spider pulsars they are usually known to high precision from optical observations [e.g., from the Gaia DR2 Catalog; 122], so no grid is required. For other pulsars where the sky position is less constrained, a grid may be needed.

The search ranges for the orbital parameters are very large, and without further knowledge a partially informed search is not possible. On the other hand, some searches are possible if the pulsar’s companion is visible in optical/X-ray observations, which constrains the search parameters. In the next section, we discuss a gamma-ray pulsar search design for 4FGL J1653.6–0158, which is thought to be an MSP in a circular orbit binary [67, 111].

2.3.2 Search design for circular binary

This section shows how to reduce the binary pulsar search parameter space by exploiting orbital constraints from the companions.

We use the gamma-ray source 4FGL J1653.6–0158, which is predicted to be a spider pulsar [111, 112], as an example. In previous LAT source catalogs the gamma-ray source is named 3FGL J1653.6–0158 and 2FGL J1653.6–0159. It was ranked second in Saz Parkinson et al.’s list (published 2016, [67]) of the most sig-

nificant *Fermi*-LAT Third Source Catalog (3FGL) unassociated sources predicted to be pulsars. The list also classifies it as a likely MSP. The gamma-ray source 4FGL J1653.6–0158 shows typical pulsar properties: a time-stable photon flux and a spectrum described by an exponential cutoff power law.

The search ranges in spin frequency f and spin-down parameter \dot{f} are guided by the known pulsar population and computational constraints. The search range is divided into YPs, with lower frequencies ($f < 44$ Hz), and MSPs, with higher frequencies ($44 \text{ Hz} < f < 1500 \text{ Hz}$)⁷. Correspondingly, the spin-down lies between 0 and $-10^{-10} \text{ Hz s}^{-1}$, for YPs and between 0 and $-10^{-13} \text{ Hz s}^{-1}$ for MSPs.

The constraints for f and \dot{f} define a region in parameter space that has to be searched. The frequency dimension can be efficiently scanned using the FFT algorithm [120] as described by Pletsch and Clark [80], and Clark et al. [78, 79] for isolated pulsars. The \dot{f} -dimension can be covered by a uniformly spaced lattice. Special treatment for these parameters is possible: since their metric components are independent of the other parameters, so is the spacing.

In practice, the FFTs are computed in frequency intervals of bandwidth $f_{\text{BW}} = 8$ Hz. These have $f_{\text{BW}}T_{\text{coh}}$ frequency grid points, with frequency spacing $1/T_{\text{coh}}$. In the semicoherent stage, for two points separated by half the grid spacing, this gives a worst-case metric mismatch $m = \pi^2/24 \approx 0.411$. (As discussed in Appendix 2.7.2 following Eq. (2.106), this can be reduced by interpolation to a worst-case value of $m = 0.14$, at no significant cost.) Thus, for one f_{BW} interval, the computing cost is the product of the cost of a single FFT multiplied by the number of parameter space grid points in the other dimensions.

The sky position is tightly constrained because a likely optical and X-ray counterpart with significant light-curve modulation was found [111, 112, 123] and proposed to be an irradiated pulsar companion. At the time, the best estimate for the position of the likely optical counterpart was from the USNO-B1.0 Catalog [124]. Using this instead of the 3FGL position makes it possible to search 3σ ranges of the sky parameters with only one semicoherent sky grid point. At high frequencies extra sky grid points are needed only in the follow-up stages. The computing costs of these are negligible compared to the semicoherent stage. The same optical source can now be identified in the Gaia DR2 Catalog [122]; see Table 2.1. For this, the uncertainty in sky position is small enough that even at $f = 1.5 \text{ kHz}$ no extra sky points are needed.

The orbital parameters Ω_{orb} and t_{asc} are directly constrained by Romani, Filippenko, and Cenko [111] using optical observations of the companion. As shown in Table 2.1, they found a significant modulation at a period of $P_{\text{orb}} = 0.05194469 \pm 1.0 \times 10^{-7} \text{ d}$, with epoch of ascending node $t_{\text{asc}} = \text{MJD } 56513.48078 \pm 5.2 \times 10^{-4}$.

Additional observations allow the third orbital parameter, the projected semi-major axis of the pulsar $x = a_1 \sin i/c$ (in units of light travel time), to be constrained. Here we denote the neutron star with subscript “1” and the companion with subscript “2”. Measurements of the companion’s velocity amplitude $K_2 = 666.9 \pm 7.5 \text{ km s}^{-1}$, together with the orbital period, imply that the pulsar mass function has the value

$$f(M_1, M_2) = \frac{P_{\text{orb}} K_2^3}{2\pi G} = \frac{M_1 \sin^3 i}{(1+q)^2} = 1.60 \pm 0.05 M_{\odot}, \quad (2.56)$$

⁷The high-frequency limit is around the second harmonic of the fastest known pulsar.

Table 2.1: Parameters and constraints for 4FGL J1653.6–0158

Parameter	Value
Range of observational data (MJD)	54682 – 58300
Reference epoch (MJD)	56500
Initial companion location from USNO-B1.0 catalog	
R.A., α (J2000.0)	16 ^h 53 ^m 38 ^s .07(10)
Decl., δ (J2000.0)	−01°58′36″.7(2)
Precise companion location from Gaia catalog	
R.A., α (J2000.0)	16 ^h 53 ^m 38 ^s .05381(5)
Decl., δ (J2000.0)	−01°58′36″.8930(5)
Constraints from probable counterpart [111]	
Ascending node epoch, t_{asc} (MJD)	$56513.48078 \pm 5.2 \times 10^{-4}$
Companion velocity, K_2 (km s ^{−1})	666.9 ± 7.5
Orbital period, P_{orb} (d)	$0.05194469 \pm 1.0 \times 10^{-7}$
equivalent to	
Orbital frequency, Ω_{orb} (10 ^{−3} Hz)	$1.3999901 \pm 2.7 \times 10^{-6}$
Derived search range	
Projected semimajor axis ^a , x (s)	0 – 0.2

Notes. — The JPL DE405 solar system ephemeris has been used, and times refer to TDB.

^a Assuming a mass ratio of $q < 0.25$; see text following Eq. (2.56).

where the mass ratio is $q = M_2/M_1$. This implies that the neutron star has mass $M_1 > 1.60 \pm 0.05 M_\odot$. Since redback companions often have masses $M_2 \lesssim 0.4 M_\odot$ [53, 54], and black widow companions are even lighter, we assume $q < 0.25$. The extremely short orbital period supports this, since evolutionary models would suggest a black widow companion [125]. From Eq. (2.56), a mass ratio of $q = 0.25$ allows neutron star masses up to $2.5 M_\odot$ for $i = 90^\circ$. (This is reassuringly conservative, since the most massive known neutron star [126] has mass $2.14 M_\odot$.) Combining the mass function with Kepler’s third law $(a_1 + a_2)^3 = G(M_1 + M_2)(P_{\text{orb}}/2\pi)^2$ and the center-of-mass definition $a_1 M_1 = a_2 M_2$ gives

$$x = \frac{q K_2 P_{\text{orb}}}{2\pi c}. \quad (2.57)$$

The upper limit for q then implies an upper limit $x \lesssim 0.2$ s.

It is challenging to build a search grid that covers the three-dimensional orbital parameter space with as few points as possible. This is because (as can be seen from the metric) the orbital parameter space is not flat, so a constant-spacing lattice is not optimal. A solution to this is presented by Fehrmann and Pletsch [127], starting with “stochastic search grids” [128, 129]. A stochastic grid is built by placing grid points with a random distribution that follows the expected distribution of metric distances, while ensuring a preset minimum distance between them. The resulting grid is then optimized by nudging grid points toward regions where neighboring grid points have higher-than-average separation. The resulting search grid is efficient

and has a *well-behaved* mismatch distribution, which simplifies the S/N distribution in the absence of signals.

The minimum number of grid points needed to cover the orbital parameter search space at mismatch m can be estimated from the proper 3-volume

$$N_{\text{orb}} \approx m^{-3/2} \int \sqrt{\det \bar{g}} \, d\boldsymbol{\lambda}_{\text{orb}}. \quad (2.58)$$

Here the integral is over the relevant range of orbital parameter space, g denotes the orbital metric from Eq. (2.50), and numerical factors of order unity related to the efficiency (technically “thickness”; see e.g. [157]⁸) of the grid lattice have been dropped. To understand how this depends on parameters, note that the integral is proportional to

$$N_{\text{orb}} \propto f^3 T_{\text{obs}} (x_{\text{max}}^3 - x_{\text{min}}^3) \Omega_{\text{orb}} \Delta \Omega_{\text{orb}} \Delta t_{\text{asc}}, \quad (2.59)$$

where the search range for x is $[x_{\text{min}}, x_{\text{max}}]$. $\Delta \Omega_{\text{orb}}$ and Δt_{asc} are the search ranges around the values of Ω_{orb} and t_{asc} estimated from the optical modeling. Furthermore, we make the assumption that $\Delta \Omega_{\text{orb}} \ll \Omega_{\text{orb}}$. The strong dependency of N_{orb} on x_{max} and f means that searches for YPs (smaller f) in tight binary orbits (smaller x_{max}) are computationally much cheaper than searches for MSPs in wide orbits. The latter are only possible if the orbital constraints are very narrow.

If the parameter space is small in a particular direction, this reduces the effective dimension of the parameter space and changes the formulae above. For example, denote the range of x by $[x_{\text{min}}, x_{\text{max}}]$. Now consider the case where $\Delta x = x_{\text{max}} - x_{\text{min}}$ is small enough that $g_{xx} \Delta x^2 \ll m$. Then, only a single grid point is needed in the x -direction, and Eq. (2.58) must be replaced with a two-dimensional integral, and the exponent on m must be replaced with -1 . Since the orbital metric components in Eq. (2.55) depend on the parameters, for example, $g_{xx} = 2\pi^2 f^2$, this reduction in dimension can take place for certain ranges of parameters (here small frequency f) and not for others.

We can estimate the computing cost of a search for 4FGL J1653.6–0158 by computing the number of grid points in parameter space. We take $f \in [0, 44]$ Hz and $\dot{f} \in [-10^{-10}, 0]$ Hz s⁻¹ for the YP search and $f \in [44, 1500]$ Hz and $\dot{f} \in [-10^{-13}, 0]$ Hz s⁻¹ for the MSP search from early in this section. The remaining parameter space search ranges are taken from Table 2.1 (no grid is needed over sky location). The frequency range is gridded in intervals of bandwidth $f_{\text{BW}} = 8$ Hz as discussed earlier in this section. The total computing cost is obtained by multiplying the cost of one FFT, the number of \dot{f} grid points, and the number of orbital grid points (which depends on the f interval) and then summing over the f intervals. Since the orbital grid depends on frequency, a new search grid is constructed for each frequency interval, using the metric at the maximum frequency of that interval.

A convenient way to express the computing cost is in terms of search duration on *Einstein@Home*, where we assume that the project provides 25,000 GPU-hr/week. This is shown in Figure 2.2 as a function of the maximum frequency searched. Searching up to $f = 1500$ Hz requires less than 80 d. Note that the search cost in

⁸In the published paper, this wrongly points to Appendix 2.7.2 instead, confusing the thickness with the geometrical factor linking the average to the maximum mismatch. In general, the thickness is the inverse ratio between the multi-dimensional volume closest to one grid point and the volume of a sphere centered on the grid point just covering the aforementioned volume.




Figure 2.2: Predicted days on *Einstein@Home* needed to search 4FGL J1653.6–0158, assuming a circular orbit. The left green curve shows the cumulative duration of a YP search from 0 Hz up to maximum frequency f . The right green curve shows the cumulative duration of an MSP search from 44 Hz up to maximum frequency f . Their slopes are $\propto f^4$ because they are an integral over the number of orbital templates in Eq. (2.59). The blue curve shows the sum: the cumulative duration of a combined YP and MSP search.

one frequency step is proportional to the number of orbital grid points. To search 3σ ranges in t_{asc} and Ω_{orb} within a reasonable amount of time, either the maximum f or x needs to be reduced.

We can also give a general estimate for the MSP search duration. Since the semimajor axis is typically not well constrained, we assume $x_{\text{min}} = 0$. We evaluate Eq. (2.59), using Kepler’s third law to replace x_{max} with the corresponding maximum searched mass ratio q_{max} , obtaining

$$N_{\text{orb}} \propto \frac{T_{\text{obs}} G M_1}{4\pi^2 c^3} f^3 \frac{q_{\text{max}}^3}{(1 + q_{\text{max}})^2} \left(\frac{\Delta\Omega_{\text{orb}}}{\Omega_{\text{orb}}} \right) \Delta t_{\text{asc}}, \quad (2.60)$$

where M_1 is the neutron star mass. As before, we assume $\dot{f} \in [-10^{-13}, 0] \text{ Hz s}^{-1}$ for an MSP search. The search duration up to a maximum frequency f_{max} is then

$$A \left(\frac{B(q_{\text{max}})}{0.01} \right) \left(\frac{f_{\text{max}}}{1 \text{ kHz}} \right)^4 \left(\frac{\Delta P_{\text{orb}}/P_{\text{orb}}}{10^{-6}} \right) \left(\frac{\Delta t_{\text{asc}}}{1 \text{ min}} \right), \quad (2.61)$$

where the dimensionless parenthetical factors are of order unity for typical systems of interest, and

$$B(q_{\text{max}}) = \frac{q_{\text{max}}^3}{(1 + q_{\text{max}})^2}. \quad (2.62)$$

For redbacks (typically: $q < 0.3$) one has $B(q) < 0.02$, whereas for black widows ($q < 0.08$) one has $B(q) < 4 \times 10^{-4}$. The time A depends on the details of the

search and the available computing resources. A typical *Einstein@Home* search as described in this section has $A \sim 10$ d.

In summary, this section has shown how the circular orbit binary pulsar search for 4FGL J1653.6–0158 can be carried out. It is computationally expensive, but by exploiting the orbital constraints, it is feasible, even for high MSP frequencies. In practice, a search would start at low frequencies, gradually working up to 1.5 kHz. To further reduce cost, the search should be stopped if a pulsar is found.

While here we have considered one specific example, these methods are more broadly applicable. With them, circular orbit binary pulsar searches are practical if there are good orbital constraints from optically visible companion stars and if the pulsar’s projected semimajor axis is not too large.

2.4 Search Method: Eccentric Binary Orbits

For pulsars in eccentric binary orbits, the photon arrival times have to be corrected for the line-of-sight motion $r_{z,\text{ell}}(t)$, which is the projection of the eccentric orbit in the line-of-sight direction. In analogy with Eq. (2.47), we can approximate the photon emission time at the pulsar as

$$t_{\text{psr}} \approx t + r_{z,\text{sky}}(t) - r_{z,\text{ell}}(t) \quad (2.63)$$

up to $\mathcal{O}(x\Omega_{\text{orb}})$. Compared with the circular case, two extra parameters are needed to describe the projected line-of-sight motion, $r_{z,\text{ell}}(t)$. For now, we take these to be the orbital eccentricity e and the angle ω between the ascending node and the pericenter.

We note that the approximation to $\mathcal{O}(x\Omega_{\text{orb}})$ is sufficient for the elliptical example source considered in this paper. If the value of x were larger, a higher-order approximation in x would also be required [96].

YPs with main-sequence stars as companions can have very eccentric orbits. For small orbits the pulsars tidally deform the companion, which dissipates energy. This tidally locks the companion, so that the same side of the companion faces the pulsar and over time circularizes the orbit [130]. This explains why old, spun-up MSPs are usually found in binaries with small or unobservable eccentricity. Only a few exceptions are known [131].

If the energy loss in a spider system is small for each orbit, the pulsar moves around a smaller ellipse and the companion around a larger ellipse. The fixed center of mass is a focus of both ellipses, and the separation vector between pulsar and companion also traces an ellipse.

The line-of-sight variation due to the elliptical motion, $r_{z,\text{ell}}(t)$, was derived by Blandford and Teukolsky [132] and can be written as

$$r_{z,\text{BT}}(t) = x \left[\sin \omega (\cos E - e) + \sqrt{1 - e^2} \cos \omega \sin E \right]. \quad (2.64)$$

In this formula the label “ell” is replaced by “BT” to denote that this is the Blandford and Teukolsky model.

The eccentric anomaly E is a parameter along the pulsar path that increases with time. If ψ is the angular position of the pulsar measured from the center of the ellipse, then $\tan \psi = (1 - e^2)^{1/2} \tan E$. Equivalently, project the pulsar’s position

parallel to the semiminor axis, onto a circle whose radius is the semimajor axis, and whose center is the center of the ellipse. Then, E is the angular position of that projected point on the circle. E obeys Kepler's equation

$$M = E - e \sin E, \quad (2.65)$$

where M is the mean anomaly. This is a linear function

$$M = \Omega_{\text{orb}}(t - T_0), \quad (2.66)$$

where $T_0 = t_{\text{asc}} + \omega/\Omega_{\text{orb}}$ is the epoch of pericenter passage.

Unfortunately, there are some problems with the BT model and this parameterization. Kepler's equation (2.65) cannot be solved in closed form to find E as a function of t . Furthermore, in small-eccentricity orbits, the pericenter is not well defined and the mismatch arising from offsets in T_0 and ω does not take the simplest possible form. For these reasons, we shift to an uncorrelated set of parameters and Taylor-expand $r_{z,\text{BT}}$ as function of e .

A new set of parameters was suggested by Lange et al. [133]. These are the time of ascending node t_{asc} and two Laplace-Lagrangian parameters ϵ_1 and ϵ_2 defined via

$$t_{\text{asc}} = T_0 - \omega/\Omega_{\text{orb}}, \quad (2.67)$$

$$\epsilon_1 = e \sin \omega, \quad (2.68)$$

$$\epsilon_2 = e \cos \omega. \quad (2.69)$$

The parameters $\{T_0, e, \omega\}$ are given by

$$T_0 = t_{\text{asc}} + \Omega_{\text{orb}}^{-1} \arctan(\epsilon_1/\epsilon_2), \quad (2.70)$$

$$e = (\epsilon_1^2 + \epsilon_2^2)^{1/2}, \quad (2.71)$$

$$\omega = \arctan(\epsilon_1/\epsilon_2). \quad (2.72)$$

With the old parameters, the region of constant mismatch around a grid point is an ellipsoid whose principal directions are not parallel to the $\{T_0, e, \omega\}$ axes. In the next section, we show that with the new parameters the region of constant mismatch is a sphere. This simplifies the code used to optimize grid point locations.

The Rømer delay $r_{z,\text{BT}}$ for the pulsar's motion can be expanded to first order in e . Following convention, we use the label "ELL1" for this linear-in- e model: $r_{z,\text{BT}} = r_{z,\text{ELL1}} + \mathcal{O}(e^2)$. This can be described using the parameters $\{T_0, e, \omega\}$ or the parameters $\{t_{\text{asc}}, \epsilon_1, \epsilon_2\}$ as

$$r_{z,\text{ELL1}}(t) = x \left[\sin(M + \omega) + \frac{e}{2} \sin(2M + \omega) - \frac{3e}{2} \sin \omega \right] \quad (2.73)$$

$$= x \left[\sin \phi + \frac{\epsilon_2}{2} \sin 2\phi - \frac{\epsilon_1}{2} \cos 2\phi - \frac{3}{2} \epsilon_1 \right]. \quad (2.74)$$

We have introduced

$$\phi = \Omega_{\text{orb}}(t - t_{\text{asc}}), \quad (2.75)$$

which is similar to M in Eq. (2.66) but shifted from pericenter to ascending node. (Note that the term $-3e \sin \omega/2 = -3\epsilon_1/2$ is typically dropped, as it is time independent.)

The ELL1 approximation to the BT model can accurately track the pulsar's rotational phase for eccentricities e below some threshold value. In Appendix 2.7.3, we show how this threshold depends on the spin frequency f and semimajor axis x .

Later in the paper, in Section 2.4.2, we design a search for 4FGL J0523.3–2527, which is a gamma-ray source predicted to harbor a redback pulsar in an eccentric orbit. For that case, the ELL1 model is insufficient and a third-order-in- e model is needed. In Appendix 2.7.3, we derive higher-order-in- e approximations to $r_{z,\text{BT}}$, and demonstrate how they improve the match (decrease the mismatch) to the BT model.

2.4.1 Parameter space metrics

In this section, we calculate the coherent and semicoherent parameter space metric for the ELL1 model. Compared to the circular case, the parameter space has two extra dimensions.

Since the ELL1 model differs at first order in e from the circular model, the coherent metric also differs at first order. However, for the $\{f, \dot{f}, n_x, n_y, \Omega_{\text{orb}}, x\}$ metric components, the first-order terms are of $\mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})$ and can be neglected; the dominant difference is second order in e . Thus, the coherent metric components given in Eqs. (2.28) and (2.50) remain valid to first order in e .

For the ELL1 model in Eq. (2.73), the dominant components for the parameters $\{T_0, e, \omega\}$ are

$$\begin{aligned} g_{T_0 T_0} &= 2\pi^2 f^2 x^2 \Omega_{\text{orb}}^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})], \\ g_{ee} &= \frac{1}{2}\pi^2 f^2 x^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})], \\ g_{\omega\omega} &= 2\pi^2 f^2 x^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})], \\ g_{T_0\omega} &= -2\pi^2 f^2 x^2 \Omega_{\text{orb}} [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})]. \end{aligned} \tag{2.76}$$

Note that the off-diagonal component $g_{T_0\omega}$ does not vanish. As described in the previous section, this complicates the form of the mismatch.

We now change to the parameters $\{t_{\text{asc}}, \epsilon_1, \epsilon_2\}$, for which it is convenient to use Eq. (2.74). For these, the diagonal components are

$$\begin{aligned} g_{t_{\text{asc}} t_{\text{asc}}} &= 2\pi^2 f^2 x^2 \Omega_{\text{orb}}^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})], \\ g_{\epsilon_1 \epsilon_1} &= \frac{1}{2}\pi^2 f^2 x^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})], \\ g_{\epsilon_2 \epsilon_2} &= \frac{1}{2}\pi^2 f^2 x^2 [1 + \mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})]. \end{aligned} \tag{2.77}$$

These diagonal metric components are of $\mathcal{O}(e^0)$. The terms that are linear in e are of $\mathcal{O}(1/\Omega_{\text{orb}}T_{\text{obs}})$, and can be neglected. Thus, the dominant diagonal e -dependent terms are of $\mathcal{O}(e^2)$. However, there are off-diagonal terms of $\mathcal{O}(e^1)$.

For small eccentricities e , the dominant metric components are given above. For

completeness, we list the $\mathcal{O}(e^1)$ -corrections, which are all off-diagonal:

$$\begin{aligned}
g_{x\epsilon_1} &= \frac{1}{2}\pi^2 f^2 x \epsilon_1 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})], \\
g_{x\epsilon_2} &= \frac{1}{2}\pi^2 f^2 x \epsilon_2 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})], \\
g_{\Omega_{\text{orb}}\epsilon_1} &= -\pi^2 f^2 x^2 \epsilon_2 T_{\text{obs}} \left[\frac{(t_0 - t_{\text{asc}})}{T_{\text{obs}}} + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}}) \right], \\
g_{\Omega_{\text{orb}}\epsilon_2} &= \pi^2 f^2 x^2 \epsilon_1 T_{\text{obs}} \left[\frac{(t_0 - t_{\text{asc}})}{T_{\text{obs}}} + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}}) \right], \\
g_{t_{\text{asc}}\epsilon_1} &= \pi^2 f^2 x^2 \Omega_{\text{orb}} \epsilon_2 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})], \\
g_{t_{\text{asc}}\epsilon_2} &= -\pi^2 f^2 x^2 \Omega_{\text{orb}} \epsilon_1 [1 + \mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})].
\end{aligned} \tag{2.78}$$

The remaining off-diagonal components of the orbital metric are of $\mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})$.

These metric components have been found to be a good approximation even for higher eccentricities where the ELL1 model is not sufficient to track the rotational phase in a search and higher-order models need to be used. This might be because many of the linear-in- e terms vanish from the metric.

The semicoherent metric components are very similar to the coherent ones. The components associated with the noneccentric parameters $\{f, \dot{f}, n_x, n_y, \Omega_{\text{orb}}, x\}$, calculated in the circular case in Eqs. (2.42) and (2.55), remain valid; they have only second-order corrections in e . For the remaining orbital parameters $\{t_{\text{asc}}, \epsilon_1, \epsilon_2\}$ the semicoherent metric components are the same as in the coherent case (this follows from Eq. (2.53)). Thus, the diagonal components for $\{t_{\text{asc}}, \epsilon_1, \epsilon_2\}$ are

$$\begin{aligned}
\bar{g}_{t_{\text{asc}}t_{\text{asc}}} &= 2\pi^2 f^2 x^2 \Omega_{\text{orb}}^2, \\
\bar{g}_{\epsilon_1\epsilon_1} &= \frac{1}{2}\pi^2 f^2 x^2, \\
\bar{g}_{\epsilon_2\epsilon_2} &= \frac{1}{2}\pi^2 f^2 x^2,
\end{aligned} \tag{2.79}$$

where we omit terms of $\mathcal{O}(1/\Omega_{\text{orb}} T_{\text{obs}})$. Thus, the semicoherent metric for the ELL1 model simply adds the components above to the semicoherent metric for the circular model.

In Figure 2.3, the mismatch and its coherent metric approximation are compared for small parameter offsets, for a realistic simulated pulsar. Apart from different f - and \dot{f} -scales, the corresponding plot for the semicoherent mismatch looks very similar. The mismatch agrees well with its metric approximation for mismatch $m \leq 0.5$, which is typical: in Appendix 2.7.2, we show that the highest sensitivity at given computing cost for an elliptical search is obtained with an average mismatch $\hat{m} = 0.471$ (see Table 2.3).

The sky position parameters $\{n_x, n_y\}$ are not shown in Figure 2.3 because we assume that for spider pulsars they are known to high precision from optical observations.

A partially informed search for binary pulsars in elliptic orbits, without exact information about the sky position and constraints on the orbital parameters, is computationally impossible. There are too many parameter space dimensions — even for circular orbits with reasonable parameter ranges, the grid has too many points. To make a search possible, one needs tight constraints derived from optical/X-ray




Figure 2.3: Comparison of the coherent metric approximation to the actual mismatch, for parameters of a simulated eccentric orbit binary pulsar in 4FGL J0523.3–2527. Blue contours show the actual mismatch and red contours the metric approximation, at $m = 0.25$ and 0.5 .

observations of the pulsar’s companion star. In the next section, we will discuss constraints and the search design for the probable eccentric orbit binary gamma-ray pulsar in 4FGL J0523.3–2527 [55, 67].

2.4.2 Search design for low-eccentricity binary

In this section, we discuss how to reduce the search parameter space using orbital constraints for the gamma-ray source 4FGL J0523.3–2527, presumed to be a pulsar in an eccentric binary orbit. The source is named 3FGL J0523.3–2528, 2FGL J0523.3–2530, or 1FGL J0523.5–2529 in previous LAT source catalogs. This is similar to the circular example of Section 2.3.2.

The gamma-ray source itself was investigated by Saz Parkinson et al. [67] and ranked ninth highest in a list of most significant 3FGL unassociated sources predicted to be pulsars. It shows typical pulsar-like properties: the photon flux is stable over time, and the spectrum is fit by an exponential cutoff power law. The source is not in the Galactic disk, which increases the odds that it hosts an MSP.

Earlier optical observations identified a likely companion and indicate an orbit with small, but not negligible, eccentricity of $e = 0.04$ [55]. In contrast to the previous paragraph, this suggests that the pulsar is a YP, because binary MSPs

tend to have rather circular orbits [130].

The frequency and spin-down search ranges are chosen following the logic of the previous search design (Section 2.3.2). For YPs we search $f \in [0, 44]$ Hz and $\dot{f} \in [-10^{-10}, 0]$ Hz/s. For MSPs we search $f \in [44, 1500]$ Hz and $\dot{f} \in [-10^{-13}, 0]$ Hz/s. The f -dimension is efficiently searched using FFTs with bandwidth $f_{\text{BW}} = 8$ Hz, and the \dot{f} -dimension is covered by a uniformly spaced lattice.

The sky position search range of the probable pulsar within 4FGL J0523.3–2527 is tightly constrained from the X-ray and optical observations of the likely companion discussed above [55]. At the time, the best estimate for the optical position was from the USNO-B1.0 Catalog [124]. It is now also identified in the Gaia DR2 Catalog [122], whose pointing is so precise (see Table 2.2) that even at $f = 1.5$ kHz no search over sky position is required.

The orbital parameter search ranges shown in Table 2.2 come from the [55] analysis of the photometric and spectroscopic optical data. The orbital period and eccentricity parameters are constrained by the periodic optical flux modulation. They assume that this arises from viewing a tidally locked and deformed (ellipsoidal) companion at different aspect angles. Hence, the orbital period is twice the observed modulation period. (Another possible explanation for the modulation would be irradiation, but spectroscopic data do not show the orbital-phase-dependent temperature change that would be expected.) The orbital period is constrained to $P_{\text{orb}} = 0.688134 \pm 0.000028$ d at epoch of superior conjunction $T_{0.5} = 56577.14636 \pm 0.0037$ MJD. The eccentric parameters $\{e, \omega\}$ fall in the ranges $e = 0.040 \pm 0.006$ and $\omega = 214 \pm 10$ deg.

The semimajor axis x is constrained using Eq. (2.57). This is similar to our previous example in Section 2.3.2, but requires fewer assumptions because the mass ratio $q = M_2/M_1$ is directly bounded from the observations. To do this, [55] estimate the rotational velocity of the companion’s Roche lobe from high-quality optical spectra. Combined with the companion’s radial velocity $K_2 = 190.3 \pm 1.1$ km s $^{-1}$, this constrains the mass ratio $q = 0.61 \pm 0.06$. Returning to Eq. (2.57), this gives $x = 3.66 \pm 0.38$.

The parameters $\{e, \omega\}$ can be converted directly to the quantities $\{\epsilon_1, \epsilon_2\}$ needed for our search, using Eqs. (2.68) and (2.69).

For our search, we also need the epoch of ascending node t_{asc} . However, the results of Strader et al. [55] are given in terms of the epoch of superior conjunction $T_{0.5}$. For circular orbits, $T_{0.5}$ and t_{asc} differ by $P_{\text{orb}}/4$, but for eccentric orbits the relation is more complicated. To second order in e , it is

$$t_{\text{asc}} = T_{0.5} + P_{\text{orb}} \left(\frac{1}{4} - \frac{\epsilon_2}{\pi} - \frac{3\epsilon_1\epsilon_2}{4\pi} + \mathcal{O}(e^3) \right). \quad (2.80)$$

For 4FGL J0523.3–2527 with $e = 0.04$, this $\mathcal{O}(e^2)$ approximation is more accurate than the uncertainties in the measured quantities on the right-hand side. (Higher-order approximations in e would be required for pulsars in binary orbits with larger eccentricities or longer orbital periods.) The resulting t_{asc} is given in Table 2.2.

A search for a pulsar in an eccentric orbit is very similar to one for a pulsar in a circular orbit. The only differences are that a more general model for the Rømer delay is required to track the pulsar phase, and the orbital grids need to cover five orbital dimensions. While the latter is much more complex, it can be done with the same optimized stochastic search grid construction methods that are used in the circular case.

Table 2.2: Parameters and constraints for 4FGL J0523.3–2527

Parameter	Value
Range of observational data (MJD)	54682 – 58300
Reference epoch (MJD)	56500
Initial companion location from USNO-B1.0 catalog	
R.A., α (J2000.0)	05 ^h 23 ^m 16 ^s .925(4)
Decl., δ (J2000.0)	–25°27′36″.92(6)
Precise companion location from Gaia catalog	
R.A., α (J2000.0)	05 ^h 23 ^m 16 ^s .931203(2)
Decl., δ (J2000.0)	–25°27′37″.12468(4)
Constraints from probable counterpart [55]	
Superior conjunction epoch, $T_{0.5}$ (MJD)	56577.14636 ± 0.0037
Companion velocity, K_2 (km s ^{–1})	190.3 ± 1.1
Mass ratio, $q = M_2/M_1$	0.61 ± 0.06
Eccentricity, e	0.040 ± 0.006
Longitude of pericenter, ω (deg)	214 ± 10
Orbital period, P_{orb} (d)	0.688134 ± 0.000028
equivalent to	
Orbital frequency, Ω_{orb} (Hz)	0.0001056801 ± 4.3 × 10 ^{–9}
Derived search parameters and corresponding uncertainties	
Projected semimajor axis, x (s)	3.66 ± 0.38
Ascending node epoch, t_{asc} (MJD)	56577.32553 ± 0.00567
First Lagrange parameter, ϵ_1	–0.0224 ± 0.0091
Second Lagrange parameter, ϵ_2	–0.0332 ± 0.0089

Notes. — The JPL DE405 solar system ephemeris has been used, and times refer to TDB.

To accurately track the rotational phase of the pulsar requires a higher-order-in- e approximation to $r_{z,\text{BT}}$ than the ELL1 model, unless the eccentricity is very small. Such approximations are computed in Appendix 2.7.3. There, we also determine which order in e is sufficient.

For the case of 4FGL J0523.3–2527, a model of $\mathcal{O}(e^3)$ is sufficient. In Figure 2.6, we show that the rotational phase error is negligible for the constrained parameter ranges given above.

Analogously to Eq. (2.58), the minimum number of grid points for the orbital parameter space can be computed from the proper 5-volume

$$N_{\text{orb}} \approx m^{-5/2} \int \sqrt{\det \bar{g}} d\lambda_{\text{orb}}. \quad (2.81)$$

Here the metric has the five dimensions $\{x, \Omega_{\text{orb}}, t_{\text{asc}}, \epsilon_1, \epsilon_2\}$. This integral is proportional to

$$N_{\text{orb}} \propto f^5 T_{\text{obs}} (x_{\text{max}}^5 - x_{\text{min}}^5) \Omega_{\text{orb}} \Delta \Omega_{\text{orb}} \Delta t_{\text{asc}} \Delta \epsilon_1 \Delta \epsilon_2, \quad (2.82)$$

where $x \in [x_{\text{min}}, x_{\text{max}}]$. $\Delta \Omega_{\text{orb}}$, Δt_{asc} , $\Delta \epsilon_1$, and $\Delta \epsilon_2$ are the search ranges for the corresponding parameters, and we made the assumption that $\Delta \Omega_{\text{orb}} \ll \Omega_{\text{orb}}$. The




Figure 2.4: Predicted days on *Einstein@Home* needed to search 4FGL J0523.3–2527, assuming a circular (green/blue) or elliptical (red/orange) orbit. The left curves show the cumulative duration of a YP search from 0 Hz up to maximum frequency f . The right curves show the cumulative duration of an MSP search from 44 Hz up to maximum frequency f . Their slopes are $\propto f^4$ and $\propto f^6$; they are integrals over the number of orbital templates. The larger slope for the elliptical search arises from the two extra dimensions of search parameter space. The blue and orange curves show the sums: the cumulative duration of a combined YP and MSP search.

number of orbital grid points and subsequently the computing cost depend even more strongly on f and x in an eccentric search than in a circular one.

The computing cost of a search for 4FGL J0523.3–2527 is estimated based on the number of grid points. We assume search ranges in f and \dot{f} as given earlier in this section. The remaining parameter space ranges are given in Table 2.2. The required total computing cost of the search is estimated by multiplying the cost of one FFT by the number of \dot{f} -grid points and the f -dependent number of orbital grid points and then summing over the f intervals.

To exemplify the computing cost of a search for 4FGL J0523.3–2527, we express it in terms of search duration on *Einstein@Home*, assuming that the project provides 25,000 GPU-hr per week. This is shown in Figure 2.4 as a function of the maximum searched frequency. For comparison, we also show the search duration for a circular binary search, i.e. setting $e = 0$ and not searching over $\{\epsilon_1, \epsilon_2\}$. An eccentric MSP search up to 1.5 kHz would take more than 100 million years on *Einstein@Home*, and even a YP search would take more than 100 years. Circular searches for YPs or MSPs up to 400 Hz would take a few hundreds days. Note that the search ranges are still the 1σ ranges, so searches within the 3σ range would be more computing intensive.

In summary, this section has shown how computing intensive a search for 4FGL J0523.3–2527 would be. An eccentric MSP search even to low frequencies ~ 100 Hz is not feasible with the current constraints, and a YP search would be

very expensive. In the optical data, [55] do not see evidence for a “false” eccentricity, but a circular search would be much less computing intensive than an eccentric one. With slightly tighter constraints, searches up to 800 Hz could be feasible.

2.5 Comparison with other Methods

Similar and alternative methods are used to search for binary pulsars in data from radio telescopes and gravitational-wave detectors. In this section, we will review these, compare them to the methods presented here, and discuss their applicability to searches for binary gamma-ray pulsars.

In addition to coming from diverse messengers and frequencies, the data have other key differences. The gamma-ray data are similar to the gravitational-wave data: the length of the data sets is months to years, and the instruments simultaneously detect signals from a substantial fraction of the sky. In contrast, typical radio surveys collect data in stretches of minutes from tiny fractions of the sky. While gamma-ray data consist of discrete photon arrival times, radio and gravitational-wave data are continuous. Therefore, it is not surprising that some pulsation search methods might work for one kind of data but not for the other.

For these other data sources, many methods have been employed by many individuals and groups. Here we are guided by reviews from Lorimer and Kramer [10] for radio search methods and Messenger et al. [134] for gravitational-wave methods. We exclude methods that require data from two detectors.

2.5.1 Acceleration search

Time-domain “acceleration searches” have been very successful in finding new radio pulsars in binaries with orbital periods shorter than a day [see, e.g., 135]. Fourier-domain acceleration searches have also been successfully used to discover binary radio pulsars [see, e.g., 136, 137]. A similar approach to search for continuous gravitational waves is the “polynomial search” [138].

These searches do not use a model that describes periodic orbital motion. Instead, they assume constant acceleration along a straight line [see also 139]. This accurately describes an orbiting system only if the data set is much shorter than one orbital period. Since the LAT data set is more than a decade long, acceleration searches would only find binary gamma-ray pulsars whose orbital periods were decades or longer.

It is straightforward to quantify the range of orbital periods an acceleration search is sensitive to. Assume that the data set is less than $\sim 10\%$ of the orbital period and is near the superior or inferior conjunction, where the velocity is changing linearly with time [139]. An acceleration a along the line of sight (“los”) toward Earth contributes an amount

$$\dot{f}_{\text{los}} = \frac{fa}{c} \quad (2.83)$$

to the observed spin frequency derivative. The maximum acceleration at inferior or superior conjunction is for a circular orbit $a = cx \Omega_{\text{orb}}^2$, and for an eccentric orbit $a = cx \Omega_{\text{orb}}^2 (1 + e)/(1 - e)$. Therefore, searches would be sensitive if the sum of the intrinsic pulsar spin-down and this line-of-sight contribution to the spin-down were within the search range. Since the intrinsic spin-down is usually negative, this is

most likely if the acceleration toward Earth is positive, i.e. if the pulsar is near the superior conjunction.

Current partially informed search surveys for isolated gamma-ray pulsars are a form of acceleration search because they scan over spin-down [79]. For YPs they search down to $\dot{f} = -10^{-9} \text{ Hz s}^{-1}$ and for MSPs down to $\dot{f} = -10^{-13} \text{ Hz s}^{-1}$. In principle, these searches are sensitive to pulsars like the young ($f \approx 7 \text{ Hz}$) binary pulsar PSR J2032+4127, which is in a 45–50 yr orbit around its companion [140]. It was found in an isolated gamma-ray search [73], and only afterward was it discovered to be in a binary system [45]. The orbit is highly eccentric ($e \approx 0.93 - 0.99$) with $x \approx 7,000 - 20,000 \text{ s}$. The maximum spin-down contribution should therefore be of order $|\max\{\dot{f}_{\text{los}}\}| = 10^{-10} \text{ Hz s}^{-1}$. This is in the search range if the pulsar is near superior conjunction during the mission time.

Searches that assume linear acceleration, i.e., that search over constant \dot{f} , are only sensitive to binary pulsars with $P_{\text{orb}} \gtrsim 10T_{\text{obs}}$. To become sensitive to shorter orbital periods, higher-order frequency derivatives must be searched. “Jerk” searches, which include the second-order frequency derivative \ddot{f} , improve the sensitivity for pulsars with orbital periods in the range $P_{\text{orb}} \in [7T_{\text{obs}}, 20T_{\text{obs}}]$ and have been successfully used in a radio pulsar search [137]. Alternatively, the full orbital motion may be taken into account, as in [85].

2.5.2 Stack/slide search

The “stack/slide” method has been used in radio pulsar searches like the Parkes Multibeam Pulsar survey to account for binary motion [141]. This led to the discovery of the double neutron star system PSR J1756–2251 with an orbital period of 7.7 hr [142]. (The words “stack/slide” are used in continuous gravitational-wave searches, not to account for binary pulsar motion but rather to remove the effects of Earth rotation and motion around the SSB [143, 144]. That is also the case for the semicoherent searches we describe in this paper to account for the LAT’s motion around the SSB.)

In a stack/slide search the data set is broken into subsets of length T_{coh} , corresponding to frequency bins of width $\Delta f = 1/T_{\text{coh}}$. T_{coh} is chosen to be small enough that the Doppler modulation induced by motion of the detector around the SSB, or of the pulsar around the binary center of mass, remains within a single bin. For circular binary motion, provided that T_{coh} is a factor of a few smaller than P_{orb} , this implies

$$fx\Omega_{\text{orb}}^2 T_{\text{coh}} < 1/T_{\text{coh}}. \quad (2.84)$$

Each of these subsets is then Fourier transformed. The resulting power spectra are added (stacked) together after the Doppler modulation is compensated by shifting the frequency (slide) in each of the spectra; sources give rise to peaks in the stacked spectra. This technique is only sensitive if the subsets are much shorter than the Doppler modulation period.

This technique is useless for spider gamma-ray pulsars because detection statistics are constructed from the differences of photon arrival times. Spider pulsars have typical orbital periods of $P_{\text{orb}} \lesssim 1 \text{ d}$, so data subsets would have to be shorter than a few hours. Most data subsets would contain no photons. A few would contain one photon. Almost none would contain enough photons to compute the differences of photon arrival times.

Stack/slide could be used for gamma-ray pulsars in orbits where P_{orb} is too small for an acceleration search but is much larger than the $T_{\text{coh}} \approx 24$ d used in this paper. Using Kepler’s third law, the condition of Eq. (2.84) can be written as

$$\frac{GM_1}{c^3} \frac{q^3}{(1+q)^2} f^3 T_{\text{coh}}^6 \Omega_{\text{orb}}^4 < 1, \quad (2.85)$$

where M_1 is the pulsar mass and $q = M_2/M_1$ is the mass ratio. (In fact, this applies provided that $T_{\text{coh}} \lesssim P_{\text{orb}}$.) This shows that with our choice of T_{coh} , stack/slide methods might be able to find gamma-ray pulsars with planetary companions, with orbital periods longer than ~ 1 yr and masses up to $\mathcal{O}(10)$ Earth masses.

2.5.3 Power spectrum search

The basic assumption of a “power spectrum search” is that the data set can be broken into subsets short enough that the observed spin frequency is constant in each one. This is the same assumption as in a stack/slide search. That technique is based on visual inspection and has been used to discover binary radio pulsars [see, e.g., 145].

To carry out the search, power spectra are computed for each subset. The spectra are binned in frequency and plotted with a frequency-versus-time color map. The colors show the power and make it easy to visually identify peaks in the power spectrum. A binary pulsar signal appears as a peak whose frequency varies sinusoidally with time.

The method “TwoSpect” uses a similar method to perform all-sky searches for continuous gravitational waves from sources in binary systems. The visual inspection is replaced by a second Fourier transform [hence the name TwoSpect; 146]. While no continuous gravitational waves have been detected, this technique has been used to put upper limits on continuous gravitational-wave emission from the low-mass X-ray binary Scorpius X-1 [147].

The power spectrum search is not suitable for detecting gamma-ray spider pulsars for the same reasons as the stack/slide method.

2.5.4 Sideband search

“Sideband searches” have found many binary radio pulsars within globular clusters [10]. The method has also been adapted to search for continuous gravitational waves from sources in binary systems [148, 149]. One first carries out a search for isolated systems, as if there were no binary motion, and then looks for a characteristic structure in the results of that isolated search.

If a binary is present, orbital motion produces sidebands around a central peak at the spin frequency of the pulsar [150]. Since the isolated search does not remove the effects of the binary motion, a pulsar’s power is spread over many Fourier bins (also called sidebands). This reduces the sensitivity compared to a matched-filter search.

The method is particularly useful for tight orbit binary pulsars where the orbital period is much smaller than the observation time span, which is the case of interest for spider pulsars. After detecting a signal, the binary parameters can be inferred from the locations and magnitudes of the sidebands and the central peak.

To see how this works, we compute the S/N of the coherent detection statistic P_n for an isolated pulsar template, with parameters $\{\nu, \dot{f}, n_x, n_y, 0, 0, 0\}$, arising from a circular binary pulsar with parameters $\{f, \dot{f}, n_x, n_y, x, f_{\text{orb}}, t_{\text{asc}}\}$, where $f_{\text{orb}} = \Omega_{\text{orb}}/2\pi$. This S/N is given by Eq. (2.17), which depends on the rotational phase difference due to the parameter mismatch:

$$\Delta\Phi(t) = 2\pi(\nu - f)(t - t_{\text{ref}}) + 2\pi f x \sin[2\pi f_{\text{orb}}(t - t_{\text{asc}})]. \quad (2.86)$$

One can think of ν as denoting the pulsar frequency in the isolated search. Our derivation closely follows [150].

To compute the detection statistic P_n , we evaluate Eq. (2.17) with the phase mismatch (2.86). We first reexpress $e^{in\Delta\Phi}$ using the Jacobi-Anger expansion

$$e^{iz \sin \vartheta} = \sum_{m=-\infty}^{\infty} J_m(z) e^{im\vartheta}, \quad (2.87)$$

with $z = 2\pi n f x$ and $\vartheta = 2\pi f_{\text{orb}}(t - t_{\text{asc}})$, where J_m is a Bessel function of the first kind. Multiplying this by $e^{i2\pi n(\nu-f)(t-t_{\text{ref}})}$ gives

$$e^{in\Delta\Phi} = \sum_{m=-\infty}^{\infty} J_m(2\pi n f x) e^{i2\pi[n(\nu-f)+mf_{\text{orb}}](t-t_{\text{asc}})}, \quad (2.88)$$

where, without loss of generality, we have set $t_{\text{ref}} = t_{\text{asc}}$. Since the S/N only depends on the modulus of $e^{in\Delta\Phi}$, we may also set $t_{\text{asc}} = 0$. We assume that there are a large number of photons from the hypothetical pulsar, which have equal weights and arrive at uniformly spaced intervals in time. The double sum $\sum_{j \neq k}$ in Eq. (2.17) may then be replaced by an integral over time, since

$$\begin{aligned} \sum_{j \neq k} e^{in(\Delta\Phi_j - \Delta\Phi_k)} &\approx \left| \sum_j e^{in\Delta\Phi_j} \right|^2 \\ &\approx \left| \frac{N}{T_{\text{obs}}} \sum_{m=-\infty}^{\infty} J_m(2\pi n f x) \int_{-T_{\text{obs}}/2}^{T_{\text{obs}}/2} e^{i2\pi[n(\nu-f)+mf_{\text{orb}}]t} dt \right|^2. \end{aligned} \quad (2.89)$$

On the right-hand side we have included the diagonal $j = k$ term, which is absent on the left-hand side, but is negligible in the limit where the number of photons N is large. The integral over time is

$$\frac{1}{T_{\text{obs}}} \int_{-T_{\text{obs}}/2}^{T_{\text{obs}}/2} e^{i2\pi Ft} dt = \frac{\sin(\pi FT_{\text{obs}})}{\pi FT_{\text{obs}}}, \quad (2.90)$$

with $F = n(\nu - f) + mf_{\text{orb}}$. For observation times that include many orbits, the right-hand side of Eq. (2.90) is unity for $F = 0$ and is negligible otherwise. Thus, the only terms in Eq. (2.89) that survive are those for which $\nu = f - mf_{\text{orb}}/n$. When that is satisfied, we have

$$\sum_{j \neq k} e^{in(\Delta\Phi_j - \Delta\Phi_k)} \approx N^2 J_m^2(2\pi n f x), \quad (2.91)$$

where m is constrained by $F = 0$. Thus, the double sum in Eq. (2.89) vanishes at all frequencies ν except for the ‘‘sideband frequencies’’ $\nu = \nu_m = f - mf_{\text{orb}}/n$, where m takes on all integer values.

We now evaluate the S/N $\theta_{P_n}^2(\nu)$ from Eq. (2.17) by substituting in Eq. (2.91), assuming that the weights w_j are constant. For the reasons just given, $\theta_{P_n}^2(\nu)$ vanishes except at the discrete sideband frequencies $\nu_m = f - mf_{\text{orb}}/n$. We obtain

$$\theta_{P_n}^2(\nu) = \begin{cases} J_m^2(2\pi n f x) \theta_{P_n}^2 & \text{for } \nu = \nu_m, m \in \mathbb{Z}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.92)$$

The quantity $\theta_{P_n}^2$ that appears on the right-hand side is given by Eq. (2.11). It is the S/N that the pulsar would have in an isolated search if the binary motion were absent. It is also the S/N that the pulsar would have in a binary pulsar search at the true signal parameter values.

The structure in frequency space ν is evident from Eq.(2.92). As described by [150], the S/N is spread over equally spaced sidebands around the pulsar frequency f , whose spacing is commensurate with the orbital frequency. The sideband width is $\sim 1/T_{\text{obs}}$, as can be seen from Eq. (2.90).

In comparison with a binary pulsar search, the isolated pulsar search has lost some S/N, since $J_m^2 \leq 1$. To recover some of the lost S/N within the isolated pulsar search, we introduce a new test statistic that sums over the first m_{orb} sidebands around the central pulsar frequency. This cumulative sideband power may be written as

$$B_n(\nu) = \sum_{m=-m_{\text{orb}}}^{m_{\text{orb}}} P_n \left(\nu - \frac{mf_{\text{orb}}}{n} \right), \quad (2.93)$$

with the detection statistic $P_n(\nu)$ appropriate to an isolated pulsar search with parameters $\{\nu, \dot{f}, n_x, n_y, 0, 0, 0\}$. (A test statistic weighing the m th sideband in Eq. (2.93) by $J_m^2(2\pi n f x)$ would be more sensitive, but for simplicity it is not considered here.)

The S/N for the cumulative sideband power B_n is easily calculated. It is defined as

$$\theta_{B_n}^2 = \frac{E_p[B_n] - E_0[B_n]}{\sqrt{E_0[B_n^2] - E_0^2[B_n]}}, \quad (2.94)$$

where p is the pulsed fraction defined in Eq. (2.3). The numerator of Eq. (2.94) can be found from Eq. (2.92), which implies that $E_p[P_n] - E_0[P_n] = \text{Var}_0[P_n] \theta_{P_n}^2 J_m^2(2\pi n f x) = 4\theta_{P_n}^2 J_m^2(2\pi n f x)$. Summing this over m gives the numerator:

$$E_p[B_n] - E_0[B_n] = 2\theta_{P_n}^2 \sum_{m=-m_{\text{orb}}}^{m_{\text{orb}}} J_m^2(2\pi n f x). \quad (2.95)$$

The denominator of Eq. (2.94) is defined in the absence of a signal, with $p = 0$. It is easily calculated if the noise at the different frequencies that contribute to the sum is independent. Since Poisson noise is stationary, these contributing terms will be independent if they are spaced more than one frequency bin apart, where the bins have width $1/nT_{\text{obs}}$. Since the sideband frequencies are separated by f_{orb}/n , these different terms will be independent if there are many orbits in the observation time:

$f_{\text{orb}}T_{\text{obs}} \gg 1$. Each term in the denominator then has variance 4, so the sum yields $E_0[B_n^2] - E_0^2[B_n] = 4(2m_{\text{orb}} + 1)$. Thus, the S/N for B_n is

$$\theta_{B_n}^2 = \frac{\theta_{P_n}^2}{\sqrt{2m_{\text{orb}} + 1}} \sum_{m=-m_{\text{orb}}}^{m_{\text{orb}}} J_m^2(2\pi n f x). \quad (2.96)$$

To maximize this S/N, what is the optimal number of sidebands m_{orb} to include?

As shown by [150], the optimal number of sidebands to include depends on

$$M_{\text{orb}} = [2\pi n f x], \quad (2.97)$$

where square brackets denote ‘‘integer part’’. To see this, consider the sum that appears in Eq. (2.96):

$$\sum_{m=-m_{\text{orb}}}^{m_{\text{orb}}} J_m^2(2\pi n f x). \quad (2.98)$$

For $m_{\text{orb}} < M_{\text{orb}}$ this sum grows (approximately linearly) with increasing m_{orb} . But the addition theorem for Bessel functions ensures that Eq. (2.98) stops growing and approaches unity as soon as m_{orb} exceeds M_{orb} . Since the denominator of the S/N in Eq. (2.96) has a term that grows like $\sqrt{2m_{\text{orb}} + 1}$, the S/N is maximized for $m_{\text{orb}} = M_{\text{orb}}$. For this number of sidebands, one thus obtains

$$\theta_{B_n}^2 \approx \frac{\theta_{P_n}^2}{\sqrt{2M_{\text{orb}} + 1}} \quad (2.99)$$

for the expected S/N of the cumulative sideband power.

The behavior we have just described, considered alongside the definition (2.94) of the S/N, shows the main weakness of sideband searches. The numerator grows (approximately) linearly as we include more sidebands, meaning that we can recover all of the signal power. But, in the absence of a signal, B_n undergoes a random walk as sidebands are included, and so the denominator of Eq. (2.94) (the root-mean-squared of B_n in the absence of a signal) increases as $\sqrt{2M_{\text{orb}} + 1}$. Thus, in comparison with an optimal matched filter, the incoherent summation over sidebands loses a factor of $\sqrt{2M_{\text{orb}} + 1}$ in the S/N. This is explicit in Eq. (2.99) and makes sideband searches ineffective if there are many sidebands, as is often the case. For example, consider the potential circular binary pulsar in 4FGL J1653.6–0158 and the potential eccentric binary pulsar in 4FGL J0523.3–2527 discussed earlier in this paper. Their estimated parameter ranges in f and x give rise to large numbers of sidebands.

This means that sideband searches work best if only a few sidebands are expected, meaning that $2\pi x f$, the total rotational phase arising from the orbital modulation, is small. This is the case for black widow systems, which have very light companions. The small companion mass means that the pulsar orbits very close to the center of mass, so the projected semimajor axis x is extremely small. Note that the modulation can be small even for the high frequencies f typically found for black widows.

Figure 2.5 illustrates this, for example, for the black widow pulsar PSR J1311–3430, which would have been a candidate for a sideband search. The figure shows the expected optimal matched-filter S/N $\theta_{P_1}^2$ required to exceed a threshold

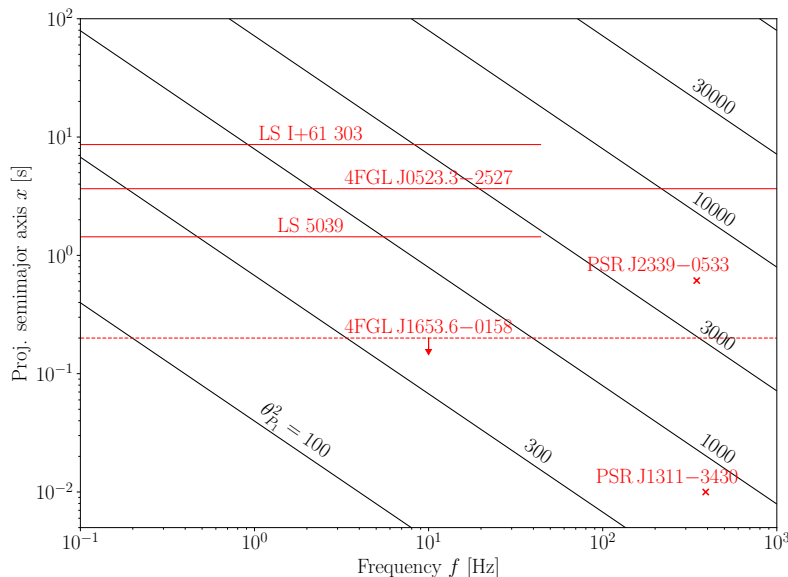


Figure 2.5: Comparison between the expected cumulative sideband S/N $\theta_{B_1}^2$ and the expected optimal matched-filter S/N $\theta_{P_1}^2$. For given frequency f and semimajor axis x , the black contours show the $\theta_{P_1}^2$ required to exceed a threshold $\theta_{B_1}^2 > 100$. The crosses are at the locations of two known pulsars: PSR J1311–3430 and PSR J2339–0533. The red lines show four potential sideband search candidates. For the YP candidates LS 5039 and LS I+61 303, and the spider candidate 4FGL J0523.3–2527, the approximate values for the semimajor axes are known. The dashed line shows the maximum semimajor axis value for 4FGL J1653.6–0158.

in the expected cumulative sideband S/N $\theta_{B_1}^2 > 100$, which is a reasonable threshold for confident detection. From Eq. (2.99), this requires $\theta_{P_1}^2$ to exceed $100\sqrt{2M_{\text{orb}} + 1}$. Hence, M_{orb} is constant on the contour lines, which therefore denote boundaries of constant fx . Since the largest observed $\theta_{P_1}^2$ values for known pulsars are ~ 1000 , the region below and to the left of the contour line corresponding to $\theta_{P_1}^2 = 1000$ might be considered for sideband searches.

Sideband searches within gamma-ray binaries like LS 5039 and LS I+61 303 would also be justified. These systems contain a compact object: a black hole or neutron star. Since both binaries are highly eccentric [$0.3 < e < 0.6$; 151], the compact objects could be YPs. These two candidate pulsars are both displayed in Figure 2.5. This is purely illustrative, since the sideband power B_n defined here is only suitable for circular binary pulsars. Eccentric pulsars will have additional sidebands [150] and thus must have an even higher pulsed fraction to be detectable in a sideband search.

This section has not discussed the implementation of a practical sideband search. We would need some constraints on the parameters f_{orb} and x to hunt for the sidebands. If those are available from optical observations, then the sky position will be known precisely. This in turn would make a fully coherent isolated pulsar search computationally feasible. The resulting test statistics could then be used to construct the sideband search statistic B_n of Eq. (2.93).

2.5.5 Discussion

The methods discussed in this section have little applicability to searches for gamma-ray pulsars in spider systems, which are the main focus of this paper. But they are of interest for other types of binary systems.

Acceleration searches could discover binary pulsars with orbital periods comparable to or longer than our observation time $T_{\text{obs}} \sim 10$ yr. These binaries have pulsars whose companions are very low mass stars or planets, in wide orbits. These pulsars might have been missed by isolated pulsar searches.

Stack/slide and power spectrum methods do not appear suitable for spider gamma-ray pulsar searches. They might potentially detect systems with orbital periods longer than our typical coherence time $T_{\text{coh}} \sim 24$ d and shorter than $T_{\text{obs}} \sim 10$ yr. However, these searches are very expensive computationally.

Sideband searches could be used to hunt for binary pulsars with low spin frequencies or in very close orbits. While these are computationally less expensive than the search methods discussed earlier in this paper, they are also considerably less sensitive.

All of these methods have a domain of applicability. Given prior knowledge and constraints on a specific target, one can investigate these different methods to determine which are feasible and to estimate which one is potentially the most sensitive.

2.6 Conclusions

This work presents computationally efficient methods to detect circular and eccentric orbit binary gamma-ray pulsars. These generalize techniques that have been previously developed to search for isolated pulsars [80].

We have presented all of the elements of this generalization. Physically, the central element is a model that accurately describes the rotational phase of a pulsar over time as would be observed at the solar system barycenter. In comparison with the isolated model, this must also account for the Rømer delay caused by the binary motion. A second key element are semicoherent and coherent test statistics, along with their expected signal-to-noise ratios. The last key element are the metrics for these statistics, which measure the “distance” in parameter space between two different rotational phase models. This metric quantifies the expected fractional loss in signal-to-noise ratio, and enables the construction of efficient parameter space grids for a search.

We have shown how these different elements can be used together to search for gamma-ray pulsars. This is analogous to the isolated pulsar case [80]: the most computationally efficient approach is a multistage search with several semicoherent and coherent stages. The computing cost is proportional to the number of points in the parameter space grid. We compute this from the metric and show how the computing cost depends on the search parameters. This in turn allows the grid spacing to be optimized, achieving the highest possible sensitivity at fixed computing cost. These methods have been very successful in discovering isolated gamma-ray pulsars [43, 75, 77–79].

Currently, a search for binary pulsars without partial information about the sky position and constraints on the orbital parameters is computationally impossible.

Because the parameter space has at least seven of the nine possible dimensions $\{f, \dot{f}, \alpha, \delta, x, P_{\text{orb}}, t_{\text{asc}}, \epsilon_1, \epsilon_2\}$, too many grid points are needed to cover it. However, in some cases, the number of dimensions can be reduced and/or the corresponding search ranges can be tightly constrained by multiwavelength observations.

This paper considers two illustrative examples of this type, drawn from potential spider pulsars. Here, analysis of optical observations constrains the orbital parameters, and we show that searches of reasonable sensitivity (in some cases limited to young pulsars) are feasible. This enables partially informed searches for binary gamma-ray pulsars that were previously not feasible. This is important because these pulsars might be impossible to detect in other wave bands.

The methods of this paper, particularly the metric in parameter space, have applications beyond partially informed searches. There are binary pulsars that are visible in radio, optical, or X-ray, for which gamma-ray pulsations have not yet been found. For recent discoveries, precise determination of their orbital and other parameters is often not possible, since it requires observations spanning several years. The methods here are useful in such cases, to carry out efficient follow-up searches to discover gamma-ray pulsations. This way, within days or weeks after radio pulsations are discovered, the pulsar's parameters can be precisely measured over the > 10 yr of elapsed LAT mission time. This approach led to the discovery of gamma-ray pulsations soon after the radio detection of the 707 Hz black widow pulsar PSR J0952–0607 [14, Chapter 3].

A significant shortcoming of this paper's methods is that the number of grid points and hence the required computing resources grow quickly with increasing frequency f and semimajor axis x . To make searches feasible, it might be necessary to balance a reduced search range (smaller maximum f and/or x) versus a reduced search sensitivity (wider grid spacing and/or shorter coherence time). Even with large computing resources like *Einstein@Home*, millisecond pulsar searches for binaries with $x \sim$ seconds are only feasible if the orbital parameters are precisely constrained.

The second significant shortcoming is that search sensitivity is lost if the pulsar's rotational phase does not match our model. This can happen for several types of pulsars and binary systems. This paper assumes that the intrinsic spin frequency f varies linearly with time. It does not include the time-dependent variations or the unpredictable frequency glitches often seen in young pulsars. This means that pulsars could be "detected" in the semicoherent stages of a search but are then discarded after the coherent stage, because they did not match the phase model well enough to produce a significant detection statistic [see, e.g., 79]. Phase model mismatch can also arise from time-dependent variations of the orbital period P_{orb} , which seems to be common in redback systems [see, e.g., 71]. For pulsars in short orbital period binaries with heavy companions, post-Keplerian gravitational corrections also have to be taken into account [see, e.g., 96, 152].

Because of these limitations, this paper also evaluates alternative search methods, which have previously been used in radio and gravitational-wave searches. While these may be applied to search for binary gamma-ray pulsars, only the sideband search methods appear to have some chance to detect tight-orbit spider pulsars, which are the main focus of this paper.

A more detailed study is necessary to make a fair sensitivity comparison between the sideband search and this paper's methods. Indeed, while the cumulative side-

band power loses a lot of signal-to-noise ratio compared to this paper’s methods, it might be improved. Since the sideband structure follows a known form, one could obtain a larger S/N by assigning weights to the sidebands before summing them, rather than using equal weights as done here.

We have implemented the new methods developed in this paper in a mixture of C and Python codes. These have been tested using simulated pulsar signals, both with our own code and with the widely used `TEMPO2` package [95]. We are confident that these codes work correctly, in part because they have discovered new spider pulsars, soon to be published.

We are currently using these codes and methods to hunt for spider pulsars in the unassociated sources of the *Fermi*-LAT Fourth Source Catalog. These partially-informed searches are guided by orbital constraints from optical observations. The orbital grids are constructed on the computing cluster ATLAS at the Albert Einstein Institute in Hannover. The first two (semicoherent) stages and the third (coherent) stage are all done on *Einstein@Home*, whose volunteers provide a massive computing pool. The final, less computation-demanding (H statistic) follow-up stage is done on ATLAS. To increase the computing power available in the initial stages of the search, we ported the search codes to work on *Einstein@Home* volunteer’s GPUs. The ATLAS cluster is also used to carry out follow-up gamma-ray searches of newly discovered radio pulsars, to refine the parameters as discussed above and in Chapter 3.

This paper has used the two gamma-ray sources 4FGL J1653.6–0158 and 4FGL J0523.3–2527 as examples, to show how a realistic search might be structured. Both of these searches are being or have been carried out, and the results will be discussed in upcoming papers.

The reader might wonder if these methods work in practice. They do, and they have already detected three spider pulsars. A preliminary version detected PSR J1311–3430 [44]. The current version successfully detected pulsations within 4FGL J2039.5–5617, by exploiting partial information [153, 154]. This confirmed that it is a redback and provides an 11 yr phase-connected rotational ephemeris (see Chapter 5). The search for 4FGL J1653.6–0158, described in Section 2.3.2, also resulted in a black widow MSP discovery (see Chapter 4).

That these methods work is not surprising: the different parts have been tested and demonstrated. *The metric approximation for the orbital parameters* was demonstrated to be a good fit to the actual mismatch for typical parameters as presented in Figures 2.1 and 2.3. The metric was used in a successful follow-up search shortly after the radio discovery of the fastest-spinning pulsar known in the Galactic field (Chapter 3). *The approximate phase model for elliptical orbits* was verified on simulated data with the results up to fifth order in eccentricity shown in Fig. 2.7. *The test statistics and the multistage search* approach have already detected more than 30 isolated pulsars [43, 79].

A topic we have not addressed is timing analysis. Following detection, this “pins down” the parameters as precisely as possible. An interesting and useful feature is that, regardless of the path to detection, if the pulsar is bright enough in gamma rays, the *Fermi*-LAT all-sky data immediately allow one to extend the ephemeris back to the launch of the *Fermi* satellite in August 2008 [52, 97]. This determines many of the pulsar’s parameters with much higher precision than is typical soon after radio/X-ray discoveries. For those, an additional campaign of timing observations

is required to infer astrophysical properties.

Gamma-ray timing analyses of LAT-discovered pulsars, which often remain undetected in radio, have led to several interesting discoveries [see, e.g., 45, 155, 156]. LAT data were used to resolve the variations in the orbital period of a binary pulsar, which was difficult to observe in radio owing to large eclipses [71]. The previously mentioned study of PSR J0952–0607 is another example.

The outlook for future searches is promising. The Gaia Catalog provides sky locations for the spider companions, which are precise enough so that no search in $\{\alpha, \delta\}$ is required. In addition, since the Large Area Telescope mission is ongoing, data sets are getting longer. Current searches use $T_{\text{obs}} \sim 11$ yr of data, compared with initial searches with $T_{\text{obs}} \sim 4$ yr. Furthermore, our available computing power is also increasing with time. This means that current searches employ $T_{\text{coh}} \sim 24$ d in the first stage, compared with initial searches with $T_{\text{coh}} \sim 12$ d. Since search sensitivity scales with $(T_{\text{coh}}T_{\text{obs}})^{1/4}$ [80], our current sensitivity has increased by more than 50%. We believe that $\mathcal{O}(10 - 30)$ of the unassociated sources in the 4FGL Catalog are undiscovered spider pulsars and that we can find some of them.

There are systems that are very likely to be spider gamma-ray pulsars for which the orbital constraints are not yet good enough to perform searches. These include the five redback pulsar candidates: 4FGL J0212.1+5321 [88, 89], 4FGL J0744.0–2525 [90], 4FGL J0838.7–2827 [91], 4FGL J0955.3–3949 [92], and the recent 4FGL J2333.1–5527 [93]. We hope that this work helps motivate additional optical observations to improve these constraints and enable new gamma-ray pulsar discoveries.

Acknowledgments

We thank Anne Archibald and Andrea Belfiore for encouraging us to look into sideband search methods. This work was supported by the Max-Planck-Gesellschaft (MPG), by the Deutsche Forschungsgemeinschaft (DFG) through Emmy Noether research grant No. PL 710/1-1 (PI: Holger J. Pletsch), and by National Science Foundation grants 1104902 and 1816904. This work was supported by an STSM Grant from COST Action CA16214. C.J.C. acknowledges support from the ERC under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 715051; Spiders).

2.7 Appendix to Chapter 2

2.7.1 Expectation values of signal statistics

Here we show how to calculate the expectation values of signal statistics. The statistics depend on the $j = 1, \dots, N$ modeled pulsar rotation phases at the photon arrival times t_j . To simplify the language and notation, we suppose that the vector of parameters $\boldsymbol{\lambda} = \{f, \dot{f}, \alpha, \delta\}$ is fixed and denote the modeled rotation phases by $\Phi_j = \Phi(t_j, \boldsymbol{\lambda}) = \Phi(t_{\text{psr}}(t_j, \alpha, \delta), f, \dot{f})$. Sums and products over j, k, ℓ run from $1, \dots, N$ unless otherwise specified. Finally, we write “the phase of the j th photon”, rather than “the modeled pulsar rotational phase associated with the j th photon”.

Our key assumption is that the phase of each photon is an independent (hence uncorrelated) random variable. This is justified because the number of photons detected is much less than one per pulsar revolution. The phase Φ_j of the j th photon is drawn from the distribution $F_j(\Phi_j)$ as given in Eq. (2.3). Thus, using Eq. (2.4), the probability distribution function of Φ_j is

$$F_j(\Phi_j) = \frac{1}{2\pi} + \frac{pw_j}{2\pi} \sum_{n=-\infty}^{\infty} \gamma_n e^{in\Phi_j}, \quad (2.100)$$

where the Fourier coefficients γ_n are defined by Eq. (2.5) for $n > 0$, by $\gamma_n = \gamma_{-n}^*$ for $n < 0$, and by $\gamma_0 = 0$ for $n = 0$.

The expectation value of any quantity $Q(\Phi_1, \dots, \Phi_N)$ is now given by

$$E[Q] = \int_0^{2\pi} d\Phi_1 F_1(\Phi_1) \cdots \int_0^{2\pi} d\Phi_N F_N(\Phi_N) Q(\Phi_1, \dots, \Phi_N), \quad (2.101)$$

where the statistical independence of the rotation phases allows the probability density to be written as a product. For example, the expected value of $\exp(-in\Phi_j)$ is

$$\begin{aligned} E[e^{-in\Phi_j}] &= \int_0^{2\pi} d\Phi_j F_j(\Phi_j) e^{-in\Phi_j} \\ &= \delta_{n0} + \frac{pw_j}{2\pi} \sum_{m=-\infty}^{\infty} \int_0^{2\pi} d\Phi \gamma_m e^{i(m-n)\Phi} \\ &= \delta_{n0} + pw_j \gamma_n, \end{aligned} \quad (2.102)$$

where δ_{nm} is the Kronecker delta, giving unity for $n = 0$.

The expected value of the coherent power signal statistic Eq. (2.6) is

$$E[P_n] = \kappa^{-2} \sum_{j,k} w_j w_k \prod_{\ell} \int_0^{2\pi} d\Phi_{\ell} F_{\ell}(\Phi_{\ell}) e^{in(\Phi_k - \Phi_j)}. \quad (2.103)$$

In the product above, only two terms are nontrivial, for which either $\ell = k$ or $\ell = j$. The integrand does not depend on the other $N - 2$ integration variables, whose corresponding integrals give unity, since the probability density is normalized. One obtains

$$\begin{aligned} E[P_n] &= \kappa^{-2} \sum_j w_j^2 + \kappa^{-2} p^2 \sum_{\substack{j,k \\ j \neq k}} w_j^2 w_k^2 |\gamma_n|^2 \\ &= 2 + 2p^2 |\gamma_n|^2 \left[\sum_j w_j^2 - \frac{\sum_j w_j^4}{\sum_j w_j^2} \right]. \end{aligned} \quad (2.104)$$

On the first line, the first sum comes from terms with $j = k$ and the second sum from terms where $j \neq k$, and we have used Eq. 2.102 to simplify both terms.

2.7.2 Maximal Sensitivity at Fixed Computing Cost

The sensitivity of a search can be quantified via the pulsed fraction p defined in Eq. (2.3). More sensitive searches can detect sources with smaller values of p .

If infinite computing power were available, we would employ the fully coherent detection statistics H or P_1 , and the sensitivity of a search would only be limited by

the data. To determine that ultimate sensitivity, consider the expected S/N $\theta_{P_1}^2$ given in Eq. (2.11). A point in parameter space where $\theta_{P_1}^2$ exceeded some threshold $\theta_{\text{threshold}}^2$ (established by the desired false-alarm and false-dismissal probabilities) would be counted as a detection. A reasonable detection threshold might be $\theta_{\text{threshold}}^2 = 50$, corresponding to pulsed fraction sensitivity $p^2 > \theta_{\text{threshold}}^2 / |\gamma_1|^2 \mu T_{\text{obs}}$. For typical values of $\mu T_{\text{obs}} = 500$ effective photons and $|\gamma_1|^2 = 0.8$, this gives an ultimate, data-limited sensitivity of $p^2 > 0.13$.

In practice, with limited computing power, we adopt the multistage hierarchical approach described in Section 2.2.3. A sensible choice is to use most of the computing power in the first, semicoherent stage. Roughly speaking, this is because a signal will only be found if it rises above the detection threshold in the first stage of the search⁹. Hence, we will assume that our sensitivity is limited by the first semicoherent search stage.

The maximum possible sensitivity of the semicoherent stage is determined by the threshold on the semicoherent S/N, whose expected value is given in Eq. (2.36). The threshold is lower than before, typically $\theta_{S_1}^2 > \theta_{\text{threshold}}^2 = 10$. Using search parameters from Eq. (2.27) and later in that section gives a minimum detectable pulsed fraction of $p^2 > \theta_{\text{threshold}}^2 / |\gamma_1|^2 \mu \sqrt{T_{\text{obs}} T_{\text{coh}}} = 0.31$. As before, this is the theoretical sensitivity that could be achieved with unlimited computing power, but employing the semicoherent statistic.

In practice, we must take the computing cost into account. This cost is proportional to the number of grid points in parameter space at which the detection statistic is calculated. Reducing the number of grid points (corresponding to a larger average mismatch) loses some S/N but the additional computing power may be used to increase the coherence time T_{coh} , which increases the S/N. What compromise maximizes the search sensitivity for a given computing cost?

To find the optimal balance between the worst-case grid mismatch m and the coherent integration time T_{coh} , we maximize the sensitivity with the constraint that the computing power is fixed, as described in [157] and [80]. What is important is the rate at which the number of grid points grows with increasing T_{coh} , which in turn depends on the dimension of the parameter space.

The number of dimensions d in the search parameter space is determined by our prior knowledge. To quantify that, we use n_{orb} (possible values 3 or 5) for the number of orbital parameters searched and n_{sky} (possible values 0 or 2) for the number of sky dimensions searched, so $d = 2 + n_{\text{orb}} + n_{\text{sky}}$. In the case of an eccentric binary with poorly known position, we have the full parameter space discussed in the main text, $\{f, \dot{f}, n_x, n_y, \Omega_{\text{orb}}, x, t_{\text{asc}}, \epsilon_1, \epsilon_2\}$, so $n_{\text{orb}} = 5$, $n_{\text{sky}} = 2$, and $d = 9$. For an eccentric binary whose position is precisely known (for example, from optical observations), $\{n_x, n_y\}$ are omitted from the search, $n_{\text{orb}} = 5$, $n_{\text{sky}} = 0$, and $d = 7$. For a circular binary whose position is precisely known, $\{\epsilon_1, \epsilon_2\}$ are also omitted, so $n_{\text{orb}} = 3$, $n_{\text{sky}} = 0$, and $d = 5$.

The smallest detectable pulsed fraction (averaged over signal location in parameter space) may be written as

$$p_{S_1}^2 = \frac{\theta_{\text{threshold}}^2}{(1 - \hat{m}) |\gamma_1|^2 \mu \sqrt{T_{\text{obs}} T_{\text{coh}}}}. \quad (2.105)$$

⁹Of course, this depends on the choice of threshold and the region of parameter space around a candidate that is searched in the subsequent stages. If the full parameter space is searched for each candidate, then the statement is false!

Here \hat{m} represents the average (over parameter space) mismatch of the grid [157].

The construction of our parameter space grid is described following Eq. (2.55); its average mismatch may be estimated as follows. Within a given 8 Hz frequency interval, the grid is the direct product of an equally spaced grid in the frequency direction, an equally spaced grid in the \dot{f} direction, a two-dimensional hexagonal lattice in sky position $\{n_x, n_y\}$, and an optimized stochastic grid in the orbital parameters. Below, we call these “subgrids”. To determine the computing cost, we need to count the number of grid points in these subgrids and multiply them together.

Because the metric has no off-diagonal terms that couple the different subgrids, the average parameter space mismatch \hat{m} can be written as

$$\hat{m} = \hat{m}_f + \hat{m}_{\dot{f}} + \hat{m}_{\text{sky}} + \hat{m}_{\text{orb}}, \quad (2.106)$$

where \hat{m}_f is the average mismatch in the frequency dimension (if all other parameters are exactly matched to the signal) and $\hat{m}_{\dot{f}}$, \hat{m}_{sky} , \hat{m}_{orb} are the corresponding average mismatches in the \dot{f} , sky, and orbital subgrids (if all other parameters are exactly matched to the signal).

The frequency dimension is searched with an FFT whose frequency spacing $df = 1/T_{\text{coh}}$. For the worst case, which is two points separated by $df/2$, the quadratic metric approximation predicts a mismatch $\bar{g}_{ff}/(2T_{\text{coh}})^2 = \pi^2/24 = 0.411$, and hence an average mismatch $\hat{m}_f = 0.14$. As is often the case, the quadratic approximation slightly overestimates the mismatch; the spherical ansatz of [121] predicts a worst-case $m = \sin^2(\sqrt{\pi^2/24}) \approx 0.36$ which agrees well with the numerically measured value given in Section 5.2 of [80]. In fact, as described before Eq. (42) of that paper, we can reduce the average mismatch to $\hat{m}_f = 0.075$ at almost no extra computational cost, by interpolating the frequency spectrum.

The \dot{f} subgrid has uniform spacing $d\dot{f}$, and is an example of a regular lattice. For regular lattices, the average mismatch \hat{m} is related to the worst-case mismatch m via $\hat{m} = \xi m$, where $\xi \in [0, 1]$ is a lattice-dependent dimensionless geometrical factor¹⁰. Here we have a (one-dimensional) hypercubic grid, for which $\xi = 1/3$, so the average mismatch $\hat{m}_{\dot{f}} = m_{\dot{f}}/3$, where $m_{\dot{f}} = \bar{g}_{\dot{f}\dot{f}}(d\dot{f}/2)^2 = \pi^2 T_{\text{coh}}^2 T_{\text{obs}}^2 d\dot{f}^2/288$ is the maximal mismatch in the \dot{f} dimension. (Since the differences are small, for simplicity we do not employ the spherical ansatz further.)

The sky subgrid is a hexagonal lattice with $\xi = 5/12 \approx 0.416$. Hence, $\hat{m}_{\text{sky}} = 0.416 m_{\text{sky}}$, where m_{sky} is the worst-case sky mismatch.

The orbital parameter grid has an average mismatch that is well estimated during the process of its construction and can be easily controlled via the parameter that determines when new points are added to the stochastic bank.

The computing cost is the product of the number of grid points in the nonfrequency dimensions with the cost of a single FFT. The number of grid points can be estimated using arguments like those given in deriving Eq. (2.58). In each of the different subgrids, the number of grid points is proportional to $\hat{m}^{-D/2}$, where \hat{m} is the average mismatch in that subgrid and D is the dimension of that subgrid. Hence, the number of grid points in the \dot{f} subgrid is proportional to $T_{\text{coh}} \hat{m}_{\dot{f}}^{-1/2}$, and the number of grid points in the sky subgrid is proportional to $T_{\text{coh}}^{n_{\text{sky}}} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2}$. The number of grid points in the orbital subgrid is proportional to $\hat{m}_{\text{orb}}^{-n_{\text{orb}}/2}$ and is

¹⁰In the published paper, this has been confused with the “thickness” [see, e.g., 157].

Table 2.3: Comparison of computationally unlimited and optimal computationally limited semicoherent searches, showing mismatches and sensitivity

Search	\hat{m}	\hat{m}_j	\hat{m}_{sky}	\hat{m}_{orb}	m_j	m_{sky}	$p_{S_1}^2$
Infinite computing cost (zero mismatch) grid	0	0	0	0	0	0	0.307
All parameters unknown ($n_{\text{sky}} = 2, n_{\text{orb}} = 5$)	0.383	0.039	0.077	0.193	0.116	0.093	0.497
Elliptical, known position ($n_{\text{sky}} = 0, n_{\text{orb}} = 5$)	0.471	0.066	0	0.330	0.198	0.159	0.580
Circular, known position ($n_{\text{sky}} = 0, n_{\text{orb}} = 3$)	0.383	0.077	0	0.231	0.231	0.185	0.497
Isolated ($n_{\text{sky}} = 2, n_{\text{orb}} = 0$)	0.221	0.049	0.097	0	0.146	0.117	0.394

Notes. — The columns show the average template bank mismatch \hat{m} , and the average mismatches in the \hat{f} , sky and orbital subgrids. (Note that the average per-dimension mismatch is constant.) Then the corresponding maximum \hat{f} and sky mismatch are listed with the (square of the) minimum detectable pulsed fraction p . The first row shows the ideal semicoherent case where the grid points are infinitesimally spaced and the computing cost is infinite. The next three rows illustrate smaller and smaller binary system parameter spaces. The final row is for an isolated pulsar with unknown sky position.

independent of T_{coh} . Since the cost of an FFT is proportional to $T_{\text{coh}} \log T_{\text{coh}}$, this gives a total computing cost C ,

$$C = C_0 \hat{m}_j^{-1/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2} T_{\text{coh}}^{2+n_{\text{sky}}}. \quad (2.107)$$

Here C_0 is a constant, and following Pletsch and Clark [80], we have omitted the slowly varying logarithmic factor from the cost of the FFT.

The method of Lagrange multipliers can be used to maximize sensitivity $p_{S_1}^{-2}$ at fixed computing cost¹¹. The quantity we extremize is

$$\begin{aligned} \mathcal{L} &= p_{S_1}^{-2} + \lambda C \\ &= c_1(1 - \hat{m})T_{\text{coh}}^{1/2} + \lambda c_2 \hat{m}_j^{-1/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2} T_{\text{coh}}^s, \end{aligned} \quad (2.108)$$

where λ is the Lagrange multiplier, $s = 2 + n_{\text{sky}}$, and c_1 and c_2 are constants (independent of the average mismatches and T_{coh}). Extremizing \mathcal{L} with respect to the coherence time and the three different average mismatches gives

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial T_{\text{coh}}} &= \frac{c_1}{2}(1 - \hat{m})T_{\text{coh}}^{-1/2} + s\lambda c_2 \hat{m}_j^{-1/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2} T_{\text{coh}}^{s-1} = 0 \\ \frac{\partial \mathcal{L}}{\partial \hat{m}_j} &= -c_1 T_{\text{coh}}^{1/2} - \frac{1}{2}\lambda c_2 \hat{m}_j^{-3/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2} T_{\text{coh}}^s = 0 \\ \frac{\partial \mathcal{L}}{\partial \hat{m}_{\text{sky}}} &= -c_1 T_{\text{coh}}^{1/2} - \frac{n_{\text{sky}}}{2}\lambda c_2 \hat{m}_j^{-1/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2-1} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2} T_{\text{coh}}^s = 0 \\ \frac{\partial \mathcal{L}}{\partial \hat{m}_{\text{orb}}} &= -c_1 T_{\text{coh}}^{1/2} - \frac{n_{\text{orb}}}{2}\lambda c_2 \hat{m}_j^{-1/2} \hat{m}_{\text{sky}}^{-n_{\text{sky}}/2} \hat{m}_{\text{orb}}^{-n_{\text{orb}}/2-1} T_{\text{coh}}^s = 0, \end{aligned}$$

where we have made use of Eq. (2.106) to evaluate the derivatives of \hat{m} .

To find the average mismatches that maximize the sensitivity at fixed computing cost, combine the first equation in turn with the second or third or fourth: T_{coh} drops

¹¹One obtains the same result by maximizing any negative power of p_{S_1} .

out, and one obtains a closed form for the corresponding average mismatch. The independence from coherence time T_{coh} in the binary pulsar case was previously shown for the isolated pulsar case by [80]. For example, to solve for \hat{m}_f , multiply the first equation by $T_{\text{coh}}^{1/2}$, multiply the second equation by $2s\hat{m}_f T_{\text{coh}}^{-1/2}$, and add them. One obtains $(1 - \hat{m})/2 - 2s\hat{m}_f = 0$, whose solution is $\hat{m}_f = (1 - \hat{m})/4s$. Doing this for all three combinations yields

$$\begin{aligned}\hat{m}_f &= \frac{1 - \hat{m}}{4(2 + n_{\text{sky}})}, \\ \hat{m}_{\text{orb}} &= \frac{1 - \hat{m}}{4(2 + n_{\text{sky}})} n_{\text{orb}}, \text{ and} \\ \hat{m}_{\text{sky}} &= \frac{1 - \hat{m}}{4(2 + n_{\text{sky}})} n_{\text{sky}}.\end{aligned}\tag{2.109}$$

Note that the optimal solution has equal average ‘‘per-dimension’’ mismatch in the non-frequency subgrids. From Eq. (2.106) it follows that $\hat{m} - \hat{m}_f$ is the sum of the three terms above, and since $d - 1 = 1 + n_{\text{orb}} + n_{\text{sky}}$, we have $\hat{m} - \hat{m}_f = (1 - \hat{m})(d - 1)/4(2 + n_{\text{sky}})$. The solution is

$$\hat{m} = \frac{1 + n_{\text{orb}} + n_{\text{sky}} + 4(2 + n_{\text{sky}})\hat{m}_f}{9 + n_{\text{orb}} + 5n_{\text{sky}}}.\tag{2.110}$$

Thus we have

$$\begin{aligned}\hat{m}_f &= \frac{1 - \hat{m}_f}{9 + n_{\text{orb}} + 5n_{\text{sky}}}, \\ \hat{m}_{\text{orb}} &= \frac{1 - \hat{m}_f}{9 + n_{\text{orb}} + 5n_{\text{sky}}} n_{\text{orb}}, \text{ and} \\ \hat{m}_{\text{sky}} &= \frac{1 - \hat{m}_f}{9 + n_{\text{orb}} + 5n_{\text{sky}}} n_{\text{sky}},\end{aligned}\tag{2.111}$$

which in turn allows us to determine the average and maximum mismatch in each of the subgrids, and the corresponding search sensitivity compared with an extremely finely spaced (but computationally very expensive) semicoherent search.

In practice, after setting the mismatch as given by this optimal point, one adjusts the coherence time T_{coh} to be as long as allowed by the available computing resources. What does this imply about the sensitivity? Above, we showed that with reasonable assumptions a semicoherent search can detect a pulsed fraction $p^2 > 0.31$ if there are infinite computing resources. With finite computing resources, this is increased by a factor of $1/(1 - \hat{m}) = (9 + n_{\text{orb}} + 5n_{\text{sky}})/4(2 + n_{\text{sky}})(1 - \hat{m}_f)$, as can be seen from Eq. (2.105). The corresponding loss of sensitivity is shown in Table 2.3. The achievable pulsed fraction sensitivity is not far from the ideal case.

This analysis extends previous work [80], which assumed a grid with fixed geometrical factor $\xi = 1/3$ in all dimensions. However, this is not the case for current searches. Here we have considered a grid that is a product of subgrids, each of which can have different geometrical properties, as used in existing searches. If we assume a fixed geometrical factor ξ , then our results and in particular the final line of Table 2.3 agree with Eq. (H2) from [80].

2.7.3 High-order phase model for elliptical binaries

The main text uses a linear-in- e “ELL1” approximation to the correct “BT” line-of-sight motion in eccentric orbits. Here we consider higher orders in the eccentricity e . The BT model is given in Eq. (2.64):

$$r_{z,\text{BT}}(t) = x \left[\sin \omega (\cos E - e) + \cos \omega \sqrt{1 - e^2} \sin E \right], \quad (2.112)$$

$$E - e \sin E = M, \quad (2.113)$$

$$M = \Omega_{\text{orb}}(t - T_0). \quad (2.114)$$

We express this as

$$r_{z,\text{BT}}(t) = x \left[\sin \omega \sum_{n=0}^{\infty} \alpha_n(e) \cos(n M) + \cos \omega \sum_{n=1}^{\infty} \beta_n(e) \sin(n M) \right], \quad (2.115)$$

where $\alpha_n(e)$ and $\beta_n(e)$ are power series in e . The goal here is to find these functions, and to determine the appropriate order needed for our searches. ([158] gives an expansion of $\sin E$ and $\cos E$ in powers of e , but does not give a similar expansion for the line-of-sight motion.)

For the derivation of the power series we introduce the Bessel functions and some of their properties. For positive integers n the Bessel function can be expressed as the power series

$$J_n(x) = \sum_{m=0}^{\infty} \frac{(-1)^m \left(\frac{x}{2}\right)^{2m+n}}{m!(n+m)!} \quad (2.116)$$

or in integral form as

$$J_n(x) = \frac{1}{2\pi} \int_0^{2\pi} \cos(n\theta - x \sin \theta) d\theta. \quad (2.117)$$

The relation

$$J_{n-1}(x) + J_{n+1}(x) = \frac{2n}{x} J_n(x) \quad (2.118)$$

is also needed.

Following [158], we start with the Fourier expansion of $\cos E$:

$$\cos E = \frac{\hat{\alpha}_0}{2} + \sum_{n=1}^{\infty} \hat{\alpha}_n \cos(n M). \quad (2.119)$$

It has Fourier coefficients

$$\hat{\alpha}_n = \frac{1}{\pi} \int_0^{2\pi} \cos E \cos(n M) dM. \quad (2.120)$$

Using Kepler’s equation (2.113) to write M and dM as functions of E , along with the integral form above, one obtains

$$\hat{\alpha}_n = \frac{2}{n} \frac{dJ_n(n e)}{d(n e)}. \quad (2.121)$$

Using the power series above, this may be written as

$$\hat{\alpha}_n(e) = \begin{cases} \sum_{m=0}^{\infty} \frac{(-1)^m (2m+n)}{n m! (m+n)!} \left(\frac{n}{2}\right)^{2m+n-1} e^{2m+n-1} & , n \geq 1 \\ -\frac{1}{2}e & , n = 0. \end{cases} \quad (2.122)$$

The analogous calculation for $\sin E$ gives

$$\sin E = \sum_{n=1}^{\infty} \hat{\beta}_n \sin(n M) \quad (2.123)$$

with coefficients

$$\hat{\beta}_n = \frac{1}{\pi} \int_0^{2\pi} \sin E \sin(n M) dM, \quad (2.124)$$

where

$$\begin{aligned} \hat{\beta}_n &= \frac{2}{n e} J_n(n e) \\ &= \sum_{m=0}^{\infty} \frac{(-1)^m \left(\frac{n e}{2}\right)^{2m+n-1}}{m! (n+m)!} e^{2m+n-1} \end{aligned} \quad (2.125)$$

is obtained using the recursion relation above.

To obtain β from $\sqrt{1-e^2}\hat{\beta}$, we first express

$$\sqrt{1-e^2} = \sum_{k=0}^{\infty} e^{2k} \left(\sum_{l=0}^{2k} (-1)^l \binom{1/2}{l} \binom{1/2}{2k-l} \right), \quad (2.126)$$

where we have introduced the generalized binomial coefficient

$$\binom{r}{k} = \frac{r \cdot (r-1) \cdots (r-(k-1))}{k!}. \quad (2.127)$$

The Cauchy product of $\sqrt{1-e^2}$ and $\hat{\beta}$ gives

$$\begin{aligned} \beta_n(e) &= \sum_{m=0}^{\infty} e^{2m+n-1} \sum_{k=0}^m \frac{\left(\frac{n}{2}\right)^{(2m-2k+n-1)}}{(m-k)! (n+m-k)!} \\ &\quad \times \sum_{l=0}^{2k} (-1)^{m-k+l} \binom{1/2}{l} \binom{1/2}{2k-l}. \end{aligned} \quad (2.128)$$

The α_n follow directly from $\hat{\alpha}_n$, and differ only for $n = 0$.

We list the results to 11th order. (A similar calculation [159] gives the coefficients




Figure 2.6: Mismatch between the BT model and models truncated at orders e^0, e^1, e^2 , and e^3 , for the source 4FGL J0523.3–2527 with $e = 0.04$. This is computed on a grid of 100×100 simulated pulsar signals, with equally spaced $\log_{10} f/\text{Hz} \in [0, 3]$, and $\log_{10} x/\text{s} \in [-2, 1]$. The gray dashed line indicates the semi-major axis $x = 3.66$ of the likely pulsar in 4FGL J0523.3–2527. The slopes of the constant-mismatch contours are the same for different models because e is fixed.




Figure 2.7: Same as Figure 2.6, but varying the eccentricity e with fixed frequency $f = 1$ kHz, and going up to e^5 . The mismatch is computed on a grid of 100×100 simulated pulsar signals, with equally spaced $\log_{10} e \in [-3, 0]$, and $\log_{10} x/s \in [-2, 1]$.

to seventh order, but without a general formula.) The α values are given by

$$\alpha_0 = -\frac{3}{2}e, \quad (2.129a)$$

$$\alpha_1 = 1 - \frac{3}{8}e^2 + \frac{5}{192}e^4 - \frac{7}{9216}e^6 + \frac{1}{81920}e^8 - \frac{11}{88473600}e^{10} \quad (2.129b)$$

$$\alpha_2 = \frac{1}{2}e - \frac{1}{3}e^3 + \frac{1}{16}e^5 - \frac{1}{180}e^7 + \frac{1}{3456}e^9, \quad (2.129c)$$

$$\alpha_3 = \frac{3}{8}e^2 - \frac{45}{128}e^4 + \frac{576}{5120}e^6 - \frac{729}{40960}e^8 + \frac{8019}{4587520}e^{10} \quad (2.129d)$$

$$\alpha_4 = \frac{1}{3}e^3 - \frac{2}{5}e^5 + \frac{8}{45}e^7 - \frac{8}{189}e^9, \quad (2.129e)$$

$$\alpha_5 = \frac{125}{384}e^4 - \frac{4375}{9216}e^6 + \frac{15625}{57344}e^8 - \frac{4296875}{49545216}e^{10}, \quad (2.129f)$$

$$\alpha_6 = \frac{27}{80}e^5 - \frac{81}{140}e^7 + \frac{729}{1792}e^9, \quad (2.129g)$$

$$\alpha_7 = \frac{16807}{46080}e^6 - \frac{117649}{163840}e^8 + \frac{63412811}{106168320}e^{10}, \quad (2.129h)$$

$$\alpha_8 = \frac{128}{315}e^7 - \frac{512}{567}e^9, \quad (2.129i)$$

$$\alpha_9 = \frac{531441}{1146880}e^8 - \frac{52612659}{45875200}e^{10}, \quad (2.129j)$$

$$\alpha_{10} = \frac{78125}{145152}e^9, \quad (2.129k)$$

$$\alpha_{11} = \frac{2357947691}{3715891200}e^{10}. \quad (2.129l)$$

The β values are given by

$$\beta_1 = 1 - \frac{5}{8}e^2 - \frac{11}{192}e^4 - \frac{457}{9216}e^6 - \frac{23479}{737280}e^8 - \frac{2014861}{88473600}e^{10} \quad (2.130a)$$

$$\beta_2 = \frac{1}{2}e - \frac{5}{12}e^3 + \frac{1}{24}e^5 - \frac{1}{45}e^7 - \frac{379}{34560}e^9, \quad (2.130b)$$

$$\beta_3 = \frac{3}{8}e^2 - \frac{51}{128}e^4 + \frac{543}{5120}e^6 - \frac{219}{8192}e^8 - \frac{18141}{4587520}e^{10} \quad (2.130c)$$

$$\beta_4 = \frac{1}{3}e^3 - \frac{13}{30}e^5 + \frac{13}{72}e^7 - \frac{739}{15120}e^9, \quad (2.130d)$$

$$\beta_5 = \frac{125}{384}e^4 - \frac{4625}{9216}e^6 + \frac{144625}{516096}e^8 - \frac{4611125}{49545216}e^{10}, \quad (2.130e)$$

$$\beta_6 = \frac{27}{80}e^5 - \frac{135}{224}e^7 + \frac{3753}{8960}e^9, \quad (2.130f)$$

$$\beta_7 = \frac{16807}{46080}e^6 - \frac{218491}{294912}e^8 + \frac{65160739}{106168320}e^{10}, \quad (2.130g)$$

$$\beta_8 = \frac{128}{315}e^7 - \frac{2624}{2835}e^9, \quad (2.130h)$$

$$\beta_9 = \frac{531441}{1146880}e^8 - \frac{53675541}{45875200}e^{10}, \quad (2.130i)$$

$$\beta_{10} = \frac{78125}{145152}e^9, \quad (2.130j)$$

$$\beta_{11} = \frac{2357947691}{3715891200}e^{10}. \quad (2.130k)$$

A line-of-sight model accurate to $\mathcal{O}(e^k)$ requires retaining terms up to and including α_{k+1} and β_{k+1} . Depending on the search parameters $\{f, x, e\}$, different orders of these Taylor series will be required.

Consider the source 4FGL J0523.3–2527. The expected eccentricity is $e \sim 0.04$. To find the appropriate order in e , we simulated 10,000 realizations of a pulsar in 4FGL J0523.3–2527, with different spin frequencies f and semimajor axes x . Figure 2.6 shows the mismatches that arise from using approximations of different orders in e , compared to the full BT model. For high frequencies the mismatch m is significant, $m \sim 0.3$ (S/N loss of up to 30%), for the $\mathcal{O}(e^2)$ -model. A sensible choice is the $\mathcal{O}(e^3)$ -model, for which the mismatch is below 1% for frequencies $f < 1$ kHz.

For systems with different eccentricities, we can also provide guidance. Since most of the known spider pulsars are MSPs, we simulated 10,000 realizations of a 1 kHz pulsar with different semimajor axes x and eccentricities e . Figure 2.7 shows the mismatches that arise up to sixth order in e .

Detection and Timing of Gamma-ray Pulsations from the 707 Hz Pulsar J0952–0607

Published as Nieder, L., Clark, C. J., Bassa, C. G., et al. 2019, ApJ, 883, 42.

DOI: [10.3847/1538-4357/ab357e](https://doi.org/10.3847/1538-4357/ab357e)

© 2019. The American Astronomical Society.

Abstract

The Low-Frequency Array radio telescope discovered the 707 Hz binary millisecond pulsar (MSP) J0952–0607 in a targeted radio pulsation search of an unidentified *Fermi* gamma-ray source. This source shows a weak energy flux of $F_\gamma = 2.6 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ in the energy range between 100 MeV and 100 GeV. Here we report the detection of pulsed gamma-ray emission from PSR J0952–0607 in a very sensitive gamma-ray pulsation search. The pulsar’s rotational, binary, and astrometric properties are measured over seven years of *Fermi*-Large Area Telescope data. For this we take into account the uncertainty on the shape of the gamma-ray pulse profile. We present an updated radio-timing solution now spanning more than two years and show results from optical modeling of the black-widow-type companion based on new multi-band photometric data taken with HiPERCAM on the Gran Telescopio Canarias on La Palma and ULTRACAM on the New Technology Telescope at ESO La Silla¹. PSR J0952–0607 is now the fastest-spinning pulsar for which the intrinsic spin-down rate has been reliably constrained ($\dot{P}_{\text{int}} \lesssim 4.6 \times 10^{-21} \text{ s s}^{-1}$). The inferred surface magnetic field strength of $B_{\text{surf}} \lesssim 8.2 \times 10^7 \text{ G}$ is among the ten lowest of all known pulsars. This discovery is another example of an extremely fast spinning black-widow pulsar hiding within an unidentified *Fermi* gamma-ray source. In the future such systems might help to pin down the maximum spin frequency and the minimum surface magnetic field strength of MSPs.

¹Based on observations collected at the European Southern Observatory, Chile; programme 0101.D-0925, PI: Clark, C. J. .

3.1 Introduction

The LAT on board the *Fermi Gamma-ray Space Telescope* [32] has proven itself to be a powerful instrument in gamma-ray pulsar astronomy. Since its 2008 launch the LAT has been operating in an all-sky survey mode. LAT data are used to identify promising pulsar candidates for deep, targeted radio searches and find gamma-ray pulsations in blind or follow-up searches [for a review see, e.g., 23]. The 10 year time span of the all-sky LAT data is also useful for establishing precise pulsar-timing ephemerides of new discoveries.

Radio pulsar searches targeting the sky positions of LAT sources have been very successful in finding isolated and binary millisecond pulsars [MSPs; e.g., 56]. The targeted sources are typically chosen to have three properties: (a) They are “unassociated”, which means that the source has no plausible counterpart belonging to a known gamma-ray-emitting source class [e.g., 37]. (b) They have curved spectra. This is parametrized in the *Fermi*-LAT source catalogs by the curvature significance, determined by the difference in log-likelihood between spectral models with curved spectra (e.g. a log parabola or exponentially cutoff power law) versus power-law spectra [36]. For most gamma-ray pulsars, curved spectra are preferred with $> 95\%$ confidence [e.g., 28]. (c) They show only little variability in brightness over time, which is indicated in the *Fermi* LAT source catalogs by the variability index, the chi-squared of the monthly flux with respect to the average flux. In the *Fermi*-LAT Third Source Catalog [3FGL; 37], only 2 out of 136 pulsars had variability indices corresponding to significant variability above the 99% confidence level. Combined, the last two properties are good indicators for gamma-ray pulsars. However, we note that the transitional MSPs [for a review see, e.g., 160] are an important exception, with significant changes in gamma-ray flux associated with transitions between accretion- and rotation-powered states [22, 161].

Searches following this approach continue to find pulsars by using radio observing frequencies ν above 300 MHz. Pulsar surveys around 350 MHz are run by the Green Bank Telescope [GBT; 62] and the Arecibo telescope [57]. The Giant Metrewave Radio Telescope searches around 607 MHz [60]. Another survey around 820 MHz is run by the GBT [61]. Finally Parkes [65], Nançay [64] and Effelsberg [58] search around 1.4 GHz. Radio observations at higher frequencies suffer less from dispersion (dispersion delay $t_d \propto \nu^{-2}$) and scattering [scattering timescale $\tau_s \propto \nu^{-4.4}$; 162] but a pulsar’s radio luminosity falls rapidly with observing frequency [radio flux density $S_\nu \propto \nu^\alpha$ with spectral index $-3.0 < \alpha < -0.5$ for most known pulsars; 163]. At observing frequencies above 1.4 GHz scattering becomes negligible away from the Galactic Center and pulsars that are bright above this frequency can be useful for Pulsar Timing Arrays [e.g., 164, 165].

However, there might be a population of steep-spectrum ($\alpha < -2.5$) radio pulsars that are most easily detectable at frequencies below 300 MHz. Searches by Frail et al. [166] for steep-spectrum sources within the localization regions of unidentified *Fermi*-LAT sources in continuum images from the Giant Metrewave Radio Telescope all-sky survey at 150 MHz led to the discovery of six new MSPs and one normal pulsar. These detections suggest that many steep-spectrum pulsars may have been missed by high-frequency radio surveys, which favor pulsars with flatter spectra [167]. Additionally, some emission models suggest that pulsars’ radio beams are wider at low frequencies [e.g., 168], making pulsars whose radio beams

miss our line of sight at GHz frequencies potentially detectable at lower frequencies. Low-frequency radio observations of gamma-ray pulsars can therefore provide an additional test of the viewing-angle explanation for the large number of radio-quiet pulsars discovered by the LAT [e.g. 73, 169]. Indeed, one emission model for the recently discovered radio-quiet MSP PSR J1744–7619 [43] suggests that radio pulsations may only be detectable at low radio frequencies.

Pleunis et al. [63] performed very-low-frequency pulsar searches at 115–155 MHz with the Low-Frequency Array [LOFAR; 170, 171]. This was possible due to new semi-coherent de-dispersion techniques that mitigate the smearing due to dispersion [172]. The searches targeted unassociated sources from the 3FGL catalog [37]. An isolated MSP, PSR J1552+5437, was detected first in radio and subsequently in gamma rays [63].

Bassa et al. [14] conducted another LOFAR survey using the same observing configuration. The 23 targets were unassociated gamma-ray sources selected from a *Fermi*-LAT source list constructed from seven years of “Pass 8” LAT data [see 173].

In this survey they discovered PSR J0952–0607, a binary radio MSP with a spin frequency of 707 Hz [14]. It is in a binary system with a very-low-mass companion star ($M_c \sim 0.02 M_\odot$) with an orbital period of 6.42 hr. PSR J0952–0607 is the fastest-spinning known neutron star outside of a globular cluster: The only pulsar spinning faster (716 Hz) is PSR J1748–2446ad, which is located in the globular cluster Terzan 5 [13]. In contrast to pulsars in globular clusters, which experience significant but unknown acceleration due to the gravitational potential within the cluster [174], the intrinsic spin-down rate of PSR J0952–0607 can be measured directly. From this, pulsar properties like the dipole surface magnetic field strength and spin-down power can be inferred. These factors are thought to govern the poorly understood accretion and ablation processes through which binary systems containing a pulsar evolve [125]. Measurements of the magnetic fields of rapidly spinning pulsars are important because the origin of the low magnetic field strength of MSPs is currently unexplained, with one popular theory being that the accreted matter buries the surface magnetic field. On the other hand recent work questions if this mechanism is effective enough [175].

To determine the pulsar properties requires precise timing solutions from frequent observations of a pulsar over several years. For some pulsar parameters (e.g. the spin frequency and spin-frequency derivative) the measurement uncertainty is directly related to the total span of observations. Furthermore, time spans shorter than one year cover less than a full cycle of the annual Roemer delay, introducing degeneracies between the spin frequency, spin-frequency derivative, and sky position. The radio-timing solution of PSR J0952–0607 reported by Bassa et al. [14] is based on observations spanning approximately 100 days, and thus suffers from these issues.

Radio searches targeting unassociated *Fermi*-LAT sources have been particularly successful at discovering “spider pulsars”, a class of extreme binary pulsars with semi-degenerate companion stars (i.e. not neutron stars or white dwarfs). These systems are categorized as “black widows” if the companion star has extremely low mass ($M_c \ll 0.1 M_\odot$, as is the case for PSR J0952–0607) and as “redbacks” if the companion star is heavier ($M_c \sim 0.15 - 0.4 M_\odot$) [53]. Optical light curves of these systems reveal that the pulsar emission heats the nearly Roche-lobe filling companion [176]. Observations of orbitally modulated X-ray emission shows that interactions

between the pulsar and companion star winds produce intra-binary shocks [e.g., 177].

For many spider pulsars the radio pulsations are completely absorbed by intra-binary material during parts of their orbit [e.g., 49], indicating that the companion stars are also ablated by the pulsar. At low radio frequencies these eclipses can cover a large fraction of the orbit [e.g., 19, 178, 179], complicating radio-timing campaigns. In contrast, gamma-ray pulsations are essentially unaffected by eclipses.

A unique value of the LAT data is that a pulsar’s discovery in gamma rays often enables the immediate measurement of the pulsar parameters over the 10-year span in which the LAT has been operating. LAT data have been used to find precise timing solutions for many pulsars including radio-quiet and radio-faint pulsars [52, 79, 97]. In the case of PSR J2339–0533, a strongly eclipsing redback pulsar, gamma-ray timing was essential for building a coherent timing solution, and enabled the discovery of large variations of the orbital period [71].

In this work we present the discovery and analysis of pulsed gamma-ray emission from PSR J0952–0607. The pulsar itself is very faint in gamma rays, and required novel search and timing methods with greater sensitivity. The resulting timing ephemeris extends the rotational and orbital history of PSR J0952–0607 back seven years to 2011. This allows us to determine the pulsar’s spin-down power and surface magnetic field strength, making it the fastest known pulsar for which these measurements can be made.

The paper is organized as follows. In Section 3.2 we describe the pulsation search and detection within LAT data. The timing analysis and resulting timing solution for PSR J0952–0607 are presented in Section 5.3.2. New radio and optical observations as well as a search for continuous gravitational waves are discussed in Section 3.4. Finally, in Section 7.4 we discuss the implications of the results presented and we conclude in Section 3.6.

3.2 Gamma-ray Pulsation Discovery

3.2.1 Data Preparation

The gamma-ray source targeted by [14] resulting in the detection of the radio pulsar PSR J0952–0607 and its optical counterpart (R.A. $\alpha_{J2000.0} = 09^{\text{h}}52^{\text{m}}08^{\text{s}}.319$, Decl. $\delta_{J2000.0} = -06^{\circ}07'23''.49$) was discovered using seven years of LAT data, but was too faint to be included in the 3FGL catalog [i.e. in four years of data; 37]. It is included in the successive 4FGL catalog based on eight years of data as 4FGL J0952.1–0607 [180].

To search for gamma-ray pulsations from PSR J0952–0607, we used “Pass 8” [173] LAT data recorded between 2008 August 4 and 2017 January 19, consisting of SOURCE-class photons above 500 MeV instead of the standard 100 MeV. Since the LAT’s angular resolution for photons improves with energy (~ 3.6 times higher angular resolution at 500 MeV compared to 100 MeV), we conservatively used 500 MeV to avoid potential contamination by other nearby sources not included in the 3FGL catalog². The photons were selected using `gtselect` from the *Fermi* Science Tools³ if

²https://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone.LAT_IRFs/IRF_PSF.html

³<https://fermi.gsfc.nasa.gov/ssc/data/analysis/software>

they were within 10° of the celestial position of the optical counterpart to PSR J0952–0607, with a maximum zenith angle of 90° . Photons were only used if the LAT was in nominal science mode and if the rocking angle was below 52° . After these cuts 114706 LAT photons remained for further analysis. The analysis was performed using the `P8R2_SOURCE_V6` instrument response functions (IRFs).

The sensitivity of a pulsation search can be greatly improved by weighting the contribution of each photon by its probability of having originated from the candidate pulsar [113, 114]. The weights are computed based on the LAT response function and a spectral model of a point source. They are used in the search and the timing analysis for background suppression without the need for arbitrary position or stronger energy cuts.

To produce the necessary spectral model we performed a binned spectral analysis with `gtlike`. We added a putative pulsar source with an exponentially cutoff power law to represent its spectrum [36] fixed to the position of the pulsar’s optical counterpart reported by Bassa et al. [14]. We used the templates `gll_iem_v06.fits` for the Galactic diffuse emission [181] and `iso_P8R2_SOURCE_V6_v06.txt`⁴ for the isotropic diffuse background. The spectral analysis included all 3FGL sources within 15° of the pulsar position and the spectral parameters for point sources within 5° of the target were allowed to vary.

For each photon within 5° of the pulsar’s optical position a probability weight w_j was calculated with `gtsrcprob`. To reduce the computing cost of the search, we only included photons with $w_j > 3.1\%$. This weight cutoff value was chosen such that only 1% of the expected pulsation S/N would be lost. After applying the cutoff $N = 1354$ actual or $\sum w_j = 193.7$ “effective” photons remain.

Upon the detection of PSR J0952–0607, we performed a dedicated spectral analysis with an extended dataset in order to enhance the pulsation significance and to model its spectral characteristics more precisely. We used the same event selection and IRFs (see above) but accepted photons without cuts on the rocking angle as this cut was found to be overly conservative⁵. We extended the dataset to include photons between 2008 August 4 and 2018 June 21. We lowered the threshold of photon energies down to 100 MeV to further constrain the spectral characteristics. We used the Preliminary LAT 8-year Point Source List⁶ (FL8Y) to construct our source model. The FL8Y source associated with the pulsar, FL8Y J0952.2–0608, was replaced by a point source fixed to the position of the detected gamma-ray pulsar. All FL8Y sources within 15° of the pulsar position were included and the spectral parameters for point sources within 5° of the pulsar were allowed to vary.

We computed the residual TS map to search for non-cataloged weak gamma-ray sources in the vicinity of the pulsar. The test statistic $TS = 2(\log \mathcal{L}(\text{source}) - \log \mathcal{L}(\text{no source}))$ quantifies how significant a source emerges from the background, where the likelihood \mathcal{L} of a model with and without a source is compared [36, 37]. Six uncataloged sources with $TS > 10$ ($\sim 3\sigma$) within 5° of the pulsar position were found and added to the source model. Using this new source model we reran the analysis. The result of the spectral analysis for PSR J0952–0607 is shown in Table 3.1. Here, we also give TS_{cut} which is computed like TS but comparing an

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

⁵https://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone_Likelihood/Exposure.html

⁶<https://fermi.gsfc.nasa.gov/ssc/data/access/lat/fl8y/>

Table 3.1: Spectral Parameters of PSR J0952–0607.

Parameter	Value
Test statistic, TS	147.77
TS of exponential cutoff, TS_{cut}	23.9
Photon index, Γ	$0.95 \pm 0.40 \pm 0.05$
Cutoff energy, E_c (GeV)	$1.62 \pm 0.55 \pm 0.01$
Photon flux ($10^{-9} \text{ cm}^{-2} \text{ s}^{-1}$)	$2.25 \pm 0.77 \pm 0.34$
Energy flux F_γ ($10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$)	$2.60 \pm 0.38 \pm 0.16$

Notes. — Gamma-ray spectrum based on LAT data between MJD 54,682–58,289 over the standard energy range from 100 MeV to 100 GeV. The first reported uncertainties are statistical, while the second uncertainties are systematic, determined by re-analyzing the data with bracketing IRFs and artificially changing the normalization of the Galactic diffuse model by $\pm 6\%$, as described by Abdo et al. [28].

exponentially cutoff power-law model and a power-law model without cutoff [28].

In the timing analysis we used all photons with weights $w_j > 1.5\%$, which is chosen as in the search such that 99% of the S/N remains. This leaves $N = 4642$ actual or $\sum w_j = 331.4$ effective photons.

3.2.2 Search

For many pulsars, LAT data covering several years of observation time are needed for significant pulsation detection [e.g. 182]. Searching for pulsations requires assigning every gamma-ray photon with the pulsar’s rotational phase Φ (defined in rotations throughout the paper) at the time of emission. To do this a phase model $\Phi(t, \boldsymbol{\lambda})$ is used that depends on time t and (for circular-binary pulsars) on a set of at least seven parameters $\boldsymbol{\lambda} = (f, \dot{f}, \alpha, \delta, P_{\text{orb}}, x, t_{\text{asc}})$. These parameters are needed to: (1) Correct the photon arrival times for the LAT’s movement with respect to the Solar System Barycenter (sky position α and δ). (2) In the case of a circular binary, account for the pulsar’s movement around the center of mass (orbital period P_{orb} , projected semi-major axis x , and epoch of ascending node t_{asc}). (3) Describe the pulsar’s rotation over time (spin frequency f and spin-frequency derivative \dot{f}).

The ephemeris obtained by timing a radio pulsar over a short interval T_{obs} often does not determine the parameters precisely enough to coherently fold the multiple years of LAT data. For $T_{\text{obs}} < 1 \text{ yr}$ the spin and position parameters of the pulsar are strongly correlated (i.e., degenerate). Over longer T_{obs} the uncertainties in the spin parameters scale with negative powers of T_{obs} . The uncertainty in the orbital period scales with T_{obs}^{-1} if $T_{\text{obs}} \gg P_{\text{orb}}$.

Searches for binary gamma-ray pulsars are therefore computationally expensive, as a multi-dimensional parameter space must be searched with a dense grid [44]. The radio detection and timing are crucial to constrain the relevant parameter space that has to be searched to find the gamma-ray pulsations.

Using the radio data [14] found that PSR J0952–0607 is in a circular-binary orbit. Furthermore, they measured α and δ by identifying the companion star

using optical data taken with the Wide Field Camera (WFC) on the 2.5m Isaac Newton Telescope on La Palma. Barycentering the radio data according to α and δ obtained from the optical data resulted in an upper limit on \dot{f} and determined f more accurately. Furthermore the radio timing constrained the orbital parameters P_{orb} , x , and t_{asc} .

The gamma-ray pulsation search exploited preliminary constraints from radio timing of the pulsar combined with the optical position.

In the gamma-ray pulsation search we used the H statistic [119]. It combines the Fourier power from several harmonics incoherently by maximizing over the first M harmonics via

$$H = \max_{1 \leq M \leq M_{\text{max}}} \left(4 - 4M + \sum_{n=1}^M \mathcal{P}_n \right), \quad (3.1)$$

with $M_{\text{max}} = 20$ as suggested by [119]. The Fourier power in the n th harmonic is given by

$$\mathcal{P}_n = \frac{1}{\kappa^2} \left| \sum_{j=1}^N w_j e^{-2\pi i n \Phi(t_j)} \right|^2, \quad (3.2)$$

with the normalization constant

$$\kappa^2 = \frac{1}{2} \sum_{j=1}^N w_j^2. \quad (3.3)$$

The construction of a grid for this search was done using a distance “metric” on the parameter space [117, 118]. This is a second-order Taylor approximation of the fractional loss in squared S/N due to an offset from the parameters of a given signal. The metric allows one to compute analytically the density of an optimally spaced grid. This method was successfully used in the blind search (i.e., a search for a previously undetected pulsar) for the black widow PSR J1311–3430 [44].

The metric components for the parameters of an isolated pulsar are given by Pletsch and Clark [80], and the additional components required to search for a binary pulsar will be described in an upcoming paper (Chapter 2). The grid point density computed with the metric varies throughout the parameter space. The grid density in α and δ increases as f increases. This is also the case for the orbital parameters. In addition, for P_{orb} and t_{asc} the grid point density increases with the projected semi-major axis, x . The small x typical for black-widow pulsars with their low-mass companions therefore greatly reduces the required density.

In addition, when performing a harmonic-summing search, any parameter offset results in a phase offset at the n th harmonic that is a factor of n larger than at the fundamental. To avoid this, the search grid density must be increased by a factor of M_{max} in each parameter. Fortunately, known gamma-ray pulsars have the most power in the first few harmonics [80]. We therefore designed the search grid to lose at most 1% of the Fourier power in the fifth harmonic in each dimension. The harmonic summing was also truncated at $M_{\text{max}} = 5$ to reduce computing cost. The required number of points in the search grid was reduced this way by a factor of 4^5 (≈ 1000) compared to a grid built for $M_{\text{max}} = 20$. This search grid was designed to be very dense since the pulsar signal was expected to be weak due to the small number of photons.

Based on the distance metric we built a hypercubic grid covering the relevant parameter space in f , \dot{f} , α , δ , and P_{orb} . This means that the parameter space is broken down into smaller cells. The edges of these cells are parallel to the parameter axes and of equal length in each dimension as computed by the distance metric. We note that a simple hypercubic grid is sufficient because the metric is nearly diagonal (off-diagonal terms are small; Chapter 2), and the dimensionality is low. For higher dimensional parameter spaces hypercubic grids become extremely wasteful. The projected semi-major axis and the epoch of the ascending node were known precisely enough from the radio ephemeris that no search over these parameters was necessary. In summary, we performed a grid-based search over five parameters (f , \dot{f} , α , δ , and P_{orb}), while keeping two parameters (x and t_{asc}) fixed to the values from the radio-timing solution.

The search used 2×10^5 CPU-core hours, meaning that the search would have taken 24 years to compute on a single core. Therefore, we distributed the work in chunks over 8000 CPU cores of the ATLAS computing cluster [183], and the search took only 2 days.

3.2.3 Detection

To ensure that the signal was inside the covered parameter space we searched over wide ranges in the highly correlated f (4σ), α , and δ (5σ each), where σ is the parameter uncertainty obtained from preliminary radio and optical observations. The chosen search range for P_{orb} (3σ) was smaller because the radio-timing-derived P_{orb} was not degenerate with the other parameters.

Surprisingly, the largest H statistic appeared close to the edge of our search range in f and with a significant offset in α and δ . The latter was determined to be due to an error in the initial astrometric calibration of the optical images of the optical counterpart. After the discovery of this error only the corrected α and δ values were published by [14]. The offset in f arose from the strong correlation with α and δ . Therefore we started another search with the same settings starting from the highest f covered in the first search. The largest H statistic was $H_m = 86.7$ (without refining the parameters any further) and lay well within the combined search parameter space.

While this H statistic was far larger than any other found in our search, it is not easy to estimate the statistical significance (or false-alarm probability) of the maximum value found in a dense, multi-dimensional H statistic search (see Appendix 3.7.1). We therefore applied a ‘‘bootstrapping’’ procedure (described in Appendix 3.7.1) to estimate the detection significance from the search results themselves, finding a trials-corrected false-alarm probability of $P_{\text{FA}} \approx 3.3 \times 10^{-3}$. After extending our data set to cover the extra year of data as explained in Section 3.2.1, and without using a weight cut (which is only introduced for computational reasons), we found that the H statistic value increased to $H = 102.9$ without further refinement (i.e., in a single trial). Since no additional trials have been performed in this step, we can multiply our false-alarm probability estimate by the known single-trial false-alarm probability [114] for this increase ($P_{\text{FA}} = \exp(-0.3984 \Delta H_m) = 1.6 \times 10^{-3}$), giving an overall false-alarm probability of $P_{\text{FA}} \approx 5.3 \times 10^{-6}$ in the extended data set, confirming the detection.

3.3 Gamma-ray Timing

3.3.1 Methods

We performed a timing analysis to measure precisely the parameters describing the pulsar’s evolution over the observation time. We also allowed additional parameters to vary to test for measurable orbital eccentricity and proper motion of the binary. Instead of using a fixed search grid we use a Monte Carlo sampling algorithm to explore the parameter space around the signal parameters detected in the search. The general timing methods are also described by [77, 79], extending the methods developed by Ray et al. [52] and Kerr et al. [97]. We enhanced these methods with the option to marginalize over the parameters of the template pulse profile as described in detail later in this section.

The starting point for the timing procedure is the construction of a template pulse profile, $\hat{g}(\Phi)$, for which we used a combination of N_p symmetrical Gaussian peaks [28]

$$\hat{g}(\Phi) = \left(1 - \sum_{i=1}^{N_p} a_i\right) + \sum_{i=1}^{N_p} a_i g(\Phi, \mu_i, \sigma_i). \quad (3.4)$$

The term $a g(\Phi, \mu, \sigma)$ denotes a wrapped Gaussian peak with amplitude a , peaked at phase μ with width σ :

$$g(\Phi, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{(\Phi + k - \mu)^2}{2\sigma^2}\right). \quad (3.5)$$

The phase at the first peak μ_1 is chosen to be the reference phase for the template. Phases of any other peak i are measured relative to the first peak as phase offset $\mu_i - \mu_1$ to avoid correlation with the overall phase. The template is fit to the weighted pulse profile obtained from the phase-folded data by maximizing over the likelihood

$$\mathcal{L}(\hat{g}, \boldsymbol{\lambda}) = \prod_{j=1}^N [w_j \hat{g}(\Phi(t_j, \boldsymbol{\lambda})) + (1 - w_j)]. \quad (3.6)$$

The Bayesian Information Criterion (BIC) [184] is used to choose the number of peaks by minimizing

$$\text{BIC} = -2 \log(\mathcal{L}(\hat{g}, \boldsymbol{\lambda})) + k \log\left(\sum_{j=1}^N w_j\right), \quad (3.7)$$

where the number of free parameters in the model is denoted by k . Thus, adding a new parameter is penalized by $\log(\sum_{j=1}^N w_j)$ to avoid overfitting. The penalty factor for adding more Gaussian peaks to the template pulse profile scales with $k = 3 \times N_p$ as each peak is described by three parameters.

As described by Clark et al. [79], this template pulse profile is used to explore the multi-dimensional likelihood surface by varying the pulsar parameters with the goal to find the parameter combination that gives the maximum likelihood. We use our own implementation of the Affine Invariant Monte Carlo method described by Goodman and Weare [185] to run many Monte Carlo chains in parallel for the

exploration and the efficient parallelization scheme described by Foreman-Mackey et al. [186]. The computations are distributed over several CPU cores.

This is repeated iteratively. Whenever a new best combination of parameters is found the template is updated using the new timing solution’s phase-folded data. Usually this converges after a few iterations. Additional parameters (e.g., eccentricity) are added one after the other and the described timing procedure is restarted each time. Here again the BIC is used to decide whether the addition of a new parameter significantly improves the pulsar ephemeris. For the timing of bright pulsars [e.g., 79] this iterative approach is sufficient.

For faint pulsars like PSR J0952–0607, the uncertainty in the gamma-ray pulse profile is not negligible. Using a fixed pulse profile template for weak pulsars could lead to systematic biases and underestimated uncertainties in the timing parameters. We therefore treated the template parameters in the same way as the pulsar parameters and let them vary jointly [as also done in 187].

Joint variation of pulsar and template parameters results in larger but more realistic uncertainties on the pulsar parameters but should be used with a caveat. Varying pulsar parameters will always line up photons as close as possible to the same rotational phases to maximize the log-likelihood. The Monte Carlo algorithm finds combinations of parameters that lead to some photons being closer to the maximum of a peak and thus to a higher and narrower peak. But if these parameters do not describe the actual pulsar well, other photons will be shifted to phases outside the range of the peak, leading to a penalty preventing the acceptance of these parameter combinations. The joint variation of pulsar and template parameters however raises the chances of combinations that do not describe the actual pulsar well, as the peak position shifts to the phase where a combination of pulsar parameters leads to a narrow peak. This is a problem for a faint pulsar like PSR J0952–0607 as the penalty factor is weaker due to the smaller amount of photons. Furthermore for a pulsar like PSR J0952–0607 with two close peaks the penalty factor can be reduced by having one broader peak and one very narrow peak.

To address this problem we adjusted our priors on the template parameters. As for the pulsar parameters we used uniform priors for most template parameters. For the width parameters we used log-uniform priors and constrained them to peaks broader than 5% of a rotation, to disfavor extremely narrow peaks which only cover few photons, and narrower than half a rotation (full-width at half maximum $\text{FWHM}_i = 2\sqrt{2\log(2)}\sigma_i$ in the range $0.05 < \text{FWHM}_i < 0.5$). This led to a steadier rise in H statistic over time and a pulse profile similar to what we get when folding the gamma-ray data with the updated radio-timing solution (see Section 3.4.1) reported in Table 4.1. In Figure 3.1 we show 100 pulse profile templates randomly picked from the resulting template parameter distribution.

3.3.2 Solution

Our timing solution is shown in Table 4.1. We did not find clear pulsations in the beginning of the *Fermi* mission at MJD 54,682 and therefore our timing solution starts at MJD 55,750 (see Figure 3.1). We discuss the absence of pulsations prior to MJD 55,750 below.

The gamma-ray pulse profile is likely double peaked as the double-peaked template is favored by the BIC over the single-peaked template. The template param-




Figure 3.1: Integrated pulse profile after MJD 55,750 and phase-time diagram of PSR J0952–0607, showing two identical rotations for clarity. Top: The orange curve indicates the template with the highest BIC. The transparent black curves illustrate 100 representative templates randomly selected from the Monte Carlo samples after the chain stabilized. The histogram shows the weighted photon counts with 30 bins per rotation. The dashed blue line shows the estimated background level. Bottom: Each point represents the rotational phase of a detected gamma-ray photon and its gray scale indicates the probability weight. The dashed-dotted green line denotes the start of our timing solution at MJD 55,750.

eters leading to the highest likelihood are given in Table 4.1.

All of the measured parameters are consistent with the initial published radio solution. The published values and uncertainties on α and δ from the optical counterpart are consistent and comparable to the ones in the gamma-ray timing solution [14]. As expected from the much longer timing baseline the uncertainties on f and P_{orb} are much smaller than in the initial radio-timing solution. Furthermore, it is possible to measure the spin-frequency derivative, $\dot{f}_{\text{obs}} = -2.382(8) \times 10^{-15} \text{ Hz s}^{-1}$. A second spin-frequency derivative, \ddot{f} , is clearly disfavored by the BIC. The gamma-ray timing solution is consistent with an updated radio ephemeris based on radio data spanning 796 days, and the parameter uncertainties are comparable or smaller (see Section 3.4.1 and Table 4.1).

Table 3.2: Properties of PSR J0952–0607 from gamma-ray and radio timing.

Parameter	Gamma-ray	Radio
Span of timing data (MJD)	55750 ^a – 58289	57759 – 58555
Reference epoch (MJD)	57980	57980
Timing Parameters		
R.A., α (J2000.0)	09 ^h 52 ^m 08 ^s .322(2)	09 ^h 52 ^m 08 ^s .32141(5)
Decl., δ (J2000.0)	–06°07′23″.51(4)	–06°07′23″.490(2)
Spin frequency, f (Hz)	707.3144458307(7)	707.31444583103(6)
Spin-frequency derivative, \dot{f}_{obs} (Hz s ^{–1})	–2.382(8) × 10 ^{–15}	–2.388(4) × 10 ^{–15}
Dispersion measure, DM (pc cm ^{–3})		22.411533(11)
Orbital period, P_{orb} (day)	0.267461034(7)	0.2674610347(5)
Projected semi-major axis, x (lt-s)	0.0626670 ^b	0.0626670(9)
Epoch of ascending node, t_{asc} (MJD)	57980.4479516 ^b	57980.4479516(5)
Template Pulse Profile Parameters		
Amplitude of first peak, α_1	0.65(18)	
Phase of first peak, μ_1	0.431(39)	
Width of first peak, σ_1	0.064(23)	
Amplitude of second peak, α_2	0.35(24)	
Phase offset of second to first peak, $\mu_2 - \mu_1$	0.198(27)	
Width of second peak, σ_2	0.040(52)	
Derived Properties (combined results)		
Spin period, P_{obs} (ms)		1.414
Spin-period derivative ^c , \dot{P}_{int} (s s ^{–1})		4.6 × 10 ^{–21}
Characteristic age ^d , τ_c (Gyr)		4.9
Spin-down power ^d , \dot{E} (erg s ^{–1})		6.4 × 10 ³⁴
Surface B -field ^d , B_{surf} (G)		8.2 × 10 ⁷
Light-cylinder B -field ^d , B_{LC} (G)		2.7 × 10 ⁵
Galactic longitude, l (°)		243.65
Galactic latitude, b (°)		+35.38
NE2001 distance, (kpc)		0.97 ^{+1.16} _{–0.53}
YMW16 distance, (kpc)		1.74 ^{+1.57} _{–0.82}
Optical distance, (kpc)		5.64 ^{+0.98} _{–0.91}
Gamma-ray luminosity ^e , L_γ (erg s ^{–1})		3.1 × 10 ³² × ($d/1$ kpc) ²

Notes. — Numbers in parentheses are statistical 1σ uncertainties. The JPL DE405 solar system ephemeris has been used and times refer to TDB. Phase 0 is defined for a photon emitted at the pulsar system barycenter and arriving at the Solar System Barycenter at the reference epoch MJD 57,980.

^a Validity range of timing solution when the data starts at MJD 54,682.

^b Fixed to values from radio-timing solution.

^c Assuming no proper motion, see Section 7.4.

^d Properties are derived as described by Abdo et al. [28] on the basis of the estimated intrinsic spin-period derivative \dot{P}_{int} .

^e Assuming no beaming and distance $d = 1$ kpc.

The upper limit on proper motion corresponds to a transverse velocity of $v_t = \mu_t d = 120 \text{ km s}^{-1} \times (d/1 \text{ kpc})$. This results in high, but not unrealistic transverse velocities when using the distances inferred from the dispersion measure ($d = 0.97 \text{ kpc}$ [188] or $d = 1.74 \text{ kpc}$ [189]). As 90% of the known MSPs in the ATNF Pulsar Catalogue⁷ [110] show transverse velocities below 200 km s^{-1} the proper motion upper limit is unrealistic for the higher distances predicted by the optical observations (4.7–6.6 kpc; see Section 7.4).

Unsurprisingly, we were unable to detect a significant timing parallax. The maximum parallax time delay for the above-mentioned distance estimates is $\Delta t_{\pi, \text{max}} \approx (500 \text{ lt-s})^2 / (2d) \sim 1 \mu\text{s}$. In comparison the resolution with which we can measure the arrival time of the pulse is $\Delta \mu_1 / f \approx 61 \mu\text{s}$.

A circular orbit is clearly favored over an eccentric orbit by the BIC. The 95% upper limit on eccentricity is set to $e < 0.004$.

The missing pulsations before MJD 55,750 seem odd as the tracks are clearly visible later in the mission (Figure 3.1). As the pulsar is not very bright one explanation might be Poisson variations in the flux leading to the loss of pulsations for a few hundred days. Possible pulsations before this period might be too weak to be picked up again as the phase uncertainty grows quickly outside the timing span. At the start of the mission (MJD 54,682) the phase uncertainty is ~ 0.6 rotations, which could be a plausible explanation for loss of coherence.

In order to understand the nature of the non-detection of gamma-ray pulsations before MJD 55750, we measured the gamma-ray flux of PSR J0952–0607 over time by sliding a 750-day-long window in steps of 50 days over the LAT data. In each of these steps we calculated the gamma-ray flux of PSR J0952–0607 over the 750 days width of the window, which allowed us to measure the spectral parameters with reasonable precision. We found that the flux of PSR J0952–0607 is lower in the beginning of the *Fermi* mission but the lower fluxes agree with the flux uncertainties from the full time span. The TS values follow the same trend as the gamma-ray fluxes in the sliding windows.

The gamma-ray source is too faint to test it unambiguously for variability. The windows need to cover 750 days to keep statistical precision. But that leaves only five independent time bins to calculate the variability index with Equation (4) from [35]. The variability index computed with these five bins is 7.18 with 4 degrees of freedom, which is below the 99% confidence level of 13.277.

We also checked if the smaller 35° rocking angle used during the first year of the *Fermi* mission decreases the pulsation significance. However, the small rocking angle is actually favorable as the exposure for PSR J0952–0607 is $\sim 20\%$ higher in the beginning of the mission.

Variations of the orbital period might be another reasonable explanation for the loss of clear pulsations. Such orbital-period variations have been measured for several spider pulsars, e.g. for the original black-widow pulsar PSR B1957+20 [98]. Nevertheless the penalty for adding orbital-period derivatives led to an increase in the BIC. Similarly, no significant semi-major-axis derivative was found.

⁷<http://www.atnf.csiro.au/research/pulsar/psrcat>

3.4 Multi wavelength

3.4.1 Updated Radio Timing

Observations of PSR J0952–0607 with LOFAR have been ongoing using an identical observational setup as in [14], namely a single tied-array beam formed from the high-band antennas (HBAs) of the central 23 LOFAR [171] core stations, using 78 MHz of bandwidth at a central frequency of 149 MHz. Before 2018 May, several 5 min integrations were obtained at each observing epoch; after that the integration times were increased to 20 min. These observations were obtained at a roughly monthly cadence. As described in [14], these observations were coherently de-dispersed, folded with DSPSR [190], and analyzed using tools in the PSRCHIVE software suite [191] and the TEMPO2 pulsar-timing software [95, 96].

The phase-connected timing solution from Bassa et al. [14] was improved by using all LOFAR HBA observations that used 78 MHz of bandwidth (hence excluding the discovery and initial follow-up observations which used half the bandwidth). Pulse time-of-arrival (TOA) measurements were obtained by referencing pulse profiles of eight frequency channels per observation to a single analytic pulse profile template. This procedure presumes that our data are not sensitive to pulse profile shape variations with frequency, which was double-checked through inspection of the difference profiles of the top and bottom parts of the bandpass: no significant structures were detected. The analytic pulse profile was created using the PSRCHIVE [192] package PAAS and was constructed from five von Mises functions that were fitted to the integrated body of observations and fully modeled any detectable pulse shapes. The resulting timing solution extends the timing baseline to 2.2 years and breaks the degeneracy between the astrometric and rotational parameters (see Table 4.1). Upon inspection of the data, a new covariance was detected, namely, between a significant ($> 4\sigma$) decrease in the dispersion measure of this pulsar (which was found to be decreasing by $5 \times 10^{-5} \text{ pc cm}^{-3} \text{ yr}^{-1}$) and the spin period. Notwithstanding the significance of this decrease, the strong anticorrelation of this parameter with the pulse period suggests an underestimate of its measurement significance, which is commonly found in pulsar-timing analyses [e.g., 193], particularly in non-periodic parameters such as linear gradients in dispersion measure. Consequently this decrease was not included in our present analysis, but future monitoring to allow more robust disentanglement of the spin period and the dispersion measure variability is warranted. We find no evidence for radio eclipses in the six LOFAR observations with orbital phases between $0.15 < \phi_{\text{orb}} < 0.35$. Using the TOAs from this orbital phase range we set a 3σ upper limit on time delays due to additional dispersion of $\Delta t < 2.3 \mu\text{s}$, and hence $\Delta\text{DM} < 1.2 \times 10^{-5} \text{ pc cm}^{-3}$.

3.4.2 Optical Photometry

Bassa et al. [14] presented an r' -band light curve of the optical companion to PSR J0952–0607 taken by the WFC on the 2.5m Isaac Newton Telescope on La Palma. The orbital light curve features a single maximum peaking at $r' \approx 22$ at the pulsar’s inferior conjunction, interpreted as being due to the pulsar heating the inside face (the “dayside”) of a tidally locked companion. Bassa et al. [14] modeled this light curve with the Icarus package [194], finding that PSR J0952–0607 is likely to have an inclination angle $i \sim 40^\circ$, but the lack of color information precluded a

Table 3.3: New optical photometry of the companion of PSR J0952–0607.

Night beginning	Instrument+Telescope	Filters	ϕ_{orb}	Airmass	Seeing	Photometric
2018 June 03	ULTRACAM+NTT	u_s, g_s, i_s	0.64–1.09	1.1–2.1	1.0''–2.0''	yes
2018 June 04	ULTRACAM+NTT	u_s, g_s, i_s	0.37–0.71	1.1–1.6	1.0''–3.0''	no
2019 January 12	HiPERCAM+GTC	u_s, g_s, r_s, i_s, z_s	0.77–0.92	1.25–2.0	< 1.5''	yes
2019 January 13 ^a	HiPERCAM+GTC	u_s, g_s, r_s, i_s, z_s	0.37–0.72	1.25–2.0	1.5''–2.0''	no
2019 March 02 ^b	ULTRACAM+NTT	u_s, g_s, i_s	0.91–1.29	1.1–1.6	0.8''–1.2''	no
2019 March 03	ULTRACAM+NTT	u_s, g_s, i_s	0.72–0.88	1.2–1.4	1.2''–2.4''	no
			1.16–1.72	1.1–1.9		

Notes. — Orbital phases are in fractions of an orbit, with $\phi_{\text{orb}} = 0$ corresponding to the pulsar’s ascending node. The ULTRACAM data from 2018 were taken as a series of 20 s exposures in g_s and i_s , and 60 s in u_s . The 2019 ULTRACAM observations were taken with 10 s exposures in g_s and i_s , and 30 s in u_s . The HiPERCAM data cover u_s, g_s, r_s, i_s , and z_s simultaneously with exposure times of 60 s in u_s, g_s, r_s , and 30 s in i_s and z_s .

^a During an episode around $\phi_{\text{orb}} = 0.6$ seeing reached over 2.3'' and 20 exposures had to be removed.

^b We removed several frames due to intermittent clouds during the observations when the transmission dropped to nearly zero.

robust estimate of other system parameters (e.g. companion temperature, heating, companion radius).

To more fully investigate the optical counterpart to PSR J0952–0607, we obtained multi-color photometry using ULTRACAM [195] on the 3.58m New Technology Telescope (NTT) at ESO La Silla, and HiPERCAM [196, 197] on the 10.4m Gran Telescopio Canarias (GTC) on La Palma. The observation specifics are given in Table 3.3.

These data were calibrated and reduced using the ULTRACAM⁸ and HiPERCAM⁹ software pipelines. Standard CCD calibration procedures were applied using bias and flat field frames taken during each run.

We extracted instrumental magnitudes using aperture photometry, and performed “ensemble photometry” [198] to correct for airmass effects and varying transparency. Magnitudes in g_s, r_s, i_s , and z_s ¹⁰ were calibrated using comparison stars chosen from the Pan-STARRS1 [199] catalog, after fitting for a color term accounting for differences between our filter sets and the Pan-STARRS1 filters. The HiPERCAM u_s observations were flux calibrated using zero-points derived from observations of two Sloan Digital Sky Survey (SDSS) standard stars [200] taken on 2019 January 11. The resulting HiPERCAM magnitudes for three nearby stars to PSR J0952–0607 were used to flux calibrate the ULTRACAM u_s data. Finally, the airmass- and ensemble-corrected count rates (C) were converted to AB flux densities according to our measured zero-point counts in each frame (C_0) by $S_{\text{AB}} = 3631 (C/C_0)$ Jy.

⁸<http://deneb.astro.warwick.ac.uk/phsaap/software/ultracam/html/>

⁹<http://deneb.astro.warwick.ac.uk/phsaap/hipercam/docs/html/>

¹⁰ULTRACAM and HiPERCAM use identical higher-throughput versions of the SDSS filter set, which we refer to as *Super-SDSS filters*: u_s, g_s, r_s, i_s , and z_s [197].

3.4.3 Optical Light-Curve Modeling

As in [14], the *Icarus* software was used to estimate parameters of the binary system. To do this, we performed a Bayesian parameter estimation using the nested sampling algorithm `MultiNest` [201] via the Python package `PyMultiNest` [202]. *Icarus* produces model light curves by computing a grid of surface elements covering the companion star, and calculating and summing the projected line-of-sight flux from each element. Here the flux from each surface element was computed by integrating spectra from the Göttingen Spectral Library models of Husser et al. [203].

In these fits we assumed that the companion star is tidally locked to the pulsar, and varied the following parameters: the companion star’s “nightside” temperature (T_n); the “irradiating temperature” (T_{irr} defined such that the dayside temperature $T_d^4 = T_{\text{irr}}^4 + T_n^4$, under the assumption that the pulsar’s irradiating flux is immediately thermalized and re-radiated, and therefore simply adds to the companion star’s intrinsic flux at each point on the surface, as in [176]); the binary inclination angle (i); the Roche-lobe filling factor (f_{RL} , defined as the ratio between the companion’s radius towards the pulsar and the inner Lagrange point (L1) radius); the distance modulus ($\mu = 5 \log_{10}(d) - 5$), with distance d in pc; and the mass of the pulsar (M_{psr}). At each point, the companion mass (M_c) and mass ratio ($q = M_{\text{psr}}/M_c$) were derived from the binary mass function according to the timing measurements of P_{orb} and x presented in Table 4.1. We also marginalize over interstellar extinction and reddening, parameterized by the $E(B - V)$ of Green et al. [204], scaled using the coefficients given therein for Pan-STARRS1 filter bands. We adopted a Gaussian prior for $E(B - V)$ (truncated at zero), using the value from Green et al. [204] for $d > 1$ kpc in the direction of PSR J0952–0607, $E(B - V) = 0.065 \pm 0.02$, found by fitting the line-of-sight dust distribution using the apparent magnitudes of nearby main-sequence stars in the Pan-STARRS1 catalog. We adopted uniform priors on the remaining parameters (and uniform in $\cos i$), with M_{psr} and f_{RL} limited to lie within $1.2 < M_{\text{psr}} < 2.5 M_{\odot}$, and $0.1 < f_{\text{RL}} < 1$. Temperatures T_n and T_d were constrained to lie within the range covered by the atmosphere models, $2300 < T < 12000$ K.

At each point in the sampling, *Icarus* computed model light curves in each band. To account for remaining systematic uncertainties in the flux calibration, extinction, and atmosphere models, the model light curve in each band was re-scaled at each parameter location to maximize the penalized chi-squared log-likelihood. Overall calibration offsets were allowed for each band, and penalized by a zero-mean Gaussian prior on the scaling factor in each band with a width of 0.1 mag (a conservative estimate based on our calibration to the Pan-STARRS1 magnitudes). We also allowed small offsets between the calibrations for each ULTRACAM run and the HiPERCAM observations, which we penalized with an additional Gaussian prior with width 0.05 mag (also a conservative estimate from the differences in magnitudes of comparison stars in the field of view on each night). In initial fits, our best-fitting model resulted in a reduced chi-squared greater than unity. We therefore also re-scaled the uncertainties in each band to maximize the (re-normalized) log-likelihood at each point in the sampling. We also found that the fit improved substantially when we fit for a small orbital phase offset. Such orbital phase offsets are often seen in the optical light curves of black-widow pulsars and have been interpreted as being due to asymmetric heating from the pulsar, which could be caused by reprocessing of the pulsar wind by an intra-binary shock [e.g., 205].




Figure 3.2: Optical light curve of the companion to PSR J0952–0607, phased using the gamma-ray timing ephemeris. For clarity, the HiPERCAM and ULTRACAM fluxes have been combined into 180 and 300 s time bins, respectively, via weighted average. The unbinned data were used for the light-curve model fitting. Dashed and solid curves show the flux in each band as predicted by the best-fitting *Icarus* model before and after allowing for uncertainties in the flux calibrations (see text), respectively.

The best-fitting light-curve model is shown in Figure 3.2, with posterior distributions for the fit parameters shown in Figure 3.3.

3.4.4 Search for Continuous Gravitational Waves

We carried out a search for near-monochromatic continuous gravitational waves phase locked at twice the pulsar rotation phase for the source PSR J0952–0607 using data from the first and second runs (O1¹¹ and O2¹²) of the two Advanced LIGO detectors [206]. The observation period spans 707 days from 2015 September to 2017 August and comprises 183 days (169 days) of data from the Hanford (Livingston) detector.

We employ the coherent multi-detector detection statistic $2\mathcal{F}$ [207, 208] that we implemented in the LIGO-LALSUITE library¹³. $2\mathcal{F}$ is the log-likelihood maximized over the amplitude parameters h_0 , $\cos \iota$, ψ and Φ_0 for a near-monochromatic¹⁴ gravitational wave signal with given frequency and frequency-derivative values, from a source in a binary at a given sky position and with given orbital parameters, in Gaussian noise. h_0 is the intrinsic gravitational wave amplitude at the detector, ι the angle between the total angular momentum of the pulsar and the line of sight to it from Earth, ψ is the gravitational wave polarization angle and Φ_0 the signal phase at a nominal reference time. In this search we assume the gravitational wave frequency and frequency derivatives equal to twice the values measured for the pulsar rotation frequency and its derivatives. In Gaussian noise the detection statistic $2\mathcal{F}$ follows a χ^2 -distribution with 4 degrees of freedom and non-centrality parameter

¹¹<https://doi.org/10.7935/K57P8W9D>

¹²<https://doi.org/10.7935/CA75-FM95>

¹³<https://git.ligo.org/lscsoft/lalsuite/>

¹⁴The signal is not strictly monochromatic because of the measured non-zero spin-frequency derivative.




Figure 3.3: Posterior distributions for optical light-curve modeling parameters. The last three parameters (companion mass M_c , volume-averaged density ρ and heating efficiency ε) were derived from the values of the other fit parameters and the gamma-ray timing ephemeris. Dashed vertical lines on histograms indicate the posterior mean and 95% confidence interval. Where non-uniform priors were assumed, these are shown by red curves on the one-dimensional histograms. Contour lines indicate 1σ and 2σ confidence regions, with individual samples outside these areas shown as points weighted by their posterior probability.

equal to 0: the expected value is $\mu = 4.0$, and the standard deviation is $\sigma = 2\sqrt{2}$. If a signal is present, the non-centrality parameter is proportional to the square of the intrinsic gravitational wave amplitude at the detector, h_0 , and to the total observation time.

The search yields the value $2\mathcal{F} = 9.9$, which is well within the bulk of the distribution consistent with a null result. Based on the measured value of the detection statistic, we set a frequentist 95% upper limit on the intrinsic gravitational wave amplitude, $h_0^{95\%}$, following a now standard procedure first developed by some of us [209]. $h_0^{95\%}$ is the smallest intrinsic gravitational wave amplitude such that 95% of the population of signals that could be emitted by PSR J0952–0607¹⁵ would yield a detection statistic value greater than the measured one, $2\mathcal{F} = 9.9$. We find $h_0^{95\%} = 6.6 \times 10^{-26}$. The uncertainty on this upper limit is $\sim \pm 14\%$, including instrument calibration errors [210].

3.5 Discussion

The pulsar’s spin period is defined as $P = 1/f$ and the spin-period derivative is $\dot{P} = -\dot{f}/f^2$. The observed spin period for PSR J0952–0607 from gamma-ray and radio timing is $P_{\text{obs}} = 1.414$ ms and the observed spin-period derivative is $\dot{P}_{\text{obs}} = 4.76 \times 10^{-21} \text{ s s}^{-1}$.

The intrinsic spin-period derivative \dot{P}_{int} can be estimated from the observed value $\dot{P}_{\text{obs}} = \dot{P}_{\text{int}} + \dot{P}_{\text{Gal}} + \dot{P}_{\text{Shk}}$. \dot{P}_{Gal} represents the part of the spin-period derivative caused by the relative Galactic acceleration [differential Galactic rotation and acceleration due to the Galactic gravitational potential; e.g., 211, 212], while \dot{P}_{Shk} accounts for the Shklovskii effect due to non-zero proper motion [213]. Both contributions, \dot{P}_{Gal} and \dot{P}_{Shk} , depend on the distance d to the pulsar.

The distance to PSR J0952–0607 is uncertain. The measured dispersion measure (DM) can be used to estimate the distance using Galactic electron-density models. The NE2001 model predicts $0.97_{-0.53}^{+1.16}$ kpc, while the YMW16 model predicts $1.74_{-0.82}^{+1.57}$ kpc. The uncertainties represent the 95% confidence regions [189]. The model predictions of the DM as a function of d in the direction of the pulsar’s sky position are shown in Figure 3.4. The models saturate at DM values that differ by $\sim 30\%$ indicating the challenge and difficulty modeling the Galactic electron density. Still the distance predictions are consistent within the large uncertainty. On the other hand, the distance derived from optical modeling is $5.64_{-0.91}^{+0.98}$ kpc. This disagrees strongly with both DM distances and suggests that both DM models are overestimating the electron density in the direction of PSR J0952–0607. The distance discrepancy is discussed in more detail below.

The estimated Galactic contribution is $\dot{P}_{\text{Gal}} = (1.7, 2.2, 3.6) \times 10^{-22} \text{ s s}^{-1}$ for the distance estimates $d = (0.97, 1.74, 5.64)$ kpc. For the Shklovskii effect we then find the 95% confidence region to $\dot{P}_{\text{Shk}} \in ([0, 2.1], [0, 3.8]) \times 10^{-21} \text{ s s}^{-1}$ from the proper motion 95% confidence region (see Section 3.3.2) and for the (NE2001, YMW16) distances. The resulting 95% confidence region on \dot{P}_{Shk} for the optical distance exceeds past \dot{P}_{obs} . Thus we only constrain the intrinsic spin-frequency derivative (at 95% confidence) to $\dot{P}_{\text{int}} \in [2.44, 4.59] \times 10^{-21} \text{ s s}^{-1}$ for the NE2001 model and $\dot{P}_{\text{int}} \in [0.69, 4.54] \times 10^{-21} \text{ s s}^{-1}$ for the YMW16 model. In the following, we conservatively

¹⁵The possible signals span uniformly distributed values of $-1 \leq \cos \iota \leq 1$ and of $0 \leq \psi \leq 2\pi$.




Figure 3.4: Dispersion measure versus distance from the NE2001 and YMW16 models at the sky position of PSR J0952–0607. For the measured $DM = 22.4 \text{ pc cm}^{-3}$ (black, horizontal line) the NE2001 model (dotted, blue line) and the YMW16 model (dashed, orange line) predict distances of 0.97 kpc and 1.74 kpc, respectively. The 95% confidence regions around those values are calculated as 120% (NE2001) and 90% (YMW16) “relative” errors on the predicted values [189]. To illustrate the discrepancy with these distance predictions, the 95% confidence region from the optical modeling is shown. The vertical, dashed-dotted line indicates the distance favored by the optical modeling.

assume zero proper motion (i.e. $\dot{P}_{\text{Shk}} = 0$) and used the fastest possible spin-down rate, $\dot{P}_{\text{int}} = 4.6 \times 10^{-21} \text{ s s}^{-1}$.

In Figure 4.3, PSR J0952–0607 is shown in a P - \dot{P} diagram with the known pulsar population outside of globular clusters. The spin parameters of the more than 2000 radio pulsars are taken from the ATNF Pulsar Catalogue² [110].

Furthermore we estimated the characteristic age τ_c , the spin-down power \dot{E} , the surface magnetic field strength B_{surf} and the magnetic field strength at the light cylinder B_{LC} (see Table 4.1). To calculate these values we assumed the pulsar to be a magnetic dipole with a canonical radius $r_{\text{psr}} = 10 \text{ km}$ and moment of inertia $I_{\text{psr}} = 10^{45} \text{ g cm}^2$ [e.g., 28]. The same assumptions were used to plot the contour lines in Figure 4.3.

Despite spinning so rapidly, the gamma-ray energy flux of PSR J0952–0607 is on the fainter end of the gamma-ray MSP population. There are several reasons why gamma-ray pulsars might appear faint, including large distance, high background, or low luminosity [182]. PSR J0952–0607 is not in a high-background region. The large distance derived from the optical modeling could be a possible explanation but disagrees with the distance estimates derived from the dispersion measure, $d = (0.97, 1.74) \text{ kpc}$ (NE2001, YMW16). The inferred gamma-ray luminosity is $L_\gamma = 4\pi d^2 F_\gamma f_\Omega \approx 3.1 \times 10^{32} \times (d/1 \text{ kpc})^2 \text{ erg s}^{-1}$. The measured LAT energy flux F_γ is given in Table 3.1 and we assumed no beaming (i.e., $f_\Omega = 1$). The gamma-ray efficiency is $\eta_\gamma = L_\gamma / \dot{E} \approx 0.5\% \times (d/1 \text{ kpc})^2$. At the optical distance, $\eta_\gamma \approx 16\%$ is typical of gamma-ray MSPs [28], while at the DM-derived distance, $\eta_\gamma \sim 1\%$ would be unusually low.




Figure 3.5: Spin period P and spin-period derivative \dot{P} of the known pulsar population outside of globular clusters. The inset shows a zoomed-in view of the known MSP population. Isolated radio pulsars (light-gray pluses), binary radio pulsars (dark-gray squares), isolated gamma-ray pulsars (light-green crosses) and binary gamma-ray pulsars (dark-green circles) are shown. The subject of this paper, the gamma-ray pulsar PSR J0952–0607, is marked by an orange star. The lines denote constant characteristic age τ_c (dotted), spin-down power \dot{E} (dashed) and surface magnetic field strength B_{surf} (dashed-dotted).

Due to the non-detection of PSR J0952–0607 in X-rays [$F_X < 1.1 \times 10^{-13} \text{ erg s}^{-1} \text{ cm}^{-2}$, 14] we can only give a lower limit for the gamma-ray-to-X-ray flux ratio $F_\gamma/F_X > 20$. This limit is at the lower end of the observed distribution but still consistent with the literature [28, 90, 214, 215].

The peak of the observed optical light curve is fairly broad in orbital phase. This requires either low inclination such that part of the heated face of the companion is visible over a large range of orbital phases, or for the companion to be close to filling its Roche lobe, such that the tidal deformation results in an “ellipsoidal” component peaking at $\phi_{\text{orb}} = 0.5$ and $\phi_{\text{orb}} = 1.0$ (with $\phi_{\text{orb}} = 0$ corresponding to the pulsar’s ascending node) where the visible surface area of the companion is largest. Our best-fitting Icarus model favors the latter explanation, with $f_{\text{RL}} \approx 88\%$ and $i \approx 61^\circ$. However, high filling factors imply a larger and hence more luminous companion, and therefore require greater distance, with our model having $d \sim 4.7\text{--}6.6$ kpc.

We tried to re-fit the optical light curve with the distance fixed at the YMW16 distance of $d = 1.74$ kpc, but the resulting model has a significantly worse fit, and the low filling factor required results in an extremely high volume-averaged density for the companion (ρ) in excess of 100 g cm^{-3} . For comparison, the densest known

black-widow companions have densities of around 50 g cm^{-3} [e.g., PSR J0636+5128, 216], with the record being that of the black-widow candidate 3FGL J1653.6–0158 in a 75-min orbit [111] where $\rho \gtrsim 70 \text{ g cm}^{-3}$. These objects have been proposed to be the descendants of ultra-compact X-ray binaries, but this origin is unlikely for PSR J0952–0607 given its much longer orbital period [217]. If the DM distances are assumed, the required density suggests that the companion star consists mostly of degenerate matter. A low filling factor may also explain the absence of radio eclipses seen from PSR J0952–0607. Alternatively, the low-density, large-distance solution has $\rho \sim 2.75 \text{ g cm}^{-3}$, close to the density of brown dwarfs of similar mass and temperature given by the model considered in Kaplan et al. [216].

We note that similar discrepancies in model distances were seen by Sanchez and Romani [205] when using a direct-heating model. Romani and Sanchez [218] and Sanchez and Romani [205] considered models that additionally include a contribution from reprocessing of the pulsar wind by an intra-binary shock, which can wrap around the companion star. This can produce broader light curves for lower filling factors as some heating flux is re-directed further around the sides of the companion star, and can also explain the small phase offset required for our direct-heating model by asymmetry in the shock front. Such a model may improve the fit for lower distances and filling factors, although an extremely high companion density would still be required to match the YMW16 distance. A likely explanation therefore could be that some heating flux is reprocessed by a shock, and the system has a moderate distance and filling factor, somewhat larger than required by the YMW16 value, but below those predicted by our direct-heating model. While more complex irradiation models [e.g., 218] may be required to address this issue, a full investigation of alternative models is beyond the scope of this study.

In both the small and large distance cases, we find that the nightside temperature of the companion is $T_n \approx 3000 \pm 250 \text{ K}$ at 95% confidence. We also find a well-constrained irradiating temperature of $T_{\text{irr}} = 6100 \pm 350 \text{ K}$, higher than that found from the single-band fit performed in Bassa et al. [14]. This heating parameter can be compared to the total energy budget of the pulsar by calculating the “efficiency”, ε , of conversion between spin-down power (\dot{E}) and heating flux [176]

$$\varepsilon = \frac{4\pi A^2 \sigma T_{\text{irr}}^4}{\dot{E}}, \quad (3.8)$$

with $\varepsilon \sim 20\%$ being typical for black-widow systems. The efficiency is also shown in Figure 3.3, calculated from T_{irr} and from the orbital separation ($A = x(1+q)/\sin i$) at each point. We find that heating represents a larger fraction of the pulsar’s total energy budget ($\varepsilon \sim 22\%$ to 48% with 95% confidence) than the observed gamma-ray emission $\eta_\gamma \approx 0.5\% \times (d/1\text{kpc})^2$. This estimate assumes that the pulsar’s heating flux is emitted isotropically. As pointed out by Draghis and Romani [219], some models of pulsar gamma-ray emission predict stronger beaming towards the pulsar’s rotational equator, and an MSP’s rotation should be aligned with the orbital plane as a result of the spin-up process. The actual gamma-ray luminosity directed towards the companion may therefore be higher than we observe. Our optical fits suggest a relatively face-on inclination (further evidenced by the lack of eclipses observed in radio observations, which often occur far outside the companion’s Roche lobe), and so the comparative faintness of the pulsar’s observed gamma-ray emission could be explained by the large viewing angle, and the fact that flux is preferentially

emitted in the equatorial plane. A full modeling of the pulsar’s phase-aligned radio and gamma-ray pulse profiles would provide an additional test of this scenario by estimating the viewing and magnetic inclination angles, and the relative beaming factors along our line of sight and in the equatorial plane. So far this is inhibited by the low significance of the gamma-ray light curve but with the continuing LAT mission this might be possible with more gamma-ray data in the future.

Alternatively, the difference between the heating flux and gamma-ray emission may suggest that another mechanism, e.g., the pulsar wind or intra-binary shock heating [218, 220], is responsible for heating the companion. Indeed, there is evidence for this being the case for the transitional PSR J1023–0038 where the optical heating is apparently unchanged between the MSP and low-mass-X-ray-binary (LMXB) states [221] despite a $5\times$ increase in the gamma-ray flux [22].

As the optical counterpart to PSR J0952–0607 is faint (peaking at $r' \approx 22$), it will be difficult to improve upon this picture of the system. While it may be possible to improve upon the dayside temperature measurement with optical spectroscopy in the future, the companion is effectively undetectable at minimum ($r' > 25.0$), precluding optical spectroscopic measurements of the companion’s nightside temperature. We are also unable to constrain the mass of PSR J0952–0607 using the optical data. Constraining the pulsar mass would require a precise measurement of the binary mass ratio, which can be obtained for black-widow systems by comparing the radial velocities of the pulsar and companion. Unfortunately, the optical counterpart of PSR J0952–0607 is too faint ($r' \sim 23$ at quadrature when the radial velocity is highest) for spectroscopic radial velocity measurements to be feasible even with 10 m class telescopes.

The gamma-ray source shows no significant variability as all flux measurements are consistent with the mean flux level. The calculated variability index also indicates a non-varying source. Here it is important to note that due to the low flux of the source the time bins had to be 750 days long to keep statistical precision. Therefore the variability index was calculated from only five independent time bins. Variations on shorter timescales can also not be found this way.

The gamma-ray pulse profile of PSR J0952–0607 shows two peaks that are separated by $\mu_2 - \mu_1 \approx 0.2$ rotations. This is typical for gamma-ray MSPs. More than half of them are double peaked with a peak separation of 0.2 – 0.5 rotations [28]. The radio pulse profile also shows two peaks with similar separation, with the radio pulse slightly leading the gamma-ray pulse (see Figure 3.6). The phase lag between the gamma-ray and radio pulse profile seems to be ~ 0.15 [the majority of two-peaked MSPs show phase lags of 0.1 – 0.3; 28]. Due to a covariance between f and dispersion measure (see Section 3.4.1) we were not able to measure significant variations in the dispersion measure. A change in dispersion measure of $10^{-3} \text{ pc cm}^{-3}$ over the course of the *Fermi* mission would lead to an error in the phase offset of 13%.

Gamma-ray pulsars are a good way to identify the maximum spin frequency of neutron stars. Among the ten fastest Galactic field pulsars only one pulsar has not been detected in gamma rays. Until the discovery of the 707 Hz pulsar PSR J0952–0607, the first MSP, PSR B1937+21, and the first black-widow pulsar, PSR B1957+20, were the fastest-spinning gamma-ray pulsars known [108]. Still, the mass-shedding spin limit for neutron stars is typically placed much higher at around 1200 Hz [222, 223]. One mechanism that could prevent neutron stars from spinning




Figure 3.6: Aligned integrated gamma-ray and radio pulse profiles of PSR J0952–0607 over two identical rotations. The black curve shows the weighted LAT photon counts after MJD 55,750 in a histogram with 30 bins per rotation. The green error bars show the phase uncertainty of the gamma-ray pulse profile. The estimated background level is indicated by the dashed blue line. The radio profile as seen by the LOFAR telescope in a 78 MHz band centered at 149 MHz is drawn in red. The error bars drawn in dark red indicate the possible phase shift of the radio pulse profile due to a dispersion measure variation of 10^{-3} pc cm $^{-3}$ over the time span of the *Fermi* mission.

up to higher frequencies is the emission of gravitational waves [for a recent work on this subject see, e.g., 224]. Another option could be that the spin-up torque might be smaller for faster pulsars with lower magnetic field strengths [225, 226].

The estimated intrinsic spin-period derivative implies a very low surface magnetic field of 8.2×10^7 G for PSR J0952–0607. Assuming non-zero proper motion would result in an even lower surface magnetic field estimate. Just nine pulsars, including the gamma-ray pulsar with the lowest surface magnetic field in the ATNF Pulsar Catalogue² [110], PSR J1544+4937 [60], show lower inferred surface magnetic fields (Figure 3.7). The surface B -field of the other recent LOFAR-detected pulsar, PSR J1552+5437, is only slightly stronger [63]. This might be a hint that pulsars with low B -fields also have steeper radio spectra.

The pulsar distribution in Figure 3.7 indicates a lower limit on the magnetic field strength independent of the spin frequency. The equilibrium spin period as predicted by [16] is $P_{\text{eq}} \propto B_{\text{surf}}^{6/7} R_{\text{psr}}^{18/7} M_{\text{psr}}^{-5/7} \dot{M}_{\text{accr}}^{-3/7}$ with pulsar radius R_{psr} , mass M_{psr} , and accretion rate \dot{M}_{accr} , which indicates that the lowest spin periods can be reached for low magnetic field strengths and high accretion rates. Nevertheless high accretion rates lead to a rapid decrease of the magnetic field strength and for low




Figure 3.7: Frequency f and surface magnetic field strength B_{surf} of the known MSP population outside of globular clusters. The surface magnetic field of PSR J0952–0607 is computed assuming $\dot{P}_{\text{shk}} = 0$ and thus represents an upper limit. The horizontal dashed blue line represents a possible minimum magnetic field strength. The three red lines are so-called spin-up lines for different accretion rates. Left panel: The markers are defined as in Figure 4.3. Right panel: Histogram with 40 bins between 3.3×10^7 G and 9×10^8 G, showing the inferred surface magnetic field strengths for the known MSP population (gray) and also the subset of LAT pulsars (green).

magnetic field strengths the angular momentum transfer is slower [226]. In order to spin up to millisecond periods a limiting magnetic field strength and accretion rate can be set as a result of the amount of time a neutron star can spend accreting matter being limited by the age of the universe [227]. For a neutron star with a mass of $1.4 M_{\odot}$, a radius of 10 km and a minimum accretion rate of $7.26 \times 10^{-11} M_{\odot} \text{ yr}^{-1}$ we get a minimum magnetic field strength of $B_{\text{surf}} \gtrsim 3.3 \times 10^7$ G, which is consistent with the observed pulsar population.

No continuous gravitational waves are detected from PSR J0952–0607, which is to date the fastest-spinning pulsar targeted for gravitational wave emission. The 95% upper limit on the intrinsic gravitational wave amplitude is set to $h_0^{95\%} = 6.6 \times 10^{-26}$. The corresponding upper limit on the ellipticity is $\epsilon^{95\%} = 3.1 \times 10^{-8} \times (d/1 \text{ kpc}) \times (10^{45} \text{ g cm}^2/I)$, where I is the principal moment of inertia of the pulsar. The intrinsic gravitational wave amplitude at the detector needed to account for all of the spin-down energy lost due to gravitational wave

emission is $h_0^{\text{sd}} = 1.5 \times 10^{-27} \times (1 \text{ kpc}/d) \times (I/10^{45} \text{ g cm}^2)^{1/2}$, corresponding to an ellipticity of $\epsilon^{\text{sd}} = 7.0 \times 10^{-10} \times (1 \text{ kpc}/d)$.

As for many other high-frequency pulsars, the indirect spin-down upper limit on h_0 is smaller and more constraining than our measured gravitational wave upper limit, in this case by a factor of ≈ 45 at 1 kpc. For a more likely larger distance the factor would be even greater, so it is not surprising that a signal was not detected [228]. The quoted spin-down upper limit could be inaccurate if the measured spin down were affected by radial motions, if the distance were smaller than estimated or if the moment of inertia of the pulsar were different than the fiducial value of 10^{45} g cm^2 . In the case of PSR J0952–0607 it is unlikely that all these effects could bridge a gap of nearly two orders of magnitude, but in line with the “eyes-wide-open” spirit of previous searches for gravitational waves from known pulsars (see [147, 228, 229] and references therein) we all the same perform the search.

3.6 Conclusions

Using a sensitive, fully coherent pulsation search technique, we detected gamma-ray pulsations from the radio pulsar PSR J0952–0607 in a search around the parameters reported by [14]. New timing methods were developed to cope with the low signal strength, allowing us to measure the spin rate, sky position, and orbital period with high precision, and in agreement with the updated radio-timing ephemeris. Furthermore thanks to the longer gamma-ray time span we reliably constrained the intrinsic spin-period derivative $\dot{P}_{\text{int}} \lesssim 4.6 \times 10^{-21} \text{ s s}^{-1}$. This measurement provides estimates of physical parameters such as the spin-down luminosity ($\dot{E} \lesssim 6.4 \times 10^{34} \text{ erg s}^{-1}$), and a surface magnetic field ($B_{\text{surf}} \lesssim 8.2 \times 10^7 \text{ G}$) among the lowest of any detected gamma-ray pulsar. Although the resulting timing solution spans 7 years to the present data, we were unable to extend this to cover data earlier than MJD 55,750. We investigated several possible reasons. Flux variations could lead to the loss of pulsations. A time-varying orbital period as seen in several spider pulsars would cause a loss of phase coherence. With our current data we are not able to ascertain the true reason. In the absence of orbital-period variations or state changes, improved timing precision from additional data should help determine the cause.

We also obtained new multi-band photometry of the pulsar’s optical counterpart, and modeled the resulting light curve. To explain the observed optical flux, our models require either a much larger distance ($\sim 5 \text{ kpc}$) than the DM-distance estimates of 0.97 kpc (NE2001) to 1.74 kpc (YMW16), or a small and extremely dense companion $\rho \gg 100 \text{ g cm}^{-3}$. More complex optical models including intra-binary shocks might help to solve this discrepancy, but a full investigation of other models is beyond the scope of this work. We found that the pulsar flux heating the companion star accounts for a much larger fraction of the pulsar’s spin-down power ($\sim 50\%$) than is converted to observed gamma-ray emission (0.5% at 1 kpc), although this difference is reduced if our larger distance estimate is adopted.

Despite the extensive analysis of PSR J0952–0607 and its companion, the study of this pulsar has not ended as some questions remain unanswered. The LAT and LOFAR continue to take gamma-ray and radio data on this source, and we plan to obtain more optical data.

LAT gamma-ray data has helped to find many new MSPs by providing promising

candidates [56]. Sophisticated methods to identify more pulsar candidates within LAT sources have been developed [e.g., 66, 67]. For instance, [230] identified 11 promising MSP candidates by checking for steep-spectrum radio sources coincident with LAT sources. With the approach successfully used in this paper, new binary MSP candidates can be searched for pulsations and upon detection the pulsar can be precisely timed within months after its discovery. Identifying more of the rapidly rotating spider pulsars will be helpful to study further the observed neutron star parameter limits like the maximum spin frequency and the minimum surface magnetic field strength.

Acknowledgments

We thank the referee for pointing out the uncertainties in the DM/distance models, and suggesting the arguments given around Figure 3.4. This work was supported by the Max-Planck-Gesellschaft (MPG) and the ATLAS cluster computing team at AEI Hannover. C.J.C., R.P.B., and D.M.-S. acknowledge support from the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 715051; Spiders). C.G.B. and J.W.T.H. acknowledge support from the European Research Council (ERC) under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement No. 337062 (DRAGNET; PI: Hessels). This work was supported by an STSM Grant from COST Action CA16214. M.R.K. is funded through a Newton International Fellowship provided by the Royal Society. Work at NRL is supported by NASA.

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. This work performed in part under DOE Contract DE-AC02-76SF00515.

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.

Part of this work is based on data obtained with the international LOFAR Telescope (ILT) under project codes LC7_018, DDT7_002, LT5_003, LC9_041, and LT10_004. LOFAR (van Haarlem et al., 2013) is the Low Frequency Array designed and constructed by ASTRON. It has observing, data processing, and data storage facilities in several countries, that are owned by various parties (each with their own funding sources) and that are collectively operated by the ILT foundation under a joint scientific policy. The ILT resources have benefited from the following recent major funding sources: CNRS-INSU, Observatoire de Paris and Université d’Orléans, France; BMBF, MIWF-NRW, MPG, Germany; Science Foundation Ireland (SFI), Department of Business, Enterprise and Innovation (DBEI), Ireland;

NWO, The Netherlands; The Science and Technology Facilities Council, UK.

HiPERCAM and V.S.D. are funded by the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) under ERC-2013-ADG grant agreement number 340040 (HiPERCAM). ULTRACAM and V.S.D. are funded by the UK Science and Technology Facilities Council. This work is based on observations made with the Gran Telescopio Canarias (GTC), installed in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias, on the island of La Palma, and on observations made with ESO Telescopes at the La Silla Paranal Observatory.

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

3.7 Appendix to Chapter 3

3.7.1 Estimating the false-alarm probability for a multi-dimensional H statistic search

It is important to estimate the false-alarm probability P_{FA} to know if the gamma-ray detection is real. As described in Section 3.2.3, there is no known analytical expression for the false-alarm probability of the maximum value from an H statistic search over a dense, multi-dimensional parameter grid. Deriving the probability distribution for the maximum value of a multi-dimensional “random field” is difficult and approximate solutions are only known for simple cases such as Gaussian or chi-squared random fields [231]. While the power in a single harmonic does follow a chi-squared random field in the presence of random noise, the known solutions cannot be applied in this case due to the maximization over summed harmonics and penalty factors defining the H statistic, and the fact that the metric density varies between different summed harmonics. Even for chi-squared random fields, there is no simple “trials factor” that can be applied to the single-trial false-alarm probability (which for the H statistic was derived by [114]): the false-alarm probability depends on the volume, shape, and dimensionality of the search space [231]. A full discussion of this is beyond the scope of this work. Below, we show empirically that a simple trials factor approach over-estimates the detection significance, and describe the “bootstrapping” method that we used to overcome this.

The false-alarm probability for a single H statistic trial is

$$P_{\text{FA}}(H_m | a) = e^{-aH_m}, \quad (3.9)$$

with scaling factor $a = 0.3984$ [114, 232]. This formula can be used to estimate the significance of the maximum H statistic value after n independent trials

$$P_{\text{FA}}(H_m | a, n) = 1 - [1 - e^{-aH_m}]^n. \quad (3.10)$$

We assume at first that our search contained a number of “effective” independent trials (N_{eff}) that is some unknown fraction of the number of actual trials (i.e. the

number of grid points at which we evaluated the H statistic). We then estimated N_{eff} from the results of our search as follows. We divided our parameter space into $n_{\text{seg}} = 2 \times 17 \times 13 = 442$ segments in f , \dot{f} , and P_{orb} respectively. The number of segments in f and \dot{f} is determined by the parameter space volumes, which were searched in parallel, as only the highest H statistic values from each were stored. To ensure that all segments were independent from the pulsar signal, we removed all grid points within those segments which were close (according to the parameter space metric; see Section 3.2.2) to the pulsar parameters.

The highest H statistic of each of the segments is plotted in the normalized histogram in Figure 3.8. We fit for the effective number of trials [as done by, e.g., 233] by maximizing the likelihood,

$$L(n, a | H_{m,i}) = \prod_i p(H_{m,i} | a, n) \quad (3.11)$$

for our set of H statistic values, according to the probability density function for H_m after n trials (the derivative of Equation (3.10)),

$$p(H_m | a, n) = a n [1 - e^{-a H_m}]^{n-1} \exp(-a H_m). \quad (3.12)$$

However, as shown in Figure 3.8, the tail of the best fitting distribution is significantly under-estimated, leading to *over-estimated* significances for large H statistic values. This demonstrates that there is no simple effective trials factor that can be applied to estimate the overall significance.

To overcome this, we performed a second fit, maximizing over the likelihood for both n and a . The resulting best-fitting distribution is also shown in Figure 3.8. We found the best-fitting scaling factor to be $\hat{a} \approx 0.284$, meaning the probability density function is flatter and gives a more conservative estimate for the significance. We note that this should not apply in general, and will depend, amongst other factors, on the dimensionality of the search space and the number of harmonics summed.

Finally, we use \hat{a} and multiply the best-fitting n (the best-fitting per-segment trials factor) by n_{seg} , and apply Equation (3.10) to obtain an approximation to the false-alarm probability for the maximum H statistic value. For the candidate pulsar signal, this was $P_{\text{FA}} = 0.33\%$. For comparison the candidate with the largest H statistic from a segment of the search not affected by the pulsar signal had $P_{\text{FA}} = 56\%$.




Figure 3.8: Normalized histogram showing the highest H statistics for 442 subsets of our search space after excluding results affected by the pulsar signal. The dotted green and dashed blue curves show normalized probability density functions for the maximum H statistic obtained after n effective trials. The curves gave maximum likelihood after varying over n with fixed single-trial scaling factor $a = 0.3984$ (dotted green) and after varying a and n jointly (dashed blue). The maximum H statistic for the pulsar $H_m = 86.7$ is marked by the vertical orange line. The red line (dashed-dotted) shows the false-alarm probability depending on H_m computed with Equation (3.10) with a and n from the joint variation.

Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated Einstein@Home

*Accepted as Nieder, L., Clark, C. J., Kandel, D., et al. 2020, ApJL.
arXiv: 2009.01513*

Abstract

We report the discovery of 1.97 ms-period gamma-ray pulsations from the 75-minute orbital-period binary pulsar now named PSR J1653–0158. The associated *Fermi*-Large Area Telescope gamma-ray source 4FGL J1653.6–0158 has long been expected to harbor a binary millisecond pulsar. Despite the pulsar-like gamma-ray spectrum and candidate optical/X-ray associations – whose periodic brightness modulations suggested an orbit – no radio pulsations had been found in many searches. The pulsar was discovered by directly searching the gamma-ray data using the GPU-accelerated *Einstein@Home* distributed volunteer computing system. The multi-dimensional parameter space was bounded by positional and orbital constraints obtained from the optical counterpart. More sensitive analyses of archival and new radio data using knowledge of the pulsar timing solution yield very stringent upper limits on radio emission. Any radio emission is thus either exceptionally weak, or eclipsed for a large fraction of the time. The pulsar has one of the three lowest inferred surface magnetic-field strengths of any known pulsar with $B_{\text{surf}} \approx 4 \times 10^7$ G. The resulting mass function, combined with models of the companion star’s optical light curve and spectra, suggest a pulsar mass $\gtrsim 2 M_{\odot}$. The companion is light-weight with mass $\sim 0.01 M_{\odot}$, and the orbital period is the shortest known for any rotation-powered binary pulsar. This discovery demonstrates the *Fermi*-Large Area Telescope’s potential to discover extreme pulsars that would otherwise remain undetected.

4.1 Introduction

The *Fermi*LAT source 4FGL J1653.6–0158 is a bright gamma-ray source, and the brightest remaining unassociated source [67]. It was first seen by the Energetic Gamma Ray Experiment Telescope [EGRET, 234], and was also listed in the LAT Bright Gamma-ray source list [34] more than a decade ago. While pulsars were discovered in several other sources from this list [see, e.g., 61], the origin of 4FGL J1653.6–0158 remained unidentified. The detection of a variable X-ray and optical candidate counterpart with 75 min-period consistent with the gamma-ray position of 4FGL J1653.6–0158 provided strong evidence of it being a binary gamma-ray pulsar [111, 112].

To identify the neutron star in 4FGL J1653.6–0158, we carried out a binary-pulsar search of the gamma rays, using the powerful GPU-accelerated distributed volunteer computing system *Einstein@Home*. Such searches are very computationally demanding, and would take decades to centuries on a single computer while still taking weeks or months on *Einstein@Home*. Thus, the search methods are specifically designed to ensure efficiency (Chapter 2). One key element is the use of constraints derived from optical observations. The companion’s pulsar-facing side is heated by the pulsar wind, leading to a periodically varying optical light curve. This permits the orbital period P_{orb} and other orbital parameters to be tightly constrained (for a feasible search the uncertainty ΔP_{orb} needs to be less than a few milliseconds). In addition, because the sky position of the optical source is typically known to high precision (sub-milliarcsecond level), a search over position parameters is not needed.

Here we present the discovery and analysis of gamma-ray pulsations from PSR J1653–0158 in 4FGL J1653.6–0158. The pulsar is spinning very rapidly, at a rotational frequency of 508 Hz. The inferred surface magnetic-field strength is one of the lowest of all known pulsars. The discovery also confirms the 75 min orbital period. This very short orbital period raises interesting questions about the evolutionary path which created the system.

The paper is organized as follows. In Section 7.3, we describe the gamma-ray search, detection and analysis within LAT data. The optical analysis of the pulsar’s companion, radio pulsation searches, and a continuous gravitational-wave follow-up search are presented in Section 4.3. We discuss the results and conclude in Section 4.4.

4.2 Gamma-ray pulsations

4.2.1 Data preparation

We searched for gamma-ray pulsations in the arrival times of photons observed by the *Fermi* LAT [32] between August 3, 2008 and April 16, 2018 (MJDs 54,681 and 58,224). We included SOURCE-class photons according to the P8R2_SOURCE_V6 [173] instrument response functions (IRFs)¹, with reconstructed incidence angles within a 5° region-of-interest (RoI) around the putative pulsar position, energies above 100 MeV and zenith angles below 90°. Here, we used the presumptive companion’s

¹See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_essentials.html

position as reported in the *Gaia* DR2 Catalog [122]; hereafter Gaia catalog. The celestial parameters (J2000.0) are $\alpha = 16^{\text{h}}53^{\text{m}}38^{\text{s}}.05381(5)$ and $\delta = -01^{\circ}58'36''.8930(5)$, with 1σ uncertainties on the last digits reported in parentheses.

Using the photon incidence angles and energies, we constructed a probability or weight for each photon, $w_j \in [0, 1]$, where j labels the photon: w_j is the probability that the j th photon originated from the posited source, as opposed to a fore- or background source. These weights were computed by `gtsrcprob`, using the preliminary *Fermi*-LAT 8-year source catalog² as a model for the flux within the RoI without performing a full spectral fit. Weighting the contribution of each photon to a detection statistic in this way greatly increases the search sensitivity [114], and the distribution of weights can be used to predict expected signal-to-noise ratios (Chapter 2).

The data set used here consisted of $N = 354,009$ photons, collected over a period of 3,542 days. The properties of the detection statistics (semicoherent power S_1 , coherent power P_1 , and H statistic) depend upon the lowest moments of the weights, which are

$$\sum_{j=1}^N w_j \approx 10266, \quad \sum_{j=1}^N w_j^2 \approx 2464, \quad \text{and} \quad \sum_{j=1}^N w_j^4 \approx 931.$$

These moments determine the ultimate sensitivity to a particular pulse profile and pulsed fraction, as given in Eq. (11) in Chapter 2.

Following the pulsar discovery, we extended this dataset to February 23, 2020 (MJD 58,902), using the latest `P8R3_SOURCE_V2` IRFs [235], a larger maximum zenith angle of 105° , and using the *Fermi*-LAT Fourth Source Catalog [hereafter 4FGL, 38] as the RoI model for the photon probability weight computations.

4.2.2 Search

The binary-pulsar search methods are described in Chapter 2, which are a generalization and extension of the isolated-pulsar search methods from Pletsch and Clark [80].

The searched ranges are guided by the known MSP population in the ATNF Pulsar Catalogue³ [110]. For the spin frequency, we searched $f \in [0, 1500]$ Hz⁴. The spin-frequency derivative was expected to be in the range $\dot{f} \in [-10^{-13}, 0]$ Hz s⁻¹.

The sky position of the candidate optical counterpart is constrained to high precision in the Gaia catalog, so no astrometric search is required. The proper motion measured by Gaia for the optical counterpart was ignored for the search.

4.2.2.1 Orbital Constraints from Optical Observations

The orbital period estimate of [111] was derived from Southern Astrophysical Research (SOAR), WIYN and Catalina Sky Survey (CSS) observations. These were augmented by new 350s SOAR Goodman High Throughput Spectrograph (GHTS)

²<https://fermi.gsfc.nasa.gov/ssc/data/access/lat/f18y/>

³<http://www.atnf.csiro.au/research/pulsar/psrcat>

⁴The upper limit has been chosen to be sensitive to pulsars spinning at up to 750 Hz, which have two-peaked pulse profiles where the peaks are half a rotation apart [see also, 80]. Note that the current record spin frequency is 716 Hz [13].

g' , r' , i' exposures (63 g' , 75 r' , 42 i') from MJD 56,514.074 – 56,516.184, and with the 300 s g' , r' and i' exposures obtained by Kong et al. [112] using the Wide Field camera (WFC) on the 2.5m Isaac Newton Telescope (INT) on La Palma. For these two data sets, the scatter about the light curve trends was appreciably larger than the very small statistical errors; we thus add 0.03 mag in quadrature to account for unmodeled fast variability and/or photometry systematics. To further refine the orbital period uncertainty, we obtained additional observations in u' , g' and i' using the high-speed multi-band imager ULTRACAM [195] on the 4.2m William Herschel Telescope (WHT) on two nights (MJDs 57,170 and 57,195), covering six and three orbits of the binary system, respectively, with a series of 20 s exposures. Conditions were very poor on the first night with seeing > 5 arcsec, particularly at the beginning of the observation. We therefore only used the second night’s data for the optical light curve modeling in Section 4.3.1, adding the latter half of the first night’s observations for orbital period estimation. Finally, we obtained further INT+WFC exposures (23 g' , 151 r' , 45 i') on MJD 57,988 – 57,991. The g' , r' , i' filter fluxes were referenced to in-field PanSTARRS catalog sources, and then converted to the SDSS scale. The u' photometry was calibrated against an SDSS standard star observed on MJD 57,170. We estimate ~ 0.05 mag systematic uncertainties in g' , r' and i' , with uncertainties as large as ~ 0.1 mag in u' .

We constrained the orbital period using the multi-band Lomb Scargle periodogram method [236, excluding the u' ULTRACAM data, as the modulation has very low signal-to-noise ratio in this band]. To infer reasonable statistical uncertainties, we fit for and removed constant magnitude offsets, consistent with our estimated calibration uncertainties, between each night’s observations in each band, and additionally rescaled the magnitude uncertainties to obtain a reduced chi-square of unity. This constrained the orbital period to $P_{\text{orb}} = 0.0519447518 \pm 6.0 \times 10^{-9}$ days, where the quoted uncertainty is the 1σ statistical uncertainty. For the pulsation search, we chose to search the 3σ range around this value.

In Romani, Filippenko, and Cenko [111], the time of the pulsar’s ascending node, T_{asc} , was estimated from the photometric light curve. However, the optical maximum is distinctly asymmetric (see Section 4.3.1), which can bias orbital phase estimates. We therefore used the spectroscopic radial velocity measurements from Romani, Filippenko, and Cenko [111], folded at the orbital period obtained above, and fit the phase of a sinusoidal radial velocity curve, finding $t_{\text{asc}} = \text{MJD } 56513.47981 \pm 2.1 \times 10^{-4}$. However, as radial velocities may still be slightly biased by asymmetric heating, we elected to search a wide range around this value, corresponding to $\pm 8\sigma$.

For the projected semimajor-axis parameter $x = a_1 \sin i/c$, we decided to start searching $x \in [0, 0.1]$ s, with the intention to go to larger values in the case of no detection. For a pulsar mass of $1.6 M_{\odot}$, this would cover the companion mass range up to $0.2 M_{\odot}$ and would include companion masses of all known “black-widow” systems as well as some of the lower mass “redback” systems [53, 54]. Here, a_1 is the pulsar’s semimajor axis, i denotes the inclination angle, and c the speed of light. As described in Chapter 2, we expected $x \in [0, 0.2]$ s based on the companion’s velocity amplitude reported by Romani, Filippenko, and Cenko [111] and the masses expected for “spider” companions, i.e. black-widow or redback companions.

4.2.2.2 Search grids

To cover the relevant orbital-parameter space in $\{x, P_{\text{orb}}, t_{\text{asc}}\}$, we use *optimized grids* [127]. These grids use as few points as possible still ensuring that a signal within the relevant space should be detected. Furthermore, they are able to cover the orbital-parameter space efficiently even though the required density depends on one of the orbital parameters, x .

Key to building an optimized grid is to know how the signal-to-noise ratio drops due to offsets from the true pulsar parameters. This is estimated using a *distance metric* on the orbital-parameter space (Chapter 2). In our case, the three-dimensional grid was designed to have a worst-case mismatch $\bar{m} = 0.2$, i.e. not more than 20% of the (semicoherent or coherent) signal power should be lost due to orbital-parameter offsets. Of most relevance is that 99% of randomly injected orbital-parameter points have a mismatch below $\bar{m} = 0.04$ to the closest grid point.

Due to the f -dependency of the required grid-point density, we search f in steps, and build the corresponding orbital grids prior to the start of the search on the computing cluster ATLAS in Hannover [84].

4.2.2.3 Einstein@Home

Searching the 5-dimensional parameter space $\{f, \dot{f}, x, P_{\text{orb}}, t_{\text{asc}}\}$ is a huge computational task with over 10^{17} trials. Thus, the first (computing-intensive) search stages were performed on *Einstein@Home*, a distributed volunteer computing system [85]. As done for radio pulsar searches previously, the search code utilizes the approximately 10,000 GPUs active on *Einstein@Home* for a computing speed-up of ~ 10 , comparing the runtimes on CPUs and GPUs.

The parameter space is divided into more than one million regions. Searching one of these is called a “work unit”. These work units are sent to computers participating in *Einstein@Home*, and are searched when the computer is otherwise idle. Depending on the system, searching a work unit takes between half an hour and up to a few hours of computational time. In total, the search would have taken more than 50 years on a single computer, but using *Einstein@Home* it took less than two weeks.

4.2.2.4 Gamma-ray detection

The search process involves multiple stages in which semicoherent statistics are constructed, and the most significant candidates are passed on to fully coherent follow-up stages (for full details of the search pipeline and signal-to-noise ratio definitions, see Chapter 2). In the last semicoherent stage, a candidate found at a frequency of 1016 Hz had signal-to-noise ratio $S_1 = 8.6$, which we now associate with PSR J1653–0158. This was not the strongest candidate or far above the background of noise, but was among the ten most significant candidates in its work unit, and therefore passed on to the coherent stage. In the coherent stage, it was very significant, with signal-to-noise ratio $P_1/2 = 94$.

The search follow-ups confirmed significant pulsations with period $P \approx 1.97$ ms (or $f \approx 508$ Hz), while the actual search revealed an alias at twice the pulsar frequency. This may be because the signal has significant power in the second harmonic.

Note that the signal was found outside the 3σ range in t_{asc} from the constraints reported in this work, and outside the 3σ range given by [111]. This can be caused by asymmetric heating (see Section 4.2.2.1).

4.2.3 Timing

The parameters used in the phase model to describe the pulsar’s rotation are measured in a timing analysis. We use the timing methods as explained in [79], which are an extension of the methods by [97]. The basic principle is that the parameter space around the discovery parameters is explored using a Monte Carlo sampling algorithm with a template pulse profile.

To marginalize over the pulse profile template, we vary the template parameters as described in Chapter 3. In the case of PSR J1653–0158, we used a template consisting of two symmetrical, wrapped Gaussian peaks. We used constraints on the peaks’ full-width at half maximum, such that the peaks must be broader than 5% of a rotation, and narrower than half a rotation.

Our timing solution over 11 years of LAT data is shown in Table 4.1. The folded gamma-ray data and the pulse profile are portrayed in Fig. 4.1.

The observed spin-down \dot{P} is one of the lowest of all known pulsars. To estimate the intrinsic \dot{P} we account for the Shklovskii effect [213], and the Galactic acceleration [see, e.g., 211]. The results are summarized in Table 4.1. The observed contribution due to the difference in Galactic acceleration of the Sun and the pulsar is computed with $R_{\text{Sun}} = 8.21$ kpc, $z_{\text{Sun}} = 14$ pc, and the Galactic potential model `PJM17_best.Tpot` [237], as implemented in their code⁵. For PSR J1653–0158, we used $R_{\text{J1653}} = 7.48$ kpc, and $z_{\text{J1653}} = 367$ pc, assuming $d = 840$ pc (see Table 4.2). The contributions parallel and perpendicular to the Galactic disk nearly cancel each other, so that the choice of the potential and its relevant parameters have a seemingly large effect on the actual small value of \dot{P}_{Gal} , and can even change the sign. However, the overall kinematic contribution to the observed \dot{P} is dominated by the Shklovskii term, and its uncertainty by the uncertainty in the distance estimate. The estimated intrinsic spin-down is $\dot{P}_{\text{int}} = 8.5 \times 10^{-22} \text{ s s}^{-1}$ for distance $d = 840$ pc.

4.3 Multiwavelength & Multimessenger

4.3.1 Optical Light Curve Modeling and System Masses

By modeling the optical light curves and radial velocities we can constrain the binary mass and distance and the system viewing angle. Comparing the individual filters between nights suggest small $\delta m \approx 0.05$ shifts in zero points, consistent with the systematic estimates above. Correcting to match the individual filters, we then re-binned the light curve, placing the photometry on a regular grid with points spaced by $\delta\phi = 0.004$, using the Python package `Lightkurve`; after excision of a few obviously discrepant points, we retain 248 u' , 239 g' , 220 r' and 245 i' points for light curve fitting (Fig. 4.2). This fitting is done with a version of the `Icarus` code of Breton et al. [238] modified to include the effect of hot spots on the companion surface, likely generated by precipitation of particles from the intrabinary

⁵<https://github.com/PaulMcMillan-Astro/GalPot>

Table 4.1: Timing solution for PSR J1653–0158.

Parameter	Value
Range of observational data (MJD)	54682 – 58902
Reference epoch (MJD)	56100.0
Celestial parameters from Gaia catalog	
R.A., α (J2000.0)	16 ^h 53 ^m 38 ^s .05381(5)
Decl., δ (J2000.0)	–01°58′36″.8930(5)
Positional epoch (MJD)	57205.875
Proper motion in R.A., $\mu_\alpha \cos \delta$ (mas yr ^{–1})	–19.62 ± 1.86
Proper motion in Dec., μ_δ (mas yr ^{–1})	–3.74 ± 1.12
Parallax ^a , ϖ (mas)	1.88 ± 1.01
Timing parameters	
Spin frequency, f (Hz)	508.21219457426(6)
Spin-frequency derivative, \dot{f} (Hz s ^{–1})	–6.204(8) × 10 ^{–16}
Spin period, P (ms)	1.9676820247057(2)
Spin-period derivative, \dot{P} (s s ^{–1})	2.402(3) × 10 ^{–21}
Proj. semimajor axis, x (s)	0.01071(1)
Orbital period, P_{orb} (days)	0.0519447575(4)
Epoch of ascending node, t_{asc} (MJD)	56513.479171(8)
Derived parameters for distance $d = 840$ pc	
Shklovskii spin down, \dot{P}_{Shk} (s s ^{–1})	1.6 × 10 ^{–21}
Galactic acceleration spin down, \dot{P}_{Gal} (s s ^{–1})	–4.8 × 10 ^{–23}
Spin-down power, \dot{E} (erg s ^{–1})	4.4 × 10 ³³
Surface B -field, B_{surf} (G)	4.1 × 10 ⁷
Light-cylinder B -field, B_{LC} (G)	5.0 × 10 ⁴
Characteristic age, τ_c (Gyr)	37
Gamma-ray luminosity ^b , L_γ (erg s ^{–1})	2.9 × 10 ³³
Gamma-ray efficiency, $n_\gamma = L_\gamma / \dot{E}$	0.66

Notes. — The JPL DE405 solar system ephemeris has been used, and times refer to TDB.

^a Corresponds to a model independent distance $d = 533_{-187}^{+625}$ pc, but for the derived parameters the consistent distance $d = 840_{-40}^{+40}$ pc derived from optical modeling is used (see Table 4.2).

^b Taken from 4FGL Source Catalog [38].

shock (IBS) to companion magnetic poles [205]. All parameter values and errors are determined by MCMC modeling.

The very shallow modulation of these light curves might normally be interpreted as indicating a small inclination i . However given the large companion radial velocity amplitude $K = 666.9 \pm 7.5$ km s^{–1}, implying a mass function $f(M) = 1.60 \pm 0.05 M_\odot$, measured by Romani, Filippenko, and Cenke [111], a small inclination would give an unphysical, large neutron star mass. As noted in that paper, the light curves and spectra show that a strong blue non-thermal veiling flux dominates at orbital minimum. With increasingly shallow modulation for the bluer colors, this is also evident in the present photometry. Thus the minimal model for this pulsar must include a non-thermal veiling flux. Although this is likely associated with the IBS,




Figure 4.1: Integrated pulse profile and phase-time diagram of PSR J1653–0158, showing two identical rotations. Top: The histogram shows the weighted counts for 50 bins. The orange curve indicates the pulse-profile template with the highest signal power, and the transparent black curves represent 100 templates randomly selected from the Monte Carlo samples after the chain stabilized, to indicate the uncertainty on the profile. The dashed blue line denotes the source background. Bottom: Each point represents the pulsar’s rotational phase at emission of a photon, with the intensity indicating the photon’s probability weight. Note that PSR J1653–0158 received more exposure between MJDs 56,600 and 57,000 when the LAT pointed more often towards the Galactic center.

we model it here as a simple power law with form $f_\nu = f_A(\nu/10^{14} \text{ Hz})^{-p}$. This flux is nearly constant through the orbit, although there are hints of phase structure, e.g. in r' and i' at $\phi_B = 0.72$ (see Fig. 4.2). Any model without such a power law component is completely unacceptable. These fits prefer an A_V slightly higher than, but consistent with the maximum in this direction [obtained by $\sim 300 \text{ pc}$; 239]⁶.

In Fig. 4.2, one notices that the orbital maximum is slightly delayed from $\phi_B = 0.75$, especially in the bluer colors. Such asymmetric heating is most easily modeled adding a polar hot spot with location (θ_c, ϕ_c) and local temperature increase A_c in a Gaussian pattern of width σ_c ; when we include such a component, the fit improves

⁶<https://doi.org/10.7910/DVN/2EJ9TX>




Figure 4.2: u' , g' , r' , and i' light curves for PSR J1653–0158, with the best-fit model curves. Note the flat minima and decreasing modulation for bluer colors, a consequence of the hard spectrum veiling flux. Two identical cycles are shown for clarity.

greatly, with $\Delta\chi^2/\text{DoF} = -0.34$. The Akaike Information Criterion comparison of the two models indicates that the model with a hot spot is preferred at the 10^{-18} level, despite the extra degrees of freedom. We give the fit parameters for both models in Table 4.2. Note that with the fine structure near maximum, the model is not yet fully acceptable ($\chi^2/\text{DoF} \sim 1.4$). More detailed models, including direct emission from the IBS or possibly the effects of companion global winds [240], may be needed to fully model the light curves. Such modeling would be greatly helped by light curves over an even broader spectral range, with IBS effects increasingly dominant in the UV, and low temperature companion emission better constrained in the IR. With many cycles we could also assess the reality (and stability) of the apparent fine structure and test for hot spot motion.

Our fit distance may be cross-checked with two other quantities. (1) With the 4FGL energy flux $f_\gamma = 3.5 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ between 100 MeV and 100 GeV, our fit distance gives an isotropic gamma-ray luminosity $L_\gamma = 3 \times 10^{33} \text{ erg s}^{-1}$, in good agreement with the $L_\gamma \approx (10^{33} \text{ erg s}^{-1} \dot{E})^{1/2}$ heuristic luminosity law [28], as a function of the spin-down power \dot{E} . This luminosity is consistent with the model for direct radiative heating of the companion. (2) Our fit distance is also consistent with the model-independent, but lower accuracy, distance from the Gaia parallax. Thus, the 840 pc distance seems reliable, although systematic effects probably dominate over the rather small ~ 50 pc statistical errors.

Armed with the fits, we can estimate the companion masses, correcting the observed radial velocity amplitude (fit with a K-star template) for the temperature-dependent weighting of the absorption lines across the companion face as in Kandel and Romani [240]. The results indicate substantial mass accretion, as expected for these ultra-short period systems. With the preferred Veiled+HS model the mass significantly exceeds $2.0 M_\odot$, adding to the growing list of spider binaries in this mass range. Note that the inclination i uncertainty dominates the error in this mass determination. Broader range photometric studies, with better constraint on the heating pattern, can reduce the i uncertainty.

Table 4.2: Light curve fit results for PSR J1653–0158.

Parameters	Veiled	Veiled+HS
Inclination, i (deg)	$79.4^{+5.7}_{-6.8}$	$72.3^{+5.0}_{-4.9}$
Filling factor, f_c	$0.97^{+0.02}_{-0.02}$	$0.88^{+0.03}_{-0.03}$
Heating luminosity, L_P (10^{33} erg s $^{-1}$)	$3.33^{+0.39}_{-0.34}$	$3.15^{+0.26}_{-0.27}$
Night-side temperature, T_N (K)	3250^{+243}_{-331}	3295^{+227}_{-300}
V-band extinction, A_V	$1.06^{+0.08}_{-0.10}$	$1.06^{+0.07}_{-0.09}$
Distance, d (pc)	830^{+50}_{-50}	840^{+40}_{-40}
Veiling flux norm, f_A (μ Jy)	$101.7^{+11.4}_{-11.1}$	$99.9^{+11.7}_{-11.4}$
Veiling flux index, p	$0.50^{+0.05}_{-0.03}$	$0.49^{+0.03}_{-0.03}$
Spot azimuth, θ_c (deg)	...	$286.8^{+5.8}_{-6.9}$
Spot co-latitude, ϕ_c (deg)	...	$-50.5^{+9.2}_{-8.4}$
Gaussian spot width, σ_c (deg)	...	$25.2^{+5.0}_{-4.9}$
Spot temperature increase, A_c	...	$0.66^{+0.21}_{-0.21}$
Neutron-star mass, M_{NS} (M_\odot)	$1.99^{+0.18}_{-0.08}$	$2.17^{+0.21}_{-0.15}$
Companion mass, M_c (M_\odot)	$0.013^{+0.001}_{-0.001}$	$0.014^{+0.001}_{-0.001}$
χ^2/DoF	1.72	1.38

Notes. — Parameters from the best-fit light curve/radial velocity models, with and without a surface hot spot, including MCMC errors.

4.3.2 Radio pulsation searches

The pulsar position has been observed in radio multiple times. Several searches were performed before the gamma-ray pulsation discovery, and a few very sensitive follow-up searches afterwards. Despite the more than 20 observations with eight of the most sensitive radio telescopes, no radio pulsations have been found.

The results of the radio searches are given in Table 4.3. Observations are spread over 11 years, with observing frequencies ranging from 100 MHz up to 5 GHz. All orbital phases have been covered by most of the telescopes. Since there was no detection, the table also gives upper limits derived from the observations. For all but LOFAR, the data (both archival and recent) were folded with the gamma-ray-derived ephemeris, and searched only over dispersion measure.

The strictest upper limits on pulsed radio emission are $8 \mu\text{Jy}$ at 1.4 GHz, and $20 \mu\text{Jy}$ at 4.9 GHz. This is fainter than the threshold of $30 \mu\text{Jy}$ that [28] use to define a pulsar to be “radio-quiet”. Note, that for the calculation of the limits we included the parts of the orbit where eclipses might be expected for spider pulsars. Thus, the limit constrains the maximum emission of the system, and not the maximum emission from the pulsar alone.

4.3.3 Continuous gravitational waves

We search for nearly monochromatic, continuous gravitational waves (GWs) from PSR J1653–0158, using data from the first⁷ and second⁸ observing runs of the Advanced LIGO detectors [245]. We assume that GWs are emitted at the first and

⁷<https://doi.org/10.7935/K57P8W9D>

⁸<https://doi.org/10.7935/CA75-FM95>

Table 4.3: Summary of radio searches for PSR J1653–0158. The columns show the telescope used, the observed frequency range, the start time and data span, the range of orbital phases covered, the resulting limit on a pulsed component, and a reference with relevant details.

Telescope	Frequency (MHz)	Data start (UTC)	Data span (s)	Orbital phase	Limit (μJy)	Reference / Survey
Effelsberg	1210–1510	2010 May 26, 21:33	1920	0.88–1.31	63	Barr et al. [58]
Effelsberg	1210–1510	2014 Aug 26, 20:27	4600	0.15–1.17	41	
Effelsberg	4608–5108	2014 Aug 29, 18:52	4600	0.62–1.65	33	
Effelsberg	4608–5108	2020 Jun 18, 22:09	11820	0.85–3.48	20	
FAST	1050–1450	2020 Jun 04, 16:30	2036	0.80–1.25	8	Li et al. [241]
GBT	720–920	2009 Sep 20, 00:49	3200	0.93–1.65	51	
GBT	720–920	2010 Dec 13, 21:04	1300	0.91–1.20	80	
GBT	720–920	2011 Dec 22, 12:11	2400	0.74–1.27	59	Sanpa-arsa [242]
GBT	305–395	2012 Feb 22, 14:31	1700	0.27–0.65	301	
GBT	1700–2300	2014 Nov 18, 14:28	1200	0.36–0.63	43	
GBT	1700–2300	2014 Nov 20, 13:56	2400	0.44–0.98	30	
GBT	1700–2300	2014 Nov 21, 22:38	1800	0.66–1.07	35	
GBT	720–920	2017 Jan 28, 13:20	1200	0.97–1.24	83	
GMRT	591–623	2011 Feb 02, 02:32	1800	0.94–1.34	730	Bhattacharyya et al.
GMRT	306–338	2012 May 15, 22:31	1800	0.54–1.06	990	([60], 2020 in prep.)
GMRT	306–338	2012 Jun 11, 17:49	1800	0.55–0.95	990	"
GMRT	591–623	2014 Aug 19, 13:44	1800	0.00–0.54	270	"
GMRT	591–623	2014 Aug 30, 11:17	1800	0.80–1.38	270	"
GMRT	591–623	2015 Dec 28, 03:55	1800	0.73–1.13	270	"
LOFAR	110–180	2017 Mar 15, 04:18	15×320	Full orbit	6,200	Bassa et al. [14]
LOFAR	110–180	2017 Apr 15, 02:20	15×320	Full orbit	6,200	"
Lovell	1332–1732	2019 Mar 15, 01:34	5400	0.57–1.77	82	
Lovell	1332–1732	2019 Mar 16, 02:53	5400	0.87–2.08	82	
Lovell	1332–1732	2019 Mar 17, 01:47	5400	0.25–1.45	82	
Nançay	1230–1742	2014 Aug 20, 18:33	1850	0.12–0.53	77	Desvignes et al. [243]
Parkes	1241–1497	2016 Nov 05, 06:17	3586	0.26–1.06	178	Camilo et al. [244]

Notes. — The orbital phase is given in orbits, and ranges > 1 indicate that more than one orbit has been observed. The considered maximum dispersion measure varies with the observing frequency from $\text{DM} = 80 \text{ pc cm}^{-3}$ at the lowest frequencies to $\text{DM} = 350 \text{ pc cm}^{-3}$ at the highest frequencies. To estimate the limit on the pulsed component, we used Eq. (6) from [52] assuming a pulse width of $0.25 P$, and a threshold signal-to-noise ratio $S/N_{\text{min}} = 7$.

second harmonic of the neutron star’s rotational frequency, as would occur if the spin axis is misaligned with the principal axes of the moment of inertia tensor [246, 247].

We employ two different analysis procedures, which yield consistent results. The first is frequentist, based on the multi-detector maximum-likelihood \mathcal{F} -statistic introduced by Cutler and Schutz [208]. The second is the Bayesian time-domain method [248] as detailed by Pitkin et al. [249], with triaxial non-aligned priors [250]. Both methods coherently combine data from the two detectors, taking into account their antenna patterns and the GW polarization. The \mathcal{F} -statistic search excludes data taken during times when the relevant frequency bands are excessively noisy.

The results are consistent with no GW emission. At twice the rotation frequency, the \mathcal{F} -statistic 95% confidence upper limit on the intrinsic GW amplitude h_0 is 4.4×10^{-26} . The 95%-credible interval upper limit from the Bayesian analysis on $h_0 = 2C_{22}$ is 3.0×10^{-26} . At the rotation frequency (only checked with the Bayesian method) the 95% confidence upper limit on the amplitude C_{21} is 6.6×10^{-26} .

Since the dominant GW frequency might be mismatched from twice the rotation frequency [251], we performed an \mathcal{F} -statistic search in a $\pm 1 \text{ Hz}$ band around this, with an extended \dot{f} -range. This yields larger upper limits on h_0 , with mean value of 1.3×10^{-25} in 10 mHz-wide bands. Full details are given in the supplementary materials.

Our upper limits on h_0 at twice the rotation frequency may also be expressed as upper limits on the ellipticity ϵ of the pulsar [252]. This is $\epsilon = 3.9 \times 10^{-8} \times$

$(h_0/5 \times 10^{-26}) \times (10^{45} \text{g cm}^3/I_{zz}) \times (840 \text{ pc}/d)$, where I_{zz} is the moment of inertia about the spin axis, and d is the distance.

As is the case for most known pulsars, it is unlikely that our searches would have detected a GW signal. In fact, suppose that all of the rotational kinetic-energy losses associated with the intrinsic spin-down is via GW emission. Then assuming the canonical $I_{zz} = 10^{45} \text{g cm}^3$, this would imply a “spin-down” ellipticity $\epsilon^{\text{sd}} = 4.7 \times 10^{-10}$, which is a factor ~ 80 below our upper limit.

4.4 Discussion & Conclusions

PSR J1653–0158 is the second binary pulsar [44], and the fourth MSP [43] to be discovered through periodicity searches of gamma rays. This pulsar is remarkable in many ways. It is only the second rotationally-powered MSP from which no radio pulsations have been detected. It is among the fastest-rotating known pulsars with spin frequency $f = 508 \text{ Hz}$. The 75 min orbital period is shorter than for any other known rotation-powered pulsar, with the previous record being PSR J1311–3430 with a 93 min orbit [44]. The inferred surface magnetic field is possibly the weakest, depending on the Shklovskii correction.

The discovery was enabled by constraints on the sky-position and orbital parameters from optical observations, together with efficient search techniques and the large computing power of the distributed volunteer computing system *Einstein@Home*. The detection proves that the optically variable candidate counterpart [111, 112] is indeed the black-widow-type binary companion to PSR J1653–0158, and it conclusively resolves the nature of the brightest remaining unidentified gamma-ray source, first found more than two decades ago [234].

The distance to PSR J1653–0158, and its proper motion are well constrained. Gaia measurements of the parallax, $\varpi = 1.88 \pm 1.01 \text{ mas}$, imply a distance $d = 530_{-200}^{+470} \text{ pc}$. A consistent, but tighter constraint is given by our optical modeling with $d = 840_{-40}^{+40} \text{ pc}$. The proper motion (see Table 4.1) is also measured with good precision (Gaia and our timing are in agreement).

PSR J1653–0158 has one of the lowest observed spin-period derivatives of all known pulsars ($\dot{P} = 2.4 \times 10^{-21} \text{ s s}^{-1}$). The intrinsic $\dot{P} = 8.5 \times 10^{-22} \text{ s s}^{-1}$ (accounting for Galactic acceleration and Shklovskii effects) is even smaller. In Fig. 4.3, PSR J1653–0158 is shown in a P - \dot{P} diagram, alongside the known radio and gamma-ray pulsar population outside of globular clusters.

The intrinsic \dot{P} can be used to estimate the pulsar’s spin-down power \dot{E} , surface magnetic-field strength B_{surf} , magnetic-field strength at the light cylinder B_{LC} , and characteristic age τ_c . These are given in Table 4.1 for $d = 840 \text{ pc}$. Constant lines of \dot{E} , B_{surf} , and τ_c are displayed in Fig. 4.3 to show the distance-dependent ranges.

Spider pulsars in very-short-period orbits are difficult to discover with traditional radio searches. Even though we can now fold the radio data with the exact parameters, PSR J1653–0158 is still not visible. There are two simple explanations for the non-detection of radio pulsations. (1) Radio emission is blocked by material produced by the pulsar evaporating its companion. Eclipses for large fractions of the orbit would be expected, since they have been seen for many spider pulsars [see, e.g., 19, 49, 253]. This is further supported by the observed extremely compact orbit and the strong IBS. Radio imaging observations could be used to check whether there is any continuum radio flux at the sky position of PSR J1653–0158, but pre-




Figure 4.3: Newly detected PSR J1653–0158 on a P - \dot{P} diagram of the known pulsar population outside of globular clusters. The MSP population is shown magnified in the inset. LAT pulsars are marked in green (isolated by a cross and binary by a circle). Non-LAT pulsars in the ATNF are marked in gray (isolated by a plus and binary by a square). The lines show constant surface magnetic-field strength (dashed-dotted), characteristic age (dotted), and spin-down power (dashed). The spin period and intrinsic spin-period derivative of PSR J1653–0158 are marked by the orange star. The transparent stars indicate the (distance-dependent) maximum and minimum intrinsic spin-period derivative according to the distance estimated from our optical models.

vious experience is not encouraging. The eclipses of a few other spider systems have been imaged at low frequencies, showing that, during the eclipse, the continuum flux from the pulsar disappears in tandem with the pulsed flux [179, 254]. (2) PSR J1653–0158 is intrinsically radio-quiet, in that its radio beam does not cross the line-of-sight, or it has a very low luminosity. There is one other radio-quiet MSP known [43].

The minimum average density of the companion 64 g cm^{-3} is very high, assuming a filled Roche lobe [255]. Using the filling factor from optical modeling, the average companion density 73 g cm^{-3} is even higher. The high density and the compact orbit suggest that the companion may be a helium white dwarf remnant, and that the system may have evolved from an ultracompact X-ray binary [216, 256]. In addition, simulations predict evolved ultracompact X-ray binaries to have orbital periods of around 70 – 80 min [257], consistent with the 75 min orbital period from PSR J1653–0158. Future analysis of optical spectroscopic data may give additional insight into the evolution and composition of the companion.

The discovery of PSR J1653–0158 is the result of a multiwavelength campaign. The pulsar-like gamma-ray spectrum, and the non-detection of radio pulsations, motivated the search for a visible companion. This was subsequently discovered in

optical and X-ray observations. Further optical observations provided constraints on the orbital parameters which were precise enough to enable a successful gamma-ray pulsation search.

Acknowledgments

We are deeply grateful to the thousands of volunteers who donated their computing time to *Einstein@Home*, and to those whose computers first detected PSR J1653–0158: Yi-Sheng Wu of Taoyuan, Taiwan; and Daniel Scott of Ankeny, Iowa, USA.

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: Holger J. Pletsch) and by National Science Foundation grants 1104902 and 1816904. L.N. was supported by an STSM Grant from COST Action CA16214. C.J.C. and R.P.B. acknowledge support from the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 715051; Spiders). V.S.D. and ULTRACAM are supported by the STFC. R.W.R. and D.K. were supported in part by NASA grant 80NSSC17K0024. S.M.R. is a CIFAR Fellow and is supported by the NSF Physics Frontiers Center award 1430284 and the NASA Fermi GO Award NNX16AR55G. Fermi research at NRL is funded by NASA. J.W.T.H. is an NWO Vici fellow.

The ULTRACAM photometry was obtained as part of program WHT/2015A/35. The William Herschel Telescope is operated on the island of La Palma by the Isaac Newton Group of Telescopes in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias. Based on observations made with the Isaac Newton Telescope (program I17BN005) operated on the island of La Palma by the Isaac Newton Group of Telescopes in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias. This paper makes use of data obtained from the Isaac Newton Group of Telescopes Archive which is maintained as part of the CASU Astronomical Data Centre at the Institute of Astronomy, Cambridge.

We acknowledge support of the Department of Atomic Energy, Government of India, under project no. 12-R&D-TFR-5.02-0700 for the GMRT observations. The GMRT is run by the National Centre for Radio Astrophysics of the Tata Institute of Fundamental Research, India. The Nançay Radio Observatory is operated by the Paris Observatory, associated with the French Centre National de la Recherche Scientifique (CNRS). We acknowledge financial support from the “Programme National Hautes Energies” (PNHE) of CNRS/INSU, France. This paper is based (in part) on data obtained with the International LOFAR Telescope (ILT) under project code LC7_018. LOFAR [171] is the Low Frequency Array designed and constructed by ASTRON. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. The Green Bank Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. FAST is a Chinese national mega-science facility, built and operated by NAOC. Partly based on observations with the 100-m telescope of the MPIfR (Max-Planck-Institut für Radioastronomie) at Effelsberg.

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. This work performed in part under DOE Contract DE-AC02-76SF00515.

The authors thank the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. This research has made use of data, software and/or web tools obtained from the LIGO Open Science Center (<https://losc.ligo.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration, to which the authors have also contributed. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

Einstein@Home Discovery of the Gamma-ray Millisecond Pulsar PSR J2039–5617 Confirms its Predicted Redback Nature

*Accepted as Clark, C. J., Nieder, L., Voisin, G., et al. 2020, MNRAS.
arXiv: 2007.14849*

Abstract

The *Fermi* Large Area Telescope gamma-ray source 3FGL J2039.6–5618 contains a periodic optical and X-ray source that was predicted to be a “redback” millisecond pulsar (MSP) binary system. However, the conclusive identification required the detection of pulsations from the putative MSP. To better constrain the orbital parameters for a directed search for gamma-ray pulsations, we obtained new optical light curves in 2017 and 2018, which revealed long-term variability from the companion star. The resulting orbital parameter constraints were used to perform a targeted gamma-ray pulsation search using the *Einstein@Home* distributed volunteer computing system. This search discovered pulsations with a period of 2.65 ms, confirming the source as a binary MSP now known as PSR J2039–5617. Optical light curve modelling is complicated and likely biased by asymmetric heating on the companion star and long-term variability, but we find an inclination $i \gtrsim 60^\circ$, for a low pulsar mass between $1.1 M_\odot < M_{\text{psr}} < 1.6 M_\odot$, and a companion mass of $0.15 - 0.22 M_\odot$, confirming the redback classification. Timing the gamma-ray pulsations also revealed significant variability in the orbital period, which we find to be consistent with quadrupole moment variations in the companion star, suggestive of convective activity. We also find that the pulsed flux is modulated at the orbital period, potentially due to inverse Compton scattering between high-energy leptons in the pulsar wind and the companion star’s optical photon field.

5.1 Introduction

Millisecond pulsars (MSPs) are old neutron stars that have been spun-up to millisecond rotation periods by the accretion of matter from an orbiting companion star [16]. The most compelling evidence for this “recycling” scenario comes from the discovery of three transitional MSPs, which have been seen to switch between rotationally powered MSP and accretion-powered low-mass X-ray binary (LMXB) states [19–22]. In their rotationally powered states, these transitional systems all belong to a class of interacting binary MSPs known as “redbacks”, which are systems containing an MSP in orbit with a low-mass ($0.1 M_{\odot} \lesssim M \lesssim 0.4 M_{\odot}$) non-degenerate companion star [53]. Redbacks, and the closely related “black widows” (which have partially degenerate companions with $M \lesssim 0.05 M_{\odot}$), are named after species of spiders in which the heavy females have been observed to consume the smaller males after mating, reflecting the fact that the lighter companion stars are being destroyed by the pulsar’s particle wind and/or intense high-energy radiation.

Until recently, only a handful of these “spider” systems had been found in radio pulsar surveys of the Galactic field. This is most likely due to the ablation phenomenon which gives redbacks and black widows their nicknames: plasma from the companion can eclipse, scatter and disperse the MSP’s radio pulsations for large fractions of an orbit [e.g., 100, 258], causing these pulsars to be easily missed in radio pulsar surveys. In addition, traditional “acceleration” search methods for binary pulsars [259] are only optimal when the integration time is $\lesssim 10\%$ of the orbital period, leading to an additional sensitivity loss to spiders, which often have orbital periods of just a few hours.

Fortunately, gamma ray emission from an MSP does not suffer from strong propagation effects from intrabinary plasma structures. A new route for binary MSP discoveries therefore appeared with the launch of the *Fermi Gamma-ray Space Telescope* in 2008. The on-board Large Area Telescope (LAT) discovered gamma-ray pulsations from a number of known MSPs shortly after launch [260]. Targeted radio observations of unidentified, but pulsar-like *Fermi*-LAT sources have since discovered more than 90 new MSPs, more than a quarter of all known MSPs in the Galactic field¹. A disproportionately large fraction of these are spiders that had been missed by previous radio surveys [56].

In addition to the large number of radio-detected spiders found in *Fermi*-LAT sources, a growing number of candidate spiders have been discovered through searches for optical and X-ray counterparts to gamma-ray sources [e.g., 55, 90–92, 111]. In a few cases, the MSP nature of these sources was confirmed by the detection of radio or gamma-ray pulsations [44, 68], however most of these candidates remain unconfirmed.

To overcome the difficulties in detecting spider MSPs in radio pulsation searches, it is possible to directly search for gamma-ray pulsations in the LAT data. In contrast to searches for isolated MSPs, which can be detected by truly “blind” gamma-ray pulsation searches [43], gamma-ray pulsation searches for binary MSPs require tight constraints on the orbital parameters of the candidate binary system to account for the orbital Doppler shift [44], which would smear out the pulsed signal if not corrected for. This in turn requires long-term monitoring of the companion star’s optical light curve to measure the orbital period with sufficient precision, and

¹<http://astro.phys.wvu.edu/GalacticMSPs/GalacticMSPs.txt>

spectroscopic radial velocity measurements and/or light curve modelling to tie the photometric light curve to the pulsar’s kinematic orbital phase. Prior to this work, such searches have been successful only twice [44, Chapter 4], with both MSPs being extremely compact black widow systems with small orbital Doppler modulations.

Salveti et al. [154] and Romani [153] discovered a high-confidence candidate redback system in the bright, pulsar-like gamma-ray source 3FGL J2039.6–5618 [37]. This source is now known as 4FGL J2039.5–5617 in the latest *Fermi*-LAT Fourth Source Catalog [38]; hereafter 4FGL. This system (which we refer to hereafter as J2039) contains a periodic X-ray and optical source with orbital period $P_{\text{orb}} \approx 5.5$ hr. The optical light curve exhibits two “ellipsoidal” peaks, interpreted as a tidally distorted companion star in an intense gravitational field being viewed from the side, where its projected surface area is highest. These peaks have unequal amplitudes, indicating a temperature difference between the leading and trailing sides of the star. Despite the high likelihood of this source being a redback system, the pulsar remained undetected in repeated observations attempting to detect its radio pulsations by Camilo et al. [65].

On 2017 June 18, we took new observations of J2039 with the ULTRACAM [195] high-speed multi-band imager on the 3.5m New Technology Telescope (NTT) at ESO La Silla. The goal of these observations was to refine the orbital period uncertainty by phase-aligning a new orbital light curve with the 2014 GROND observations from Salvetti et al. [154]. However, we found that the optical light curve had changed significantly. Further observations obtained on 2018 June 02 also found a light curve that differed from the first two. This variability, similar to that discovered recently in other redback pulsars [261, 262], poses challenges for obtaining reliable estimates of the physical properties such as the binary inclination angle and pulsar mass via optical light curve modelling [e.g., 194, 263].

Using constraints on the pulsar’s orbital period and epoch of ascending node from preliminary models fit to the optical data, we performed a gamma-ray pulsation search using the *Einstein@Home* distributed volunteer computing system [85, 264], which finally identified the millisecond pulsar, now named PSR J2039–5617, at the heart of the system.

In this paper, we present the detection and timing of gamma-ray pulsations from PSR J2039–5617, and our new optical observations of the system. The paper is organised as follows: in Section 5.2 we review the literature on recent observations of the system to update our knowledge of its properties; Section 5.3 presents updated analysis of *Fermi*-LAT gamma-ray observations of 4FGL J2039.5–5617, and describes the gamma-ray pulsation search, discovery and timing of PSR J2039–5617; in Section 5.4 we describe the newly obtained optical data, and model the optical light curves to estimate physical properties of the system and investigate the observed variability; in Section 7.4 we discuss the newly clarified picture of PSR J2039–5617 in the context of recent observations of redback systems; and finally a brief summary of our results is given in Section 3.6.

Shortly after the discovery of gamma-ray pulsations reported in this paper, the initial timing ephemeris was used to fold existing radio observations taken by the CSIRO Parkes radio telescope. The resulting detections of radio pulsations and orbital eclipses are presented in a companion paper [105], hereafter Paper II.

5.2 Summary of previous literature

The periodic optical counterpart to 4FGL J2039.5–5617 was discovered by Salvetti et al. [154] and Romani [153] in photometric observations of the gamma-ray source region taken over three nights on 2014 June 16–18 with GROND [265] on the ESO/MPG 2.2m telescope on La Silla. These observations covered SDSS g' , r' , i' , and z' optical filters in simultaneous 115 s exposures, and H , J , and K near infrared filters in simultaneous 10 s exposures. For consistency with the new optical light curves presented in this paper, we re-reduced the optical observations but chose not to include the infrared observations, which were not compatible with our reduction pipeline. These observations revealed a double-peaked light curve typical of red-back systems, but with the peak corresponding to the companion's ascending node brighter and bluer than that of the descending node. This requires the trailing side of the star to be hotter than the leading side, perhaps due to heating flux being redirected by an asymmetric intra-binary shock [e.g., 218], or due to the presence of cold spots on the leading edge [e.g., 261].

Salvetti et al. [154] and Romani [153] also analyzed X-ray observations of J2039 taken by *XMM-Newton*. These data had insufficient time resolution to test for millisecond X-ray pulsations, but did reveal a periodic (~ 5.5 hr) modulation in the X-ray flux, which the authors identified as likely being due to synchrotron emission from particles accelerated along an intra-binary shock, commonly seen in black widow and redback systems. However, without long-term timing to precisely measure the orbital period the authors were unable to unambiguously phase-align the optical and X-ray light curves. The Catalina Surveys Southern Periodic Variable Star Catalogue [266] includes 223 photometric observations of J2039 between 2005 and 2010. While the uncertainties on these unfiltered data are too large for a detailed study of the light curve over these 5 years, the underlying periodicity is clearly recovered by a 2-harmonic Lomb Scargle periodogram, which reveals a significant signal with an orbital period of $P_{\text{orb}} = 0.227980(1)$ d with no significant aliases. Folding at this period shows that the X-ray modulation peaks at the putative pulsar's inferior conjunction, indicating that the shock wraps around the pulsar. This scenario requires the companion's outflowing wind to overpower the pulsar wind [see e.g. 218, 220].

Using 9.5 yr of *Fermi*-LAT data, Ng et al. [267] discovered that the gamma-ray emission from J2039 contains a component below 3 GeV that is modulated at the orbital period, peaking around the companion star's inferior conjunction, i.e. half an orbit out of phase with the X-ray modulation. This phase offset rules out synchrotron emission from particles accelerated along the shock front as an origin for the gamma-ray flux, as such a component would occur at the same orbital phase as the X-ray modulation. Instead, Ng et al. [267] propose that this component is produced by inverse Compton scattering between the pulsar's high-energy particle wind and the companion star's optical photon flux. Such a component would be strongest if our line of sight to the pulsar passes close to the limb of the companion star, suggesting an intermediate inclination angle $i \sim 80^\circ$.

Strader et al. [54] obtained spectroscopic observations with the Goodman Spectrograph [268] on the SOAR telescope. The spectra suggest a mid-G-type companion star, with temperature $T \approx 5500$ K and variations of up to ± 200 K across the orbit attributed to heating from the pulsar. The spectroscopy also revealed a single-line

radial velocity curve whose semi-amplitude of $K_c = 324 \pm 5 \text{ km s}^{-1}$ implies an unseen primary with a minimum mass $M > 0.8M_\odot$. Strader et al. [54] modelled the GROND light curve, incorporating two large cold spots on the outer face of the companion star to account for the light curve asymmetry, and found an inclination angle $i \sim 55^\circ$, from which they deduce a heavy neutron star primary with $M \gtrsim 1.8M_\odot$.

The optical counterpart is also covered in the Second *Gaia* Data Release [DR2, 122, 269]. Using Equation (2) of Jordi et al. [270], the *Gaia* DR2 colour $G_{\text{BP}} - G_{\text{RP}} = 1.02$ implies an effective temperature of $T_{\text{eff}} = 5423 \pm 249 \text{ K}$, consistent with the spectroscopic temperature measured by Strader et al. [54]. The *Gaia* DR2 also provides a marginal parallax detection ($\varpi = 0.40 \pm 0.23 \text{ mas}$) for a minimum (95% confidence) distance of $d > 1.2 \text{ kpc}$, and a total proper motion of $\mu = 15.51 \pm 0.26 \text{ mas yr}^{-1}$, corresponding to a distance-dependent transverse velocity of $v(d) \approx 75 (d/1 \text{ kpc}) \text{ km s}^{-1}$. The systemic velocity (the radial velocity of the binary center of mass) measured from optical spectroscopy by Strader et al. [54] is just $6 \pm 3 \text{ km s}^{-1}$ indicating that the 3D velocity vector is almost entirely transverse.

5.3 Gamma-ray Observations

To update the gamma-ray analysis of J2039 from previous works [154, 267], we selected SOURCE-class gamma-ray photons detected by the *Fermi* LAT between 2008 August 04 and 2019 September 12. Photons were included from within a 15° region of interest (RoI) around J2039, with energies greater than 100 MeV, and with a maximum zenith angle of 90° , according to the “Pass 8” P8R3_SOURCE_V2 [173, 235] instrument response functions (IRFs) ².

We first investigated the gamma-ray spectral properties of 4FGL J2039.5–5617. We used the 4FGL catalog as an initial model for the RoI, and used the `gll_iem_v07.fits` and `iso_P8R3_SOURCE_V2_v1.txt` models to describe the Galactic and isotropic diffuse emission, respectively. We replaced 4FGL J2039.5–5617 in the RoI model with a point source at the *Gaia* DR2 position of the optical source. To model the source spectrum, we used a subexponentially-cutoff power-law spectrum typical for gamma-ray pulsars (4FGL),

$$\frac{dN}{dE} = K \left(\frac{E}{E_0} \right)^{-\Gamma} \exp \left(-a \left(\frac{E}{1 \text{ MeV}} \right)^b \right), \quad (5.1)$$

where the parameters $E_0 = 1 \text{ GeV}$ (“pivot energy”) and $b = 2/3$ (exponential index) were fixed at their 4FGL values, while the parameters K (normalisation), Γ (low-energy spectral index) and a (exponential factor) were free to vary during fitting. We performed a binned likelihood analysis using `fermipy` [271] version 0.18.0, with $0.05^\circ \times 0.05^\circ$ bins and 10 logarithmic energy bins per decade. For this analysis we utilized the “PSF” event types and corresponding IRFs, which partition the LAT data into quartiles based on the quality of the reconstructed photon arrival directions. All 4FGL sources within 25° of the optical counterpart position were included in the model. Using the “optimize” function of `fermipy`, the parameters of all sources in the region were updated from their 4FGL values one at a time to find a good starting point. We then performed a full fit for the region surrounding

²See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_essentials.html

J2039. The spectral parameters of all sources within 10° were free to vary in the fitting, as were the normalizations of the diffuse models and the spectral index of the Galactic diffuse model.

The gamma-ray source at the location of the optical counterpart is detected with test statistic $TS = 2167$ (the TS is defined as twice the increase in log-likelihood when the source is added to the model). The spectrum has a photon power-law index of $\Gamma = 1.4 \pm 0.1$ and an exponential factor of $a = (7 \pm 1) \times 10^{-3}$. The total energy flux above 100 MeV is $G_\gamma = (1.46 \pm 0.06) \times 10^{-11}$ erg cm $^{-2}$ s $^{-1}$. At an assumed distance of 1.7 kpc (from our optical light-curve modelling in Section 5.4.2), this gives a gamma-ray luminosity of $L_\gamma = (5.0 \pm 0.6) \times 10^{33}$ erg s $^{-1}$, assuming isotropic emission.

In gamma-ray pulsation analyses, photon weights are used to weight the contribution of each photon to a pulsation detection statistic to increase its sensitivity, and avoid the need for hard cuts on photon energy and incidence angle [114]. A weight w_j represents the probability that the j th photon was emitted by a target source, as opposed to by a fore/background source, based on the reconstructed photon energy and arrival direction, and a model for gamma-ray sources within the RoI. We computed these weights for photons whose arrival directions were within 5° of J2039 using `gtsrcprob`, again making use of the PSF event types. Within this region, there were 181,813 photons in total, with $\sum_j w_j = 3850$ “effective” photons. To speed up our timing analyses (Section 5.3.2) we additionally removed photons with $w < 0.1$, leaving 6571 photons which account for 93% of the expected pulsation signal-to-noise ratio [which is proportional to $\sum_j w_j^2$, 79].

The data set described above was used for the timing (Section 5.3.2) and orbital modulation analyses (Section 5.3.3) presented in this paper. For the pulsation search described in Section 5.3.1, we used an earlier data set which only covered data up to 2019 January 10 and used spectral parameters from a preliminary version³ of the 4FGL catalog when computing photon weights.

5.3.1 Gamma-ray Pulsation Search

Using the hierarchical search methods described by Pletsch and Clark [80], extended to provide sensitivity to binary pulsars (Chapter 2), we performed a search for gamma-ray pulsations in the weighted *Fermi*-LAT photon arrival times.

For this, it was necessary to search for an unknown spin frequency ν , spin-down rate $\dot{\nu}$, as well as the orbital period P_{orb} , pulsar’s time of ascending node T_{asc} , and pulsar’s projected semimajor axis $x = A_{\text{psr}} \sin i$, where A_{psr} is the (non-projected) semimajor axis, and i is the binary inclination angle. We did not search a range of sky positions as we used the precise Gaia position of the optical counterpart.

This 5-dimensional parameter volume is extremely large, and requires large computing resources and efficient algorithms to cover. To meet the large computational cost of the searches, we utilized the distributed volunteer computing system, *Einstein@Home* [85, 264]. Under this system, the parameter space is split into millions of smaller chunks which can be searched by a typical personal computer within a few hours. These “work units” are computed while volunteer’s computers are otherwise idle. We also ported our *Einstein@Home* search code from CPUs to GPUs, which has previously been done for radio pulsar searches [85]. The approximately

³<https://fermi.gsfc.nasa.gov/ssc/data/access/lat/fl18y/>

10,000 GPUs active on *Einstein@Home* increase the computing speed by an order of magnitude.

Despite this large computational resource, major efficiency gains and compromises are required to ensure that the computational effort of the search remains feasible. Key to improved efficiency is ensuring that the parameter space is covered by a grid of search locations that is as sparse as possible, yet sufficiently covers the volume to avoid missing signals. The required density is described by a *distance metric* – a function relating parameter space offsets to a corresponding expected loss in signal strength. This metric is described in Chapter 2.

In the binary pulsar search, the spin parameters are searched in the same way as they are in isolated pulsar searches [see, e.g., 79]. ν is searched via fast Fourier transforms (FFTs). The relevant range in ν is covered by a frequency-independent lattice.

The computational effort to search the orbital parameters depends linearly on the number of grid points. Searching the orbital parameters in a uniformly-spaced grid would be inefficient because the required metric spacing depends strongly on ν and x , i.e. at higher values for ν and x the grid needs to be denser (Chapter 2). To deal with the ν -dependency, we break down the search into discrete 8 Hz bands which are searched separately, and in each band the grid over the orbital parameters is designed to be dense enough for the maximum frequency in the band.

The orbital grid would be *optimal* if it has the lowest number of grid points such that each point in the parameter space is “covered”. A location in the parameter space is covered if the distance to the closest grid point is less than a chosen maximum. In inhomogeneous parameter spaces, the optimal grid is unknown. However, the required number of grid points N_{opt} for such a grid can be estimated using the distance metric under the assumption that locally the parameter space is sufficiently flat.

To search the inhomogeneous (x -dependent) orbital-parameter space efficiently, *optimized grids* are used [127]. These are built from *stochastic grids*, which are grids where grid points are placed stochastically while no two grid points are allowed to be closer than a minimum distance [129]. We create a stochastic grid with N_{opt} grid points and optimize it by nudging the position of each grid point one by one towards “uncovered space” using a neighboring cell algorithm [127]. After a few nudging iterations over all grid points the covering is typically sufficient for the search.

Using preliminary results from our optical modelling (see Section 5.4.2), obtained prior to the publication by Strader et al. [54] of spectroscopic radial velocities which better constrain T_{asc} , we constrained our orbital search space to $P_{\text{orb}} = 0.2279799(3)$ d and $T_{\text{asc}} = \text{MJD } 56884.9678(8)$. The range of expected x values was not well constrained by this model, and as the computing cost increases with x^3 we chose to initially search up to $x = 0.5$ lt-s, with the intention of searching to higher values should the search be unsuccessful.

The search revealed a signal with $\nu \approx 377$ Hz that was highly significant in both the initial semi-coherent and fully coherent follow-up search stages. The signal had $x \approx 0.47$ lt-s, which along with the companion’s radial velocity measurements by Strader et al. [54] gives a mass ratio of $q = M_{\text{psr}}/M_c = K_c P_{\text{orb}}/2\pi x \approx 7.2$, and a minimum companion mass of $M_c > 0.15 M_\odot$ assuming $i = 90^\circ$. These features conclusively confirm that the source is indeed a redback millisecond pulsar system,

which can now be named PSR J2039–5617.

5.3.2 Gamma-ray Timing

Following the discovery of gamma-ray pulsations, we used the *Fermi*-LAT data set to obtain a rotational ephemeris spanning 11 years. To do so, we followed the principles described by Kerr et al. [97], in which a template pulse profile $F(\phi)$ is produced, and the parameters $\boldsymbol{\lambda}$ of a phase model $\phi(t | \boldsymbol{\lambda})$, are fit to maximize the Poisson log-likelihood of the unbinned photon phases. Assuming that the weights derived in Section 5.3 represent the probability that each photon was emitted by the pulsar, then the contribution to the pulsation log-likelihood from the j th photon, with weight w_j , is a mixture model between a constant (i.e. uniform in phase) background rate and the template pulse profile, with mixture weights of $1 - w_j$ and w_j respectively. Hence, the overall log-likelihood is

$$\log L(\boldsymbol{\lambda} | t_j, w_j, F) = \sum_{j=1}^N \log [w_j F(\phi(t_j | \boldsymbol{\lambda})) + (1 - w_j)] \quad (5.2)$$

where t_j denotes the measured arrival time of the j th detected gamma-ray photon.

Folding the LAT data with the initial discovery ephemeris showed that the signal was not phase-connected over the entire data span, with the pulse profile drifting in and out of focus, indicative of a varying orbital period. Such effects are common among redback pulsars, and are attributed to variations in the quadrupole moment of the companion star coupling with the orbital angular momentum [e.g., 71, 98, 99]. These effects significantly complicate efforts to time redbacks over more than a few months [e.g., 100].

In previous works, these effects have been accounted for by adding a Taylor series expansion of the orbital frequency perturbations to the constant-period orbital phase model, where the derivatives of the orbital angular frequency become additional parameters in the timing model. However, this parameterization has a number of drawbacks. Large correlations between the orbital frequency derivatives greatly increase the time required for a Markov Chain Monte Carlo (MCMC) sampling procedure, which suffers from inefficient sampling and exploration during the “burn-in” phase for highly correlated parameter spaces. The Taylor series model also has poor predictive power as the orbital phase model “blows up” when extrapolating beyond the fit interval, making it difficult to extend an existing timing solution to incorporate new data. An astrophysical interpretation of the resulting timing solution is also not straightforward, as the measured orbital frequency derivatives depend on one’s choice of reference epoch (T_{asc}), and are not representative of long-term trends in P_{orb} due to e.g. mass loss from the system.

These problems are very similar to those encountered when timing young pulsars with strong “timing noise”: unpredictable variations in the spin frequency over time. To address these issues, modern timing analyses treat timing noise as a stationary noise process, i.e. a random process with a constant correlation function, on top of the long-term spin-down due to the pulsar’s braking [193].

Kerr et al. [97] used this method to time gamma-ray pulsars using *Fermi*-LAT data. To do this, a template pulse profile is constructed and cross correlated with the photon phases within weeks- or months-long segments to obtain a discrete

pulse phase measurement, or “time-of-arrival” (TOA), for each segment, and the stochastic noise process is fit to these phase measurements. Timing parameters can then be fit analytically to minimize the chi-square log-likelihood of the covariance-transformed TOA residuals including a Bayesian penalty factor for the required timing noise process. However, this procedure has the drawback that for faint pulsars the segment length required to obtain a significant TOA measurement can become very long, and phase variations due to timing noise *within* each segment can no longer be neglected. Of course, the timing noise within a segment cannot be accounted for without a description of the noise process, which in turn cannot be obtained without the TOAs, creating a circular problem.

While this circular problem can be partially overcome by fitting iteratively, we have developed a new method to fit the noise process using every individual photon, rather than obtaining and fitting discrete TOAs. To obtain this best-fitting function and its uncertainty, we apply the *sparse online Gaussian process* (SOGP) procedure developed by Csató and Oppen [272]. For purely Gaussian likelihoods, the Gaussian process framework would allow an exact posterior distribution for the noise process to be computed analytically [273]. In our case, however, the likelihood for each photon phase in Equation (5.2) is a mixture model of Gaussian peaks describing the template pulse profile with a constant background level. Seiferth et al. [274] describe how to apply the SOGP procedure to obtain an optimal Gaussian approximation to the posterior distribution for a stationary process with Gaussian mixture model likelihoods, and we use this formulation to derive our timing solution.

For J2039, we require a timing model which accounts for variations in the orbital phase, which we treat as a stationary random process. The overall goal is therefore to find the best-fitting continuous function describing the phase deviations from a constant orbital period model, given a prior covariance function ($C_0(t_1, t_2)$).

Before fitting, we must choose the form of the prior covariance function, and hyperparameters controlling its properties. Here we assumed a Matérn covariance function [273],

$$C_0(t_1, t_2) = \frac{h^2 2^{1-n}}{\Gamma(n)} \left(\frac{\sqrt{2n}}{\ell} |t_1 - t_2| \right)^n K_n \left(\frac{\sqrt{2n}}{\ell} |t_1 - t_2| \right), \quad (5.3)$$

where K_n is the modified Bessel function. The hyperparameters are the length scale, ℓ , controlling the time span over which the orbital period remains correlated, an amplitude parameter, h , which describes the expected magnitude of the orbital phase variations, and the degree n , which controls the smoothness of the noise process. In the limit of $n \rightarrow \infty$, this reduces to the simpler squared-exponential covariance function,

$$C_0(t_1, t_2) = h^2 \exp \left(-\frac{|t_1 - t_2|^2}{2\ell^2} \right). \quad (5.4)$$

In the frequency domain, a noise process with the Matérn covariance function of Eq. (5.3) has a power spectral density,

$$P(f) \propto h^2 \left(1 + \left(\frac{f}{f_c} \right)^2 \right)^{-(n+1/2)}, \quad (5.5)$$

i.e. constant below a “corner frequency” of $f_c = \sqrt{n}/\sqrt{2\pi\ell}$, and breaking smoothly to a power-law process with index $-(2n + 1)$ at higher frequencies.

With our chosen covariance function, we obtain a timing solution by varying the timing parameters λ and hyperparameters (ℓ, h, n) using the `emcee` Markov-Chain Monte Carlo (MCMC) algorithm [186]. At each MCMC sample, we use the PINT software package [275] to phase-fold the gamma-ray data according to the timing parameters, and then apply the SOGP method to find the best-fitting Gaussian approximation to the posterior distribution of the continuous function describing the orbital phase variations. This posterior is marginalized analytically, and the log marginal likelihood passed to the MCMC algorithm. This allows the MCMC process to optimize both the timing parameters and the hyperparameters of the prior covariance function simultaneously.

Using the best-fitting timing solution, we then re-fold the photon arrival times, and update the template pulse profile. This process is applied iteratively until the timing parameters and template pulse profile converges. For J2039, this required three iterations. The results from our timing analyses of J2039 are shown in Figure 5.1 and the resulting parameter estimates are given in Table 5.1.

We also show the amplitude spectra of the orbital phase variations and our best fitting covariance model in Figure 5.2. This spectrum was estimated by measuring the orbital phase in discrete segments of data, and performing the Cholesky least-squares spectral estimation method of Coles et al. [193]. This is only used to illustrate the later discussion (Section 5.5.4), while statements about the measured hyperparameter values are from the full unbinned timing procedure described above.

We have extended `TEMP02` [96] with a function that interpolates orbital phase variations between those specified at user-defined epochs. This allows gamma-ray or radio data to be phase-folded using the ephemerides that result from our Gaussian process model for orbital period variations.

5.3.3 Gamma-ray Variability

The subset of transitional redback systems has been seen to transition to and from long-lasting accretion-powered states, in which their gamma-ray flux is significantly enhanced [22, 161, 276]. To check for such behavior from J2039, we investigated potential gamma-ray variability over the course of the *Fermi*-LAT data span. In 4FGL, J2039 has two-month and one-year variability indices (chi-squared variability tests applied to the gamma-ray flux measured in discrete time intervals) of 44 with 48 degrees of freedom, and 13 with 7 degrees of freedom, respectively. Although the one-year variability index is slightly higher than expected for a steady source, we note that the gamma-ray light curves in Ng et al. [267] indicate that a flare from a nearby variable blazar candidate, 4FGL J2052.2–5533, may have contaminated the estimated flux from J2039 around MJD 57,100. The true variability is therefore likely lower than suggested by the slightly elevated annual variability index, and indeed the two-month variability index is consistent with a non-variable source.

We also checked for a potential gamma-ray eclipse, which may occur if the binary inclination angle is high enough that the pulsar passes behind the companion star around superior conjunction, as has been observed in the transitional MSP candidate 4FGL J0427.8–6704 [277, 278]. For J2039, this would occur for inclinations $i \gtrsim 78^\circ$, and could last for up to 7% of an orbital period, assuming a Roche-lobe filling

Table 5.1: Timing solution for PSR J2039–5617.

Parameter	Value
Astrometric Parameters ^a	
R.A. (J2000), α	$20^{\text{h}}39^{\text{m}}34^{\text{s}}.9681(1)$
Decl. (J2000), δ	$-56^{\circ}17'09''.268(1)$
Proper motion in R.A., $\mu_{\alpha} \cos \delta$ (mas yr ⁻¹)	4.2(3)
Proper motion in Decl., μ_{δ} (mas yr ⁻¹)	-14.9(3)
Parallax, ϖ (mas)	0.40(23)
Position reference epoch (MJD)	57205.875
Timing Parameters	
Solar System Ephemeris	DE430
Data span (MJD)	54682–58738
Spin frequency reference epoch, t_{ref} (MJD)	56100
Spin frequency, ν (Hz)	$377.22936337986(5)$
Spin-down rate, $\dot{\nu}$ (Hz s ⁻¹)	$-2.0155(6) \times 10^{-15}$
Spin period, P (ms)	$2.6509071060648(5)$
Spin period derivative, \dot{P}	$1.4164(4) \times 10^{-20}$
Pulsar’s semimajor axis, x (lt s)	0.47105(1)
Epoch of pulsar’s ascending node, T_{asc} (MJD)	56884.96698(2)
Orbital period, P_{orb} (d)	$0.227979805(3)$
Orbital period derivative, \dot{P}_{orb}	$8(5) \times 10^{-12}$
Amplitude of orbital phase noise ^b , h (s)	$3.9^{+2.2}_{-1.1}$
Correlation timescale ^b , ℓ (d)	156^{+127}_{-41}
Matérn function degree ^c , n	> 1.5
Derived properties ^d	
Shklovskii spin down, $\dot{\nu}_{\text{Shk}}$ (Hz s ⁻¹)	$(-0.37 \pm 0.02) \times 10^{-15}$
Galactic acceleration spin down, $\dot{\nu}_{\text{acc}}$ (Hz s ⁻¹)	1.2×10^{-17}
Spin-down power, \dot{E} (erg s ⁻¹)	2.5×10^{34}
Surface magnetic field strength, B_{S} (G)	2×10^8
Light cylinder magnetic field strength, B_{LC} (G)	8.8×10^4
Characteristic age, τ_{c} (yr)	4×10^9
Gamma-ray luminosity, L_{γ} (erg s ⁻¹)	$(5.0 \pm 0.6) \times 10^{33}$
Gamma-ray efficiency, $\eta_{\gamma} = L_{\gamma}/\dot{E}$	0.21

Notes. — Epochs and units are in the Barycentric Dynamical Time (TDB) system. The numerical values for timing parameters are the mean values of the MCMC samples, with 1σ uncertainties on the final digits quoted in brackets.

^a Astrometric parameters are taken from Gaia Collaboration et al. [122].

^b The hyperparameters h and ℓ have asymmetric posterior distributions, and so we report the mean value and 95% confidence interval limits in super- and subscripts.

^c The Matérn function degree n is poorly constrained by the data; we report only a 95% confidence lower limit.

^d Derived properties are order-of-magnitude estimates calculated using the following expressions [e.g., 28], which assume a dipolar magnetic field, and canonical values for the neutron-star moment of inertia, $I = 10^{45}$ g cm² and radius, $R = 10$ km: $\dot{E} = -4\pi^2 I \nu \dot{\nu}$; $B_{\text{S}} = \sqrt{-1.5 I c^3 \dot{\nu} \nu^{-3}} / (2\pi R^3)$; $B_{\text{LC}} = 4\pi^2 \sqrt{-I \dot{\nu} \nu^3} / c^3$; $\tau_{\text{c}} = \nu / 2\dot{\nu}$. The corrections to $\dot{\nu}$ due to transverse motion (the Shklovskii effect) and radial acceleration in the Galactic potential were applied prior to computing other derived properties, assuming $d = 1.7$ kpc from optical light curve modelling described in Section 5.4.2.




Figure 5.1: Results from gamma-ray timing analysis. Left panel: photon phases after folding with the original discovery ephemeris (with a constant orbital period). The intensity of each point represents the corresponding probability weight for that photon. The apparent loss of signal around MJDs 55,500 and 56,800 is due to the varying orbital period. Although present throughout the entire data set, the deviations between the true orbital phase and that predicted by the constant-orbital-period folding model are at their largest at these epochs. Center left panel: offset in the time of the pulsar’s ascending node from the initial constant orbital period ephemeris. In the timing procedure we fit for an “average” orbital phase, period and first frequency derivative, and model the orbital phase variations as a Gaussian process on top of this base model. Variations requiring a Gaussian process with a larger amplitude or more complexity suffer a Bayesian penalty factor. Black and red lines show the best-fitting orbital phase variations and the underlying “average” orbital model, respectively, for randomly selected samples from the MCMC process. Green and blue curves show the samples with the highest log marginal likelihood. The epochs of our optical observations are marked by horizontal dashed lines with the same color as the corresponding light curves in Section 5.4. Center right panel: as before but for the orbital period (i.e. derivatives of the curves in the previous panel). Right panel: photon phases after correcting for the orbital phase variations using the best-fitting parameter values.

companion. We modelled the eclipse as a simple “top-hat” function, in which the flux drops to zero within the eclipse, and used the methods described by Kerr [279], and applied to the eclipse of 4FGL J0427.8–6704 by Kennedy et al. [278], to evaluate the log-likelihood of this model given the observed photon orbital phases. We find that an eclipse lasting longer than 0.1% of an orbit is ruled out by the gamma-ray data with 95% confidence. We interpret this as evidence that the pulsar is not eclipsed, and will use this to constrain the binary inclination while modelling the




Figure 5.2: Power spectral density of the orbital phase noise process. The top panel shows the power spectral densities for the orbital phase variations. The green horizontal dot-dashed line shows the estimated measurement uncertainty level, to which the power spectrum breaks at high frequencies. The solid blue and red curves show the best-fitting Matérn covariance function model, and those of random samples from the MCMC process, respectively. The dashed blue line additionally includes the measurement noise level. The lower panels show the power spectra of the orbital phase residuals after whitening using the Cholesky decomposition of the model covariance matrix (i.e. accounting for the blue curve in the upper panel). The horizontal dotted line shows the estimated level which the noise power in 95% of independent trials should be below. The vertical line in both panels shows the time span covered by the *Fermi*-LAT data – noise power close to and below this frequency is suppressed by our inclusion of a single orbital period derivative in the timing model.

optical light curves in Section 5.4.2.

5.3.3.1 Gamma-ray Orbital Modulation

As noted previously, Ng et al. [267] discovered an orbitally modulated component in the gamma-ray flux from 4FGL J2039.5–5617. Using the now precisely determined gamma-ray timing ephemeris (see Section 5.3.2) we computed the orbital Fourier power of the weighted photon arrival times, finding $P = 29.7$ for a slightly more significant single-trial false-alarm probability of $p_{\text{FA}} = e^{-P/2} \approx 4 \times 10^{-7}$ compared to that found by Ng et al. [267]. Those authors found the modulation was not detected after MJD 57,040 and speculated that this could be due to changes in the relative strengths of the pulsar wind and companion wind/magnetosphere. We do see a slight leveling-off in the rate of increase of P with time; however it picks up again after MJD 58,100. Variations in the slope of this function due to statistical

Table 5.2: Gamma-ray spectral parameters in two orbital phase regions.

Parameter	$0 < \Phi \leq 0.5$	$0.5 < \Phi \leq 1$
Photon index, Γ	1.25 ± 0.13	1.42 ± 0.14
Exponential factor, a (10^{-3})	9.0 ± 1.3	5.7 ± 1.2
Photon flux ($10^{-8} \text{ cm}^{-2} \text{ s}^{-1}$)	1.8 ± 0.2	1.3 ± 0.2
Energy flux, G_γ ($10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$)	1.7 ± 0.1	1.3 ± 0.1

Notes. — Photon and energy fluxes are integrated over photon energies $E > 100 \text{ MeV}$. Uncertainties are at the 1σ level.

(Poisson) fluctuations can appear large when the overall detection significance is low [42], and so we do not consider this to be compelling evidence for long-term flux variability from the system.

The gamma-ray and X-ray orbital light curves are shown in Figure 5.3. We also find no power at higher harmonics of the orbital period, indicating an essentially sinusoidal profile. The gamma-ray flux peaks at orbital phase $\Phi = 0.25 \pm 0.03$ (pulsar superior conjunction), almost exactly half an orbit away from the X-ray peak, and has an energy-averaged pulsed fraction of $24 \pm 5\%$ [using the definition from Equation (14) of 79]. As noted by Ng et al. [267], this phasing might suggest an inverse Compton scattering (ICS) origin, as opposed to being the high-energy tail of the population responsible for X-ray synchrotron emission from the intra-binary shock, for example, which would be phase-aligned with the X-ray modulation.

To further investigate this modulation, we performed a second spectral analysis, using the same procedure as above, but additionally separating the photons into “maximum” ($0.0 < \Phi \leq 0.5$) and “minimum” ($0.5 < \Phi_{\text{orb}} \leq 1.0$) orbital phases. We fit the spectral parameters of J2039 separately in each component, while the parameters of other nearby sources and of the diffuse background were not allowed to vary between the two components. The results are given in Table 5.2 and the resulting spectral energy distributions shown in Figure 5.4. Subtracting the “minimum” spectrum from the “maximum” spectrum, we find an additional component peaking at around 1 GeV, and decaying quickly above that, whose total energy flux is around 30% of the flux at the orbital minimum. This model has a significant log-likelihood increase of $\Delta \log L = 14$ (TS = 28 for a false-alarm probability of 5×10^{-6} given 3 degrees of freedom) compared to our earlier model where the gamma-ray flux is constant with orbital phase.

Similar orbital modulation has been observed from a handful of other spider systems [187, 280–282]. In two of these systems the gamma-ray flux peaks at the same orbital phase as is seen here from J2039, and importantly, from the redback PSR J2339–0533 the orbitally modulated component appears to be pulsed in phase with the “normal” intrinsic gamma-ray pulses.

Using the timing solution from Section 5.3.2, we can now investigate any rotational phase dependence of the orbitally-modulated component. In Figure 5.5 we show the gamma-ray pulse profile, split into two equal orbital phase regions around the pulsar superior ($0 < \Phi \leq 0.5$) and inferior conjunctions ($0.5 < \Phi \leq 1$). We find that the estimated background levels, calculated independently in each phase region from the photon weights as $b = \sum_j w_j - w_j^2$ [28], are very similar between the two orbital phase selections, that the pulse profile drops to the background level




Figure 5.3: Orbital light curves of J2039 from *XMM-Newton* (lower panel) and *Fermi-LAT* (upper panel) observations. Data have been folded using the pulsar timing ephemeris from Section 5.3.2. The dashed red horizontal line on the gamma-ray light curve indicates the expected background level computed from the distribution of photon weights.

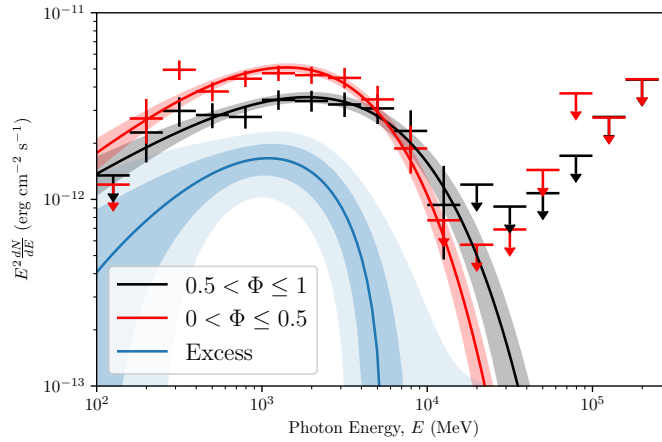


Figure 5.4: Gamma-ray spectral energy distributions (SEDs) for PSR J2039–5617, measured in two discrete orbital phase ranges around pulsar superior ($0.0 < \Phi \leq 0.5$) and inferior ($0.5 < \Phi \leq 1.0$) conjunctions. Error bars are derived by fitting the normalization of a power-law spectrum with index 2 to the flux measured in five discrete logarithmically spaced energy bands per decade. The deviating points at low energies are likely due to source confusion, as seen in the SEDs of several sources in 4FGL. The curved lines and shaded regions illustrate the best-fitting spectral models and one-sigma uncertainties in each phase interval. The blue curve and shaded regions show the difference between the spectral models measured in the two phase intervals.

in both, and that the gamma-ray pulse is significantly brighter around the pulsar




Figure 5.5: The gamma-ray pulse profile of PSR J2039–5617 measured in data taken in two equally-sized orbital phase regions around the pulsar superior (left) and inferior (right) conjunctions. The red dashed line indicates the background level, estimated independently in each orbital phase region using the distribution of photon probability weights. The gamma-ray pulse profile is clearly enhanced around superior conjunction, and there is no evidence for an unpulsed component in either orbital phase region.

superior conjunction. There is therefore no evidence for an unpulsed component to the gamma-ray flux from J2039, and the extra flux at the companion inferior conjunction is in fact pulsed and in phase with the pulsar’s intrinsic pulsed gamma-ray emission.

We consider two possible explanations for this orbitally-modulated excess. In these models, charged particles are accelerated in an inclined, fan-like current sheet at the magnetic equator that rotates with the pulsar. The intrinsic pulsed gamma-ray emission is curvature radiation seen when the current sheet crosses the line of sight. In the first scenario, the additional component is ICS from relativistic leptons upscattering the optical photon field surrounding the companion star. In the second, these leptons emit synchrotron radiation in the companion’s magnetosphere. These processes cause the normally unseen flux of relativistic leptons that is beamed towards the observer when the current sheet crosses the line of sight to become detectable as an additional pulsed gamma-ray flux that is coherent in phase with the intrinsic emission. We shall defer a full treatment of this additional emission component to a future work (Voisin, G. et al. 2020, in prep), and instead discuss some broad implications of the detection.

In the ICS scenario, it appears unlikely that the ICS population and the population responsible for the intrinsic (curvature) emission share the same energy. Indeed, the typical energy of the scattered photons, about $E_s \sim 1$ GeV, suggests scattering in the Thomson regime (for leptons) with $E_s \sim \gamma_s^2 E_b$, where γ_s is the typical Lorentz factor of the scatterer and $E_b \sim 1$ eV is the energy of soft photons coming from the companion star. This implies $\gamma_s \sim 3 \times 10^4$ which fulfills the condition $E_s \ll \gamma_s mc^2$ necessary for Thomson regime scattering. On the other hand, the Lorentz factor required to produce intrinsic gamma rays at an energy $E_i \sim 2$ GeV is about $\gamma_i \sim 10^7$ assuming the mechanism is curvature radiation [as is favoured by 283]. We assumed a curvature radius equal to the light-cylinder radius $r_{LC} = 126$ km and a magnetic field intensity equal to $B_{LC} = 7 \times 10^4$ G in these estimates. Thus, the ICS scenario requires two energetically distinct populations of leptons in order to explain the orbital enhancement. Under this interpretation, the more relativistic curvature-

emitting population would also produce an ICS component peaking around 10 TeV, which may be detectable by future ground-based Cherenkov telescopes.

The synchrotron scenario, on the other hand, allows for the possibility that the same particle population responsible for intrinsic pulsed gamma-ray (curvature) emission can produce the orbital flux enhancement, provided the companion magnetic field strength is on the order of 10^3 G [284]. The synchrotron critical frequency in a 10^3 G field of the companion magnetosphere is ~ 1 GeV for a Lorentz factor of $\gamma_i = 10^7$, while the cooling timescale is about $10^{-5} - 10^{-4}$ s, i.e. leptons cool almost immediately after crossing the shock, and phase coherence can be maintained. Moreover, the particles are energetic enough to traverse the shock without being greatly influenced, and would emit in less than a single gyroperiod, so emission would likely be beamed in the same direction as the intrinsic curvature radiation.

For the pulsed orbital modulation in PSR J2339–0533, An et al. [282] also consider an alternative scenario in which intrinsic pulsed emission is absorbed around the pulsar’s inferior conjunction. This model explains the softer spectrum around the maximum, as leptons in the pulsar wind have a higher scattering cross section for low-energy gamma rays. However, they conclude that the pair density within the pulsar wind is far too low to provide sufficient optical depth.

5.4 Optical observations and Modelling

5.4.1 New optical observations

We performed optical photometry of J2039 with the high-speed triple-beam CCD camera ULTRACAM [195] on the NTT on 2017 June 18, 2018 June 02 and 2019 July 07. The first two observations each covered just over one full orbital period, while the third was affected by intermittent cloud cover throughout before being interrupted by thick clouds after 70% of an orbit had been observed. We observed simultaneously in u_s , g_s and i_s ⁴, with 13 s exposures (65 s in u_s) and negligible dead time between frames. Each image was calibrated using a bias frame taken on the same night and a flat-field frame taken during the same observing run.

All reduction and calibration was performed using the ULTRACAM software pipeline⁵ (GROND images were first converted to the ULTRACAM pipeline’s data format). Instrumental magnitudes were extracted using aperture photometry, with each star’s local per-pixel background count rate being estimated from a surrounding annulus and subtracted from the target aperture.

To calibrate the photometry, we took ULTRACAM observations of two Southern SDSS standard fields (Smith, J.A., et al. 2007, AJ, submitted)⁶ on 2018 June 01 and 2018 June 04. The resulting zeropoints were used to calibrate the ULTRACAM observations of J2039. Zeropoint offsets between 2017, 2018 and 2019 observations and frame-to-frame transparency variations were corrected via “ensemble photometry” [198] using a set of 15 stars that were present in all ULTRACAM and GROND images of J2039.

⁴ULTRACAM uses higher-throughput versions of the SDSS filter set, which we refer to as *Super-SDSS filters*: u_s , g_s , r_s , i_s , and z_s [197].

⁵<http://deneb.astro.warwick.ac.uk/phsaap/software/ultracam/html/>

⁶<http://www-star.fnal.gov/>

To calibrate the archival GROND data, we computed average magnitudes for five comparison stars that were covered in g_s and i_s by the ULTRACAM observations, and fit for a linear color term between the GROND and ULTRACAM filter sets. Neither r' nor z' were covered by ULTRACAM. In r' we therefore used magnitudes of four stars from the APASS catalogue [285]. No catalogs contained calibrated z' magnitudes for stars within the GROND images. We therefore adopted the reference GROND zeropoint⁷ in this band. The g' , r' and i' the GROND calibrations agreed with these reference zeropoints to within 0.07 mag. As a cross-check we derived alternative zeropoints using a set of stars in the images which have magnitudes listed in the APASS catalogue. For both GROND and ULTRACAM the APASS-derived zeropoints agree with the calibrations using the ULTRACAM standard-derived zeropoints to within 0.06 mag in both g_s and i_s .

The resulting light curves in the g_s and i_s bands (the only two bands covered by all four observations) are shown in Figure 5.6. The long-term changes in the light curve are clearly visible, with ~ 0.2 mag variability in the second maximum (near the companion star’s descending node) and ~ 0.1 mag variations in the minimum at the companion’s inferior conjunction. The apparent variations around the first maximum (companion’s ascending node), are closer to our systematic uncertainty in the relative flux calibrations.

To estimate the level of variability that can be attributed to our flux calibration, we checked the recovered mean magnitudes of the ensemble stars used to flux-calibrate the data. These all varied by less than 0.05 mag across all sets of observations.

5.4.2 Light curve modelling

To estimate physical properties of the binary system, we fit a model of the binary system to the observed light curves using the *Icarus* binary light curve synthesis software [194].

Icarus assumes point masses at the location of the pulsar (with mass M_{psr}) and companion star centre-of-masses, and solves for the size and shape of the companion star’s Roche lobe, according to an assumed binary mass ratio $q \equiv M_{\text{psr}}/M_c$, inclination angle i , projected velocity semi-amplitude K_c and orbital period P_{orb} . These parameters are linked through the binary mass function,

$$f(M_{\text{psr}}) = \frac{M_{\text{psr}} \sin^3 i}{(1 + 1/q)^2} = \frac{K_c^3 P_{\text{orb}}}{2\pi G}, \quad (5.6)$$

and hence only four out of these five values are independent. With the pulsation detection, we have an extremely precise timing measurement of P_{orb} , and the pulsar’s projected semimajor axis (x), which further fixes $q = K_c P_{\text{orb}}/2\pi x$. We therefore chose to fit for i and K_c , and derive q and M_{psr} from these. For i , we adopted a prior that is uniform in $\cos i$ (to ensure that the prior distribution for the orbital angular momentum direction is uniform over the sphere). Since no evidence is seen for a gamma-ray eclipse (see Section 5.3.3), we assume that the pulsar is not occluded by the companion star, which provides an upper limit on the inclination of $i \lesssim 79^\circ$ (the precise limit additionally depends on the size of the star, and was computed “on-the-fly” by *Icarus* while fitting). We additionally assumed a conservative lower

⁷<http://www.mpe.mpg.de/~jcg/GROND/calibration.html>




Figure 5.6: Folded orbital light curves of J2039 across four observing runs (2014 August 16–18 with GROND, 2017 June 18, 2018 June 02 and 2019 July 07 with ULTRACAM). For clarity, the ULTRACAM data points have been combined into 250s integrations. The folded light curves are repeated twice, with uncertainties shown only in the first cycle. These are mostly smaller than the corresponding markers. Here, and throughout this paper, orbital phase zero corresponds to the pulsar’s ascending node. The GROND light curves have been corrected to the ULTRACAM magnitude system via color corrections computed from the magnitudes of comparison stars in the field.

limit of $i > 40^\circ$, since lower inclinations would require an unrealistically high pulsar mass ($> 4 M_\odot$).

The size and shape of the star within the Roche lobe is parameterized by the Roche lobe filling factor f_{RL} , defined as the ratio between the radius from the star’s center-of-mass in the direction towards the pulsar and the distance between the star’s center-of-mass and the Lagrange L1 point.

Once the shape of the star has been calculated, the surface temperature of the companion star is defined by another set of parameters. The temperature model starts with the “night” side temperature of the star, T_n , which is the base temperature at the pole of the star prior to irradiation. We assumed a Gaussian prior on T_n with mean 5423 K and width ± 249 K taken from the *Gaia* color–temperature relation [270]. To account for gravity darkening, we modify the surface temperature at a given location for the local effective gravitational acceleration by $T_g = T_n (g/g_{\text{pole}})^\beta$, where g_{pole} is the effective gravitational acceleration at the pole. We used a fixed value of $\beta = 0.08$, which assumes that the companion star has a convective envelope [286].

We account for the effect of heating from the pulsar by modelling it as an isotropically emitting point source of heating flux, with luminosity L_{irr} (although note that the pulsar’s beam is generally more concentrated towards the equator, see [263] who account for this when fitting black-widow light curves). In *Icarus*, heating is parameterized by the “irradiation temperature” $T_{\text{irr}} = L_{\text{irr}} / (4\pi\sigma A^2)$, where σ is the Stefan-Boltzmann constant and $A = x(1 + q) / \sin i$ is the orbital separation. In the later discussion, we will compare this luminosity with the pulsar’s total

spin-down power via the heating efficiency, $\epsilon = L_{\text{irr}}/\dot{E}$ [176] which absorbs several unknown quantities such as the stellar albedo, and the “beaming factor” accounting for the pulsar’s non-isotropic emission. A location on the stellar surface which is a distance r from the pulsar, and whose normal vector is at an angle χ from the vector pointing to the pulsar, receives heating power of $\sigma T_{\text{irr}}^4 \cos \chi A^2/r^2$ per unit area. We assume that the star remains in thermal equilibrium, and so this flux is entirely re-radiated, and hence the surface temperature at this location is raised to $T = (T_g^4 + \cos \chi T_{\text{irr}}^4 A^2/r^2)^{1/4}$. To account for the light curve asymmetry and variability, we require additional parameters describing deviations from this direct-heating temperature model; these will be discussed below.

Given this set of parameters, Icarus computes model light curves in each band by solving for the stellar equipotential surface, generating a grid of elements covering this surface, calculating the temperature of each element as above, and simulating the projected flux (including limb darkening) from every surface element at a given inclination angle and at the required orbital phases. For the flux simulation, we used the model spectra from the Göttingen Spectral Library⁸ [203] produced by the PHOENIX [287] stellar atmosphere code. We integrated these model spectra over the transmission curves of the observing setups to obtain flux models in the ULTRACAM and GROND filters.

The flux was rescaled in each band for a distance d and reddening due to interstellar extinction, parameterized by the V-band extinction, A_V , for which we assumed a uniform prior between $0.0 < A_V < 0.14$, with the (conservative) upper limit being twice that found by Romani [153] from fits to the X-ray spectrum. Since the *Gaia* parallax measurement is marginal, we followed the recommendations of Luri et al. [288] to derive a probability distribution for the distance by multiplying the Gaussian likelihood of the parallax measurement, $p(\varpi|d)$, by an astrophysically motivated distance prior for MSPs. For this, we take the density of the Galactic MSP population along the line of sight to J2039 according to the model of Levin et al. [289]. This model has a Gaussian profile in radial distance from the Galactic center (r) with width $r_0 = 4.5$ kpc, and an exponential decay with height z above the Galactic plane, with scale height $z_0 = 0.5$ kpc. The transverse velocity distribution for binary MSPs in the ATNF Pulsar Catalogue [110] is well approximated by an exponential distribution with mean $v_0 = 100$ km s⁻¹, which we apply as an additional distance prior. In total, the distance prior is,

$$p(d) \propto p(\varpi|d) d^2 \exp \left[-\frac{v(d)}{v_0} - \frac{z(d)}{z_0} - \frac{1}{2} \left(\frac{r(d)}{r_0} \right)^2 \right], \quad (5.7)$$

where the d^2 term arises from integrating the Galactic MSP density model at each distance over the 2D area defined by the *Gaia* localization region. Finally, we used the radio dispersion measure, $\text{DM} = 24.6$ pc cm⁻³ (see Paper II) as an additional distance constraint. The Galactic electron density model of Yao, Manchester, and Wang [189] (hereafter YMW16) gives an estimated distance of $d = 1.7$ kpc, with nominal fractional uncertainties of $\pm 45\%$. We therefore multiplied the distance prior by a log-normal distribution with this mean value and width. This overall prior gives a 95% confidence interval of $1.2 \text{ kpc} < d < 3.0 \text{ kpc}$, with expectation value $\bar{d} = 1.9$ kpc.

⁸<http://phoenix.astro.physik.uni-goettingen.de/>

In our preliminary *Icarus* models, constructed prior to the spectroscopic observations by Strader et al. [54] and the pulsation detection presented here, we jointly fit all three light curves, and additionally fit for P_{orb} and T_{asc} . For this we used a Gaussian prior on P_{orb} according to the best-fitting period and uncertainty from the Catalina Surveys Southern periodic variable star catalogue [266, see Section 5.2], and refolded the optical observations appropriately. The resulting posterior distributions on P_{orb} , T_{asc} and on x were used to constrain the parameter space for the gamma-ray pulsation search in Section 5.3.1.

In these preliminary models, we accounted for the light curve asymmetry and variability by describing the surface temperature of the star using an empirical spherical harmonic decomposition whose coefficients could vary between the three epochs. While this model served our initial goal of phase-aligning the light curves to constrain the orbital parameters, the spherical harmonic temperature parameterization suffered from several deficiencies. Firstly, the decomposition had to include at least the quadrupole ($l = 2$) order to obtain a satisfactory fit. Several of these coefficients were highly correlated with one another, and polar terms ($m = 0$) are poorly constrained as the system is only viewed from one inclination angle, leading to very poor sampling efficiency. Secondly, the quadrupole term naturally adds power into the second harmonic of the light curve, changing the amplitude of the two peaks in the light curve. In the base model, this amplitude depends only on the inclination and Roche-lobe filling factor, and so the extra contribution of the quadrupole term made these parameters highly uncertain.

To try to obtain more realistic parameter estimates, we instead modelled the asymmetry and variability by adding a cold spot to the surface temperature of the star. While cool star spots caused by magnetic activity are a plausible explanation for variability and asymmetry in the optical light curves [261], other mechanisms such as asymmetric heating from the pulsar [205, 218], or heat re-distribution due to convective flows on the stellar surface [240, 290], may also explain this. Our choice to model the light curves using a cool spot came from this being a convenient parameterization for a temperature variation on the surface of the star, rather than from assuming that variability is due to magnetic star spot activity.

In our model, this spot subtracts from the gravity-darkened temperature of the star, with a temperature difference of τ at the center of the spot, which falls off with a 2D Gaussian profile with width parameter ρ in angular distance (Δ) from the center of the spot. The spot location on the surface of the star is parameterized by the polar coordinates (θ, ϕ) , with $\theta = 0$ aligned with the orbital angular momentum, $\phi = 0$ pointing towards the pulsar and $\phi = 90^\circ$ aligned with the companion’s direction of motion. We assumed a sinusoidal prior on θ to ensure our priors covered the surface of the star approximately uniformly (the approximation would be exact for a spherical, i.e. non-rotating and non-tidally distorted star). The spot width was confined to be $5^\circ < \rho < 30^\circ$. The lower limit prevents very small and very cold spots, while the upper limit ensures that the effects of spots do not extend over much more than one hemisphere.

To prevent over-fitting, we added an extra penalty factor on the total (bolometric) difference in flux that the spot adds to the model. This is, approximately, proportional to $I = \iint_S \tau^4 e^{-(\Delta(\theta, \phi)^2/2\rho^2)} dS$ where S is the surface of the star. In our fits we adopted a Gaussian prior on I , centred on $I = 0$ with width parameter $\sigma_I = 6.25 \times 10^{10} \text{ K}^4 \text{ sr}$, corresponding to a -500K spot covering 1 steradian of the

star’s surface. Noting from Figure 5.6 that the first peak (at the pulsar’s ascending node) is always larger than the second, and that the variability seems to be strongest around the second peak, we assumed in our model that the light curve asymmetry is due to a variable cold spot ($\tau < 0\text{ K}$) on the leading edge of the companion star, and confined $0^\circ < \phi < 180^\circ$.

To investigate the light curve variability and understand what effect this has on our inference of the binary parameters, we chose to fit each light curve separately. Here we only model the three complete light curves from 2014, 2017 and 2018. The partial 2019 light curve is missing the first peak, and hence models fit only to the data around the variable second peak would have very high uncertainties on the fit parameters, making this of limited use compared to the other three light curves. We included the ULTRACAM u_s data in our model fitting, since they were obtained simultaneously with the i_s and g_s data without requiring additional observing time, and provide an additional color for temperature estimation. However, as the signal-to-noise is much lower in this band, we do not expect it to have had a large effect on the results.

To account for uncertainties in our atmosphere models, extinction, or photometric calibration, we allowed for constant offsets in the magnitudes in each band, penalizing the chi-squared log-likelihoods using a Gaussian prior on the magnitude offset with a width of 0.05 mag. As the resulting reduced chi-squared was greater than unity, we also applied rescaling factors to the uncertainties in each band. Both the band calibration offsets and uncertainty rescaling factors were computed to maximize the penalized log-likelihood.

At each sampled location in the parameter space, `Icarus` additionally computed the projected velocity of every surface element, and averaged these weighting by their r' flux, to obtain a simulated radial velocity curve. This filter band was chosen as it covers the sodium absorption line seen in Strader et al. [54]. The simulated radial velocity curve was compared to the measured radial velocities from Strader et al. [54], additionally fitting for a constant systemic radial velocity, and the resulting chi-squared term added to the overall log-likelihood.

The model fits were performed using the `pymultinest` Python interface [202] to the `Multinest` nested sampling algorithm [201]. The best fitting models and light curves are shown in Figure 5.7, the posterior distributions for our model parameters are shown in Figures 5.8, 5.9 and 5.10, and numerical results are given in Table 5.3. While the inferred posterior distributions from each epoch generally overlap with each other (except for the spot and heating parameters encapsulating variability), in the following discussion we take the full range covered by the 95% confidence intervals of the three posterior distributions as estimates for the model uncertainty, in the hope that biases due to variability are contained within that range.

5.5 Results and Discussion

5.5.1 Binary Inclination and Component Masses

Perhaps one of the more important questions is whether or not we are able to obtain a reliable measurement of the mass of the neutron star in the system. The maximum neutron star mass is a crucial unknown quantity which can discriminate between different nuclear equations-of-state [see 291, and references therein]. Recent works

Table 5.3: Icarus fit results.

Parameter	2014 June 16–18 (GROND)	2017 June 18 (ULTRACAM)	2018 June 02 (ULTRACAM)
χ^2 (degrees of freedom)	958.6 (824)	3930.3 (3529)	3699.5 (3330)
Icarus fit parameters			
Systemic velocity (km s ⁻¹)	6 ^{+6.6} _{-6.5}	5.9 ^{+6.4} _{-6.4}	6.4 ^{+6.7} _{-6.6}
Companion’s projected radial velocity, K_c (km s ⁻¹)	327.3 ^{+8.9} _{-8.8}	329.6 ^{+8.6} _{-8.7}	325.4 ⁺⁹ _{-8.9}
Distance, d (kpc)	1.7 ^{+0.16} _{-0.14}	1.69 ^{+0.12} _{-0.11}	1.8 ^{+0.21} _{-0.17}
V-band extinction, A_V	0.072 ^{+0.064} _{-0.069}	0.091 ^{+0.047} _{-0.082}	0.095 ^{+0.043} _{-0.086}
Inclination, i (°)	74.5 ^{+3.7} _{-5.1}	77.1 ^{+1.2} _{-3.8}	69 ^{+8.8} ₋₈
Roche-lobe filling factor, f_{RL}	0.835 ^{+0.021} _{-0.014}	0.821 ^{+0.012} _{-0.008}	0.839 ^{+0.048} _{-0.033}
Base temperature, T_b (K)	5451 ⁺⁸² ₋₅₇	5395 ⁺⁸⁷ ₋₇₄	5471 ⁺¹²⁴ ₋₁₀₉
Irradiating temperature, T_{irr} (K)	3456 ⁺⁹⁶ ₋₇₂	3807 ⁺⁹⁴ ₋₈₅	3700 ⁺¹⁶⁰ ₋₁₄₁
Spot central temperature difference, τ (K)	-540 ⁺¹⁷⁰ ₋₁₅₀	-620 ⁺³⁴⁰ ₋₂₆₀	-600 ⁺⁷⁰ ₋₈₀
Spot Gaussian width parameter, ρ (°)	23 ^{+5.7} _{-3.9}	14.4 ^{+10.9} _{-3.7}	28 ^{+1.9} _{-3.9}
Spot co-latitude, θ (°)	42.7 ^{+9.1} _{-7.5}	33.1 ^{+16.1} _{-10.5}	96.7 ^{+13.2} _{-13.1}
Spot longitude, ϕ (°)	77.9 ⁺³ _{-4.6}	73.6 ^{+6.2} _{-12.4}	56.3 ^{+3.8} _{-4.6}
Derived parameters			
Heating efficiency, ϵ	0.061 ^{+0.011} _{-0.007}	0.089 ^{+0.012} _{-0.01}	0.086 ^{+0.03} _{-0.019}
Pulsar mass, M_{psr} (M_\odot)	1.2 ^{+0.14} _{-0.11}	1.18 ^{+0.1} _{-0.09}	1.3 ^{+0.31} _{-0.2}
Companion mass, M_c (M_\odot)	0.165 ^{+0.017} _{-0.012}	0.162 ^{+0.011} _{-0.008}	0.18 ^{+0.041} _{-0.026}
Mass ratio, $q \equiv M_{psr}/M_c$	7.27 ^{+0.2} _{-0.2}	7.32 ^{+0.19} _{-0.19}	7.22 ^{+0.2} _{-0.2}
Volume-averaged companion density (g cm ⁻³)	4.38 ^{+0.1} _{-0.14}	4.48 ^{+0.07} _{-0.09}	4.35 ^{+0.25} _{-0.27}
Spot integral, I (10^{10} K ⁴ sr)	8.4 ^{+8.8} _{-5.7}	6.0 ^{+9.4} _{-5.3}	18.1 ^{+8.3} _{-6.5}

Notes. — Numerical values are the median of the marginalized posterior distributions output by `Multinest`, with the 95% confidence regions shown in sub- and superscript.




Figure 5.7: Best-fitting optical light curve models for J2039. Each row shows the best-fitting model for a given epoch. Left panels show the observed light curves and uncertainties (colored error bars) in each band. The fluxes in each band predicted by the best-fitting *Icarus* model are shown as dashed curves. When fitting these models we allowed for small offsets in the flux calibration of the observed light curves. The solid curves show the model light curves after applying these calibration offsets. Right panels show the *Icarus* model according to the best-fitting parameters at the pulsar’s ascending (A) and descending (D) nodes, and superior (S) and inferior (I) conjunctions, marked on the top right panel, with the direction of motion shown by an arrow. The pulsar’s position at each phase is shown by a black dot. Phase zero in the light curves corresponds to the pulsar’s ascending node. The axes are in units of the orbital separation (A). The surface temperature of the companion star is shown by the color bar.

[292, 293] have found very heavy pulsar masses for spider pulsars, and there are hints that these systems may in general contain heavier neutron stars than e.g. double neutron star systems [54].

Combining the radial velocity curve measured via spectroscopy by Strader et al. [54] with our pulsar timing measurement of the pulsar’s projected semi-major axis constrains the mass ratio to $q = 7.3 \pm 0.2$. Hence, all parameters in the binary




Figure 5.8: Posterior distributions for the Icarus model parameters. Red, blue and orange histograms and contours show the posterior distributions from fits to the GROND, ULTRACAM (2017) and ULTRACAM (2018) light curves, respectively. Contour lines are shown at 1σ and 2σ levels. Where a non-uniform prior is assumed, this is shown as a black curve on the corresponding parameter’s one-dimensional histogram.

mass function (Equation 5.6) are relatively well measured, with the exception of the binary inclination angle. Measuring this by modelling the optical light curves was therefore a key goal for our study of this system. The observed asymmetry and variability in the light curve are significant complicating factors for this, as estimates for the inclination angle are determined by the amplitude of the ellipsoidal peaks. In J2039, these do not have equal amplitudes, and vary over time.

Our optical model fits to all three complete orbital light curves consistently preferred high inclinations, hitting the upper limit of ($i \lesssim 79^\circ$) imposed by our assertion that the pulsar is not eclipsed at superior conjunction. The second ULTRACAM light curve results in the widest 95% confidence interval, with $61^\circ < i < 78^\circ$. Marginalizing over the uncertainty in the radial velocity amplitude, the corresponding pulsar mass range is $1.1M_\odot < M_{\text{psr}} < 1.6M_\odot$, with a median of $M_{\text{psr}} \approx 1.3M_\odot$,




Figure 5.9: Posterior distributions from Icarus model fitting, as in Figure 5.8, but for the binary inclination, pulsar and companion masses.




Figure 5.10: Posterior distributions from Icarus model fitting, as in Figure 5.8, but for the parameters of the cold spot added to the companion's surface. These parameters account for the significant variability observed in the light curves, hence the rather different values recovered.

but the models for the other two epochs give narrower ranges $1.1 M_{\odot} < M_{\text{psr}} < 1.35 M_{\odot}$. The posterior distributions on these parameters are shown in Figure 5.9.

Inclination angles derived from optical light curve fits are highly dependent on

the chosen temperature and irradiation models and priors. In particular, we caution that there is likely to be a large (but unknown) systematic uncertainty underlying our inclination estimates, caused by our simplifying assumption that the variability and asymmetry can be modelled by one cold spot on the leading face of the star. Other models for light curve asymmetry, e.g. intra-binary shock heating models or models featuring convective winds on the stellar surface [218, 240, 290], may give different results. Pulsar masses derived from optical light curve modelling should therefore be treated with caution, as the results can be highly model-dependent. For instance, if part of the asymmetry is caused by excess heating on the trailing face of the star rather than a cool spot on the leading face, then the leading peak of the light curve will be larger than predicted by the direct-heating model, and the model’s inclination angle will increase to compensate. Nevertheless, our results suggest that a high inclination and fairly low pulsar mass is compatible with the observed light curves.

Our resulting mass is rather lower than those inferred from other redback systems, which Strader et al. [54] found to cluster around $1.8 M_{\odot}$, but has a range similar to that found for PSR J1723–2837 ($M_{\text{psr}} < 1.4 M_{\odot}$) by van Staden and Antoniadis [261]. While some of the redback masses compiled by Strader et al. [54] do have strict lower limits (i.e. for edge-on orbits) that are above our inferred mass range, it is possible that unmodelled asymmetries and variability may be systematically biasing optical-modelling based inclination measurements to lower values, and hence biasing the redback pulsar mass distribution towards higher values.

By generating and fitting a flux-averaged radial velocity curve, our binary system model additionally corrects for possible biases in the observed radial velocity curve due to a difference between the center of mass of the companion star and the position on the surface where spectral lines contribute most strongly to the observed spectra [e.g., 293]. For J2039, heating has a fairly small effect on the light curve, and the resulting correction to the radial velocity curve is small: the epoch with the largest inferred center-of-mass radial velocity amplitude (2017 June 18) has $K_2 = 330 \pm 5 \text{ km s}^{-1}$, compared to $K_2 = 324 \pm 5 \text{ km s}^{-1}$ that Strader et al. [54] found from a simple sinusoidal fit. This implies that the required K_2 -correction is only $\Delta K_2/K_2 \lesssim 2\%$, and only increases the inferred pulsar mass by $\Delta M_{\text{psr}}/M_{\text{psr}} \lesssim 6\%$. While here this additional bias is far lower than that caused by our uncertainty on the inclination, this is not true in general for other redback systems. Large changes in the heating of redback companions have been observed [262], and so reliable center-of-light corrections require photometry observations to be taken as close in time as possible to spectroscopic radial velocity measurements to mitigate possible errors due to variations in heating.

Our pulsar mass range is lower than that estimated by Strader et al. [54] ($M_{\text{psr}} > 1.8 M_{\odot}$) from similar fits to the GROND light curve. Prior to our pulsation detection, the binary mass ratio was unconstrained, and so this was an additional free parameter in their model. The authors used two large cold spots in their model, which were both found to lie towards the unheated side of the star. These spots will affect the amplitudes of both ellipsoidal peaks, and therefore will affect the estimation of the inclination angle, filling factor and mass ratio that are constrained by these amplitudes. Their fits found a much lower mass ratio than is obtained from the pulsar’s semimajor axis measurement ($q < 5.3$ vs. $q = 7.3 \pm 0.2$ here) and a nearly Roche-lobe filling companion $f_{\text{RL}} \approx 95\%$. Both of these parameter differences will

increase the amplitude of the ellipsoidal modulations, allowing for a more face-on inclination and thus a heavier pulsar, explaining our disagreement.

The inferred inclination angle is also (qualitatively) consistent with the observed gamma-ray pulse profile. Since the pulsar has been spun-up via accretion its spin axis should be aligned to the orbital axis, and hence the pulsar viewing angle (the angle between the line-of-sight and the pulsar’s spin axis) will match the orbital inclination. The gamma-ray pulse profile features one broad main peak, with a smaller trailing peak. This therefore rules out an equatorial viewing angle to the pulsar, and hence an edge-on orbital inclination $i \sim 90^\circ$ as in that case the gamma-ray pulse should exhibit two similar peaks approximately half a rotation apart. The detection of radio pulsations enables a full investigation of this, fitting both the gamma-ray pulse profile shape and its phase relative to the radio pulse using theoretical pulse emission models. This will be described in detail in Paper II, but we note here that these models suggest a lower viewing angle of $i \sim 67^\circ$, for a pulsar mass of $M_{\text{psr}} \sim 1.4 M_\odot$.

For the companion mass, we find $0.15M_\odot < M_c < 0.22M_\odot$. Our *Icarus* model fits gave the companion star base temperature $T_n \approx 5400$ K and volume-averaged radius $R_c \approx 0.4 R_\odot$. These are both significantly larger than would be expected for a main-sequence star of the same mass. Indeed, this is not surprising, as the accretion required to recycle the MSP will have stripped the majority of the stellar envelope, while tidal forces and heating from the pulsar continue to add additional energy into the companion star [101], causing a further departure from ordinary stellar evolution.

5.5.2 Distance and Energetics

The *Icarus* fits to our three light curves all returned consistent distance estimates around $d = 1.7_{-0.1}^{+0.3}$ kpc, consistent with the *Gaia* parallax and YMW16 DM distance estimates. Assuming a fiducial distance of $d = 1.7$ kpc, the *Gaia* proper motion measurement implies a transverse velocity of $v_T \approx 125$ km s⁻¹. This transverse velocity will induce an apparent linear decrease in both the spin and orbital frequencies due to the increasing radial component of the initially transverse velocity [hereafter referred to as the Shklovskii effect after 213]. This effect accounts for around 20% of the observed spin-down rate. An additional contribution to the observed spin-down rate comes from the pulsar’s relative acceleration due to the Galactic rotation and gravitational potential. Using the formula given by Matthews et al. [294] and references therein, we estimate this accounts for less than 1% of the observed spin-down rate. At the fiducial distance the gamma-ray flux corresponds to a luminosity of $L_\gamma = 5 \times 10^{33}$ erg s⁻¹, or a Shklovskii-corrected gamma-ray efficiency of $\eta_\gamma = L_\gamma/\dot{E} = 21\%$, which is typical for gamma-ray MSPs [28]. Recently, Kalapotharakos et al. [283] discovered a “fundamental plane” linking pulsars’ gamma-ray luminosities to their spin-down powers, magnetic field strengths and spectral cut-off energies [283]. For J2039, this predicts $L_{\gamma,\text{FP}} = 1.3 \times 10^{34}$ erg s⁻¹, or 0.4 dex above the observed value, consistent with the scatter about the fundamental plane seen by Kalapotharakos et al. [283].

In our *Icarus* model, we assume that the inner side of the companion star is heated directly by flux from the pulsar. For PSR J2039–5617, our optical models hint that the heating flux reaching the companion star may be variable, and is on

the order of a few percent of the total spin-down luminosity of the pulsar, with $\epsilon = L_{\text{irr}}/\dot{E} \sim 0.05$ to 0.12 . This is a somewhat lower efficiency than is typically observed in spider systems, where heating normally accounts for around 20% of the pulsar’s spin-down power [176, 263].

The precise nature of the mechanism by which redback and black-widow pulsars heat their companions is currently unclear. For J2039, the inferred gamma-ray luminosity is larger than the heating power, and so we may infer that gamma rays are a sufficient heating mechanism in this case. For other spiders, this is not always true, with heating powers found to be much larger than gamma-ray luminosities (e.g., Chapter 3). Some discrepancy between the two can be explained by underestimated distances, or beamed (i.e. non-isotropic) gamma-ray flux that is preferentially emitted in the equatorial plane, although heating efficiencies and gamma-ray efficiencies remain only loosely correlated even with these corrections [263]. This may indicate that another mechanism, e.g. high-energy leptons in the pulsar wind, is responsible for heating the companion star. Note that both η_γ and ϵ are fractions of \dot{E} , so while \dot{E} is an order-of-magnitude estimate dependent on the chosen value for the pulsar moment of inertia, the ratio between η_γ and ϵ is independent of this.

5.5.3 Optical light curve asymmetry and variability

In the above heating efficiency calculation, we only included *direct* heating i.e. flux from the pulsar that is immediately thermalized and re-radiated from the surface of the companion star at the location on which it impinges. For J2039 the asymmetry of the light curve, and relative lack of variability on the leading peak may suggest that some heating is being re-directed toward the trailing face of the companion star, keeping this side at a more constant temperature. However, with only three optical light curves covering this orbital phase this is purely speculative, and requires additional optical monitoring to check for variability in the leading peak.

Nevertheless, similar light curve asymmetry, with the leading peak typically appearing as the brighter of the two, seems to be common in many types of close binary systems (e.g. cataclysmic variables (CVs) and W UMa-type eclipsing binaries), where it is often referred to as the *O’Connell effect* [after 295]. Several processes have been proposed to explain this in general, and in redbacks in particular, but so far without consensus. Possible processes include: reprocessing of the pulsar wind by a swept-back asymmetric intra-binary shock [218]; channeling of charged particles in the pulsar wind onto the poles of a companion’s misaligned dipolar magnetic field [205, 284]; or heat redistribution due to fluid motion in the outer layers of the star [240, 290, 296].

For J2039, the presence of an intra-binary shock wrapping around the pulsar is required to explain the observed orbital modulation of X-rays. Following the model of Romani and Sanchez [218], it therefore seems plausible that extra heating flux could be directed at the trailing face of the companion star, and could at least partially explain the observed light-curve asymmetry. We are then left to explain the variability in the light curve. Cho, Halpern, and Bogdanov [262] observe similar variability in the light curves of three other redback systems, attributing this to variability in the stellar wind and hence in the intra-binary shock.

An alternative explanation for redback variability is that magnetic activity in the companion leads to large cool star spots on the stellar surface, which migrate around

the star and may appear and disappear over time. This star-spot interpretation has been invoked to explain the similar optical variability seen in long-term monitoring of the redback system PSR J1723–2837 [261]. A periodogram analysis of these light curves found a component with a period slightly shorter than the known orbital period, which the authors interpret as being due to asynchronous (i.e. non-tidally locked) rotation of the companion star. Alternatively, this could also be due to differential rotation of the stellar surface, as seen in sun spots, and observed e.g. in CV secondaries via Roche tomography [e.g., 297]. Given the year-long time intervals between our ULTRACAM light curves of J2039 we cannot perform the same analysis to track a single variable component over time to confirm this picture, but this may be possible in the future with sufficient monitoring. Another interesting question that may be addressed with additional monitoring is whether or not the optical variability correlates with the variations in the orbital period, as both may be linked through magnetic cycles in the stellar interior.

To create our binary system models, we used a toy model for the stellar surface temperature that included a variable cold spot to account for the asymmetry and variability. The posterior distributions on the parameters of these spots are shown in Figure 5.10. This model is certainly an over-simplification of the truth, and so we will avoid placing much emphasis on the numerical results for these parameters, noting that our goal was instead to marginalize over the variability to retrieve estimates for more tangible quantities such as the inclination and filling factor. Our chosen prior, which aims to minimize the bolometric flux $\propto \tau^4 \sigma^2$ subtracted by the cool spot, penalizes small but very cold spots over larger and warmer spots. This prevents our model reaching the very cold spot temperatures ($\tau \sim -2000$ K) that have been observed in well-studied main-sequence stars [298]. Instead, our model prefers large spots (close to our upper limit of $\rho = 30^\circ$) with a central temperature difference between $\tau \sim -300$ K to $\tau \sim -700$ K. While such a temperature reduction could be plausibly explained by magnetic star spot activity, we are hesitant to interpret these as “true” star spots, but rather consider them to be areas of decreased temperature due to unknown variable effects, e.g. asymmetric heating from the pulsar, or heat re-distribution due to convective flows on the stellar surface. Continued photometric monitoring of J2039 to test the star-spot explanation may reveal evidence that these cool areas migrate across the surface of the star, as they do in PSR J1723–2837 [261]. We discuss this possibility further below. Furthermore, a dedicated study of the spectra observed by Strader et al. [54] may be able to detect the presence of spectral lines associated with cooler temperatures to further investigate the star-spot hypothesis.

We also note that a better understanding of variability in rotationally powered redback systems may offer insight into some of the most extreme behavior exhibited by binary MSP systems: the sudden (dis)appearance of accretion discs in transitional MSP systems [tMSPs, 19–22]. To provide material to power a tMSP’s accretion state, the companion star must be overflowing its Roche lobe. However, optical modelling of PSR J1023+0038 somewhat surprisingly suggests a companion that significantly underfills its Roche lobe [299, 300]. This therefore requires a significant change in the radius of the companion star, and the timescale on which this takes place is currently unknown. For J2039, we also find that the companion star is significantly smaller than its Roche lobe ($f_{\text{RL}} \approx 0.83$), and do not find any evidence for variations in the stellar radius over the three light curves.

5.5.4 Orbital Period Variability

In Section 5.3.2 we measured the orbital period of J2039, finding significant deviations in the orbital phase from a constant-period model. Such variations are common among redback systems [e.g., 19, 71, 100]. This phenomenon has been attributed to the Applegate mechanism [101, 102], originally invoked to explain period variations in eclipsing Algol-type and CV binaries, in which periodic magnetic activity cycles in the convective zone of the companion star introduce a varying quadrupole moment, which couples with the orbital angular momentum to manifest as variations in the orbital period.

Using our new Gaussian process description for the orbital phase variations, we can hope to quantify the required changes in the quadrupole moment using the best-fitting values for the hyperparameters of the Gaussian process used to model the orbital phase variations in Section 5.3.2.

Under the Applegate model, the change in orbital period is directly related to the change in the companion star’s gravitational quadrupole moment Q [102],

$$\frac{\Delta P_{\text{orb}}}{P_{\text{orb}}} = -9 \frac{\Delta Q}{M_c A^2}, \quad (5.8)$$

where $A = x(1 + q)/\sin i$ is the orbital separation. For comparison, the total quadrupole moment induced by the spin of the companion star and the tidal distortion in the pulsar’s gravitational field is [103]

$$\frac{Q}{M_c A^2} = -\frac{2}{9} k_2 \left(\frac{R_c}{A} \right)^5 (4q + 1), \quad (5.9)$$

where R_c is the radius of the companion star and k_2 is the apsidal motion constant, a parameter describing the deformability of the companion star [301]. For solar-type stars $k_2 \sim 0.035$ [302], while if we assume that redback companions are akin to the companions in CV systems whose outer envelopes have also been stripped through accretion then we may expect a smaller value $k_2 \sim 10^{-3}$ [303]. For J2039, the hyperparameter $h = 3.9_{-1.2}^{+2.2}$ s corresponds to the typical fractional amplitude for the variations in orbital phase. Taking the simpler squared exponential covariance function of Equation (5.4) corresponding to $n \rightarrow \infty$ then the deviations in orbital period have covariance function,

$$\begin{aligned} K_{\Delta P_{\text{orb}}/P_{\text{orb}}}(t_1, t_2) &= \frac{\partial^2 K}{\partial t_1 \partial t_2} \\ &= \frac{h^2}{l^2} \exp\left(-\frac{(t_1 - t_2)^2}{2\ell^2}\right) \left(1 - \frac{(t_1 - t_2)^2}{\ell^4}\right). \end{aligned} \quad (5.10)$$

The typical (fractional) amplitude of the orbital period variations is therefore $\Delta P_{\text{orb}}/P_{\text{orb}} \sim h/\ell = (3 \pm 1) \times 10^{-7}$, corresponding to $\Delta Q/Q \sim 3 \times 10^{-5} k_2^{-1}$. The time-varying component to the gravitational quadrupole moment is therefore required to be of order a few percent of the total expected quadrupole moment at most to explain the observed orbital period variations. From this, it seems plausible that the observed period variations can be powered by quadrupole moment changes, without requiring that a large fraction of the star be involved in the process. The required fractional quadrupole moment changes are very similar to those recently

calculated for the companion to the black widow PSR J2051–0827 by Voisin et al. [104], despite the large difference in their masses.

For our assumed Matérn covariance function, the parameter n is related to how smooth the noise process is: random walks in orbital phase, period or period derivative would manifest as noise processes with $n = 1/2, 3/2$ or $5/2$ respectively [97]. This hyperparameter may therefore encode information about the source of the orbital period variation. If the quadrupole moment exhibits random walk behavior (i.e. the stellar structure switches rapidly between different states), we would expect to see a random walk in orbital period ($n = 3/2$). Alternatively, if the system is affected by a variable torque (e.g. variable mass loss, or magnetic braking) then this would manifest as a random walk in the orbital period derivative or higher orders ($n \gtrsim 5/2$).

Unfortunately, Figure 5.2 illustrates that we are insensitive to the value of n , as the variability quickly falls below the measurement uncertainty level for periods shorter than $\ell \approx 130$ d, preventing measurement of the power-law slope above the corner frequency. We find only that a very shallow power-law spectrum $n < 1$ is ruled out with 95% confidence, but models with finite $n > 1.5$ fit the data equally well as the squared exponential kernel corresponding to $n \rightarrow \infty$.

We also find marginal evidence for an excess in the noise power at periods longer than the 11-years of *Fermi*-LAT data. This is not well accounted for by a longer correlation timescale ℓ and shallower spectral index, as this leaves excess power at intermediate frequencies, and we do find that a break in the spectrum is preferred, with $\ell \ll T_{\text{obs}}$. One explanation could be that instead of breaking to a constant power level at low frequencies, the noise process breaks to a shallower power law. However, with only a handful of independent frequencies below the corner frequency, this slope is hard to probe, although this may be worth revisiting as the timing baseline grows.

Alternatively, this low-frequency excess could be explained by a steadily increasing orbital period, which would introduce a quadratic term in the orbital phase that would appear as noise power at a period longer than the observation timespan. In Section 5.3.2, we accounted for this by including a constant orbital period derivative $\dot{P}_{\text{orb}} = (8 \pm 5) \times 10^{-12}$.

While there are several physical processes which could lead to a long term increase in the orbital period on top of the Applegate-style stochastic variability, the magnitude of the effect here is hard to reconcile. For example, the Shklovskii-induced orbital period derivative is $\dot{P}_{\text{orb,Shk}} = v_{\text{T}}^2 P_{\text{orb}} / cd = 2 \times 10^{-14}$, almost three orders of magnitude smaller than the measured value. Other incompatible explanations for an apparent period derivative include acceleration in the Galactic potential ($\dot{P}_{\text{orb,acc}} = -6 \times 10^{-16}$), or loss of angular momentum due to gravitational wave emission, which would decrease the orbital period and hence has the wrong sign here.

In principle, a long-term increase in the orbital period could be explained by steady mass loss from the system. Under this model, the inferred mass-loss rate would be $\dot{M} = -0.5 (M_{\text{c}} + M_{\text{psr}}) \dot{P}_{\text{orb}} / P_{\text{orb}} = -8 \times 10^{-9} M_{\odot} \text{yr}^{-1}$. This is an extremely high rate, and implies that the companion star would be completely ablated after just 19 Myr, assuming a constant mass-loss rate. If we assume that such a mass loss is driven by material ablated from the companion star by the pulsar, then the total power budget available for this process should be similar to the spin-down

power of the pulsar. Centrifugal effects from the orbital motion reduce the gravitational potential difference which must be overcome for matter to escape the system. Denoting the potential at the stellar surface, and the maximum potential within the system as φ_c and φ_{esc} , respectively, then an estimate for the maximum possible mass-loss rate due to ablation, assuming 100% efficiency and isotropic emission from the pulsar, is,

$$\dot{M}_{\dot{E}} = \frac{\dot{E} R_c^2}{4A^2} \frac{1}{\varphi_c - \varphi_{\text{esc}}}. \quad (5.11)$$

Calculating φ_c and φ_{esc} using *Icarus*, we find for J2039 $|\dot{M}_{\dot{E}}| \lesssim 1.5 \times 10^{-8} M_{\odot} \text{ yr}^{-1}$, and so at first glance it seems that mass-loss through ablation by the pulsar may be sufficient to explain the observed \dot{P}_{orb} . However, studies of radio eclipses in redback and black widow systems, in which radio pulsations are absorbed, dispersed and scattered by diffuse plasma in an extended region outside the companion star's Roche lobe, typically infer mass-loss rates on the order of $\dot{M} \sim 10^{-12} M_{\odot} \text{ yr}^{-1}$ or lower [e.g., 179, 304]. These mass-loss rates are therefore clearly incompatible with a mass-loss interpretation for the potential long-term period increase. Radio eclipses have been observed from J2039, and these will be investigated in Paper II.

Another alternative mechanism that could lead to a significant \dot{P}_{orb} is that considered by van Staden and Antoniadis [261], in which asynchronous rotation of the companion star leads to a tidal force that transfers angular momentum from the star to the orbit. If the star spins down at a constant rate, $\dot{\Omega}_c$, then conserving total angular momentum gives

$$\dot{P}_{\text{orb}} = -3I_c \dot{\Omega}_c M_{\text{psr}} M_c \left(\frac{2\pi (M_{\text{psr}} + M_c) P^2}{G^2} \right)^{1/3}, \quad (5.12)$$

where I_c is the companion star's rotational moment of inertia. Following Zahn [305], if the star rotates with an angular frequency $\Delta\Omega_c$ larger than the synchronous frequency $\Omega_0 \simeq \Omega$, which we can approximate by the orbital frequency Ω due to the much larger angular momentum of the orbit compared to the spin of the star, then tidal forces will reduce $\Delta\Omega_c$ to zero over the *synchronisation timescale*,

$$t_{\text{sync}} = -\frac{\Delta\Omega_c}{\dot{\Omega}_c} = \frac{I_c}{6k_2 q^2 M_c R_c^2} \left(\frac{M_c R_c^2}{L_c} \right)^{1/3} \left(\frac{A}{R_c} \right)^6, \quad (5.13)$$

where L_c is the star's luminosity. This expression assumes that the star has a large convective envelope. Rearranging for \dot{P}_{orb} , we find $\dot{P}_{\text{orb}} \sim 10^{-8} (k_2/10^{-3}) (\Delta\Omega_c/\Omega)$. For the asynchronous rotation of the companion star to the redback PSR J1723–2837, van Staden and Antoniadis [261] found $(\Delta\Omega_c/\Omega) \approx 3 \times 10^{-3}$. Adopting a similar asynchronicity here results in an expected $\dot{P}_{\text{orb}} \sim 3 \times 10^{-11}$, of similar magnitude to that observed from J2039.

Unfortunately, we are unable to measure an asynchronicity of this magnitude with existing data. Assuming $\Delta\Omega_c/\Omega = 3 \times 10^{-3}$, the star will undergo one extra rotation with respect to the co-rotating orbital frame every 75 days, therefore completing several cycles between each of our light curves. Since the numbers of such cycles in between our observations are unknown and unmeasurable, we cannot unambiguously measure the asynchronous rotation by comparing the spot longitudes in each of our light curves. Indeed, we cannot even be sure that the same spot is

present on each of our observing epochs. The asynchronous rotation also cannot be measured from the orbital period variations seen in Section 5.3.2, as these effects are caused by variations in the internal structure, and can occur even in the case of a tidally-locked companion star.

Thus, a tempting (although highly speculative!) picture emerges, in line with that proposed by van Staden and Antoniadis [261], where many of the variable phenomena seen from J2039 are due to magnetic activity and asynchronous rotation in the companion star. In this picture, the magnetic activity leads to large star spots, explaining the asymmetry in the optical light curve, and to quadrupole moment variations in the stellar envelope, explaining the short-term orbital period variations. Asynchronous rotation of the spotted surface then leads to the observed optical variability, and introduces a tidal force that is responsible for the putative long-term increase in the orbital period. A large stellar magnetic field would also be consistent with the synchrotron explanation for the orbital gamma-ray modulation described in Section 5.3.3.1. Investigating this picture in the future will require several more years of timing measurements to confirm or refute the long-term \dot{P}_{orb} , and high-cadence (e.g. monthly) optical monitoring to test for evidence for asynchronous rotation in the form of a periodicity in the optical variability.

5.5.5 Prospects for binary gamma-ray pulsar searches

Over the course of the *Fermi*-LAT mission, a number of candidate redback systems, similar to J2039, have been discovered within unidentified LAT sources [e.g., 55, 88–93]. Our detection of gamma-ray pulsations from J2039 shows that, with sufficiently precise orbital constraints, gamma-ray pulsation searches are a viable method to confirm their redback natures.

However, the orbital period variations common in redbacks, and present here for J2039, will make them more difficult to detect. Indeed, J2039 is the first spider MSP exhibiting rapid orbital period variations to have been discovered in a gamma-ray search. Due to the low photon flux from a typical pulsar, multiple years of LAT gamma-ray data are required for a discovery in a directed search. On such timescales the pulsar’s ascending node can shift back and forth by more than 15s for some redback pulsars [see e.g., 71, 100]. In a pulsation search, we are forced to assume a constant orbital period. The fact that orbital period variations are common in redbacks therefore has two major implications, which we illustrate in Figure 5.11.

Firstly, the recovered signal-to-noise ratio drops significantly. As shown in Figure 5.11, the maximum signal strength found in the search for J2039 was 66% smaller than the signal strength obtained after accounting for the orbital period variations in our timing analysis. The reason for this is visible from the left panel of Figure 5.1, where it can be seen that the signal becomes clearer as the offset in T_{asc} decreases, while at epochs where the offset is largest ($\Delta T_{\text{asc}} \sim \pm 10\text{s}$), the signal disappears entirely. Despite this, J2039 was still easily detected above the statistical noise level, but for fainter pulsars in future searches this could reduce the signal strength below detectable levels.

Secondly, the signal is spread over a larger parameter volume compared to a signal from a pulsar with a constant orbital period. This could actually be beneficial to future searches: assuming the signal is strong enough to remain detectable over




Figure 5.11: Fraction of maximum signal power recovered as a function of offsets in P_{orb} and T_{asc} from the timing solution in Table 5.1. The origin is the point in parameter space giving the highest signal power for J2039 using a model for orbital motion with constant period. Red contour lines show the expected fraction of signal power recovered according to the metric approximation used to construct the search grid, which assumes that the signal has a constant orbital period. The effect of the orbital period variations is to reduce the maximum signal power, and to spread it over a larger region of the parameter space. The black contour lines show the actual recovered signal power as a function of P_{orb} and T_{asc} . Blue crosses show parameter space locations at which a significant signal was detected in our search. Note that the 90% ellipse was used for the grid generation in the search described in this work. A search grid designed to take into account the smearing effect of the orbital period variations could feasibly have been several times sparser in these parameters without missing the signal.

small portions of the LAT data, the orbital period variations may actually allow pulsations to be detected over a larger range in P_{orb} and T_{asc} , as can be seen in Figure 5.11. This could be exploited to reduce the computing cost of future searches by using less dense grids over the orbital parameter space. Another option could be to search the results for a clustering that indicates a wider-than-expected spread of a signal. We intend to investigate both options.

5.6 Conclusions

Using a directed search for gamma-ray pulsations running on the distributed volunteer computing system *Einstein@Home*, we have confirmed the reback nature of the candidate binary system within 4FGL J2039.5–5617. This is the first reback pulsar to be discovered through its gamma-ray pulsations, providing hope that a number of similar reback candidates identified in *Fermi*-LAT sources might be confirmed in this way in the near future, even though their orbital periods display large variability.

Optical observations of variations in the orbital light curve, and gamma-ray

timing observations of its changing orbital period, add another example to a growing body of evidence that redback companions have activity on super-orbital timescales. A better understanding of variable phenomena in redback companions is required both to ensure that the properties inferred from optical light-curve modelling (e.g. inclination angles, pulsar and companion mass estimates) are reliable, and to better understand their evolution.

The origin of the light curve variability remains unclear, but requires temperature variations of a few hundred K over a reasonably large fraction of the visible surface of the star. We speculate on a few possible origins for these temperature variations, including reprocessing of the pulsar’s heating flux in a variable intra-binary shock, variable convective flows on the stellar surface, or magnetic star spot activity. The latter picture fits well with the interpretation of orbital period variations being caused by quadrupole moment variations driven by magnetic activity in the companion star.

To quantify the orbital period variations, we have developed a new gamma-ray pulsation timing method that treats the orbital phase as a stochastically varying function, and provides statistical estimates of the amplitude and characteristic timescale of the variability. We find that the magnitude of the orbital period variations requires only a small fractional change (a few percent) in the stellar quadrupole moment, suggesting that this is indeed a plausible scenario. However, due to the sparsity of optical observations, we are so far unable to probe correlations between the optical light curve variability and changes in the orbital period. Based on these phenomena, we are pursuing long-term monitoring of redback companions to reveal whether or not optical variability is correlated with quadrupole moment variations. Future light curves, ideally obtained with a denser (e.g. monthly) observing cadence, may also allow us to track variable light curve components as they evolve over time, providing a probe for possible asynchronous rotation.

We modelled the optical light curves, using the new timing measurement of the projected semi-major axis of the pulsar’s orbit to constrain the binary mass function. Although our modelling is complicated and likely biased by the unexplained variability and light-curve asymmetry, the gamma-ray data significantly rule out any substantial eclipse and set a maximum inclination of $i \lesssim 78^\circ$, and we find that an inclination of $i \sim 75^\circ$ provides a consistent fit to all light curves with an inclination as low as $i \sim 60^\circ$ being consistent with one single-epoch light curve. This implies a fairly low pulsar mass $1.1M_\odot < M_{\text{psr}} < 1.6M_\odot$, and companion mass $0.15M_\odot < M_c < 0.22M_\odot$. Additional light curves may help to reduce the inclination and mass uncertainties by “marginalizing out” the variability, although these will remain highly model-dependent and subject to systematic biases without an independent method to validate with. The joint radio and gamma-ray pulse profile modelling described in Paper II is one such method, and broadly agrees with our results here.

We also find that an orbitally modulated component to the gamma-ray flux is in fact pulsed emission in phase with the magnetospheric gamma-ray pulses, rather than being an additional unpulsed component. We speculate that this could be due to inverse Compton scattering or synchrotron radiation from the high-energy pulsar wind. This is the second redback from which such an effect has been detected, and this may prove to be a valuable probe of the pulsars’ high-energy winds in future studies.

Acknowledgments

We are very grateful to the thousands of volunteers who donated computing time to *Einstein@Home*, and to those whose computers first detected PSR J2039–5617: J. Bencin of Cleveland, OH, USA; and an anonymous volunteer whose username is “Peter”.

C.J.C. would like to thank Steven G. Parsons for guidance and useful discussions while observing with ULTRACAM.

C.J.C., G.V., R.P.B., M.R.K., D.M.-S. and J.S., acknowledge support from the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 715051; Spiders). V.S.D. and ULTRACAM acknowledge the support of the STFC. Z.W. is supported by the NASA postdoctoral program. *Einstein@Home* is supported by NSF grants 1104902 and 1816904.

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. This work performed in part under DOE Contract DE-AC02-76SF00515

Based on observations collected at the European Southern Observatory under ESO programme 0101.D-0925(B).

This paper uses data taken from the XMM-Newton Science Archive (Observation ID: 0720750301) and produced using the Pipeline Processing System. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement. This research was made possible through the use of the AAVSO Photometric All-Sky Survey (APASS), funded by the Robert Martin Ayers Sciences Fund and NSF AST-1412587.

Data Availability

The *Fermi*-LAT data are available from the Fermi Science Support Center (<http://fermi.gsfc.nasa.gov/ssc>). The *XMM-Newton* data are available from the XMM-Newton Science Archive (<http://nxsas.esac.esa.int>). The ULTRACAM data are available in Zenodo (<https://doi.org/10.5281/zenodo.3964204>). The GROND

data are available from the ESO Science Archive Facility (<http://archive.eso.org/>).

Gamma-ray Follow-up Detection and Timing of PSR J1555–2908

*This chapter is a contribution to Ray, P. S., Nieder, L., Clark, C. J., et al.
The paper is in draft stage and is intended to be submitted to ApJ soon.*

6.1 Introduction

This section gives the context for Section 6.2, which will be the author’s contribution towards a multiwavelength paper, concerning the discovery and study of the radio and gamma-ray black widow pulsar PSR J1555–2908.

The pulsar was detected first in radio observations at 820 MHz by the Green Bank Telescope (GBT). The search was part of a larger survey of the Pulsar Search Consortium (PSC) [56], targeting pulsar candidates identified by Frail et al. [166]. The discovery revealed a 559 Hz binary millisecond pulsar (MSP) in a 5.6 hr orbit with an $0.05 M_{\odot}$ companion. The low mass and the observed radio eclipses suggest a black-widow system. Subsequently, PSR J1555–2908 was detected in a gamma-ray follow-up search, enabling a study of this pulsar over the 12 years of *Fermi*-LAT data.

In the paper, the pulsar and its companion will be discussed based on multiwavelength observations. Apart from the radio and gamma-ray analyses, some optical observations and an X-ray pulsation search will be presented. Furthermore, two companion papers are planned: one about a dedicated optical analysis of the companion and one about a gamma-ray timing analysis with an alternative phase model (Chapter 7).

6.2 Gamma-ray Observations

The initial (not phase-connected) radio timing parameters provided the seed necessary to make a computationally-tractable pulsation search in the gamma-ray data, as we describe here.

6.2.1 *Fermi*-LAT Data Preparation

To search for gamma-ray pulsations from PSR J1555–2908, we selected SOURCE-class gamma-ray photons detected by the *Fermi* LAT between 2008 August 03 and 2018 April 19. Photons were included from within a 5 deg region of interest (RoI) around the radio position, with energies greater than 100 MeV, and with a maximum zenith angle of 90 deg, according to the “Pass 8” P8R2_SOURCE_V6 [173] instrument response functions (IRFs)¹.

To increase the sensitivity of the pulsation search, and avoid the need for hard cuts on photon energies and incidence angles, we weighted the contributions of each photon to the pulsation detection statistic [114]. The weights represent the probability of each photon having been emitted by the targeted gamma-ray source, rather than by a nearby point source, or by the diffuse Galactic or isotropic gamma-ray components. The weights were computed with `gtsrcprob`, using the preliminary “FL8Y” version of the *Fermi*-LAT Fourth Source Catalog, and corresponding diffuse and isotropic emission templates, as the input model for the gamma-ray flux within the RoI.

6.2.2 Gamma-ray Pulsation Search and Detection

We performed a gamma-ray pulsation search around the parameters from a preliminary radio timing solution based on data spanning 50 days. A search was necessary as the parameters in the preliminary radio timing solution were not measured precisely enough to safely extrapolate over multiple years of LAT data, and the pulsar’s gamma-ray photon flux is too weak to show significant pulsations over such a short time span. For several parameters (spin frequency, orbital period) the resolution scales inversely with the length of the data span. At this stage the spin frequency derivative had not been measured.

Apart from the radio timing solution, the search parameter space was constrained by a *Gaia* DR2 sky position [122, 269] and the distribution of spin frequency derivatives for known MSPs. It was assumed that a *Gaia* source with magnitude $G = 20.46 \pm 0.36$ [122], coincident with the radio pulsar’s sky position, is the pulsar’s companion star. The spin frequency derivative was searched in the range $\dot{f} \in [-2 \times 10^{-14}, 0] \text{ Hz s}^{-1}$ as more than 95% of the known MSPs fall into this range [ATNF Pulsar Catalogue²; 110]. In the case of a non-detection, this range would have been extended in steps.

The sensitive H statistic [114, 119] was utilized to search for gamma-ray pulsations. This statistic incoherently combines the Fourier power of the lowest M harmonics. Typically most power is found in the lowest 5 harmonics [80]. Hence for computational reasons, in the search we truncated the harmonic summing at $M = 5$ as successfully done before in the follow-up search of PSR J0952–0607 (Chapter 3).

Building an efficient and dense grid covering the parameter space was key to the detection of gamma-ray pulsations. To build such a grid we exploited the distance “metric”, which is a second order approximation of the expected fractional loss in squared signal-to-noise ratio due to offsets in the signal parameters [117, 118]. The metric components for the binary pulsar parameters are presented in Chapter 2.

¹See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_essentials.html

²<http://www.atnf.csiro.au/research/pulsar/psrcat>

Denser grids are required for higher harmonics and thus the grid is build for the highest harmonic, i.e. $M = 5$.

The search space was split into smaller parts and carried out in parallel on the ATLAS computing cluster in Hannover [84]. On one single computer the search would have taken ~ 70 days. Distributing the work over 7,170 CPU cores the search only took ~ 15 minutes.

Significant gamma-ray pulsations were detected over most of the 10-year LAT-data span in this search. The rotational phases assigned to the detected photons showed pulse arrival times varying slowly over the observation time span indicating that additional spin-frequency derivatives might be required to model the pulsar more accurately. Still, the maximum H statistic detected in the search was $H_m = 276.3$. Conservatively assuming that all 6×10^{11} trials were independent, the false-alarm probability is $P_{\text{FA}} = 9.6 \times 10^{-37}$, which confirms the detection of gamma-ray pulsations.

6.2.3 Gamma-ray Pulsation Timing

Following the detection of gamma-ray pulsations, we extended the observation span to cover the most recent *Fermi*-LAT data, up to 2020 August 05, and using the more recent P8R3_SOURCE_V2 IRFs [235]. In the *Fermi*-LAT Fourth Source Catalog, the gamma-ray spectrum of 4FGL J1555.7–2908 is modelled with a simple power-law, rather than the curved sub-exponentially-cutoff power-law spectrum typical for gamma-ray pulsars. This is because the curved spectrum did not provide a significantly better fit for the observed gamma-ray flux in 4FGL. This is likely due to a combination of a low overall photon flux, and high uncertainties in the low-energy flux due to the contribution from the diffuse Galactic interstellar emission. Bruel [115] developed a method to obtain optimized photon weights by adjusting the spectral parameters to maximize the resulting weighted H -test. Adopting this technique to optimize the photon probability weights, we found a pulsar-like sub-exponentially-cutoff power-law spectrum that resulted in a much more significant pulsation detection ($H = 860.5$) than was obtained using the simple power-law spectral model. We therefore adopted these weights for the follow-up timing analyses presented below.

In the timing analysis the pulsar is analyzed precisely using the likelihood, \mathcal{L} , and the Bayesian Information Criterion [BIC; 184]. To measure the pulsar parameters, \mathcal{L} is maximized by fitting a pulse profile to a template pulse profile [see, e.g., 52, 79, 97] by marginalizing over the pulsar parameters and the template parameters jointly as described in Chapter 3. The template is a sum of Gaussian peaks. For the width, we used a log-uniform prior and constrained the range to allow only peaks broader than 1% of a rotation and narrower than half a rotation. All other parameters used a uniform prior. The analysis is also performed with more parameters than searched, but these are only accepted if they lead to a decrease of the BIC.

To account for small phase variations over the full data span, additional spin-frequency derivatives are needed. While only the first derivative was included in the search, three additional derivatives were favored by the BIC throughout the timing analysis. Additional higher-order derivatives were disfavored by the BIC.

It was also tested whether it is preferable to fix some parameters to the values obtained with radio timing. Subsequently, the parameters x , t_{asc} , ϵ_1 , and ϵ_2 were

Table 6.1: Timing solution for PSR J1555–2908.

Parameter	Value
Range of observational data (MJD)	54681 – 59066
Reference epoch (MJD)	57800.0
Timing parameters	
R.A., α (J2000.0)	15 ^h 55 ^m 40 ^s .6586(10)
Decl., δ (J2000.0)	–29°08′28″.426(13)
Spin frequency, f (Hz)	559.44000642613(6)
1st spin-frequency derivative, \dot{f} (Hz s ^{–1})	–1.39430(1) × 10 ^{–14}
2nd spin-frequency derivative, \ddot{f} (Hz s ^{–2})	–1.0(3) × 10 ^{–25}
3rd spin-frequency derivative, \dddot{f} (Hz s ^{–3})	7(1) × 10 ^{–33}
4th spin-frequency derivative, $\overset{\dots}{f}$ (Hz s ^{–4})	8(1) × 10 ^{–41}
Proj. semimajor axis ^a , x (s)	0.1514468(1)
Orbital period, P_{orb} (days)	0.23350026854(11)
Epoch of ascending node ^a , t_{asc} (MJD)	57785.53936387(3)
1st Laplace-Lagrange parameter ^a , ϵ_1	2(2) × 10 ^{–6}
2nd Laplace-Lagrange parameter ^a , ϵ_2	–4(1) × 10 ^{–6}
Derived parameters ^b	
Spin period, P (ms)	1.788
Spin-period derivative, \dot{P} (s s ^{–1})	4.5 × 10 ^{–20}
Spin-down power, \dot{E} (erg s ^{–1})	3.1 × 10 ³⁵
Surface B -field, B_{surf} (G)	2.9 × 10 ⁸
Light-cylinder B -field, B_{LC} (G)	4.6 × 10 ⁵
Characteristic age, τ_c (Gyr)	0.6

Notes. — Numbers in parentheses are statistical 1σ uncertainties on the final digits. The JPL DE405 solar system ephemeris has been used, and times refer to TDB.

^a Parameter fixed to radio timing solution.

^a Not corrected for Shklovskii and Galactic acceleration effects due to highly uncertain distance measurement. The non-detection of proper motion suggests that these estimates should be accurate to a few percent.

fixed. This is not because the timing would result in a higher likelihood \mathcal{L} , which it does not in this case, but because the BIC prefers it for the lower number of free parameters.

Our full timing solution is shown in Table 6.1. The gamma-ray pulse profile and a phase-time diagram are plotted alongside the radio profile in Figure 6.1.

We tested for the presence of additional effects, including proper motion. These parameters were found to be consistent with zero and disfavored by the BIC. For those parameters, the timing analysis sets 95% confidence upper limits. The timing analysis clearly favors zero total proper motion $\mu_t = \sqrt{\mu_\alpha^2 \cos^2 \delta + \mu_\delta^2}$, setting the upper limit $\mu_t^{95\%} \leq 6.0 \text{ mas yr}^{-1}$. The 95% confidence interval on a variable orbital period is $-6.5 \times 10^{-13} \text{ s s}^{-1} < \dot{P}_{\text{orb}} < 6.1 \times 10^{-13} \text{ s s}^{-1}$.




Figure 6.1: Integrated pulse profile and phase-time diagram of PSR J1555–2908, showing two identical rotations for clarity. Top: The orange curve indicates the template with the highest BIC. The transparent black curves illustrate 100 representative templates randomly selected from the Monte Carlo samples after the chain stabilized. The histogram shows the weighted photon counts with 50 bins per rotation. The dashed blue line shows the estimated background level. Bottom: Each point represents the rotational phase of a detected gamma-ray photon and the gray scale indicates the probability weight.

6.3 Discussion

The measured spin parameters of PSR J1555–2908 are remarkable compared to other pulsars in the ATNF Pulsar Catalogue³ [110]. Its 559 Hz spin frequency is one of the fastest, with only twelve pulsars rotating more rapidly. The observed spin-period derivative, and the inferred spin-down luminosity, are among the ten largest within the known MSP population.

The observation of higher-order nonzero spin-frequency derivatives is extraordinary. Only a few MSPs have a nonzero second derivative, and even higher orders are unusual. The timing residuals for excluding \dot{f} , \ddot{f} , and \dddot{f} would be larger than $100 \mu\text{s}$. While this would be attributed to “timing noise” in young pulsars, the timing noise in MSPs typically has amplitudes of order $1 \mu\text{s}$ [97]. Hence, PSR J1555–2908 is either the first MSP to display timing noise this large, or a different explanation is required. Such a possible explanation is presented in Chapter 7, describing the available gamma-ray pulsations equally good as the higher-order spin-frequency model used here.

³<http://www.atnf.csiro.au/research/pulsar/psrcat>

Does the Black Widow Pulsar PSR J1555–2908 have an additional Planetary-mass Companion?

This chapter is an adaptation of a paper in draft stage intended to be submitted soon to ApJL as Nieder, L., Clark, C. J., Ray, P. S., et al.

Abstract

As presented in Chapter 6, the 559 Hz black-widow pulsar PSR J1555–2908, originally discovered in radio, is also visible in *Fermi*-Large Area Telescope gamma-ray data. The pulsar shows long-term variations in its spin frequency via timing analysis that are much larger than is observed from other millisecond pulsars. This could be due to intrinsic variability and/or timing noise. Here, we consider an alternative explanation: the variations arise from the presence of a light-weight third object in a wide multi-year orbit around the neutron star and its low-mass companion. With current data, and without increasing the number of free parameters, this hierarchical-triple-system model describes the pulsar’s rotation as accurately as the timing-noise model. Future observations will show if this alternative explanation is correct.

7.1 Introduction

PSR J1555–2908 is a neutron star spinning rapidly at 559 Hz, in a tight-orbit binary system with a low-mass companion. This millisecond pulsar (MSP) was first detected in a Green Bank Telescope (GBT) pulsar survey (P. S. Ray et al. 2020, in prep., hereafter Chapter 6). The searches targeted steep-spectrum radio sources [166] within the localization region of *Fermi*-LAT gamma-ray sources. After the radio detection, the pulsar was found in gamma rays, allowing the timing measurement of the system parameters over the 12-year *Fermi* mission time span.

Timing analysis using the multi-year LAT gamma-ray data reveals variations of the spin frequency that are larger than is typical for MSPs [97]. Such variations are often seen in young gamma-ray pulsars, where they are labeled as “timing noise”. However, this is rare in MSPs, for which the intrinsic rotational phase is generally well described by a quadratic function of time. In contrast, for PSR J1555–2908 three additional terms in the rotational phase model are required (making it a quintic function of time).

In this paper, we discuss an alternative explanation: that the variations arise from the presence of a third body in the system. This additional object is in a wide, multi-year orbit around the closely orbiting neutron star and its low-mass companion. This would make it similar to the hierarchical-triple-system pulsar PSR J0337+1715 [306].

With the currently available data, this model describes the pulsar as well as the timing-noise model, but provides a simple and clear physical explanation for the frequency variations.

7.2 Rotational phase model

To precisely track the rotational phase, the photon arrival times need to be corrected for the line-of-sight motion of the pulsar, if it is in a on orbit around one or more companions. For a simple hierarchical triple system (HTS), we assume that the gravitational interaction between the two companions can be neglected, and the third body orbits the center of mass of the pulsar and its close companion. The photon arrival times at the barycenter of the triple system, t_{tb} , can be expressed as a function of the photon’s emission time, t_{em} , as

$$t_{\text{tb}} = t_{\text{em}} + x_{\text{p,I}}\Delta_{\text{R,I}}(\Omega_{\text{b,I}}t_{\text{em}}) + x_{\text{p,O}}\Delta_{\text{R,O}}(\Omega_{\text{b,O}}t_{\text{em}}). \quad (7.1)$$

Here, $x_{\text{p,I}}$ and $x_{\text{p,O}}$ are the times light needs to travel the radius of the pulsar’s orbit around the inner (I) and the outer (O) companion projected onto the line of sight. $\Delta_{\text{R,I}}$ and $\Delta_{\text{R,O}}$ are dimensionless functions that describe the respective modulations depending on the orbital phases, $\Omega_{\text{b,I}}t_{\text{em}}$ and $\Omega_{\text{b,O}}t_{\text{em}}$.

As shown in [96], we can make $x_{\text{p,I}}\Delta_{\text{R,I}}(\Omega_{\text{b,I}}t_{\text{em}}) + x_{\text{p,O}}\Delta_{\text{R,O}}(\Omega_{\text{b,O}}t_{\text{em}})$ a function of t_{tb} by Taylor expansion. To second order (i.e. including terms of order $(x_{\text{p,I}}\Omega_{\text{b,I}})^2$,

$(x_{p,O}\Omega_{b,O})^2$, $x_{p,I}x_{p,O}\Omega_{b,I}^2$, and $x_{p,I}x_{p,O}\Omega_{b,O}^2$), we find

$$\begin{aligned}
& x_{p,I}\Delta_{R,I}(\Omega_{b,I}t_{em}) + x_{p,O}\Delta_{R,O}(\Omega_{b,O}t_{em}) \\
&= (x_{p,I}\Delta_{R,I}(\Omega_{b,I}t_{tb}) + x_{p,O}\Delta_{R,O}(\Omega_{b,O}t_{tb})) \\
&\quad \times \left[1 - (x_{p,I}\Omega_{b,I}\Delta'_{R,I}(\Omega_{b,I}t_{tb}) + x_{p,O}\Omega_{b,O}\Delta'_{R,O}(\Omega_{b,O}t_{tb})) \right. \\
&\quad \quad + (x_{p,I}\Omega_{b,I}\Delta'_{R,I}(\Omega_{b,I}t_{tb}) + x_{p,O}\Omega_{b,O}\Delta'_{R,O}(\Omega_{b,O}t_{tb}))^2 \\
&\quad \quad + \frac{1}{2} (x_{p,I}\Omega_{b,I}^2\Delta''_{R,I}(\Omega_{b,I}t_{tb}) + x_{p,O}\Omega_{b,O}^2\Delta''_{R,O}(\Omega_{b,O}t_{tb})) \\
&\quad \quad \left. \times (x_{p,I}\Delta_{R,I}(\Omega_{b,I}t_{tb}) + x_{p,O}\Delta_{R,O}(\Omega_{b,O}t_{tb})) \right],
\end{aligned} \tag{7.2}$$

where Δ'_R and Δ''_R denote the first and second derivative of Δ_R with respect to $\Omega_b t$.

In the case of a circular orbit, Δ_R takes the simple form of a sinusoid, $\Delta_R(t) = \sin(\Omega_b(t - T_{asc}))$. Hence, $\Delta'_R = \cos(\Omega_b(t - T_{asc}))$ and $\Delta''_R = -\sin(\Omega_b(t - T_{asc}))$. In Section 7.3, we show that PSR J1555–2908 can be well described as an HTS with both orbits being circular. However, Eq. (7.2) can also be applied to the widely used small-eccentricity-orbit “ELL1 model” [133], or the larger-eccentricity models presented in Chapter 2.

For a timing analysis, it is helpful to have a good starting point in the multi-dimensional parameter space. The parameters describing the potential outer orbit can be roughly estimated from the timing solution presented in Chapter 6, which requires four frequency derivatives (“FX model”).

First, we need to derive how the observed frequency changes in the case of an outer orbit. Assuming only a non-zero spin frequency f and a non-zero first time-derivative \dot{f} , the additional evolution of the spin frequency over time can be written in two ways:

$$\begin{aligned}
& f(t) - f^* - \dot{f}^*(t - t_0) \\
&= \frac{1}{2}\ddot{f}^*(t - t_0)^2 + \frac{1}{6}\ddot{\dot{f}}^*(t - t_0)^3 + \frac{1}{24}\ddot{\dot{\dot{f}}}^*(t - t_0)^4
\end{aligned} \tag{7.3}$$

$$= \Delta f^* + \Delta \dot{f}^*(t - t_0) - (f^* + \Delta f^*)x_{p,O}\Omega_{b,O} \cos(\Omega_{b,O}(t - T_{asc,O})). \tag{7.4}$$

Here, we denote the frequency parameters measured with the FX model with an asterisk.

Second, since all parameters in Eq. (7.3) are measured, we fit Eq. (7.4) to Eq. (7.3) over the valid time span of the pulsar timing solution. To do this, we used `curve_fit` from `scipy`. This gives initial values for the five free parameters $\{\Delta f^*, \Delta \dot{f}^*, x_{p,O}, \Omega_{b,O}, T_{asc,O}\}$ from Eq. (7.4).

Encouragingly, the fitted values are of reasonable sizes. The pulsar’s potential movement would have a radius of a few tens of kilometers with a orbital period of ~ 10 yr, indicating a low-mass companion in a wide orbit. The changes in f would be very small, and \dot{f} would change by $\sim 20\%$.

7.3 Gamma-ray timing analysis

In the timing analysis here, we use the LAT data prepared for the timing in Section 6.2.3 of Chapter 6. These data included SOURCE-class photons detected by LAT

Table 7.1: Timing solutions for PSR J1555–2908 with FX and HTS models.

Parameter	FX model	HTS model
Range of observational data (MJD)	54681 – 59066	
Reference epoch (MJD)	57800.0	
Timing parameters		
R.A., α (J2000.0)	15 ^h 55 ^m 40 ^s .6586(10)	15 ^h 55 ^m 40 ^s .6587(10)
Decl., δ (J2000.0)	–29°08′28″.424(13)	–29°08′28″.421(12)
Spin frequency, f (Hz)	559.44000642611(5)	559.44000642603(12)
1st spin-frequency derivative, \dot{f} (Hz s ^{–1})	$-1.39431(12) \times 10^{-14}$	$-1.39194(7) \times 10^{-14}$
2nd spin-frequency derivative, \ddot{f} (Hz s ^{–2})	$-1.0(3) \times 10^{-25}$	
3rd spin-frequency derivative, \dddot{f} (Hz s ^{–3})	$7(1) \times 10^{-33}$	
4th spin-frequency derivative, $\ddot{\ddot{f}}$ (Hz s ^{–4})	$8(2) \times 10^{-41}$	
Inner-orbit binary parameters		
Proj. semimajor axis, $x_{p,I}$ (s)	0.151445(3)	0.151446(3)
Orbital period, $P_{b,I}$ (days)	0.23350026826(13)	0.23350026823(13)
Epoch of ascending node, $T_{asc,I}$ (MJD)	57785.5393614(9)	57785.5393608(9)
Outer-orbit binary parameters		
Proj. semimajor axis, $x_{p,O}$ (s)		0.000124(19)
Orbital period, $P_{b,O}$ (days)		$3.9(3) \times 10^3$
Epoch of ascending node, $T_{asc,O}$ (MJD)		54860(170)

Notes. — Numbers in parentheses are statistical 1σ uncertainties on the final digits. The JPL DE405 solar system ephemeris has been used, and times refer to TDB.

between 2008 August 03 and 2020 August 05 within a 5 deg region of interest (RoI) around the pulsar position, with energies greater than 100 MeV, and with a maximum zenith angle of 90 deg, according to the recent P8R3_SOURCE_V2 instrument response functions (IRFs)¹ [235].

The timing analysis is done twice, very similarly to that in Chapter 6. Firstly, we redid the timing analysis with the FX model, varying over all parameters (i.e. not fixing some parameters to the radio solution). Secondly, we amended the timing code with the HTS model given in Equations (7.1) and (7.2). To judge the timing results, we used the H statistic [114, 119], the likelihood \mathcal{L} [79], and the Bayesian Information Criterion (BIC) [184].

The highest test statistics, H and \mathcal{L} , found with both models are nearly identical, only marginally favoring the HTS model. Since both models require the same number of parameters, three additional spin-frequency derivatives for the FX model and three additional orbital parameters for the HTS model, the BIC can also not discriminate. Our final timing solutions over 12 years of LAT data using the FX model and the HTS model are shown in Table 7.1.

7.4 Discussion

The PSR J1555–2908 displays spin-frequency variations (Chapter 6) that would be unusually large in amplitude for timing noise of an MSP [97]. Here, we have pre-

¹See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_essentials.html




Figure 7.1: Mass-mass plot showing the potential outer companion mass for different inclination angles (measured with respect to line-of-sight to Earth) and total masses of the inner binary system. For comparison, the masses of Moon, Mercury, and Mars are indicated by the dashed lines.

sented an alternative explanation, assuming another companion orbiting the inner binary system. To account for the additional body in the system, we have developed a simple, but sufficient rotational phase model. In a timing analysis with current data, the results with both models are indistinguishable. Within the timing range the cumulative phase offset between both models is $< 1\%$ of one rotation.

It is predictable that the two models would not differ much. The variations are measurable with the LAT data, but the amplitude of the cumulative phase offset is still only 14% of one rotation over the course of 5.5 yr. Furthermore, the potential outer companion would have completed roughly one orbit during the *Fermi* mission which can be well approximated with the five spin parameters.

The binary mass function may be used to estimate the mass of the outer companion (see Fig. 7.1). For common pulsar masses $1.4 M_\odot < M_p < 2.0 M_\odot$, a inner companion mass $M_{c,I} = 0.05 M_\odot$ (Chapter 6), and inclination angles $20^\circ - 90^\circ$, the outer companion would have roughly a Mercury-like mass $\sim (0.5 - 2.5) \times 10^{-7} M_\odot$.

Other HTSs which include a pulsar are known. PSR J0337+1715 is in an orbit with two white dwarf companions [306]. Carefully studying this system led to one of the most stringent tests of General Relativity’s predicted universality of free fall [307]. A binary system consisting of the MSP PSR B1620–20 and a white dwarf is orbited by a planetary companion of $\sim 2.5 M_{\text{Jupiter}} (= 2.4 \times 10^{-3} M_\odot)$ [308, 309].

Apart from the two tested explanations, more are possible. For example, the timing variations of the MSP PSR B1937+21 could be explained with an asteroid belt surrounding the pulsar [310].

With current data, both investigated models track the pulsar’s rotation equally well. However, the ongoing *Fermi* mission will continue to collect data, and eventually the nature of this system will be resolved. Only to emphasize the power of pulsar timing with LAT data, we want to note that if PSR J1555–2908 has an additional planetary-mass companion, then it was discovered because over the course

of 100 billion pulsar rotations, 14% of one single rotation are missing.

Acknowledgments

This work was supported by the Max-Planck-Gesellschaft (MPG).

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'Énergie 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. This work performed in part under DOE Contract DE-AC02-76SF00515.

Conclusion and Outlook

The launch of the *Fermi Gamma-ray Space Telescope* in 2008 has led to a large increase in the numbers of known MSPs [see, e.g., 23, 43] and spider pulsars [see, e.g., 53, 56]. Most of these were discovered by radio telescopes pointing at *Fermi*-Large Area Telescope (LAT) gamma-ray sources with pulsar-like properties¹.

Observing spider pulsars in gamma rays has two major advantages. Firstly, the gamma rays are essentially unaffected by the material ablated from the pulsar's companion. In comparison, radio observations are often eclipsed or scattered for large fractions of the binary orbit in such systems [see, e.g., 19, 49, 100, 178, 179, 258]. Secondly, the ongoing all-sky survey of the *Fermi*-LAT has been providing gamma-ray data since its launch in 2008 for any point in the sky. This allows measurement of the variations of the orbital period over time that are often seen in spider systems, especially those containing redback pulsars [4, 71]. On the other hand, a spider pulsar discovered in radio would need to be observed for several years with high cadence to track these variations [98, 100].

However, there is one downside: The gamma-ray data are very sparse. For most pulsars the LAT detects fewer than one photon per million rotations. To significantly detect gamma-ray pulsations, months and years of data need to be analyzed, which results in computationally extremely demanding searches if the pulsar parameters are not known in advance (e.g. from radio). Prior to this thesis, only the spider pulsar PSR J1311–3430 was discovered by directly searching the gamma-ray data [44].

Within this thesis, new search methods for binary gamma-ray pulsars were developed. Using these, new spider gamma-ray pulsars were discovered in computationally demanding, partially informed searches and in follow-up searches of radio pulsars with only preliminary ephemerides.

The binary search methods presented in Chapter 2 can be seen as an extension and generalization of previous thesis concerning isolated pulsar searches by Pletsch and Clark [80]. For the application of partially informed searches, a pipeline has been set up running the most computing intensive stages on the volunteer computing project *Einstein@Home*, and the last, most sensitive stage on the ATLAS computing

¹<https://tinyurl.com/fermipulsars>

cluster at AEI Hannover [84]. For the application of a radio follow-up search, a pipeline was set up that can quickly perform a search, exploiting the parameter constraints from preliminary radio ephemerides.

The detection of gamma-ray pulsations from the black widow PSR J1653–0158 by *Einstein@Home* finally presents the long expected solution to an unidentified source that has been known for more than two decades. In Chapter 4, the discovery and gamma-ray timing are presented along with an optical analysis of the companion, as well as searches for radio pulsations and continuous gravitational waves. The pulsar shows one of the lowest spin-period derivatives and inferred surface magnetic field strengths. Its mass possibly exceeds $2 M_{\odot}$, and despite numerous sensitive searches at various observing frequencies the pulsar remains undetected in radio. The companion has a high average density of 73 g cm^{-3} , and might be a helium white dwarf remnant. This suggests an even denser, degenerate interior with density decreasing with radial distance.

The discovery of the redback pulsar PSR J2039–5617 via a partially informed gamma-ray search is a significant step forward. This is presented in Chapter 5 and shows that redbacks can be discovered through their gamma-ray pulsations despite their often large variabilities of the orbital period. In Section 5.5.5, we found that these orbital-period variations lower the signal-to-noise ratio, while the signal is smeared out over a larger volume in parameter space. In Chapter 5, PSR J2039–5617 and its companion are studied in detail using optical and gamma-ray observations, while the subsequent detection of radio pulsations is presented in a companion paper [105].

PSR J0952–0607 rotates extremely rapidly at 707 Hz, has first been discovered through low-frequency radio observations with the Low-Frequency Array (LOFAR), and its gamma-ray pulsations were revealed in a follow-up search on ATLAS. The long-term gamma-ray observations with the *Fermi*-LAT allowed measuring the pulsar parameters precisely and inferring one of lowest surface magnetic field strengths ever seen for a pulsar. The gamma-ray and optical study, a radio update, and a continuous gravitational wave search are presented in Chapter 3. Due to its faintness, some puzzles about this source remain. The distances inferred from the dispersion measure and the optical-modelling distance are incompatible, which either suggests an extremely dense companion or an overestimated electron density in the direction of PSR J0952–0607. The non-detection of gamma-ray pulsations before MJD 57,500 has been investigated, but remains unsolved. Additional data in the future might help to find the cause.

Another follow-up discovery of a gamma-ray pulsar is the black widow pulsar PSR J1555–2908, first discovered in a radio search with the Green Bank Telescope (GBT). The gamma-ray search and discovery of the narrow gamma-ray pulse profile are described in Chapter 6. The gamma-ray timing analysis measured four significant spin-frequency derivatives, which is quite unusual since timing noise is expected to be smaller with MSPs. A different physical explanation is given in Chapter 7. Here, the black-widow system is orbited by a planetary-mass companion in a multi-year orbit. With current data, this model shows as strong a signal as the model with four spin-frequency derivatives requiring the same number of parameters. Fortunately, with the ongoing *Fermi* mission, this riddle should be resolved within the next few years.

Of the over 5,000 gamma-ray sources listed in the *Fermi*-LAT Fourth Source

Catalog more than a quarter are still unassociated. A few of them are suspected to be spider pulsars detectable in partially informed searches (see, e.g., the overview in Chapter 2). Currently, a circular-binary search for a redback pulsar within 4FGL J0523.3–2527 is running on *Einstein@Home*, as described in Chapter 2. More searches for other spider systems will be started when tighter constraints are available.

So far, gamma-ray pulsations have only been seen from one young binary pulsar. The eccentric-orbit gamma-ray binary LS 5039 is predicted to host a compact object, potentially a neutron star. Using the parameter space metric from Chapter 2, H. Fehrmann recently developed a new code to build the five-dimensional stochastic grids needed to efficiently search the orbital parameter space for a young, binary gamma-ray pulsar within LS 5039. A search with *Einstein@Home* is planned in the near future.

Apart from finding new pulsars, the all-sky LAT data spanning 11 years can also be used to measure the pulsar system parameters more precisely, and track time-varying parameters like the spin-frequency changes in many younger pulsars or the orbital period changes in many spider pulsars. In Chapter 5 a new way to track orbital period variations is presented. A systematic study of the known handful of gamma-ray spider pulsars which experience these variations is planned using this tool.

This thesis presented novel gamma-ray pulsar search methods and the discoveries of four highly interesting spider pulsars. Hopefully the pulsars that have been and will be discovered with the new methods will lead to a better understanding of the nature and evolution of binary pulsars.

Bibliography

- [1] L. Nieder et al. “Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-Ray Pulsars”. In: *ApJ* 901.2 (Apr. 2020), arXiv:2004.11740. DOI: 10.3847/1538-4357/abaf53. arXiv: 2004.11740 [astro-ph.HE].
- [2] L. Nieder et al. “Detection and Timing of Gamma-Ray Pulsations from the 707 Hz Pulsar J0952-0607”. In: *ApJ* 883.1, 42 (2019), p. 42. DOI: 10.3847/1538-4357/ab357e. arXiv: 1905.11352 [astro-ph.HE].
- [3] L. Nieder et al. “Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated Einstein@Home”. In: *arXiv e-prints*, arXiv:2009.01513 (Sept. 2020). arXiv: 2009.01513 [astro-ph.HE].
- [4] C. J. Clark et al. “Einstein@Home Discovery of Gamma-ray Pulsations Confirms the Redback Nature of 3FGL J2039.6-5618”. In: *arXiv e-prints*, arXiv:2007.14849 (July 2020). arXiv: 2007.14849 [astro-ph.HE].
- [5] W. Baade and F. Zwicky. “Remarks on Super-Novae and Cosmic Rays”. In: *Physical Review* 46.1 (July 1934), pp. 76–77. DOI: 10.1103/PhysRev.46.76.2.
- [6] A. Hewish et al. “Observation of a Rapidly Pulsating Radio Source”. In: *Nature* 217.5130 (Feb. 1968), pp. 709–713. DOI: 10.1038/217709a0.
- [7] J. D. H. Pilkington et al. “Observations of some further Pulsed Radio Sources”. In: *Nature* 218.5137 (Apr. 1968), pp. 126–129. DOI: 10.1038/218126a0.
- [8] F. Pacini. “Energy Emission from a Neutron Star”. In: *Nature* 216.5115 (Nov. 1967), pp. 567–568. DOI: 10.1038/216567a0.
- [9] T. Gold. “Rotating Neutron Stars as the Origin of the Pulsating Radio Sources”. In: *Nature* 218.5143 (May 1968), pp. 731–732. DOI: 10.1038/218731a0.
- [10] D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Dec. 2004.
- [11] Andrew G. Lyne and Francis Graham-Smith. *Pulsar Astronomy*. 2006.
- [12] D. C. Backer et al. “A millisecond pulsar”. In: *Nature* 300.5893 (Dec. 1982), pp. 615–618. DOI: 10.1038/300615a0.

- [13] J. W. T. Hessels et al. “A Radio Pulsar Spinning at 716 Hz”. In: *Science* 311 (Mar. 2006), pp. 1901–1904. DOI: 10.1126/science.1123430. eprint: astro-ph/0601337.
- [14] C. G. Bassa et al. “LOFAR Discovery of the Fastest-spinning Millisecond Pulsar in the Galactic Field”. In: *ApJ* 846, L20 (Sept. 2017), p. L20. DOI: 10.3847/2041-8213/aa8400. arXiv: 1709.01453 [astro-ph.HE].
- [15] V. Morello et al. “The High Time Resolution Universe survey - XIV. Discovery of 23 pulsars through GPU-accelerated reprocessing”. In: *MNRAS* 483.3 (Mar. 2019), pp. 3673–3685. DOI: 10.1093/mnras/sty3328. arXiv: 1811.04929 [astro-ph.IM].
- [16] M. A. Alpar et al. “A new class of radio pulsars”. In: *Nature* 300 (Dec. 1982), pp. 728–730. DOI: 10.1038/300728a0.
- [17] D. Bhattacharya and E. P. J. van den Heuvel. “Formation and evolution of binary and millisecond radio pulsars”. In: *Phys. Rep.* 203.1-2 (Jan. 1991), pp. 1–124. DOI: 10.1016/0370-1573(91)90064-S.
- [18] Juhan Frank, Andrew King, and Derek J. Raine. *Accretion Power in Astrophysics: Third Edition*. 2002.
- [19] A. M. Archibald et al. “A Radio Pulsar/X-ray Binary Link”. In: *Science* 324 (June 2009), p. 1411. DOI: 10.1126/science.1172740. arXiv: 0905.3397 [astro-ph.HE].
- [20] A. Papitto et al. “Swings between rotation and accretion power in a binary millisecond pulsar”. In: *Nature* 501 (Sept. 2013), pp. 517–520. DOI: 10.1038/nature12470. arXiv: 1305.3884 [astro-ph.HE].
- [21] C. G. Bassa et al. “A state change in the low-mass X-ray binary XSS J12270-4859”. In: *MNRAS* 441 (June 2014), pp. 1825–1830. DOI: 10.1093/mnras/stu708. arXiv: 1402.0765 [astro-ph.HE].
- [22] B. W. Stappers et al. “A State Change in the Missing Link Binary Pulsar System PSR J1023+0038”. In: *ApJ* 790, 39 (July 2014), p. 39. DOI: 10.1088/0004-637X/790/1/39. arXiv: 1311.7506 [astro-ph.HE].
- [23] P. A. Caraveo. “Gamma-Ray Pulsar Revolution”. In: *ARA&A* 52 (Aug. 2014), pp. 211–250. DOI: 10.1146/annurev-astro-081913-035948. arXiv: 1312.2913 [astro-ph.HE].
- [24] D. A. Kniffen et al. “Gamma radiation from the Crab Nebula above 35 MeV”. In: *Nature* 251.5474 (Oct. 1974), pp. 397–399. DOI: 10.1038/251397a0.
- [25] C. E. Fichtel et al. “High-energy gamma-ray results from the second Small Astronomy Satellite.” In: *ApJ* 198 (May 1975), pp. 163–182. DOI: 10.1086/153590.
- [26] J. P. Halpern and S. S. Holt. “Discovery of soft X-ray pulsations from the γ -ray source Geminga”. In: *Nature* 357.6375 (May 1992), pp. 222–224. DOI: 10.1038/357222a0.
- [27] G. F. Bignami et al. “The COS-B experiment for gamma-ray astronomy.” In: *Space Science Instrumentation* 1 (Aug. 1975), pp. 245–268.

- [28] A. A. Abdo et al. “The Second Fermi Large Area Telescope Catalog of Gamma-Ray Pulsars”. In: *ApJS* 208, 17 (Oct. 2013), p. 17. DOI: 10.1088/0067-0049/208/2/17. arXiv: 1305.4385 [astro-ph.HE].
- [29] G. F. Bignami and P. A. Caraveo. “Geminga: new period, old γ -rays”. In: *Nature* 357.6376 (May 1992), p. 287. DOI: 10.1038/357287a0.
- [30] G. Kanbach et al. “The project EGRET (energetic gamma-ray experiment telescope) on NASA’s Gamma-Ray Observatory GRO”. In: *Space Sci. Rev.* 49.1-2 (Jan. 1989), pp. 69–84. DOI: 10.1007/BF00173744.
- [31] M. Tavani et al. “The AGILE Mission”. In: *A&A* 502.3 (Aug. 2009), pp. 995–1013. DOI: 10.1051/0004-6361/200810527. arXiv: 0807.4254 [astro-ph].
- [32] W. B. Atwood et al. “The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission”. In: *ApJ* 697 (June 2009), pp. 1071–1102. DOI: 10.1088/0004-637X/697/2/1071. arXiv: 0902.1089 [astro-ph.IM].
- [33] Charles Meegan et al. “The Fermi Gamma-ray Burst Monitor”. In: *ApJ* 702.1 (Sept. 2009), pp. 791–804. DOI: 10.1088/0004-637X/702/1/791. arXiv: 0908.0450 [astro-ph.IM].
- [34] A. A. Abdo et al. “Fermi/Large Area Telescope Bright Gamma-Ray Source List”. In: *ApJS* 183.1 (July 2009), pp. 46–66. DOI: 10.1088/0067-0049/183/1/46. arXiv: 0902.1340 [astro-ph.HE].
- [35] A. A. Abdo et al. “Fermi Large Area Telescope First Source Catalog”. In: *The Astrophysical Journal Supplement Series* 188 (June 2010), pp. 405–436. DOI: 10.1088/0067-0049/188/2/405. arXiv: 1002.2280 [astro-ph.HE].
- [36] P. L. Nolan et al. “Fermi Large Area Telescope Second Source Catalog”. In: *ApJS* 199, 31 (Apr. 2012), p. 31. DOI: 10.1088/0067-0049/199/2/31. arXiv: 1108.1435 [astro-ph.HE].
- [37] F. Acero et al. “Fermi Large Area Telescope Third Source Catalog”. In: *ApJS* 218, 23 (June 2015), p. 23. DOI: 10.1088/0067-0049/218/2/23. arXiv: 1501.02003 [astro-ph.HE].
- [38] S. Abdollahi et al. “Fermi Large Area Telescope Fourth Source Catalog”. In: *ApJS* 247.1, 33 (Mar. 2020), p. 33. DOI: 10.3847/1538-4365/ab6bcb. arXiv: 1902.10045 [astro-ph.HE].
- [39] A. A. Abdo et al. “The First Fermi Large Area Telescope Catalog of Gamma-ray Pulsars”. In: *ApJS* 187.2 (Apr. 2010), pp. 460–494. DOI: 10.1088/0067-0049/187/2/460. arXiv: 0910.1608 [astro-ph.HE].
- [40] Ren-Bo Wang and Kouichi Hirotani. “Death Line of Gamma-Ray Pulsars with Outer Gaps”. In: *ApJ* 736.2, 127 (Aug. 2011), p. 127. DOI: 10.1088/0004-637X/736/2/127. arXiv: 1105.3030 [astro-ph.HE].
- [41] L. Guillemot et al. “The gamma-ray millisecond pulsar deathline, revisited. New velocity and distance measurements”. In: *A&A* 587, A109 (Mar. 2016), A109. DOI: 10.1051/0004-6361/201527847. arXiv: 1601.05987 [astro-ph.HE].
- [42] D. A. Smith et al. “Searching a Thousand Radio Pulsars for Gamma-Ray Emission”. In: *ApJ* 871.1, 78 (2019), p. 78. DOI: 10.3847/1538-4357/aaf57d. arXiv: 1812.00719 [astro-ph.HE].

- [43] Colin J. Clark et al. “Einstein@Home discovers a radio-quiet gamma-ray millisecond pulsar”. In: *Science Advances* 4 (2018), eaao7228. DOI: 10.1126/sciadv.aao7228. arXiv: 1803.06855 [astro-ph.HE].
- [44] H. J. Pletsch et al. “Binary Millisecond Pulsar Discovery via Gamma-Ray Pulsations”. In: *Science* 338 (Dec. 2012), p. 1314. DOI: 10.1126/science.1229054. arXiv: 1211.1385 [astro-ph.HE].
- [45] A. G. Lyne et al. “The binary nature of PSR J2032+4127”. In: *MNRAS* 451.1 (July 2015), pp. 581–587. DOI: 10.1093/mnras/stv236. arXiv: 1502.01465 [astro-ph.HE].
- [46] Isabelle A. Grenier and Alice K. Harding. “Gamma-ray pulsars: A gold mine”. In: *Comptes Rendus Physique* 16.6-7 (Aug. 2015), pp. 641–660. DOI: 10.1016/j.crhy.2015.08.013. arXiv: 1509.08823 [astro-ph.HE].
- [47] B. Cerutti and A. Beloborodov. “Electrodynamics of pulsar magnetospheres”. In: *ArXiv e-prints* (Nov. 2016). arXiv: 1611.04331 [astro-ph.HE].
- [48] Alice K. Harding. “Pulsar Emission Physics: The First Fifty Years”. In: *Pulsar Astrophysics the Next Fifty Years*. Ed. by P. Weltevrede et al. Vol. 337. Aug. 2018, pp. 52–57. DOI: 10.1017/S1743921318000509. arXiv: 1712.02409 [astro-ph.HE].
- [49] A. S. Fruchter, D. R. Stinebring, and J. H. Taylor. “A millisecond pulsar in an eclipsing binary”. In: *Nature* 333 (May 1988), pp. 237–239. DOI: 10.1038/333237a0.
- [50] A. S. Fruchter et al. “The Eclipsing Millisecond Pulsar PSR 1957+20”. In: *ApJ* 351 (Mar. 1990), p. 642. DOI: 10.1086/168502.
- [51] Mallory S. E. Roberts. “New Black Widows and Redbacks in the Galactic Field”. In: *American Institute of Physics Conference Series*. Ed. by Marta Burgay et al. Vol. 1357. American Institute of Physics Conference Series. Aug. 2011, pp. 127–130. DOI: 10.1063/1.3615095. arXiv: 1103.0819 [astro-ph.HE].
- [52] P. S. Ray et al. “Precise γ -ray Timing and Radio Observations of 17 Fermi γ -ray Pulsars”. In: *ApJS* 194, 17 (June 2011), p. 17. DOI: 10.1088/0067-0049/194/2/17. arXiv: 1011.2468 [astro-ph.HE].
- [53] Mallory S. E. Roberts. “Surrounded by spiders! New black widows and redbacks in the Galactic field”. In: *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*. Ed. by Joeri van Leeuwen. Vol. 291. IAU Symposium. 2013, pp. 127–132. DOI: 10.1017/S174392131202337X. arXiv: 1210.6903 [astro-ph.HE].
- [54] Jay Strader et al. “Optical Spectroscopy and Demographics of Redback Millisecond Pulsar Binaries”. In: *ApJ* 872.1, 42 (Feb. 2019), p. 42. DOI: 10.3847/1538-4357/aafbba. arXiv: 1812.04626 [astro-ph.HE].
- [55] J. Strader et al. “1FGL J0523.5-2529: A New Probable Gamma-Ray Pulsar Binary”. In: *ApJ* 788, L27 (June 2014), p. L27. DOI: 10.1088/2041-8205/788/2/L27. arXiv: 1405.5533 [astro-ph.HE].
- [56] P. S. Ray et al. “Radio Searches of Fermi LAT Sources and Blind Search Pulsars: The Fermi Pulsar Search Consortium”. In: *ArXiv e-prints* (May 2012). arXiv: 1205.3089 [astro-ph.HE].

- [57] H. T. Cromartie et al. “Six New Millisecond Pulsars from Arecibo Searches of Fermi Gamma-Ray Sources”. In: *ApJ* 819, 34 (Mar. 2016), p. 34. DOI: 10.3847/0004-637X/819/1/34. arXiv: 1601.05343 [astro-ph.HE].
- [58] E. D. Barr et al. “Pulsar searches of Fermi unassociated sources with the Effelsberg telescope”. In: *MNRAS* 429 (2013), pp. 1633–1642. DOI: 10.1093/mnras/sts449. arXiv: 1301.0359 [astro-ph.HE].
- [59] Pei Wang et al. “FAST’s Discovery of a New Millisecond Pulsar (MSP) toward the Fermi-LAT unassociated source 3FGL J0318.1+0252”. In: *The Astronomer’s Telegram* 11584 (Apr. 2018), p. 1.
- [60] B. Bhattacharyya et al. “GMRT Discovery of PSR J1544+4937: An Eclipsing Black-widow Pulsar Identified with a Fermi-LAT Source”. In: *ApJ* 773, L12 (Aug. 2013), p. L12. DOI: 10.1088/2041-8205/773/1/L12. arXiv: 1304.7101 [astro-ph.HE].
- [61] S. M. Ransom et al. “Three Millisecond Pulsars in Fermi LAT Unassociated Bright Sources”. In: *ApJ* 727, L16 (Jan. 2011), p. L16. DOI: 10.1088/2041-8205/727/1/L16. arXiv: 1012.2862 [astro-ph.HE].
- [62] K. Stovall et al. “The Green Bank Northern Celestial Cap Pulsar Survey. I. Survey Description, Data Analysis, and Initial Results”. In: *ApJ* 791, 67 (Aug. 2014), p. 67. DOI: 10.1088/0004-637X/791/1/67. arXiv: 1406.5214 [astro-ph.HE].
- [63] Z. Pleunis et al. “A Millisecond Pulsar Discovery in a Survey of Unidentified Fermi γ -Ray Sources with LOFAR”. In: *ApJ* 846, L19 (Sept. 2017), p. L19. DOI: 10.3847/2041-8213/aa83ff. arXiv: 1709.01452 [astro-ph.HE].
- [64] I. Cognard et al. “Discovery of Two Millisecond Pulsars in Fermi Sources with the Nançay Radio Telescope”. In: *ApJ* 732, 47 (May 2011), p. 47. DOI: 10.1088/0004-637X/732/1/47. arXiv: 1102.4192 [astro-ph.HE].
- [65] F. Camilo et al. “Parkes Radio Searches of Fermi Gamma-Ray Sources and Millisecond Pulsar Discoveries”. In: *ApJ* 810, 85 (Sept. 2015), p. 85. DOI: 10.1088/0004-637X/810/2/85. arXiv: 1507.04451 [astro-ph.HE].
- [66] K. J. Lee et al. “Application of the Gaussian mixture model in pulsar astronomy - pulsar classification and candidates ranking for the Fermi 2FGL catalogue”. In: *MNRAS* 424 (Aug. 2012), pp. 2832–2840. DOI: 10.1111/j.1365-2966.2012.21413.x. arXiv: 1205.6221 [astro-ph.IM].
- [67] P. M. Saz Parkinson et al. “Classification and Ranking of Fermi LAT Gamma-ray Sources from the 3FGL Catalog using Machine Learning Techniques”. In: *ApJ* 820, 8 (Mar. 2016), p. 8. DOI: 10.3847/0004-637X/820/1/8. arXiv: 1602.00385 [astro-ph.HE].
- [68] Paul S. Ray et al. “Radio Discovery of and Gamma-Ray Pulsations from PSR J2339-0533”. In: *Research Notes of the AAS* 4.3 (2020), p. 37. DOI: 10.3847/2515-5172/ab7eb5. URL: <https://doi.org/10.3847/2515-5172/ab7eb5>.
- [69] Roger W. Romani and Michael S. Shaw. “The Orbit and Companion of Probable γ -Ray Pulsar J2339-0533”. In: *ApJ* 743.2, L26 (Dec. 2011), p. L26. DOI: 10.1088/2041-8205/743/2/L26. arXiv: 1111.3074 [astro-ph.HE].

- [70] A. K. H. Kong et al. “Discovery of an Unidentified Fermi Object as a Black Widow-like Millisecond Pulsar”. In: *ApJ* 747.1, L3 (Mar. 2012), p. L3. DOI: 10.1088/2041-8205/747/1/L3. arXiv: 1201.3629 [astro-ph.HE].
- [71] H. J. Pletsch and C. J. Clark. “Gamma-Ray Timing of Redback PSR J2339-0533: Hints for Gravitational Quadrupole Moment Changes”. In: *ApJ* 807, 18 (July 2015), p. 18. DOI: 10.1088/0004-637X/807/1/18. arXiv: 1504.07466 [astro-ph.HE].
- [72] A. A. Abdo et al. “The Fermi Gamma-Ray Space Telescope Discovers the Pulsar in the Young Galactic Supernova Remnant CTA 1”. In: *Science* 322.5905 (Nov. 2008), p. 1218. DOI: 10.1126/science.1165572. arXiv: 0810.3562 [astro-ph].
- [73] A. A. Abdo et al. “Detection of 16 Gamma-Ray Pulsars Through Blind Frequency Searches Using the Fermi LAT”. In: *Science* 325 (Aug. 2009), p. 840. DOI: 10.1126/science.1175558. arXiv: 1009.0748 [astro-ph.GA].
- [74] P. M. Saz Parkinson et al. “Eight γ -ray Pulsars Discovered in Blind Frequency Searches of Fermi LAT Data”. In: *ApJ* 725 (Dec. 2010), pp. 571–584. DOI: 10.1088/0004-637X/725/1/571. arXiv: 1006.2134 [astro-ph.HE].
- [75] H. J. Pletsch et al. “Discovery of Nine Gamma-Ray Pulsars in Fermi Large Area Telescope Data Using a New Blind Search Method”. In: *ApJ* 744, 105 (Jan. 2012), p. 105. DOI: 10.1088/0004-637X/744/2/105. arXiv: 1111.0523 [astro-ph.HE].
- [76] H. J. Pletsch et al. “Einstein@Home Discovery of Four Young Gamma-Ray Pulsars in Fermi LAT Data”. In: *ApJ* 779.1, L11 (Dec. 2013), p. L11. DOI: 10.1088/2041-8205/779/1/L11. arXiv: 1311.6427 [astro-ph.HE].
- [77] C. J. Clark et al. “PSR J1906+0722: An Elusive Gamma-Ray Pulsar”. In: *ApJ* 809, L2 (Aug. 2015), p. L2. DOI: 10.1088/2041-8205/809/1/L2. arXiv: 1508.00779 [astro-ph.HE].
- [78] C. J. Clark et al. “The Braking Index of a Radio-quiet Gamma-Ray Pulsar”. In: *ApJ* 832, L15 (Nov. 2016), p. L15. DOI: 10.3847/2041-8205/832/1/L15. arXiv: 1611.01292 [astro-ph.HE].
- [79] C. J. Clark et al. “The Einstein@Home Gamma-ray Pulsar Survey. I. Search Methods, Sensitivity, and Discovery of New Young Gamma-Ray Pulsars”. In: *ApJ* 834, 106 (Jan. 2017), p. 106. DOI: 10.3847/1538-4357/834/2/106. arXiv: 1611.01015 [astro-ph.HE].
- [80] H. J. Pletsch and C. J. Clark. “Optimized Blind Gamma-Ray Pulsar Searches at Fixed Computing Budget”. In: *ApJ* 795, 75 (Nov. 2014), p. 75. DOI: 10.1088/0004-637X/795/1/75. arXiv: 1408.6962 [astro-ph.HE].
- [81] Holger J. Pletsch and Bruce Allen. “Exploiting Large-Scale Correlations to Detect Continuous Gravitational Waves”. In: *Phys. Rev. Lett.* 103.18, 181102 (Oct. 2009), p. 181102. DOI: 10.1103/PhysRevLett.103.181102. arXiv: 0906.0023 [gr-qc].
- [82] Holger J. Pletsch. “Parameter-space metric of semicoherent searches for continuous gravitational waves”. In: *Phys. Rev. D* 82.4, 042002 (Aug. 2010), p. 042002. DOI: 10.1103/PhysRevD.82.042002. arXiv: 1005.0395 [gr-qc].

- [83] Holger J. Pletsch. “Sliding coherence window technique for hierarchical detection of continuous gravitational waves”. In: *Phys. Rev. D* 83.12, 122003 (June 2011), p. 122003. DOI: 10.1103/PhysRevD.83.122003. arXiv: 1101.5396 [gr-qc].
- [84] C. Aulbert and H. Fehrmann. “Searching gravitational waves with one of the world’s fastest super computers”. In: *Forschungsbericht 2008 - Max-Planck-Institut für Gravitationsphysik, Teilinstitut Hannover*, <https://www.mpg.de/308429/forschungsSchwerpunkt> (2008).
- [85] B. Allen et al. “The Einstein@Home Search for Radio Pulsars and PSR J2007+2722 Discovery”. In: *ApJ* 773, 91 (Aug. 2013), p. 91. DOI: 10.1088/0004-637X/773/2/91. arXiv: 1303.0028 [astro-ph.IM].
- [86] David P Anderson. “Boinc: A system for public-resource computing and storage”. In: *Fifth IEEE/ACM international workshop on grid computing*. IEEE, 2004, pp. 4–10.
- [87] Roger W. Romani. “2FGL J1311.7-3429 Joins the Black Widow Club”. In: *ApJ* 754.2, L25 (Aug. 2012), p. L25. DOI: 10.1088/2041-8205/754/2/L25. arXiv: 1207.1736 [astro-ph.HE].
- [88] K.-L. Li et al. “Discovery of a Redback Millisecond Pulsar Candidate: 3FGL J0212.1+5320”. In: *ApJ* 833, 143 (Dec. 2016), p. 143. DOI: 10.3847/1538-4357/833/2/143. arXiv: 1609.02951 [astro-ph.HE].
- [89] Manuel Linares et al. “A millisecond pulsar candidate in a 21-h orbit: 3FGL J0212.1+5320”. In: *MNRAS* 465.4 (Mar. 2017), pp. 4602–4610. DOI: 10.1093/mnras/stw3057. arXiv: 1609.02232 [astro-ph.HE].
- [90] D. Salvetti et al. “A multiwavelength investigation of candidate millisecond pulsars in unassociated γ -ray sources”. In: *MNRAS* 470 (Sept. 2017), pp. 466–480. DOI: 10.1093/mnras/stx1247. arXiv: 1702.00474 [astro-ph.HE].
- [91] J. P. Halpern, J. Strader, and M. Li. “A Likely Redback Millisecond Pulsar Counterpart of 3FGL J0838.8-2829”. In: *ApJ* 844, 150 (Aug. 2017), p. 150. DOI: 10.3847/1538-4357/aa7cff. arXiv: 1706.09511 [astro-ph.HE].
- [92] Kwan-Lok Li et al. “Multiwavelength Observations of a New Redback Millisecond Pulsar Candidate: 3FGL J0954.8-3948”. In: *ApJ* 863.2, 194 (Aug. 2018), p. 194. DOI: 10.3847/1538-4357/aad243. arXiv: 1807.02508 [astro-ph.HE].
- [93] Samuel J. Swihart et al. “A New Likely Redback Millisecond Pulsar Binary with a Massive Neutron Star: 4FGL J2333.1-5527”. In: *ApJ* 892.1, 21 (Mar. 2020), p. 21. DOI: 10.3847/1538-4357/ab77ba. arXiv: 1912.02264 [astro-ph.HE].
- [94] C. Braglia et al. “A multiwavelength search for black widow and redback counterparts of candidate γ -ray millisecond pulsars”. In: *MNRAS* 497.4 (Aug. 2020), pp. 5364–5382. DOI: 10.1093/mnras/staa2339. arXiv: 2007.00442 [astro-ph.HE].
- [95] G. B. Hobbs, R. T. Edwards, and R. N. Manchester. “TEMPO2, a new pulsar-timing package - I. An overview”. In: *MNRAS* 369 (June 2006), pp. 655–672. DOI: 10.1111/j.1365-2966.2006.10302.x. eprint: astro-ph/0603381.

- [96] R. T. Edwards, G. B. Hobbs, and R. N. Manchester. “TEMPO2, a new pulsar timing package - II. The timing model and precision estimates”. In: MNRAS 372 (Nov. 2006), pp. 1549–1574. DOI: 10.1111/j.1365-2966.2006.10870.x. eprint: astro-ph/0607664.
- [97] M. Kerr et al. “Timing Gamma-ray Pulsars with the Fermi Large Area Telescope: Timing Noise and Astrometry”. In: ApJ 814, 128 (Dec. 2015), p. 128. DOI: 10.1088/0004-637X/814/2/128. arXiv: 1510.05099 [astro-ph.HE].
- [98] Z. Arzoumanian, A. S. Fruchter, and J. H. Taylor. “Orbital variability in the eclipsing pulsar binary PSR B1957+20”. In: ApJ 426 (May 1994), pp. 85–88. DOI: 10.1086/187346. eprint: astro-ph/9312032.
- [99] K. Lazaridis et al. “Evidence for gravitational quadrupole moment variations in the companion of PSR J2051-0827”. In: MNRAS 414.4 (2011), pp. 3134–3144. DOI: 10.1111/j.1365-2966.2011.18610.x. arXiv: 1102.5646 [astro-ph.GA].
- [100] J. S. Deneva et al. “Multiwavelength Observations of the Redback Millisecond Pulsar J1048+2339”. In: ApJ 823, 105 (June 2016), p. 105. DOI: 10.3847/0004-637X/823/2/105. arXiv: 1601.03681 [astro-ph.HE].
- [101] James H. Applegate and Jacob Shaham. “Orbital Period Variability in the Eclipsing Pulsar Binary PSR B1957+20: Evidence for a Tidally Powered Star”. In: ApJ 436 (1994), p. 312. DOI: 10.1086/174906.
- [102] James H. Applegate and Joseph Patterson. “Magnetic Activity, Tides, and Orbital Period Changes in Close Binaries”. In: ApJ 322 (1987), p. L99. DOI: 10.1086/185044.
- [103] Guillaume Voisin, René P. Breton, and Charlotte Summers. “A spider timing model: accounting for quadrupole deformations and relativity in close pulsar binaries”. In: MNRAS 492.2 (2020), pp. 1550–1565. DOI: 10.1093/mnras/stz3430. arXiv: 1912.03935 [astro-ph.HE].
- [104] Guillaume Voisin et al. “First measurement of the total gravitational quadrupole moment of a black widow companion”. In: MNRAS 494.3 (Apr. 2020), pp. 4448–4453. DOI: 10.1093/mnras/staa953. arXiv: 2004.01564 [astro-ph.HE].
- [105] A. Corongiu et al. “Radio pulsations from the γ -ray millisecond pulsar PSR J2039-5617”. In: *arXiv e-prints* (July 2020). arXiv: 2007.14889 [astro-ph.HE].
- [106] A. A. Abdo et al. “Discovery of Pulsed γ -Rays from the Young Radio Pulsar PSR J1028-5819 with the Fermi Large Area Telescope”. In: ApJ 695 (Apr. 2009), pp. L72–L77. DOI: 10.1088/0004-637X/695/1/L72. arXiv: 0903.1602 [astro-ph.HE].
- [107] A. A. Abdo et al. “Pulsed Gamma Rays from the Millisecond Pulsar J0030+0451 with the Fermi Large Area Telescope”. In: ApJ 699 (July 2009), pp. 1171–1177. DOI: 10.1088/0004-637X/699/2/1171. arXiv: 0904.4377 [astro-ph.HE].
- [108] L. Guillemot et al. “Pulsed Gamma Rays from the Original Millisecond and Black Widow Pulsars: A Case for Caustic Radio Emission?” In: ApJ 744, 33 (Jan. 2012), p. 33. DOI: 10.1088/0004-637X/744/1/33. arXiv: 1110.1271 [astro-ph.HE].

- [109] D. A. Smith et al. “Gamma-ray pulsars with Fermi”. In: *ArXiv e-prints* (June 2017). arXiv: 1706.03592 [astro-ph.HE].
- [110] R. N. Manchester et al. “The Australia Telescope National Facility Pulsar Catalogue”. In: *AJ* 129 (Apr. 2005), pp. 1993–2006. DOI: 10.1086/428488. eprint: astro-ph/0412641.
- [111] R. W. Romani, A. V. Filippenko, and S. B. Cenko. “2FGL J1653.6-0159: A New Low in Evaporating Pulsar Binary Periods”. In: *ApJ* 793, L20 (Sept. 2014), p. L20. DOI: 10.1088/2041-8205/793/1/L20. arXiv: 1408.2886 [astro-ph.HE].
- [112] A. K. H. Kong et al. “Discovery of an Ultracompact Gamma-Ray Millisecond Pulsar Binary Candidate”. In: *ApJ* 794, L22 (Oct. 2014), p. L22. DOI: 10.1088/2041-8205/794/2/L22. arXiv: 1408.5162 [astro-ph.HE].
- [113] P. Bickel, B. Kleijn, and J. Rice. “Event-Weighted Tests for Detecting Periodicity in Photon Arrival Times”. In: *ApJ* 685, 384-389 (Sept. 2008), pp. 384–389. DOI: 10.1086/590399. arXiv: 0706.4108 [stat.ME].
- [114] M. Kerr. “Improving Sensitivity to Weak Pulsations with Photon Probability Weighting”. In: *ApJ* 732, 38 (May 2011), p. 38. DOI: 10.1088/0004-637X/732/1/38. arXiv: 1103.2128 [astro-ph.IM].
- [115] P. Bruel. “Extending the event-weighted pulsation search to very faint gamma-ray sources”. In: *A&A* 622, A108 (Feb. 2019), A108. DOI: 10.1051/0004-6361/201834555. arXiv: 1812.06681 [astro-ph.IM].
- [116] N. Meinshausen, P. Bickel, and J. Rice. “Efficient blind search: Optimal power of detection under computational cost constraints”. In: *Ann. Appl. Stat.* 3.1 (2009), pp. 38–60. DOI: 10.1214/08-AOAS180. URL: <http://dx.doi.org/10.1214/08-AOAS180>.
- [117] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. “Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters”. In: *Phys. Rev. D* 53 (Mar. 1996), pp. 3033–3055. DOI: 10.1103/PhysRevD.53.3033. eprint: gr-qc/9508011.
- [118] B. J. Owen. “Search templates for gravitational waves from inspiraling binaries: Choice of template spacing”. In: *Phys. Rev. D* 53 (June 1996), pp. 6749–6761. DOI: 10.1103/PhysRevD.53.6749. eprint: gr-qc/9511032.
- [119] O. C. de Jager, B. C. Raubenheimer, and J. W. H. Swanepoel. “A powerful test for weak periodic signals with unknown light curve shape in sparse data”. In: *A&A* 221 (Aug. 1989), pp. 180–190.
- [120] M. Frigo and S. G. Johnson. “The Design and Implementation of FFTW3”. In: *Proceedings of the IEEE* 93.2 (2005), pp. 216–231. ISSN: 0018-9219. DOI: 10.1109/JPROC.2004.840301.
- [121] Bruce Allen. “Spherical ansatz for parameter-space metrics”. In: *Phys. Rev. D* 100.12, 124004 (Dec. 2019), p. 124004. DOI: 10.1103/PhysRevD.100.124004. arXiv: 1906.01352 [gr-qc].
- [122] Gaia Collaboration et al. “Gaia Data Release 2. Summary of the contents and survey properties”. In: *A&A* 616, A1 (Aug. 2018), A1. DOI: 10.1051/0004-6361/201833051. arXiv: 1804.09365.

- [123] C. Y. Hui et al. “Searches for Millisecond Pulsar Candidates among the Unidentified Fermi Objects”. In: *ApJ* 809, 68 (Aug. 2015), p. 68. DOI: 10.1088/0004-637X/809/1/68. arXiv: 1507.02604 [astro-ph.HE].
- [124] D. G. Monet et al. “The USNO-B Catalog”. In: *AJ* 125 (Feb. 2003), pp. 984–993. DOI: 10.1086/345888. eprint: astro-ph/0210694.
- [125] H.-L. Chen et al. “Formation of Black Widows and Redbacks - Two Distinct Populations of Eclipsing Binary Millisecond Pulsars”. In: *ApJ* 775, 27 (Sept. 2013), p. 27. DOI: 10.1088/0004-637X/775/1/27. arXiv: 1308.4107 [astro-ph.SR].
- [126] H. T. Cromartie et al. “Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar”. In: *Nature Astronomy* 4 (Jan. 2020), pp. 72–76. DOI: 10.1038/s41550-019-0880-2. arXiv: 1904.06759 [astro-ph.HE].
- [127] H. Fehrmann and H. J. Pletsch. “Efficient generation and optimization of stochastic template banks by a neighboring cell algorithm”. In: *Phys. Rev. D* 90.12, 124049 (Dec. 2014), p. 124049. DOI: 10.1103/PhysRevD.90.124049. arXiv: 1411.3899 [astro-ph.IM].
- [128] Stanislav Babak. “Building a stochastic template bank for detecting massive black hole binaries”. In: *Classical and Quantum Gravity* 25.19, 195011 (Oct. 2008), p. 195011. DOI: 10.1088/0264-9381/25/19/195011. arXiv: 0801.4070 [gr-qc].
- [129] I. W. Harry, B. Allen, and B. S. Sathyaprakash. “Stochastic template placement algorithm for gravitational wave data analysis”. In: *Phys. Rev. D* 80.10, 104014 (Nov. 2009), p. 104014. DOI: 10.1103/PhysRevD.80.104014. arXiv: 0908.2090 [gr-qc].
- [130] E. S. Phinney. “Pulsars as Probes of Newtonian Dynamical Systems”. In: *Philosophical Transactions of the Royal Society of London Series A* 341 (Oct. 1992), pp. 39–75. DOI: 10.1098/rsta.1992.0084.
- [131] B. Knispel et al. “Einstein@Home Discovery of a PALFA Millisecond Pulsar in an Eccentric Binary Orbit”. In: *ApJ* 806, 140 (June 2015), p. 140. DOI: 10.1088/0004-637X/806/1/140. arXiv: 1504.03684 [astro-ph.HE].
- [132] R. Blandford and S. A. Teukolsky. “Arrival-time analysis for a pulsar in a binary system.” In: *ApJ* 205 (Apr. 1976), pp. 580–591. DOI: 10.1086/154315.
- [133] C. Lange et al. “Precision timing measurements of PSR J1012+5307”. In: *MNRAS* 326 (Sept. 2001), pp. 274–282. DOI: 10.1046/j.1365-8711.2001.04606.x. eprint: astro-ph/0102309.
- [134] C. Messenger et al. “Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors”. In: *Phys. Rev. D* 92.2, 023006 (July 2015), p. 023006. DOI: 10.1103/PhysRevD.92.023006. arXiv: 1504.05889 [gr-qc].
- [135] F. Camilo et al. “Observations of 20 Millisecond Pulsars in 47 Tucanae at 20 Centimeters”. In: *ApJ* 535 (June 2000), pp. 975–990. DOI: 10.1086/308859. eprint: astro-ph/9911234.
- [136] S. M. Ransom et al. “A Binary Millisecond Pulsar in Globular Cluster NGC 6544”. In: *ApJ* 546 (Jan. 2001), pp. L25–L28. DOI: 10.1086/318062. eprint: astro-ph/0010243.

- [137] Bridget C. Andersen and Scott M. Ransom. “A Fourier Domain “Jerk” Search for Binary Pulsars”. In: *ApJ* 863.1, L13 (Aug. 2018), p. L13. DOI: 10.3847/2041-8213/aad59f. arXiv: 1807.07900 [astro-ph.HE].
- [138] S. van der Putten et al. “Searching for gravitational waves from pulsars in binary systems: An all-sky search”. In: *Journal of Physics Conference Series*. Vol. 228. Journal of Physics Conference Series. May 2010, p. 012005. DOI: 10.1088/1742-6596/228/1/012005.
- [139] H. M. Johnston and S. R. Kulkarni. “On the detectability of pulsars in close binary systems”. In: *ApJ* 368 (Feb. 1991), pp. 504–514. DOI: 10.1086/169715.
- [140] W. C. G. Ho et al. “Multiwavelength monitoring and X-ray brightening of Be X-ray binary PSR J2032+4127/MT91 213 on its approach to periastron”. In: *MNRAS* 464 (Jan. 2017), pp. 1211–1219. DOI: 10.1093/mnras/stw2420. arXiv: 1609.06328 [astro-ph.SR].
- [141] A. J. Faulkner et al. “The Parkes Multibeam Pulsar Survey - V. Finding binary and millisecond pulsars”. In: *MNRAS* 355 (Nov. 2004), pp. 147–158. DOI: 10.1111/j.1365-2966.2004.08310.x. eprint: astro-ph/0408228.
- [142] A. J. Faulkner et al. “PSR J1756-2251: A New Relativistic Double Neutron Star System”. In: *ApJ* 618 (Jan. 2005), pp. L119–L122. DOI: 10.1086/427776. eprint: astro-ph/0411796.
- [143] P. R. Brady and T. Creighton. “Searching for periodic sources with LIGO. II. Hierarchical searches”. In: *Phys. Rev. D* 61.8, 082001 (Apr. 2000), p. 082001. DOI: 10.1103/PhysRevD.61.082001. eprint: gr-qc/9812014.
- [144] Keith Riles. “Recent searches for continuous gravitational waves”. In: *Modern Physics Letters A* 32.39, 1730035-685 (Dec. 2017), pp. 1730035–685. DOI: 10.1142/S021773231730035X. arXiv: 1712.05897 [gr-qc].
- [145] A. G. Lyne et al. “Radio pulsars in Terzan 5”. In: *MNRAS* 316 (Aug. 2000), pp. 491–493. DOI: 10.1046/j.1365-8711.2000.03517.x. eprint: astro-ph/0003070.
- [146] E. Goetz and K. Riles. “An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems”. In: *Classical and Quantum Gravity* 28.21, 215006 (Nov. 2011), p. 215006. DOI: 10.1088/0264-9381/28/21/215006. arXiv: 1103.1301 [gr-qc].
- [147] J. Aasi et al. “Gravitational Waves from Known Pulsars: Results from the Initial Detector Era”. In: *ApJ* 785, 119 (Apr. 2014), p. 119. DOI: 10.1088/0004-637X/785/2/119. arXiv: 1309.4027 [astro-ph.HE].
- [148] C. Messenger and G. Woan. “A fast search strategy for gravitational waves from low-mass x-ray binaries”. In: *Classical and Quantum Gravity* 24 (Oct. 2007), S469–S480. DOI: 10.1088/0264-9381/24/19/S10. eprint: gr-qc/0703155.
- [149] L. Sammut et al. “Implementation of the frequency-modulated sideband search method for gravitational waves from low mass x-ray binaries”. In: *Phys. Rev. D* 89.4, 043001 (Feb. 2014), p. 043001. DOI: 10.1103/PhysRevD.89.043001. arXiv: 1311.1379 [gr-qc].

- [150] S. M. Ransom, J. M. Cordes, and S. S. Eikenberry. “A New Search Technique for Short Orbital Period Binary Pulsars”. In: *ApJ* 589 (June 2003), pp. 911–920. DOI: 10.1086/374806. eprint: astro-ph/0210010.
- [151] C. Aragona et al. “The Orbits of the γ -Ray Binaries LS I +61 303 and LS 5039”. In: *ApJ* 698 (June 2009), pp. 514–518. DOI: 10.1088/0004-637X/698/1/514. arXiv: 0902.4015 [astro-ph.HE].
- [152] T. Damour and N. Deruelle. “General relativistic celestial mechanics of binary systems. II. The post-Newtonian timing formula.” In: *Ann. Inst. Henri Poincaré Phys. Théor* 44.3 (Jan. 1986), pp. 263–292.
- [153] R. W. Romani. “A Likely Millisecond Pulsar Binary Counterpart for Fermi Source 2FGL J2039.6-5620”. In: *ApJ* 812, L24 (Oct. 2015), p. L24. DOI: 10.1088/2041-8205/812/2/L24.
- [154] D. Salvetti et al. “Multi-wavelength Observations of 3FGL J2039.6-5618: A Candidate Redback Millisecond Pulsar”. In: *ApJ* 814, 88 (Dec. 2015), p. 88. DOI: 10.1088/0004-637X/814/2/88. arXiv: 1509.07474 [astro-ph.HE].
- [155] A. Allafort et al. “PSR J2021+4026 in the Gamma Cygni Region: The First Variable γ -Ray Pulsar Seen by the Fermi LAT”. In: *ApJ* 777, L2 (Nov. 2013), p. L2. DOI: 10.1088/2041-8205/777/1/L2. arXiv: 1308.0358 [astro-ph.HE].
- [156] F. K. Schinzel et al. “The Tail of PSR J0002+6216 and the Supernova Remnant CTB 1”. In: *ApJ* 876.1, L17 (May 2019), p. L17. DOI: 10.3847/2041-8213/ab18f7. arXiv: 1904.07993 [astro-ph.HE].
- [157] R. Prix and M. Shaltev. “Search for continuous gravitational waves: Optimal StackSlide method at fixed computing cost”. In: *Phys. Rev. D* 85, 084010 (2012), p. 084010. DOI: 10.1103/PhysRevD.85.084010. arXiv: 1201.4321 [gr-qc].
- [158] L. G. Taff. *Celestial mechanics: A computational guide for the practitioner*. 1985.
- [159] S. V. Dhurandhar and A. Vecchio. “Searching for continuous gravitational wave sources in binary systems”. In: *Phys. Rev. D* 63.12, 122001 (June 2001), p. 122001. DOI: 10.1103/PhysRevD.63.122001. eprint: gr-qc/0011085.
- [160] Amruta Jaodand, Jason W. T. Hessels, and Anne Archibald. “A decade of transitional millisecond pulsars”. In: *Pulsar Astrophysics the Next Fifty Years*. Ed. by P. Weltevrede et al. Vol. 337. IAU Symposium. 2018, pp. 47–51. DOI: 10.1017/S1743921317010407. arXiv: 1711.10565 [astro-ph.HE].
- [161] T. J. Johnson et al. “Discovery of Gamma-Ray Pulsations from the Transitional Redback PSR J1227-4853”. In: *ApJ* 806, 91 (June 2015), p. 91. DOI: 10.1088/0004-637X/806/1/91. arXiv: 1502.06862 [astro-ph.HE].
- [162] L. Levin et al. “The NANOGrav Nine-year Data Set: Monitoring Interstellar Scattering Delays”. In: *ApJ* 818, 166 (Feb. 2016), p. 166. DOI: 10.3847/0004-637X/818/2/166. arXiv: 1601.04490 [astro-ph.HE].
- [163] D. A. Frail et al. “Known Pulsars Identified in the GMRT 150 MHz All-sky Survey”. In: *ApJ* 829, 119 (Oct. 2016), p. 119. DOI: 10.3847/0004-637X/829/2/119. arXiv: 1606.00449 [astro-ph.HE].

- [164] J. P. W. Verbiest et al. “The International Pulsar Timing Array: First data release”. In: MNRAS 458.2 (2016), pp. 1267–1288. DOI: 10.1093/mnras/stw347. arXiv: 1602.03640 [astro-ph.IM].
- [165] Caterina Tiburzi. “Pulsars Probe the Low-Frequency Gravitational Sky: Pulsar Timing Arrays Basics and Recent Results”. In: PASA 35, e013 (2018), e013. DOI: 10.1017/pasa.2018.7. arXiv: 1802.05076 [astro-ph.IM].
- [166] D. A. Frail et al. “An image-based search for pulsars among Fermi unassociated LAT sources”. In: MNRAS 475 (Mar. 2018), pp. 942–954. DOI: 10.1093/mnras/stx3281. arXiv: 1712.06609 [astro-ph.HE].
- [167] S. D. Bates, D. R. Lorimer, and J. P. W. Verbiest. “The pulsar spectral index distribution”. In: MNRAS 431 (May 2013), pp. 1352–1358. DOI: 10.1093/mnras/stt257. arXiv: 1302.2053 [astro-ph.SR].
- [168] Sarah A. Story, Peter L. Gonthier, and Alice K. Harding. “Population Synthesis of Radio and γ -Ray Millisecond Pulsars from the Galactic Disk”. In: ApJ 671 (2007), pp. 713–726. DOI: 10.1086/521016. arXiv: 0706.3041 [astro-ph].
- [169] J. Wu et al. “The Einstein@Home Gamma-ray Pulsar Survey. II. Source Selection, Spectral Analysis, and Multiwavelength Follow-up”. In: ApJ 854, 99 (Feb. 2018), p. 99. DOI: 10.3847/1538-4357/aaa411. arXiv: 1712.05395 [astro-ph.HE].
- [170] B. W. Stappers et al. “Observing pulsars and fast transients with LOFAR”. In: A&A 530, A80 (June 2011), A80. DOI: 10.1051/0004-6361/201116681. arXiv: 1104.1577 [astro-ph.IM].
- [171] M. P. van Haarlem et al. “LOFAR: The LOw-Frequency ARray”. In: A&A 556, A2 (Aug. 2013), A2. DOI: 10.1051/0004-6361/201220873. arXiv: 1305.3550 [astro-ph.IM].
- [172] C. G. Bassa, Z. Pleunis, and J. W. T. Hessels. “Enabling pulsar and fast transient searches using coherent dedispersion”. In: *Astronomy and Computing* 18 (Jan. 2017), pp. 40–46. DOI: 10.1016/j.ascom.2017.01.004. arXiv: 1607.00909 [astro-ph.IM].
- [173] W. Atwood et al. “Pass 8: Toward the Full Realization of the Fermi-LAT Scientific Potential”. In: *arXiv e-prints*, arXiv:1303.3514 (2013), arXiv:1303.3514. arXiv: 1303.3514 [astro-ph.IM].
- [174] B. J. Prager et al. “Using Long-term Millisecond Pulsar Timing to Obtain Physical Characteristics of the Bulge Globular Cluster Terzan 5”. In: ApJ 845, 148 (Aug. 2017), p. 148. DOI: 10.3847/1538-4357/aa7ed7. arXiv: 1612.04395 [astro-ph.SR].
- [175] D. Mukherjee. “Revisiting Field Burial by Accretion onto Neutron Stars”. In: *Journal of Astrophysics and Astronomy* 38, 48 (Sept. 2017), p. 48. DOI: 10.1007/s12036-017-9465-6. arXiv: 1709.07332 [astro-ph.HE].
- [176] R. P. Breton et al. “Discovery of the Optical Counterparts to Four Energetic Fermi Millisecond Pulsars”. In: ApJ 769, 108 (June 2013), p. 108. DOI: 10.1088/0004-637X/769/2/108. arXiv: 1302.1790 [astro-ph.HE].

- [177] M. S. E. Roberts et al. “Intrabinary shock emission from “black widows“ and “redbacks””. In: *Astronomische Nachrichten* 335.3 (2014), pp. 313–317. DOI: 10.1002/asna.201312038. arXiv: 1402.5507 [astro-ph.HE].
- [178] B. W. Stappers et al. “Probing the Eclipse Region of a Binary Millisecond Pulsar”. In: *ApJ* 465 (July 1996), p. L119. DOI: 10.1086/310148.
- [179] E. J. Polzin et al. “The low-frequency radio eclipses of the black widow pulsar J1810+1744”. In: *MNRAS* 476.2 (2018), pp. 1968–1981. DOI: 10.1093/mnras/sty349. arXiv: 1802.02594 [astro-ph.HE].
- [180] The Fermi-LAT collaboration. “Fermi Large Area Telescope Fourth Source Catalog”. In: *arXiv e-prints*, arXiv:1902.10045 (2019), arXiv:1902.10045. arXiv: 1902.10045 [astro-ph.HE].
- [181] F. Acero et al. “Development of the Model of Galactic Interstellar Emission for Standard Point-source Analysis of Fermi Large Area Telescope Data”. In: *ApJS* 223, 26 (Apr. 2016), p. 26. DOI: 10.3847/0067-0049/223/2/26. arXiv: 1602.07246 [astro-ph.HE].
- [182] X. Hou et al. “Six faint gamma-ray pulsars seen with the Fermi Large Area Telescope. Towards a sample blending into the background”. In: *A&A* 570, A44 (Oct. 2014), A44. DOI: 10.1051/0004-6361/201424294. arXiv: 1407.6271 [astro-ph.HE].
- [183] C. Aulbert and H. Fehrmann. “Gravitationswellensuche mit einem der weltweit größten Supercomputer”. In: *Forschungsbericht 2009 - Max-Planck-Institut für Gravitationsphysik, Teilinstitut Hannover*, <https://www.mpg.de/308429/forschungsSchwerpunkt> (2009).
- [184] G. Schwarz. “Estimating the Dimension of a Model”. In: *Annals of Statistics* 6 (July 1978), pp. 461–464.
- [185] J. Goodman and J. Weare. “Ensemble samplers with affine invariance”. In: *Communications in Applied Mathematics and Computational Science, Vol. 5, No. 1, p. 65-80, 2010* 5 (2010), pp. 65–80. DOI: 10.2140/camcos.2010.5.65.
- [186] D. Foreman-Mackey et al. “emcee: The MCMC Hammer”. In: *PASP* 125 (Mar. 2013), p. 306. DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].
- [187] H. An et al. “High-energy Variability of PSR J1311-3430”. In: *ApJ* 850, 100 (Nov. 2017), p. 100. DOI: 10.3847/1538-4357/aa947f. arXiv: 1710.06097 [astro-ph.HE].
- [188] J. M. Cordes and T. J. W. Lazio. “NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations”. In: *arXiv e-prints*, astro-ph/0207156 (2002), astro-ph/0207156. arXiv: astro-ph/0207156 [astro-ph].
- [189] J. M. Yao, R. N. Manchester, and N. Wang. “A New Electron-density Model for Estimation of Pulsar and FRB Distances”. In: *ApJ* 835, 29 (Jan. 2017), p. 29. DOI: 10.3847/1538-4357/835/1/29. arXiv: 1610.09448.
- [190] W. van Straten and M. Bailes. “DSPSR: Digital Signal Processing Software for Pulsar Astronomy”. In: *PASA* 28 (Jan. 2011), pp. 1–14. DOI: 10.1071/AS10021. arXiv: 1008.3973 [astro-ph.IM].

- [191] A. W. Hotan, W. van Straten, and R. N. Manchester. “PSRCHIVE and PSRFITS: An Open Approach to Radio Pulsar Data Storage and Analysis”. In: PASA 21 (2004), pp. 302–309. DOI: 10.1071/AS04022. eprint: astro-ph/0404549.
- [192] Willem van Straten, Paul Demorest, and Stefan Osłowski. “Pulsar Data Analysis with PSRCHIVE”. In: *Astronomical Research and Technology* 9.3 (2012), pp. 237–256. arXiv: 1205.6276 [astro-ph.IM].
- [193] W. Coles et al. “Pulsar timing analysis in the presence of correlated noise”. In: MNRAS 418.1 (2011), pp. 561–570. DOI: 10.1111/j.1365-2966.2011.19505.x. arXiv: 1107.5366 [astro-ph.IM].
- [194] R. P. Breton et al. “KOI 1224: A Fourth Bloated Hot White Dwarf Companion Found with Kepler”. In: ApJ 748, 115 (Apr. 2012), p. 115. DOI: 10.1088/0004-637X/748/2/115. arXiv: 1109.6847 [astro-ph.SR].
- [195] V. S. Dhillon et al. “ULTRACAM: an ultrafast, triple-beam CCD camera for high-speed astrophysics”. In: MNRAS 378 (July 2007), pp. 825–840. DOI: 10.1111/j.1365-2966.2007.11881.x. arXiv: 0704.2557.
- [196] V. S. Dhillon et al. “HiPERCAM: a high-speed quintuple-beam CCD camera for the study of rapid variability in the universe”. In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Vol. 9908. Proc. SPIE. Aug. 2016, 99080Y. DOI: 10.1117/12.2229055. arXiv: 1606.09214 [astro-ph.IM].
- [197] Vik Dhillon et al. “First light with HiPERCAM on the GTC”. In: *Ground-based and Airborne Instrumentation for Astronomy VII*. Vol. 10702. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. 2018, p. 107020L. DOI: 10.1117/12.2312041. arXiv: 1807.00557 [astro-ph.IM].
- [198] R. K. Honeycutt. “CCD ensemble photometry on an inhomogeneous set of exposures”. In: PASP 104 (June 1992), pp. 435–440. DOI: 10.1086/133015.
- [199] K. C. Chambers et al. “The Pan-STARRS1 Surveys”. In: *arXiv e-prints*, arXiv:1612.05560 (2016), arXiv:1612.05560. arXiv: 1612.05560 [astro-ph.IM].
- [200] J. A. Smith et al. “The u’g’r’i’z’ Standard-Star System”. In: AJ 123 (Apr. 2002), pp. 2121–2144. DOI: 10.1086/339311. eprint: astro-ph/0201143.
- [201] F. Feroz et al. “Importance Nested Sampling and the MultiNest Algorithm”. In: *ArXiv e-prints* (June 2013). arXiv: 1306.2144 [astro-ph.IM].
- [202] J. Buchner et al. “X-ray spectral modelling of the AGN obscuring region in the CDFS: Bayesian model selection and catalogue”. In: A&A 564, A125 (Apr. 2014), A125. DOI: 10.1051/0004-6361/201322971. arXiv: 1402.0004 [astro-ph.HE].
- [203] T.-O. Husser et al. “A new extensive library of PHOENIX stellar atmospheres and synthetic spectra”. In: A&A 553, A6 (May 2013), A6. DOI: 10.1051/0004-6361/201219058. arXiv: 1303.5632 [astro-ph.SR].
- [204] G. M. Green et al. “Galactic reddening in 3D from stellar photometry - an improved map”. In: MNRAS 478 (July 2018), pp. 651–666. DOI: 10.1093/mnras/sty1008. arXiv: 1801.03555.

- [205] Nicolas Sanchez and Roger W. Romani. “B-ducted Heating of Black Widow Companions”. In: *ApJ* 845, 42 (2017), p. 42. DOI: 10.3847/1538-4357/aa7a02. arXiv: 1706.05467 [astro-ph.HE].
- [206] Michele Vallisneri et al. “The LIGO Open Science Center”. In: *Journal of Physics Conference Series*. Vol. 610. Journal of Physics Conference Series. 2015, p. 012021. DOI: 10.1088/1742-6596/610/1/012021. arXiv: 1410.4839 [gr-qc].
- [207] P. Jaranowski, A. Królak, and B. F. Schutz. “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection”. In: *Phys. Rev. D* 58.6, 063001 (Sept. 1998), p. 063001. DOI: 10.1103/PhysRevD.58.063001. eprint: gr-qc/9804014.
- [208] C. Cutler and B. F. Schutz. “Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars”. In: *Phys. Rev. D* 72.6, 063006 (Sept. 2005), p. 063006. DOI: 10.1103/PhysRevD.72.063006. eprint: gr-qc/0504011.
- [209] B. Abbott et al. “Setting upper limits on the strength of periodic gravitational waves from PSR J1939+2134 using the first science data from the GEO 600 and LIGO detectors”. In: *Phys. Rev. D* 69.8, 082004 (Apr. 2004), p. 082004. DOI: 10.1103/PhysRevD.69.082004. eprint: gr-qc/0308050.
- [210] Craig Cahillane et al. “Calibration uncertainty for Advanced LIGO’s first and second observing runs”. In: *Phys. Rev. D* 96.10, 102001 (2017), p. 102001. DOI: 10.1103/PhysRevD.96.102001. arXiv: 1708.03023 [astro-ph.IM].
- [211] T. Damour and J. H. Taylor. “On the orbital period change of the binary pulsar PSR 1913 + 16”. In: *ApJ* 366 (Jan. 1991), pp. 501–511. DOI: 10.1086/169585.
- [212] D. J. Nice and J. H. Taylor. “PSR J2019+2425 and PSR J2322+2057 and the proper motions of millisecond pulsars”. In: *ApJ* 441 (Mar. 1995), pp. 429–435. DOI: 10.1086/175367.
- [213] I. S. Shklovskii. “Possible Causes of the Secular Increase in Pulsar Periods.” In: *Soviet Ast.* 13 (Feb. 1970), p. 562.
- [214] M. Marelli, A. De Luca, and P. A. Caraveo. “A Multiwavelength Study on the High-energy Behavior of the Fermi/LAT Pulsars”. In: *ApJ* 733, 82 (June 2011), p. 82. DOI: 10.1088/0004-637X/733/2/82. arXiv: 1103.0572 [astro-ph.HE].
- [215] M. Marelli et al. “Radio-quiet and Radio-loud Pulsars: Similar in Gamma-Rays but Different in X-Rays”. In: *ApJ* 802, 78 (Apr. 2015), p. 78. DOI: 10.1088/0004-637X/802/2/78. arXiv: 1501.06215 [astro-ph.HE].
- [216] D. L. Kaplan et al. “A Dense Companion to the Short-period Millisecond Pulsar Binary PSR J0636+5128”. In: *ApJ* 864.1, 15 (2018), p. 15. DOI: 10.3847/1538-4357/aad54c. arXiv: 1807.04610 [astro-ph.HE].
- [217] L. M. van Haaften et al. “The evolution of ultracompact X-ray binaries”. In: *A&A* 537, A104 (Jan. 2012), A104. DOI: 10.1051/0004-6361/201117880. arXiv: 1111.5978 [astro-ph.SR].

- [218] R. W. Romani and N. Sanchez. “Intra-binary Shock Heating of Black Widow Companions”. In: *ApJ* 828, 7 (Sept. 2016), p. 7. DOI: 10.3847/0004-637X/828/1/7. arXiv: 1606.03518 [astro-ph.HE].
- [219] P. Draghis and R. W. Romani. “PSR J0636+5128: A Heated Companion in a Tight Orbit”. In: *ApJ* 862, L6 (July 2018), p. L6. DOI: 10.3847/2041-8213/aad2db. arXiv: 1807.04249 [astro-ph.HE].
- [220] Zorawar Wadiasingh et al. “Constraining Relativistic Bow Shock Properties in Rotation-powered Millisecond Pulsar Binaries”. In: *ApJ* 839.2, 80 (2017), p. 80. DOI: 10.3847/1538-4357/aa69bf. arXiv: 1703.09560 [astro-ph.HE].
- [221] M. R. Kennedy et al. “Kepler K2 observations of the transitional millisecond pulsar PSR J1023+0038”. In: *MNRAS* 477 (June 2018), pp. 1120–1132. DOI: 10.1093/mnras/sty731. arXiv: 1801.10609 [astro-ph.HE].
- [222] Gregory B. Cook, Stuart L. Shapiro, and Saul A. Teukolsky. “Rapidly Rotating Neutron Stars in General Relativity: Realistic Equations of State”. In: *ApJ* 424 (Apr. 1994), p. 823. DOI: 10.1086/173934.
- [223] J. M. Lattimer and M. Prakash. “The Physics of Neutron Stars”. In: *Science* 304 (Apr. 2004), pp. 536–542. DOI: 10.1126/science.1090720. arXiv: astro-ph/0405262 [astro-ph].
- [224] Fabian Gittins and Nils Andersson. “Population Synthesis of Accreting Neutron Stars Emitting Gravitational Waves”. In: *ArXiv e-prints*, arXiv:1811.00550 (Nov. 2018), arXiv:1811.00550. arXiv: 1811.00550 [astro-ph.HE].
- [225] Alessandro Patruno, Brynmor Haskell, and Caroline D’Angelo. “Gravitational Waves and the Maximum Spin Frequency of Neutron Stars”. In: *ApJ* 746, 9 (Feb. 2012), p. 9. DOI: 10.1088/0004-637X/746/1/9. arXiv: 1109.0536 [astro-ph.HE].
- [226] A. Bonanno and V. Urpin. “The accretion rate and minimum spin period of accreting pulsars”. In: *MNRAS* 451 (Aug. 2015), pp. 2117–2122. DOI: 10.1093/mnras/stv1112.
- [227] Y. Y. Pan et al. “The minimum magnetic field of millisecond pulsars calculated according to accretion: application to the X-ray neutron star SAX J1808.4-3658 in a low-mass X-ray binary”. In: *MNRAS* 480 (Oct. 2018), pp. 692–696. DOI: 10.1093/mnras/sty1851.
- [228] B. P. Abbott et al. “Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data”. In: *arXiv e-prints*, arXiv:1902.08507 (2019), arXiv:1902.08507. arXiv: 1902.08507 [astro-ph.HE].
- [229] B. P. Abbott et al. “First Search for Gravitational Waves from Known Pulsars with Advanced LIGO”. In: *ApJ* 839, 12 (Apr. 2017), p. 12. DOI: 10.3847/1538-4357/aa677f. arXiv: 1701.07709 [astro-ph.HE].
- [230] D. A. Frail et al. “Pulsar candidates towards Fermi unassociated sources”. In: *MNRAS* 461 (Sept. 2016), pp. 1062–1067. DOI: 10.1093/mnras/stw1390. arXiv: 1606.03450 [astro-ph.HE].
- [231] Robert J. Adler and Jonathan E. Taylor. *Random Fields and Geometry*. eng. Springer Monographs in Mathematics. New York, NY: Springer New York, 2007. ISBN: 1-281-98633-X.

- [232] O. C. de Jager and I. Büsching. “The H-test probability distribution revisited: improved sensitivity”. In: *A&A* 517, L9 (July 2010), p. L9. DOI: 10.1051/0004-6361/201014362. arXiv: 1005.4867 [astro-ph.HE].
- [233] A. T. Kruger, T. J. Lored, and I. Wasserman. “Search for High-Frequency Periodicities in Time-tagged Event Data from Gamma-Ray Bursts and Soft Gamma Repeaters”. In: *ApJ* 576 (Sept. 2002), pp. 932–941. DOI: 10.1086/341541. eprint: astro-ph/0112192.
- [234] R. C. Hartman et al. “The Third EGRET Catalog of High-Energy Gamma-Ray Sources”. In: *ApJS* 123.1 (July 1999), pp. 79–202. DOI: 10.1086/313231.
- [235] P. Bruel et al. “Fermi-LAT improved Pass⁸ event selection”. In: *arXiv e-prints*, arXiv:1810.11394 (Oct. 2018), arXiv:1810.11394. arXiv: 1810.11394 [astro-ph.IM].
- [236] Jacob T. VanderPlas and Željko Ivezić. “Periodograms for Multiband Astronomical Time Series”. In: *ApJ* 812.1, 18 (Oct. 2015), p. 18. DOI: 10.1088/0004-637X/812/1/18. arXiv: 1502.01344 [astro-ph.IM].
- [237] Paul J. McMillan. “The mass distribution and gravitational potential of the Milky Way”. In: *MNRAS* 465.1 (Feb. 2017), pp. 76–94. DOI: 10.1093/mnras/stw2759. arXiv: 1608.00971 [astro-ph.GA].
- [238] R. P. Breton et al. “Discovery of the Optical Counterparts to Four Energetic Fermi Millisecond Pulsars”. In: *ApJ* 769, 108 (June 2013), p. 108. DOI: 10.1088/0004-637X/769/2/108. arXiv: 1302.1790 [astro-ph.HE].
- [239] Gregory M. Green et al. “A 3D Dust Map Based on Gaia, Pan-STARRS 1, and 2MASS”. In: *ApJ* 887.1, 93 (Dec. 2019), p. 93. DOI: 10.3847/1538-4357/ab5362. arXiv: 1905.02734 [astro-ph.GA].
- [240] D. Kandel and Roger W. Romani. “Atmospheric Circulation on Black Widow Companions”. In: *ApJ* 892.2, 101 (Apr. 2020), p. 101. DOI: 10.3847/1538-4357/ab7b62. arXiv: 2002.12483 [astro-ph.HE].
- [241] Di Li et al. “FAST in Space: Considerations for a Multibeam, Multipurpose Survey Using China’s 500-m Aperture Spherical Radio Telescope (FAST)”. In: *IEEE Microwave Magazine* 19.3 (Apr. 2018), pp. 112–119. DOI: 10.1109/MMM.2018.2802178. arXiv: 1802.03709 [astro-ph.IM].
- [242] Siraprapa Sanpa-arsa. “Searching for New Millisecond Pulsars with the GBT in Fermi Unassociated Sources”. PhD thesis. Graduate School of Arts and Sciences, University of Virginia, 2016. DOI: 10.18130/V36K7P.
- [243] Gregory Desvignes et al. “SPAN512: A new mid-latitude pulsar survey with the Nançay Radio Telescope”. In: *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*. Ed. by Joeri van Leeuwen. Vol. 291. IAU Symposium. Mar. 2013, pp. 375–377. DOI: 10.1017/S1743921312024179. arXiv: 1211.3936 [astro-ph.HE].
- [244] F. Camilo et al. “Discovery of a Millisecond Pulsar in the 5.4 day Binary 3FGL J1417.5-4402: Observing the Late Phase of Pulsar Recycling”. In: *ApJ* 820, 6 (Mar. 2016), p. 6. DOI: 10.3847/0004-637X/820/1/6. arXiv: 1603.00142 [astro-ph.SR].

- [245] The LIGO Scientific Collaboration et al. “Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo”. In: *arXiv e-prints*, arXiv:1912.11716 (Dec. 2019), arXiv:1912.11716. arXiv: 1912.11716 [gr-qc].
- [246] D. I. Jones. “Gravitational wave emission from rotating superfluid neutron stars”. In: MNRAS 402.4 (Mar. 2010), pp. 2503–2519. DOI: 10.1111/j.1365-2966.2009.16059.x. arXiv: 0909.4035 [astro-ph.SR].
- [247] D. I. Jones. “Parameter choices and ranges for continuous gravitational wave searches for steadily spinning neutron stars”. In: MNRAS 453.1 (Oct. 2015), pp. 53–66. DOI: 10.1093/mnras/stv1584. arXiv: 1501.05832 [gr-qc].
- [248] Réjean J. Dupuis and Graham Woan. “Bayesian estimation of pulsar parameters from gravitational wave data”. In: Phys. Rev. D 72.10, 102002 (Nov. 2005), p. 102002. DOI: 10.1103/PhysRevD.72.102002. arXiv: gr-qc/0508096 [gr-qc].
- [249] Matthew Pitkin et al. “A nested sampling code for targeted searches for continuous gravitational waves from pulsars”. In: *arXiv e-prints*, arXiv:1705.08978 (May 2017), arXiv:1705.08978. arXiv: 1705.08978 [gr-qc].
- [250] M. Pitkin et al. “First results and future prospects for dual-harmonic searches for gravitational waves from spinning neutron stars”. In: MNRAS 453.4 (Nov. 2015), pp. 4399–4420. DOI: 10.1093/mnras/stv1931. arXiv: 1508.00416 [astro-ph.HE].
- [251] B. P. Abbott et al. “Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run”. In: Phys. Rev. D 99.12, 122002 (June 2019), p. 122002. DOI: 10.1103/PhysRevD.99.122002. arXiv: 1902.08442 [gr-qc].
- [252] B. P. Abbott et al. “Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data”. In: ApJ 879.1, 10 (July 2019), p. 10. DOI: 10.3847/1538-4357/ab20cb. arXiv: 1902.08507 [astro-ph.HE].
- [253] E. J. Polzin et al. “Study of spider pulsar binary eclipses and discovery of an eclipse mechanism transition”. In: MNRAS 494.2 (Mar. 2020), pp. 2948–2968. DOI: 10.1093/mnras/staa596. arXiv: 2003.02335 [astro-ph.HE].
- [254] J. W. Broderick et al. “Low-radio-frequency eclipses of the redback pulsar J2215+5135 observed in the image plane with LOFAR”. In: MNRAS 459.3 (July 2016), pp. 2681–2689. DOI: 10.1093/mnras/stw794. arXiv: 1604.05722 [astro-ph.HE].
- [255] P. P. Eggleton. “Approximations to the radii of Roche lobes.” In: ApJ 268 (May 1983), pp. 368–369. DOI: 10.1086/160960.
- [256] Rahul Sengar et al. “Novel modelling of ultracompact X-ray binary evolution - stable mass transfer from white dwarfs to neutron stars”. In: MNRAS 470.1 (Sept. 2017), pp. L6–L10. DOI: 10.1093/mnrasl/slx064. arXiv: 1704.08260 [astro-ph.SR].
- [257] L. M. van Haften et al. “The evolution of ultracompact X-ray binaries”. In: A&A 537, A104 (Jan. 2012), A104. DOI: 10.1051/0004-6361/201117880. arXiv: 1111.5978 [astro-ph.SR].

- [258] P. S. Ray et al. “Radio Detection of the Fermi-LAT Blind Search Millisecond Pulsar J1311-3430”. In: *ApJ* 763, L13 (Jan. 2013), p. L13. DOI: 10.1088/2041-8205/763/1/L13. arXiv: 1210.6676 [astro-ph.HE].
- [259] Scott M. Ransom, Stephen S. Eikenberry, and John Middleditch. “Fourier Techniques for Very Long Astrophysical Time-Series Analysis”. In: *AJ* 124.3 (Sept. 2002), pp. 1788–1809. DOI: 10.1086/342285. arXiv: astro-ph/0204349 [astro-ph].
- [260] A. A. Abdo et al. “A Population of Gamma-Ray Millisecond Pulsars Seen with the Fermi Large Area Telescope”. In: *Science* 325 (Aug. 2009), p. 848. DOI: 10.1126/science.1176113.
- [261] A. D. van Staden and J. Antoniadis. “An Active, Asynchronous Companion to a Redback Millisecond Pulsar”. In: *ApJ* 833, L12 (Dec. 2016), p. L12. DOI: 10.3847/2041-8213/833/1/L12. arXiv: 1609.08024 [astro-ph.SR].
- [262] Patricia B. Cho, Jules P. Halpern, and Slavko Bogdanov. “Variable Heating and Flaring of Three Redback Millisecond Pulsar Companions”. In: *ApJ* 866.1, 71 (2018), p. 71. DOI: 10.3847/1538-4357/aade92. arXiv: 1809.00215 [astro-ph.HE].
- [263] Paul Draghis et al. “Multiband Optical Light Curves of Black-Widow Pulsars”. In: *arXiv e-prints*, arXiv:1908.00992 (2019), arXiv:1908.00992. arXiv: 1908.00992 [astro-ph.HE].
- [264] B. Knispel et al. “Pulsar Discovery by Global Volunteer Computing”. In: *Science* 329 (Sept. 2010), p. 1305. DOI: 10.1126/science.1195253. arXiv: 1008.2172 [astro-ph.GA].
- [265] J. Greiner et al. “GROND—a 7-Channel Imager”. In: *PASP* 120.866 (Apr. 2008), p. 405. DOI: 10.1086/587032. arXiv: 0801.4801 [astro-ph].
- [266] A. J. Drake et al. “The Catalina Surveys Southern periodic variable star catalogue”. In: *MNRAS* 469 (Aug. 2017), pp. 3688–3712. DOI: 10.1093/mnras/stx1085.
- [267] C. W. Ng et al. “Evidence on the Orbital Modulated Gamma-Ray Emissions from the Redback Candidate 3FGL J2039.6-5618”. In: *ApJ* 867.2, 90 (Nov. 2018), p. 90. DOI: 10.3847/1538-4357/aae308.
- [268] J. Christopher Clemens, J. Adam Crain, and Robert Anderson. “The Goodman spectrograph”. In: *Ground-based Instrumentation for Astronomy*. Ed. by Alan F. M. Moorwood and Masanori Iye. Vol. 5492. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Sept. 2004, pp. 331–340. DOI: 10.1117/12.550069.
- [269] Gaia Collaboration et al. “The Gaia mission”. In: *A&A* 595, A1 (Nov. 2016), A1. DOI: 10.1051/0004-6361/201629272. arXiv: 1609.04153 [astro-ph.IM].
- [270] C. Jordi et al. “Gaia broad band photometry”. In: *A&A* 523, A48 (Nov. 2010), A48. DOI: 10.1051/0004-6361/201015441. arXiv: 1008.0815 [astro-ph.IM].
- [271] M. Wood et al. “Fermipy: An open-source Python package for analysis of Fermi-LAT Data”. In: *International Cosmic Ray Conference* 35, 824 (Jan. 2017), p. 824. arXiv: 1707.09551 [astro-ph.IM].

- [272] Lehel Csató and Manfred Opper. “Sparse On-Line Gaussian Processes”. In: *Neural Computation* 14.3 (2002), pp. 641–668. DOI: 10.1162/089976602317250933. eprint: <https://doi.org/10.1162/089976602317250933>. URL: <https://doi.org/10.1162/089976602317250933>.
- [273] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN: 026218253X.
- [274] D. Seiferth et al. “Online Gaussian Process regression with non-Gaussian likelihood”. In: *2017 American Control Conference (ACC)*. 2017, pp. 3134–3140. DOI: 10.23919/ACC.2017.7963429.
- [275] Jing Luo et al. “PINT, A Modern Software Package for Pulsar Timing”. In: *American Astronomical Society Meeting Abstracts #231*. Vol. 231. American Astronomical Society Meeting Abstracts. Jan. 2018, p. 453.09.
- [276] Diego F. Torres et al. “A Search for Transitions between States in Redbacks and Black Widows Using Seven Years of Fermi-LAT Observations”. In: *ApJ* 836, 68 (Feb. 2017), p. 68. DOI: 10.3847/1538-4357/836/1/68. arXiv: 1612.07083 [astro-ph.HE].
- [277] Jay Strader et al. “A New γ -Ray Loud, Eclipsing Low-mass X-Ray Binary”. In: *ApJ* 831.1, 89 (2016), p. 89. DOI: 10.3847/0004-637X/831/1/89. arXiv: 1608.02583 [astro-ph.HE].
- [278] M. R. Kennedy et al. “Optical, X-ray, and γ -ray observations of the candidate transitional millisecond pulsar 4FGL J0427.8-6704”. In: *MNRAS* 494.3 (Apr. 2020), pp. 3912–3926. DOI: 10.1093/mnras/staa912. arXiv: 2003.13718 [astro-ph.HE].
- [279] M. Kerr. “Multiscale Time- and Frequency-domain Likelihood Analysis with Photon Weights”. In: *ApJ* 885.1, 92 (2019), p. 92. DOI: 10.3847/1538-4357/ab459f. arXiv: 1910.00140 [astro-ph.HE].
- [280] E. M. H. Wu et al. “Orbital-phase-dependent γ -Ray Emissions from the Black Widow Pulsar”. In: *ApJ* 761, 181 (Dec. 2012), p. 181. DOI: 10.1088/0004-637X/761/2/181. arXiv: 1210.7209 [astro-ph.HE].
- [281] Hongjun An, Roger W. Romani, and Matthew Kerr. “Signatures of Intra-binary Shock Emission in the Black Widow Pulsar Binary PSR J2241-5236”. In: *ApJ* 868.1, L8 (2018), p. L8. DOI: 10.3847/2041-8213/aaedaf. arXiv: 1810.13055 [astro-ph.HE].
- [282] Hongjun An et al. “Orbital Modulation of Gamma Rays from PSR J2339-0533”. In: *ApJ* 897.1, 52 (July 2020), p. 52. DOI: 10.3847/1538-4357/ab93ba. arXiv: 2005.07888 [astro-ph.HE].
- [283] Constantinos Kalapotharakos et al. “A Fundamental Plane for Gamma-Ray Pulsars”. In: *ApJ* 883.1, L4 (Sept. 2019), p. L4. DOI: 10.3847/2041-8213/ab3e0a. arXiv: 1904.01765 [astro-ph.HE].
- [284] Zorawar Wadiasingh et al. “Pressure Balance and Intra-binary Shock Stability in Rotation-powered-state Redback and Transitional Millisecond Pulsar Binary Systems”. In: *ApJ* 869.2, 120 (Dec. 2018), p. 120. DOI: 10.3847/1538-4357/aaed43. arXiv: 1810.12958 [astro-ph.HE].

- [285] Arne A. Henden et al. “APASS Data Release 10”. In: *American Astronomical Society Meeting Abstracts #232*. Vol. 232. American Astronomical Society Meeting Abstracts. 2018, p. 223.06.
- [286] L. B. Lucy. “Gravity-Darkening for Stars with Convective Envelopes”. In: *ZAp* 65 (1967), p. 89.
- [287] Peter H. Hauschildt, France Allard, and E. Baron. “The NextGen Model Atmosphere Grid for $3000 < T_{eff} < 10,000$ K”. In: *ApJ* 512.1 (1999), pp. 377–385. DOI: 10.1086/306745. arXiv: astro-ph/9807286 [astro-ph].
- [288] X. Luri et al. “Gaia Data Release 2. Using Gaia parallaxes”. In: *A&A* 616, A9 (Aug. 2018), A9. DOI: 10.1051/0004-6361/201832964. arXiv: 1804.09376 [astro-ph.IM].
- [289] L. Levin et al. “The High Time Resolution Universe Pulsar Survey -VIII. The Galactic millisecond pulsar population”. In: *MNRAS* 434 (Sept. 2013), pp. 1387–1397. DOI: 10.1093/mnras/stt1103. arXiv: 1306.4190 [astro-ph.SR].
- [290] G. Voisin et al. “A model for redistributing heat over the surface of irradiated spider companions”. In: *arXiv e-prints*, arXiv:2006.06552 (June 2020), arXiv:2006.06552. arXiv: 2006.06552 [astro-ph.HE].
- [291] Feryal Özel and Paulo Freire. “Masses, Radii, and the Equation of State of Neutron Stars”. In: *ARA&A* 54 (2016), pp. 401–440. DOI: 10.1146/annurev-astro-081915-023322. arXiv: 1603.02698 [astro-ph.HE].
- [292] M. H. van Kerkwijk, R. P. Breton, and S. R. Kulkarni. “Evidence for a Massive Neutron Star from a Radial-velocity Study of the Companion to the Black-widow Pulsar PSR B1957+20”. In: *ApJ* 728.2, 95 (2011), p. 95. DOI: 10.1088/0004-637X/728/2/95. arXiv: 1009.5427 [astro-ph.HE].
- [293] M. Linares, T. Shahbaz, and J. Casares. “Peering into the Dark Side: Magnesium Lines Establish a Massive Neutron Star in PSR J2215+5135”. In: *ApJ* 859.1, 54 (2018), p. 54. DOI: 10.3847/1538-4357/aabde6. arXiv: 1805.08799 [astro-ph.HE].
- [294] Allison M. Matthews et al. “The NANOGrav Nine-year Data Set: Astrometric Measurements of 37 Millisecond Pulsars”. In: *ApJ* 818.1, 92 (Feb. 2016), p. 92. DOI: 10.3847/0004-637X/818/1/92. arXiv: 1509.08982 [astro-ph.GA].
- [295] D. J. K. O’Connell. “The so-called periastron effect in close eclipsing binaries ; New variable stars (fifth list)”. In: *Publications of the Riverview College Observatory* 2 (Aug. 1951), pp. 85–100.
- [296] T. J. Martin and S. C. Davey. “Application of smoothed particle hydrodynamics to atmospheric motions in interacting binary systems”. In: *MNRAS* 275.1 (1995), pp. 31–42. DOI: 10.1093/mnras/275.1.31.
- [297] C. A. Hill et al. “Roche tomography of cataclysmic variables - VI. Differential rotation of AE Aqr - not tidally locked!” In: *MNRAS* 444.1 (Oct. 2014), pp. 192–207. DOI: 10.1093/mnras/stu1460. arXiv: 1407.6148 [astro-ph.SR].
- [298] Svetlana V. Berdyugina. “Starspots: A Key to the Stellar Dynamo”. In: *Living Reviews in Solar Physics* 2.1, 8 (2005), p. 8. DOI: 10.12942/lrsp-2005-8.

- [299] O. McConnell et al. “Roche lobe underfilling of the secondary star in PSR J102347.6+003841?” In: MNRAS 451.4 (2015), pp. 3468–3472. DOI: 10.1093/mnras/stv1197. arXiv: 1507.07710 [astro-ph.HE].
- [300] T. Shahbaz et al. “The binary millisecond pulsar PSR J1023+0038 - II. Optical spectroscopy”. In: MNRAS 488.1 (2019), pp. 198–212. DOI: 10.1093/mnras/stz1652. arXiv: 1906.04524 [astro-ph.SR].
- [301] T. E. Sterne. “Apsidal motion in binary stars”. In: MNRAS 99 (1939), pp. 451–462. DOI: 10.1093/mnras/99.5.451.
- [302] Gordon I. Ogilvie. “Tidal Dissipation in Stars and Giant Planets”. In: ARA&A 52 (2014), pp. 171–210. DOI: 10.1146/annurev-astro-081913-035941. arXiv: 1406.2207 [astro-ph.SR].
- [303] J. U. Cisneros-Parra. “Apsidal motion in close binaries with and without mass exchange.” In: A&A 8 (1970), p. 141.
- [304] E. J. Polzin et al. “Long-term variability of a black widow’s eclipses - A decade of PSR J2051-0827”. In: MNRAS 490.1 (Nov. 2019), pp. 889–908. DOI: 10.1093/mnras/stz2579. arXiv: 1909.06130 [astro-ph.HE].
- [305] J. P. Zahn. “Reprint of 1977A&A...57..383Z. Tidal friction in close binary stars.” In: A&A 500 (1977), pp. 121–132.
- [306] S. M. Ransom et al. “A millisecond pulsar in a stellar triple system”. In: Nature 505.7484 (Jan. 2014), pp. 520–524. DOI: 10.1038/nature12917. arXiv: 1401.0535 [astro-ph.SR].
- [307] Anne M. Archibald et al. “Universality of free fall from the orbital motion of a pulsar in a stellar triple system”. In: Nature 559.7712 (July 2018), pp. 73–76. DOI: 10.1038/s41586-018-0265-1. arXiv: 1807.02059 [astro-ph.HE].
- [308] S. E. Thorsett, Z. Arzoumanian, and J. H. Taylor. “PSR B1620-26: A Binary Radio Pulsar with a Planetary Companion?” In: ApJ 412 (July 1993), p. L33. DOI: 10.1086/186933.
- [309] Steinn Sigurdsson et al. “A Young White Dwarf Companion to Pulsar B1620-26: Evidence for Early Planet Formation”. In: *Science* 301.5630 (July 2003), pp. 193–196. DOI: 10.1126/science.1086326. arXiv: astro-ph/0307339 [astro-ph].
- [310] R. M. Shannon et al. “An Asteroid Belt Interpretation for the Timing Variations of the Millisecond Pulsar B1937+21”. In: ApJ 766.1, 5 (Mar. 2013), p. 5. DOI: 10.1088/0004-637X/766/1/5. arXiv: 1301.6429 [astro-ph.SR].

Acknowledgments

I would like to thank a lot of people for their help, support, and advice throughout the time of this thesis. It has been exciting to work on this project and I am looking forward to continuing the research with these people.

I am very grateful to Bruce Allen for being my advisor, for leading such a great institute and group, and for being so excited about science. I enjoyed the long discussions about photon weights and test statistics, and I think that these discussions really improved the thesis. I would also like to thank Holger Pletsch for being my supervisor for the first half year of this thesis.

I would also like to thank Colin Clark. He introduced me to the gamma-ray pulsar group after supervising a seminar talk of mine. During my master thesis we was the person to answer all of my questions, and that continued throughout the PhD thesis, especially for the first two years when we shared the office. After moving to Manchester, Colin worked on optical analyses of spider companions and provided the orbital constraints leading to the two *Einstein@Home* discoveries. I am really looking forward to our study of some nasty reback pulsars.

I would like to thank the *Einstein@Home* and ATLAS staff for keeping these two amazing supercomputers running. Especially, I would like to thank Bernd, Carsten, Henning, and Oliver for their help whenever a problem occurred and the discussions about new projects.

I am grateful to the people in the *Fermi*-LAT collaboration for being so nice and curious towards new members. While there are many more people to mention within the collaboration, I would like to especially thank a few close collaborators: David, Elizabeth, Jason, Lucas, Matthew, Paul, and Pablo. I would also like to thank the many collaborators who contributed to the papers – I am very grateful for your contributions!

It has been a pleasure to work on a PhD thesis at this institute in part because the people are all open, kind and in a good way crazy. It has been fun with all of you to go on lecture weeks, christmas parties, defenses, and to play the football faculty cup!

Finally, I would like to thank my friends and family for their love, support, and encouragement throughout the years.

Personal Information

Name: Lars Nieder
Date of Birth: 15.12.1989
Place of Birth: Hannover (Germany)

Education

- 2015–2020 **PhD studies in Physics**, *Leibniz Universität Hannover*
- Doctoral thesis at the Max Planck Institute for Gravitational Physics, supervisor Prof. Dr. Allen: “Binary Gamma-ray Pulsars”
 - Affiliated member of the *Fermi*-LAT collaboration (since 2016)
 - PhD student representative (2016–2019)
 - Supervisor of student lab “Data Analysis” (2017–2018)
 - Correction of student exercises and exams (2016–2017)
- 2013–2015 **Master studies in Physics**, *Leibniz Universität Hannover*
- Master thesis at the Max Planck Institute for Gravitational Physics, supervisor Dr. Pletsch: “Method Development and Search for an unknown Gamma-ray Pulsar in an Eccentric Binary Orbit”
- 2013–2013 **Exchange semester**, *Chalmers University of Technology*, Gothenburg
- 2009–2013 **Bachelor studies in Physics**, *Leibniz Universität Hannover*
- Bachelor thesis at the Institute for Theoretical Physics, supervisor Prof. Dr. Giuliani: “Estimates on the Influence of the Solar Corona on Light Propagation in view of proposed Shapiro Time-Delay Experiments”
- 2002–2009 **Abitur**, *Humboldtschule*, Hannover

Publication List

- Nieder, L., Clark, C. J., Kandel, D., et al., 2020, *Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated Einstein@Home*, accepted for publication in The Astrophysical Journal Letters.
- Corongiu, A., Mignani, R. P., Seyffert, A. S., et al., 2020, *Radio Pulsations from the Gamma-ray Millisecond Pulsar PSR J2039–5617*, submitted to Monthly Notices of the Royal Astronomical Society.
- Clark, C. J., Nieder, L., Voisin, G., et al., 2020, *Einstein@Home Discovery of the Gamma-ray Millisecond Pulsar PSR J2039–5617 Confirms its Predicted Redback Nature*, accepted for publication in Monthly Notices of the Royal Astronomical Society.
- Nieder, L., Allen, B., Clark, C. J., Pletsch, H. J., 2020, *Exploiting Orbital Constraints from Optical Data to Detect Binary Gamma-ray Pulsars*, The Astrophysical Journal, 901, 2.
- Nieder, L., Clark, C. J., Bassa, C. G., et al., 2019, *Detection and Timing of Gamma-ray Pulsations from the 707 Hz Pulsar J0952–0607*, The Astrophysical Journal, 883, 42.
- Clark, C. J., Pletsch, H. J., Wu, J., et al., 2018, *Einstein@Home Discovers a Radio-quiet Gamma-ray Millisecond Pulsar*, Science Advances, 4, eaao7228.
- Bassa, C. G., Pleunis, Z., Hessels, J. W. T., et al., 2017, *LOFAR Discovery of the Fastest-spinning Millisecond Pulsar in the Galactic Field*, The Astrophysical Journal Letters, 846, L20.
- 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.