Large-Scale Surveys for Continuous Gravitational Waves: From Data Preparation to Multi-Stage Hierarchical Follow-Ups

Von der QUEST-Leibniz-Forschungsschule
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
Dr. rer. nat.

genehmigte Dissertation von

M.Sc. Benjamin Steltner

2023

Referent:

Prof. Dr. Maria Alessandra Papa

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Hannover, Germany
University of Wisconsin Milwaukee
Leibniz Universität Hannover, Hannover, Germany

Korreferenten:

Prof. Dr. Badri Krishnan

Institute for Mathematics, Astrophysics and Particle Physics, Radboud University, Nijmegen, The Netherlands
Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Hannover, Germany
Leibniz Universität Hannover, Hannover, Germany

Asst. Prof. Dr. Irene Di Palma

Università di Roma "La Sapienza", Roma, Italy

Abstract

The gravitational wave event GW150914 was the first direct detection of gravitational waves roughly 100 years after their prediction by Albert Einstein. The detection was a breakthrough, opening another channel to observe the Universe. Since then over 90 detections of merging compact objects have been made, most of them coalescences of binary black holes of different masses. There have been two black hole-neutron star, and two binary neutron-star mergers. Another breakthrough was the first binary neutron-star merger, GW170817, associated with a slew of electromagnetic observations, including a gamma-ray burst $1.7\,$s after the merger.

Compact binary coalescence events are cataclysmic events in which multiple solar masses are emitted in gravitational waves in $\sim$ seconds. Still, their gravitational wave detection requires sophisticated measuring devices: kilometer-scale laser interferometers.

Another not yet detected form of gravitational radiation are continuous gravitational waves from e.g., but not limited to, fast-spinning neutron stars nonaxisymmetric relatively to their rotational axis. The gravitational wave amplitude on Earth is orders of magnitude weaker than the compact binary coalescence events, but, in the case of the nonaxisymmetric neutron star, is emitted as long as the neutron star is spinning and sustaining the deformation, which may be months to years.

The gravitational wave is mostly emitted at twice the rotational frequency, with a possible frequency evolution (spin-down) due to the energy emitted by gravitational waves, as well as other braking mechanisms. This nearly monochromatic continuous wave is received by observers on Earth Doppler modulated by Earth's orbit and spin.

Although the waveform is seemingly simple, the detection problem for signals from unknown sources is very challenging. The all-sky search for unknown neutron stars in our galaxy detailed in this work used the volunteer distributed computing project Einstein@Home and the ATLAS supercomputer for several months, taking tens of thousands of total CPU-time years to complete. In this work I describe the full-scale data analysis procedure, including data preparation, search set-up optimization and post-processing of search results, whose design and implementation is the core of my doctoral research work. I also present a number of observational results that demonstrate the real-world application of the methodologies that I designed.

Keywords: neutron stars, continuous gravitational waves, data analysis

Zusammenfassung

Das Gravitationswellenereignis GW150914 war der erste direkte Nachweis von Gravitationswellen rund 100 Jahre nach deren Vorhersage durch Albert Einstein. Die Entdeckung war ein Durchbruch und eröffnete einen weiteren Kanal zur Beobachtung des Universums. Seitdem wurden über 90 weitere verschmelzende kompakte Objekte entdeckt, die meisten binäre schwarze Löcher unterschiedlicher Masse, aber auch zweimal verschmelzende Schwarze Löcher mit Neutronensternen und zwei Verschmelzungen von binären Neutronensternen. Ein weiterer Durchbruch war die Beobachtung der ersten Verschmelzung zweier Neutronensterne, GW170817, die mit einer Reihe von elektromagnetischen Beobachtungen einherging, darunter ein Gammastrahlenausbruch 1.7 s nach der Verschmelzung.

Bei der Verschmelzung kompakter Objekte handelt es sich um kataklysmische Ereignisse, bei denen innerhalb von $\sim$ Sekunden mehrere Sonnenmassen in Form von Gravitationswellen ausgestoßen werden. Ihr Nachweis erfordert jedoch hochentwickelte Messgeräte: Laserinterferometer im Kilometermaßstab.

Eine weitere, noch nicht nachgewiesene Form der Gravitationsstrahlung sind kontinuierliche Gravitationswellen, die z.B., aber nicht nur, von schnell rotierenden Neutronensternen ausgehen, die relativ zu ihrer Rotationsachse nicht achsensymmetrisch sind. Die Amplitude der kontinuierlichen Gravitationswellen auf der Erde ist um Größenordnungen schwächer als die der verschmelzenden kompakten Objekte, wird aber im Fall des nicht achsensymmetrischen Neutronensterns so lange abgestrahlt, wie der Neutronenstern rotiert und die Deformation aufrechterhält, was Monate bis Jahre sein können.

Die Gravitationswelle wird meist mit der doppelten Rotationsfrequenz ausgestrahlt, wobei eine Frequenzentwicklung (Spin-down) aufgrund der von Gravitationswellen ausgesandten Energie, sowie anderer Bremsmechanismen möglich ist. Diese nahezu monochromatische, kontinuierliche Welle wird von einem Beobachter auf der Erde Doppler-moduliert durch die Erdumlaufbahn und die Erddrehung empfangen.

Obwohl die Wellenform scheinbar einfach ist, ist das Problem des Nachweises von Signalen aus unbekannten Quellen eine große Herausforderung. Die in dieser Arbeit beschriebene Suche nach unbekannten Neutronensternen in unserer Galaxie über den kompletten Himmel verwendete über mehrere Monate hinweg das Volunteer-Computing-Projekt Einstein@Home und den ATLAS-Supercomputer und benötigte insgesamt Zehntausende von Jahren an Rechenzeit. In dieser Arbeit beschreibe ich das vollständige Datenanalyseverfahren einschließlich der Datenvorbereitung, der Optimierung der Suchparameter und der Nachbearbeitung der Suchergebnisse, dessen Entwurf und Implementierung das Kernstück meiner Doktorarbeit darstellt. Außerdem stelle ich eine Reihe von Beobachtungsergebnissen vor, welche die praktische Anwendung der von mir entwickelten Methoden demonstrieren.

Schlagworte: Neutronensterne, kontinuierliche Gravitationswellen, Datenanalyse

Contents

Introduction

The works in this thesis aim to find continuous gravitational waves. Large scale surveys for continuous gravitational waves require a phenomenal amount of computing power: $\approx$ PFLOPS (10^{15} floating-point operations per second) sustained 24/7 over multiple months. For this a given data set is usually only searched once. It is therefore of outmost importance to optimize that single search in every possible way.

A lot of my research has been exactly this: finding ways to make the searches more sensitive. The adaptive data conditioning methodologies described in Chapter 2, the novel candidate clustering methodologies of Chapter 3 and the multi-stage all-sky hierarchical search of Chapter 4, all aim at improving the chances of detecting a continuous wave signal.

All publications are given as self-contained original manuscripts, i.e. as published by either the American Physical Society or the American Astronomical Society. Therefore the references in each article refer to the articles' equations, figures, tables and bibliography, i.e. not those of this thesis.

Chapter 10 concludes this thesis with a summary and discussion of the results presented in this thesis, and an outlook on the future.

In this chapter the background to this thesis is presented.

1.1 General Relativity and Gravitational Waves

In 1687 Isaac Netwon published his Philosophiæ Naturalis Principia Mathematica (Mathematical Principles of Natural Philosophy) establishing classical mechanics, and among others, Newton's law of gravitation. However, certain effects like the small deviations in Mercury's orbit (the perihelion precession), could not be explained by it.

In 1915 Albert Einstein published his groundbreaking General Theory of Relativity [43], replacing Newtonian gravity in the regime of large masses and relativistic velocities. Instead of a gravitational force attracting two bodies towards each other, Einstein postulated that gravity is rather an effect of a curved spacetime. Masses follow the curvature of spacetime, and masses (rather energy and momentum, to be precise) curve spacetime. Over 100 years of extraordinary scrutiny have been applied to this theory, but so far no experiment has proven Einstein wrong [10, 59, 103].

One prediction of General Relativity are gravitational waves (GWs) – solutions of Einstein's equations. Even before Einstein's theory of General Relativity an analogous wave to electromagnetic waves was considered for gravity. In Newtonian

Gravity the force acts instantly between two masses, indifferent to their distance. With General Relativity information of local changes in gravity propagates at finite speed, the speed of light, mediated by gravitational waves. The existence of gravitational waves, however, was controversial at first, with Einstein changing his stance on their existence multiple times [34].

There were also doubts whether gravitational waves could ever be measured due to their tiny effects [34]. Like electromagnetic waves, gravitational waves are also transverse: they stretch and compress spacetime perpendicular to their travel direction [89] (see Fig. 1.1). They have two polarizations, the "plus" + and "cross" × polarization. When a free-falling particle initially at rest is passed by a gravitational wave, it is not accelerated, it stays at rest. The proper distance between two particles, however, can change if a gravitational wave passes. It can be shown that the relative distance change due to a gravitational wave traveling in z-direction, between two particles separated by an initial distance L, is [89]

$$\frac{\Delta L}{L} = h, \tag{1.1}$$

where h is the gravitational wave amplitude. The largest gravitational wave amplitudes received on Earth are on the order of $h \lesssim 10^{-21}$ or way less [87], so all in all the effective distance change is extremely small: $\sim 0.15\,\mathrm{nm}$ for a test particle on Earth and one on the Sun, or $\sim 0.04\,\mathrm{mm}$ between Earth and the nearest star, Alpha Centauri.
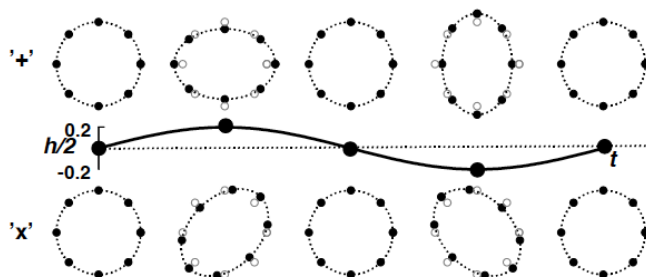

Figure 1.1: Effect of a gravitational wave traveling in z-direction on a circle of free-falling particles in the x-y-plane. The upper row shows the + polarization, the lower the × polarization and in the middle the associated amplitude of the wave is shown (Credit: Fig. 1.1 in [89]).

This effect is used in laser interferometer gravitational wave detectors, such as the two Advanced LIGO (Laser Interferometer Gravitational Wave Observatory) detectors. These are highly advanced Michelson interferometers: a laser shines a beam of light onto a beam splitter, which splits it 50/50 in two 4 kilometers long arms. At the end of the arms a mirror reflects the light back to the beam splitter, where it interferes and is split again, to a photodetector. Depending on the relative phase between the two arms, it interferes constructively or destructively, leading to more or less optical power on the photodetector. A (sufficiently well oriented) gravitational wave compresses one arm, and stretches the other. The relative light travel times of the beams in the two arms changes, leading to a different relative phase and thus to a change of optical power on the photodetector. In fact the

gravitational wave is measured by the necessary adjustments to keep a destructive interference at the photodetector, the DC readout scheme [31, 32, 45, 52].

1.2 The first and following gravitational wave detections

Although gravitational waves have long been controversial, today, we do know they exist. The first indirect detection was the observation of the shrinkage of the orbit of the Hulse-Taylor-Pulsar due to the loss of kinetic energy radiated in gravitational waves [101].

The first direct detection of a gravitational wave (GW150914) on the Sept. 14, 2015 by the two Advanced LIGO gravitational wave detectors opened a new window on the Universe [4]. Colloquially sometimes referred to as "listening" to the Universe due to the frequencies of the observed signals being in the "audio-band", the metaphor of a new sense is quite fitting: gravitational waves bring new insights into aspects of the Universe not observable by electromagnetic, cosmic ray or neutrino probes. Gravitational waves are one of the pillars of multi-messenger astronomy [7, 8, 9].

The breakthrough detection of GW150914 involved two black holes with masses of $\sim 36M_\odot$ and $\sim 29M_\odot$ ($M_\odot$ solar masses) spiraling into each other and finally merging. This cataclysmic event emitted $\sim 3M_\odot$ in gravitational waves, leaving a final black hole of $62M_\odot$. Black holes are objects so compact, that spacetime is so strongly curved that light cannot escape. They are thus very hard to observe, and GW150914 was the first observation of a black hole [4].

Since then over 90 more gravitational wave detections were made in the first three observation runs (O1, O2, O3) of Advanced LIGO and Advanced Virgo [14, 73]. All of them compact binary coalescence (CBC) events, and while most of them were binary black hole mergers, there have been two binary neutron star mergers (GW170817 [7], GW190425 [11]) and two neutron star-black hole (GW200105 [15], GW200115 [15]) mergers. GW170817 was extraordinary as it was observed on multiple channels, including a gamma-ray burst detected by the Fermi Gamma-Ray Burst Monitor $\sim 1.7\,\mathrm{s}$ after the merger [7, 8]. An extensive observation campaign by a plethora of different observatories all around Earth gained new insights into the so-called kilonova and the formation of heavy elements by r-process nucleosynthesis [8]. Observations of gravitational waves have been used to shed light on stellar evolution, to test General Relativity, and to measure the Hubble constant [5, 6, 10, 14, 42, 97].

1.3 Continuous gravitational waves and neutron stars

Coalescences of compact binaries are not the sole astronomical source of detectable gravitational waves. Among others, continuous gravitational waves are expected, but not yet detected. Probably the most promising source of continuous gravitational waves are nonaxisymmetric fast-spinning neutron stars, which continuously emit gravitational waves as long as the neutron star keeps rotating while supporting the asymmetry, possibly years [57]. The involved masses are much smaller than for
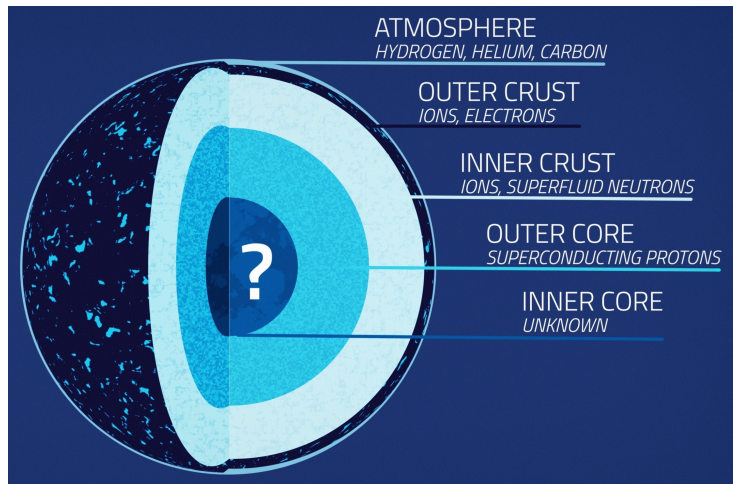
Figure 1.2: Cross-section trough a neutron star (not to scale). Credits: NASA's Goddard Space Flight Center/Conceptual Image Lab.

binary inspirals, making signals from continuous gravitational waves (CWs) orders of magnitude weaker than the violent merger events.

Neutron stars are the remnants of massive stars. They are incredibly dense with a mass of around 1.4 to $2M_\odot$ but a radius of only about $10\,\mathrm{km}$, featuring surface densities of some $\mathrm{g\,cm^{-3}}$ and $\geqslant 10^{15}\mathrm{g\,cm^{-3}}$ at the center of its core [49]. The relation between the pressure and the density is given by the, largely unknown, equation of state (EOS) [49, 60, 74]. A description is difficult since the density spans orders of magnitudes, and in particular, densities not achievable in laboratories and thus of matter in largely unknown states.

The current understanding of the neutron star structure is that the neutron star interior can roughly be described with five distinct regions [49] (see also Fig. 1.2).

- The surface of the neutron star is covered by an atmosphere only a few centimeters thick, likely consisting of the lighter elements hydrogen, helium[90], maybe up to carbon [53]. This is where the thermal radiation detectable by electromagnetic observers is formed [49].

- The outer crust is expected to be around half a kilometer thick, and of roughly equal composition as a white dwarf, with heavy nuclei and a strongly degenerate free electron gas [49].

- The inner crust is made up of neutron-rich atomic nuclei, neutrons and electrons. It is roughly a kilometer thick [49].

- The outer and inner core are several kilometers thick and make up most of the star. The outer core consist of mainly neutrons with some strongly degenerate protons, electrons and maybe muons. The electrons and muons are assumed to be an almost ideal Fermi gas, the neutrons and protons a Fermi liquid [49]. Neutrons and protons could be in a superfluid state and are assumed to be superconductive [63]. The composition of the inner core is largely unknown and very model dependent [49].
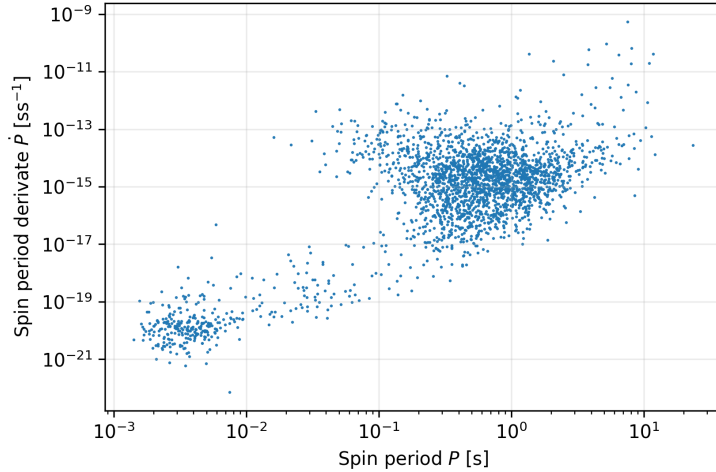
Figure 1.3: The population of known pulsars. The spin period P and spin period derivative $\dot{P}$ is shown (omitting negative $\dot{P}$). The spin period derivative is related to the spin-down by $\dot{f} = -\dot{P}/P^2$. There are two distinct populations, the millisecond pulsars (MSPs, lower left) and the normal pulsars (upper right). MSPs are belived to be recycled pulsars, i.e. neutron stars spun-up by accretion of matter from a close-by companion [26]. Pulsar data taken from the ATNF catalog [33, 65].
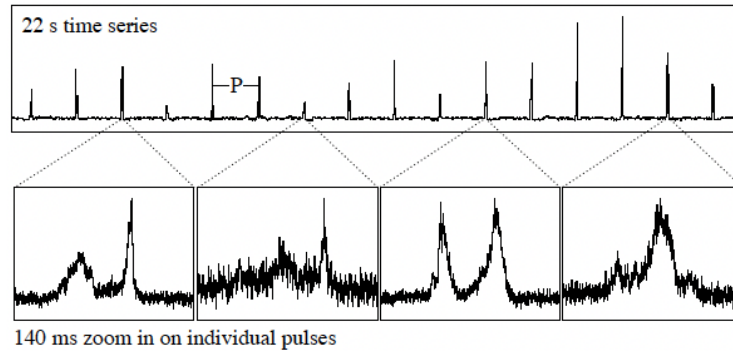

Figure 1.4: Pulsations of the pulsar PSR B0301+19 over 22 s. Credit: [59, Fig. 1.1]

1.4 Neutrons stars as sources of continuous waves

There are multiple emission mechanisms with which a neutron star could emit continuous gravitational waves. Possibly the most promising are fast-spinning neutron stars with a small deformation not aligned with their rotational axis. These are commonly referred to as "mountains" but are of only mm or cm size. These may be sustained by the neutron star crust or the strong magnetic field and formed during their birth, by starquakes or accretion [50, 62, 69]. Accretion can happen when the neutron star has a companion: The neutron star tidally disrupts the companion and strips them of their mass, possibly spinning them up. In the case where the mass flow from the companion onto the neutron stars happens asymmetrically, they could emit continuous gravitational waves, powered as long as the accretion continues [89]. Such a nonaxisymmetric neutron star radiates gravitational waves at twice the rotational frequency of the neutron star [107, 108].

Around 3300 neutron stars are known due to electromagnetic observations [33, 65] (see Fig. 1.3). A highly magnetized spinning neutron star can emit electro-

magnetic radiation in a cone out of its magnetic poles. When the cone is directed towards Earth as the star rotates, we observe light as very consistent electromagnetic pulsations, much like a light house [48, 64, 75] (see Fig. 1.4). First detected by Bell Burnell and Hewish [51], these pulsars are intriguing targets for continuous wave searches: their sky location, frequency and frequency evolution (spin-down) are often precisely measured over years or decades [41, 65].

1.5 The continuous wave signal model

The nearly monochromatic continuous wave signal at the output of the gravitational wave detector can be described by [55],

$$h(t) = F_+ h_+(t) + F_\times h_\times(t), \tag{1.2}$$

with the waveforms h_+ and $h_\times$

$$
\begin{aligned}
h_+(t) &= A_+ \cos \Phi(t) \\
h_\times(t) &= A_\times \sin \Phi(t),
\end{aligned}
\tag{1.3}
$$

with the amplitudes of the two polarizations "+" and "×"

$$
\begin{aligned}
A_+ &= \frac{1}{2} h_0 (1 + \cos^2 \iota) \\
A_\times &= h_0 \cos \iota,
\end{aligned}
\tag{1.4}
$$

and the detector beam-pattern functions for the "+" and "×" polarizations, F_+ and $F_\times$. In fact F_+ and $F_\times$ are functions of $(\alpha, \delta, \psi; t)$, the sky location of the gravitational wave source α, δ (right ascension and declination), and the polarization angle ψ of the wave [55]. ι is the angle between the rotation axis and the line of sight from the neutron star towards Earth, $\Phi(t)$ is the phase of the gravitational wave at time t and h_0 is the intrinsic gravitational wave amplitude.

The continuous gravitational wave amplitude h_0 of a rotating neutron star takes the form [107, 108]:

$$h_0 = \frac{4\pi^2 G}{c^4} \frac{I_{zz} f^2 \varepsilon}{d}, \tag{1.5}$$

where G is Newton's constant, f the frequency of the gravitational wave and d the distance between Earth and the neutron star. It is expected that $f_\text{gw} = 2 f_\text{rot}$. ε is the neutron stars equatorial ellipticity, the deviation from being axisymmetric with respect to the rotation axis. It is defined by

$$\varepsilon = \frac{|I_{xx} - I_{yy}|}{I_{zz}}, \tag{1.6}$$

where I_{zz} is the moment of inertia along the spin-axis (a principle axis), and I_{xx} and I_{yy} are the moments of inertia with respect to the two other principal axis, orthogonal to the spin axis [55].

There is much uncertainty on the actual values of the ellipticity ε. According to Johnson-McDaniel and Owen [56] the neutron star crust could support ellipticities of around 10^{-5}, or in the 10^{-6} regime according to [69, 98], whereas other models

find much lower ellipticities around 10^{-9} [47]. Woan et al. [104] have proposed millisecond pulsars to have a *minimum* ellipticity of $\sim 10^{-9}$. Current searches for continuous gravitational waves can place upper limits on detectable gravitational waves based on non-detections, and probe ranges of $\varepsilon \sim 10^{-5}$ to $\sim 10^{-8}$ (frequency, distant and search dependent).

With distances of say $1\,\mathrm{kpc}$, a gravitational wave frequency of $100\,\mathrm{Hz}$ and moment of inertia of $I_{zz} = 10^{45}\mathrm{g\,cm^2}$ [55], one arrives with Eq. (1.5) at continuous gravitational wave amplitudes values on the order of

$$
\begin{aligned}
& h_0 \sim 10^{-25} \text{ for ellipticities around } \varepsilon \sim 10^{-5} \text{or} \\
& h_0 \sim 10^{-29} \text{ for ellipticities around } \varepsilon \sim 10^{-9}.
\end{aligned} \tag{1.7}
$$

The maximum gravitational wave amplitude for a signal at a given frequency and frequency derivative is the spin-down limit amplitude. One assumes that all loss in rotational energy is converted solely to gravitational radiation. For a neutron star emitting continuous waves at twice the rotational frequency, this is [82]

$$
h_0 = \frac{1}{d}\sqrt{\frac{5GI_{zz}}{2c^3}\frac{\left|\dot{f}\right|}{f}}. \tag{1.8}
$$

For most pulsars the spin-down limits are at a level of $\sim 10^{-28} \sim 10^{-25}$, and thus for most still lower than current searches are able to detect [17, 18]. For some pulsars the spin-down limit has been beaten, most notably for the Vela pulsar [1] or the Crab pulsar [2]. Latest results show no gravitational wave emission at amplitudes a factor of ~ 100 (Crab) and ~ 20 (Vela) *lower* than the spin-down upper limits [17]. This implies that $\leqslant 1\%$ of the Crab's rotational energy is converted and radiated as gravitational waves. It is however difficult to generalize these, rather bleak, results to the largely unknown neutron star population. On the other hand one could be optimistic as the data and the searches become so sensitive as to detect percentile-losses in rotational energies.

Compared with the gravitational wave amplitudes of compact binary coalescences, $h_0 \leqslant 10^{-21}$, the expected continuous wave amplitudes are orders of magnitude weaker. Special search techniques have to be applied to recover the signal buried in the detector noise.

1.6 Matched filtering

Matched filtering is such a search technique, first introduced in the continuous wave data analysis context by the famous "JKS" paper [55], by Jaranowski, Królak and Schutz, 1998. Assuming the noise $n(t)$ in the detector to be Gaussian, additive, stationary and a zero-mean continuous random process, the data $x(t)$, observed over some time $t \in [-T_{\mathrm{obs}}/2, T_{\mathrm{obs}}/2]$, can be described by [55]:

$$
\begin{aligned}
H_0: \quad & x(t) = n(t) && \text{, if no signal is present} & (1.9) \\
H_1: \quad & x(t) = n(t) + h(t) && \text{, if a signal } h(t) \text{ is present.} & (1.10)
\end{aligned}
$$

These are two hypotheses, the null hypothesis H_0 that there is no signal in the data $x(t)$, and H_1 that there is the signal $h(t)$ present. Whether to decide for H_0 or

H_1 is decided by a test statistic, also detection statistic, $\Lambda(x(t))$, which is a function of the data $x(t)$ [27]. Depending on the value of Λ, H_0 or H_1 are rejected.

The Neymann-Pearson-Lemma [70] states that any detection statistic with the same isosurfaces as the likelihood is an optimal detection statistic in the sense that it maximizes the detection probability at any fixed alarm rate. The log-likelihood is also an optimal statistic and it is usually more convenient. This takes the form [54]

$$\ln \Lambda(x) = (x|h) - \frac{1}{2}(h|h), \tag{1.11}$$

where $(x|y)$ is the scalar product defined as

$$(x|y) = 4\mathfrak{Re} \int_0^\infty \frac{\tilde{x}(f)\tilde{y}^*(f)}{S_n(f)} df, \tag{1.12}$$

where $\tilde{x}$ denotes the Fourier transform of x, x^* the complex conjugate of x and $S_n(f)$ the one-sided power spectral density of the detector.

Eq. 1.12 describes how to combine the data to optimally detect a signal with shape $h(t)$. The optimal estimation is a convolution in time of the signal waveform with the data. When the data and the signal "match", the value of the detection statistic is maximum.

In the case of a continuous wave signal the optimal detection statistic, analytically maximized over the parameters h_0, ι, ψ, ϕ_0, is called the $\mathcal{F}$-statistic (ϕ_0 is is the initial phase defining the position of Earth at a fiducial time $t = 0$ [55]). In Gaussian noise the $\mathcal{F}$-statistic follows a chi-square distribution with 4 degrees of freedom, $\chi_4^2(x)$ [55].

In the presence of a signal the $\mathcal{F}$-statistic has a non-zero noncentrality parameter equal to the optimal signal-to-noise ratio ρ^2 . It can be shown that in the presence of a signal the noncentrality parameter is

$$\rho^2 = (h|h) \tag{1.13}$$

and in Gaussian stationary and approximately white noise

$$\rho^2 \propto \frac{h_0^2 T_{\text{obs}}}{S_n(f_0)}, \tag{1.14}$$

where f_0 is the frequency of the continuous gravitational wave signal and T_{obs} the data duration [55].

1.7 Targeted, directed and all-sky searches

Albeit the seemingly easy signal model for a continuous gravitational wave, searches for them are computationally difficult. The nearly monochromatic signal, which may spin-down over time due to energy loss, is Doppler modulated by Earth's rotation and orbit around the Sun. So the same intrinsic wave shape is seen as different for sources at different sky locations. This means that different matched filters need to be computed. While not a problem for a known source, i.e. known frequency, spin-down and sky location, it quickly becomes computationally challenging for unknown sources.

There are two different categories of continuous wave search methods: *coherent* and *semi-coherent*. In a coherent search the data of the whole observation time is taken as a whole, and the signal template is applied to it. In a semi-coherent search the data is split into shorter segments, which are analyzed coherently. The resulting detection statistics are combined incoherently, i.e. there is no phase coherence required between the segments. This looses a bit of the sensitivity compared to coherent searches, but is much faster. The next Section 1.9 will explain this further. There are variations and different approaches, however, we will focus on coherent and semi-coherent searches in the $\mathcal{F}$-statistic context. For an overview see Riles [85].

Around 3300 neutron stars are known due to electromagnetic observations [33, 65]. Some pulsars are observed for decades, i.e. the spin-frequency, frequency-evolution and sky position are known to a high certainty [41]. For these one can in principle run inexpensive single-template, so called *targeted searches*, although one usually searches a small range around the parameters to account for uncertainties in the pulsar parameters or uncertainties in the relationship between spin-frequency and gravitational wave frequency. These are then called *narrow-band searches* [18]. Both searches are relatively inexpensive and can be done coherently. The searches presented in Chapters 5, 6 and 7 are targeted searches for pulsars not previously searched for.

For some sources a sky location is known, but there are high uncertainties in its frequency, or the frequency is completely unknown. In this case *directed searches* are done. The range of frequencies and spin-downs searched depends on the information available and typically is astrophysically motivated. Possible targets are suspected neutron stars, i.e. where no periodic pulsations to determine the gravitational wave frequency are seen, like e.g. Formalhaut b [58] (later ruled out to be a planet [38, 46]), or supernova remnants like Cassiopeia A , Vela Jr. or G347.3-0.5 [68]. Other possible targets are regions where a high density of neutron stars is expected, like the Galactic Center [16] or globular clusters (see Chap. 9) [40]. Searches targeting such regions are sometimes also referred to as *spotlight searches*. Directed searches are computationally challenging due to the high number of search templates necessary to cover the parameter space, thus they are usually done semi-coherently. Chapter 8 describes a directed search for the supernova remnant G347.3-0.5 and Chapter 9 for a spotlight search for continuous gravitational wave emission from the globular cluster Terzan 5 and the Galactic Center.

Based on population studies of stars the number of neutron stars in our galaxy is estimated to be around $10^8 \sim 10^9$ [84, 86]. Many of them will be old and spun-down as to not be detectable anymore, however, some of them can be "recycled" by accreting matter of a companion and be spun-up due to this [98]. If just a small fraction of that billion is young and energetic, but they either have no beam, or the beam does not sweep over Earth, or it is blocked by cosmic dust, gravitational waves could be our only chance to observe them.

This warrants *blind* or *all-sky searches*, which set only very loose constrains on frequency and spin-down, and search all over the sky. These are the most computationally expensive searches for gravitational waves and require a huge amount of waveforms to cover the parameter space, since each sky point requires a different demodulation. These searches are done semi-coherently. Chapter 4 describes the Einstein@Home all-sky search on data of the second observation run of Advanced LIGO (O2).

Aforementioned searches can also be done for non-isolated neutron stars, i.e. in binary or more-body systems where the additional Doppler modulation due to the sources orbital motion too has to be accounted for, making them even more expensive to run. Recent years showed an emergence of more binary all-sky searches [13, 36].

Electromagnetic observations show, especially young and energetic, neutron stars to glitch. A glitch is described by an instantaneous jump in frequency with some relaxation over time back to nearly the previous frequency [19]. The process why this happens is not yet understood, possible explanations are starquakes or the breaking of the neutron star crust. These glitches break the phase coherence of the signal, hence special searches for glitching neutron stars can be done. Due to the additional parameters of time and strength of the glitch, these searches are expensive, too [44].

1.8 Template banks and the mismatch

We search a range of frequencies, spin-downs and sky positions, and discretize it to a number of discrete values. These are often referred to as *template banks*, defined by the *search grids* on the different signal parameters.

We consider a set of waveform templates covering the signal parameter range that we want to investigate. Each template is described by a set of signal/source parameter values, for example sky position, frequency and frequency derivative(s) (usually only the first or second). The set of templates is hence described by the different combinations of parameters.

The distance between the values is chosen in such a way that the loss of signal-to-noise ratio due to the signal-template mismatch, if the signal did not lie on one of the grid points, would on average not exceed some predefined value.

If we indicate the optimal signal-to-noise ratio at a certain template $\vec{\lambda}_0$ as $\rho^2(\vec{\lambda}_0)$, then we indicate the mismatch $\mu(\vec{\lambda}_0, \Delta\vec{\lambda}_0)$ at $\vec{\lambda}_0$ due to a signal-template distance $\Delta\vec{\lambda}_0$ as:

$$\mu(\vec{\lambda}_0, \Delta\vec{\lambda}_0) = 1 - \frac{\rho^2(\vec{\lambda}_0, \Delta\vec{\lambda}_0)}{\rho^2(\vec{\lambda}_0)}, \tag{1.15}$$

where

$$\rho^2(\vec{\lambda}_0) = \rho^2(\mu(\vec{\lambda}_0, \vec{\lambda}_0)). \tag{1.16}$$

For small $\Delta\vec{\lambda}$ the mismatch can be approximated as a quadratic form of the distance in parameter space

$$\mu(\vec{\lambda}_0) = g_{ij}(\vec{\lambda}_0)\Delta\lambda^i\Delta\lambda^j \tag{1.17}$$

[79], and more precisely [20] by

$$\mu(\vec{\lambda}_0) = \sin^2\sqrt{g_{ij}(\vec{\lambda}_0)\Delta\lambda^i\Delta\lambda^j}. \tag{1.18}$$

Typically the grids used in broad continuous wave searches are much coarser than the scales at which the metric approximation of Eq. (1.17) and (1.18) holds. This is good because in fact the metric approximation overestimates the mismatch and therefore we do not need as fine grids as it prescribes. On the other hands this means that one has to run large scale Monte Carlo simulations to find the optimal

search grids within the available computational bounds. For further reading we refer to Ming et al. [67] describing the process for directed searches, but the same process is applicable and was applied to all-sky searches as well.

1.9 Semi-coherent searches

Large scale surveys typically require $\geqslant 10^{17}$ template waveforms [12, 66, 68, 93]. The observation time T_{Obs} of an observation run is usually between half a year to a year [1]. A coherent search, i.e. a coherent waveform matched against the full time span of observational data, becomes computationally infeasible quickly. This is due to the fact that the resolution in parameter space grows for $f \propto T_{\mathrm{obs}}$, $\dot{f} \propto T_{\mathrm{obs}}^2$ and $\ddot{f} \propto T_{\mathrm{obs}}^3$ [3]. The spacings in sky (α, δ) are proportional to the frequency squared f^2 and approximately proportional to T_{obs}^2 (in the range between a few days to about one year) [3].

Directed searches have at most only a few sky points but the signal model includes $\ddot{f}$, and thus the resolution usually scales with $\propto T_{\mathrm{obs}}^6$. All-sky searches typically only consider signals with at most a resolved first order frequency derivative, so in the end the resolution still scales $\propto T_{\mathrm{obs}}^6$, as for directed searches. Therefore fully coherent all-sky and directed searches over months of data are unfeasible.

Brady and Creighton [29] proposed *semi-coherent* searches. In essence the data is divided into N_{seg} segments spanning $T_{\mathrm{seg}} < T_{\mathrm{obs}}$, each segment is analyzed coherently and the per-segment detection statistics are combined incoherently. Hence to first order the computational cost is reduced from $\sim T_{\mathrm{obs}}^6$ to $\sim N_{\mathrm{seg}} T_{\mathrm{seg}}^6$. With N_{seg} usually on the order of ten to hundred the computational cost is reduced by a factor[2] of 10^5 to 10^{10}.

While the sensitivity of a semi-coherent search is generally lower than a fully coherent search on a single template, Brady and Creighton [29] first showed that over a large parameter space, and at fixed computational resources, semi-coherent approaches yield better sensitivities than what could be achieved with the same computing budget with a coherent search.

There are a number of different semi-coherent approaches and variations: The Time-Domain $\mathcal{F}$-statistic, Five-vector, Sky- and Frequency Hough, StackSlide, PowerFlux, Viterbi, Sideband and other methods. For an overview see [85] and references therein.

In this thesis we will restrict ourselves to the StackSlide-$\mathcal{F}$ method as described in Pletsch [80], Pletsch and Allen [81]. The available detector data is split into half-hour chunks and, after some time- and frequency-domain conditioning of the data to remove spectral contamination and time-domain disturbances (see Chapter 2) , it is Fourier transformed into $1\,800\,\mathrm{s}$ time-baseline Short Fourier Transforms (SFTs). The coherent segments usually span $60\,\mathrm{h}$ to $120\,\mathrm{h}$ and are constructed from the SFTs.

The StackSlide-$\mathcal{F}$ method produces an average $2\overline{\mathcal{F}}$ [80, 81]

---

[1]O1 was four months, O2 nine and O3a + O3b eleven months, while the earlier "science"-runs S5 and S6 where 23 and $15\frac{1}{2}$ months respectively.

[2]This is: $T_{\mathrm{seg}} \cong T_{\mathrm{obs}}/N_{\mathrm{seg}}$ and $N_{\mathrm{seg}} \left(\frac{T_{\mathrm{obs}}}{N_{\mathrm{seg}}}\right)^6 = \frac{T_{\mathrm{obs}}^6}{N_{\mathrm{seg}}^5}$

$$2\overline{\mathcal{F}} = \frac{1}{N_{\text{seg}}} \sum_{i=1}^{N_{\text{seg}}} 2\mathcal{F}_i, \tag{1.19}$$

where the $2\mathcal{F}_i$ are the coherent per-segment $\mathcal{F}$-statistic values, usually computed at different templates per segment, tracking the phase evolution of the signal. The resulting $2\overline{\mathcal{F}}$ approximates the $2\overline{\mathcal{F}}$ computed at a point on a finer grid than the one used for the N_{seg} coherent searches.

In the searches usually the N_{cand} candidates with the highest detection statistic value are saved into a *top list*, and in the end the detection statistic is recalculated at the exact fine-grid point for the top list candidate. The result is usually denoted by a subscript r, i.e. $2\mathcal{F}_r$.

The average $2\mathcal{F}_r$ has a noncentrality parameter which is proportional to T_{seg} and the distribution of $N_{\text{seg}}2\overline{\mathcal{F}}$ follows a chi square distribution with $4N_{\text{seg}}$ degrees of freedom [83, 102].

The summed $2\overline{\mathcal{F}}_{\text{sum}}$ have the expectation value and variance [83, 102]

$$E\left[2\overline{\mathcal{F}}\text{sum}\right] = 4N_{\text{seg}} + \rho^2, \tag{1.20}$$

$$\sigma^2\left[2\overline{\mathcal{F}}\text{sum}\right] = 8N_{\text{seg}} + 4\rho^2. \tag{1.21}$$

1.10 Hierarchical searches

An additional technique to increase the sensitivity of semi-coherent searches are hierarchical searches, first introduced as two stages [29, 30], later extended to multi-stage searches [39, 76, 88], and first implemented in Papa et al. [78].

The first stage, commonly Stage 0, allows for a high mismatch of usually around 50% and a high false-alarm level. In the subsequent follow-up stages, the high number of false alarms are steadily reduced, while being cautions that no signals are falsely discarded. This is achieved with Monte Carlo studies on thousands of fake signals from the target population, and making sure that the false-dismissal probability is very low, usually $\leqslant 0.1\%$. In each follow-up stage the resolution is increased, and/or the coherence time T_{seg} is increased. This is computationally feasible because (a) only a smaller volume around the candidates is being searched, and (b) candidates are rejected.

Candidates are rejected by comparing their $2\mathcal{F}_r$ increase against the increase in signal-to-noise ratio expected for signals. Candidates due to noise and signals are $\chi^2_{4N_{\text{seg}}}$-distributed, but in the case of noise the distribution is central, and in the case of a signal it has a noncentrality parameter $\rho^2 \propto T_{\text{seg}}$. With increasing coherence time T_{seg} the variance in the summed $2\overline{\mathcal{F}}_{\text{sum}}$ decreases and the noncentrality parameter increases, hence the distribution of signals and noise separate. By setting a threshold on the lowest increase in $2\mathcal{F}_r$ from stage to stage, further and further candidates are rejected.

Candidates which survive the fully-coherent follow-up stage—outliers—are interesting continuous wave candidates. One can verify a candidate in other detectors and/or using data from different observation runs. If nothing is found this hints at a noise outlier or at a signal whose waveform model is different than the assumed one. If there are coincident results, this would add significantly to the possibility to

be from an astrophysical source. At that stage further observations from different channels could clearly add to the significance of a detection, i.e. electromagnetic observation for radio-, x-ray or gamma-ray pulsations or emissions (see Papa et al. [77]).

Usually the number of candidates from Stage 0 is on the order of billions, making a follow-up computationally infeasible. A crucial step in the post-processing, i.e. after Stage 0, is a method called *clustering*. Due to the way searches are set up, one disturbance or signal can give rise to few thousand nearby candidates, while noise fluctuations usually produce less clumped candidates. Clustering aims at finding all candidates due to one cause. It bundles—clusters—them and picks one representative candidate, the seed, which is then followed up instead of the bunch. This drastically reduces the number of candidates and make a follow-up possible. There have been different approaches to clustering, and the difficulty lies in balancing a high noise rejection of 95% to 99.9% while simultaneously loosing as few signals as possible. Chapter 3 describes the clustering method that I designed and then used in the Einstein@Home O2 and O3 all-sky searches.

1.11 Computational aspects

I am grateful to have two immensely useful computing facilities at hand, the ATLAS computing cluster at the Albert Einstein Institute in Hannover, and Einstein@Home, a volunteer distributed computing project build on the BOINC infrastructure [21, 22, 28]. With Einstein@Home volunteers (hosts) can spend their idle computational power to not only advance our science case, looking for continuous wave, but also to search for radio or gamma-ray signals of pulsars.

Usually the large Stage 0 searches are run on Einstein@Home. The template placement is carefully chosen by minimizing the average mismatch of the setup (Eq. (1.15)) while staying within the wall-clock runtime of usually around 6 months. The whole parameter space is split into evenly long-running jobs, so-called work units (WU). Work units are not dependent on each other, thus can be run in an *embarrassingly parallel* way. Einstein@Home-hosts are a very heterogeneous mix of hardware and software, making runtime estimations challenging. A fraction of work units are calculated twice and validated against each other to identify hosts that are not operating reliably. Historically well-behaved hosts are not checked as often as new hosts.

Before a search can run on Einstein@Home it has to be set up, and results depend on the host completing the calculations, so the return time is not instant. Ad-hoc or time-critical projects or investigations that need huge amounts of data, are more efficiently performed on the ATLAS high-throughput computing cluster. In particular the data preparation, mismatch- and Monte Carlo studies, post-processing, follow-up searches and upper limits presented in this thesis were performed on the ATLAS cluster. ATLAS consists of currently $\sim 50\,000$ CPU cores on ~ 3000 machines, and ~ 2350 GPUs, making it able to calculate more than 2 PFLOP/s (10^{15} floating-point operations per second) [35], and an extraordinary resource to have available.

Identification and removal of non-Gaussian noise transients for gravitational wave searches

This work describes a method to prepare Advanced LIGO gravitational wave observatory data for (continuous) gravitational wave searches.

The work was initiated as the author's master thesis as an adaptation of a gating method used for transient signal searches [99] to make it more suitable for continuous gravitational wave searches. The author ported the first python implementation to C, into a local LALSuite-framework [61]. During the master thesis the author prepared data of the second observation run (O2) of Advanced LIGO, which was used for the Einstein@Home O2 all-sky search. In the author's PhD studies the methodology was significantly enhanced and generalized. The author developed a framework around this application, with which data (SFTs) of any observation run can be easily produced, with an arbitrary time-baseline, i.e. shorter for binary searches or longer to inspect narrow line features.

The end product is a general purpose tool that is publicly available (it was released open source).

This has become the standard tool used to prepare data for most of the group's continuous wave searches. In fact the author prepared the O2, O3a and O3b data sets for isolated and binary searches, like employed in [25, 37, 44, 68, 71, 72, 77, 91, 105, 106].

The methodological development work was carried out by the author with advice from the supervisor. The programming work benefited from the guidance, especially in the initial phases, of H.-B. Eggenstein, who is a software engineer.

The paper was written by the author under the guidance of the supervisor M. A. Papa. All figures were produced by the author.

Identification and removal of non-Gaussian noise transients for gravitational-wave searches

Benjamin Steltner[1,2,*] Maria Alessandra Papa[1,3,2,†] and Heinz-Bernd Eggenstein[1,2,‡]

[1]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute),*
Callinstrasse 38, 30167 Hannover, Germany
[2]*Leibniz Universität Hannover, D-30167 Hannover, Germany*
[3]*University of Wisconsin Milwaukee, 3135 N Maryland Avenue, Milwaukee, Wisconsin 53211, USA*

We present a new *gating* method to remove non-Gaussian noise transients in gravitational-wave data. The method does not rely on any *a priori* knowledge on the amplitude or duration of the transient events. In light of the character of the newly released LIGO O3a data, glitch-identification is particularly relevant for searches using this data. Our method preserves more data than previously achieved, while obtaining the same, if not higher, noise reduction. We achieve a $\approx$2-fold reduction in zeroed-out data with respect to the gates released by LIGO on the O3a data. We describe the method and characterise its performance. While developed in the context of searches for continuous signals, this method can be used to prepare gravitational-wave data for any search. As the cadence of compact-binary inspiral detections increases and the lower noise level of the instruments unveils new glitches, excising disturbances effectively, precisely, and in a timely manner, becomes more important. Our method does this. We release the source code associated with this new technique and the gates for the newly released O3 data.

I. INTRODUCTION

While many loud gravitational-wave signals have been detected, much of the high precision science and new discoveries in the nascent field of gravitational-wave astronomy will benefit from noise-characterization and noise-mitigation techniques [1–8].

The data of gravitational-wave detectors is dominated by noise. This noise is by and large Gaussian with a stable spectrum, but $\lesssim 10\%$ of it may be infested by high-powered short-lived disturbances (glitches) and by nearly monochromatic coherent spectral artefacts (lines) in a variety of amplitudes, from extremely large to extremely weak.

Typically the short-lived glitches affect the sensitivity of short-lived signal searches while the coherent lines affect the sensitivity of searches for persistent signals. But when a short-lived glitch is powerful enough, it can also temporarily degrade the sensitivity of searches for long-lived

signals, by increasing the average noise-floor level in a broad frequency range for its duration.

Two noise-mitigating techniques are typically used to prepare the gravitational-wave data for searches: *gating*, performed in the time-domain and *line-cleaning*, performed in the frequency domain. Broadly speaking, the former is used to remove loud glitches and the latter to remove spectral lines. The latter is usually only used in searches for persistent signals or stochastic backgrounds [9–11].

In this paper we illustrate a new gating application, `gatestrain`, which enables a more precise removal of glitches compared with other widely used and publicly available gating methods, discarding significantly less data. Furthermore our gating procedure does not rely on any single fixed threshold that establishes what data should be gated, but rather it adjusts the threshold based on the achieved noise reduction. These are important features when the glitches vary much from dataset to dataset, and within the same dataset, because the method does not require time-intensive tuning of ad-hoc parameters.

We publish the gates found with our new method on the public O3 data of the Advanced LIGO detectors as well as the new gating application in the Supplemental Material [12].

The paper is organized as follows. In Sec. II we describe the noise disturbances in Advanced LIGO data, which are particularly detrimental to continuous-wave searches, and the typical mitigation techniques used to prepare the data before performing such searches. In Sec. III we present the

*benjamin.steltner@aei.mpg.de
†maria.alessandra.papa@aei.mpg.de
‡heinz-bernd.eggenstein@aei.mpg.de

16

idea of time-domain gating and explain our new method `gatestrain`. The performance of `gatestrain` on Advanced LIGO data from the first, second and third observation runs (O1, O2 and O3a) is presented in Sec. IV. In Sec. V we discuss our results.

II. NOISE AND MITIGATION TECHNIQUES

The present generation of gravitational-wave detectors operates in a noise-dominated regime. The noise is primarily Gaussian with two main types of deviations: short-lived non-Gaussian transients—*glitches*—and long-lived nearly monochromatic coherent artefacts—*lines*.

Lines are instrumental or environmental disturbances manifesting as narrow spectral artefacts, sometimes visible as lines in the frequency domain of the raw data. These disturbances lead to false candidates in continuous-wave searches and in searches for stochastic backgrounds [13,14]. A standard way to deal with lines has been to replace the affected frequency bins with Gaussian noise in the data input to the search, not allowing the excess power to "spread" to many signal-frequency results. This method is called *line cleaning*. Line cleaning relies on knowing where the spectral contamination occurs and hence on detector-characterization studies such as [15] that produce the so-called "lines lists" that LIGO releases together with its data.

Loud glitches impact the sensitivity of transient signal searches by contributing to the background distribution used to estimate the significance of any finding. But they also degrade persistent-signal searches in two ways: (1) a high-power glitch directly increases the noise floor in a broad frequency range and (2) a loud glitch invalidates one of the assumptions of the line-cleaning method, and introduces artefacts in the cleaned data. We will discuss this latter point in Sec. IV B.

A typical mitigation technique for glitches is *gating*: the time-domain data affected by a glitch is simply removed. Gating is a standard step of the compact-binary coalescence search pipeline `pyCBC` [16], but continuous-wave search pipelines also use it [8,17,18]. In fact the O1 data Einstein@Home search for continuous waves from Cassiopeia A, Vela Jr. and G347.3 [10,19] used the gating on its data and specifically used the `pyCBC` gating module because of its ease of use and prompt availability.

III. TIME-DOMAIN GATING

The core idea of gating is to detect and remove glitches in the time-domain. The different applications differ mostly in how the glitch detection is done and what classifies as being part of the glitch. This has implications on the effectiveness of the gating and/or on how much "tinkering" and tuning is necessary to achieve optimal gating in any specific dataset.

Typically gating in preparation for transient signal searches tends to be less aggressive in removing data than the gating in preparation for persistent signal searches, by using shorter gates and higher thresholds. With this approach a number of glitches survive, increasing the false alarm rate, but transient signals that may happen near glitches are not discarded together with the glitch. Persistent signals on the other hand are always-on, thus even with the most aggressive cleaning, only a small fraction of signal is removed.

In this section we describe our gating method that presents two novelties with respect to publicly available gating methods: (1) the adaptive determination of the gate duration (2) the self-adjusting amplitude threshold for gate-identification, based on a iterative data-quality check of the gated data.

In most gating methods employed in gravitational-wave searches, the glitch detection does not happen on the raw data. Our gating procedure uses the same initial signal-processing steps as `pyCBC`, up to the actual detection of glitches, when the two methods differ. The data is divided in chunks with durations on the order of few to tens of minutes. In [10,19] we used chunks that are 1800 s long. For each chunk the following steps are taken:

(i) high-pass the data with an 8th-order Butterworth filter at 10 Hz. Since the released Advanced LIGO data is not to be used for astrophysical searches below 10 Hz, we refer to this as the input data.

(ii) let $\mathcal{P}_r(f)$ be the power spectral density noise floor (with units [1/Hz]) estimated from this data. Since the data is not stationary, in order to produce this "reference" power spectral density we divide the chunk in O(100) segments, compute the noise spectrum on each of these and, bin per bin, take the median over all the realizations.

(iii) take the Fourier-transform $\tilde{h}(f)$, whiten and obtain: $\tilde{h}_w(f) = \frac{\tilde{h}(f)}{\sqrt{\mathcal{P}_r(f)}}$

(iv) inverse-Fourier-transform and obtain $h_w(t)$ ($[\sqrt{\text{Hz}}]$)

(v) gate the $h_w(t)$ time series

This whitening process produces a $h_w(t)$ time series that in Gaussian noise has a mean $\mu = 0$ and standard deviation $\sigma = 1.0$ [3], with similar contributions from all frequencies. When a glitch happens, it is more visible in $h_w(t)$ than in the original $h(t)$, because its contribution is not "hidden" by loud noise from the low frequencies. This is shown clearly in Fig. 2.

Periods that harbor glitches are identified in $h_w(t)$. The data in these periods is set to zero with a Tukey taper on either side of the period. With the expression gate, g_i, we refer to each set of neighboring data points whose original value has been set to zero and their time stamps: $g_i = (\{t\}_i, \{h_w^g\}_i)$.

The `pyCBC` gates are established based on two parameters: a threshold H and a duration τ_{dur}. All times t_k are recorded where $|h_w(t_k)| > H$. The process of constructing the gates is strictly sequential. It starts with the first point t_1.

This identifies the points lying within a distance $\tau_{\text{dur}}/2$ of t_1. The time of the largest h_w among these, $t_1^\star$, is taken as the center of the first gate. All points $\in [t_1^\star - \tau_{\text{dur}}/2, t_1^\star + \tau_{\text{dur}}/2]$ are zeroed out. A taper is also applied to either side of this interval to prevent discontinuities [16]. The second point t_2 is the next t_k that has not been affected by the gating at the previous iteration. The process repeats for t_2 as for t_1 and a second gate is identified. The process ends when there are no more t_ks. We note that gates may overlap. The typical values for transient signal searches, e.g., in the first gravitational-wave transient catalogs [20–24], are $H = 100\sqrt{\text{Hz}}$ ($25\sqrt{\text{Hz}}$), $\tau_{\text{dur}} = 0.25$ s (0.125 s) and a Tukey taper of 0.25 s (0.125 s) (for the latest catalogs). In [19] we used $H = 50\sqrt{\text{Hz}}$, $\tau_{\text{dur}} = 16$ s and a Tukey taper of 0.25 s.

Figure 1 shows that glitches can last from fractions of a second to a few tens of seconds, with more than 60% of the glitches in the O2 data lasting less than 1s. It also shows that there is a great variability in glitch duration, depending on the detector and on the run. For instance, in O1 $\sim 50\%$ of glitches in either detectors last less than 1 s, whereas in O2 $\sim 50\%$ of the L1 glitches last less than 0.3 s. This variability is hard to capture with simple glitch-detection schemes: for example for glitches having "long tails" that do not make it above the single threshold, those tails remain undetected and are excluded from the gates.

We develop a more generic glitch-identification and -removal scheme, with a varying gate size, estimated on the data itself. This is particularly relevant when other data, for example from environmental monitors around the detectors, is not available, as for the gravitational-wave data releases.
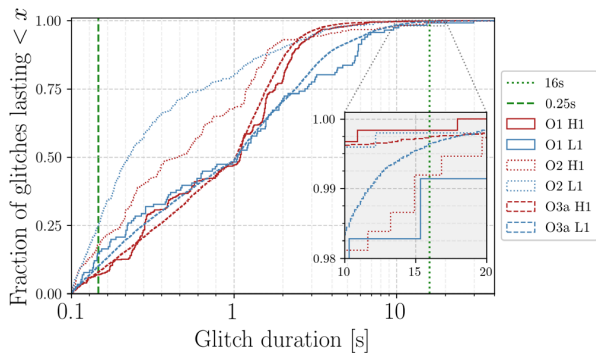

FIG. 1. Cumulative histogram of measured glitch durations. Depending on run and detector 50% to 80% of glitches last less than one second, while less than a few percent last longer than O(10 s). The maximum glitch duration is a few tens of seconds for H1 in O1 and O2 data, and L1 in O1 data, and reaches nearly 300 s in L1 O2 data. The glitch durations are measured with `gatestrain`. Note that the x-axis is displayed in symlog, i.e., linear between 0–1 s and logscale above.

Our method uses three parameters: a high threshold H^{high}, a low threshold H^{low} and a duration parameter τ_{dur}. The gates are constructed as follows:

(i) All times t_k^h are recorded where $|h_w(t_k^h)| > H^{\text{high}}$.
(ii) All times t_j^ℓ are recorded where $|h_w(t_j^\ell)| > H^{\text{low}}$.
(iii) The t_j^ℓ are then divided in groups such that for each member of the group there exists at least another member closer than τ_{dur}. When there are no nearby points, a single-member group is created.
(iv) We only keep those groups such that there exists at least a t_k^h closer than τ_{dur} to at least one member of the group.
(v) For each of the surviving groups: all timestamps between the earliest and latest, plus a Tukey taper to either side, constitute a gate.

The low threshold is set as $H^{\text{low}} = n^\ell \sigma$, where σ is an estimate of the standard deviation of well-behaved parts of the data. We have used the harmonic mean of the standard deviation of $h_w(t)$ from shorter duration chunks, say ≈ 10 s long, out of the 30 minute segment under consideration. We use the harmonic mean so that σ is not affected by the presence of disturbances. We set n^ℓ to be high enough that Gaussian noise fluctuations at such level are rare, typically $n^\ell \approx 5.5$.

The high threshold is crucial because whether a glitch is identified, hinges on there being $|h_w(t)|$ values above H^{high}. A too low H^{high} leads to too many unnecessary gates and thus wasted data, while a too high H^{high} leads to missed glitches. We use an iterative lowering of H^{high} and evaluate the performance of the gating at each threshold. We stop lowering the threshold when it has reached a preset minimum value or when the measured performance is satisfactory.

As an indicator of the performance of the gating we take the quantity

$$R = \frac{1}{N_f} \sum_{f_i}^{N_f} \frac{\mathcal{P}(f_i)}{\mathcal{P}_r(f_i)}, \tag{1}$$

where $\mathcal{P}$ is the power spectral density from the gated data and $\mathcal{P}_r$ is the reference power spectral density described in Sec. III.[1] The sum is over frequency bins f_i. Experience has shown that using a ≈ 5–10 Hz band between 25 Hz and 40 Hz, depending on the run, without loud lines or disturbances is suitable to identify most glitches. The reason for this lies in the character of the LIGO data, with most glitches having spectral content at lower frequencies. A value of $R \approx 1$ indicates that the gated time series is

---

[1]In `gatestrain` it is alternatively possible to specify an external file which holds the reference PSD. This could e.g., be calculated by taking the harmonic mean over the full run. Since the detector changes on various timescales, this method is not recommended.
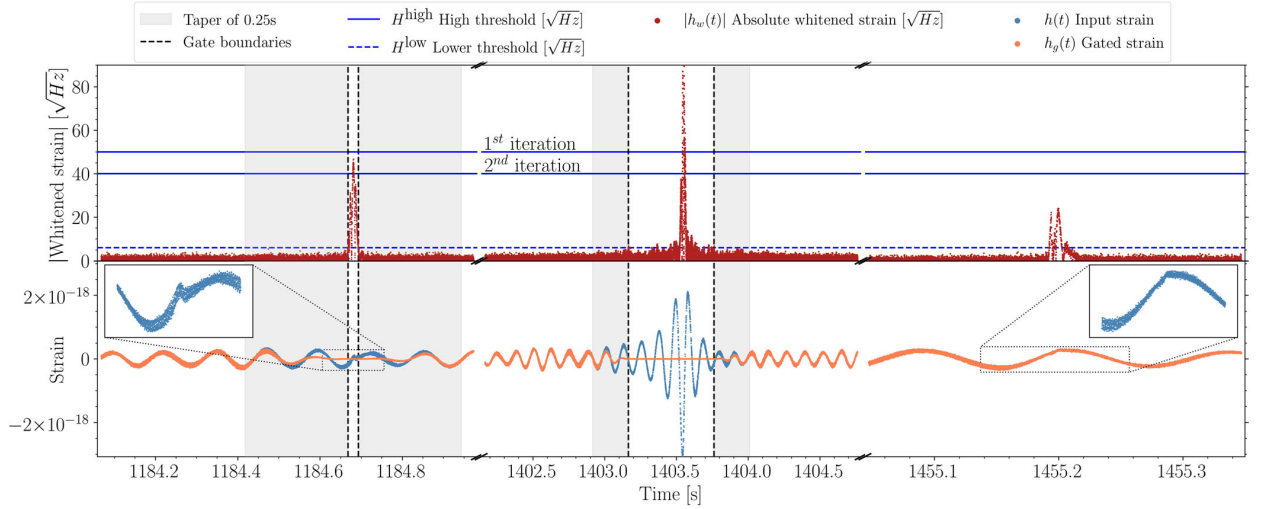
FIG. 2. Example of how the gating procedure works using three snippets of data spanning ≈270 s. The lower panel shows the original strain data $h(t)$ (blue) and the gated $h_g(t)$ (orange). The insets magnify the input strain to show where the glitch occurs. The three snippets of data present glitches of different size. The upper panel shows the absolute value of the whitened time series $|h_w(t)|$ (red), which is the quantity used to detect glitches, as explained in the main text. The glitch-detection threshold is the (blue) horizontal lines, solid and dashed.

not / no more affected by glitches. Value of $R > 1$ indicates the presence of glitches in the gated data.

At the first iteration we use a high threshold, $H^{high} \approx 50\sigma$, gate the data and compute R. If this ratio exceeds a threshold R^{th}, we reduce H^{high}, gate the data and check R again. We continue until either H^{high} reaches a minimum value or R becomes small enough. In the O1 and O2 data we found that decreasing the H^{high} threshold by 10 at each iteration, and setting the R and the H^{high} thresholds to 1.05 and 20, respectively, achieved stable and good performance. On O3 data the same choices gave very good performance. We note that a reasonable choice for the H^{high} could be to set it equal to H^{low}, but in O1 and O2 data this leads to sacrificing a lot more data, for a very small decrease in noise level. For this reason we leave it as a free parameter.

An example of this process with $\tau_{dur} = 3$ s and a Tukey window of 0.25 s, is shown in Fig. 2 (time-domain) and Fig. 3 (frequency-domain). Three glitches are clearly seen in the time-domain plot. In the first iteration, with the highest H^{high} threshold, only the middle peak is detected and removed. The resulting amplitude spectral density (ASD, equal to $\sqrt{\mathcal{P}}$) is shown in purple in Fig. 3. The comparison with the reference $\mathcal{P}_r$ yields $R \geq 1.32$ and indicates that there could be more glitches, so the process continues with a lower values of H^{high}. In the second iteration $H^{high} = 40$ and the second glitch is included. After removing the second glitch $R \leq 1.05$ and this concludes the gating procedure. The third peak is thus not gated as it has not enough impact on the sensitivity. For comparison the ASD after removal of the third peak is also shown.
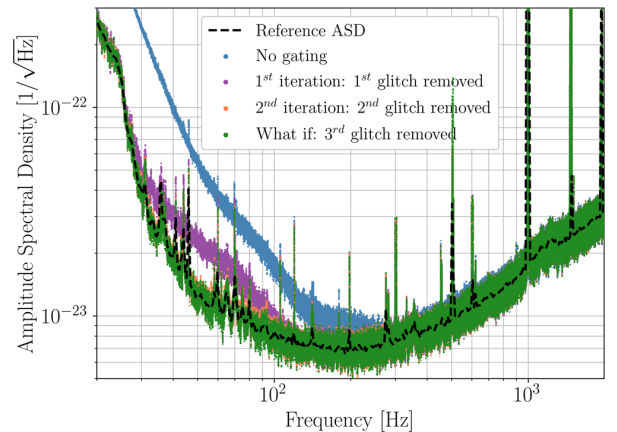

FIG. 3. Amplitude spectral density (ASD) of the data during the 1800 s from which the snippets of Fig. 2 are taken. We show the ASD of the data at different stages of the gating process. The noise floor of the original data (top curve, blue) is significantly higher than the reference ASD (dashed line). In the first iteration (purple) the glitch from the central data snippet is removed, resulting in a vastly improved ASD, but the comparison to the reference ASD shows potential for further improvement below ∼100 Hz. In the second iteration the left data-snippet glitch of the previous figure is found and removed, and this lowers the noise floor the low-frequency region (orange). At this point gating ends because the reference power spectral density $\mathcal{P}_r$ is matched. The glitch in the right hand-side data snippet is not large enough to make a difference and it is left ungated. For reference, the bottom line (green) shows the ASD after removing this third glitch, which our procedure does not do.

TABLE I.   Data that we used the gating procedure on.

Data	SFTs	Data [h]	Data [d]
O1-H1	3124	1562	65.1
O1-L1	2120	1060	44.2
O2-H1	5066	2533	105.5
O2-L1	4984	2492	103.8
O3a-H1	5977	2988.5	124.5
O3a-L1	6377	3188.5	132.9

We note that very loud glitches can lead to a ringing of the whitening filter and thus an overestimation of the glitch duration: the full glitch is removed, but the method's efficiency in saving data is degraded.

IV. RESULTS

We consider public data from the O1, O2 and O3a Advanced LIGO runs [25]. We produce half-hour baseline (Short) Fourier transforms, SFTs [26], as summarized in Table I. We prepare different datasets, depending on the run, to compare the performance of our method with existing ones. Table II summarizes how much data was gated by the different procedures.

A. Gating the O1 and O2 LIGO data

We prepare three different sets, one without gating, one with the pyCBC gating procedure used in [19], with τ_{dur} conservatively set to 16 s, and one with our new gating gatestrain.

Table II shows how many gates were used and how much time was zeroed-out by each procedure. While the number of gates of our procedure is similar or larger than the

TABLE II.   Total amount and duration of gates for each detector / observing run produced by gatestrain and pyCBC as used in [19] or LIGO's self-gating procedure used on O3 data [27]. Since pyCBC gates may overlap, their total duration is less than the total number of gates times 16 s.

Data	Method	SFTs w/gates	Gates	How much data zeroed-out [s]	[h]
O1-H1	pyCBC	667	884	12360.62	3.43
O1-H1	gatestrain	686	799	827.20	0.23
O1-L1	pyCBC	173	222	3110.29	0.86
O1-L1	gatestrain	183	205	271.05	0.08
O2-H1	pyCBC	708	784	12453.41	3.46
O2-H1	gatestrain	852	980	479.29	0.13
O2-L1	pyCBC	620	692	10603.56	2.95
O2-L1	gatestrain	981	1151	723.94	0.20
O3a-H1	self-gating	5695	20205	141070.9	39.2
O3a-H1	gatestrain	4885	11581	38236.8	10.6
O3a-L1	self-gating	6366	49653	1441915.8	400.5
O3a-L1	gatestrain	5825	21525	742985.6	206.4

number of gates identified by pyCBC, overall our gate-strain removes much less data: ~4–9% of what is removed by pyCBC. For instance of the 3124 O1-H1 SFTs, 686 are affected by one or more glitches which are gated with a total of 799 gates and 0.23 h of time-domain data lost. In comparison, pyCBC gating results in a slightly lower number of affected SFTs (667) but 3.43 h of lost data. For the L1 detector both methods find roughly ~180 O1 SFTs to be affected by glitches, where pyCBC gating removes 0.86 h of time-domain data while gatestrain removes 0.08 h.

The noise floor of gated data is lower than that of ungated data. Figure 4 compares the amplitude spectral density (ASD) of gated data and ungated data.

The actual improvements differ between detectors and runs: an improvement of a factor greater than 3 is seen in O1-H1 data, in the highest sensitivity region in frequency, and of ~4% in O1-L1 data. In O2 data gating significantly decreases the noise floor below 60 Hz in both detectors, and for H1 yields an appreciable decrease in the 100–450 Hz range.

We demonstrate the gain in sensitivity in continuous-wave searches with a Monte Carlo simulation where we consider 500 simulated continuous-wave signals with frequency between 20–1000 Hz distributed log-uniformly. The amplitude of the signals is such that they are clearly visible in the search results. The signals are added to the real data in the time domain. The data is then treated as it would be treated for a search, i.e., it is gated and Fourier-transformed in chunks to yield the SFTs. We perform a perfectly matched single-template $\mathcal{F}$-statistic search [28] using these SFTs, from non-gated and gated data. We compare the results in Fig. 5. An overall positive effect of gating can be seen, with a relative increase in detection statistic of up to 33% for H1 O1 data. Since gating lowers the noise level more in the low-frequency region, the signal-recovery improves more for low-frequency signals than for higher-frequency signals. The pyCBC -gating results are comparable to the gatestrain results, so in Figure 5 we only show the gatestrain results.

B. Gating and line-cleaning in the presence of glitches

Gating also mitigates artefacts introduced by glitches at the frequencies cleaned-out in the frequency domain. The line-cleaning procedure used in many continuous-wave searches substitutes the data at frequency bins that have been flagged to harbor disturbances, with Gaussian noise. In these bins fake SFT data is created with a standard deviation consistent with the noise level estimated based on the real data, in nearby-frequency bins. If the data in these nearby bins is quite Gaussian, the fake noise will look like a realization of noise from the nearby bins. But if the nearby noise has significant non-gaussian contributions, the fake noise will not look at all like the noise in the nearby bins, and in the presence of loud glitches, it will be higher. The
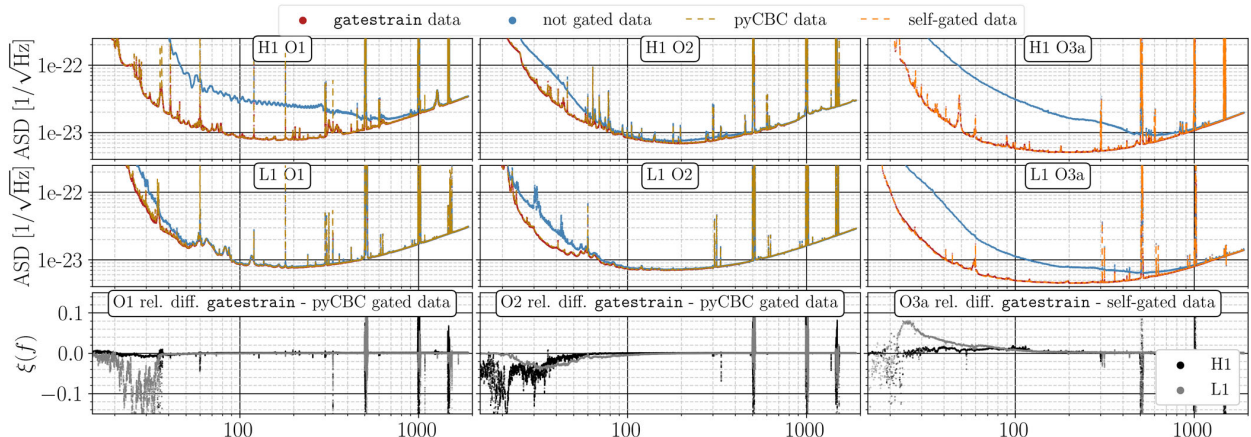
20

FIG. 4. The upper rows show the average amplitude spectral densities (ASDs) of H1/L1 data from the O1, O2 and O3a runs, before and after removal of glitches with our `gatestrain`-method. The ASD is estimated as the square root of the arithmetic mean over all timestamps, bin per bin, of the power spectral density. For O1 and O2 we also plot the pyCBC-gated data in gold, and for O3 the ASD of the `self-gated` data [27] in orange. `gatestrain` achieves a noise floor level comparable to other methods, apart for L1 O3a data between 20 Hz and 50 Hz. The relative difference of `gatestrain`-gated data and `pyCBC`-/ `self-gated` data $\xi(f) = (ASD_{\texttt{gatestrain}} - ASD_{\texttt{self-gated,pyCBC}})/ASD_{\texttt{self-gated,pyCBC}}$ is shown in the lower row respectively.

gating removes these non-Gaussian contributions and, with them, this type of problem. This is illustrated in Fig. 6 that shows how the noise floor of the cleaned data is greatly reduced and that the gating before the line cleaning allows for the lines to be removed without producing other spectral artefacts.

C. Gating the O3 LIGO data

The first six months of the O3 data (O3a) were publicly released shortly before the initial submission of this paper. This data presents multiple families of glitches that have required substantial effort by LIGO in order to work-around with an ad-hoc gating procedure [29]. The basic algorithm

is called `self-gating` and it is described in [27]. The resulting gates are released with [27].

We apply our `gatestrain` to the O3a public data with minimal changes in parameters with respect to the O1/O2 data, in order to allow for a slightly more aggressive gating. This is justified because the O3 data is significantly more glitchy than the data from the two previous Advanced
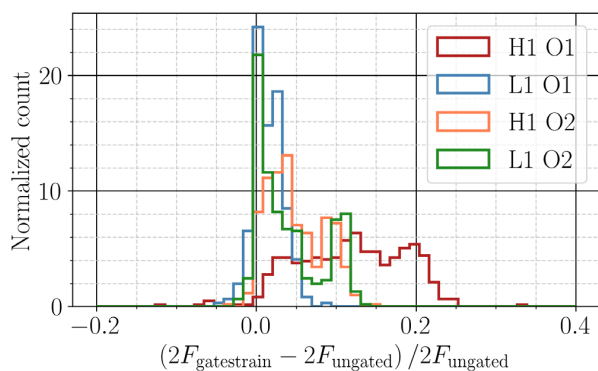

FIG. 5. Relative differences in the detection statistic $2F$ of 500 recovered simulated signals in `gatestrain`-gated versus non-gated (left panel) SFTs. It can be seen that gating has an overall positive effect which varies depending on detector and observation run.
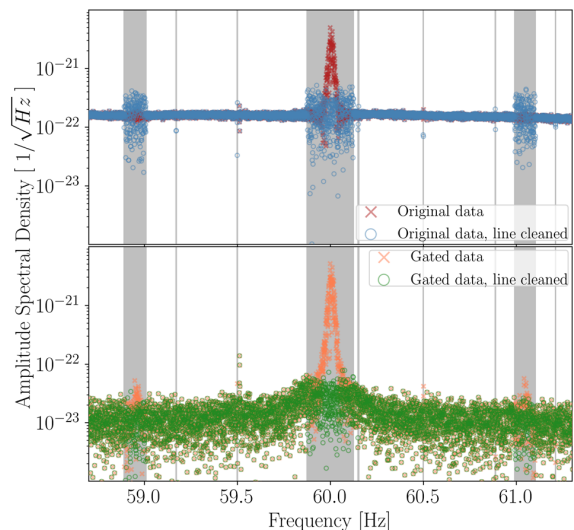

FIG. 6. Amplitude spectral density (ASD) of an SFT disturbed by a loud glitch before and after gating. The top panel shows the ASD of the *original data* before and after line cleaning. The lower panel shows the ASD of the *gated data* with and without line cleaning. The noise floor is greatly reduced thanks to the gating and then line cleaning removes the peaks without the discontinuities evident in the upper plot.

LIGO runs. At every iteration we change H^{high} by 11, rather than 10; we set the smallest H^{high} threshold to be H^{low}; we reduce the threshold for the gating result-check $R^{th} : 1.05 \rightarrow 1.01$.

Table II shows how many SFTs are affected by gates, and how many gates the `self-gating` method and our method produce. It also shows how much data is zeroed-out as a result of these gates. There is a caveat: [27] exclude SFTs with a total gate duration longer than 30 s. The data zeroed-out in this way is included in the count given in the last two columns of Table II. In order to make a fair comparison we also adopt this criterium and zero-out the entire SFT when our gate duration is longer than 30 s. We include this data in the zeroed-out count in the last columns of Table II for the `gatestrain` method. With this convention `gatestrain` preserves 222 h of data that the `self-gating` gates instead remove.

`gatestrain` achieves a more precise removal: Only 14(402) SFTs in H1 (L1) are excluded due to long gate duration when using `gatestrain`, while [27] exclude 61(773) in H1 (L1), respectively. On "good data," i.e., on SFTs that are free of long-duration gates, the gain is less dramatic, but it is still significant: `gatestrain` removes 3.4 h (10.1 h) in H1 (L1), whereas [27] exclude 8.2 h (13.5 h).

In the frequency range $\approx$20–50 Hz, the amplitude spectral density of the L1 data gated with `self-gating` is lower by $\leq 9\%$ with respect to the L1 data gated with our method. Elsewhere the performance of the two methods is comparable, as shown in Fig. 4.

We also recover the fifteen isolated continuous-wave hardware injections below 2 kHz [30] using the $\mathcal{F}$-statistic with comparable efficiency in both datasets. Like [27] we too could not recover the signal at 12.34 Hz due to the high noise in this frequency range. The relative gain in detection statistic with respect to [27] is a few percent.

We publish our O3a gates in the Supplemental Materials and at [12].

V. DISCUSSION

In this paper we present a new method to remove non-Gaussian noise transients in an overall fairly well-behaved noise background. The method extends the gating implementation by [16] with two main novelties: (i) the duration of glitches is measured and only the data affected by the glitch is removed (ii) the amplitude threshold that defines glitch-data is not fixed, but rather it is adaptive and is iteratively changed during the gating procedure.

As shown in Fig. 1, there is no single typical glitch duration. Glitches come in various sizes and durations, and the glitch populations change from one run to the next. While the `pyCBC` method of [16] is proven perfectly adequate for compact-binary-coalescence searches in O1

and O2, for continuous-wave searches our `gatestrain` method keeps more data untouched. For both H1 and L1 the number of gated SFTs increases slightly compared to the 16 s fixed-gate-duration of `pyCBC` (as used in [19]). On the other hand, `gatestrain` removes less than 10% of the data that `pyCBC` removes.

We note that the `pyCBC` gate duration $\tau_{dur} = 16$ s that we used in [19] is cautiously long and thus unsurprisingly more data is lost. In Appendix we show what happens with $\tau_{dur} = 3$ s and $\tau_{dur} = 0.25$ s. While shorter `pyCBC` gate durations lead to a decrease in the amount of lost data, the noise-level decrease may be adversely affected.

Below ~ 1 kHz the recovery of hardware and software injections shows SNR improvements compared to not gating and consistent with what is observed with `pyCBC` gating. No negative effect are recorded apart from regions of loud lines, e.g., violin modes, calibrations lines and power mains. These are cleaned anyway.

Unlike `pyCBC` which is written in Python, `gatestrain` is written in C. It is developed as part of LALSuite and leverages LALSuite functions and methods; it can be found in the LALSuite fork [12] under the name of `lalapps_gateStrain_v1`. It takes $\sim$30–40 s for an instance of `gatestrain` to produce a gated SFT in the frequency range 10 to 2000 Hz, lasting 1800s. The input data are gravitational-wave frame (gwf) files of the public data release. The outputs are gated SFTs or gated gravitational-wave frame (gwf) files and optionally ungated SFTs.

Although other gating methods exist [18], they are utilized within specific frameworks, the software is not publicly available and the input data is not the standard gravitational-wave frame (gwf) format. Our `lalapps_gateStrain_v1` works within the general LIGO Algorithm Library framework, and could be in fact be merged in the official LALSuite repository.

We present the results in the context of continuous gravitational-wave searches, however the method is valid for other searches, e.g., transient searches and stochastic background searches. We analyze the data around the eleven compact-binary coalescence gravitational-wave events of the GWTC-1 catalog. Not one was gated with out-of-the-box `gatestrain`-method. The glitch near GW170817 was automatically detected starting 1.09s before the event, 7.41 ms off of LIGO's gate mid-time. Our `gatestrain` applied a 92.65 ms gate with 0.25 s taper to each side in comparison to LIGO's gate with 0.2 s duration and 0.5 s taper [25,31]. The loud GW150914 signal produces a peak in $h_w(t) \sim 1.5\sqrt{Hz}$ which is well below the lowest H^{high} threshold of $\approx 6\sqrt{Hz}$. The weakest GW150914-like signal that would trigger our gating in O1 data is more than four times stronger than GW150914. This means that even though with the threshold settings described, the gating is very unlikely to remove a signal, as the detectors become more sensitive, and depending on the

type of search carried out, the threshold levels need to be evaluated. On the other hand, as the rate of detectable short-duration signals increases, an efficient method that automatically excises the disturbed portions of the data, becomes even more important.

Borne out of the desire to generalize the methodology that we had used on O1 data in continuous wave searches, and successfully applied to O2 data, our gating procedure successfully gates the O3 data, achieving a much smaller data loss than the LIGO gates with the same spectral noise improvement. Our procedure removes less than half the data compared to the *ad hoc* `self-gating` procedure of [27].

Since the gated data remains a small fraction of the total dataset, the impact of the more efficient gating on the detection statistic is small and the very loud hardware-injected fake signals present in the LIGO data for validation purposes, are recovered with comparable values of the detection statistic in both `gatestrain` data and `self-gated` data. The benefits of our method for gating is that it does not require ad-hoc time-consuming studies and careful tuning for every new dataset and every new family of glitches that appears.

Thanks to its adaptive algorithm, with practically no tuning, we were able to determine the O3a gates in less than a week. We make our tool available together with the O3a gates in the Supplemental Material [12], for others to employ in their analyses of LIGO O3 data. [12] will be updated with the O3b gates, as soon as that data becomes public.

APPENDIX: `pyCBC` GATING WITH A 3 s AND 0.25 s GATE DURATION

In this paper we compare the performance of our `gatestrain` with that of the `pyCBC` method with a fixed gate duration of 16 s. The reason is that, in absence of nonfixed duration gating procedures, 16 s is the `pyCBC` gate duration that was used in previous continuous-wave searches [19]. The typical `pyCBC` gate duration for

TABLE III. Amount of gated O1 data with different `pyCBC` gate-duration values. 0.25 s is the parameter value used in recent compact-binary-coalescence searches; 16 s is the value that was used in previous continuous-wave searches [19].

	DURATION OF GATED DATA			
	`pyCBC` 16 s [h]	`pyCBC` 3 s [h]	`pyCBC` 0.25 s [h]	`gatestrain` [h]
H1	3.43	0.68	0.06	0.23
L1	0.86	0.17	0.02	0.08

compact-binary-coalescence searches is 0.25 s [20,22,23] (or 0.125 s in O3 [21,24]). So, while our 16 s choice removed more data than a transient signal search would remove, it still removed a very small portion of the data and did not impact the continuous-wave search sensitivity. We compare performance with `pyCBC`-gating with 3 s or 0.25 s gate-duration. This is shown in Table III for the O1. Not surprisingly less data is lost with the 3 s gate duration and even fewer with 0.25 s gate duration. With 3 s gate-duration `pyCBC` loses ∼3 times more data in comparison to `gatestrain` while a 0.25 s gate duration leads to a loss four times smaller with `pyCBC` than with `gatestrain`. However, as shown in Fig. 7, the noise level does not change above 40 Hz, but below 40 Hz in L1 it increases significantly with respect to the data gated with 3 s gate duration. This indicates that there are families of glitches, with long tails, whose structure is not well captured by a fixed-duration gating procedure.
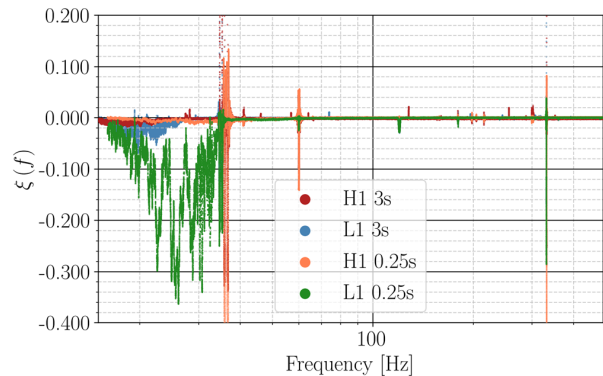

FIG. 7. $\xi(f)_{0.25\text{ s},3\text{ s}} = (\text{PSD}_{16\text{ s}} - \text{PSD}_{0.25\text{ s},3\text{ s}})/\text{PSD}_{0.25\text{ s},3\text{ s}}$, relative difference in the PSD when using data gated with a 0.25 s or 3 s fixed-duration gate and with a 16 s gate duration (`pyCBC`). The 16 s gate duration produces overall a slightly lower noise floor and up to 7% lower in the lowest frequency range. We note that the outliers are associated with spectral lines and these are cleaned-out with a separate procedure. The gating procedure aims at lowering the noise floor.

[1] D. Davis *et al.* (LIGO Collaboration), LIGO detector characterization in the second and third observing runs, Classical Quant. Grav. **38,** 135014 (2021).

[2] F. Robinet, N. Arnaud, N. Leroy, A. Lundgren, D. Macleod, and J. McIver, Omicron: A tool to characterize transient noise in gravitational-wave detectors, SoftwareX **12,** 100620 (2020).

[3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, Classical Quant. Grav. **37,** 055002 (2020).

[4] C. Pankow, K. Chatziioannou, E. A. Chase, T. B. Littenberg, M. Evans, J. McIver, N. J. Cornish, C. J. Haster, J. Kanner, V. Raymond *et al.*, Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817, Phys. Rev. D **98,** 084016 (2018).

[5] J. C. Driggers *et al.* (LIGO Scientific Collaboration), Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO, Phys. Rev. D **99,** 042001 (2019).

[6] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, Classical Quant. Grav. **33,** 134001 (2016).

[7] J. L. McIver, The impact of terrestrial noise on the detectability and reconstruction of gravitational wave signals from core-collapse supernovae, UMASS-539 (2015), 10.7275/7537749.0.

[8] O. J. Piccinni, P. Astone, S. D'Antonio, S. Frasca, G. Intini, P. Leaci, S. Mastrogiovanni, A. Miller, C. Palomba, and A. Singhal, A new data analysis framework for the search of continuous gravitational wave signals, Classical Quant. Grav. **36,** 015008 (2019).

[9] R. Abbott *et al.* (LIGO Scientific, Virgo and KAGRA Collaborations), Upper limits on the isotropic gravitational-wave background from Advanced LIGO's and Advanced Virgo's third observing run, Phys. Rev. D **104,** 022004 (2021).

[10] M. A. Papa, J. Ming, E. V. Gotthelf, B. Allen, R. Prix, V. Dergachev, H. B. Eggenstein, A. Singh, and S. J. Zhu, Search for continuous gravitational waves from the central compact objects in supernova remnants Cassiopeia A, Vela Jr., and G347.3–0.5, Astrophys. J. **897,** 22 (2020).

[11] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, Einstein@Home all-sky search for continuous gravitational waves in LIGO O2 public data, Astrophys. J. **909,** 79 (2021).

[12] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevD.105.022005, for the source code of the gating application and the timestamps of the O3 gates in machine-readable format.

[13] Y. Zhang, M. A. Papa, B. Krishnan, and A. L. Watts, Search for continuous gravitational waves from Scorpius X-1 in LIGO O2 data, Astrophys. J. Lett. **906,** L14 (2021).

[14] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search in early O3 LIGO data for continuous

gravitational-wave signals from unknown neutron stars in binary systems, Phys. Rev. D **103,** 064017 (2021).

[15] P. B. Covas *et al.* (LSC Collaboration), Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO, Phys. Rev. D **97,** 082002 (2018).

[16] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, Classical Quant. Grav. **33,** 215004 (2016).

[17] P. Astone, S. Frasca, and C. Palomba, The short FFT database and the peak map for the hierarchical search of periodic sources, Classical Quant. Grav. **22,** S1197 (2005).

[18] P. Leaci, P. Astone, M. A. Papa, and S. Frasca, Using a cleaning technique for the search of continuous gravitational waves in LIGO data, J. Phys. Conf. Ser. **228,** 012006 (2010).

[19] J. Ming, M. A. Papa, A. Singh, H. B. Eggenstein, S. J. Zhu, V. Dergachev, Y. M. Hu, R. Prix, B. Machenschalk, C. Beer *et al.*, Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr. and G347.3, Phys. Rev. D **100,** 024063 (2019).

[20] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, Phys. Rev. X **9,** 031040 (2019).

[21] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run, Phys. Rev. X **11,** 021053 (2021).

[22] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, Astrophys. J. **872,** 195 (2019).

[23] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, Astrophys. J. **891,** 123 (2020).

[24] A. H. Nitz, C. D. Capano, S. Kumar, Y. F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, Astrophys. J. **922,** 76 (2021).

[25] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, SoftwareX **13,** 100658 (2021).

[26] B. Allen and G. Mendell, SFT Data Format Version 2Specification (2004), https://dcc.ligo.org/LIGO-T040164/public.

[27] J. Zweizig and K. Riles, Information on self-gating of h(t) used in O3 continuous-wave searches (2021), Technical Document T2000384 available at https://dcc.ligo.org/LIGO-T2000384/public.

[28] P. Jaranowski, A. Krolak, and B. F. Schutz, Data analysis of gravitational—wave signals from spinning neutron stars. 1. The signal and its detection, Phys. Rev. D **58**, 063001 (1998).

[29] J. Wang and K. Riles, LIGO Scientific Collaboration, A Semi-Coherent Directed Search for Continuous Gravitational Waves from Supernova Remnants in the LIGO O3 Data Set, presentation at the April 2021 APS meeting, session Y16.00007.

[30] See table of hardware injections in O3a at https://www.gw-openscience.org/O3/O3April1_injection_parameters.

[31] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, Phys. Rev. Lett. **119**, 161101 (2017).

Density-clustering of continuous gravitational wave candidates from large surveys

This publication describes a new method in the first crucial post-processing process called *clustering*.

The project started with the tuning of the adaptive clustering method [92] to cluster the initial candidates from the Einstein@Home O2 all-sky search [93] with reasonably satisfactory results. Simultaneously S. Walsh, at the time a postdoctoral researcher in the group, was looking into clustering with machine learning, which the author at some point took over. In investigations on the density of signals, versus the density of background noise, the author identified the density of candidates in parameter space as a possibly better detection criterion than the detection statistic. Based on this idea, the author designed the core methodology. The initial method was thus developed ad-hoc by the author. The author implemented an improved method into the LALSuite-framework. The final implementation benefited from input from the other authors of this paper, including a number of computational-optimization aspects.

The optimization of density clustering and adaptive clustering for the paper, and the optimization for the O2 and O3 all-sky search were done by the author [93][94, in preparation]).

The paper was written by the author with helpful comments from M. A. Papa. All figures were produced by the author.

Density-clustering of continuous gravitational wave candidates from large surveys

B. Steltner[1,2,*] T. Menne,[1,2] M. A. Papa[1,2,3] and H.-B. Eggenstein[1,2]

[1]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute),*
Callinstrasse 38, 30167 Hannover, Germany
[2]*Leibniz Universität Hannover, D-30167 Hannover, Germany*
[3]*University of Wisconsin Milwaukee, 3135 N Maryland Avenue, Milwaukee, Wisconsin 53211, USA*

Searches for continuous gravitational waves target nearly monochromatic gravitational wave emission from, e.g., nonaxisymmetric fast-spinning neutron stars. Broad surveys often require us to explicitly search for a very large number of different waveforms, easily exceeding $\sim 10^{17}$ templates. In such cases, for practical reasons, only the top, say $\sim 10^{10}$, results are saved and followed up through a hierarchy of stages. Most of these candidates are not completely independent of neighboring ones, but arise due to some common cause: a fluctuation, a signal, or a disturbance. By judiciously clustering together candidates stemming from the same root cause, the subsequent follow-ups become more effective. A number of clustering algorithms have been employed in past searches based on iteratively finding symmetric and compact overdensities around candidates with high detection statistic values. The new clustering method presented in this paper is a significant improvement over previous methods: it is agnostic about the shape of the overdensities, is very efficient and it is effective: at a very high detection efficiency, it has a noise rejection of 99.99%, is capable of clustering two orders of magnitude more candidates than attainable before and, at fixed sensitivity it enables more than a factor of 30 faster follow-ups. We also demonstrate how to optimally choose the clustering parameters.

I. INTRODUCTION

Continuous gravitational waves are long-lasting signals that may come from fast-spinning nonaxisymmetric neutron stars, unstable r-modes [1,2], the fast inspiral of dark-matter objects [3,4] or emission from clouds of axionlike particles around black holes [5,6]. Unlike the short-lived signals stemming from the mergers of compact binary objects [7–13], continuous gravitational wave have thus far eluded any detection, due to their strength being orders of magnitudes smaller than that of binary merger signals. The detection of continuous gravitational waves will open a new field of gravitational wave astronomy, may probe the fundamental nature of gravity [14,15] and unlock unprecedented information on neutron star interiors [16–18]. For these reasons researchers tirelessly search for continuous gravitational wave signals [19]. Broad surveys using

months of data pose phenomenal challenges. We present here a new efficient method to identify the most promising candidates from broad parameter-space continuous waves surveys.

Independently of the emission mechanism, continuous gravitational waves are expected to be nearly monochromatic signals at the source, that due to the relative motion with respect to the source, appear to us on Earth to be frequency- and amplitude- modulated. Searches employ template parametrized by signal frequency, frequency derivatives, and source position, with $\sim 10^{17}$ template waveforms for observations lasting months. For template banks that are this big, typically only the top results are saved—say the $\sim 10^{10}$ results with the highest detection statistic values. Even though at this stage most of the results are not statistically significant, they are referred to as "candidates."

The candidates are followed up with a series of searches at increasing sensitivity. The signal-to-noise ratio of a signal increases from one stage to the next in a well-defined way, whereas noise does not, and this allows us to weed out noise candidates in the follow-ups [20–22]. Each follow-up search considers not only the candidates' parameters but a parameter-space region around each candidate. So if every candidate were to be followed up independently, the points in parameter space around nearby candidates

*[*]benjamin.steltner@aei.mpg.de

would be searched more than once, resulting in a waste of computing resources and aggravating an already challenging problem. The core idea of clustering is to avoid this by identifying candidates likely due to the same root cause, bundling (*clustering*) them and considering them as a single entity in follow-up studies. Clustering is hence an important step in the postprocessing of the results because it organizes and reduces the $\sim 10^{10}$ candidates to a more useful and manageable set of $\approx$ independent $\sim 10^{6}$ candidates.

Each cluster is represented by the parameters of the so-called *seed* candidate and by a *containment region*. The latter measures how far from the seed associated with a signal, the true signal parameters are. In follow-up studies the entire containment region around each seed is surveyed. The containment region is the same for all seeds and it is determined statistically, such that it holds for a very large fraction ($> 99\%$) of signals, across the parameter space.

It has also been observed that a threshold on the minimum number of candidates in a cluster is effective at discarding noise-clusters. With a fixed computing budget for follow-ups, fewer candidates means that freed-up computational capacity can be used on additional, lower significance candidates which translates in deeper and more sensitive searches.

The most compute-intensive continuous waves searches have been carried out since the mid 2000s using idle cycles donated by the general public, through the volunteer distributed computing project Einstein@Home[1] [23–25]. The massive computational power that we can harvest today amounts to several Pflops, sustained 24×7, and enables us to investigate over 10^{19} waveforms.

For this reason clustering procedures have been in use for a long time: One of the first nontrivial clustering procedures is box-clustering [26,27], which dates back to nearly a decade ago. More recently a more flexible adaptive clustering technique has been used [28] which however does not converge fast enough when used on many data points. This is a significant drawback, as we want to set lower thresholds, which means considering more candidates in the follow-ups. Attempts to use machine-learning for clustering have been successful for directed searches, but not for all-sky searches [29,30].

We present here the new *density clustering* algorithm, able to process orders of magnitude more candidates than previous clustering strategies at comparable, if not lower, computing cost. We show how to choose the clustering parameters, and demonstrate its performance on real data. We concentrate on clustering results from very large template banks—with over 10^{16} points—and hence refer to the Einstein@Home results, but this method can also be employed in less challenging environments.

The paper is organized as follows: In Sec. II we describe the input data; in Sec. III the method itself; in Sec. IV the

choice of the clustering parameters; in Sec. V the implementation; in Sec. VI the method is compared with adaptive clustering under realistic conditions, i.e., by applying it to the data of the Stage 0 results of the Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data of the second observation run (O2) [22,31].

II. INPUT DATA TO CLUSTERING

Clustering works on a set of candidates, i.e., selected results from a search. A candidate is described by the values of the template that produced the detection statistic result, and the detection statistic result. For an all-sky search including up to second-order spin-down parameters, a generic candidate i is of the form

$$(f_i, \dot{f}_i, \ddot{f}_i, \alpha_i, \delta_i, \chi_i), \tag{1}$$

where f indicates the signal-template frequency, α, δ the source sky position and χ the value of the detection statistic used for the original candidate ranking.

We illustrate clustering for these 5 dimensions; fewer or more dimensions are treated analogously.

Since continuous waves are modulated by the Earth's rotation and orbit around the Sun, the sky grids are set up in sky coordinates projected on the ecliptic plane, x_{ecl}, y_{ecl}. Therefore for clustering we convert for the candidates $(\alpha_i, \delta_i) \rightarrow (x_{\mathrm{ecl}\,i}, y_{\mathrm{ecl}\,i})$—see Eqs. (14) and (15) in [28] for the conversion between $(\alpha, \delta) \rightarrow (x_{\mathrm{ecl}}, y_{\mathrm{ecl}})$.

The sky grids are approximately uniform hexagonal grids on the ecliptic plane and are defined by the hexagon edge length d:

$$d(m_{\mathrm{sky}}) = \frac{1}{f} \frac{\sqrt{m_{\mathrm{sky}}}}{\pi \tau_E}, \tag{2}$$

with $\tau_E \simeq 0.021$ s being half of the light travel-time across the Earth and m_{sky} a constant which controls the resolution of the sky grid [22]. From Eq. (2) it is clear that the sky-grid density increases with frequency f.

III. DENSITY CLUSTERING

We bin the parameter space in equally spaced cells of size

$$\delta b = (\delta f, \delta \dot{f}, \delta \ddot{f}, \delta x_{\mathrm{ecl}}, \delta y_{\mathrm{ecl}}) \tag{3}$$

in each dimension. The $\delta f, \delta \dot{f}, \delta \ddot{f}$ are each an integer multiple of the search grid spacing. The sky grid has a hexagonal tiling, so the square tiling of the bins above does not match it. The bins are usually chosen to be large enough that this does not matter and the square covering greatly simplifies the binning and the identification of neighboring bins. The bin size is always a multiple of the hexagon side, so the bins shrink with increasing frequency as the sky-grid pixels, keeping the average number of candidates per bin the same.

We only consider candidates with detection statistic values above a threshold Γ_L. In each bin j we count the number of candidates $N_{occ,j}$ with parameters in that bin. Bins with $N_{occ,j} \leq N_{occ,\min}$ are discarded. $N_{occ,\min}$ is one of the clustering parameters and its optimal value depends on the search setup and on the bin size.

Among the surviving bins, we cluster together nearby ones, to create a cluster. The basic notion of vicinity is controlled by two parameters: N^j and N_c. A bin b_a is a neighbor of bin b_c if the distances k^j in integer bin spacings

$$b_a - b_c = (k^1 \delta f, k^2 \delta \dot{f}, k^3 \delta \ddot{f}, k^4 \delta x_{\text{ecl}}, k^5 \delta y_{\text{ecl}}) \qquad (4)$$

satisfy the following conditions:

$$\begin{cases} k^j \leq N^j \text{ with } j = 1, \ldots, M \\ \sum_{j=1}^{M} k^j \leq N_c, \end{cases} \qquad (5)$$

where M is the number of dimensions. The first condition sets the maximum distance in every dimension, whereas the second condition sets an overall maximum distance. With $M = 3$, $N_c = 1$ means that the two nearby bins have to share a face, $N_c = 2$ that they have to share an edge and $N_c = 3$ that they have to share a vertex. Default values are $N^j = 1$, equal for all j, and $N_c = M$.

Among the clusters from the previous step, we remove the ones with too few bins: $N_{\text{bins}} \leq N_{\text{bins,min}}$.

For each remaining cluster a representative candidate becomes the seed. The seed is by default the candidate with the highest detection statistic value (the loudest) of all candidates in the cluster. In noisier data it may make sense to look at the loudest candidate in the bin with the most candidates (densest bin) or the loudest candidate in the bin with the highest average over all detection statistic values of the candidates within that bin (loudest bin).

Finally all clusters with a seed with detection statistic value smaller than Γ_S are discarded. The process is illustrated in Fig. 1 for two dimensional, higher dimensions follow analogously.

An additional parameter can be used to mitigate binning effects: an overdensity of candidates may not be perfectly contained within one bin, but may extend across bin boundaries. For faint signals with just enough candidates to surpass the occupancy threshold $N_{occ,\min}$, this effect can make the difference between recovering a signal or not. Boundary effects can be partly mitigated by smoothing over bins, e.g., adding bin counts over neighboring bins or adding bin counts weighted with a Gaussian kernel. The overall impact of using smoothing procedures should be evaluated within the general framework of choosing the optimal clustering parameters, as described in the next section, but we will not explicitly consider it here.

IV. CHOOSING THE PARAMETERS OF THE CLUSTERING PROCEDURE

A number of parameters define the density clustering algorithm, and they are summarized in Table I. We choose the parameter values such that at fixed computational cost for the follow-up of the resulting seeds, the sensitivity of the clustering procedure is maximized. Below we describe how this optimization, yielding the values of the clustering parameters of Table I, is carried out.

The sensitivity of the clustering procedure is measured by the gravitational wave signal amplitude $h_0^{90\%}$ at which the detection efficiency ϵ of the clustering procedure is 90%, for signals with parameters in the search range. $h_0^{90\%}$ depends on the signal frequency like the amplitude spectral density of the noise $\sqrt{S_h(f)}$, so we maximize the quantity $\mathcal{D}^{90\%} = \sqrt{S_h(f)}/h_0^{90\%}(f)$, instead, that does not depend on frequency. $\mathcal{D}$ is also known as the sensitivity depth [27].

Since there is no way to predict the detection efficiency of the clustering procedure, we measure it with a Monte Carlo. We add fake signals from our target population to the real data, with amplitudes corresponding to a given value of $\mathcal{D}$. For each signal we perform the same search as the actual search, we cluster the results and produce seeds. If one of the seeds comes from the added signal, we consider the signal detected by the clustering procedure. The fraction of detected signals to total signals gives the detection efficiency at that sensitivity depth: $\epsilon(\mathcal{D})$. $\mathcal{D}^{90\%}$ is then

$$\epsilon(\mathcal{D}^{90\%}) = 90\%. \qquad (6)$$

For each clustering setup we estimate
(i) $\mathcal{D}^{90\%}$
(ii) the containment region (see Sec. I).

TABLE I. Parameters of density clustering in the order that they are employed.

Parameter	Function
Input threshold Γ_L	Discards candidates with detection statistic $\leq \Gamma_L$. Filters input candidates
Bin sizes δb	Binning
Smoothing	Smooth histogram or not
Occupancy threshold $N_{occ,\min}$	Discard bins with $N_{occ} \leq N_{occ,\min}$ candidates
Neighbor criterion, N^j and N_c	Defines what a neighbor is
Cluster-size threshold N_{bins}	Discard clusters with $N_{\text{bins}} \leq N_{\text{bins,min}}$ bins
Seed criterion	Loudest candidate in cluster, loudest in most-populated bin or in bin with highest average detection statistic
Output threshold Γ_S	Discards cluster whose seed has detection statistic $\leq \Gamma_S$. Reduces false alarms

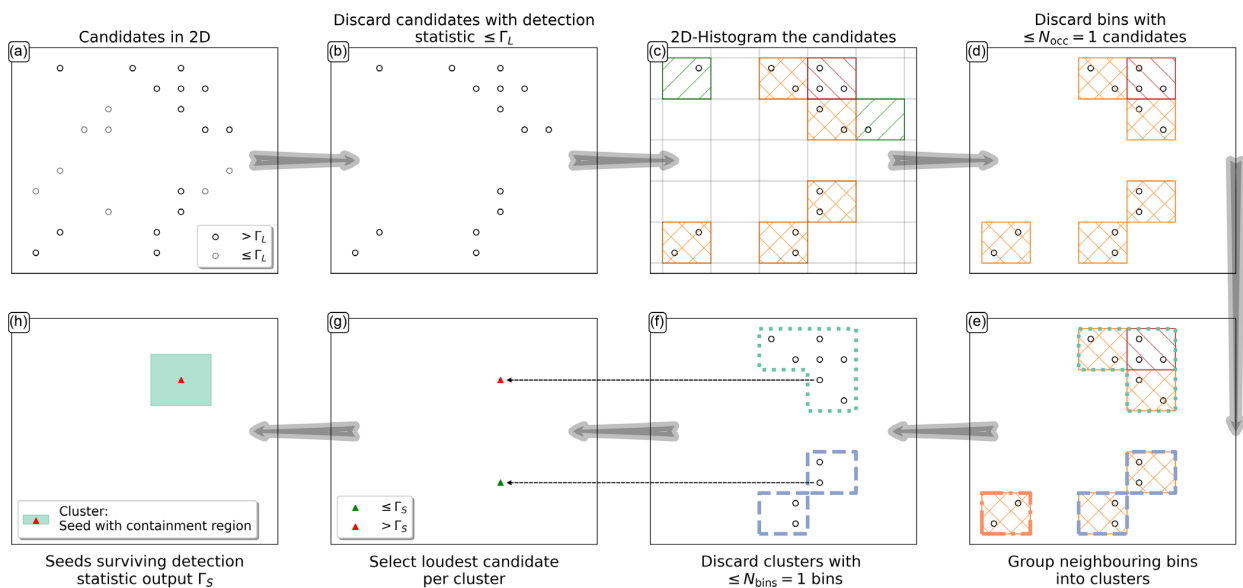FIG. 1. Schematic illustration of the main steps of density clustering.

(iii) the false alarm rate. This is done by running the clustering on a subset on the search results, at different frequencies. Since we operate in the regime of very rare signals, we take this as a measure of the false alarm.
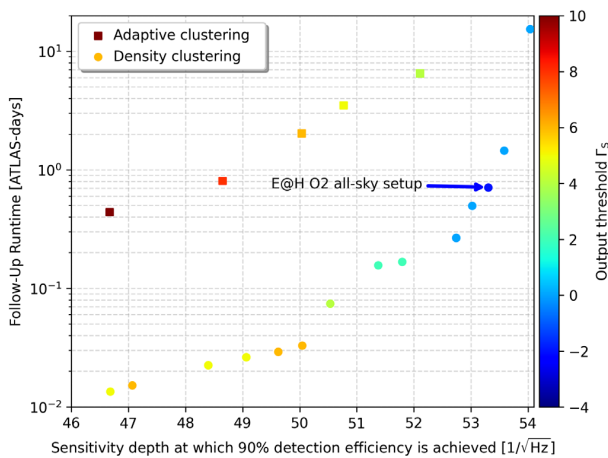

FIG. 2. Performance comparison between the previous clustering method, adaptive clustering, and density clustering. Each point represents a different clustering setup, used on the results of the Einstein@Home search [22]. To avoid excessive clutter we do not show all considered setups, but rather only those with runtime close to the smallest runtime at each $\mathcal{D}^{90\%}$. The color encodes the Γ_S threshold parameter value which illustrates the need to consider more candidates to achieve better sensitivities. The arrow indicates the density clustering setup chosen for the follow-up analysis reported in [22], which was the optimal for the clustering procedure at the time, under the constraint of maximum 1 ATLAS-day computing time for the first-stage follow-up.

For each clustering setup, from the number of expected seeds, the containment region and the timing model of our software [32], we estimate the computing cost of the first follow-up stage. This is illustrated in Fig. 2 for the results of the Stage-0 Einstein@Home search [22]. We can now identify the clustering setup that yields the highest $\mathcal{D}^{90\%}$, within the computing budget.

In principle one could optimize the follow-up search setup for each clustering setup. This would, however, be extremely expensive, and experience has shown that a setup choice guided by the sensitivity gain with respect to the previous stage, at accessible computing cost, lands a choice not significantly far from optimum. So we assume here that the follow-up setup is fixed.

V. IMPLEMENTATION

In the previous section we have described how the optimal combination of clustering parameters is identified. As we have seen, this requires a Monte Carlo in order to measure the false alarm and 90% detection-efficiency signal-amplitude $\mathcal{D}^{90\%}$, for every clustering setup.

For each setup we cluster $\gtrsim 2000$ result files corresponding to data with different fake signals—this is to determine $\mathcal{D}^{90\%}$. We cluster $\gtrsim 500$ search result files with no fake signals, in order to estimate the false alarm. These operations can be quite time-consuming, so we describe here how to reduce the computing cost of this step.

Einstein@Home search results typically come in files that cover a 50 mHz range of template frequencies, with size varying between a few MB to few GB, due to the different sky resolutions in the range 50–600 Hz. Each clustering instance uses as input one of these 50 mHz

results files. Since the time to cluster is $\ll$ the time it takes to load such a file, it is faster to load a results file, keep it in memory, and test different clustering setups.

Further savings are obtained by reusing intermediate results:

(i) we compute a histogram for a choice of Γ_L and δb, and reuse it to produce the bin counts for different values of $N_{occ,\min}$

(ii) similarly, for a choice of Γ_L, δb and $N_{occ,\min}$ from the bin counts we produce different clusters for different values of $N_{\mathrm{bins,min}}$

(iii) for each cluster different seeds are produced, based on different seed-selection criteria, e.g., the loudest cluster candidate, the loudest in the densest bin, the loudest in the bin with the highest average detection statistic (we call this the "loudest bin").

(iv) finally, each seed, and with it the whole cluster, may be discarded depending on the value of Γ_S.

With this scheme, testing a single clustering setup costs (on average, over many setups) just under a second, with more than half the time spent on the initial histogram and clustering. In order to reduce memory usage, the candidates are internally addressed only by an id. For thresholding on Γ_S and for computing the containment region, the actual seed parameters must be retrieved. This operation accounts for another 20% of the computing time. The remaining time is due to fluctuations in these estimates due to varying number of seeds and the initial overhead.

Given the computing-load profile described above, we parallelize the work among different independent processors, with each processor working only with a single results file and several $(\Gamma_L, \delta b)$-combinations. Say we have 2500 result files, 1000 $(\Gamma_L, \delta b)$-combinations and 1000 combinations of the remaining parameters, each processor analyzes 100 $(\Gamma_L, \delta b)$-combinations, exhausting all 1000 combinations of the remaining parameters. Hence, with 10 processes per result file, 25000 processes are spawned in total.

Using the large-capacity and fast-loading hdf5 and FITS file formats, and a HDD-raid configuration results file server, testing a single $(\Gamma_L, \delta b)$-combination and all 1000 combinations of the remaining parameters, takes ≈ 0.26 h. Thus one processor exhausting 100 $(\Gamma_L, \delta b)$-combinations takes $\approx$ a day. On the ATLAS cluster[2] using 25000 parallel processes the full testing of 1000×1000 setups is carried out in a day.

VI. PERFORMANCE ON E@H O2 ALL-SKY

We compare our density clustering with the adaptive clustering [28] on the results of the Stage-0 Einstein@ Home O2 all-sky search [22].

---

[2]ATLAS is the super-computer cluster at the MPI for Gravitational Physics in Hannover: https://www.atlas.aei.uni-hannover.de/.

We characterize the detection efficiency on a set of $\sim$2900 fake signals from the target source population of the search: signals with spin-frequencies uniformly distributed; spin-downs log-uniform distributed and all other parameters distributed uniformly: orientation $\cos \iota \in [-1, 1)$, polarization angle $\psi \leq |\pi/4|$, sky position $0 \leq \alpha \leq 2\pi$ and $-1 \leq \sin \delta \leq 1$. The signal amplitude h_0 ranges from loud to faint signals with $\sim$1000 signals too faint to be detectable by either method.

The results of the procedure described in the previous section in order to identify the optimal density clustering parameters, are shown in Fig. 2. We compare with the results for the optimal parameter choice for adaptive clustering.

The density clustering setup chosen in [22] with a first-stage follow-up runtime-cost of ≤ 1 ATLAS-day is $\approx$10% more sensitive than the adaptive clustering setup at the same computing cost. In continuous gravitational wave searches a 10% improvement, solely due to a better search method, is a big gain.

Perhaps more immediately impressive is the fact that at fixed sensitivity, density clustering enables follow-ups that are a factor of $\gtrsim$30 faster than previous methods.

This gain can be reinvested in deeper follow-ups by using a lower Γ_S, albeit the gain in practice is limited by the steep increase in computing cost for $\Gamma_S \lesssim 4$. With a threshold $\Gamma_S = -3.7$, density clustering is able to process two orders of magnitude more candidates than with a threshold $\Gamma_S = 4$, whereas adaptive clustering could not be used at all.

The performance of the adaptive clustering was characterized in [28] by the detection efficiency and the noise rejection NR defined as

$$\mathrm{NR} := 1 - \frac{N_{\mathrm{out}}}{N_{\mathrm{in}}}, \qquad (7)$$

where N_{in} is the number of candidates above the threshold Γ_S and N_{out} is the number of seeds produced by the clustering procedure.

With a threshold of $\Gamma_S \geq 4$ adaptive clustering and density clustering achieve similar performance with NR $\geq$ 99% and detection efficiencies above 98%. At lower thresholds, adaptive clustering does not converge in weeks of runtime, indicating that the method struggles to identify over densities due to faint signals. Density clustering, instead, can probe threshold values as low as -3.7, still achieving NR $\geq$ 99.99% and attaining a very respectable detection efficiency (now at the 85% level) on a set that includes very faint signals with detection statistic values $\in [-3.7, 4]$, which are much harder to find.

The density clustering setup chosen in [22] has a binning of 300 and 290 search bins in frequency and spindown respectively, and the square lattice on the projected ecliptic plane with edge lengths of $5 \cdot d(m_{\mathrm{sky}} = 0.008)$. A cluster is

formed if there is at least one bin with more than three candidates. Of all candidates in the cluster the loudest becomes the seed.

VII. CONCLUSION AND OUTLOOK

We have presented a new, fast, and efficient clustering method—density clustering—for continuous gravitational wave search postprocessing.

Density clustering works by identifying overdensities of candidates in parameter space: clusters are purely build on candidates' closeness to each other and the detection statistic value is nearly irrelevant. This result may be somewhat surprising because the detection statistic ranks results based on the likeliness of they originating from a signal. However, some of our faintest—but still recoverable—signals show detection statistic values at which there are thousands to millions of louder candidates purely from noise. Our results show that in this regime overdensities are a better detection criterion than the significance given by the detection statistic value alone, even in Gaussian noise. This is probably due to the fact that the search is a semicoherent search.

The overdensities are uncovered by binning the parameter space and this is performed in one pass instead of the previously employed slower iterative procedures. The clustering step is thus largely independent of the number of input candidates, and this allows to process orders of magnitude more candidates with comparable computing resources, probing deeper into the noise.

Until now Einstein@Home searches have returned about $\mathcal{O}(10^4)$ candidates per work unit (e.g., [22]), which was more than adequate for what previous clustering algorithms could process. Density clustering can cluster orders of magnitude more candidates, which means that more results can be inspected, allowing to recover fainter signals in upcoming searches.

The previous clustering method, adaptive clustering, assumes compact overdensities, whereas signals typically present X- or Y-shaped overdensities which are hard to capture (and practically impossible to predict). Density clustering is agnostic about the shape of the overdensities and for this reason it is significantly more effective at identifying even very weak signals.

A different approach of using machine learning for clustering was developed and applied to the Einstein@ Home O2 all-sky dataset in [29,30]. They cluster in $f, \dot{f}$ and achieve better sensitivity depths at fixed false alarms, but lack in sky localization to the point of clustering together candidates from "seemingly unrelated sky positions" [30]. This means that a follow-up would entail searching over the whole sky, whereas density clustering restricts the sky position to a patch of ∼9% to 0.01% of the full sky,

depending on the frequency, between ∼20 Hz to 600 Hz, respectively. Even with the smaller uncertainties in $f, \dot{f}$ and only half the false alarms [30], the computational cost of their approach is higher by one order of magnitude compared to density clustering. They propose to generalize to include sky, and the results will be interesting to see.

Clustering is not a problem unique to gravitational wave astronomy, and a number of generic clustering methods exist. For example k-means [33] is a clustering method widely used in a variety of applications including signal-, image- and text-processing, health, cybersecurity, machine learning and big data [34]. It works based on minimizing the cluster-occupants' distance to the cluster center. Limitations of k-means are that the number of clusters must be known *a priori* and clusters are assumed to be roughly spherical and similar size. Density-based clustering applications exist: for example DBSCAN [35,36] and its many generalizations, like, e.g., OPTICS [37] or HDBSCAN [38], identify overdensities generated by a minimum number of points within a given volume. They are, however, not suitable for the large number of points in our results, and they are not as efficient as density clustering on our data.

A major advantage of our approach is the versatility of the method. Density clustering can cluster in any combination of dimensions, so it is easily extendable to, e.g., third/higher order spindowns $\ddot{f}, \ldots$ or to the 5 additional orbital parameters for searches for neutron stars in binary systems. In these searches signal-template offsets in orbital parameters can be to some extent compensated by offsets in frequency- and derivative(s). This translates into correlations between different templates and results in more candidates due to the same root cause [39], making clustering all the more important. All-sky binary searches are computationally extremely expensive and so are the follow-ups. A first test of density clustering on the results-data from [40] showed promising results within a few hours of clustering in 6 dimensions $(f, \alpha, \delta, \tau_{\mathrm{asc}}, P_b, a)$, showcasing the flexibility and ease of use of the method presented here.

[1] P. D. Lasky, Gravitational waves from neutron stars: A review, Pub. Astron. Soc. Aust. **32,** e034 (2015).

[2] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Gravitational waves from hot young rapidly rotating neutron stars, Phys. Rev. D **58,** 084020 (1998).

[3] C. J. Horowitz and S. Reddy, Gravitational Waves from Compact Dark Objects in Neutron Stars, Phys. Rev. Lett. **122,** 071102 (2019).

[4] C. Horowitz, M. Papa, and S. Reddy, Search for compact dark matter objects in the solar system with LIGO data, Phys. Lett. B **800,** 135072 (2020).

[5] A. Arvanitaki, M. Baryakhtar, and X. Huang, Discovering the QCD axion with black holes and gravitational waves, Phys. Rev. D **91,** 084011 (2015).

[6] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, Characterizing the continuous gravitational-wave signal from boson clouds around Galactic isolated black holes, Phys. Rev. D **102,** 063020 (2020).

[7] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, Phys. Rev. X **9,** 031040 (2019).

[8] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, Phys. Rev. X **11,** 021053 (2021).

[9] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, arXiv:2111.03606.

[10] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, Astrophys. J. **872,** 195 (2019).

[11] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, Astrophys. J. **891,** 123 (2020).

[12] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, Astrophys. J. **922,** 76 (2021).

[13] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact-binary mergers, arXiv:2112.06878.

[14] M. Isi, A. J. Weinstein, C. Mead, and M. Pitkin, Detecting beyond-Einstein polarizations of continuous gravitational waves, Phys. Rev. D **91,** 082002 (2015).

[15] M. Isi, M. Pitkin, and A. J. Weinstein, Probing dynamical gravity with the polarization of continuous gravitational waves, Phys. Rev. D **96,** 042001 (2017).

[16] F. Gittins and N. Andersson, Modelling neutron star mountains in relativity, Mon. Not. R. Astron. Soc. **507,** 116 (2021).

[17] F. Gittins, N. Andersson, and D. I. Jones, Modelling neutron star mountains, Mon. Not. R. Astron. Soc. **500,** 5570 (2020).

[18] J. A. Morales and C. J. Horowitz, Neutron star crust can support a large ellipticity, arXiv:2209.03222.

[19] K. Riles, Searches for continuous-wave gravitational radiation, arXiv:2206.06447.

[20] M. A. Papa *et al.*, Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data, Phys. Rev. D **94,** 122006 (2016).

[21] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data, Phys. Rev. D **96,** 122004 (2017).

[22] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, Einstein@Home all-sky search for continuous gravitational waves in LIGO O2 public data, Astrophys. J. **909,** 79 (2021).

[23] BOINC, http://boinc.berkeley.edu/ (2020).

[24] D. P. Anderson, BOINC: A system for public-resource computing and storage, in *Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing (GRID04)* (IEEE Computer Society, Washington, DC, 2004), pp. 4–10.

[25] D. P. Anderson, C. Christensen, and B. Allen, Designing a runtime system for volunteer computing, in *Proceedings of the 2006 ACM/IEEE conference on Supercomputing* (Association for Computing Machinery, New York, NY, 2006), pp. 126–136.

[26] J. Aasi *et al.* (LIGO Scientific and VIRGO Collaborations), Directed search for continuous gravitational waves from the Galactic center, Phys. Rev. D **88,** 102002 (2013).

[27] B. Behnke, M. A. Papa, and R. Prix, Postprocessing methods used in the search for continuous gravitational-wave signals from the Galactic Center, Phys. Rev. D **91,** 064007 (2015).

[28] A. Singh, M. A. Papa, H.-B. Eggenstein, and S. Walsh, Adaptive clustering procedure for continuous gravitational wave searches, Phys. Rev. D **96,** 082003 (2017).

[29] B. Beheshtipour and M. A. Papa, Deep learning for clustering of continuous gravitational wave candidates, Phys. Rev. D **101,** 064009 (2020).

[30] B. Beheshtipour and M. A. Papa, Deep learning for clustering of continuous gravitational wave candidates II: Identification of low-SNR candidates, Phys. Rev. D **103,** 064027 (2021).

[31] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, The LIGO open science center, J. Phys. Conf. Ser. **610,** 012021 (2015).

[32] J. Ming, M. A. Papa, B. Krishnan, R. Prix, C. Beer, S. J. Zhu, H.-B. Eggenstein, O. Bock, and B. Machenschalk, Optimally setting up directed searches for continuous gravitational waves in Advanced LIGO O1 data, Phys. Rev. D **97,** 024051 (2018).

[33] S. Lloyd, Least squares quantization in PCM, IEEE Trans. Inf. Theory **28,** 129 (1982).

[34] M. Ahmed, R. Seraj, and S. M. S. Islam, The k-means algorithm: A comprehensive survey and performance evaluation, Electronics **9,** 1295 (2020).

34

[35] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96 (AAAI Press, Portland, USA, 1996), p. 226–231.

[36] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, Dbscan revisited, revisited: Why and how you should (still) use dbscan, ACM Transactions on Database Systems **42**, 1 (2017).

[37] M. Ankerst, M. M. Breunig, H.-P. Kriegel, and J. Sander, Optics: Ordering points to identify the clustering structure, SIGMOD Record **28**, 49 (1999).

[38] R. J. G. B. Campello, D. Moulavi, and J. Sander, Density-based clustering based on hierarchical density estimates, in *Advances in Knowledge Discovery and Data Mining*, edited by J. Pei, V. S. Tseng, L. Cao, H. Motoda, and G. Xu (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013), pp. 160–172.

[39] A. Singh, M. A. Papa, and V. Dergachev, Characterizing the sensitivity of isolated continuous gravitational wave searches to binary orbits, Phys. Rev. D **100**, 024058 (2019).

[40] P. B. Covas, M. A. Papa, R. Prix, and B. J. Owen, Constraints on r-modes and mountains on millisecond neutron stars in binary systems, Astrophys. J. Lett. **929**, L19 (2022).

[41] https://www.gw-openscience.org/

Einstein@Home all-sky search for continuous gravitational waves in LIGO O2 public data

The results of the Einstein@Home all-sky search in the public O2 Advanced LIGO data are presented.

The search was set up and run on Einstein@Home based on the work of others, but the post-processing of the results was led in all parts by the author, who also developed all necessary software.

Specifically the author identified the disturbed bands, carried out the clustering as detailed in Chapter 3, characterized all the follow-up steps and carried out the follow-up searches.

The author set up the new upper limit procedure, which, at the cost of being more expensive to run, produces very accurate upper limits over every half-Hz, factoring in all systematic errors, like e.g. the effect of line cleaning on a signal. The author then implemented the method and produced the upper limits on the gravitational wave amplitude and neutron star ellipticity. The methodological aspects were developed in collaboration with the supervisor.

The paper was written by M. A. Papa and the author. All figures but Fig. 6 and Fig. 7 were produced by the author.

The Stage 0 results and upper limit method of this search were used by A. Singh [91] to constrain the gravitational wave amplitude over a large and unexplored range of binary orbital parameters, namely objects in long-periodicity binaries. This would still appear as a signal in Stage 0 of the Einstein@Home search, albeit an isolated search.

CrossMark

Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data

B. Steltner[1,2] , M. A. Papa[1,2,3] , H.-B. Eggenstein[1,2] , B. Allen[1,2,3] , V. Dergachev[1,2], R. Prix[1,2], B. Machenschalk[1,2],
S. Walsh[1,2,3], S. J. Zhu[1,2,4], O. Behnke[1,2], and S. Kwang[3]
[1] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany; benjamin.steltner@aei.mpg.de,
maria.alessandra.papa@aei.mpg.de
[2] Leibniz Universität Hannover, D-30167 Hannover, Germany
[3] University of Wisconsin Milwaukee, 3135 N Maryland Avenue, Milwaukee, WI 53211, USA
[4] DESY, D-15738 Zeuthen, Germany

Abstract

We conduct an all-sky search for continuous gravitational waves in the LIGO O2 data from the Hanford and Livingston detectors. We search for nearly monochromatic signals with frequency $20.0 \text{ Hz} \leqslant f \leqslant 585.15 \text{ Hz}$ and spin-down $-2.6 \times 10^{-9} \text{ Hz s}^{-1} \leqslant \dot{f} \leqslant 2.6 \times 10^{-10} \text{ Hz s}^{-1}$. We deploy the search on the Einstein@Home volunteer-computing project and follow-up the waveforms associated with the most significant results with eight further search stages, reaching the best sensitivity ever achieved by an all-sky survey up to 500 Hz. Six of the inspected waveforms pass all the stages but they are all associated with hardware injections, which are fake signals simulated at the LIGO detector for validation purposes. We recover all these fake signals with consistent parameters. No other waveform survives, so we find no evidence of a continuous gravitational wave signal at the detectability level of our search. We constrain the h_0 amplitude of continuous gravitational waves at the detector as a function of the signal frequency, in half-Hz bins. The most constraining upper limit at 163.0 Hz is $h_0 = 1.3 \times 10^{-25}$, at the 90% confidence level. Our results exclude neutron stars rotating faster than 5 ms with equatorial ellipticities larger than 10^{-7} closer than 100 pc. These are deformations that neutron star crusts could easily support, according to some models.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Neutron stars (1108); Compact objects (288); Astronomy data analysis (1858)

Supporting material: tar.gz file

1. Introduction

Continuous gravitational waves are expected in a variety of astrophysical scenarios: from rotating neutrons stars if they present some sort of asymmetry with respect to their rotation axis or through the excitation of unstable r-modes (Owen et al. 1998; Lasky 2015); from the fast inspiral of dark-matter objects (Horowitz & Reddy 2019; Horowitz et al. 2020); through superradiant emission of axion-like particles around black holes (Arvanitaki et al. 2015; Zhu et al. 2020).

The expected gravitational wave amplitude at the Earth is several orders of magnitude smaller than that of signals from compact binary inspirals, but because the signal is long-lasting one can integrate it over many months and increase the signal-to-noise ratio (S/N) very significantly.

The most challenging searches for this type of signal are the all-sky surveys, where one looks for a signal from a source that is not known. The main challenge of these searches is that the number of waveforms that can be resolved over months of observation is very large, and so the sensitivity of the search is limited by its computational cost.

In this paper we present the results from an all-sky search for continuous gravitational wave signals with frequency f between 20.0 Hz and 585.15 Hz and spin-down $-2.6 \times 10^{-9} \text{Hz s}^{-1} \leqslant \dot{f} \leqslant 2.6 \times 10^{-10} \text{ Hz s}^{-1}$, carried out thanks to

the computing power donated by the volunteers of the Einstein@Home project.

The results from the Einstein@Home search are further processed using a hierarchy of eight follow-up searches, similarly to what was previously done for recent Einstein@-Home searches (Abbott et al. 2017; Ming et al. 2019; Papa et al. 2020).

We use LIGO O2 public data (Vallisneri et al. 2015; LIGO 2019; Abbott et al. 2021) and, thanks to a much longer coherent-search baseline, achieve a significantly higher sensitivity than the LIGO Collaboration O2 results in the same frequency range (Abbott et al. 2019; Palomba et al. 2019). Our results complement those of the high-frequency Falcon search (Dergachev & Papa 2020a, 2020b), which cover the range from 500–1700 Hz.

The plan of the paper is the following: we introduce the signal model and generalities about the search in Sections 2 and 3, respectively. In Sections 4 and 5 we detail the Einstein@-Home search and the follow-up searches. Constraints on the gravitational wave amplitude and on the ellipticity of neutron stars are obtained in Section 6, and conclusions are drawn in Section 7.

2. The Signal

The search described in this paper targets nearly monochromatic gravitational wave signals of the form described, for example, in Section II of Jaranowski et al. (1998). At the output

1

Steltner et al.

of a gravitational wave detector the signal has the form

$$h(t) = F_+ (\alpha, \delta, \psi; t) h_+ (t) + F_\times (\alpha, \delta, \psi; t) h_\times (t). \quad (1)$$

$F_+ (\alpha, \delta, \psi; t)$ and $F_\times (\alpha, \delta, \psi; t)$ are the detector beam-pattern functions for the "+" and "×" polarizations, (α, δ) are the R.A. and decl. of the source, ψ is the polarization angle, and t is the time at the detector. The waveforms $h_+ (t)$ and $h_\times (t)$ take the form

$$h_+(t) = A_+ \cos \Phi(t)$$
$$h_\times(t) = A_\times \sin \Phi(t), \quad (2)$$

with the "+" and "×" amplitudes

$$A_+ = \frac{1}{2} h_0 (1 + \cos^2 \iota)$$
$$A_\times = h_0 \cos \iota. \quad (3)$$

The angle between the total angular momentum of the star and the line of sight is $0 \leqslant \iota \leqslant \pi$ and $h_0 \geqslant 0$ is the intrinsic gravitational wave amplitude. $\Phi(t)$ of Equation (2) is the phase of the gravitational wave signal at time t. If τ_{SSB} is the arrival time of the wave with phase $\Phi(t)$ at the solar system barycenter, then $\Phi(t) = \Phi(\tau_{\mathrm{SSB}}(t))$. The gravitational wave phase as a function of τ_{SSB} is assumed to be

$$\Phi(\tau_{\mathrm{SSB}}) = \Phi_0 + 2\pi [f(\tau_{\mathrm{SSB}} - \tau_{0\mathrm{SSB}}) +$$
$$\frac{1}{2} \dot{f} (\tau_{\mathrm{SSB}} - \tau_{0\mathrm{SSB}})^2]. \quad (4)$$

We take $\tau_{0\mathrm{SSB}} = 1177858472.0$ (TDB in GPS seconds) as a reference time.

3. Generalities of the Searches
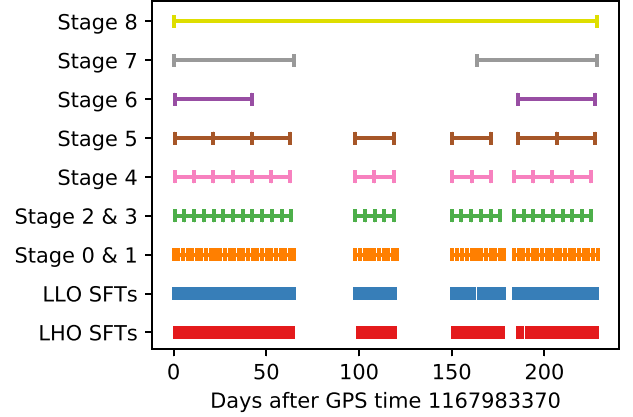
3.1. The Data

We use LIGO O2 public data from the Hanford (LHO) and the Livingston (LLO) detectors between GPS time 1167983370 (2017 January 9) and 1187731774 (2017 August 25). This data has been treated to remove spurious noise due to the LIGO laser beam jitter, calibration lines and the main power lines (Davis et al. 2019).

We additionally remove very loud short-duration glitches (B. Steltner et al. 2021, in preparation) and substitute Gaussian noise at frequency bins affected by line contamination (Covas et al. 2018). This is a procedure common to all Einstein@Home searches and it prevents spectral contamination from spreading to many nearby signal frequencies. The list of cleaned frequency bins can be found in the online tar.gz file.

As is customary, the input to our searches is in the form of short time-baseline (30 minutes) Fourier transforms (SFTs). These are grouped in segments of variable duration that correspond to the coherent time baselines of the various searches, as shown in Figure 1.

3.2. The Detection Statistics

For each search we partition the data in N_{seg} segments, with each segment spanning a duration T_{coh}. The data of both detectors from each segment i are combined coherently to construct a matched-filter detection statistic, the $\mathcal{F}$-statistic (Cutler & Schutz 2005). The $\mathcal{F}$-statistic depends on the template waveform that is being tested for consistency with the data. The $\mathcal{F}$-statistic at a given template point is the log-

Figure 1. Segmentation of the data used for the Einstein@Home search and the follow-up stages. The lower two bars show the input SFTs. The first gap in the data—starting at $\lesssim 70$ days—is due to spectral contamination in LHO, based on which we decided to exclude this period from the analysis. The second large gap—starting at ≈ 120 days—is due to an interruption of the science run for detector commissioning.

likelihood ratio of the data containing Gaussian noise plus a signal with the shape given by the template, to the data being purely Gaussian noise.

The $\mathcal{F}$-statistic values are summed, one per segment ($\mathcal{F}_i$), and, after dividing by N_{seg}, this yields our core detection statistic (Pletsch & Allen 2009; Pletsch 2010):

$$\overline{\mathcal{F}} := \frac{1}{N_{\mathrm{seg}}} \sum_{i=1}^{N_{\mathrm{seg}}} \mathcal{F}_i. \quad (5)$$

The $\overline{\mathcal{F}}$ is the average of $\mathcal{F}$ over segments, in general computed at different templates for every segment. The resulting $\overline{\mathcal{F}}$ is an approximation to the detection statistic at some template "in-between" the ones used to compute the single-segment $\mathcal{F}_i$. In fact these "in-between" templates constitute a finer grid based on which the summations of Equation (5) are performed.

The most significant Einstein@Home results are saved in the top list that the volunteer computer (host) returns to the Einstein@Home server. For these results the host also recomputes $\overline{\mathcal{F}}$ at the exact fine-grid template point. We indicate the recomputed statistic with a subscript "r," as, for example, in $\overline{\mathcal{F}_r}$.

In Gaussian noise $N_{\mathrm{seg}} \times 2\overline{\mathcal{F}}$ follows a chi-squared distribution with $4N_{\mathrm{seg}}$ degrees of freedom, $\chi^2_{4N_{\mathrm{seg}}}(\rho^2)$. The noncentrality parameter $\rho^2 \propto h_0^2 T_{\mathrm{data}}/S_h$, where T_{data} is the duration of time for which data is available and S_h is the strain power spectral density of the noise. The expected S/N squared is equal to ρ^2 (Jaranowski et al. 1998). For simplicity, in the rest of the paper, when we refer to the S/N we mean S/N squared.

If the noise contains some coherent instrumental or environmental signal, it is very likely that for some of the templates the distribution of $\overline{\mathcal{F}}$ will have a nonzero noncentrality parameter, even though there is no astrophysical signal. The reason is that in this case the data looks more like a noise+signal than like pure Gaussian noise.

It is possible to identify a nonastrophysical signal if it presents features that distinguish it from the astrophysical signals that the search is targeting, for example, if it is present only in one of the two detectors, or if it is present only for part

2

39

of the observation time. In the past we have used these signatures to construct ad hoc vetoes, such as the $\mathcal{F}$-stat consistency veto (Aasi et al. 2013a) and the permanence veto (Aasi et al. 2013b; Behnke et al. 2015). These vetoes are still widely used although with different names: the "single interferometer veto" in Sun et al. (2020) and Jones & Sun (2020) as well as the "persistency veto" of Abbott et al. (2019) and Astone et al. (2014).

We incorporated the ideas of the $\mathcal{F}$-stat consistency veto and of the permanence veto in the design of a new detection statistic, $\hat{\beta}_{\mathrm{S/GLtL}}$. The new detection statistic is an odds ratio that tests the signal hypothesis against a noise model, which in addition to Gaussian noise also includes single-detector continuous or transient spectral lines (Keitel et al. 2014; Keitel 2016). The subscript "L" in $\hat{\beta}_{\mathrm{S/GLtL}}$ stands for line, "G" for Gaussian, and "tL" for transient-line. We use this detection statistic to rank the Einstein@Home results. In this way we limit the number of results that make it in the top list but that would later be discarded by the vetoes. This frees up space on the top list for other, more interesting, results.

3.3. The Search Grids

For a rotating isolated neutron star, the template waveform is defined by the signal frequency, the spin-down, and the source sky-position. The range searched in each of these variables is gridded in such a way that the fractional loss in S/N, or mismatch, due to a signal falling in-between grid points is on average 0.5.

The grids in frequency and spin-down are each described by a single parameter, the grid spacing, which is constant over the search range. The sky grid is approximately uniform on the celestial sphere orthogonally projected on the ecliptic plane. The tiling is a hexagonal covering of the unit circle with hexagon edge length d:

$$d(m_{\mathrm{sky}}) = \frac{1}{f}\frac{\sqrt{m_{\mathrm{sky}}}}{\pi \tau_{\mathrm{E}}}, \qquad (6)$$

with $\tau_{\mathrm{E}} \simeq 0.021$ s being half of the light travel-time across the Earth and m_{sky} a constant that controls the resolution of the sky grid. The sky grids are constant over 5 Hz bands and the spacings are the ones associated through Equation (6) to the highest frequency in each 5 Hz. The resulting number of templates used to search 50 mHz bands as a function of frequency is shown in Figure 2. The grid spacings and m_{sky} are given in Table 1.

3.4. The Monte Carlos and the Assumed Signal Population

The loss in S/N $\mu(\vec{\lambda}_0)$ due to the parameters $\vec{\lambda}_0$ of a signal not perfectly matching the parameters $\vec{\lambda}_0 \pm \Delta\vec{\lambda}$ of the template can be described by a quadratic form, as long as the signal and the template parameters are fairly close, i.e., as long as $\Delta\vec{\lambda}$ is small:

$$\mu(\vec{\lambda}_0) = g_{ij}(\vec{\lambda}_0)\Delta\lambda^i\Delta\lambda^j. \qquad (7)$$

The metric g_{ij} for the search at hand can be estimated, at least numerically.

Setting up a search is a matter of deciding what loss in S/N one is willing to accept (fixing the mismatch), picking a tiling method and setting up a grid accordingly. Once the search grid
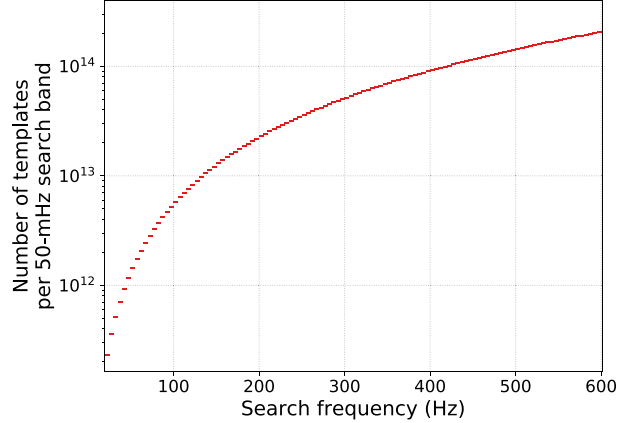
Figure 2. Number of searched templates per 50 mHz band as a function of frequency. The sky resolution increases with frequency causing the increase in the number of templates. The number of templates in frequency and spin-down is 14,970 and 8855 respectively.

is established, one can determine the computational cost of the search. If that is found to be too high, one must decide whether to reduce the T_{coh} or to increase the mismatch, and repeat the procedure. Ultimately the best operating point is a compromise between computational cost and sensitivity.

It turns out that for the first stages of all-sky surveys with T_{coh} of at least several hours, the optimal grids are typically ones with spacings $\gg$ the ones at which the metric approximation of Equation (7) holds. In particular we find that the metric mismatch overestimates the actual mismatch. This is good because it means that in order to achieve a certain maximum mismatch level, we need fewer templates than what the metric predicts. On the other hand it means that we cannot predict the mismatch analytically using Equation (7). Instead we must resort to simulating signals, searching for them with a given grid and measuring the loss in S/N with respect to a perfectly matched template. And we have to do this many times to probe different signals ($\vec{\lambda}_0$ values) and different random offsets between the template grid and the signal parameters.

This is the basic reason why in this paper we often refer to Monte Carlo studies. In all these studies the choice of signal parameters $\vec{\lambda}_0$ represents our target source population, which we assume to be uniformly distributed in spin frequency, log-uniformly distributed in spin-down, with orientation $\cos\iota$ uniformly distributed between -1 and 1, polarization angle ψ uniformly distributed in $|\psi| \leqslant \pi/4$ and source position uniformly distributed on the sky (uniform in $0 \leqslant \alpha < 2\pi$ and in $-1 \leqslant \sin\delta \leqslant 1$). The log-uniform distribution of spin-down values reflects our ignorance of the actual spin-down distribution of the sources over our large target range.

The Monte Carlo studies make the results robust and simple to interpret: all systematic effects in the analysis, both known and unknown, are automatically incorporated.

We note that since this analysis was carried out, a new metric ansatz was suggested (Allen 2019), which shows that the metric mismatch generically overestimates the actual mismatch, and shows how to extend the range of validity of the metric approximation. This might mitigate the need for such extensive Monte Carlo studies.

3

Steltner et al.

Table 1
Overview of the Searches

Search	T_{coh} (hr)	N_{seg}	δf (μHz)	$\delta\dot{f}$ (10^{-14} Hz s^{-1})	m_{sky}	$\langle\mu\rangle$	Δf (μHz)	$\Delta\dot{f}$ (10^{-14} Hz s^{-1})	$\dfrac{r_{sky}}{d(8.0\times10^{-3})}$	R^a	N_{in}	N_{out}
Stage 0	60	64	3.34	32.747 9	8.0×10^{-3}	0.5	full range	full range	all-sky	⋯	7.9×10^{17}	350 145
Stage 1	60	64	3.34	20	5.0×10^{-4}	0.3	850.0	1.2×10^{-10}	5.0	0.75	350 145	101 001
Stage 2	126	29	1	2	1.0×10^{-5}	0.09	130.0	2.0×10^{-11}	0.75	1.99	101 001	11 915
Stage 3	126	29	0.19	2	1.0×10^{-7}	0.002	10.0	2.0×10^{-12}	0.1	2.2	11 915	6 128
Stage 4	250	14	0.025	2	2.5×10^{-8}	0.001	0.4	3.2×10^{-13}	0.02	4.3	6 128	33
Stage 5	500	7	0.01	1	1.0×10^{-8}	0.001	0.17	1.45×10^{-13}	0.008	6.0	33	21
Stage 6	1000	2	0.001	0.1	1.0×10^{-9}	0.000 2	0.067	6.4×10^{-14}	0.003 7	10.0	21	18
Stage 7	1 563	2	0.001	0.1	5.0×10^{-10}	0.000 1	0.05	8.0×10^{-14}	0.005	15.0	18	8
Stage 8	$\approx$5 486	1	0.001	0.1	1.0×10^{-10}	0.000 7	0.032 5	4.25×10^{-14}	0.002 5	50.0	8	6

Note. We show the values of the following parameters: the number of segments N_{seg} and the coherent time baseline of each segment T_{coh}; the grid spacings δf, $\delta\dot{f}$, and m_{sky}; the average mismatch $\langle\mu\rangle$; the parameter space volume searched around each candidate, $\pm\Delta f$, $\pm\Delta\dot{f}$, and r_{sky} expressed in units of the side of the hexagon sky-grid tile of the Stage 0 search (Equation (6)); the threshold value R^a used to veto candidates of Stage a (Equation (8)); the number of templates searched (N_{in}) and how many of those survive and make it to the next stage (N_{out}). The first search, Stage 0, is the Einstein@Home search, hence the searched volume is the entire parameter space. The other searches are the follow-up stages.

4. The Einstein@Home Search

4.1. The Distribution of the Computational Load on Einstein@Home

This search leverages the computing power of the Einstein@Home project. This is built upon the BOINC (Berkeley Open Infrastructure for Network Computing) architecture (Anderson 2004; Anderson et al. 2006; BOINC 2020): a system that uses the idle time on volunteer computers to solve scientific problems that require large amounts of computing power.

The total number of templates that we searched with Einstein@Home is 7.9×10^{17}. The search is split into work units (WUs) sized to keep the average Einstein@Home volunteer computer busy for about 8 CPU hours. A total of 8 million WUs are necessary to cover the entire parameter space, representing of order 10,000 CPU years of computing.

Each WU searches 9.8×10^{10} templates, and covers 50 mHz, the entire spin-down range and a portion of the sky. Out of the detection statistic values computed for the 9.8×10^{10} templates, the WU-search returns to the Einstein@Home server only the information of the highest 7500 $\hat{\beta}_{S/GLtL}$ results.

This search ran on Einstein@Home between 2018 April and 2019 July, with an interruption of 8 months at the request of the LIGO/Virgo Collaboration, after the authors left the Collaboration.

### 4.2. Postprocessing of the Einstein@Home Search

We refer to a waveform template and the associated search results as a "candidate." All in all the Einstein@Home search returns 6.0×10^{10} candidates: the top 7500 candidates per WU $\times$ 8 million WUs. This is where the postprocessing begins.

The postprocessing consists of three steps:

1. *Banding:* as described in the previous section, each volunteer computer searches for signals with frequency within a given 50 mHz band, with spin-down between -2.6×10^{-9} and 2.6×10^{-10} Hz s^{-1} and a portion of the sky. The first step of the postprocessing is to gather together all results that pertain to the same 50 mHz band. We compute some basic statistics from these results and produce a series of diagnostic plots, that we can conveniently access through a GUI (graphical user interface) tool that we have developed for this purpose. This provides an overview of the result-set in any 50 mHz band.

2. *Identification of disturbed bands:* as done in previous Einstein@Home searches (Singh et al. 2016; Zhu et al. 2016; Abbott et al. 2017; Ming et al. 2019; Papa et al. 2020) we identify bands that present very significant deviations of the detection statistics from what we expect from a reasonably clean noise background. Such deviations can arise due to spectral disturbances or to extremely loud signals. We do not exclude these bands from further inspection, but we do flag them as this information is necessary when we set upper limits. We mark 273 50-mHz bands as disturbed.

3. *Clustering:* in this step we identify clusters of candidates that are close enough in parameter space that they are likely due to the same root cause. We associate with each cluster the template values of the candidate with the highest $\hat{\beta}_{S/GLtL\,r}$, which we also refer to as cluster seed. We use a new clustering method (B. Steltner et al. 2021, in preparation) that identifies regions in frequency-spin-down-sky-position that harbor an overdensity of candidates—a typical signal signature. This method achieves a lower false dismissal of signals at a fixed false alarm rate, with respect to the previous clustering (Singh et al. 2017) by tracing the S/N reduction function with no assumption on its profile in parameter space. An occupancy veto is also applied, requiring at least four candidates to be associated with a cluster. Most candidates have *fewer* than three nearby partners, so this clustering procedure greatly reduces the number of candidates, namely from 6.0×10^{10} to 350,145.

4. *Follow-up searches:* after the clustering we have 350,145 candidates, shown in Figure 3. Of these, 1352 come from bands that have been marked as disturbed. We follow all of them up as detailed in Section 5. The list of the disturbed 50 mHz bands is provided in the online tar.gz file.

4

Steltner et al.
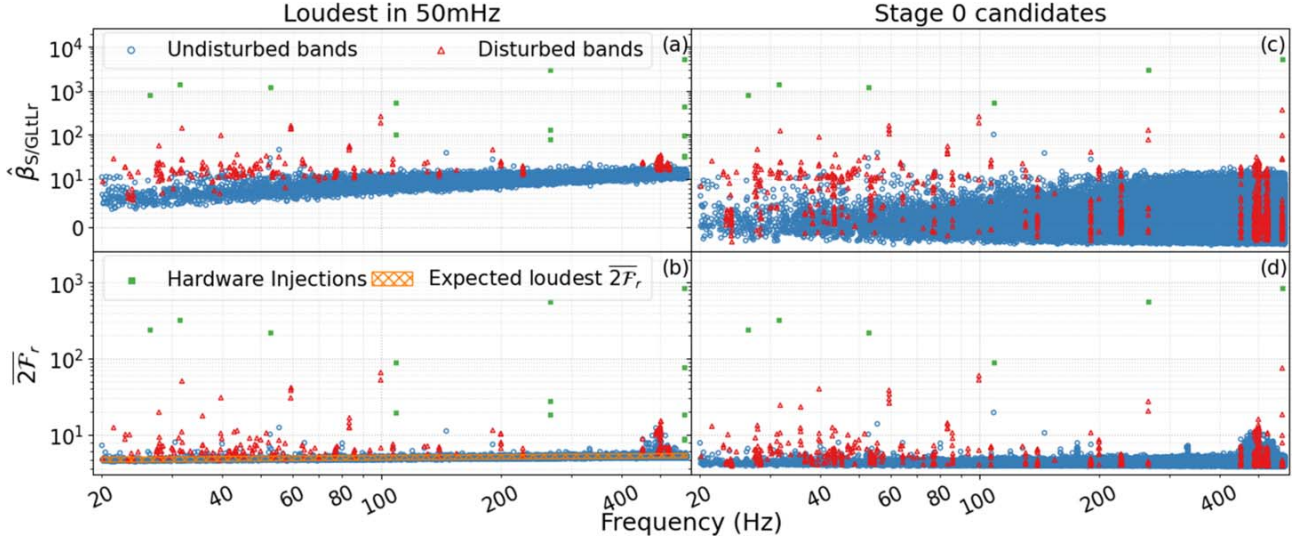
Figure 3. Detection statistics values of candidates as a function of frequency. The candidates coming from undisturbed bands are blue circles, those from disturbed bands are red triangles, and those from hardware injections are green squares. An unconventional vertical scale is used in all plots, which is linear below 10 and $\log_{10}$ elsewhere. Left panels: $\hat{\beta}_{\mathrm{S/GLtL}\,r}$ and $\overline{2\mathcal{F}_r}$ value of the loudest candidate (the candidate with the highest $\hat{\beta}_{\mathrm{S/GLtL}\,r}$) over 50 mHz, the entire sky and the full spin-down range, out of the Einstein@Home search. The increase in detection statistics with frequency is due to the number of searched templates increasing with frequency, as shown in Figure 2. The orange gridded area in the lower left panel indicates the 3σ expected range in Gaussian noise. Right panels: detection statistics values of the 350,145 candidates that are followed-up. By comparing the right and left panels one can see how we "dig" below the level of the loudest 50 mHz candidate with our follow-up stages.

In order to give a sense of the overall set of Einstein@Home results, in the left panels of Figure 3 we show the detection statistic value of the most significant result from every 50 mHz band. The large majority of the results falls within the expected range for noise-only. Most of the highest detection statistic values stem from hardware injections or from disturbed bands and are due to spectral contamination, i.e., signals (as opposed to noise fluctuations) of nonastrophysical origin.

5. The Follow-up Searches

Each stage takes as input the candidates that have survived the previous stage. Waveforms around the nominal candidate parameters are searched, so that if the candidate were due to a signal it would not be missed in the follow-up. The extent of the volume to search is based on the results of injection-and-recovery Monte Carlo studies and is broad enough to contain the true signal parameters for $\gtrsim 99.8\%$ of the signal population. For this reason we also refer to this volume as the "signal-containment region."[5] The containment region in the sky is a circle in the orthogonally projected ecliptic plane with radius r_{sky}.

The search setups for Stages 1-8 are chosen so that the S/N of a signal would increase from one stage to the next. This is achieved in two ways: by increasing the T_{coh} of the search and/ or by using a finer grid and hence by decreasing the average mismatch. The mismatch distributions of the various searches are shown in Figure 4. We note that even though average mismatch of Stage 8 is larger than that of the previous two stages, this does not imply that the expected S/N for a signal out of Stage 8 is smaller. In fact, because of the larger T_{coh} used
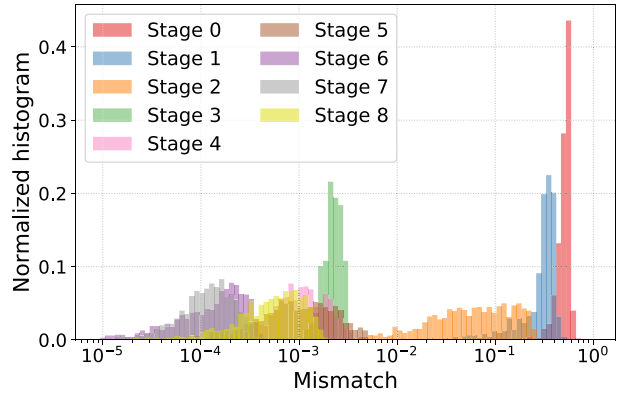

Figure 4. Mismatch distributions for the various follow-up searches based on 1000 injection-and-recovery Monte Carlos. The search setups are chosen so that the S/N of a signal increases from one stage to the next. This is achieved either by increasing the T_{coh} of the search and/or by decreasing the mismatch. We note that even though the average mismatch of Stage 8 is larger than that of the previous two stages, this does not imply that the expected S/N for a signal out of Stage 8 is smaller.

in Stage 8, the expected S/N for a signal out of Stage 8 is larger than that of the same signal out of Stage 7 or 6. This can be seen by comparing the values of R^8, R^7, and R^6, in Table 1 (the quantity R^a is defined below in Equation (8) and is related to the expected S/N increase at Stage a with respect to Stage 0).

We cluster the results of each search and consider the most significant cluster. We associate to the cluster the parameters of the member with the highest detection statistic value, and refer to this as the candidate from that follow-up stage.

We veto candidates at stage a whose S/N does not increase as expected for signals, with respect to Stage 0. We do this by

---

[5] The Monte Carlos were performed with 1839 signals, of which in Stage 1 the chosen containment region contained 1836. For the other stages all the signals were recovered within the chosen containment regions.
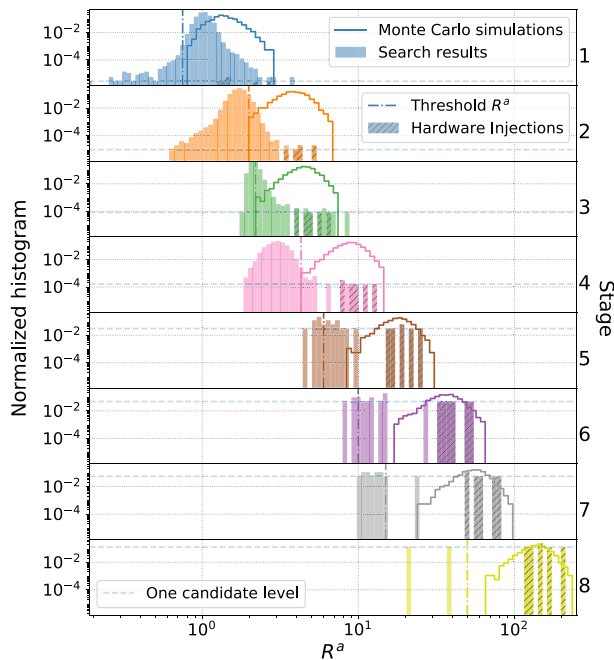
Figure 5. Distributions of R^a of candidates from signal injection-and-recovery Monte Carlos (solid lines) and from the actual search (shaded areas). The dashed-shaded areas show the R^a bins associated with the hardware injections. The dashed vertical lines mark the R^a threshold values. The dashed horizontal lines mark the one-candidate level in the search results.

setting a threshold on the quantity

$$R^a = \frac{2\overline{\mathcal{F}}_r^{\text{Stage a}} - 4}{2\overline{\mathcal{F}}_r^{\text{Stage 0}} - 4}. \tag{8}$$

The threshold is set based on signal injection-and-recovery Monte Carlos, as shown in Figure 5. The values are given in Table 1. Because of the large number of candidates in the first four follow-up stages, the R^a thresholds for $a = 1 \cdots 4$ are stricter than those used for the last four stages.

All the parameters relative to the searches, as well as the number of candidates surviving each stage, are shown in Table 1.

Only six candidates are left at the output of Stage 8. They are due to fake signals present in the data stream for validation purposes, the so-called hardware injections. In fact there are six hardware injections with parameters that fall in our search volume, those with ID 0, 2, 3, 5, 10, 11 (LIGO & Virgo 2018). We recover them all with consistent parameters.

6. Upper Limits

Based on our null result we set 90% confidence frequentist upper limits on the gravitational wave amplitude h_0 in half-Hz bands. The upper limit value is the smallest signal amplitude that would have produced a signal above the sensitivity level of our search for 90% of the signals of our target population (see Section 3.4). We establish the detectability of signals based on injection-and-recovery Monte Carlos. The upper limits are shown in Figure 6 and provided in machine-readable format in the online tar.gz file.

Our upper limits do not hold in some 50-mHz bands, namely those marked as disturbed and those associated with the
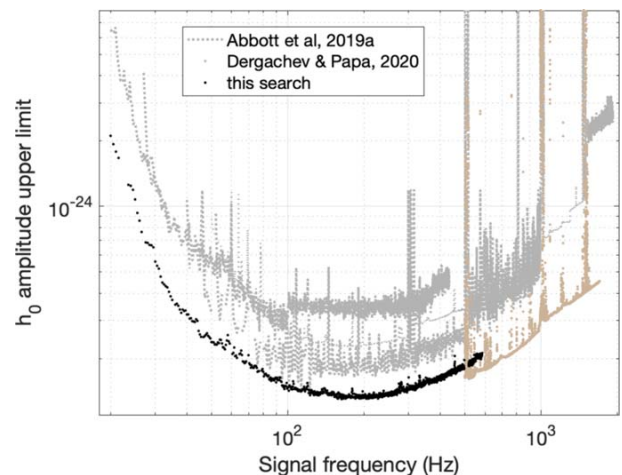

Figure 6. Smallest gravitational wave amplitude h_0 that we can exclude from the assumed population of signals (see Section 3.4). We compare our results with the latest literature: the Falcon search (Dergachev & Papa 2020a) and the LIGO results (Abbott et al. 2019) on the same data. There are multiple curves associated with the LIGO results because they used different analysis pipelines.

hardware injections. Even though we have followed-up candidates from these bands, we cannot exclude that a signal with strength below the disturbance but above the detection threshold—and hence above the upper limit—could be hidden by the loud disturbance, for example, by being associated with its large noise cluster. Another reason why we cannot guarantee that our upper limit holds in the presence of a disturbance is the saturation in the Einstein@Home top list that a loud disturbance produces. This prevents candidates from quieter parameter space regions in that band from being recorded. Given how loud the hardware injections are, for similar reasons, we also exclude the 50-mHz bands associated with these. The 50-mHz bands where the upper limits do not hold are provided in the online tar.gz file.

Upper limits are also not given in some half-Hz bands. This happens for two reasons: (1) if all 50-mHz bands in a half-Hz band are disturbed (2) due to the bin-cleaning procedure. In Section 3.1 we explained that we remove contaminated frequency bins and substitute them with Gaussian noise. If a signal were present in the cleaned-out bins, it too, would be removed. So in the half-Hz bands affected by cleaning, the upper limit Monte Carlos include the cleaning step *after* the signal has been added to the data. In this way the loss in detection efficiency due to the cleaning procedure is naturally folded into the upper limit. When a large fraction of the half-Hz bins is cleaned out, however, the detection efficiency may not reach the target 90% level. In this case we do not give an upper limit in the affected band. The list of half-Hz bands for which we do not give upper limits is given in the online tar.gz file.

Based on the upper limits, we compute the sensitivity depth $\mathcal{D}$ of the search Behnke et al. (2015) and find values between (49–56) $1/\sqrt{\text{Hz}}$. This is consistent with, and slightly better than, the previous performance of Einstein@Home searches (Dreissigacker et al. 2018). We provide the power spectral density estimate used to derive the sensitivity depth in the online tar.gz file.

We can express the h_0 upper limits as upper limits on the ellipticity ε of a source modeled as a triaxial ellipsoid spinning around a principal moment of inertia axis $\hat{I}$ at a distance D

Steltner et al.
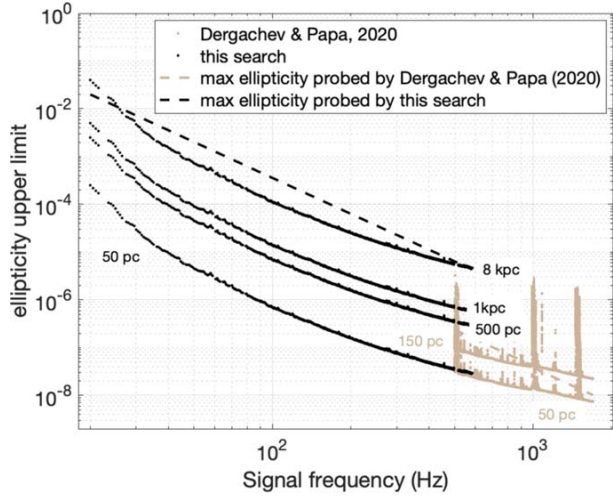
Figure 7. Upper limits on the ellipticity of a source at a certain distance (black). We also show the recent upper limits from the low ellipticity all-sky search of Dergachev & Papa (2020a). The dashed line is the spin-down ellipticity for the highest spin-down rate probed by each search.

(Jaranowski et al. 1998; Gao et al. 2020):

$$\varepsilon = 1.4 \times 10^{-6} \left(\frac{h_0}{1.4 \times 10^{-25}} \right)$$
$$\times \left(\frac{D}{1 \text{ kpc}} \right) \left(\frac{170 \text{ Hz}}{f} \right)^2 \left(\frac{10^{38} \text{ kg m}^2}{I} \right). \quad (9)$$

The ellipticity ε upper limits are plotted in Figure 7. If the spin-down of the signal were just due to the decreasing spin rate of the neutron star, then our search could not probe ellipticities higher than the spin-down limit ellipticity corresponding to the highest spin-down rate considered in the search, -2.6×10^{-9} Hz s^{-1}. This is indicated in Figure 7 by a dashed line.

Proper motion can reduce the apparent spin-down (Shklovskii 1970), so in principle we could detect a signal from a source with ellipticity above the dashed line. However, even in extreme cases (source distance 8 kpc, spin period 1 ms, large proper motion 100 mas yr^{-1} (Hobbs et al. 2005) or source distance 10 pc, spin period 1 ms and tangential velocity of 1000 km s^{-1}) the change in maximum detectable ellipticity is negligible.

7. Conclusions

We present the results from an Einstein@Home search for continuous, nearly monochromatic, gravitational waves with frequency between 20.0 and 585.15 Hz, and spin-down between -2.6×10^{-9} and 2.6×10^{-10} Hz s^{-1}. We use LIGO O2 public data and compare it against 7.9×10^{17} waveforms. We follow-up the most likely 350,145 candidates through a hierarchy of eight searches, each being more sensitive but requiring more per-template computing power than the previous one. No candidate survives all the stages.

This is the most sensitive search performed on this parameter space on O2 data, and sets the most stringent upper limits on the intrinsic gravitational wave amplitude h_0. The most constraining h_0 upper limit is 1.3×10^{-25} at 163.0 Hz, corresponding to a neutron star at, say, 100 pc, having an ellipticity of $\lesssim 5 \times 10^{-7}$ and rotating with a spin period of ≈ 12 ms. Our results thus exclude neutron stars rotating faster

than 12 ms, within 100 pc of Earth, with ellipticities in the few $\times 10^{-7}$ range and reach the 1×10^{-7} mark for spins of 5 ms.

These results probe a plausible range of pulsar ellipticity values, well within the boundaries of what the crust of a standard neutron star could support, around 10^{-5}, according to some models (Johnson-McDaniel & Owen 2013). It is hard to produce a definitive estimate of such a quantity and it may be that this maximum value is significantly lower (Gittins et al. 2020). Since the closest neutron star is expected to be at about a distance of 10 pc (Dergachev & Papa 2020a), it is likely that there are several hundreds within 100 pc. On the other hand, recent analyses of the population of known pulsars suggest that their ellipticity should lie in the 10^{-9} decade (Woan et al. 2018; Bhattacharyya 2020), which we reach only for sources rotating faster than 5 ms and within 10 pc. When the O3 LIGO data is released, its sensitivity improvement with respect to the O2 data used here (Buikema et al. 2020) will allow us to extend the reach of our search and probe ellipticities in the 10^{-9} decade, at these higher frequencies.

ORCID iDs

B. Steltner ⓘ https://orcid.org/0000-0003-1833-5493
M. A. Papa ⓘ https://orcid.org/0000-0002-1007-5298
H.-B. Eggenstein ⓘ https://orcid.org/0000-0001-5296-7035
B. Allen ⓘ https://orcid.org/0000-0003-4285-6256

References

Aasi, J., Abadie, J., Abbott, B. P., et al. 2013a, PhRvD, 87, 042001
Aasi, J., Abadie, J., Abbott, B. P., et al. 2013b, PhRvD, 88, 102002
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, PhRvD, 96, 122004
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, PhRvD, 100, 024004
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, SoftX, 13, 100658
Allen, B. 2019, PhRvD, 100, 124004
Anderson, D. P. 2004, in Proc. 5th IEEE/ACM Int. Workshop on Grid Computing (Washington, DC: IEEE Computer Society), 4
Anderson, D. P., Christensen, C., & Allen, B. 2006, in Proc. 2006 ACM/IEEE SC06 Conf. (New York: Association for Computing Machinery), 126
Arvanitaki, A., Baryakhtar, M., & Huang, X. 2015, PhRvD, 91, 084011
Astone, P., Colla, A., D'Antonio, S., Frasca, S., & Palomba, C. 2014, PhRvD, 90, 042002
Behnke, B., Papa, M. A., & Prix, R. 2015, PhRvD, 91, 064007
Bhattacharyya, S. 2020, MNRAS, 498, 728
BOINC 2020, http://boinc.berkeley.edu/
Buikema, A., Cahillane, C., Mansell, G. L., et al. 2020, PhRvD, 102, 062003
Covas, P. B., Effler, A., Goetz, E., et al. 2018, PhRvD, 97, 082002
Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006
Davis, D., Massinger, T., Lundgren, A., et al. 2019, CQGra, 36, 055011
Dergachev, V., & Papa, M. A. 2020a, PhRvL, 125, 171101
Dergachev, V., & Papa, M. A. 2020b, arXiv:2012.04232
Dreissigacker, C., Prix, R., & Wette, K. 2018, PhRvD, 98, 084058

Gao, Y., Shao, L., Xu, R., et al. 2020, MNRAS, 498, 1826
Gittins, F., Andersson, N., & Jones, D. 2020, arXiv:2009.12794
Hobbs, G., Lorimer, D., Lyne, A., & Kramer, M. 2005, MNRAS, 360, 974
Horowitz, C., Papa, M., & Reddy, S. 2020, PhLB, 800, 135072
Horowitz, C. J., & Reddy, S. 2019, PhRvL, 122, 071102
Jaranowski, P., Krolak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001
Johnson-McDaniel, N. K., & Owen, B. J. 2013, PhRvD, 88, 044004
Jones, D., & Sun, L. 2020, arXiv:2007.08732
Keitel, D. 2016, PhRvD, 93, 084024
Keitel, D., Prix, R., Papa, M. A., Leaci, P., & Siddiqi, M. 2014, PhRvD, 89, 064023
Lasky, P. D. 2015, PASA, 32, e034
LIGO 2019, The O2 Data Release, https://www.gw-openscience.org/O2/
LIGO Virgo 2018, https://www.gw-openscience.org/O2_injection_params
Ming, J., Papa, M. A., Singh, A., et al. 2019, PhRvD, 100, 024063
Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, PhRvD, 58, 084020
Palomba, C., D'Antonio, S., Astone, P., et al. 2019, PhRvL, 123, 171101
Papa, M. A., Ming, J., Gotthelf, E. V., et al. 2020, ApJ, 897, 22
Pletsch, H. J. 2010, PhRvD, 82, 042002
Pletsch, H. J., & Allen, B. 2009, PhRvL, 103, 181102
Shklovskii, I. S. 1970, SvA, 13, 562
Singh, A., Papa, M. A., Eggenstein, H.-B., et al. 2016, PhRvD, 94, 064061
Singh, A., Papa, M. A., Eggenstein, H.-B., & Walsh, S. 2017, PhRvD, 96, 082003
Sun, L., Brito, R., & Isi, M. 2020, PhRvD, 101, 063020
Vallisneri, M., Kanner, J., Williams, R., Weinstein, A., & Stephens, B. 2015, JPhCS, 610, 012021
Woan, G., Pitkin, M., Haskell, B., Jones, D., & Lasky, P. 2018, ApJL, 863, L40
Zhu, S. J., Baryakhtar, M., Papa, M. A., et al. 2020, PhRvD, 102, 063020
Zhu, S. J., Papa, M. A., Eggenstein, H.-B., et al. 2016, PhRvD, 94, 082008

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

The detection of pulsed gamma-ray emission from the pulsar PSR J0952-0607 in seven years of *Fermi*-Large Area Telescope (LAT) data is presented.

The author prepared the O1 and O2 data of Advanced LIGO for the continuous gravitational wave search with the method described in Chapter 2.

Detection and Timing of Gamma-Ray Pulsations from the 707 Hz Pulsar J0952−0607

L. Nieder[1,2] , C. J. Clark[3] , C. G. Bassa[4] , J. Wu[5] , A. Singh[1,2,6], J. Y. Donner[5,7], B. Allen[1,2,8] , R. P. Breton[3] ,
V. S. Dhillon[9,10] , H.-B. Eggenstein[1,2] , J. W. T. Hessels[4,11] , M. R. Kennedy[3] , M. Kerr[12] , S. Littlefair[9] ,
T. R. Marsh[13] , D. Mata Sánchez[3] , M. A. Papa[1,2,8] , P. S. Ray[12] , B. Steltner[1,2] , and J. P. W. Verbiest[5,7]

[1] Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), 30167 Hannover, Germany; lars.nieder@aei.mpg.de
[2] Leibniz Universität Hannover, 30167 Hannover, Germany
[3] Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, M13 9PL, UK
[4] ASTRON, The Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands
[5] Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
[6] The Geophysical Institute, Bjerknes Centre for Climate Research, University of Bergen, Bergen NO-5007, Norway
[7] Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany
[8] Department of Physics, University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA
[9] Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, UK
[10] Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain
[11] Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
[12] Space Science Division, Naval Research Laboratory, Washington, DC 20375-5352, USA
[13] Astronomy and Astrophysics Group, Department of Physics, University of Warwick, Coventry CV4 7AL, UK
Received 2019 May 24; revised 2019 July 11; accepted 2019 July 23; published 2019 September 18

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 $F_\gamma = 2.6 \times 10^{-12}$ erg cm^{-2} 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 7 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 2 years and show results from optical modeling of the black-widow-type companion based on new multiband photometric data taken with HiPERCAM on the Gran Telescopio Canarias on La Palma and ULTRACAM on the New Technology Telescope at ESO La Silla (based on observations collected at the European Southern Observatory, Chile; programme 0101.D-0925, PI: Clark, C. J.). PSR J0952−0607 is now the fastest-spinning pulsar for which the intrinsic spin-down rate has been reliably constrained ($\dot{P}_{\rm int} \lesssim 4.6 \times 10^{-21}$ s s^{-1}). The inferred surface magnetic field strength of $B_{\rm surf} \lesssim 8.2 \times 10^{7}$ G is among the 10 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.

Unified Astronomy Thesaurus concepts: Gamma-ray sources (633); Millisecond pulsars (1062)

1. Introduction

The Large Area Telescope (LAT) on board the *Fermi Gamma-ray Space Telescope* (Atwood et al. 2009) 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., Caraveo 2014). 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., Ray et al. 2012). 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., Acero et al. 2015). (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 (Nolan et al. 2012). For most gamma-ray pulsars, curved spectra are preferred with >95% confidence (e.g., Abdo et al. 2013). (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; Acero et al. 2015), 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., Jaodand et al. 2018) are an important exception, with significant changes in gamma-ray flux associated with transitions between accretion- and rotation-powered states (Stappers et al. 2014; Johnson et al. 2015).

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

1

Nieder et al.

(GBT; Stovall et al. 2014) and the Arecibo telescope (Cromartie et al. 2016). The Giant Metrewave Radio Telescope searches around 607 MHz (Bhattacharyya et al. 2013). Another survey around 820 MHz is run by the GBT (Ransom et al. 2011). Finally Parkes (Camilo et al. 2015), Nançay (Cognard et al. 2011) and Effelsberg (Barr et al. 2013) 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 $\tau_s \propto \nu^{-4.4}$; Levin et al. 2016) 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; Frail et al. 2016a). 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., Verbiest et al. 2016; Tiburzi 2018).

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. (2018) 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 (Bates et al. 2013). Additionally, some emission models suggest that pulsars' radio beams are wider at low frequencies (e.g., Story et al. 2007), 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., Abdo et al. 2009; Wu et al. 2018). Indeed, one emission model for the recently discovered radio-quiet MSP PSR J1744−7619 (Clark et al. 2018) suggests that radio pulsations may only be detectable at low radio frequencies.

Pleunis et al. (2017) performed very-low-frequency pulsar searches at 115–155 MHz with the Low-Frequency Array (LOFAR; Stappers et al. 2011; van Haarlem et al. 2013). This was possible due to new semi-coherent de-dispersion techniques that mitigate the smearing due to dispersion (Bassa et al. 2017a). The searches targeted unassociated sources from the 3FGL catalog (Acero et al. 2015). An isolated MSP, PSR J1552+5437, was detected first in radio and subsequently in gamma-rays (Pleunis et al. 2017).

Bassa et al. (2017b) 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 7 years of "Pass 8" LAT data (see Atwood et al. 2013).

In this survey they discovered PSR J0952−0607, a binary radio MSP with a spin frequency of 707 Hz (Bassa et al. 2017b). 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 (Hessels et al. 2006). In contrast to pulsars in globular clusters, which experience significant but unknown acceleration due to the gravitational potential within the cluster (Prager et al. 2017), 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 (Chen et al. 2013). 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 (Mukherjee 2017).

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 1 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. (2017b) 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 ($M_c \sim 0.15$–$0.4\,M_\odot$) (Roberts 2013). Optical light curves of these systems reveal that the pulsar emission heats the nearly Roche-lobe filling companion (Breton et al. 2013). Observations of orbitally modulated X-ray emission shows that interactions between the pulsar and companion star winds produce intra-binary shocks (e.g., Roberts et al. 2014).

For many spider pulsars the radio pulsations are completely absorbed by intra-binary material during parts of their orbit (e.g., Fruchter et al. 1988), 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., Stappers et al. 1996; Archibald et al. 2009; Polzin et al. 2018), 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 (Ray et al. 2011; Kerr et al. 2015; Clark et al. 2017). 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 (Pletsch & Clark 2015).

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 7 years to 2011. This allows us to

2

Nieder et al.

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 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 3. New radio and optical observations as well as a search for continuous gravitational waves are discussed in Section 4. Finally, in Section 5 we discuss the implications of the results presented and we conclude in Section 6.

2. Gamma-Ray Pulsation Discovery

2.1. Data Preparation

The gamma-ray source targeted by Bassa et al. (2017b) resulting in the detection of the radio pulsar PSR J0952−0607 and its optical counterpart (R.A. $\alpha_{J2000.0} = 09^h52^m08\overset{s}{.}319$, decl. $\delta_{J2000.0} = -06°07'23\overset{''}{.}49$) was discovered using 7 years of LAT data, but was too faint to be included in the 3FGL catalog (i.e., in 4 years of data; Acero et al. 2015). It is included in the successive 4FGL catalog based on 8 years of data as 4FGL J0952.1−0607 (The *Fermi*-LAT collaboration 2019).

To search for gamma-ray pulsations from PSR J0952−0607, we used "Pass 8" (Atwood et al. 2013) 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 catalog.[14] The photons were selected using gtselect from the *Fermi* Science Tools[15] if 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 (Bickel et al. 2008; Kerr 2011). 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 (Nolan et al. 2012) fixed to the position of the pulsar's optical counterpart reported by Bassa et al. (2017b). We used the templates gll_iem_v06.fits for the Galactic diffuse emission (Acero et al. 2016) and iso_P8R2_SOURCE_V6_v06.txt[16] for the isotropic diffuse background. The spectral analysis included all 3FGL sources within 15° of the pulsar position and the spectral

Table 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} cm^{-2} s^{-1})	$2.25 \pm 0.77 \pm 0.34$
Energy flux F_γ (10^{-12} erg cm^{-2} s^{-1})	$2.60 \pm 0.38 \pm 0.16$

Note. 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 $\pm6\%$, as described in Abdo et al. (2013).

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 signal-to-noise ratio (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 data set 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 conservative.[17] We extended the data set 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 List[18] (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 (Nolan et al. 2012; Acero et al. 2015). Six uncatalogued sources with TS > 10 ($\sim3\sigma$) 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 1. Here, we also give TS_{cut} which is computed like TS but comparing an exponentially cutoff power-law model and a power-law model without cutoff (Abdo et al. 2013).

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

---

[14] https://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone_LAT_IRFs/IRF_PSF.html

[15] https://fermi.gsfc.nasa.gov/ssc/data/analysis/software

[16] https://fermi.gsfc.nasa.gov/ssc/data/access/lat/BackgroundModels.html

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

[18] https://fermi.gsfc.nasa.gov/ssc/data/access/lat/fl8y/

3

Nieder et al.

the S/N remains. This leaves $N = 4642$ actual or $\sum w_j = 331.4$ effective photons.

2.2. Search

For many pulsars, LAT data covering several years of observation time are needed for significant pulsation detection (e.g., Hou et al. 2014). 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 $\Phi(t, \lambda)$ is used that depends on time t and (for circular-binary pulsars) on a set of at least seven parameters $\lambda = (f, \dot{f}, \alpha, \delta, P_{orb}, x, t_{asc})$. 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 semimajor axis x, and epoch of ascending node t_{asc}), and (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 $T_{obs} < 1$ 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 $T_{obs} \gg P_{orb}$.

Searches for binary gamma-ray pulsars are therefore computationally expensive, as a multidimensional parameter space must be searched with a dense grid (Pletsch et al. 2012). 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 Bassa et al. (2017b) 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 2.5 m 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 (de Jager et al. 1989). It combines the Fourier power from several harmonics incoherently by maximizing over the first M harmonics via

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

with $M_{max} = 20$ as suggested by de Jager et al. (1989). The Fourier power in the nth harmonic is given by

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

with the normalization constant

$$\kappa^2 = \frac{1}{2} \sum_{j=1}^{N} w_j^2. \quad (3)$$

The construction of a grid for this search was done using a distance "metric" on the parameter space (Balasubramanian et al. 1996; Owen 1996). 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 (Pletsch et al. 2012).

The metric components for the parameters of an isolated pulsar are given in Pletsch & Clark (2014), and the additional components required to search for a binary pulsar will be described in an upcoming paper (L. Nieder et al. 2019, in preparation). 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 semimajor 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 nth 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 (Pletsch & Clark 2014). 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 $M_{max} = 5$ 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 $M_{max} = 20$. 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; Nieder et al. 2019, in preparation), and the dimensionality is low. For higher dimensional parameter spaces hypercubic grids become extremely wasteful. The projected semimajor 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 (Aulbert & Fehrmann 2009), and the search took only 2 days.

4

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 Bassa et al. (2017b). 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, multidimensional H statistic search (see Appendix). We therefore applied a "bootstrapping" procedure (described in the Appendix) to estimate the detection significance from the search results themselves, finding a trials-corrected false-alarm probability of $P_{FA} \approx 3.3 \times 10^{-3}$. After extending our data set to cover the extra year of data as explained in Section 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 (Kerr 2011) for this increase ($P_{FA} = \exp(-0.3984 \,\Delta H_m) = 1.6 \times 10^{-3}$), giving an overall false-alarm probability of $P_{FA} \approx 5.3 \times 10^{-6}$ in the extended data set, confirming the detection.

3. Gamma-Ray Timing

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 Clark et al. (2015, 2017), extending the methods developed by Ray et al. (2011) and Kerr et al. (2015). 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 (Abdo et al. 2013)

$$\hat{g}(\Phi) = \left(1 - \sum_{i=1}^{N_p} a_i\right) + \sum_{i=1}^{N_p} a_i \, g(\Phi, \mu_i, \sigma_i). \quad (4)$$

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

$$g(\Phi, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{(\Phi + k - \mu)^2}{2\sigma^2}\right). \quad (5)$$

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

$$\mathcal{L}(\hat{g}, \boldsymbol{\lambda}) = \prod_{j=1}^{N} [w_j \hat{g}(\Phi(t_j, \boldsymbol{\lambda})) + (1 - w_j)]. \quad (6)$$

The Bayesian information criterion (BIC; Schwarz 1978) is used to choose the number of peaks by minimizing

$$\text{BIC} = -2\log(\mathcal{L}(\hat{g}, \boldsymbol{\lambda})) + k \log\left(\sum_{j=1}^{N} w_j\right), \quad (7)$$

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. (2017), this template pulse profile is used to explore the multidimensional 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 & Weare (2010) to run many Monte Carlo chains in parallel for the exploration and the efficient parallelization scheme described by Foreman-Mackey et al. (2013). 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., Clark et al. 2017) 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 An et al. 2017).

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

5

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 $\mathrm{FWHM}_i = 2\sqrt{2\log(2)}\,\sigma_i$ in the range $0.05 < \mathrm{FWHM}_i < 0.5$). 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 4.1) reported in Table 2. In Figure 1 we show 100 pulse profile templates randomly picked from the resulting template parameter distribution.

3.2. Solution

Our timing solution is shown in Table 2. 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 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 parameters leading to the highest likelihood are given in Table 2.

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 (Bassa et al. 2017b). 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, $\dot{f}_{\mathrm{obs}} = -2.382(8) \times 10^{-15}$ 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 4.1 and Table 2).

It is not possible for us to confidently determine the proper motion as we find hints for and against nonzero proper motion. Allowing proper motion to vary jointly with the template parameters results in a significantly improved H statistic, log-likelihood, and BIC. The timing analysis sets the 95% confidence region on proper motion to $\mu_\alpha \cos\delta \in [-27.4,\ -1.9]$ mas yr^{-1} and $\mu_\delta \in [-23.0,\ 19.1]$ mas yr^{-1}. The most likely total proper motion $\mu_{\mathrm{t}} = \sqrt{\mu_\alpha^2 \cos^2\delta + \mu_\delta^2}$ is 14.8 mas yr^{-1} with a 95% upper limit of 25.3 mas yr^{-1}. Typically, however, it is assumed that the H statistic rises linearly with exposure time and nonzero proper motion resulting from this timing analysis leads to a bumpier rise in the H statistic over time. This indicates that the

Table 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$^{\mathrm{h}}$52$^{\mathrm{m}}$08$^{\mathrm{s}}$.322(2)	09$^{\mathrm{h}}$52$^{\mathrm{m}}$08$^{\mathrm{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, $\dot{f}_{\mathrm{obs}}$ (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 semimajor 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] $\dot{P}_{\mathrm{int}}$ (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} × $(d/1\,\mathrm{kpc})^2$	

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 5.
[d] Properties are derived as described in Abdo et al. (2013) on the basis of the estimated intrinsic spin-period derivative $\dot{P}_{\mathrm{int}}$.
[e] Assuming no beaming and distance $d = 1$ kpc.

proper motion resulting from our analysis might not be correct. Keeping the template fixed to the template parameters found by folding the gamma-ray data with the radio ephemeris results in a 95% confidence region on proper motion consistent with zero. Zero proper motion is also favored by the BIC. The same is found when using a single-peaked profile in the timing analysis and varying the template parameters jointly.

Nieder et al.
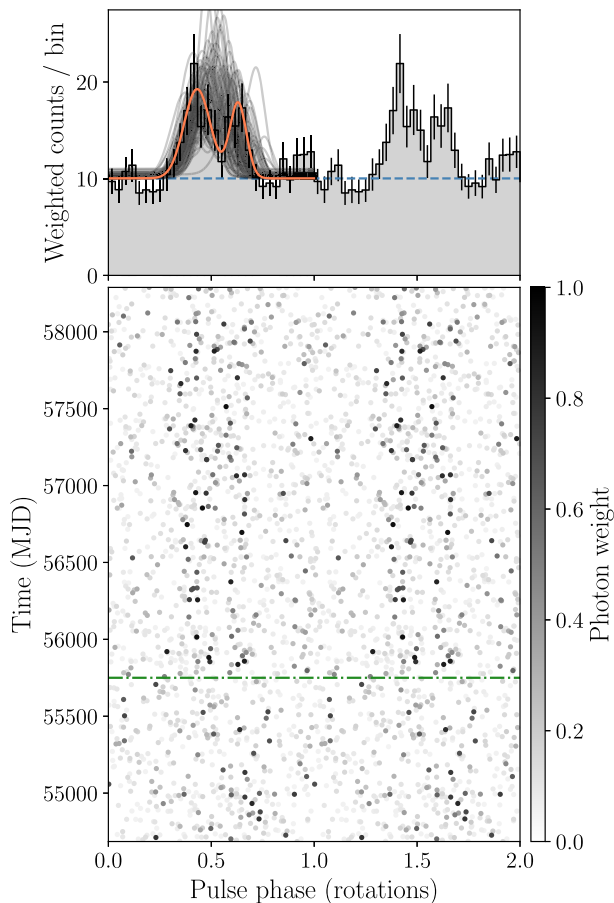
Figure 1. Integrated pulse profile after MJD 55,750 and phase-time diagram of PSR J0952−0607, showing two identical rotations for clarity. Top: the orange curve indicates the template with the highest BIC. The transparent black curves illustrate 100 representative templates randomly selected from the 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.

The upper limit on proper motion corresponds to a transverse velocity of $v_t = \mu_t d = 120\,\mathrm{km\,s^{-1}} \times (d/1\,\mathrm{kpc})$. This results in high, but not unrealistic transverse velocities when using the distances inferred from the dispersion measure ($d = 0.97$ kpc (Cordes & Lazio 2002, hereafter NE2001) or $d = 1.74$ kpc (Yao et al. 2017, hereafter YMW16)). As 90% of the known MSPs in the Australia Telescope National Facility (ATNF) Pulsar Catalogue[19] (Manchester et al. 2005) show transverse velocities below $200\,\mathrm{km\,s^{-1}}$ the proper motion upper limit is unrealistic for the higher distances predicted by the optical observations (4.7–6.6 kpc; see Section 5).

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

---

[19] http://www.atnf.csiro.au/research/pulsar/psrcat

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 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 $\sim$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 55,750, 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 Abdo et al. (2010). 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 (Arzoumanian et al. 1994). Nevertheless the penalty for adding orbital-period derivatives led to an increase in the BIC. Similarly, no significant semimajor-axis derivative was found.

4. Multiwavelength

4.1. Updated Radio Timing

Observations of PSR J0952−0607 with LOFAR have been ongoing using an identical observational setup as in Bassa et al. (2017b), namely a single tied-array beam formed from the high-band antennas (HBAs) of the central 23 LOFAR (van Haarlem et al. 2013) core stations, using 78 MHz of bandwidth at a central frequency of 149 MHz. Before 2018 May, several 5 minute integrations were obtained at each observing epoch; after that the integration times were increased to 20 minutes. These observations were obtained at a roughly monthly cadence. As described in Bassa et al. (2017b), these observations were coherently de-dispersed, folded with DSPSR (van Straten & Bailes 2011), and analyzed using tools in the PSRCHIVE software suite (Hotan et al. 2004) and the TEMPO2 pulsar-timing software (Edwards et al. 2006; Hobbs et al. 2006).

7

Nieder et al.

The phase-connected timing solution from Bassa et al. (2017b) 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 (van Straten et al. 2012) 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 2). 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 5×10^{-5} pc cm^{-3} 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., Coles et al. 2011), particularly in nonperiodic 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 $0.15 < \phi_{orb} < 0.35$. Using the TOAs from this orbital phase range we set a 3σ upper limit on time delays due to additional dispersion of $\Delta t < 2.3$ μs, and hence ΔDM $< 1.2 \times 10^{-5}$ pc cm^{-3}.

4.2. Optical Photometry

Bassa et al. (2017b) presented an r'-band light curve of the optical companion to PSR J0952−0607 taken by the WFC on the 2.5 m 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. (2017b) modeled this light curve with the Icarus package (Breton et al. 2012), finding that PSR J0952−0607 is likely to have an inclination angle $i \sim 40°$, but the lack of color information precluded a 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 multicolor photometry using ULTRACAM (Dhillon et al. 2007) on the 3.58 m New Technology Telescope (NTT) at ESO La Silla, and HiPER-CAM (Dhillon et al. 2016, 2018) on the 10.4 m Gran Telescopio Canarias (GTC) on La Palma. The observation specifics are given in Table 3.

These data were calibrated and reduced using the ULTRA-CAM[20] and HiPERCAM[21] 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" (Honeycutt 1992) to correct for airmass effects and varying transparency. Magnitudes in g_s, r_s, i_s, and z_s[22] were calibrated using comparison stars chosen from the Pan-STARRS1 (Chambers et al. 2016) 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 (Smith et al. 2002) 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 $S_{AB} = 3631(C/C_0)$ Jy.

4.3. Optical Light-curve Modeling

As in Bassa et al. (2017b), 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 (Feroz et al. 2013) via the Python package PyMultiNest (Buchner et al. 2014). 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. (2013).

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 $T_d^4 = T_{irr}^4 + T_n^4$, 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 Breton et al. 2013); the binary inclination angle (i); the Roche-lobe filling factor (f_{RL}, defined as the ratio between the companion's radius toward 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 ($q = M_{psr}/M_c$) were derived from the binary mass function according to the timing measurements of P_{orb} and x presented in Table 2. We also marginalize over interstellar extinction and reddening, parameterized by the $E(B − V)$ of Green et al. (2018), 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. (2018) 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 $1.2 < M_{psr} < 2.5\,M_\odot$, and $0.1 < f_{RL} < 1$. Temperatures T_n and T_d were

---

[20] http://deneb.astro.warwick.ac.uk/phsaap/software/ultracam/html/
[21] http://deneb.astro.warwick.ac.uk/phsaap/hipercam/docs/html/

[22] ULTRACAM 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 (Dhillon et al. 2018).
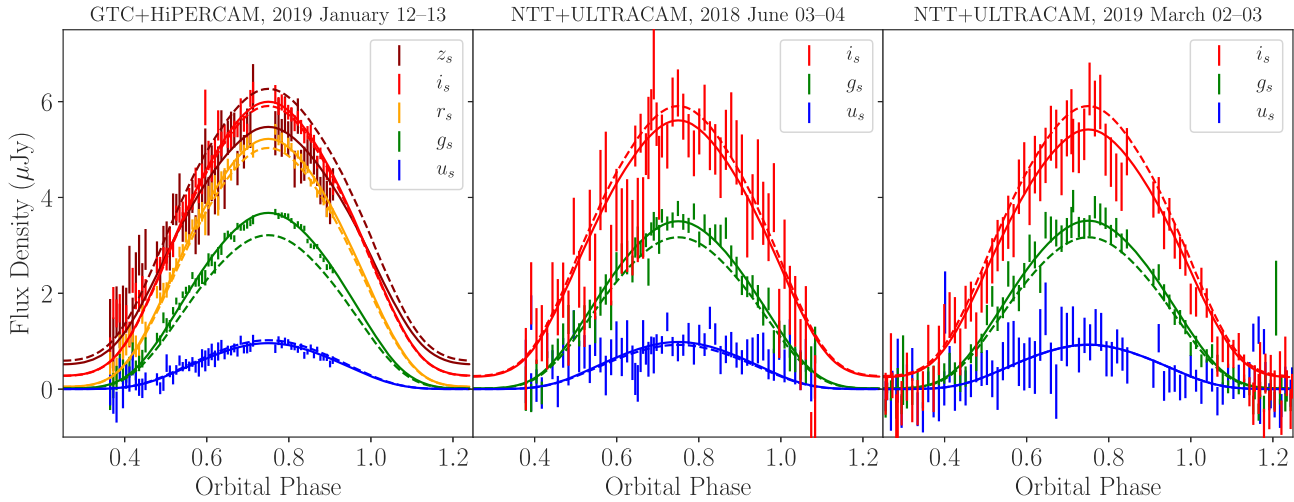
Figure 2. Optical light curve of the companion to PSR J0952−0607, phased using the gamma-ray timing ephemeris. For clarity, the HiPERCAM and ULTRACAM fluxes have been combined into 180 and 300 s time bins, respectively, via weighted average. The unbinned data were used for the light-curve model fitting. 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.

Table 3
New Optical Photometry of the Companion of PSR J0952−0607

Night Beginning	Instrument+Telescope	Filters	$\phi_{\rm orb}$	Airmass	Seeing	Photometric
2018 Jun 3	ULTRACAM+NTT	u_s, g_s, i_s	0.64–1.09	1.1–2.1	1″0–2″0	yes
2018 Jun 4	ULTRACAM+NTT	u_s, g_s, i_s	0.37–0.71	1.1–1.6	1″0–3″0	no
2019 Jan 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 Jan 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 Mar 2[b]	ULTRACAM+NTT	u_s, g_s, i_s	0.91–1.29	1.1–1.6	0″8–1″2	no
2019 Mar 3	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 $\phi_{\rm orb} = 0$ 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 $\phi_{\rm orb} = 0.6$ 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.

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., Sanchez & Romani 2017).

The best-fitting light-curve model is shown in Figure 2, with posterior distributions for the fit parameters shown in Figure 3.

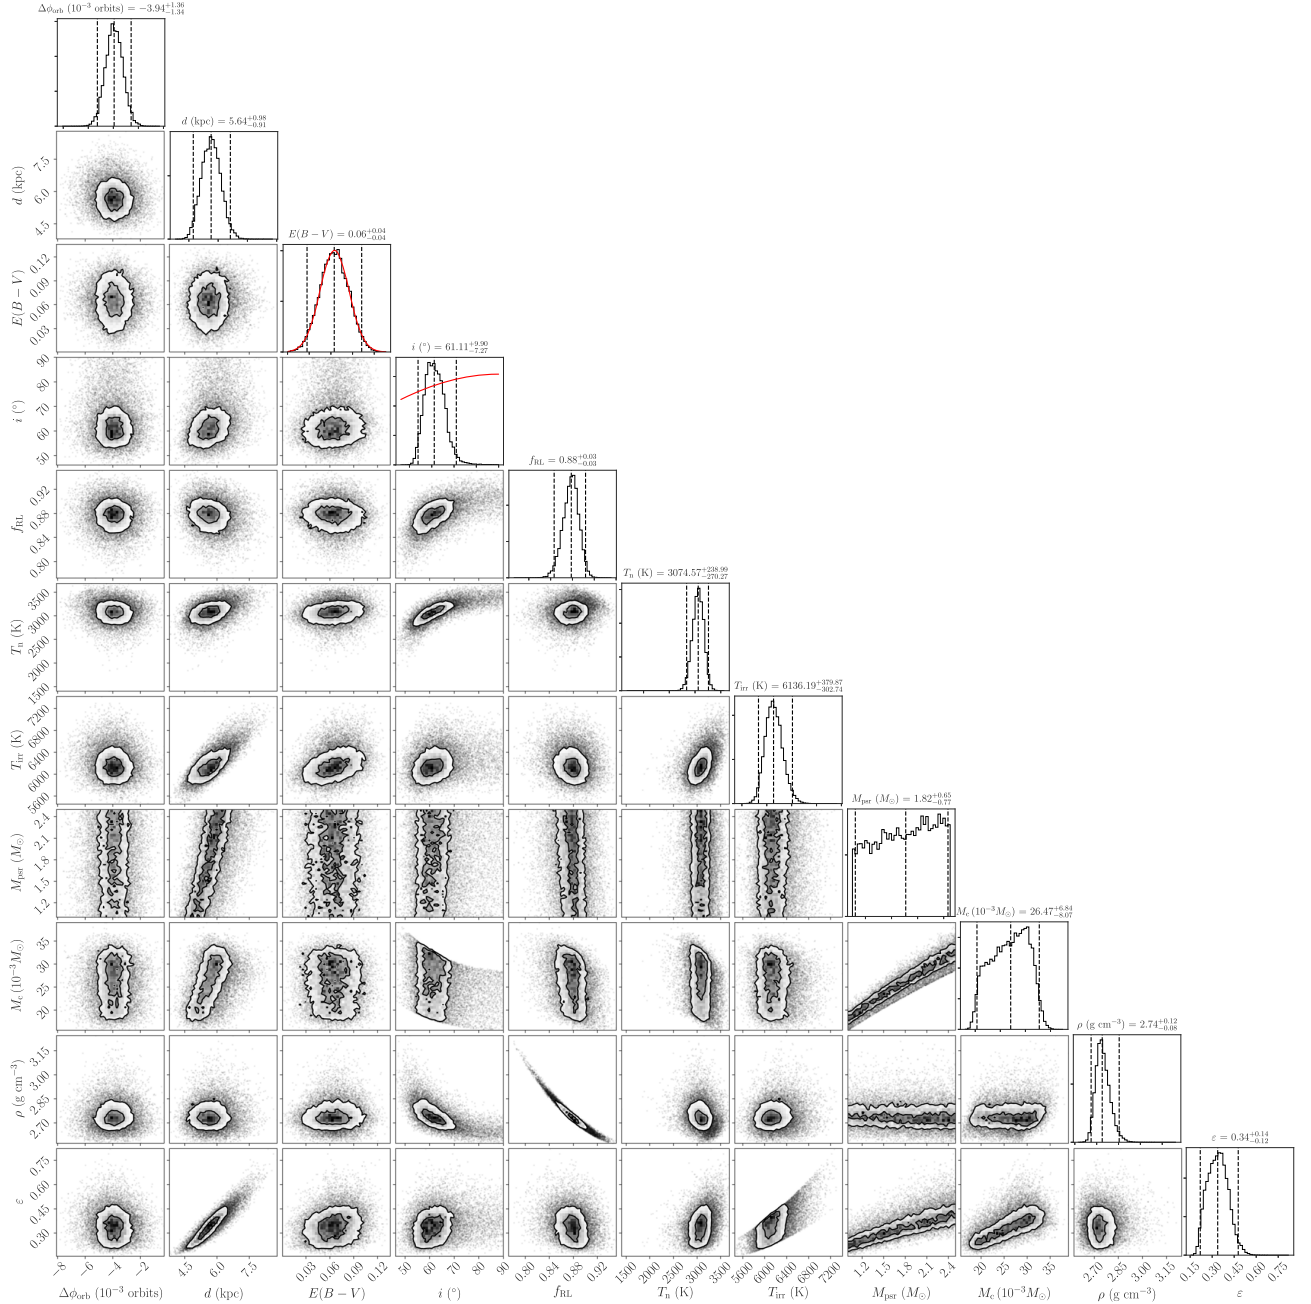### 4.4. Search for Continuous Gravitational Waves

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 (O1[23] and O2[24]) of the two Advanced LIGO detectors (Vallisneri et al. 2015). 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 $2\mathcal{F}$ (Jaranowski et al. 1998; Cutler & Schutz 2005) that we

[23] https://doi.org/10.7935/K57P8W9D
[24] https://doi.org/10.7935/CA75-FM95

9

Nieder et al.

Figure 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 nonuniform 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.

implemented in the LIGO-LALSUITE library.[25] $2\mathcal{F}$ is the log-likelihood maximized over the amplitude parameters h_0, $\cos\iota$, ψ and Φ_0 for a near-monochromatic[26] 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 equal to 0: the expected value is $\mu = 4.0$, and the

---

[25] https://git.ligo.org/lscsoft/lalsuite/

[26] The signal is not strictly monochromatic because of the measured nonzero spin-frequency derivative.

10

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 (Abbott et al. 2004). $h_0^{95\%}$ is the smallest intrinsic gravitational wave amplitude such that 95% of the population of signals that could be emitted by PSR J0952 −0607[27] 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 (Cahillane et al. 2017).
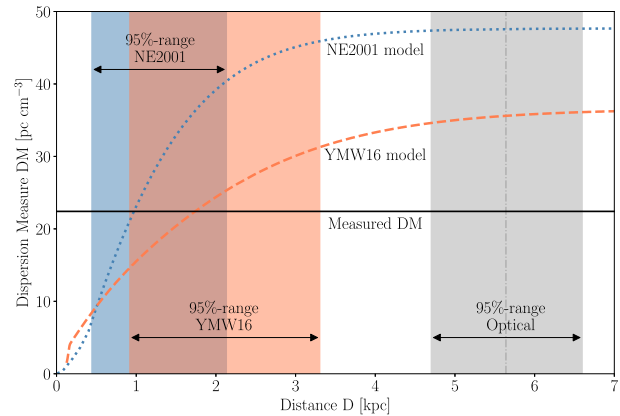
5. Discussion

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 $P_{\rm obs} = 1.414$ ms and the observed spin-period derivative is $\dot{P}_{\rm obs} = 4.76 \times 10^{-21}$ s s^{-1}.

The intrinsic spin-period derivative $\dot{P}_{\rm int}$ can be estimated from the observed value $\dot{P}_{\rm obs} = \dot{P}_{\rm int} + \dot{P}_{\rm Gal} + \dot{P}_{\rm Shk}$. $\dot{P}_{\rm Gal}$ 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., Damour & Taylor 1991; Nice & Taylor 1995), while $\dot{P}_{\rm Shk}$ accounts for the Shklovskii effect due to nonzero proper motion (Shklovskii 1970). Both contributions, $\dot{P}_{\rm Gal}$ and $\dot{P}_{\rm Shk}$, depend on the distance d to the pulsar.

The distance to PSR J0952−0607 is uncertain. The measured DM can be used to estimate the distance using Galactic electron-density models. The NE2001 model predicts $0.97^{+1.16}_{-0.53}$ kpc, while the YMW16 model predicts $1.74^{+1.57}_{-0.82}$ kpc. The uncertainties represent the 95% confidence regions (Yao et al. 2017). The model predictions of the DM as a function of d in the direction of the pulsar's sky position are shown in Figure 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 $5.64^{+0.98}_{-0.91}$ 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 $\dot{P}_{\rm Gal} = (1.7, 2.2, 3.6) \times 10^{-22}$ 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 $\dot{P}_{\rm Shk} \in ([0, 2.1], [0, 3.8]) \times 10^{-21}$ s s^{-1} from the proper motion 95% confidence region (see Section 3.2) and for the (NE2001, YMW16) distances. The resulting 95% confidence region on $\dot{P}_{\rm Shk}$ for the optical distance exceeds past $\dot{P}_{\rm obs}$. Thus we only constrain the intrinsic spin-frequency derivative (at 95% confidence) to $\dot{P}_{\rm int} \in [2.44, 4.59] \times 10^{-21}$ s s^{-1} for the NE2001 model and $\dot{P}_{\rm int} \in [0.69, 4.54] \times 10^{-21}$ s s^{-1} for the YMW16 model. In the following, we conservatively assume zero proper

---

[27] The possible signals span uniformly distributed values of $-1 \leqslant \cos\iota \leqslant 1$ and of $0 \leqslant \psi \leqslant 2\pi$.

Figure 4. Dispersion measure vs. distance from the NE2001 and YMW16 models at the sky position of PSR J0952−0607. For the measured DM = 22.4 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 (Yao et al. 2017). 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.

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

In Figure 5, 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 Catalogue (see footnote 19) (Manchester et al. 2005).

Furthermore we estimated the characteristic age $\tau_{\rm c}$, the spin-down power $\dot{E}$, the surface magnetic field strength $B_{\rm surf}$ and the magnetic field strength at the light cylinder $B_{\rm LC}$ (see Table 2). To calculate these values we assumed the pulsar to be a magnetic dipole with a canonical radius $r_{\rm psr} = 10$ km and moment of inertia $I_{\rm psr} = 10^{45}$ g cm^2 (e.g., Abdo et al. 2013). The same assumptions were used to plot the contour lines in Figure 5.

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 (Hou et al. 2014). 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, $d = (0.97, 1.74)$ kpc (NE2001, YMW16). The inferred gamma-ray luminosity is $L_\gamma = 4\pi d^2 F_\gamma\, f_\Omega \approx 3.1 \times 10^{32} \times (d/1\,{\rm kpc})^2$ erg s^{-1}. The measured LAT energy flux F_γ is given in Table 1 and we assumed no beaming (i.e., $f_\Omega = 1$). The gamma-ray efficiency is $\eta_\gamma = L_\gamma/\dot{E} \approx 0.5\% \times (d/1\,{\rm kpc})^2$. At the optical distance, $\eta_\gamma \approx 16\%$ is typical of gamma-ray MSPs (Abdo et al. 2013), while at the DM-derived distance, $\eta_\gamma \sim 1\%$ would be unusually low.

Due to the non-detection of PSR J0952−0607 in X-rays ($F_{\rm X} < 1.1 \times 10^{-13}$ erg s^{-1} cm^{-2}, Bassa et al. 2017b) we can only give a lower limit for the gamma-ray-to-X-ray flux ratio $F_\gamma/F_{\rm X} > 20$. This limit is at the lower end of the observed
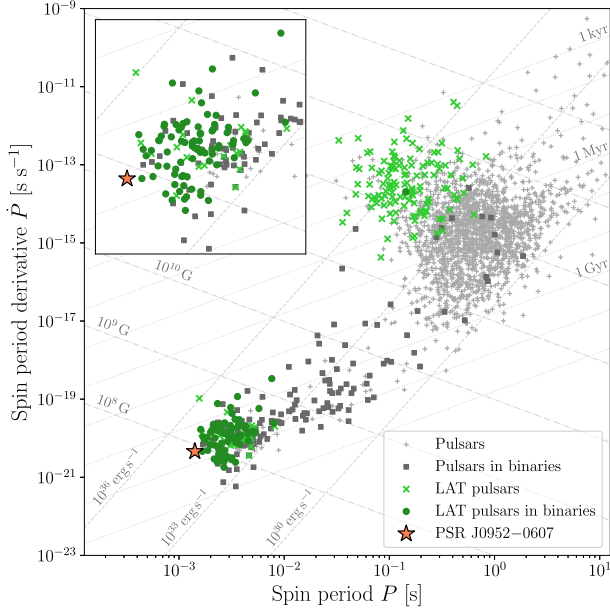
Figure 5. Spin period P and spin-period derivative $\dot{P}$ of the known pulsar population outside of globular clusters. The inset shows a zoomed-in view of the known MSP population. Isolated radio pulsars (light-gray pluses), binary radio pulsars (dark-gray squares), isolated gamma-ray pulsars (light-green 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).

distribution but still consistent with the literature (Marelli et al. 2011, 2015; Abdo et al. 2013; Salvetti et al. 2017).

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 $\phi_{orb} = 0.5$ and $\phi_{orb} = 1.0$ (with $\phi_{orb} = 0$ 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 $f_{RL} \approx 88\%$ and $i \approx 61°$. However, high filling factors imply a larger and hence more luminous companion, and therefore require greater distance, with our model having $d \sim 4.7$–6.6 kpc.

We tried to refit 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 $100 \, \mathrm{g \, cm}^{-3}$. For comparison, the densest known black-widow companions have densities of around $50 \, \mathrm{g \, cm}^{-3}$ (e.g., PSR J0636+5128 Kaplan et al. 2018), with the record being that of the black-widow candidate 3FGL J1653.6−0158 in a 75 minute orbit (Romani et al. 2014) where $\rho \gtrsim 70 \, \mathrm{g \, cm}^{-3}$. 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 (van Haaften et al. 2012). 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 $\rho \sim 2.75 \, \mathrm{g \, cm}^{-3}$, close to the density of brown dwarfs of similar mass and temperature given by the model considered in Kaplan et al. (2018).

We note that similar discrepancies in model distances were seen by Sanchez & Romani (2017) when using a direct-heating model. Romani & Sanchez (2016) and Sanchez & Romani (2017) 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 redirected 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., Romani & Sanchez 2016) 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 $T_n \approx 3000 \pm 250$ K at 95% confidence. We also find a well-constrained irradiating temperature of $T_{irr} = 6100 \pm 350$ K, higher than that found from the single-band fit performed in Bassa et al. (2017b). 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 (Breton et al. 2013)

$$\varepsilon = \frac{4\pi A^2 \sigma T_{irr}^4}{\dot{E}}, \qquad (8)$$

with $\varepsilon \sim 20\%$ being typical for black-widow systems. The efficiency is also shown in Figure 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 $\eta_\gamma \approx 0.5\% \times (d/1 \, \mathrm{kpc})^2$. This estimate assumes that the pulsar's heating flux is emitted isotropically. As pointed out by Draghis & Romani (2018), some models of pulsar gamma-ray emission predict stronger beaming toward 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 toward 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

12

Nieder et al.

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 (Romani & Sanchez 2016; Wadiasingh et al. 2017), 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 (Kennedy et al. 2018) despite a 5× increase in the gamma-ray flux (Stappers et al. 2014).

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 (Abdo et al. 2013). The radio pulse profile also shows two peaks with similar separation, with the radio pulse slightly leading the gamma-ray pulse (see Figure 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; Abdo et al. 2013). Due to a covariance between f and dispersion measure (see Section 4.1) we were not able to measure significant variations in the dispersion measure. A change in dispersion measure of 10^{-3} 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 10 fastest Galactic field pulsars only 1 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 (Guillemot et al. 2012). Still, the mass-shedding spin limit for neutron stars is typically placed much higher at around 1200 Hz (Cook et al. 1994; Lattimer & Prakash 2004). One mechanism that could prevent neutron stars from spinning up to
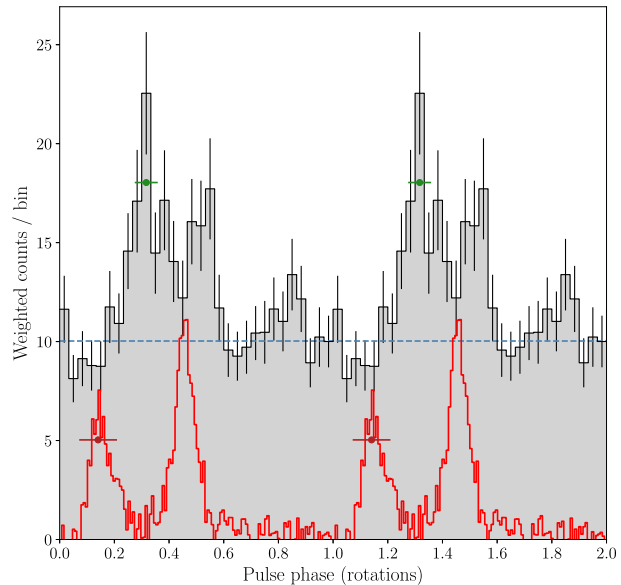

Figure 6. Aligned integrated gamma-ray and radio pulse profiles of PSR J0952 −0607 over two identical rotations. The black curve shows the weighted LAT photon counts after MJD 55,750 in a histogram with 30 bins per rotation. The green error bars show the phase uncertainty of the gamma-ray pulse profile. 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.

higher frequencies is the emission of gravitational waves (for a recent work on this subject see, e.g., Gittins & Andersson 2019). Another option could be that the spin-up torque might be smaller for faster pulsars with lower magnetic field strengths (Patruno et al. 2012; Bonanno & Urpin 2015).

The estimated intrinsic spin-period derivative implies a very low surface magnetic field of 8.2×10^7 G for PSR J0952 −0607. Assuming nonzero 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 Catalogue (see footnote 19) (Manchester et al. 2005), PSR J1544+4937 (Bhattacharyya et al. 2013), show lower inferred surface magnetic fields (Figure 7). The surface *B*-field of the other recent LOFAR-detected pulsar, PSR J1552+5437, is only slightly bigger (Pleunis et al. 2017). This might be a hint that pulsars with low *B*-fields also have steeper radio spectra.

The pulsar distribution in Figure 7 indicates a lower limit on the magnetic field strength independent of the spin frequency. The equilibrium spin period as predicted by Alpar et al. (1982) is $P_{eq} \propto B_{surf}^{6/7} R_{psr}^{18/7} M_{psr}^{-5/7} \dot{M}_{accr}^{-3/7}$ 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 magnetic field strengths the angular momentum transfer is slower (Bonanno & Urpin 2015). 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 (Pan et al. 2018). For a neutron star with a mass

13
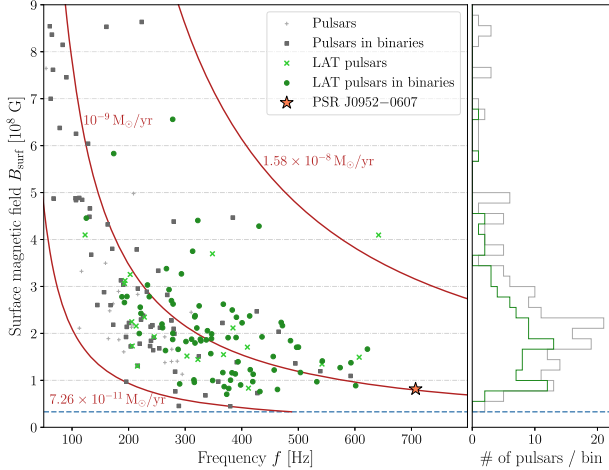
Figure 7. Frequency f and surface magnetic field strength B_{surf} of the known MSP population outside of globular clusters. The surface magnetic field of PSR J0952−0607 is computed assuming $\dot{P}_{Shk} = 0$ and thus represents an upper limit. The horizontal-dashed blue line represents a possible 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 5. 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).

of $1.4\,M_\odot$, a radius of 10 km and a minimum accretion rate of $7.26 \times 10^{-11}\,M_\odot\,\mathrm{yr}^{-1}$ we get a minimum magnetic field strength of $B_{surf} \gtrsim 3.3 \times 10^7$ 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 $\epsilon^{95\%} = 3.1 \times 10^{-8} \times (d/1\,\mathrm{kpc}) \times (10^{45}\,\mathrm{g\,cm^2}/I)$, 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 $h_0^{sd} = 1.5 \times 10^{-27} \times (1\,\mathrm{kpc}/d) \times (I/10^{45}\,\mathrm{g\,cm^2})^{1/2}$, corresponding to an ellipticity of $\epsilon^{sd} = 7.0 \times 10^{-10} \times (1\,\mathrm{kpc}/d)$.

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 (Abbott et al. 2019). 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 $10^{45}\,\mathrm{g\,cm^2}$. 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 Abbott et al. 2019, 2017; Aasi et al. 2014 and references therein) we all the same perform the search.

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 Bassa et al. (2017b). 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 $\dot{P}_{int} \lesssim 4.6 \times 10^{-21}\,\mathrm{s\,s^{-1}}$. This measurement provides estimates of physical parameters such as the spin-down luminosity ($\dot{E} \lesssim 6.4 \times 10^{34}\,\mathrm{erg\,s^{-1}}$), and a surface magnetic field ($B_{surf} \lesssim 8.2 \times 10^7$ 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 multiband 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 (~ 5 kpc) than the DM-distance estimates of 0.97 kpc (NE2001) to 1.74 kpc (YMW16), or a small and extremely dense companion $\rho \gg 100\,\mathrm{g\,cm^{-3}}$. 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 (Ray et al. 2012). Sophisticated methods to identify more pulsar candidates within LAT sources have been developed (e.g., Lee et al. 2012; Saz Parkinson et al. 2016). For instance, Frail et al. (2016b) 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.

Nieder et al.

Appendix
Estimating the False-alarm Probability for a Multidimensional 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 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, multidimensional parameter grid. Deriving the probability distribution for the maximum value of a multidimensional "random field" is difficult and approximate solutions are only known for simple cases such as Gaussian or chi-squared random fields (Adler & Taylor 2007). 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 Kerr 2011): the false-alarm probability depends on the volume, shape, and dimensionality of the search space (Adler & Taylor 2007). A full discussion of this is beyond the scope of this work. Below, we show empirically that a simple trials factor approach overestimates 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

$$P_{FA}(H_m \mid a) = e^{-a\,H_m}, \qquad (9)$$

with scaling factor $a = 0.3984$ (de Jager & Büsching 2010; Kerr 2011). This formula can be used to estimate the significance of the maximum H statistic value after n *independent* trials

$$P_{FA}(H_m \mid a, n) = 1 - [1 - e^{-a\,H_m}]^n. \qquad (10)$$

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 $n_{seg} = 2 \times 17 \times 13 = 442$ 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 2.2) to the pulsar parameters.

15
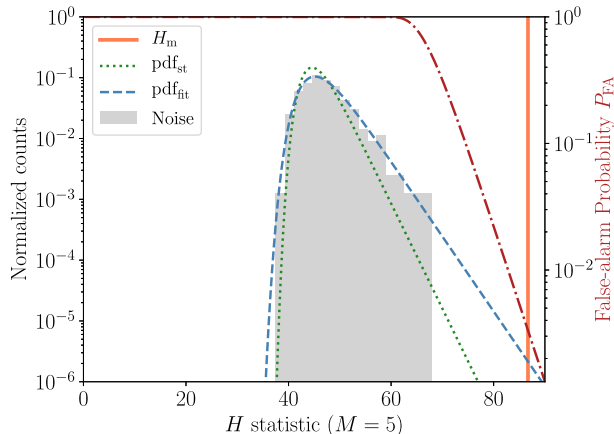
Figure 8. Normalized histogram showing the highest H statistics for 442 subsets of our search space after excluding results affected by the pulsar signal. The dotted green and dashed blue curves show normalized probability density functions for the maximum H statistic obtained after n effective trials. 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 $H_{\mathrm{m}} = 86.7$ is marked by the vertical orange line. The red line (dashed–dotted) shows the false-alarm probability depending on H_{m} computed with Equation (10) with a and n from the joint variation.

The highest H statistic of each of the segments is plotted in the normalized histogram in Figure 8. We fit for the effective number of trials (as done by, e.g., Kruger et al. 2002) by maximizing the likelihood,

$$L(n, a|H_{m,i}) = \prod_i p(H_{m,i}|a, n) \qquad (11)$$

for our set of H statistic values, according to the probability density function for H_{m} after n trials (the derivative of Equation (10))

$$p(H_{\mathrm{m}}|a, n) = a\,n\,[1 - e^{-a\,H_{\mathrm{m}}}]^{n-1} \exp(-a\,H_{\mathrm{m}}). \qquad (12)$$

However, as shown in Figure 8, the tail of the best-fitting distribution is significantly underestimated, leading to *overestimated* 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 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, among 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 (10) to obtain an approximation to the false-alarm probability for the maximum H statistic value. For the candidate pulsar signal, this was $P_{\mathrm{FA}} = 0.33\%$. For comparison the candidate with the largest H statistic from a segment of the search not affected by the pulsar signal had $P_{\mathrm{FA}} = 56\%$.

ORCID iDs

L. Nieder ⓘ https://orcid.org/0000-0002-5775-8977
C. J. Clark ⓘ https://orcid.org/0000-0003-4355-3572
C. G. Bassa ⓘ https://orcid.org/0000-0002-1429-9010
J. Wu ⓘ https://orcid.org/0000-0003-3536-4368
B. Allen ⓘ https://orcid.org/0000-0003-4285-6256
R. P. Breton ⓘ https://orcid.org/0000-0001-8522-4983
V. S. Dhillon ⓘ https://orcid.org/0000-0003-4236-9642
H.-B. Eggenstein ⓘ https://orcid.org/0000-0001-5296-7035
J. W. T. Hessels ⓘ https://orcid.org/0000-0003-2317-1446
M. R. Kennedy ⓘ https://orcid.org/0000-0001-6894-6044
M. Kerr ⓘ https://orcid.org/0000-0002-0893-4073
S. Littlefair ⓘ https://orcid.org/0000-0001-7221-855X
T. R. Marsh ⓘ https://orcid.org/0000-0002-2498-7589
D. Mata Sánchez ⓘ https://orcid.org/0000-0003-0245-9424
M. A. Papa ⓘ https://orcid.org/0000-0002-1007-5298
P. S. Ray ⓘ https://orcid.org/0000-0002-5297-5278
B. Steltner ⓘ https://orcid.org/0000-0003-1833-5493
J. P. W. Verbiest ⓘ https://orcid.org/0000-0002-4088-896X

References

Aasi, J., Abadie, J., Abbott, B. P., et al. 2014, ApJ, 785, 119
Abbott, B., Abbott, R., Adhikari, R., et al. 2004, PhRvD, 69, 082004
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, ApJ, 839, 12
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, arXiv:1902.08507
Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, Sci, 325, 840
Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, ApJS, 188, 405
Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, ApJS, 208, 17
Acero, F., Ackermann, M., Ajello, M., et al. 2015, ApJS, 218, 23
Acero, F., Ackermann, M., Ajello, M., et al. 2016, ApJS, 223, 26
Adler, R. J., & Taylor, J. E. 2007, Random Fields and Geometry, Springer Monographs in Mathematics (New York, NY: Springer)
Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, Natur, 300, 728
An, H., Romani, R. W., Johnson, T., Kerr, M., & Clark, C. J. 2017, ApJ, 850, 100
Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, Sci, 324, 1411
Arzoumanian, Z., Fruchter, A. S., & Taylor, J. H. 1994, ApJL, 426, 85
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Atwood, W., Albert, A., Baldini, L., et al. 2013, arXiv:1303.3514
Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, ApJ, 697, 1071
Aulbert, C., & Fehrmann, H. 2009, Forschungsbericht 2009—Max-Planck-Institut für Gravitationsphysik, Teilinstitut Hannover, https://www.mpg.de/308429/forschungsSchwerpunkt
Balasubramanian, R., Sathyaprakash, B. S., & Dhurandhar, S. V. 1996, PhRvD, 53, 3033
Barr, E. D., Guillemot, L., Champion, D. J., et al. 2013, MNRAS, 429, 1633
Bassa, C. G., Pleunis, Z., & Hessels, J. W. T. 2017a, A&C, 18, 40
Bassa, C. G., Pleunis, Z., Hessels, J. W. T., et al. 2017b, ApJL, 846, L20
Bates, S. D., Lorimer, D. R., & Verbiest, J. P. W. 2013, MNRAS, 431, 1352
Bhattacharyya, B., Roy, J., Ray, P. S., et al. 2013, ApJL, 773, L12
Bickel, P., Kleijn, B., & Rice, J. 2008, ApJ, 685, 384
Bonanno, A., & Urpin, V. 2015, MNRAS, 451, 2117
Breton, R. P., Rappaport, S. A., van Kerkwijk, M. H., & Carter, J. A. 2012, ApJ, 748, 115
Breton, R. P., van Kerkwijk, M. H., Roberts, M. S. E., et al. 2013, ApJ, 769, 108
Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, A&A, 564, A125
Cahillane, C., Betzwieser, J., Brown, D. A., et al. 2017, PhRvD, 96, 102001
Camilo, F., Kerr, M., Ray, P. S., et al. 2015, ApJ, 810, 85
Caraveo, P. A. 2014, ARA&A, 52, 211
Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
Chen, H.-L., Chen, X., Tauris, T. M., & Han, Z. 2013, ApJ, 775, 27
Clark, C. J., Pletsch, H. J., Wu, J., et al. 2015, ApJL, 809, L2
Clark, C. J., Pletsch, H. J., Wu, J., et al. 2018, SciA, 4, eaao7228
Clark, C. J., Wu, J., Pletsch, H. J., et al. 2017, ApJ, 834, 106
Cognard, I., Guillemot, L., Johnson, T. J., et al. 2011, ApJ, 732, 47

Coles, W., Hobbs, G., Champion, D. J., Manchester, R. N., & Verbiest, J. P. W. 2011, MNRAS, 418, 561
Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1994, ApJ, 424, 823
Cordes, J. M., & Lazio, T. J. W. 2002, arXiv:astro-ph/0207156
Cromartie, H. T., Camilo, F., Kerr, M., et al. 2016, ApJ, 819, 34
Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006
Damour, T., & Taylor, J. H. 1991, ApJ, 366, 501
de Jager, O. C., & Büsching, I. 2010, A&A, 517, L9
de Jager, O. C., Raubenheimer, B. C., & Swanepoel, J. W. H. 1989, A&A, 221, 180
Dhillon, V., Dixon, S., Gamble, T., et al. 2018, Proc. SPIE, 10702, 10702LD
Dhillon, V. S., Marsh, T. R., Bezawada, N., et al. 2016, Proc. SPIE, 9908, 99080Y
Dhillon, V. S., Marsh, T. R., Stevenson, M. J., et al. 2007, MNRAS, 378, 825
Draghis, P., & Romani, R. W. 2018, ApJL, 862, L6
Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549
Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, arXiv:1306.2144
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Frail, D. A., Jagannathan, P., Mooley, K. P., & Intema, H. T. 2016a, ApJ, 829, 119
Frail, D. A., Mooley, K. P., Jagannathan, P., & Intema, H. T. 2016b, MNRAS, 461, 1062
Frail, D. A., Ray, P. S., Mooley, K. P., et al. 2018, MNRAS, 475, 942
Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, Natur, 333, 237
Gittins, F., & Andersson, N. 2019, MNRAS, 488, 99
Goodman, J., & Weare, J. 2010, CAMCS, 5, 65
Green, G. M., Schlafly, E. F., Finkbeiner, D., et al. 2018, MNRAS, 478, 651
Guillemot, L., Johnson, T. J., Venter, C., et al. 2012, ApJ, 744, 33
Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, Sci, 311, 1901
Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, MNRAS, 369, 655
Honeycutt, R. K. 1992, PASP, 104, 435
Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, PASA, 21, 302
Hou, X., Smith, D. A., Guillemot, L., et al. 2014, A&A, 570, A44
Hunter, J. D. 2007, CSE, 9, 90
Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, A&A, 553, A6
Jaodand, A., Hessels, J. W. T., & Archibald, A. 2018, in IAU Symp. 337, Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede (Cambridge: Cambridge Univ. Press), 47
Jaranowski, P., Królak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001
Johnson, T. J., Ray, P. S., Roy, J., et al. 2015, ApJ, 806, 91
Kaplan, D. L., Stovall, K., van Kerkwijk, M. H., Fremling, C., & Istrate, A. G. 2018, ApJ, 864, 15
Kennedy, M. R., Clark, C. J., Voisin, G., & Breton, R. P. 2018, MNRAS, 477, 1120
Kerr, M. 2011, ApJ, 732, 38
Kerr, M., Ray, P. S., Johnston, S., Shannon, R. M., & Camilo, F. 2015, ApJ, 814, 128
Kruger, A. T., Loredo, T. J., & Wasserman, I. 2002, ApJ, 576, 932
Lattimer, J. M., & Prakash, M. 2004, Sci, 304, 536
Lee, K. J., Guillemot, L., Yue, Y. L., Kramer, M., & Champion, D. J. 2012, MNRAS, 424, 2832
Levin, L., McLaughlin, M. A., Jones, G., et al. 2016, ApJ, 818, 166

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993
Marelli, M., De Luca, A., & Caraveo, P. A. 2011, ApJ, 733, 82
Marelli, M., Mignani, R. P., De Luca, A., et al. 2015, ApJ, 802, 78
Mukherjee, D. 2017, JApA, 38, 48
Nice, D. J., & Taylor, J. H. 1995, ApJ, 441, 429
Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, ApJS, 199, 31
Oliphant, T. E. 2006, A Guide to NumPy, Vol. 1 (USA: Trelgol Publishing)
Owen, B. J. 1996, PhRvD, 53, 6749
Pan, Y. Y., Zhang, C. M., Song, L. M., et al. 2018, MNRAS, 480, 692
Patruno, A., Haskell, B., & D'Angelo, C. 2012, ApJ, 746, 9
Pitkin, M. 2018, JOSS, 3, 538
Pletsch, H. J., & Clark, C. J. 2014, ApJ, 795, 75
Pletsch, H. J., & Clark, C. J. 2015, ApJ, 807, 18
Pletsch, H. J., Guillemot, L., Fehrmann, H., et al. 2012, Sci, 338, 1314
Pleunis, Z., Bassa, C. G., Hessels, J. W. T., et al. 2017, ApJL, 846, L19
Polzin, E. J., Breton, R. P., Clarke, A. O., et al. 2018, MNRAS, 476, 1968
Prager, B. J., Ransom, S. M., Freire, P. C. C., et al. 2017, ApJ, 845, 148
Ransom, S. M., Ray, P. S., Camilo, F., et al. 2011, ApJL, 727, L16
Ray, P. S., Abdo, A. A., Parent, D., et al. 2012, arXiv:1205.3089
Ray, P. S., Kerr, M., Parent, D., et al. 2011, ApJS, 194, 17
Roberts, M. S. E. 2013, in IAU Symp. 291, Neutron Stars and Pulsars: Challenges and Opportunities after 80 years, ed. J. van Leeuwen (Cambridge: Cambridge Univ. Press), 127
Roberts, M. S. E., Mclaughlin, M. A., Gentile, P., et al. 2014, AN, 335, 313
Romani, R. W., Filippenko, A. V., & Cenko, S. B. 2014, ApJL, 793, L20
Romani, R. W., & Sanchez, N. 2016, ApJ, 828, 7
Salvetti, D., Mignani, R. P., De Luca, A., et al. 2017, MNRAS, 470, 466
Sanchez, N., & Romani, R. W. 2017, ApJ, 845, 42
Saz Parkinson, P. M., Xu, H., Yu, P. L. H., et al. 2016, ApJ, 820, 8
Schwarz, G. 1978, AnSta, 6, 461
Shklovskii, I. S. 1970, SvA, 13, 562
Smith, J. A., Tucker, D. L., Kent, S., et al. 2002, AJ, 123, 2121
Stappers, B. W., Archibald, A. M., Hessels, J. W. T., et al. 2014, ApJ, 790, 39
Stappers, B. W., Bailes, M., Lyne, A. G., et al. 1996, ApJL, 465, L119
Stappers, B. W., Hessels, J. W. T., Alexov, A., et al. 2011, A&A, 530, A80
Story, S. A., Gonthier, P. L., & Harding, A. K. 2007, ApJ, 671, 713
Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, ApJ, 791, 67
The Fermi-LAT Collaboration 2019, arXiv:1902.10045
Tiburzi, C. 2018, PASA, 35, e013
Vallisneri, M., Kanner, J., Williams, R., Weinstein, A., & Stephens, B. 2015, JPhCS, 610, 012021
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
van Haaften, L. M., Nelemans, G., Voss, R., Wood, M. A., & Kuijpers, J. 2012, A&A, 537, A104
van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2
van Straten, W., & Bailes, M. 2011, PASA, 28, 1
van Straten, W., Demorest, P., & Oslowski, S. 2012, AR&T, 9, 237
Verbiest, J. P. W., Lentati, L., Hobbs, G., et al. 2016, MNRAS, 458, 1267
Wadiasingh, Z., Harding, A. K., Venter, C., Böttcher, M., & Baring, M. G. 2017, ApJ, 839, 80
Wu, J., Clark, C. J., Pletsch, H. J., et al. 2018, ApJ, 854, 99
Yao, J. M., Manchester, R. N., & Wang, N. 2017, ApJ, 835, 29

Discovery of a Gamma-Ray Black Widow Pulsar by GPU-accelerated Einstein@Home

The discovery of the gamma-ray black widow pulsar PSR J1653-0158 with an Einstein@Home-search is presented.

The author prepared the data for the continuous gravitational wave search with the method described in Chapter 2.

CrossMark

Discovery of a Gamma-Ray Black Widow Pulsar by GPU-accelerated Einstein@Home

L. Nieder[1,2] , C. J. Clark[3] , D. Kandel[4] , R. W. Romani[4] , C. G. Bassa[5] , B. Allen[1,6,2] , A. Ashok[1,2] , I. Cognard[7,8] ,
H. Fehrmann[1,2] , P. Freire[9] , R. Karuppusamy[9] , M. Kramer[9,3] , D. Li[10,11] , B. Machenschalk[1,2] , Z. Pan[10] ,
M. A. Papa[1,6,2] , S. M. Ransom[12] , P. S. Ray[13] , J. Roy[14] , P. Wang[10] , J. Wu[9] , C. Aulbert[1,2] , E. D. Barr[9] ,
B. Beheshtipour[1,2] , O. Behnke[1,2] , B. Bhattacharyya[14] , R. P. Breton[3] , F. Camilo[15] , C. Choquet[16] ,
V. S. Dhillon[17,18] , E. C. Ferrara[19,20] , L. Guillemot[7,8] , J. W. T. Hessels[5,21] , M. Kerr[13] , S. A. Kwang[6] , T. R. Marsh[22] ,
M. B. Mickaliger[3] , Z. Pleunis[23,24] , H. J. Pletsch[1] , M. S. E. Roberts[25,26] , S. Sanpa-arsa[27] , and B. Steltner[1,2]
[1] Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), 30167 Hannover, Germany; lars.nieder@aei.mpg.de
[2] Leibniz Universität Hannover, 30167 Hannover, Germany
[3] Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK
[4] KIPAC/Dept. of Physics, Stanford University, Stanford, CA 94305, USA
[5] ASTRON, The Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, 7991 PD Dwingeloo, The Netherlands
[6] Department of Physics, University of Wisconsin–Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA
[7] Laboratoire de Physique et Chimie de l'Environnement et de l'Espace, Université d'Orléans/CNRS, F-45071 Orléans Cedex 02, France
[8] Station de radioastronomie de Nançay, Observatoire de Paris, CNRS/INSU, F-18330 Nançay, France
[9] Max-Planck-Institut für Radioastronomie, auf dem Hügel 69, 53121 Bonn, Germany
[10] National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100101, People's Republic of China
[11] NAOC-UKZN Computational Astrophysics Centre, University of KwaZulu-Natal, Durban 4000, South Africa
[12] National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA, 22903, USA
[13] Space Science Division, Naval Research Laboratory, Washington, DC 20375-5352, USA
[14] National Centre for Radio Astrophysics, Tata Institute of Fundamental Research, Pune 411 007, India
[15] South African Radio Astronomy Observatory, 2 Fir Street, Black River Park, Observatory 7925, South Africa
[16] Résidence Le Dauphiné, rue Jean Bleuzen, Vanves, France
[17] Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, UK
[18] Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain
[19] NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[20] Department of Astronomy, University of Maryland, College Park, MD 20742, USA
[21] Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
[22] Astronomy and Astrophysics Group, Department of Physics, University of Warwick, Coventry CV4 7AL, UK
[23] Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada
[24] McGill Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada
[25] New York University Abu Dhabi, P.O. Box 129188, Abu Dhabi, UAE
[26] Eureka Scientific, Inc., 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA
[27] National Astronomical Research Institute of Thailand (Public Organization), 260 Moo 4, T. Donkaew, A. Maerim, Chiang Mai, 50180, Thailand

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 multidimensional 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 $B_{surf} \approx 4 \times 10^7$ G. The resulting mass function, combined with models of the companion star's optical light curve and spectra, suggests a pulsar mass $\gtrsim 2 M_\odot$. The companion is lightweight 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.

Unified Astronomy Thesaurus concepts: Gamma-ray sources (633); Millisecond pulsars (1062); Neutron stars (1108); Binary pulsars (153)

Supporting material: data behind figure

1

1. Introduction

The Fermi Large Area Telescope (LAT) source 4FGL J1653.6−0158 is a bright gamma-ray source, and the brightest remaining unassociated source (Saz Parkinson et al. 2016). It was first seen by the Energetic Gamma Ray Experiment Telescope (EGRET; Hartman et al. 1999), and was also listed in the LAT Bright Gamma-ray source list (Abdo et al. 2009) more than a decade ago. While pulsars were discovered in several other sources from this list (see, e.g., Ransom et al. 2011), the origin of 4FGL J1653.6−0158 remained unidentified. The detection of a variable X-ray and optical candidate counterpart with 75 minute period consistent with the gamma-ray position of 4FGL J1653.6−0158 provided strong evidence of it being a binary gamma-ray pulsar (Kong et al. 2014; Romani et al. 2014).

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 (Nieder et al. 2020). 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 minute orbital period. This very short orbital period raises interesting questions about the evolutionary path which created the system.

This Letter is organized as follows. In Section 2, 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 3. We discuss the results and conclude in Section 4.

2. Gamma-Ray Pulsations

2.1. Data Preparation

We searched for gamma-ray pulsations in the arrival times of photons observed by the Fermi-LAT (Atwood et al. 2009) between 2008 August 3 and 2018 April 16 (MJDs 54,681 and 58,224). We included SOURCE-class photons according to the P8R2_SOURCE_V6 (Atwood et al. 2012) instrument response functions (IRFs),[28] 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 position as reported in the Gaia DR2 Catalog (hereafter Gaia catalog; Gaia Collaboration et al. 2018). The celestial parameters (J2000.0) are $\alpha = 16^{\mathrm{h}}53^{\mathrm{m}}38\overset{s}{.}05381(5)$ and $\delta = -01°58'36\overset{''}{.}8930(5)$, 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 jth 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 yr source catalog[29] 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 (Kerr 2011), and the distribution of weights can be used to predict expected signal-to-noise ratios (Nieder et al. 2020).

The data set used here consisted of $N = 354,009$ photons, collected over a period of 3542 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 Equation (11) in Nieder et al. (2020).

Following the pulsar discovery, we extended this data set to 2020 February 23 (MJD 58,902), using the latest P8R3_SOUR-CE_V2 IRFs (Bruel et al. 2018), a larger maximum zenith angle of 105°, and using the Fermi-LAT Fourth Source Catalog (hereafter 4FGL; Abdollahi et al. 2020) as the RoI model for the photon probability weight computations.

2.2. Search

The binary-pulsar search methods are described by Nieder et al. (2020), which are a generalization and extension of the isolated-pulsar search methods from Pletsch & Clark (2014).

The searched ranges are guided by the known millisecond pulsar (MSP) population in the Australia Telescope National Facility (ATNF) Pulsar Catalogue[30] (Manchester et al. 2005). For the spin frequency, we searched $f \in [0, 1500]$ Hz.[31] 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.

2.2.1. Orbital Constraints from Optical Observations

The orbital-period estimate of Romani et al. (2014) was derived from Southern Astrophysical Research (SOAR), WIYN, and Catalina Sky Survey (CSS) observations. These were augmented by new 350 s SOAR Goodman High Throughput

---

[28] See https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_essentials.html.

[29] https://fermi.gsfc.nasa.gov/ssc/data/access/lat/fl8y/

[30] http://www.atnf.csiro.au/research/pulsar/psrcat

[31] The 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 Pletsch & Clark 2014). Note that the current record spin frequency is 716 Hz (Hessels et al. 2006).

2

Spectrograph (GHTS) 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. (2014) using the Wide Field camera (WFC) on the 2.5 m 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 multiband imager ULTRA-CAM (Dhillon et al. 2007) on the 4.2 m 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″, particularly at the beginning of the observation. We therefore only used the second night's data for the optical light-curve modeling in Section 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 Sloan Digital Sky Survey (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 multiband Lomb–Scargle periodogram method (VanderPlas & Ivezić 2015, 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 $P_{\rm orb} = 0.0519447518 \pm 6.0 \times 10^{-9}$ 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 et al. (2014), the time of the pulsar's ascending node, $T_{\rm asc}$, was estimated from the photometric light curve. However, the optical maximum is distinctly asymmetric (see Section 3.1), which can bias orbital phase estimates. We therefore used the spectroscopic radial-velocity measurements from Romani et al. (2014), folded at the orbital period obtained above, and fit the phase of a sinusoidal radial-velocity curve, finding $T_{\rm asc} = $ MJD $56,513.47981 \pm 2.1 \times 10^{-4}$. 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 (Roberts 2013; Strader et al. 2019). Here, a_1 is the pulsar's semimajor axis, i denotes the inclination angle, and c is the speed of light. As described in Nieder et al. (2020), we expected $x \in [0, 0.2]$ s based on the companion's velocity amplitude reported by Romani et al. (2014) and the masses

expected for "spider" companions, i.e., black-widow or redback companions.

2.2.2. Search Grids

To cover the relevant orbital-parameter space in $\{x, P_{\rm orb}, T_{\rm asc}\}$, we use *optimized grids* (Fehrmann & Pletsch 2014). 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 (Nieder et al. 2020). 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 (semi-coherent 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 (Aulbert & Fehrmann 2008).

2.2.3. Einstein@Home

Searching the five-dimensional parameter space $\{f, \dot{f}, x, P_{\rm orb}, T_{\rm asc}\}$ 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 (Allen et al. 2013). As done for radio pulsar searches previously, the search code utilizes the approximately 10,000 GPUs active on Einstein@Home for a computing speedup 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 2 weeks.

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 Nieder et al. 2020). 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 10 most significant candidates in its work unit, and therefore passed on to the coherent stage. In the coherent stage, it was very significant, with a 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

Nieder et al.

Table 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)	$16^{\rm h}53^{\rm m}38\!\!\stackrel{\rm s}{.}05381(5)$
Decl., δ (J2000.0)	$-01°58'36\!\!\stackrel{''}{.}8930(5)$
Positional epoch (MJD)	57205.875
Proper motion in R.A., $\mu_\alpha \cos\delta$ (mas yr^{-1})	-19.62 ± 1.86
Proper motion in decl., 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})	$-6.204(8) \times 10^{-16}$
Spin period, P (ms)	1.9676820247057(2)
Spin-period derivative, $\dot{P}$ (s s^{-1})	$2.402(3) \times 10^{-21}$
Proj. semimajor axis, x (s)	0.01071(1)
Orbital period, $P_{\rm orb}$ (days)	0.0519447575(4)
Epoch of ascending node, $T_{\rm asc}$ (MJD)	56513.479171(8)
Derived Parameters for Distance $d = 840$ pc	
Shklovskii spin-down, $\dot{P}_{\rm Shk}$ (s s^{-1})	1.6×10^{-21}
Galactic acceleration spin-down, $\dot{P}_{\rm Gal}$ (s s^{-1})	-4.8×10^{-23}
Spin-down power, $\dot{E}$ (erg s^{-1})	4.4×10^{33}
Surface B-field, $B_{\rm surf}$ (G)	4.1×10^{7}
Light-cylinder B-field, $B_{\rm LC}$ (G)	5.0×10^{4}
Characteristic age, $\tau_{\rm c}$ (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 $d = 533^{+625}_{-187}$ pc, but for the derived parameters the consistent distance $d = 840^{+40}_{-40}$ pc derived from optical modeling is used (see Table 2).
[b] Taken from 4FGL Source Catalog (Abdollahi et al. 2020).

because the signal has significant power in the second harmonic.

Note that the signal was found outside the 3σ range in $T_{\rm asc}$ from the constraints reported in this work, and outside the 3σ range given by Romani et al. (2014). This can be caused by asymmetric heating (see Section 2.2.1).

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 Clark et al. (2017), which are an extension of the methods by Kerr et al. (2015). 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 Nieder et al. (2019). In the case of PSR J1653−0158, we used a template consisting of two symmetrical, wrapped Gaussian peaks. We used constraints on the peaks' FWHM, such that the peaks must be broader than 5% of a rotation, and narrower than half a rotation.
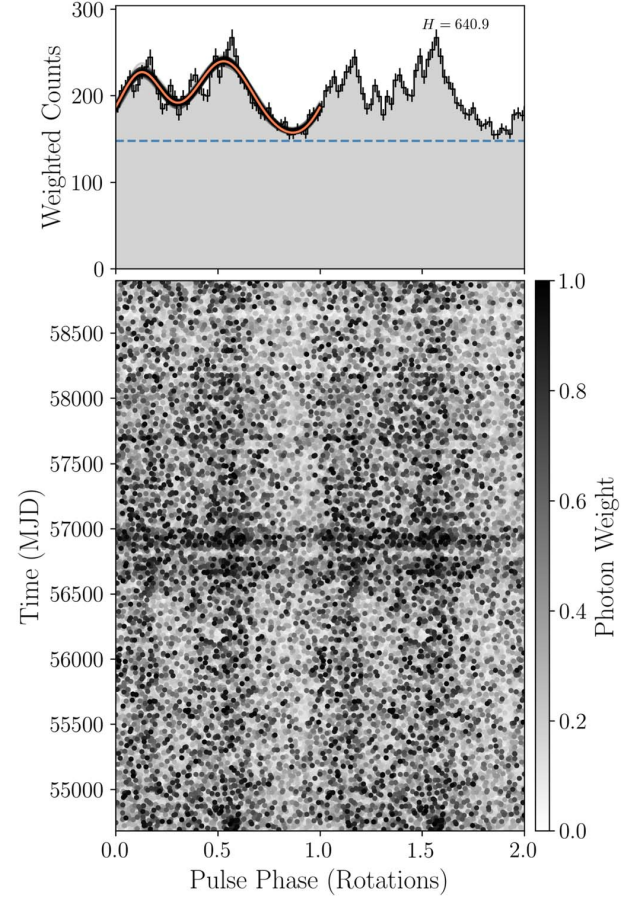
Figure 1. Integrated pulse-profile and phase-time diagram of PSR J1653 −0158, showing two identical rotations. Top: the histogram shows the weighted counts for 50 bins. The orange curve indicates the pulse-profile template with the highest signal power, and the transparent black curves represent 100 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 toward the Galactic center.

Our timing solution over 11 yr of LAT data is shown in Table 1. The folded gamma-ray data and the pulse profile are portrayed in Figure 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 (Shklovskii 1970), and the Galactic acceleration (see, e.g., Damour & Taylor 1991). The results are summarized in Table 1. The observed contribution due to the difference in Galactic acceleration of the Sun and the pulsar is computed with $R_{\rm Sun} = 8.21$ kpc, $z_{\rm Sun} = 14$ pc, and the Galactic potential model PJM17_best.Tpot (McMillan 2017), as implemented in their code.[32] For PSR J1653−0158, we used $R_{\rm J1653} = 7.48$ kpc, and $z_{\rm J1653} = 367$ pc, assuming $d = 840$ pc (see Table 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

---

[32] https://github.com/PaulMcMillan-Astro/GalPot

Table 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 (μJy)	$101.7^{+11.4}_{-11.1}$	$99.9^{+11.7}_{-11.4}$
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

Note. Parameters from the best-fit light-curve/radial-velocity models, with and without a surface hot spot, including MCMC errors.

Figure 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.

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 $\dot{P}_{int} = 8.5 \times 10^{-22}$ s s^{-1} for distance $d = 840$ pc.

3. Multiwavelength and Multimessenger

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 rebinned 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 (Figure 2). This fitting is done with a version of the `Icarus` code of Breton et al. (2013) modified to include the effect of hot spots on the companion surface, likely generated by precipitation of particles from the intrabinary shock (IBS) to companion magnetic poles (Sanchez & Romani 2017). All parameter values and errors are determined by Markov Chain Monte Carlo (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 et al. (2014), 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 nonthermal 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 nonthermal veiling flux. Although this is likely associated with the IBS, we
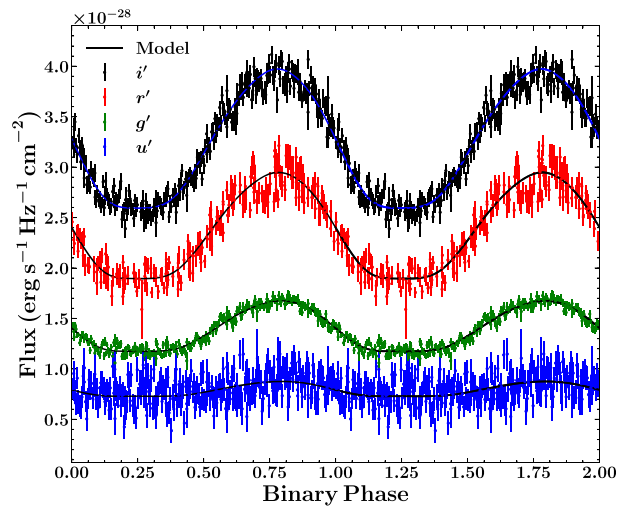
model it here as a simple power law with form $f_\nu = f_A$ $(\nu/10^{14} \, \text{Hz})^{-p}$. This flux is nearly constant through the orbit, although there are hints of phase structure, e.g., in r' and i' at $\phi_B = 0.72$ (see Figure 2). Any model without such a power-law component is completely unacceptable. These fits prefer an A_V slightly higher than, but consistent with, the maximum in this direction (obtained by ∼300 pc; Green et al. 2019).[33]

In Figure 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 greatly, with $\Delta\chi^2$/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 2. Note that with the fine structure near maximum, the model is not yet fully acceptable (χ^2/DoF ~ 1.4). More detailed models, including direct emission from the IBS or possibly the effects of companion global winds (Kandel & Romani 2020), 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 $f_\gamma = 3.5 \times 10^{-11}$ erg cm^{-2} s^{-1} between 100 MeV and 100 GeV, our fit distance gives an isotropic gamma-ray luminosity $L_\gamma = 3 \times 10^{33}$ erg s^{-1}, in good agreement with the $L_\gamma \approx (10^{33}$ erg s$^{-1}\dot{E})^{1/2}$ heuristic luminosity law (Abdo et al. 2013), 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

---

[33] https://doi.org/10.7910/DVN/2EJ9TX

5

Table 3
Summary of Radio Searches for PSR J1653−0158

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. (2013)
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. (2018)
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 (2016)
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	(2013, 2020, in preparation)
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. (2017)
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. (2013)
Parkes	1241–1497	2016 Nov 05, 06:17	3586	0.26–1.06	178	Camilo et al. (2016)

Note. 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. 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 DM = 80 pc cm^{-3} at the lowest frequencies to DM = 350 pc cm^{-3} at the highest frequencies. To estimate the limit on the pulsed component, we used Equation (6) from Ray et al. (2011) assuming a pulse width of 0.25 P, and a threshold signal-to-noise ratio S/N$_{min}$ = 7.

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 & Romani (2020). The results indicate substantial mass accretion, as expected for these ultrashort-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.

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 afterward. 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 3. Observations are spread over 11 yr, 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 μJy at 1.4 GHz, and 20 μJy at 4.9 GHz. This is fainter than the threshold of 30 μJy that Abdo et al. (2013) 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.

3.3. Continuous Gravitational Waves

We search for nearly monochromatic, continuous gravitational waves (GWs) from PSR J1653−0158, using data from the first[34] and second[35] observing runs of the Advanced LIGO detectors (The LIGO Scientific Collaboration et al. 2019). We assume that GWs are emitted at the first and 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 (Jones 2010, 2015).

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 & Schutz (2005). The second is the Bayesian time-domain method

---

[34] https://doi.org/10.7935/K57P8W9D
[35] https://doi.org/10.7935/CA75-FM95

(Dupuis & Woan 2005) as detailed by Pitkin et al. (2017), with triaxial nonaligned priors (Pitkin et al. 2015). 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 (Abbott et al. 2019a), we performed an $\mathcal{F}$-statistic search in a ± 1 Hz band around this, with an extended $\dot{f}$-range. This yields larger upper limits on h_0, with a 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 (Abbott et al. 2019b). This is $\epsilon = 3.9 \times 10^{-8} \times (h_0/5 \times 10^{-26}) \times (10^{45}\,\mathrm{g\,cm^3}/I_{zz}) \times (840\,\mathrm{pc}/d)$, 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 are via GW emission. Then assuming the canonical $I_{zz} = 10^{45}\,\mathrm{g\,cm^3}$, this would imply a "spin-down" ellipticity $\epsilon^{\mathrm{sd}} = 4.7 \times 10^{-10}$, which is a factor ~ 80 below our upper limit.

4. Discussion and Conclusions

PSR J1653−0158 is the second binary pulsar (Pletsch et al. 2012) and the fourth MSP (Clark et al. 2018) 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 $f = 508$ Hz. The 75 minute orbital period is shorter than for any other known rotation-powered pulsar, with the previous record being PSR J1311−3430 with a 93 minute orbit (Pletsch et al. 2012). 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 (Kong et al. 2014; Romani et al. 2014) 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 2 decades ago (Hartman et al. 1999).

The distance to PSR J1653−0158 and its proper motion are well constrained. Gaia measurements of the parallax, $\varpi = 1.88 \pm 1.01$ mas, imply a distance $d = 530^{+470}_{-200}$ pc. A consistent, but tighter constraint is given by our optical modeling with $d = 840^{+40}_{-40}$ pc. The proper motion (see Table 1) is also measured with good precision (Gaia and our timing are in agreement).
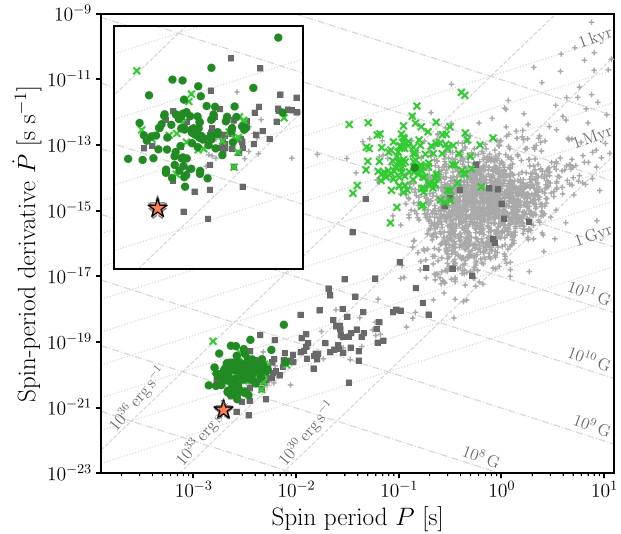
Figure 3. Newly detected PSR J1653−0158 on a P–$\dot{P}$ diagram of the known pulsar population outside of globular clusters. The MSP population is shown magnified in the inset. LAT pulsars are marked in green (isolated by a cross and binary by a circle). Non-LAT pulsars in the ATNF are marked in gray (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 derivatives according to the distance estimated from our optical models.

PSR J1653−0158 has one of the lowest observed spin-period derivatives of all known pulsars ($\dot{P} = 2.4 \times 10^{-21}\,\mathrm{s\,s^{-1}}$). The intrinsic $\dot{P} = 8.5 \times 10^{-22}\,\mathrm{s\,s^{-1}}$ (accounting for Galactic acceleration and Shklovskii effects) is even smaller. In Figure 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 1 for $d = 840$ pc. Constant lines of $\dot{E}$, B_{surf}, and τ_c are displayed in Figure 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 nondetection 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., Fruchter et al. 1988; Archibald et al. 2009; Polzin et al. 2020). 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 previous 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 (Broderick et al. 2016; Polzin et al. 2018). (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 (Clark et al. 2018).

Nieder et al.

The minimum average density of the companion 64 g cm^{-3} is very high, assuming a filled Roche lobe (Eggleton 1983). 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 (Sengar et al. 2017; Kaplan et al. 2018). In addition, simulations predict evolved ultracompact X-ray binaries to have orbital periods of around 70–80 minutes (van Haaften et al. 2012), consistent with the 75 minute 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 nondetection 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 that were precise enough to enable a successful gamma-ray pulsation search.

Appendix
Continuous Gravitational Waves

Acknowledging the possibility of mismatches between the pulsar rotation frequency and the gravitational-wave frequency,
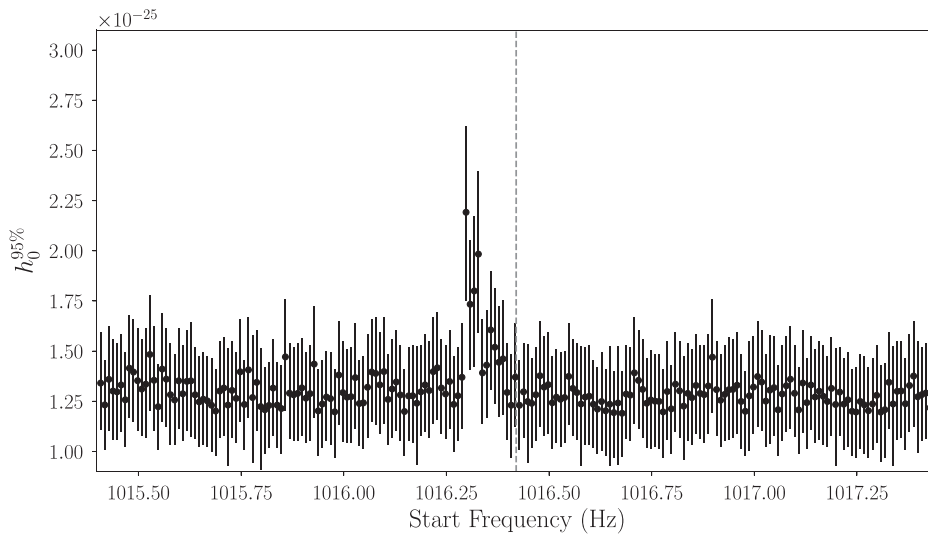
Figure 4. 95% confidence upper limits on the gravitational-wave amplitude in 10 mHz bands around twice the rotation frequency of PSR J1653−0158, which is indicated by the gray dashed line. The bars indicate a conservative estimate of the uncertainty on the upper limit values. The "spike" does not indicate a detection: it is due to a disturbance in L1 around ≈1016.32 Hz. The upper limit values are available as data behind the figure.

(The data used to create this figure are available.)

we perform an $\mathcal{F}$-statistic search in a ∼2 Hz band around twice the rotation frequency, a factor of 10^{-3} of the gravitational-wave frequency, similarly to what was done in Abbott et al. (2019a) and also extend the spin-down search to the range $2\dot{f} \in (-1.260, -1.2216) \times 10^{-15}$ Hz s^{-1}. Overall, we use 2.4×10^{9} templates resulting in an average mismatch of 1%. We examine the results in 10 mHz-wide bands. The most significant $2\mathcal{F}$ values from each band are consistent with the noise-only expectation, apart from six outliers that can be ascribed to a disturbance in L1 around ≈1016.32 Hz. We set upper limits in each band. The values are plotted in Figure 4 and are provided as data behind the figure. The mean value is 1.3×10^{-25} and it is higher than the targeted search upper limit, consistently with the larger volume of searched wave shapes.

ORCID iDs

L. Nieder ⓘ https://orcid.org/0000-0002-5775-8977
C. J. Clark ⓘ https://orcid.org/0000-0003-4355-3572
D. Kandel ⓘ https://orcid.org/0000-0002-5402-3107
R. W. Romani ⓘ https://orcid.org/0000-0001-6711-3286
C. G. Bassa ⓘ https://orcid.org/0000-0002-1429-9010
B. Allen ⓘ https://orcid.org/0000-0003-4285-6256
A. Ashok ⓘ https://orcid.org/0000-0002-8395-957X
I. Cognard ⓘ https://orcid.org/0000-0002-1775-9692
H. Fehrmann ⓘ https://orcid.org/0000-0001-8036-1882
P. Freire ⓘ https://orcid.org/0000-0003-1307-9435
R. Karuppusamy ⓘ https://orcid.org/0000-0002-5307-2919
M. Kramer ⓘ https://orcid.org/0000-0002-4175-2271
D. Li ⓘ https://orcid.org/0000-0003-3010-7661
B. Machenschalk ⓘ https://orcid.org/0000-0002-2332-0459
Z. Pan ⓘ https://orcid.org/0000-0001-7771-2864
M. A. Papa ⓘ https://orcid.org/0000-0002-1007-5298
S. M. Ransom ⓘ https://orcid.org/0000-0001-5799-9714
P. S. Ray ⓘ https://orcid.org/0000-0002-5297-5278
J. Roy ⓘ https://orcid.org/0000-0002-2892-8025

P. Wang ⓘ https://orcid.org/0000-0002-3386-7159
J. Wu ⓘ https://orcid.org/0000-0003-3536-4368
C. Aulbert ⓘ https://orcid.org/0000-0002-1481-8319
E. D. Barr ⓘ https://orcid.org/0000-0001-8715-9628
B. Beheshtipour ⓘ https://orcid.org/0000-0002-8524-1537
O. Behnke ⓘ https://orcid.org/0000-0003-0679-8562
B. Bhattacharyya ⓘ https://orcid.org/0000-0002-6287-6900
R. P. Breton ⓘ https://orcid.org/0000-0001-8522-4983
F. Camilo ⓘ https://orcid.org/0000-0002-1873-3718
C. Choquet ⓘ https://orcid.org/0000-0001-6900-1851
V. S. Dhillon ⓘ https://orcid.org/0000-0003-4236-9642
E. C. Ferrara ⓘ https://orcid.org/0000-0001-7828-7708
L. Guillemot ⓘ https://orcid.org/0000-0002-9049-8716
J. W. T. Hessels ⓘ https://orcid.org/0000-0003-2317-1446
M. Kerr ⓘ https://orcid.org/0000-0002-0893-4073
T. R. Marsh ⓘ https://orcid.org/0000-0002-2498-7589
Z. Pleunis ⓘ https://orcid.org/0000-0002-4795-697X
H. J. Pletsch ⓘ https://orcid.org/0000-0002-1164-4755
M. S. E. Roberts ⓘ https://orcid.org/0000-0002-9396-9720
B. Steltner ⓘ https://orcid.org/0000-0003-1833-5493

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, PhRvD, 99, 122002
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, ApJ, 879, 10
Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, ApJS, 183, 46
Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, ApJS, 208, 17
Abdollahi, S., Acero, F., Ackermann, M., et al. 2020, ApJS, 247, 33
Allen, B., Knispel, B., Cordes, J. M., et al. 2013, ApJ, 773, 91
Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, Sci, 324, 1411
Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Atwood, W., Albert, A., Baldini, L., et al. 2012, arXiv:1303.3514
Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, ApJ, 697, 1071
Aulbert, C., & Fehrmann, H. 2008, Forschungsbericht 2008—Max-Planck-Institut für Gravitationsphysik, Teilinstitut Hannover, https://www.mpg.de/308429/forschungsSchwerpunkt
Barr, E. D., Guillemot, L., Champion, D. J., et al. 2013, MNRAS, 429, 1633

9

Bassa, C. G., Pleunis, Z., Hessels, J. W. T., et al. 2017, ApJL, 846, L20

Bhattacharyya, B., Roy, J., Ray, P. S., et al. 2013, ApJL, 773, L12

Breton, R. P., Rappaport, S. A., van Kerkwijk, M. H., & Carter, J. A. 2012, ApJ, 748, 115

Breton, R. P., van Kerkwijk, M. H., Roberts, M. S. E., et al. 2013, ApJ, 769, 108

Broderick, J. W., Fender, R. P., Breton, R. P., et al. 2016, MNRAS, 459, 2681

Bruel, P., Burnett, T. H., Digel, S. W., et al. 2018, arXiv:1810.11394

Camilo, F., Reynolds, J. E., Ransom, S. M., et al. 2016, ApJ, 820, 6

Clark, C. J., Pletsch, H. J., Wu, J., et al. 2018, SciA, 4, eaao7228

Clark, C. J., Wu, J., Pletsch, H. J., et al. 2017, ApJ, 834, 106

Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006

Damour, T., & Taylor, J. H. 1991, ApJ, 366, 501

Desvignes, G., Cognard, I., Champion, D., et al. 2013, in IAU Symposium, Vol. 291, Neutron Stars and Pulsars: Challenges and Opportunities after 80 years, ed. J. van Leeuwen (Cambridge: Cambridge Univ. Press), 375

Dhillon, V. S., Marsh, T. R., Stevenson, M. J., et al. 2007, MNRAS, 378, 825

Dupuis, R. J., & Woan, G. 2005, PhRvD, 72, 102002

Eggleton, P. P. 1983, ApJ, 268, 368

Fehrmann, H., & Pletsch, H. J. 2014, PhRvD, 90, 124049

Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2019, OJAp, 2, 10

Fruchter, A. S., Stinebring, D. R., & Taylor, J. H. 1988, Natur, 333, 237

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, A&A, 616, A1

Green, G. M., Schlafly, E., Zucker, C., Speagle, J. S., & Finkbeiner, D. 2019, ApJ, 887, 93

Hartman, R. C., Bertsch, D. L., Bloom, S. D., et al. 1999, ApJS, 123, 79

Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, Sci, 311, 1901

Hunter, J. D. 2007, CSE, 9, 90

Jones, D. I. 2010, MNRAS, 402, 2503

Jones, D. I. 2015, MNRAS, 453, 53

Kandel, D., & Romani, R. W. 2020, ApJ, 892, 101

Kaplan, D. L., Stovall, K., van Kerkwijk, M. H., Fremling, C., & Istrate, A. G. 2018, ApJ, 864, 15

Kerr, M. 2011, ApJ, 732, 38

Kerr, M., Ray, P. S., Johnston, S., Shannon, R. M., & Camilo, F. 2015, ApJ, 814, 128

Kong, A. K. H., Jin, R., Yen, T. C., et al. 2014, ApJL, 794, L22

Li, D., Wang, P., Qian, L., et al. 2018, IMMag, 19, 112

Lightkurve Collaboration, Cardoso, J. V. d. M. a., Hedges, C., et al. 2018, Lightkurve: Kepler and TESS Time Series Analysis in Python, Astrophysics Source Code Library, ascl:1812.013

LIGO Scientific Collaboration 2018, LIGO Algorithm Library—LALSuite, free software (GPL), doi:10.7935/GT1W-FZ16

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, AJ, 129, 1993

McMillan, P. J. 2017, MNRAS, 465, 76

Nieder, L., Allen, B., Clark, C. J., & Pletsch, H. J. 2020, ApJ, 901, 156

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

Oliphant, T. E. 2006, A Guide to NumPy, Vol. 1 (USA: Trelgol Publishing)

Pitkin, M. 2018, JOSS, 3, 538

Pitkin, M., Gill, C., Jones, D. I., Woan, G., & Davies, G. S. 2015, MNRAS, 453, 4399

Pitkin, M., Isi, M., Veitch, J., & Woan, G. 2017, arXiv:1705.08978

Pletsch, H. J., & Clark, C. J. 2014, ApJ, 795, 75

Pletsch, H. J., Guillemot, L., Fehrmann, H., et al. 2012, Sci, 338, 1314

Polzin, E. J., Breton, R. P., Bhattacharyya, B., et al. 2020, MNRAS, 494, 2948

Polzin, E. J., Breton, R. P., Clarke, A. O., et al. 2018, MNRAS, 476, 1968

Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, AJ, 124, 1788

Ransom, S. M., Ray, P. S., Camilo, F., et al. 2011, ApJL, 727, L16

Ray, P. S., Kerr, M., Parent, D., et al. 2011, ApJS, 194, 17

Roberts, M. S. E. 2013, in IAU Symposium, Vol. 291, Neutron Stars and Pulsars: Challenges and Opportunities after 80 Years, ed. J. van Leeuwen (Cambridge: Cambridge Univ. Press), 127

Romani, R. W., Filippenko, A. V., & Cenko, S. B. 2014, ApJL, 793, L20

Sanchez, N., & Romani, R. W. 2017, ApJ, 845, 42

Sanpa-arsa, S. 2016, PhD Thesis, Graduate School of Arts and Sciences, University of Virginia

Saz Parkinson, P. M., Xu, H., Yu, P. L. H., et al. 2016, ApJ, 820, 8

Sengar, R., Tauris, T. M., Langer, N., & Istrate, A. G. 2017, MNRAS, 470, L6

Shklovskii, I. S. 1970, SvA, 13, 562

Strader, J., Swihart, S., Chomiuk, L., et al. 2019, ApJ, 872, 42

The LIGO Scientific Collaborationthe Virgo Collaboration, Abbott, R., et al. 2019, arXiv:1912.11716

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22

van Haaften, L. M., Nelemans, G., Voss, R., Wood, M. A., & Kuijpers, J. 2012, A&A, 537, A104

van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2

VanderPlas, J. T., & Ivezić, Ž 2015, ApJ, 812, 18

New Searches for Continuous Gravitational Waves from Seven Fast Pulsars

The results of seven targeted searches for continuous gravitational waves of known pulsars in the data of all three first observing runs of Advanced LIGO (O1, O2 and O3) are presented.

The data in the form of 60 s time-baseline SFTs for O1, O2 and O3 were prepared with the method of Chapter 2 by the author.

CrossMark

New Searches for Continuous Gravitational Waves from Seven Fast Pulsars

A. Ashok[1,2] , B. Beheshtipour[1,2] , M. A. Papa[1,2,3] , P. C. C. Freire[4] , B. Steltner[1,2] , B. Machenschalk[1,2], O. Behnke[1,2] ,
B. Allen[1,2,3] , and R. Prix[1,2]

[1] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany; anjana.ashok@aei.mpg.de, maria.alessandra.papa@aei.mpg.de
[2] Leibniz Universität Hannover, D-30167 Hannover, Germany
[3] University of Wisconsin Milwaukee, 3135 N Maryland Ave., Milwaukee, WI 53211, USA
[4] Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121 Bonn, Germany

Abstract

We conduct searches for continuous gravitational waves from seven pulsars that have not been targeted in continuous wave searches of Advanced LIGO data before. We target emission at exactly twice the rotation frequency of the pulsars and in a small band around such a frequency. The former search assumes that the gravitational-wave quadrupole is changing in a phase-locked manner with the rotation of the pulsar. The latter search over a range of frequencies allows for differential rotation between the component emitting the radio signal and the component emitting the gravitational waves, for example the crust or magnetosphere versus the core. Timing solutions derived from the Arecibo 327 MHz Drift-Scan Pulsar Survey observations are used. No evidence of a signal is found and upper limits are set on the gravitational-wave amplitude. For one of the pulsars we probe gravitational-wave intrinsic amplitudes just a factor of 3.8 higher than the spin-down limit, assuming a canonical moment of inertia of 10^{38} kg m^2. Our tightest ellipticity constraint is 1.5×10^{-8}, which is a value well within the range of what a neutron star crust could support.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Pulsars (1306); LIGO (920); Neutron stars (1108)

Supporting material: machine-readable tables

1. Introduction

Continuous gravitational waves are expected from rotating neutron stars if these objects present a deviation from a perfectly axisymmetric configuration (Jaranowski et al. 1998; Lasky 2015). On the whole, the expected signal is simple, consisting of one or two harmonics at the rotation frequency of the star and at twice this frequency (Jones 2015).

The sensitivity of the LIGO instruments allows for the probing of continuous gravitational-wave emission from the Galactic population of neutron stars, for deformations of a few parts in a million and smaller, depending on the search, over a broad range of frequencies. Different types of searches are carried out: "blind" all-sky surveys (Abbott et al. 2021a; Steltner et al. 2021b; Covas & Sintes 2020; Dergachev & Papa 2021, 2020; Abbott et al. 2019a), searches directed at neutron star candidates like supernova remnants and low mass X-ray binaries (Zhang et al. 2021; Abbott et al. 2021b; Papa et al. 2020; Lindblom & Owen 2020; Jones & Sun 2021; Ming et al. 2019), and targeted searches aimed at known pulsars (Abbott et al. 2019b, 2019c; Nieder et al. 2020, 2019; Fesik & Papa 2020; Abbott et al. 2021c, 2021d).

Among the different searches, the ones that target pulsars have a special place. Pulsars are believed to be neutron stars, the distance is usually known and the rotation frequency and its derivatives are also known. This has important consequences: a null measurement is directly informative on the gravitational-wave emission—there is no question about whether a source is there in the first place. The search is simple because whatever the emission mechanism is, the gravitational frequency depends on the spin frequency, which is known. A detection would therefore immediately encode information on what is generating the gravitational waves. Because there is little to no uncertainty on

the gravitational waveform from a known pulsar, the number of templates that are searched is many orders of magnitude smaller than those investigated in surveys: the O2 data all-sky search of Steltner et al. (2021b) probed $\approx 10^{17}$ more templates than a targeted search. Fewer probed waveforms make targeted searches the most sensitive: the smallest detectable signal is a few times smaller than what the most sensitive broad survey could detect at the same frequency.

In this paper we present results from searches for emission from seven new pulsars using public data from all three Advanced LIGO observing runs (O1, O2 and O3; Abbott et al. 2021e and LIGO 2019a, 2019b, 2019c).

In this paper we introduce the signal model in Section 2, we detail the targeted objects in Section 3, the gravitational-wave searches are described in Section 4, and the results are presented and discussed in Section 5.

2. The Signal

The search described in this paper targets nearly monochromatic gravitational-wave signals of the form described, for example, in Section II of Jaranowski et al. (1998). In the calibrated strain data from a gravitational-wave detector the signal has the form

$$h(t) = F_+(\alpha, \delta, \psi; t)h_+(t) + F_\times(\alpha, \delta, \psi; t)h_\times(t), \quad (1)$$

with the "+" and "×" indicating the two gravitational-wave polarizations. $F_+(\alpha, \delta, \psi; t)$ and $F_\times(\alpha, \delta, \psi; t)$ are the detector sensitivity pattern functions, which depend on relative orientation between the detector and the source, and hence on time t, on the position (α, δ) of the source, and on ψ, the polarization

1

Ashok et al.

angle. The waveforms $h_+(t)$ and $h_\times(t)$ are

$$h_+(t) = A_+ \cos \Phi(t)$$
$$h_\times(t) = A_\times \sin \Phi(t), \tag{2}$$

with

$$A_+ = \frac{1}{2}h_0(1 + \cos^2 \iota)$$
$$A_\times = h_0 \cos \iota. \tag{3}$$

The angle between the total angular momentum of the star and the line of sight is $0 \leqslant \iota \leqslant \pi$ and $h_0 \geqslant 0$ is the intrinsic gravitational-wave amplitude. $\Phi(t)$ of Equation (2) is the phase of the gravitational-wave signal at time t. If τ_{SSB} is the arrival time of the wave with phase $\Phi(t)$ at the solar system barycenter, then $\Phi(t) = \Phi(\tau_{\mathrm{SSB}}(t))$. The gravitational-wave phase as a function of τ_{SSB} is assumed to be

$$\Phi(\tau_{\mathrm{SSB}}) = \Phi_0 + 2\pi[f(\tau_{\mathrm{SSB}} - \tau_{0\mathrm{SSB}})$$
$$+ \frac{1}{2}\dot{f}(\tau_{\mathrm{SSB}} - \tau_{0\mathrm{SSB}})^2]. \tag{4}$$

We take $\tau_{0\mathrm{SSB}}$ consistently with the timing solution, and hence differently for every pulsar, as shown in Table 3.

3. The Pulsars

We target continuous gravitational-wave emission from seven recycled pulsars discovered and/or timed with data from the Arecibo 327 MHz Drift-Scan Pulsar Survey (AO327; Martinez et al. 2017, 2019): PSR J2204+2700, PSR J1411+2551, PSR J0709+0458, PSR J0824+0028, PSR J0732+2314, PSR J0509 +0856, and PSR J0154+1833. For practicality we mostly use abbreviated forms of the names of the pulsars, omitting the "PSR" prefix and the part after the "+".

These pulsars have never been searched for before for gravitational-wave emission. They represent a relatively nearby sample, with distances smaller than 2 kpc, which is typical of all-sky surveys. This makes them particularly interesting for gravitational-wave searches, the only exception being 2204 +2700, which is more distant, and also having an extremely low spin-down.

Our targets are all in binary systems except for J0154+1833, which is an isolated millisecond pulsar. Our set includes the radio pulsar in the notable double-neutron-star system PSR J1411+2551.

When available, we take the orbital inclination angle as estimate of the inclination angle ι for the determination of the constrained prior upper limit, see Section 5.1. We take the following values: $\iota_{\mathrm{J1411}} = 0.83$ rad, $\iota_{\mathrm{J0709}} = 1.30$ rad, $\iota_{\mathrm{J0824}} = 1.32$ rad, $\iota_{\mathrm{J0732}} = 0.93$ rad. For J0154, J0509, and J2204 we do not have an estimate of the inclination angle.

4. The Gravitational-wave Searches

We use LIGO public data from the Hanford (H1) and the Livingston (L1) detectors from the O1, O2, and the recently released first six months of the O3 science run (LIGO 2019a, 2019b, 2019c). The data is gated to remove loud glitches (Steltner et al. 2021a) and contiguous segments are Fourier transformed to produce the input to the search. After having excluded egregiously noisy segments in the band of each pulsar, we have ≈175 days of data from each detector
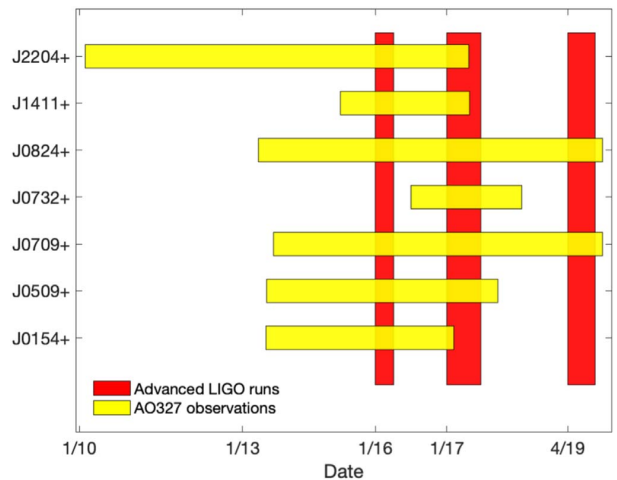
Figure 1. Time intervals corresponding to the O1, O2, and O3 LIGO runs are shown in red as vertical rectangles and the radio observation periods for each pulsar are shown in yellow as horizontal rectangles.

from the O1 and O2 runs combined, and ≈125 during the O3 run for H1 and ≈129 days for L1.

No glitch was recorded by the AO327 in any of the pulsars' spins. As Figure 1 shows, these observations do not perfectly cover all the Advanced LIGO runs so we cannot exclude the possibility of a glitch. Even though our targets are very stable pulsars and a glitch is unlikely, we perform different searches and coherently combine the O1 and O2 data, the O3 data, and also all the data that we have, O1O2O3. We use the matched-filter detection statistic—$\mathcal{F}$-statistic (Cutler & Schutz 2005)— as our detection statistic. The $\mathcal{F}$-statistic is the maximum log-likelihood ratio of the signal hypothesis to the Gaussian-noise hypothesis. The signal is described by a frequency, spin-down, sky position, and orbital parameter values, which define the template waveform and are explicitly searched over. The signal amplitude parameters $\cos \iota$, ψ, Φ_0, and h_0 are analytically maximized over.

In Gaussian noise the $2\mathcal{F}$-statistic follows a noncentral chi-squared distribution with 4 degrees of freedom, $\chi^2_4(2\mathcal{F}, \rho^2)$. The noncentrality parameter ρ^2 is the expected squared signal-to-noise ratio and it is proportional to $h_0^2 T_{\mathrm{data}}/S_h$, where T_{data} is the duration of time for which data is available and S_h is the strain power spectral density of the noise (Jaranowski et al. 1998).

For every pulsar and data set we conduct two searches: one with a single template with the gravitational-wave frequency f and spin-down $\dot{f}$ being twice the spin frequency ν and spin-down $\dot{\nu}$, and one for a range of frequencies and spin-downs around these. The parameters of the targeted searches are given in Table 3 in the Appendix.

The search at $f = 2\nu$ is appropriate if the gravitational-wave frequency is exactly locked with the observed spin frequency. Mechanisms exist, however, that could produce a small difference between the gravitational-wave frequency and twice the spin frequency: a misalignment of the rotation axis with the symmetry axis of the star, causing free precession; or the component responsible for the gravitational-wave emission— for example a solid core—not spinning as the radio-emitting component. For such cases, it has been found that $f = 2\nu(1 \pm \delta_f)$ with $\delta_f \lesssim 10^{-4}$ (Jones & Andersson 2002; Abbott et al. 2008). With this in mind, we conservatively perform

2

Ashok et al.

searches over a band $\pm 2\nu \times 2 \times 10^{-3}$ of $f = 2\nu$, and consistently for $\dot{f}$.

For the band searches we set up a template grid in frequency and spin-down with spacings of 2.6×10^{-9} Hz and 1.9×10^{-17} Hz s^{-1}, respectively. These grids yield a maximum mismatch smaller than 1% for the O1O2 and O3 searches, and smaller than 8% for the O1O2O3 searches.

We also conduct the single-template searches using a Bayesian approach. We demodulate the data according to the expected signal, we heterodyne/downsample the data, and then search over the waveform amplitude parameters with a nested-sampling algorithm. The method is exactly the same as used by Abbott et al. (2019d), with the same uniform angular priors, namely $\Phi_0 \in [0, \pi]$, $\psi \in [0, \frac{\pi}{2}]$, $\cos \iota \in [-1, 1]$. For the intrinsic amplitude we adopt the same broad uniform prior for all sources with $h_0 \in [10^{-27}, 10^{-24}]$. The data used for this search is not gated. We report the results for the combined O1O2O3 data.

5. Results

In order to evaluate the significance of the search results we compute p-values. We do this because, based on insignificant p-values, we can exclude the presence of signals that we can confidently detect. We note, however, that a very low p-value in general is not enough to claim a confident detection.

The p-value associated with the realization $2\mathcal{F}'$ of a random variable is defined as

$$p(2\mathcal{F}') = \int_{2\mathcal{F}'}^{\infty} p_0(2\mathcal{F}) \, d2\mathcal{F}, \qquad (5)$$

where p_0 is the distribution of $2\mathcal{F}$ in the presence of noise only.

If our data were Gaussian and our search pipelines were completely perfect implementations of the $\mathcal{F}$-statistic, $p_0 = \chi_4^2(2\mathcal{F}, 0)$. In reality the distribution of our search results may differ slightly from $\chi_4^2(2\mathcal{F}, 0)$, and for targeted searches we evaluate it on the actual data by running searches for fiducial sources at frequencies close to the target frequency.

None of the targeted searches yield a detection. Figure 2 shows the p-values for the targeted searches (blue circles): all the results for the targeted searches are consistent with the noise-only hypothesis. The most significant targeted-search result comes from PSR J0709 from the O3 data search, with a p-value of $\approx 23\%$. The product of the $\approx 55\%$ p-value of the O1O2 result and the O3 result is $\approx 12\%$, however the coherent O1O2O3 data search yields a totally insignificant p-value of $\approx 83\%$. The Bayesian posteriors of Figure 3 are consistent with the $\mathcal{F}$-stat results, with the only slightly off-zero posterior found for PSR J0709. Such a posterior is very broad, includes zero, and may happen just due to noise fluctuations. We also note that the target frequency for PSR J0709 is at ≈ 58 Hz, which is a highly contaminated region.

To evaluate the results from the band searches, for every pulsar we consider the most significant result in subbands that are 10 mHz wide, and compute the Gaussian-noise p-value associated with it. We do this by searching 10 mHz thousands of times, each time with a different Gaussian-noise realization. From each search we find the value of the loudest, $2\mathcal{F}_\ell$, and from the ensemble we estimate the $p_0(2\mathcal{F}_\ell)$, which we use to compute the p-value from Equation (5). Since we use Gaussian noise, this is a Gaussian-noise p-value and it is conservative (in the sense that we would not accidentally discard a potential signal) because in general it will overestimate the significance

of a result with respect to the p-value evaluated on real data. We do not use real data because each Monte Carlo realization covers 10 mHz; with thousands of independent realizations we would be considering target frequencies several hertz away from the original pulsar frequency, and at these distances there is no assurance that the noise is representative of the noise contributing to the original result. These p-values are the red circles shown in the left-hand-side plots in the first three rows of Figure 2.

We want to evaluate whether for any pulsar the band search yields a very significant result, so we consider the most significant (lowest) p-value found for each pulsar. One would then want to compare these lowest p-values but this cannot be done directly, because they do not come from the same distribution. In fact, simply due to the trials factor, the larger the band that has been searched, the lower the expected lowest p-value for that pulsar. To normalize the results and allow a direct comparison we estimate the trials factor to be equal to the number of independent 10 mHz subbands N^α searched for pulsar α, and introduce the following measure of significance for the most significant result for pulsar α:

$$s^\alpha = N^\alpha \min_{i \in [1, N^\alpha]} \{p_i^\alpha\}, \qquad (6)$$

where p_i^α is the p-value associated with the ith subband: $p(2\mathcal{F}_i^{\star \alpha})$. These quantities are shown in the right-hand-side plots of Figure 2.

A value of $s^\alpha < 1$ indicates that we would have to repeat the α-pulsar band search in random noise $1/s^\alpha$ times before we can expect a result as significant as the observed one. In this case s^α can be interpreted as a p-value. If $s^\alpha > 1$, it means that in a band search like the one conducted, in random noise, we expect that in s^α subbands the loudest values will be at least as significant as the observed result. For a detection we would need $s^\alpha \lll 1$.

The most significant result from the band searches comes again from PSR J0709 in the O1O2O3 coherent search and is at the level of $\approx 9\%$. However this significance is not confirmed in the O3 data where the lowest p-value (s^{J0709}) is 22% and it is at a different subband than the one that produced the O1O2O3 most significant result.

Figures 4 show the distributions of the most significant 10 mHz p-values and illustrate that they are consistent with the Gaussian-noise expectations for searches on all data.

5.1. Upper Limits

Based on the O1O2, O3, and O1O2O3 targeted-search results we place 95% confidence upper limits on the intrinsic gravitational-wave amplitude at the detector h_0 defined in Equation (3). We use a series of Monte Carlos where we simulate signals at a fixed amplitude in real data and measure the detection efficiency of our search. The detection criterion is that the obtained value of the detection statistic be equal or greater than the one found in the real search: if the measured detection statistic is high, a higher gravitational-wave amplitude will be needed in order for the signals to be detected. The amplitude for which 95% of the tested signals is detected is the upper limit value, $h_0^{95\%}$. With minor variations on the theme, this is the standard approach that we have taken for $\mathcal{F}$-statistic searches since the very first continuous waves search on LIGO data back in Abbott et al. (2004). The $\mathcal{F}$-stat upper limits are shown in Table 1.

3

Ashok et al.
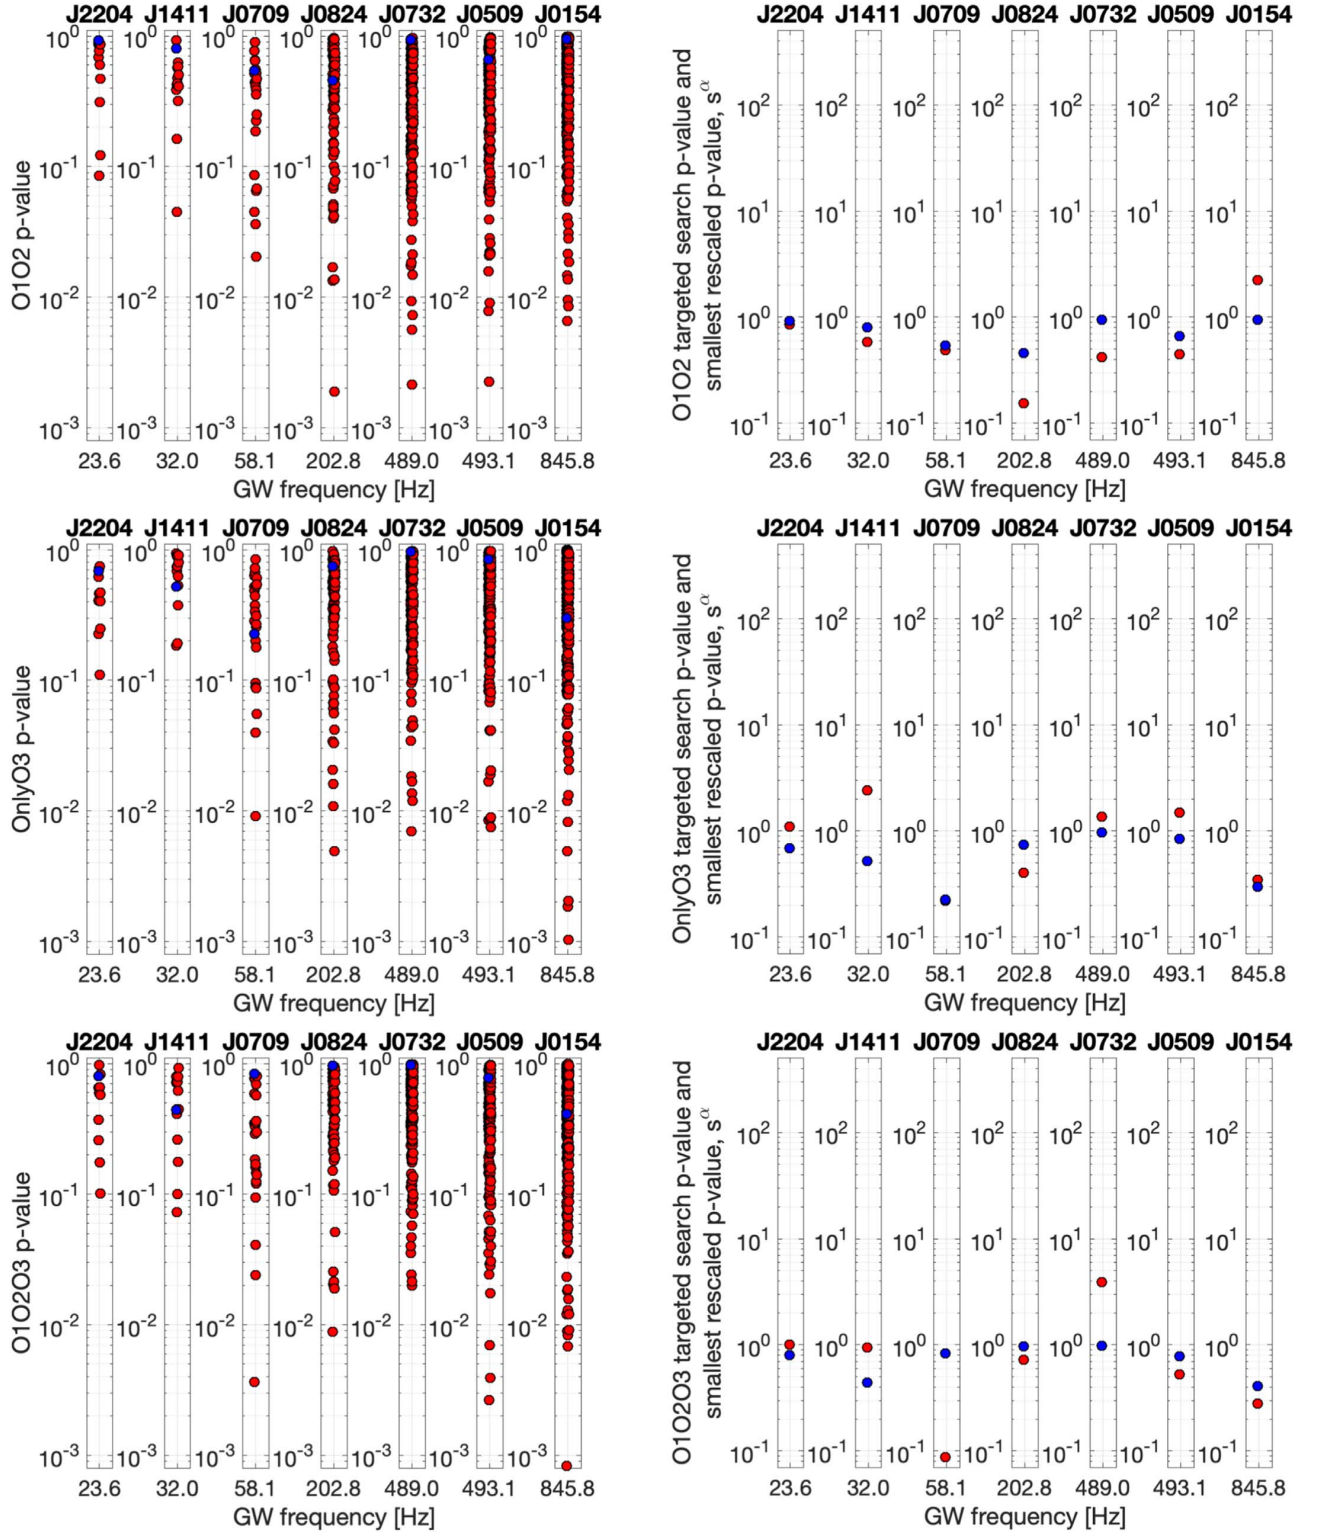
Figure 2. O1O2 (first row), O3 (second row), and O1O2O3 (third row) results. The blue circles show the *p*-values of the targeted searches. The red circles in the left-hand plots show the *p*-value of the most significant result in each 10 mHz subband of the band searches. These are generally higher than the targeted-search *p*-values because they are maxima over 10 mHz, whereas the targeted searches probe only a single waveform. The red circles in the right-hand plots show, for each pulsar, the lowest *p*-value among the subbands rescaled according to Equation (6), s^α. When $s^\alpha \leqslant 1$ it can be directly interpreted as a *p*-value. When $s^\alpha > 1$ it represents the number of 10 mHz subbands in which we would expect, in a band search of Gaussian-noise data such as that performed for pulsar α, to measure a result more significant than the most significant found in real data.

4
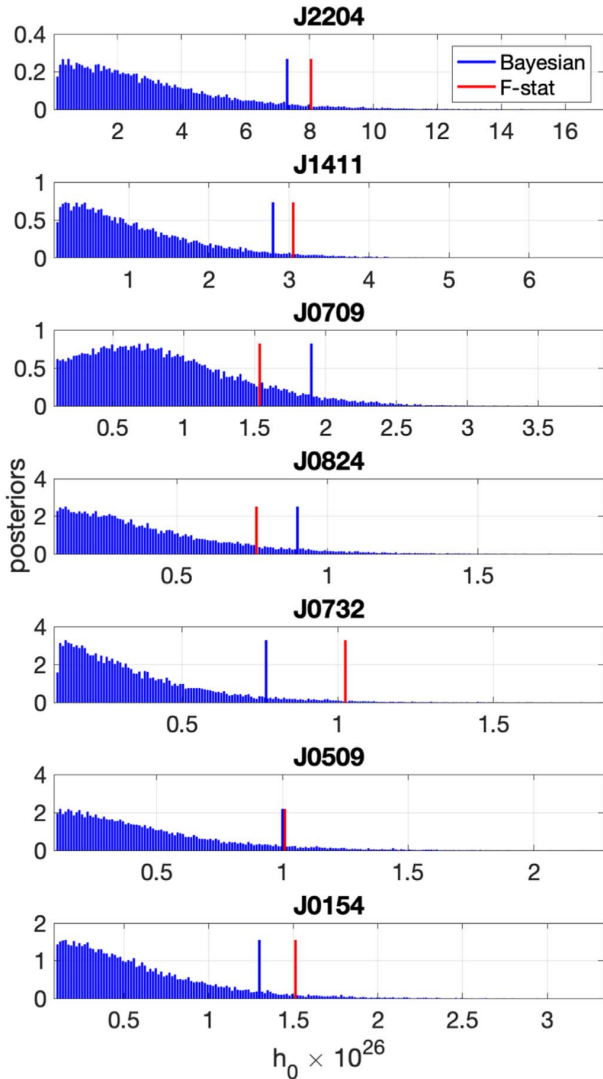
Figure 3. Bayesian posteriors for the combined O1O2O3 searches (blue) and associated 95% confidence upper limits. We also show (red) the $\mathcal{F}$-stat upper limits.

Figure 4. O1O2, O3, and O1O2O3 band-search results. For each 10 mHz frequency band searched, we show the cumulative distribution of the Gaussian p-value of the most significant result. If the data were Gaussian noise, the distribution would follow the dashed black line.

As discussed in the previous section the h_0 posteriors from the Bayesian analysis shown in Figure 3 are consistent with a null result. Because of this, the Bayesian upper limits are readily derived by integrating the posteriors up to a value such that the overall probability is the desired confidence level. Such value is the upper limit, and it represents the smallest extremum of the credible interval. The Bayesian upper limit values are shown in Table 2.

For the O1O2 band searches we divide the searched frequency range into 10 mHz subbands and take the most significant detection statistic value in that subband for our detection criteria. The subband searches probe numerous waveforms and so the loudest detection statistic value is going to be higher than for the targeted searches. Correspondingly the upper limits will also be higher, as shown in Figure 5 in the Appendix, typically by a factor of ≈ 2.7. Since this is the most computationally intense part of this work, and we do not find evidence for a signal, we do not set upper limits based on the O3 data, or on the O1O2O3 data, but we

expect that these would also be higher than the corresponding targeted ones by a factor of a few.

The populations of fake signals used to determine the detection efficiencies have polarization angle ψ and initial

Table 1
95% Confidence Upper Limits on the Gravitational-wave Amplitude for the Targeted Searches Based on No Assumptions on the Inclination Angle (Unconstrained Prior) and if the Inclination Angle is the Same as the Estimated Value of the Orbital Inclination Angle (Constrained Prior), Using Different Data

Pulsar	f (Hz)	$h_0^{95\%}$ Unconstrained Prior	$h_0^{95\%}$ Constrained Prior	$h_0^{\rm spdwn}$	$\epsilon^{95\%}$ Unconstrained Prior	$\epsilon^{95\%}$ Constrained Prior	$h_0^{95}/h_0^{\rm spdwn}$ Unconstrained Prior	$h_0^{95}/h_0^{\rm spdwn}$ Constrained Prior
O1 O2								
J0154	≈ 845.8	$2.7^{+0.5}_{-0.5} \times 10^{-26}$	...	9.0×10^{-28}	3.1×10^{-08}	...	30.3	...
J0509	≈ 493.1	$1.9^{+0.5}_{-0.4} \times 10^{-26}$	...	5.2×10^{-28}	1.1×10^{-07}	...	36.3	...
J0709	≈ 58.1	$3.3^{+0.6}_{-0.6} \times 10^{-26}$	$3.9^{+0.6}_{-0.7} \times 10^{-26}$	1.5×10^{-27}	1.6×10^{-05}	2.0×10^{-05}	21.9	26.5
J0732	≈ 489.0	$1.9^{+0.5}_{-0.4} \times 10^{-26}$	$1.4^{+0.4}_{-0.4} \times 10^{-26}$	5.8×10^{-28}	1.2×10^{-07}	9.2×10^{-08}	32.0	24.0
J0824	≈ 202.8	$1.4^{+0.3}_{-0.3} \times 10^{-26}$	$1.8^{+0.4}_{-0.4} \times 10^{-26}$	2.0×10^{-27}	4.9×10^{-07}	6.2×10^{-07}	6.9	8.8
J1411	≈ 32.0	$5.2^{+1.2}_{-1.4} \times 10^{-26}$	$3.6^{+0.9}_{-0.9} \times 10^{-26}$	1.1×10^{-27}	4.7×10^{-05}	3.2×10^{-05}	47.7	32.9
J2204	≈ 23.6	$1.5^{+0.4}_{-0.3} \times 10^{-25}$	...	4.8×10^{-28}	5.6×10^{-04}	...	319.2	...
O3								
J0154	≈ 845.8	$1.9^{+0.4}_{-0.3} \times 10^{-26}$	...	9.0×10^{-28}	2.2×10^{-08}	...	21.2	...
J0509	≈ 493.1	$1.2^{+0.2}_{-0.3} \times 10^{-26}$	...	5.2×10^{-28}	6.9×10^{-08}	...	23.7	...
J0709	≈ 58.1	$2.2^{+0.3}_{-0.4} \times 10^{-26}$	$2.4^{+0.3}_{-0.1} \times 10^{-26}$	1.5×10^{-27}	1.1×10^{-05}	1.2×10^{-05}	14.6	16.1
J0732	≈ 489.0	$1.2^{+0.3}_{-0.3} \times 10^{-26}$	$10.0^{+2.5}_{-2.5} \times 10^{-27}$	5.8×10^{-28}	7.9×10^{-08}	6.6×10^{-08}	20.6	17.1
J0824	≈ 202.8	$9.9^{+1.5}_{-1.4} \times 10^{-27}$	$1.2^{+0.2}_{-0.3} \times 10^{-26}$	2.0×10^{-27}	3.5×10^{-07}	4.2×10^{-07}	5.0	5.9
J1411	≈ 32.0	$3.9^{+1.1}_{-1.1} \times 10^{-26}$	$2.7^{+0.7}_{-0.7} \times 10^{-26}$	1.1×10^{-27}	3.5×10^{-05}	2.5×10^{-05}	35.7	25.1
J2204	≈ 23.6	$9.9^{+2.6}_{-2.4} \times 10^{-26}$	...	4.8×10^{-28}	3.6×10^{-04}	...	204.0	...
O1 O2 O3								
J0154	≈ 845.8	$1.5^{+0.3}_{-0.3} \times 10^{-26}$	...	9.0×10^{-28}	1.7×10^{-08}	...	16.7	...
J0509	≈ 493.1	$1.0^{+0.2}_{-0.3} \times 10^{-26}$	...	5.2×10^{-28}	5.7×10^{-08}	...	19.6	...
J0709	≈ 58.1	$1.5^{+0.3}_{-0.3} \times 10^{-26}$	$1.9^{+0.3}_{-0.4} \times 10^{-26}$	1.5×10^{-27}	7.7×10^{-06}	9.5×10^{-06}	10.3	12.7
J0732	≈ 489.0	$1.0^{+0.3}_{-0.3} \times 10^{-26}$	$7.8^{+2.5}_{-2.0} \times 10^{-27}$	5.8×10^{-28}	6.7×10^{-08}	5.1×10^{-08}	17.6	13.3
J0824	≈ 202.8	$7.6^{+1.6}_{-1.9} \times 10^{-27}$	$1.2^{+0.3}_{-0.2} \times 10^{-26}$	2.0×10^{-27}	2.7×10^{-07}	4.1×10^{-07}	3.8	5.8
J1411	≈ 32.0	$3.1^{+0.7}_{-0.7} \times 10^{-26}$	$2.2^{+0.4}_{-0.3} \times 10^{-26}$	1.1×10^{-27}	2.8×10^{-05}	2.0×10^{-05}	28.0	19.9
J2204	≈ 23.6	$8.1^{+2.0}_{-2.1} \times 10^{-26}$	...	4.8×10^{-28}	2.9×10^{-04}	...	166.4	...

Note. We also show the spin-down upper limit calculated for a nominal value of the moment of inertia of 10^{38} kg m^2, and the upper limits on the ellipticity of the star. The last two columns indicate how far our results are from being physically interesting: if $h_0^{95\%}/h_0^{\rm spdwn}$ is less than 1, then the upper limits are informative.

(This table is available in machine-readable form.)

Table 2
Bayesian Upper Limits of O1–O2–O3 Targeted Searches (Unconstrained cos ι Priors)

O1 O2 O3 Bayesian Pulsar	$h_0^{95\%}$	$\epsilon^{95\%}$	$h_0^{95}/h_0^{\rm spdwn}$
J0154	1.3×10^{-26}	1.5×10^{-08}	14.9
J0509	1.0×10^{-26}	5.9×10^{-08}	20.1
J0709	1.9×10^{-26}	9.4×10^{-08}	12.5
J0732	7.7×10^{-27}	5.1×10^{-08}	13.2
J0824	9.0×10^{-27}	3.2×10^{-07}	4.5
J1411	2.8×10^{-26}	2.5×10^{-05}	25.9
J2204	7.3×10^{-26}	2.7×10^{-04}	151.4

(This table is available in machine-readable form.)

phase Φ_0 uniformly distributed as described in Section 4. For the orientation angle we consider two cases: cos ι uniformly distributed in $[-1,1]$ and fixed at the value of the orbital inclination, when available from the radio observations. We refer to the resulting upper limits as *unconstrained* and *constrained*, respectively.

If we assume that the neutron star is a triaxial ellipsoid spinning around a principal moment of inertia axis I_{zz}, and that the continuous wave emission is due to an ellipticity

$$\varepsilon = \frac{I_{xx} - I_{yy}}{I_{zz}}, \qquad (7)$$

based on the intrinsic gravitational-wave amplitude upper limits $h_0^{95\%}$, we can exclude neutron star deformations above a $\varepsilon^{95\%}$ level. The ellipticity needed for a neutron star at a distance D, spinning at $f/2$, to produce continuous gravitational waves with an intrinsic amplitude on Earth of h_0 is (Jaranowski et al. 1998; Gao et al. 2020):

$$\begin{aligned}\varepsilon =\; & 2.4 \times 10^{-7} \left(\frac{h_0}{1 \times 10^{-26}} \right) \\ & \times \left(\frac{D}{1\ \text{kpc}} \right) \left(\frac{200\ \text{Hz}}{f} \right)^2 \left(\frac{10^{38}\ \text{kg m}^2}{I_{zz}} \right).\end{aligned} \qquad (8)$$

The ellipticity $\varepsilon^{95\%}$ upper limits are given in Table 1.

Ashok et al.

5.2. Discussion

We have searched for continuous gravitational waves from seven pulsars that have not been targeted before. We use all the publicly available Advanced LIGO data, namely from the O1, O2, and O3 science runs.

We find no evidence of a gravitational-wave signal at a detectable level. The posterior probability distribution for PSR J0709 is peaked slightly off-zero, but this could well be a noise fluctuation as well as due to spectral contamination. At the lower frequencies in particular, it is not uncommon to find these posteriors, see for example Figure 3 of Abbott et al. (2020) showing the results for the Vela Pulsar from the search at ≈ 22 Hz. The $\mathcal{F}$-statistic results for the coherent O1O2O3 search are insignificant, which indicates that a coherent signal during all the observations is not detected. On the other hand PSR J0709 is one of the only two pulsars for which the radio observations overlap with all the LIGO runs, so we are most confident of the used template waveform.

For more than half of the pulsar sample, our searches probe ellipticities $\lesssim 3 \times 10^{-7}$, which could be sustained by neutron star crusts (Johnson-McDaniel & Owen 2013; Bhattacharyya 2020; Gittins et al. 2021). Our tightest ellipticity bound amounts to 1.7×10^{-8} (1.5×10^{-8} from the Bayesian analysis), for PSR J0154. The remaining four pulsars are more distant and/or spin slower, this yields less-constraining upper ellipticity limits. For the pulsar PSR J0824, assuming a canonical moment of inertia of 10^{38} kg m^2, our upper limits are within a factor of 3.8 (5.8) of the spin-down upper limit, for an unconstrained and constrained $\cos \iota$ prior, respectively. The actual moment of inertia of the star may differ from the canonical one up by a factor of a few. These are physically interesting ellipticity ranges (Woan et al. 2018), and showcase the potential for this type of search.

Appendix A
Targeted Search Parameters

The parameters of the gravitational-wave templates for the pulsars in this search are provided in Table 3.

Table 3
Gravitational Waveform Parameters

Pulsar	GW Frequency f (Hz)	GW Freq. Derivative $\dot{f}$ (Hz s^{-1})	R.A. (rad)	Decl. (rad)	Epoch τ_{0SSB} (MJD)	Proj. Semimajor Axis $a \sin i$ (lt-s)	Binary Period (s)	Eccentricity	Arg. of Periastron ω (rad)	Distance (pc)
J2204	≈23.6	-3.652×10^{-17}	5.7802112	0.4715040	56,805.0	210.680632593	70,437,206.11210689	0.00152	0.1118	2150
J1411	≈32.0	-4.904×10^{-17}	3.7145600	0.4512083	57,257.864168	9.204790917	226,010.02575114663	0.16993	1.4209	977
J0709	≈58.1	-6.418×10^{-16}	1.8724740	0.0869343	56,983.893691	15.716582025	377,281.1771422851	0.00023	5.6319	1790
J0824	≈202.8	-3.021×10^{-15}	2.2009213	0.0081476	56,600.0	18.988928488	2,005,081.527370442	0.00023	0.8084	1530
J0732	≈489.0	-7.344×10^{-16}	1.9749503	0.4057610	58,000.0	10.625842295	2,611,878.6842582175	0.00001	1.1879	1660
J0509	≈493.1	-5.366×10^{-16}	1.3498838	0.1560374	57,384.0	2.458025534	424,049.2036136159	0.00002	0.5642	1450
J0154	≈845.8	-1.046×10^{-15}	0.5001010	0.3240047	56,900.0	...	...	...	...	860

8

Ashok et al.

Appendix B
Band Search Upper Limit Plots

In this Appendix, in Figure 5, we show the O1O2 band search upper limit results and for reference we also plot the targeted-search upper limits.
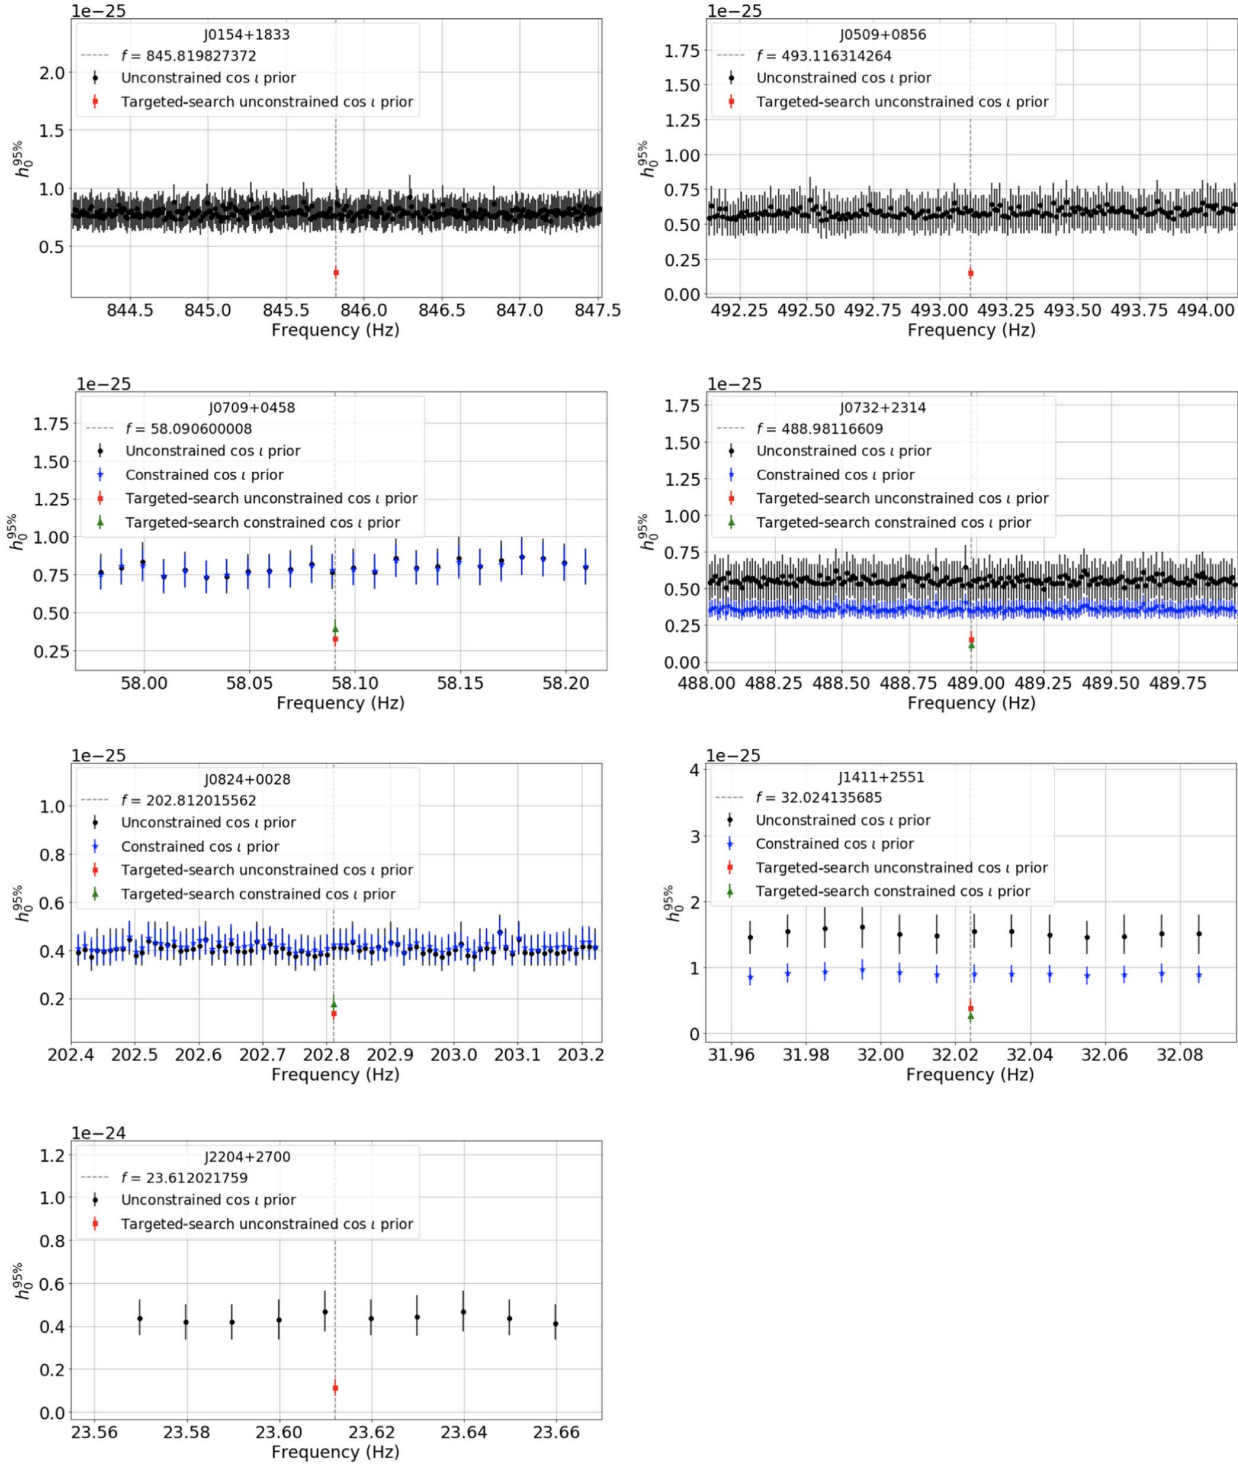
Figure 5. O1O2 data upper limits on the gravitational-wave amplitude in each 10 mHz frequency subband searched, based on the most significant result in that 10 mHz subband. We also show the targeted-search results, which are the lower points at the central frequency, which is twice the pulsar rotation frequency.

9

ORCID iDs

A. Ashok ⬤ https://orcid.org/0000-0002-8395-957X
B. Beheshtipour ⬤ https://orcid.org/0000-0002-8524-1537
M. A. Papa ⬤ https://orcid.org/0000-0002-1007-5298
P. C. C. Freire ⬤ https://orcid.org/0000-0003-1307-9435
B. Steltner ⬤ https://orcid.org/0000-0003-1833-5493
O. Behnke ⬤ https://orcid.org/0000-0003-0679-8562
B. Allen ⬤ https://orcid.org/0000-0003-4285-6256
R. Prix ⬤ https://orcid.org/0000-0002-3789-6424

References

Abbott, B., Abbott, R., Adhikari, R., et al. 2004, PhRvD, 69, 082004
Abbott, B., Abbott, R., Adhikari, R., et al. 2008, ApJL, 683, L45
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, PhRvD, 100, 024004
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, ApJ, 879, 10
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019c, PhRvD, 99, 122002
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019d, ApJ, 879, 10
Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, ApJL, 902, L21
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, PhRvD, 104, 082004
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, ApJ, 921, 80
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021c, ApJ, 913, L27
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021d, arXiv:2104.14417
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021e, SoftX, 13, 100658
Bhattacharyya, S. 2020, MNRAS, 498, 728
Covas, P. B., & Sintes, A. M. 2020, PhRvL, 124, 191102
Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006
Dergachev, V., & Papa, M. A. 2020, PhRvL, 125, 171101
Dergachev, V., & Papa, M. A. 2021, PhRvD, 103, 063019
Fesik, L., & Papa, M. A. 2020, ApJ, 895, 11
Gao, Y., Shao, L., Xu, R., et al. 2020, MNRAS, 498, 1826
Gittins, F., Andersson, N., & Jones, D. 2021, MNRAS, 500, 5570
Jaranowski, P., Krolak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001
Johnson-McDaniel, N. K., & Owen, B. J. 2013, PhRvD, 88, 044004
Jones, D., & Sun, L. 2021, PhRvD, 103, 023020
Jones, D. I. 2015, MNRAS, 453, 53
Jones, D. I., & Andersson, N. 2002, MNRAS, 331, 203
Lasky, P. D. 2015, PASA, 32, e034
LIGO 2019a, The O1 Data Release, https://www.gw-openscience.org/O1/
LIGO 2019b, The O2 Data Release, https://www.gw-openscience.org/O2/
LIGO 2019c, The O3 Data Release, https://www.gw-openscience.org/O3/O3a/
Lindblom, L., & Owen, B. J. 2020, PhRvD, 101, 083023
Martinez, J. G., Gentile, P., Freire, P. C. C., et al. 2019, ApJ, 881, 166
Martinez, J. G., Stovall, K., Freire, P. C. C., et al. 2017, ApJL, 851, L29
Ming, J., Papa, M. A., Singh, A., et al. 2019, PhRvD, 100, 024063
Nieder, L., Clark, C. J., Bassa, C. G., et al. 2019, ApJ, 883, 42
Nieder, L., Clark, C. J., Kandel, D., et al. 2020, ApJL, 902, L46
Papa, M. A., Ming, J., Gotthelf, E. V., et al. 2020, ApJ, 897, 22
Steltner, B., Papa, M. A., & Eggenstein, H.-B. 2021a, arXiv:2105.09933
Steltner, B., Papa, M. A., Eggenstein, H. B., et al. 2021b, ApJ, 909, 79
Woan, G., Pitkin, M., Haskell, B., Jones, D., & Lasky, P. 2018, ApJL, 863, L40
Zhang, Y., Papa, M. A., Krishnan, B., & Watts, A. L. 2021, ApJL, 906, L14

Results From an Einstein@Home Search for Continuous Gravitational Waves From G347.3 at Low Frequencies in LIGO O2 Data

The results of a directed Einstein@Home search for continuous gravitational waves from the supernova remnant G347.3-0.5 in Advanced LIGO data of the second observation run (O2) are presented.

The data was prepared with the method of Chapter 2 by the author.

Results From an Einstein@Home Search for Continuous Gravitational Waves From G347.3 at Low Frequencies in LIGO O2 Data

J. Ming[1,2] , M. A. Papa[1,2,3] , H.-B. Eggenstein[1,2] , B. Machenschalk[1,2], B. Steltner[1,2], R. Prix[1,2], B. Allen[1,2,3], and O. Behnke[1,2]

[1] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany; jing.ming@aei.mpg.de, maria.alessandra.papa@aei.mpg.de

[2] Leibniz Universität Hannover, D-30167 Hannover, Germany

[3] University of Wisconsin Milwaukee, 3135 N Maryland Ave, Milwaukee, WI 53211, USA

Received 2021 August 16; revised 2021 October 18; accepted 2021 November 1; published 2022 January 20

Abstract

We present results of a search for periodic gravitational wave signals with frequencies between 20 and 400 Hz from the neutron star in the supernova remnant G347.3-0.5 using LIGO O2 public data. The search is deployed on the volunteer computing project Einstein@Home, with thousands of participants donating compute cycles to make this endeavour possible. We find no significant signal candidate and set the most constraining upper limits to date on the amplitude of gravitational wave signals from the target, corresponding to deformations below 10^{-6} in a large part of the band. At the frequency of best strain sensitivity, near 166 Hz, we set 90% confidence upper limits on the gravitational wave intrinsic amplitude of $h_0^{90\%} \approx 7.0 \times 10^{-26}$. Over most of the frequency range our upper limits are a factor of 20 smaller than the indirect age-based upper limit.

Unified Astronomy Thesaurus concepts: Gravitational wave astronomy (675); Neutron stars (1108); Supernova remnants (1667)

1. Introduction

Continuous gravitational waves (CWs) are among the gravitational wave signals that have not yet been detected. Fast spinning neutron stars with non-axisymmetric deformations or with unstable r-modes are expected to emit continuous waves that lie in the high-sensitivity frequency range of ground-based interferometers (Owen et al. 1998; Owen 2010; Lasky 2015).

Although the expected waveforms are fairly simple, the search for continuous wave signals is very challenging due to their extreme weakness. A signal-to-noise ratio (S/N) is accumulated by integrating the signal over many months, and this increases our ability to resolve different waveforms. This also means that if the signal waveform is not a priori known, many different waveforms must be searched for, and the computing cost increases very significantly. In fact, when searching a broad range of waveforms, the sensitivity of continuous wave searches is usually limited by the computing power.

Since the Advanced LIGO (Abbott et al. 2015) detectors began observations, various continuous waves searches have been carried out. Among them, the searches for continuous waves from known pulsars, with known spin frequency and frequency evolution, are the most sensitive and computationally inexpensive (Abbott et al. 2019a, 2021a; Ashok et al. 2021). At the other extreme, there are the all-sky surveys with no prior information of frequency and sky location (Dergachev & Papa 2021a, 2020, 2021b; Steltner et al. 2021b; Abbott et al. 2019b, 2021b, 2021c; Covas & Sintes 2020). In between, the directed searches target locations in the sky that are known or suspected to harbor a neutron star, although a pulsation shape has generally not been observed. Searches of this type include

the Galactic Center (Dergachev et al. 2019; Piccinni et al. 2020), young supernova remnants (SNRs; Ming et al. 2019; Abbott et al. 2019c; Millhouse et al. 2020; Lindblom & Owen 2020; Papa et al. 2020; Abbott et al. 2021d), glitching pulsars (Fesik & Papa 2020; Abbott et al. 2021e), and low-mass X-ray binaries such as Scorpius X-1 (Zhang et al. 2021).

Young neutron stars are good continuous wave candidates: an indirect upper limit can be placed on the continuous gravitational wave strength that is proportional to $1/\sqrt{\tau}$, with τ being the age of the neutron star (Wette et al. 2008; Zhu et al. 2016). Fifteen young SNRs have been identified in our Galaxy that could host a young neutron star and potentially be promising targets. Recent searches probe emission from *all* of these over a broad range of waveforms (Lindblom & Owen 2020; Abbott et al. 2021d).

An alternative approach is to identify the most promising targets and concentrate the search efforts on these. In Ming et al. (2016), we propose an optimization scheme to decide how to spend the computing budget in such a way to maximize the detection probability. With a computing budget of a few months on the Einstein@Home volunteer computing project (Einstein@Home 2019), the indication is to search for emission from the neutron star in the SNRs Vela Jr. (G266.2-1.2), Cassiopeia A (G111.7-2.1), and G347.3 (G347.3-0.5). We carried out searches using O1 data, and O2 data for follow-ups, and set the most constraining upper limits on gravitational wave emission from these sources with those data (Ming et al. 2019; Papa et al. 2020).

In Papa et al. (2020), we also found a subthreshold candidate at around 369 Hz. Gravitational wave follow-ups were not completely conclusive and we found no evidence of pulsations from searches of archival X-ray data to validate this candidate, but the X-ray searches had limited sensitivity. Abbott et al. (2021d) did not find this candidate in the first half of O3 data; however the sensitivity of Abbott et al. (2021d) is lower than that of our original search. We thus prioritize a deep search for

1

Ming et al.

G347.3 in the O2 data below 400 Hz. This paper presents results from such a search.

The paper is organized as follows. In Section 2 we review the astrophysical target and the model gravitational waveform. After a brief description of the data in Section 3, in Section 4 we describe the search. The results follow in Section 5, where we explain how the $h_0^{90\%}$ intrinsic continuous gravitational wave amplitude upper limits are derived. These are also recast as upper limits on the star's ellipticity and r-mode saturation. We conclude with a discussion of the results, comparing and contrasting with existing literature in Section 6.

2. The Target

2.1. G347.3-0.5

The SNR G347.3 is suggested to be the remnant of the AD393 "guest star" (Wang et al. 1997). We therefore assume an age of 1600 years, although this estimate is not completely uncontroversial (Fesen et al. 2012). Using XMM data, Cassam-Chenaï et al. (2004) estimate its distance to be around 1.3 kpc. The position of the central compact object in the G347.3 SNR is given with subarcsecond accuracy in Mignani et al. (2008), based on Chandra data. Among the SNRs in our galaxy, G347.3 is one of the most interesting directed search targets because of its relatively young age and close distance (Ming et al. 2016).

In the deep CW search for G347.3 in O1 data (Papa et al. 2020), we find an interesting candidate at around 369 Hz. The spindown energy loss from the candidate parameters yields an unusually high value, 1.6×10^{40} erg s^{-1}, which exceeds the most energetic Crab pulsar's $\dot{E} = 4.6 \times 10^{38}$ erg s^{-1} and J0537-6910's $\dot{E} = 4.9 \times 10^{38}$ erg s^{-1}.

2.2. The Signal

We assume a standard IT2 continuous gravitational wave signal (Dergachev & Papa 2021b) produced by asymmetric rotating neutron stars that, in the detector data, has a form (Jaranowski et al. 1998):

$$h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t), \tag{1}$$

where $F_+(t)$ and $F_\times(t)$ are the antenna pattern functions of the detector for the two gravitational wave polarizations "+" and "×". They depend on the sky position of the source (defined by the R.A. and decl.), and the orientation ψ of the wave frame with respect to the detector frame. $F_+(t)$ and $F_\times(t)$ are periodic time functions with a period of one sidereal day, because the detector rotates with the Earth.

The phase $\Phi(t)$ of the signal at the solar system barycenter (SSB) frame has the form:

$$\Phi(\tau_{SSB}) = \Phi_0 + 2\pi[f(\tau_{SSB} - \tau_{0SSB}) \\ + \frac{1}{2}\dot{f}(\tau_{SSB} - \tau_{0SSB})^2 + \frac{1}{6}\ddot{f}(\tau_{SSB} - \tau_{0SSB})^3], \tag{2}$$

where f is the signal frequency and τ_{SSB} is the arrival time of the GW front at the SSB frame.

3. The Data

The LIGO O2 public data (LIGO 2019; Abbott et al. 2021f) are used in this search. The data are from the two observatories

in the USA, one in Hanford (Washington) and the other in Livingston (Louisiana). The data used in this search are between GPS time 1167983370 (2017 Jan 09) and 1187731774 (2017 Aug 25). Short Fourier transforms (SFTs) of data segments 1800 seconds long (Allen & Mendell 2004) are created as is customary for Einstein@Home searches.

Calibration lines, the mains power lines, and some other spurious noise due to the LIGO laser beam jitter are removed in the publicly released O2 data (Davis et al. 2019). Additionally we remove loud short-duration glitches with the gating procedure described in Steltner et al. (2021a) and substitute Gaussian noise in the frequency domain in disturbed bins. This is a standard procedure in Einstein@Home searches.

4. The Search

We use a "stack-slide" type of search based on the GCT (Global correlation transform) method (Pletsch 2008, 2010; Pletsch & Allen 2009). The data are partitioned in N_{seg} segments and each segment spans a duration T_{coh}. The data of both detectors from each segment i are searched with a maximum likelihood coherent method to construct the detection statistic, $\mathcal{F}$-statistic (Jaranowski et al. 1998; Cutler & Schutz 2005). The statistics $\mathcal{F}_i$ from the coherent searches of the different segments are summed, and the value of the core detection statistic $\overline{\mathcal{F}}$ is obtained:

$$\overline{\mathcal{F}} := \frac{1}{N_{seg}}\sum_{i=1}^{N_{seg}} \mathcal{F}_i. \tag{3}$$

In Gaussian noise $N_{seg} \times 2\overline{\mathcal{F}}$ follows a chi-squared distribution with $4N_{seg}$ degrees of freedom, and a noncentrality parameter ρ^2. If a signal is present, ρ^2 is proportional to $\frac{h_0^2 T_{obs}}{S_h}$, where S_h is the strain power spectral density of the noise at the frequency of the signal, and h_0 is the signal intrinsic amplitude at Earth (Jaranowski et al. 1998).

Small portions of the data exist that are not perfectly Gaussian, and despite the removal of many spectral lines of instrumental and environmental origin, some coherent disturbances persist (Covas et al. 2018). The $\overline{\mathcal{F}}$ can be affected by these coherent disturbances and present increased values. In order to mitigate these occurrences, a line robust detection statistic $\hat{\beta}_{S/GLtL}$ (Keitel et al. 2014; Keitel 2016) is computed. This statistic is the log of a Bayesian odds ratio that tests the signal hypothesis (S) versus an extended noise hypothesis. The "GLtL" noise model of this statistic consists of Gaussian noise (G) or coherent single-detector signals that are always-ON (lines, L) or transient lines (tL).

The Einstein@Home results from this search are ranked according to $\hat{\beta}_{S/GLtL}$, such that the top list contains fewer candidates produced by coherent disturbances.

The search setup, i.e., the coherent baseline T_{coh}, the template grid spacings, and the search ranges are all derived from the optimization procedure.

We search for signal waveforms with frequency and frequency derivatives as follows:

$$\begin{cases} 20\,\text{Hz} \leqslant f \leqslant 400\,\text{Hz} \\ -f/\tau \leqslant \dot{f} \leqslant 0\,\text{Hz/s} \\ 0\,\text{Hz/s}^2 \leqslant \ddot{f} \leqslant 7\,|\dot{f}|_{max}^2/f = 7f/\tau^2, \end{cases} \tag{4}$$

Table 1
Spacings on the Signal Parameters Used for the Templates in the Search

Search Setup					
$T_{\mathrm{coh}} = 1080$ hr	$N_{\mathrm{seg}} = 5$	$\delta f = 1.3 \times 10^{-7}$ Hz	$\delta\dot{f} = 1.5 \times 10^{-14}$ Hz s^{-1}	$\delta\ddot{f} = 1.2 \times 10^{-20}$ Hz s^{-2}	$T_{\mathrm{ref}} = 1177858472.0$[a]

Note.
[a] Barycentric Dynamical Time in GPS seconds.

where $\tau = 1600$ years. The ranges for $\dot{f}$ and $\ddot{f}$ correspond to different breaking index n values, namely 2 and 7. In the $\dot{f}$ equation, $n = 2$ is used to encompass the broadest range of $\dot{f}$ values. In the $\ddot{f}$ equation, $n = 7$ is used to encompass all astrophysical scenarios including the phase evolution purely due to gravitational wave emission ($n = 5$) and r-mode oscillations ($n = 7$). At 400 Hz, the $\dot{f}$ extends down to -8.0×10^{-9} Hz s^{-1} and the $\ddot{f}$ range up to 1.1×10^{-18} Hz s^{-2}.

The grid spacings in frequency and spindowns are constant over these search ranges and are given in Table 1. The number of searched templates per 1 Hz band increases as the frequency increases, as Equation (4) shows. Figure 1 shows the number of templates searched in 1 Hz bands as a function of frequency.

The search is performed on the Einstein@Home volunteer computing project. Einstein@Home is built on the BOINC (Berkeley Open Infrastructure for Network Computing) architecture (Anderson 2004; Anderson et al. 2006), which uses the idle time on volunteer computers to tackle scientific problems such as this that require large amounts of computing power.

Overall we search $\approx 5.1 \times 10^{16}$ templates, utilizing Einstein@Home for several weeks. The workload is split in work units, sized to keep the average volunteer host busy for 8 hours. The whole search task is split into about 2.5 million work units. Only information from the most promising 10,000 results from each work unit is communicated back to the central Einstein@Home server.

5. Results

After the Einstein@Home server has received all search results, the post-processing begins. In total we have 2.5 million work units times 10,000 results returned per work unit $\approx 2.5 \times 10^{10}$ search results. Each result is identified by the template waveform parameters $(f, \dot{f}, \ddot{f})$ and by the detection statistics values.

With a parameter-space clustering procedure we identify the most interesting results (Singh et al. 2017; Beheshtipour & Papa 2020, 2021; Steltner et al. 2021b). We refer to these as "candidates".

We consider the top 1 million candidates, corresponding to a detection statistic threshold $\hat{\beta}_{\mathrm{S/GLtL}r} = 1.948$. The subscript "$r$" refers to a recomputation of the detection statistic performed on all top-list candidates. This recomputation is a typical step in a semicoherent search, where the detection statistic is an approximation of the exact value for any given template. The recomputation step computes the detection statistic at the exact template. If the candidate is due to a signal, the exact template is closer to the signal template and the detection statistic on average increases.

The distribution of the detection statistic $\hat{\beta}_{\mathrm{S/GLtL}r}$ and $2\overline{\mathcal{F}}_r$ for these candidates is shown in Figure 2. We use $\hat{\beta}_{\mathrm{S/GLtL}r}$ to rank our candidates but also show $2\overline{\mathcal{F}}_r$ because its distribution in Gaussian noise is known. A detectable signal would look like
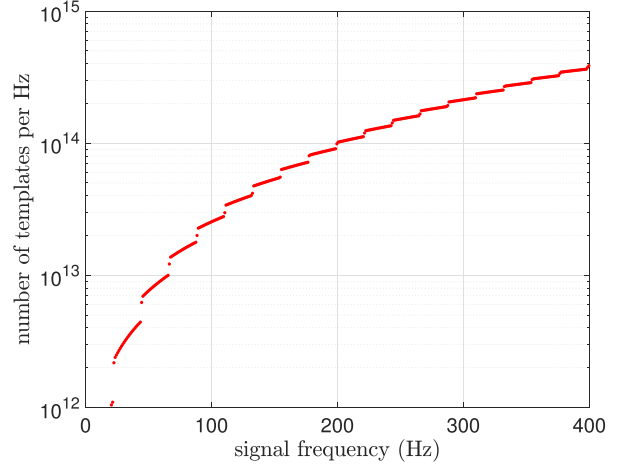
Figure 1. Number of templates searched in 1 Hz bands as a function of signal frequency.

an obvious outlier in both distributions. In Figure 2 we instead see an outlier in the $2\overline{\mathcal{F}}_r$ distribution but not in the $\hat{\beta}_{\mathrm{S/GLtL}r}$ distribution. This is an indication that a coherence in one of the two detectors is causing the high value of $2\overline{\mathcal{F}}_r$. In particular the $2\overline{\mathcal{F}}_r$ outlier has a value of 32.9, whereas its $\hat{\beta}_{\mathrm{S/GLtL}r} = 2.0$, which is in 5th percentile of lowest values. We follow up this candidate with a semicoherent search with $T_{\mathrm{coh}} = 2760$ hr. The most significant result $2\overline{\mathcal{F}}_r = 54.6$ is lower than what one would expect from a signal. The expectation for signals is formed with Monte Carlo simulations of over a thousand fake IT2 signals, added to the O2 data. The data are then processed as in the Einstein@Home search and followed up with the same semicoherent $T_{\mathrm{coh}} = 2760$ hr search as used for the candidate. *All* of the signals show an increase in the detection statistic from the Einstein@Home search to the follow-up, larger than what we find for the candidate. We hence conclude that this candidate is not consistent with our signal model. After excluding this candidate, Figures 2 shows no significant signal candidate in either $2\overline{\mathcal{F}}_r$ or $\hat{\beta}_{\mathrm{S/GLtL}r}$.

Papa et al. (2020) find a subthreshold candidate at around 369 Hz using the O1 data. This candidate has a $2\mathcal{F} = 57.0$[4] in the fully coherent follow-up search of the first half of O2 data, with a p-value of about 4%. In the presence of an IT2 signal, if the measured $2\mathcal{F}$ is close to the expected value, the S/N$^2 \simeq 57 - 4$. This leads to an expected $2\mathcal{F} \approx 110.0$ and a standard deviation of ≈ 21 in the fully coherent search of the whole O2 data set. A fully coherent search using the whole O2 data set around that candidate falls short of the expectation.

The lowest $2\overline{\mathcal{F}}_r$ value among the candidates followed up in the search presented here is 18.2. It corresponds to a $2\mathcal{F} = 75.0$ in the fully coherent search of the whole O2 data

---

[4]   In a fully coherent search $2\mathcal{F}_r = 2\mathcal{F}$ and, of course, $\overline{\mathcal{F}} = \mathcal{F}$.

3

Ming et al.
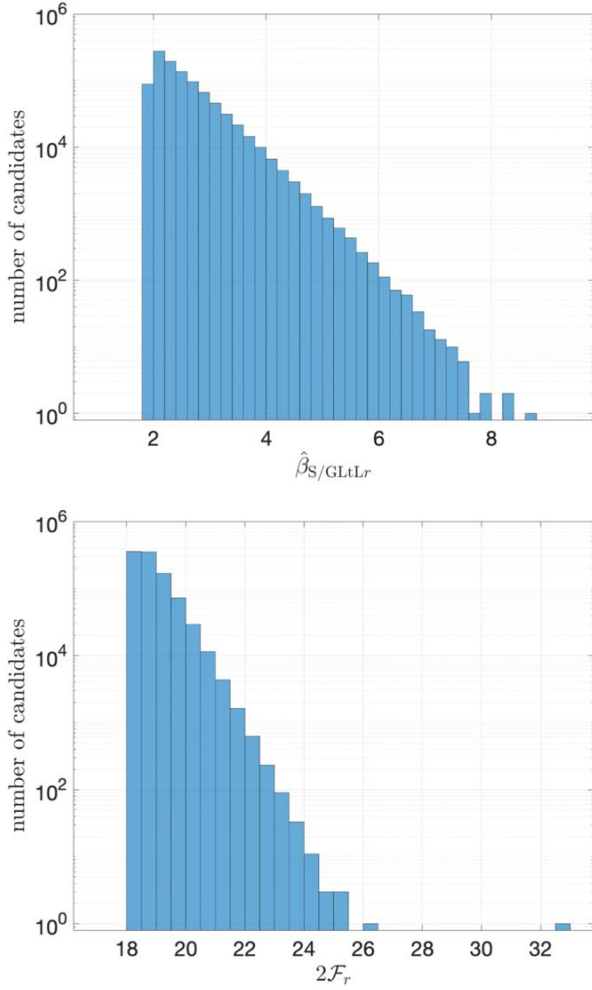
Figure 2. Distribution of the detection statistics $\hat{\beta}_{S/GLtLr}$ (top) and $2\mathcal{F}_r$ (bottom) of the top 1 million candidates ranked according to $\hat{\beta}_{S/GLtLr}$, which is the line and transient-line robust statistic.

set. This search is thus not as sensitive to IT2 signals as the fully coherent O2 follow-up that we performed in Papa et al. (2020). It does however cover the entire parameter space of the original O1 search, not just the small volume around the candidate parameters, and it is more sensitive than the original O1 search. In addition, since it is a semicoherent search, it is more robust to deviations of the signal waveform from the assumed model. So, if the original candidate was due to a signal that deviated a bit from the IT2 model, the current search might well produce a candidate whose parameters are close to, although not perfectly consistent with, those of the candidate from the previous search. This would warrant further investigations. The search however does not show any significant candidates that can be associated with that subthreshold O1 candidate.

5.1. Upper Limits

We determine the smallest h_0 that would have produced a detection statistic as high as the most significant measured in every half Hz band. We assume the source to be at the position of our target, the spindown to be in the target range, and the frequency varying in each half Hz. We set the confidence level
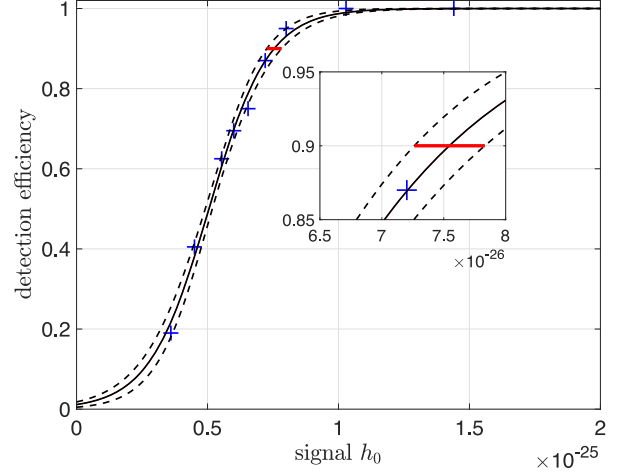
Figure 3. Blue crosses: measured detection efficiency $C(h_0)$ from search-and-recovery Monte Carlos in the frequency band 149–149.5 Hz. The solid line is the best fit and the dashed lines represent 95% confidence intervals on the fit. The red line marks the 90% detection rate, with the uncertainties introduced by this fitting procedure of 4%. The inset shows a zoom around the 90% confidence level.

at 90%, meaning that 90% of the signals in the considered range with an amplitude at the upper-limit value $h_0^{90\%}$ would yield a value of the detection statistic larger than the loudest search result from that parameter range. We use the $\hat{\beta}_{S/GLtLr}$ as our reference statistic, since it is our ranking statistic.

In each half Hz band, 200 simulated signals with a fixed value of the intrinsic amplitude h_0 are added to the real detector data. The data are then processed as the data that were searched, i.e., they are gated and line-cleaned.

The parameters of simulated signals, the frequency, inclination angle $\cos\iota$, polarization ψ, and initial phase values are uniformly randomly distributed in their respective ranges. The spindown values, $\dot{f}$ and $\ddot{f}$, are log-uniformly randomly distributed in their respective ranges.

A search is performed to recover each injection with the same grid and setup as in the original Einstein@Home search. The search is more limited than the original search to save computations, and covers the parameter space neighboring the fake signal. The fake signal is counted as recovered if the $\hat{\beta}_{S/GLtLr}$ from the search is higher than the maximum $\hat{\beta}_{S/GLtLr}$ from the Einstein@Home results in the same half Hz band.

This whole procedure is repeated for various values of h_0. For each value of h_0, the fraction of detected injections is determined in this way and by varying h_0 the confidence $C(h_0)$ curve is constructed. We use a fit with a sigmoid of the form:

$$C(h_0) = \frac{1}{1 + \exp\left(\frac{a - h_0}{b}\right)}, \qquad (5)$$

and from it we read off the h_0 amplitude that corresponds to 90% confidence, our upper limit value.

The Matlab nonlinear regression prediction confidence intervals routine `nlpredci` is used to yield the best fit for a and b values and the covariance matrix. This covariance matrix can be used to compute the 95% credible interval on the fit of $h_0^{90\%}$. Figure 3 shows the sigmoid curve fitting for the 149 −149.5 Hz band, as a representative example of the results obtained with this procedure. The best fit for $h_0^{90\%}$ in this band

4
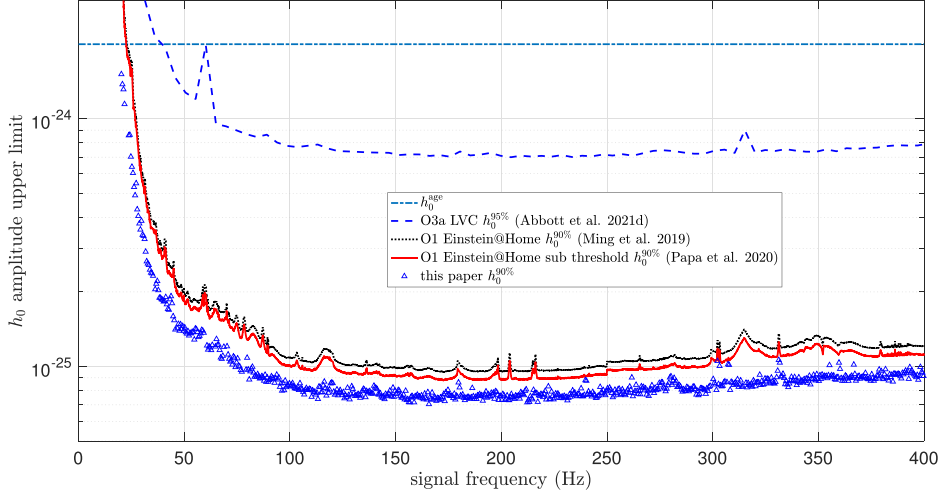
Figure 4. The 90% confidence upper limits on the gravitational wave amplitude of continuous gravitational wave signals from G347.3 for signals with frequencies between 20 and 400 Hz. The lower blue triangles are the results of this search and we compare them with results from previous searches. The blue dots are the upper limits from the LVC search of the O3a (Abbott et al. 2021d); the black dots are Einstein@Home results from O1 data (Ming et al. 2019), and the red solid line is the subthreshold search (Papa et al. 2020).

is 7.5×10^{-26}. The uncertainties introduced by this procedure are less than 4%. The total uncertainty in the upper limit is the sum of the fitting procedure uncertainty and the calibration uncertainties. We conservatively use 5% as the calibration uncertainty (Cahillane et al. 2017).

The $h_0^{90\%}$ upper limits are shown in Figure 4 and are provided in machine readable format at https://www.aei.mpg.de/continuouswaves/O2G347-DirectedSearches.

In nineteen half Hz bands (2.5% of the total) we do not set an upper limit; therefore, in the upper-limit files we have 741 entries rather than 760. The cleaning procedure substitutes disturbed frequency-domain data with Gaussian noise in order to avoid further spectral contamination from "leakage" in the search results. Those bands are consistently cleaned in the upper-limit Monte Carlos after a signal is injected, so it may happen that most of the injected signal is removed. When that happens, no matter how loud the signal is, the detection efficiency does not increase. In these bands the 90% detection rate level cannot be reached and we do not set any upper limit. This reflects the fact that, even if we had a signal there, because of the cleaning procedure, we could not detect it. In the nineteen half Hz bands where the 90% confidence level is not reached, 25% of the data had on average been cleaned away.

In other bands the cleaning procedure partly or completely removes some of the signals, depending on their frequency. So, in order to produce a detection statistic value above a given threshold, statistically, a louder signal is required than in nearby bands that are not cleaned. In those bands the upper limit is higher than what it would be if the data had not been cleaned. For example, $h_0^{90\%}$ of the band 331–331.5 Hz is about 15% larger than the $h_0^{90\%}$ of the neighboring half Hz bands. In this band, 8% of the data are Gaussian noise data.

5.2. Upper Limits on the Astrophysical Parameters

The h_0 upper limits can be converted to constraints on the equatorial ellipticity ε of the neutron star at distance D and at
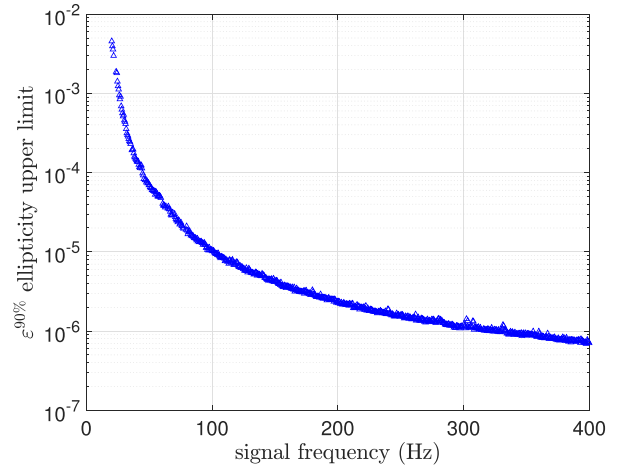
Figure 5. Upper limits on the equatorial ellipticity of G347.3. We assume a distance of 1.3 kpc.

frequency f (Zimmermann & Szedenits 1979):

$$\varepsilon = \frac{c^4}{4\pi^2 G} \frac{h_0 D}{I f^2}, \qquad (6)$$

where c is the speed of light, G is the gravitational constant, and I the principal moment of inertia of the star. Assuming a fiducial value of the principal moment of inertia of 10^{38} kg m^2 and $D = 1.3$ kpc, we convert $h_0^{90\%}(f)$ into upper limits on the ellipticity of the source G347.3. These are shown in Figure 5.

R-mode oscillations of a spinning neutron star also produce continuous gravitational waves. The amplitude h_0 for a signal with frequency f from a source at a distance D depends on the r-mode amplitude α as follows (Owen 2010):

$$\alpha = 0.028 \left(\frac{h_0}{10^{-24}} \right) \left(\frac{D}{1 \text{ kpc}} \right) \left(\frac{100 \text{ Hz}}{f} \right)^3. \qquad (7)$$

Ming et al.
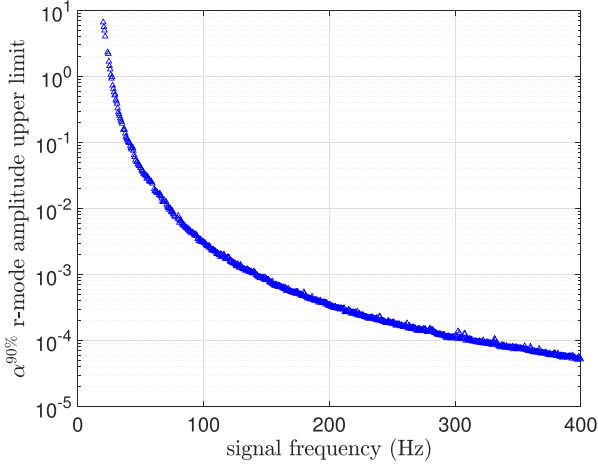
Figure 6. Upper limits on the *r*-mode amplitude.

Our $h_0^{90\%}$ upper limits can then be recast as upper limits on the *r*-mode amplitude. The result is shown in Figure 6.

6. Conclusions

In this paper we present results from the most sensitive search to date for continuous gravitational wave emission from the SNR G347.3-0.5 in the frequency range $20-400$ Hz and the broadest first and second frequency-derivative ranges. Electromagnetic pulsations have not been detected from this object, and the direct observation of continuous gravitational emission would provide the first gravitational wave pulsar timing solution.

We prioritize this target with respect to other SNRs because of a subthreshold candidate from a previous search. We do not find a signal. Our findings are consistent with those of Tenorio et al. (2021).

We constrain the amplitude of continuous gravitational wave emission at a level that is more than a factor of 20 smaller than the indirect age-based limit over most of the frequency range. The most constraining intrinsic gravitational wave amplitude upper limit is 7.0×10^{-26} near 166 Hz. This result improves over our O1 result (Ming et al. 2019) and over the extensive subthreshold O1 search (Papa et al. 2020). It is also more constraining than the recent search result of Abbott et al. (2021d) that uses the significantly more sensitive O3a data. In fairness, we note however that Abbott et al. (2021d) search a broader frequency range and their search uses a technique that is more robust to possible deviations of the signal from the IT-*n* model.

Recast in terms of equatorial ellipticity of the neutron star, our results constrain it to be below 10^{-6} at frequencies higher than ≈ 320 Hz, reaching bounds of 6.9×10^{-7} at 400 Hz. This is a physically plausible value of neutron star deformation (Johnson-McDaniel & Owen 2013; Gittins et al. 2020; Gittins & Andersson 2021). Such a limit is not matched in Abbott et al. (2021d) even at 2000 Hz.

Our spindown range is high enough to allow for braking indexes as high as 7, encompassing *r*-mode emission. Our null result can then constrain the *r*-mode amplitude and does so at a level below 10^{-4} at frequencies higher than ≈ 310 Hz. This is also a physically possible value (Haskell 2015).

This is the first O2 public data Einstein@Home search for continuous gravitational waves from SNRs and probes a physically interesting range of source parameters. Building on this, future searches will extend the parameters space and/or include more targets and/or more data, pushing further into interesting territory.

ORCID iDs

J. Ming ⬥ https://orcid.org/0000-0002-2150-3235
M. A. Papa ⬥ https://orcid.org/0000-0002-1007-5298
H.-B. Eggenstein ⬥ https://orcid.org/0000-0001-5296-7035

References

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, CQGra, 32, 074001
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, ApJ, 879, 10
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, PhRvD, 100, 024004
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019c, ApJ, 875, 122
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, ApJL, 913, L27
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021b, PhRvD, 103, 064017
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021c, PhRvD, 104, 082004
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021d, ApJ, 921, 80
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021e, ApJ, 922, 71
Abbott, R., Abbott, T. D., Abraham, S., et al. 2021f, SoftX, 13, 100658
Allen, B., & Mendell, G. 2004, https://dcc.ligo.org/LIGO-T040164/public
Anderson, D. P. 2004, in Proc. Fifth IEEE/ACM Int. Workshop on Grid Computing (GRID04), 4
Anderson, D. P., Christensen, C., & Allen, B. 2006, in Proc. 2006 ACM/IEEE Conf. on Supercomputing, 126
Ashok, A., Beheshtipour, B., Papa, M. A., et al. 2021, ApJ, 923, 85
Beheshtipour, B., & Papa, M. A. 2020, PhRvD, 101, 064009
Beheshtipour, B., & Papa, M. A. 2021, PhRvD, 103, 064027
Cahillane, C., Betzwieser, J., Brown, D. A., et al. 2017, PhRvD, 96, 102001
Cassam-Chenaï, G., Decourchelle, A., Ballet, J., et al. 2004, A&A, 427, 199

6

Ming et al.

Covas, P. B., & Sintes, A. M. 2020, PhRvL, 124, 191102
Covas, P. B., Effler, A., Goetz, E., et al. 2018, PhRvD, 97, 082002
Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006
Davis, D., Massinger, T., Lundgren, A., et al. 2019, CQGra, 36, 055011
Dergachev, V., & Papa, M. A. 2020, PhRvL, 125, 171101
Dergachev, V., & Papa, M. A. 2021a, PhRvD, 104, 043003
Dergachev, V., & Papa, M. A. 2021b, PhRvD, 103, 063019
Dergachev, V., Papa, M. A., Steltner, B., & Eggenstein, H.-B. 2019, PhRvD, 99, 084048
Einstein@Home 2019, The Volunteer Computing Project Einstein@Home, https://einsteinathome.org/
Fesen, R. A., Kremer, R., Patnaude, D., & Milisavljevic, D. 2012, AJ, 143, 27
Fesik, L., & Papa, M. A. 2020, ApJ, 895, 11
Gittins, F., & Andersson, N. 2021, MNRAS, 507, 116
Gittins, F., Andersson, N., & Jones, D. I. 2020, MNRAS, 500, 5570
Haskell, B. 2015, IJMPE, 24, 1541007
Jaranowski, P., Krolak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001
Johnson-McDaniel, N. K., & Owen, B. J. 2013, PhRvD, 88, 044004
Keitel, D. 2016, PhRvD, 93, 084024
Keitel, D., Prix, R., Papa, M. A., Leaci, P., & Siddiqi, M. 2014, PhRvD, 89, 064023
Lasky, P. D. 2015, PASA, 32, e034
LIGO 2019, The O2 Data Release, doi:10.7935/CA75-FM95

Lindblom, L., & Owen, B. J. 2020, PhRvD, 101, 083023
Mignani, R. P., Zaggia, S., de Luca, A., et al. 2008, A&A, 484, 457
Millhouse, M., Strang, L., & Melatos, A. 2020, PhRvD, 102, 083025
Ming, J., Krishnan, B., Papa, M. A., Aulbert, C., & Fehrmann, H. 2016, PhRvD, 93, 064011
Ming, J., Papa, M. A., Singh, A., et al. 2019, PhRvD, 100, 024063
Owen, B. J. 2010, PhRvD, 82, 104002
Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, PhRvD, 58, 084020
Papa, M. A., Ming, J., Gotthelf, E. V., et al. 2020, ApJ, 897, 22
Piccinni, O. J., Astone, P., D'Antonio, S., et al. 2020, PhRvD, 101, 082004
Pletsch, H. J. 2008, PhRvD, 78, 102005
Pletsch, H. J. 2010, PhRvD, 82, 042002
Pletsch, H. J., & Allen, B. 2009, PhRvL, 103, 181102
Singh, A., Papa, M. A., Eggenstein, H.-B., & Walsh, S. 2017, PhRvD, 96, 082003
Steltner, B., Papa, M. A., & Eggenstein, H.-B. 2021a, arXiv:2105.09933
Steltner, B., Papa, M. A., Eggenstein, H.-B., et al. 2021b, ApJ, 909, 79
Tenorio, R., Keitel, D., & Sintes, A. M. 2021, PhRvD, 104, 084012
Wang, Z. R., Qu, Q. Y., & Chen, Y. 1997, A&A, 318, L59
Wette, K., Owen, B. J., Allen, B., et al. 2008, CQGra, 25, 235011
Zhang, Y., Papa, M. A., Krishnan, B., & Watts, A. L. 2021, ApJL, 906, L14
Zhu, S. J., Papa, M. A., Eggenstein, H.-B., et al. 2016, PhRvD, 94, 082008
Zimmermann, M., & Szedenits, E. 1979, PhRvD, 20, 351

7

Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the Galactic Center

The results of a loosely coherent directed/spotlight search in Advanced LIGO data of the first observation run (O1) are presented.

The loosely coherent search was run by V. Dergachev. The semi-coherent follow-up of the outliers in O1 and O2 data was done by the author.

Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the Galactic Center

Vladimir Dergachev,[1,2,*] Maria Alessandra Papa,[1,2,3,†] Benjamin Steltner,[1,2,‡] and Heinz-Bernd Eggenstein[1,2,§]

[1]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute),*
Callinstrasse 38, 30167 Hannover, Germany
[2]*Leibniz Universität Hannover, D-30167 Hannover, Germany*
[3]*University of Wisconsin Milwaukee, 3135 N Maryland Ave, Milwaukee, 53211 Wisconsin, USA*

We report results of a search for continuous gravitational waves from a region covering the globular cluster Terzan 5 and the Galactic Center. Continuous gravitational waves are expected from fast-spinning, slightly nonaxisymmetric isolated neutron stars as well as more exotic objects. The regions that we target are believed to be unusually abundant in neutron stars. We use a new loosely coherent search method that allows one to reach unprecedented levels of sensitivity for this type of search. The search covers the frequency band 475–1500 Hz and frequency time derivatives in the range of $[-3.0, +0.1] \times 10^{-8}$ Hz/s, which is a parameter range not explored before with the depth reached by this search. As to be expected with only a few months of data from the same observing run, it is very difficult to make a confident detection of a continuous signal over such a large parameter space. A list of parameter space points that passed all the thresholds of this search is provided. We follow up the most significant outlier on the newly released O2 data and cannot confirm it. We provide upper limits on the gravitational wave strength of signals as a function of signal frequency.

I. INTRODUCTION

Centre National de la Recherche Scientifique

Continuous gravitational waves (CWs) are expected from fast-spinning neutron stars in a variety of circumstances, for example if they present a slight nonaxisymmetry (ellipticity). Many CW searches have been carried out on LIGO data [1], including several all-sky searches [2–5] and broadband directed searches [6]. No signals have been detected yet.

Directed searches are searches for signals from interesting targets—both specific objects or/and regions. The search presented here, targeting emission from the globular cluster Terzan 5 and the Galactic Center, falls into this category.

We use data collected during the first Advanced LIGO observing run, O1 [7–10], and employ a new medium-scale loosely coherent algorithm [11–13]. We probe a broad class of signals with frequencies between 475 and 1500 Hz, with unprecedented sensitivity. For sources at 8.5 kpc this search is sensitive to signals from neutron star deformations well within the range allowed by conventional neutron star equations of state [14].

Additionally this search was used as a pilot study of the new loosely coherent search method. The search uses a substantially longer coherence length than used before and hence presents most of the challenges and difficulties of an all-sky search, but without the substantial load of searching the whole sky. This search has exposed performance bottlenecks in the algorithms implementation and has paved the way for the first all-sky loosely coherent search [5].

The paper is organized as follows: Secs. II and III briefly introduce the LIGO detectors, the data that is used and the signal waveform that we target with this search. Section IV describes the features of the main building block of the search, the enhanced loosely coherent method, and Sec. V illustrates the pipeline, including the way the upper limits are established and the ranking of the outliers. The results are presented and discussed in Sec. VI. The Appendix contains the outlier tables.

II. LIGO INTERFEROMETERS AND THE O1 OBSERVING RUN

The LIGO gravitational wave detector consists of two 4 km dual-recycling Michelson interferometers, one in Hanford, Washington, and the other in Livingston, Louisiana, separated by a 3000-km baseline. The interferometer mirrors act as test masses, and the passage of a gravitational wave induces a differential arm length change that is proportional to the gravitational wave strain amplitude. The Advanced LIGO [9,10] interferometers came online in September 2015, after a major upgrade.

---
*vladimir.dergachev@aei.mpg.de
†maria.alessandra.papa@aei.mpg.de
‡benjamin.steltner@aei.mpg.de
§heinz-bernd.eggenstein@aei.mpg.de

The O1 run occurred between September 12, 2015, and January 19, 2016, from which approximately 77 days and 66 days of analyzable data were produced by the Hanford (H1) and Livingston (L1) interferometers, respectively.

Notable instrumental contaminants affecting the searches described here include spectral combs of narrow lines in both interferometers, many of which were identified after the run had ended and were mitigated for future runs [3,4,15]. For instance an 8-Hz comb in H1 with the even harmonics (16-Hz comb) being especially strong was ascribed to digitization roundoff error in a high-frequency excitation applied in order to servo-control the cavity length of the Output Mode Cleaner (OMC). Similarly, a set of lines found to be linear combinations of 22.7 Hz and 25.6 Hz in the L1 data was tracked down to digitization error in an OMC excitation at a still higher frequency.

Although most of these strong and narrow lines are stationary in frequency and hence do not exhibit the Doppler modulations due to the Earth's motion expected for a CW signal from most sky locations, they do degrade the sensitivity to astrophysical signals at the frequencies where they occur.

III. SIGNAL WAVEFORM

In this paper we assume a standard model of a spinning nonaxisymmetric neutron star. Such a neutron star radiates circularly polarized gravitational radiation along the rotation axis and linearly polarized radiation in the directions perpendicular to the rotation axis. For the purposes of detection and establishing upper limits the linear polarization is the worst case, as such signals contribute the smallest amount of power to the detector.

The strain signal measured by a detector is

$$h(t) = h_0 \left(F_+(t, \alpha_0, \delta_0, \psi) \frac{1 + \cos^2(\iota)}{2} \cos(\Phi(t)) \right.$$
$$\left. + F_\times(t, \alpha_0, \delta_0, \psi) \cos(\iota) \sin(\Phi(t)) \right), \quad (1)$$

where F_+ and $F_\times$ are the detector responses to signals with "+" and "×" quadrupolar polarizations [16–18], the sky location of the source is described by right ascension α_0 and declination δ_0, the inclination of the source rotation axis to the line of sight is ι, and we use ψ to denote the polarization angle (i.e., the projected source rotation axis in the sky plane).

The phase evolution of the signal is given by

$$\Phi(t) = 2\pi(f_0 \cdot (t - t_0) + f_0^{(1)} \cdot (t - t_0)^2/2) + \phi, \quad (2)$$

with f_0 being the source frequency and $f_0^{(1)}$ denoting the first frequency derivative (which, when negative, is termed the *spindown*). We use t to denote the time in the Solar System barycenter frame. The initial phase ϕ is computed relative to reference time t_0. When expressed as a function of local time of ground-based detectors, Equation (2) acquires sky-position-dependent Doppler shift terms.

Most natural "isolated" sources are expected to have negative first frequency derivative, due to the energy lost to emission of gravitational or electromagnetic radiation. The frequency derivative can be positive because of residual motions due to, for instance, a long-period orbit.

IV. THE MEDIUM SCALE LOOSELY COHERENT SEARCH

The medium scale loosely coherent search is the basic building block of this search. It is described in [13] and follows earlier loosely coherent implementations [11,12]. Here we highlight features that are useful to understand search output, in particular upper limits and outliers.

The input to the search are Hann-windowed 3600 s short Fourier transforms (SFTs) for each of the LIGO interferometers: $\{a_{tfi}\}$, indexed by time t, discrete frequencies f and interferometer index i. A value of the weighted power sum $P(f_0, \vec{p})$ is computed for every searched wave shape, parametrized by the frequency of the source f_0 and a set of values for its spindown, sky position and source orientation $\vec{p} = (\alpha, \delta, f_0^{(1)}, \iota)$.

The loosely coherent weighted power sum is a bilinear function of the SFT data:

$$P(f_0, \vec{p}) = \frac{\sum_{t_1, t_2, i_1, i_2} K(t_1, t_2, \vec{p}, f_0) a_{t_1 f'_1 i_1} \bar{a}_{t_2 f'_2 i_2}}{\sum_{t_1, t_2} W(t_1, t_2, \vec{p})}. \quad (3)$$

Here f'_1 and f'_2 are the interferometer frame signal frequencies at the detector time t_1 and t_2. The kernel $K(t_1, t_2, \vec{p}, f_0)$ is equivalent to a narrow band filter on the input data that includes phase corrections to account for the signals' Doppler shifts and relativistic effects. The weight term $W(t_1, t_2, \vec{p})$ folds in the noise level of the individual SFTs and the detectors' response to the specific source as a function of time (it is fourth order in the antenna response). The explicit expressions for these functions are very involved, not very illuminating without extensive additional information. We hence do not report them here but rather refer the interested reader to Secs. II–IV of [13].

Because the polarization coefficients are factored out of power sums [Eq. (3)], which involve thousands of SFTs, it is easy to produce separate power sums for any polarization of interest. For instance, we will provide upper limits for a population of circularly polarized signals which corresponds to the star's rotation axis pointing towards us [$\iota = 0$ or π in Eq. (1)].

The fact that we compute power sums makes it possible to set upper limits on the signal strain amplitude by estimating the power excess that we would measure from the target signals at a given strain amplitude. This estimate is computed using the universal statistics algorithm which produces statistically valid results without assumptions on

the probability distribution function of the noise—a rigorous derivation of the algorithm is given in [19]. An intuitive explanation of why this is possible is that if the expected power of the noise is bounded, then the expectation of any continuous function of the noise is also bounded over the space of all probability distributions (in mathematical terminology the space of probability distributions is compact in weak topology). If the noise is Gaussian, the implementation of the *Universal* statistic used in this search provides close-to-optimal values.

In order to bracket the range of upper limit strain values, depending on the orientation of the source, we consider the so called "worst-" and "best-" case polarization upper limits. The upper limits are given as a function of frequency and apply to 0.125 Hz signal-frequency intervals; i.e., there is a single upper limit number for every 0.125 Hz band. The "worst-case" upper limits are based on the maximum universal statistic value over the frequencies in any given band and all spindowns, sky positions and polarizations, further increased (by 7%) to account for losses due to signal-template mismatch.[1] This maximization tends to select increased universal statistic values due to disturbances in the data, when present. For this reason the worst-case upper limit curve has larger outliers than the circular polarization ("best-case" one). The "best-case" upper limits are based on the maximum universal statistic value over the frequencies in any given the band and all spindowns and sky positions, while circular polarization is assumed for incoming signals.

The computation of universal statistic [19] also computes SNR as a byproduct; this is used as a detection statistic for identifying outliers.

V. SEARCH PIPELINE

We search a disk on the sky of radius 0.06 rad (3.43°) centered on right ascension 4.65 rad (266.42°) and declination −0.46 rad (−26.35°). This search area is chosen to cover both the globular cluster Terzan 5 and Sagittarius A*, galactic regions expected to contain many neutron stars. Terzan 5, in particular, has many known radio pulsars [20–22].

The search pipeline iteratively uses the medium scale loosely coherent algorithm in a cascade of 7 different stages. The first stage employs an 8 hour coherence length. Outliers identified at this stage are followed up with more sensitive searches that utilize increasingly longer coherence lengths, as detailed in Table I. For all stages the detection statistic combines coherently over the coherent length the data from both detectors. At the last stage, the detection statistic from each detector separately is also computed and the additional requirement is set on surviving candidates that the parameters be consistent across the multidetector and single-detector statistics. The consistency condition demands that outliers

---

[1]The 7% is derived from the results of Monte Carlo simulations of this search on simulated signals [13].

TABLE I. Search pipeline. Parameters of search pipeline. As explained in the text stage 6 also features an additional consistency check between the single-detector statistics.

Stage	Coherence length (hours)	Minimum SNR
0	8	6
1	12	6.5
2	16	7
3	24	8
4	36	9
5	48	11
6	72	13

from the same sky point and spindown are no further than $5\,\mu\text{Hz}$ in frequency.

The pipeline is validated using extensive Monte Carlos that simulate signals in the real data and test the recovery efficiency of the whole pipeline. This approach is completely standard for this type of search, where the expected signals are weak, and in many frequency bands it is impossible to model the noise reliably. This procedure also validates the correctness of the upper limit values given here.

A. Outlier ranking

The likelihood of a search outlier to have astrophysical origin is commonly described by the false alarm rate—an estimate of probability that this outlier is produced by pure chance. The most obvious method of computing this rate is to repeat the search many times with different realizations of the noise and count how many similar outliers are produced. This is impractical for broad parameter searches which usually take weeks to months to complete.

A commonly used shortcut is to reuse the data from the original search but combine it differently, for instance with nonastrophysical offsets for coincidence parameters (such as time or frequency)—for a notable example see [23]. The idea is to simulate different noise realizations of the search results, by constructing "off-source" combinations of the actual search results. Unfortunately, producing an "off-source" noise realization by combining the single-detector outliers from the last stage of this pipeline is not viable because the preceding stages are based on multidetector statistics. This means that the outliers at the last stage present correlations between the frequencies of peaks in single-detector data. We want the artificially generated noise realizations (the off-source data) to also display such correlations. Unfortunately the standard methods to construct the off-source data by recombining the single-detector candidates with nonphysical offsets would destroy such correlations; hence they are not suitable.

We take here a different approach and derive an approximate analytical expression, under the assumption that underlying noise is Gaussian. This is a strong assumption that is known not to hold in many frequency bands. Thus this expression should not be used as criterion

for detection. Rather it is meant as a figure of merit to evaluate relative significance of outliers.

As the entire hierarchical 7-stage pipeline is difficult to model, we derive the false alarm rate for a hypothetical search that used the last stage of follow-up to analyze the entire parameter space. In the next paragraphs we describe the quantities that are necessary in order to estimate the false alarm rate Eq. (4). These quantities are: the total number of templates N that would have been used by the stage 6 search over the entire searched parameter space; the distribution of the detection statistic for the stage 6 search, $P_{\chi^2, k}$; the "coincidence probability" associated with the multidetector/single-detector consistency check, p_{coinc}. We derive these below.

We (over)-estimate the total number of templates N necessary to perform such a search to be 1.6×10^{27}. We arrive at this number as follows: The total number of templates in the grid for the entire search over 1025 Hz, the whole sky, polarization and spindown is 9.3×10^{21}. We however search more waveforms than these because we additionally allow the frequency to change by up to one frequency bin 11 times, equally spaced throughout the observation period. This adds robustness to our search with respect to deviations of the real signal from a strictly coherent signal model. To account for this, we increase 9.3×10^{21} by a factor of 3^{11}. This overcounts the number of independent templates. For example, two templates different only by a single jump in frequency bin in the middle of the run would be highly correlated.

Because we consider the last stage as a separate search the frequencies of outliers in individual interferometers are independent. The frequency coincidence criterion can be falsely triggered in pure noise with probability $p_{\text{coinc}} = 3.59 \times 10^{-5}$.

The last stage of the analysis uses a 3-day coherence time. As the variations in W [Eq. (3)] due to amplitude modulations average out over this time, the power sums can be modeled as a χ^2 variable with at most $k = 80$ degrees of freedom, with k expected to be smaller for frequency regions with highly contaminated data. The reason for decrease in k is that the terms in the sum [Eq. (3)] containing contaminated data are deweighted, and hence they contribute less than others to the total number of degrees of freedom. In the case of equal weighted data $k = 80$ because there are 40 3-day chunks in a 4-month run and each chunk contributes two degrees of freedom.

We take the Gaussian false alarm figure of merit for a candidate at signal-to-noise ratio value SNR, at the end of the last follow-up stage, to be

$$\log_{10}(\text{GFA}) = \log_{10}(P_{\chi^2, k}(k + \sqrt{2k} \cdot \text{SNR}))$$
$$+ \log_{10}(N) + \log_{10}(p_{\text{coinc}}), \qquad (4)$$

where SNR is defined as the ratio of the deviation of the detection statistic from its expected value to the standard deviation.

We emphasize again that the formula (4) was derived under the assumption of stationary Gaussian noise that is independent between the H1 and L1 interferometers. Since this assumption is violated in many frequency bands, this figure is not meant as a criteria for detection. For example, large negative values for outliers 1 through 8 are an indication of a presence of a signal, but these signals are known to be instrumental in origin.

VI. RESULTS

The search produces a number of outliers, the strongest of which are traced to clear instrumental artifacts. A number of unclassified outliers with smaller signal-to-noise ratios passes the follow-up pipeline. While the pipeline has been demonstrated to recover injected signals successfully even in the most heavily contaminated regions [13], the presence of noise does increase the false alarm rate. As the O1 data is highly contaminated with both stationary and nonstationary instrumental lines, classification of weak outliers is particularly difficult. This problem is made more challenging by the presence of instrumental artifacts coherent between both interferometers.

We further extend the coherent baseline of the search with *ad hoc* semicoherent follow-up searches, like the ones used in [2,3], on 352 outliers. We use three stages with coherent baselines of 210 hrs (12 segments), 500 hrs (6 segments) and 1260 hrs (2 segments), respectively. We denote the stages by FU0, FU1 and FU2. Since FU1 is rather computationally intensive we do not follow up any outlier that can be associated with a disturbance (see comment field in the tables of the Appendix. Twenty-one outliers survive all thresholds from these follow-up searches. The outlier with id 68 appears to be the most significant. On it we perform a dedicated search using the FU1 search on 480 hrs of the newly released data from the O2 run [8]. The search could not recover the candidate with detection statistic values consistent with what would have been expected if outlier 68 had been a continuous wave described by Eq. (1). The Appendix details all the outliers and indicates at what stage of these follow-ups the candidate was rejected.

The simulations described in [13] have shown that an astrophysical source adhering to the expected signal model will be recovered within 15 μHz of true frequency and within 1.5×10^{-11} Hz/s of true spindown. The sky position mismatch depends on frequency and, for outliers with frequency f is no more than $6.5 \times 10^{-4} \cdot (1 \text{ kHz}/f)$ in ecliptic distance, defined as the distance between outlier location and true injection location after projection onto the ecliptic plane.

The universal statistic algorithm allows one to set valid upper limits even in the most heavily contaminated bands.
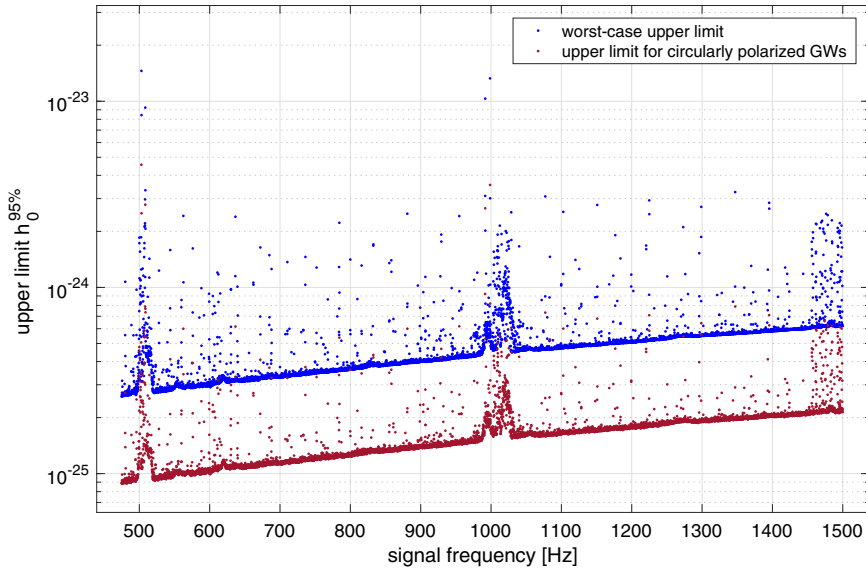
FIG. 1.   Upper limits on gravitational wave strain. The dimensionless strain h_0 (vertical axis) is plotted against signal frequency. The upper (blue) curve shows worst-case (linearly polarized) 95% confidence level (CL) upper limits as a function of frequency. The upper limits are maximized over sky and all intrinsic signal parameters for each frequency band. The lower (red) curve shows upper limits assuming a circularly polarized source. The data for this plot can be found in [24].
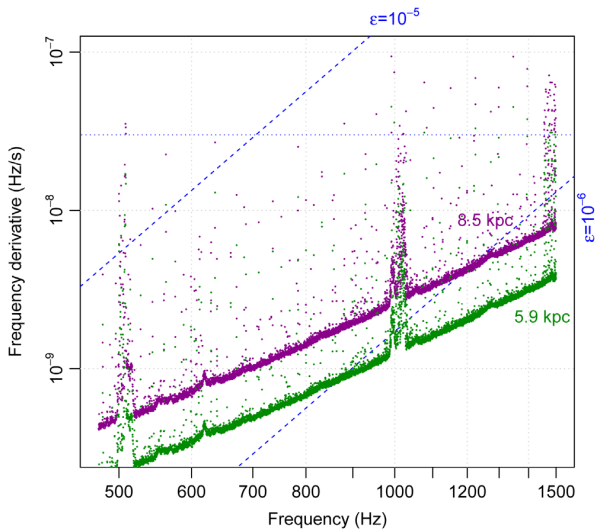
FIG. 2.   Range of the search for neutron stars spinning down solely due to gravitational radiation. This is a superposition of two contour plots. The green (bottom) and purple (top) solid markers are contours of the maximum distance at which a neutron star could be detected as a function of gravitational wave frequency f and its derivative $\dot{f}$. The dashed lines are contours of the corresponding ellipticity $\epsilon(f, \dot{f})$. The fine dotted line marks the maximum spindown searched. Together these quantities tell us the maximum range of the search in terms of various populations (see text for details). 5.9 kpc is the distance associated to many ATNF [22,25] catalogue pulsars in the direction of Terzan 5.

Figure 1 shows the best-case and worst-case 95% confidence upper limits on the signal strain in 0.125 Hz frequency intervals. At the highest frequency (1500 Hz) the worst-case upper limit on gravitational wave strain is 6.2×10^{-25}, which translates in a source with an ellipticity of 2.5×10^{-6} at 8.5 kpc. Because of the maximization procedure the confidence level of the worst-case upper limits remains 95% or higher for any subset of parameters. For example, if one picks a sky location of the Terzan 5 globular cluster, spindown of 5×10^{-9} Hz/s and a frequency of 550 Hz the worst case upper limit is 2.89×10^{-25} Hz, with a confidence level which is guaranteed to be at least 95%. The actual confidence level is likely to be larger than 95% for the specific point, because the quoted upper limit is the highest over all sampled spindowns and the wider sky area.

Figure 2 shows the astrophysical reach of the search, i.e., the maximum distance at which this search could have detected a signal of a given frequency and spin-down, under the assumption that all the lost rotational energy is emitted in gravitational waves. The search presented here is sensitive to an optimally oriented neutron star at the Galactic Center (circularly polarized signal) with ellipticity of 10^{-6} and emitting gravitational waves at a frequency of 1200 Hz. In Terzan 5 a signal at 1200 Hz from an optimally oriented source having ellipticity of $\leq 7 \times 10^{-7}$ could have been detected.

The search presented is the most sensitive to date, aimed at this interesting region of our galaxy. This is reflected in the sensitivity depth of the search which is defined as the

ratio of the upper limit value and the noise floor at nearby frequencies $\mathcal{D}(f) := \sqrt{\frac{S_h(f)}{h_0^{UL}}}$ [26]. Following [27], we estimate the noise taking the harmonic mean across the different detectors and obtain the following values of the sensitivity depth across the entire frequency range searched:

$$
\begin{cases}
\mathcal{D}_{\text{circ-pol}}(f) = 116 \; [\sqrt{\text{Hz}}]^{-1/2} \\
\mathcal{D}_{\text{worst-pol}}(f) = 42 \; [\sqrt{\text{Hz}}]^{-1/2}.
\end{cases}
\tag{5}
$$

The radiometer search [28] targeting the Galactic Center is 4 times less sensitive than our most conservative upper limit (the worst case one), achieving, on the same data, a sensitivity depth smaller than 10. This search covers a larger spindown range than any previously published all-sky search, hence probing younger sources from our search area. Furthermore even our worst-case upper limits are more constraining than any all-sky search result to date, including the state of the art paper [29] that uses the more sensitive and longer duration data set from the O2 run.

APPENDIX: OUTLIER TABLES

Outliers passing all stages of automated follow-up from 475–1500 Hz band are separated into five tables. Table VI shows outliers inside the contaminated regions 495–520 Hz and 990–1033 Hz. The rest of the outliers are split into four regions 475–900 Hz, 900–1200 Hz, 1200–1400 Hz and 1400–1500 Hz (Tables II–V).

TABLE II. Outliers below 900 Hz that passed the automated detection pipeline excluding regions heavily contaminated with violin modes. Outliers marked with "line" had strong narrow-band disturbances identified near the outlier location. We have marked outliers not consistent with the target signals at one of the semicoherent $\mathcal{F}$-statistic follow-ups with "(FU0/1/2)", depending on the stage at which they did not pass the detection thresholds. Frequencies are converted to epoch GPS 1130529362.

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{\text{J2000}}$ degrees	DEC$_{\text{J2000}}$ degrees	Description
6	23.9	−17.5	612.486 10	−9.197	267.037	−29.754	Broad large line in L1 at 612.45 Hz
8	23.5	−16.6	736.094 75	−22.991	267.808	−28.287	Sharp bin-centered lines at 736 Hz (H1) and 736.1 (L1)
13	21.2	−11.7	736.097 91	−21.997	266.439	−29.165	Sharp bin-centered lines at 736 Hz (H1) and 736.1 (L1)
17	20.3	−9.6	684.965 15	−17.607	265.576	−22.464	Sharp bin-centered line in L1 at 684.9 Hz
20	19.2	−7.5	662.183 56	−24.621	263.548	−25.681	Strong bin-centered line in L1 at 662.20 Hz
24	17.9	−4.8	710.544 65	−21.826	264.817	−27.547	Strong bin-centered line in L1 at 710.50 Hz
32	17.6	−4.0	599.193 67	−15.234	266.039	−28.665	Large broad lines in H1 near 599.14 Hz and 599.16 Hz
34	17.5	−3.8	761.755 80	−13.218	265.953	−25.889	Strong bin-centered line in L1 at 761.70 Hz (FU0)
43	16.8	−2.4	707.651 62	−10.129	266.799	−24.227	Strong bin-centered line in L1 at 707.6 Hz
44	16.6	−2.1	575.231 74	−5.519	266.534	−29.341	Hardware injected pulsar 2
45	16.6	−2.0	898.869 70	−0.491	269.397	−28.914	Large broad lines in H1
47	16.5	−1.9	898.846 67	−15.157	266.790	−26.909	Large broad lines in H1
70	15.2	0.8	659.354 18	−12.386	269.229	−26.124	Strong bin-centered line in L1 at 659.3 Hz
72	15.1	0.9	629.864 31	−29.241	267.817	−25.916	Large broad lines in L1

(Table continued)

TABLE II. *(Continued)*

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	Description
75	15.0	1.1	787.356 87	−8.004	265.084	−29.806	Sharp bin-centered line in L1 at 787.3 Hz
80	14.8	1.4	660.513 61	−20.472	268.538	−22.492	
84	14.8	1.5	829.859 46	−8.609	264.856	−25.815	Strong bin-centered line in L1 at 829.8 Hz
85	14.8	1.5	520.848 15	−3.204	267.198	−29.179	Large broad line in H1 near 520.82 Hz
89	14.7	1.7	520.848 14	−3.201	267.202	−29.359	Large broad line in H1 near 520.82 Hz
99	14.5	2.1	763.943 06	−20.331	262.960	−26.415	Hardware injected pulsar 9
102	14.5	2.1	873.267 13	−29.286	264.627	−28.210	(FU0)
103	14.5	2.1	606.636 06	−27.504	263.085	−27.581	Large broad line in H1 at 606.67 Hz
104	14.4	2.3	730.353 49	−8.153	266.654	−26.898	Sharp bin-centered line in L1 at 730.3 Hz
119	14.2	2.7	787.355 42	−8.931	264.431	−28.451	Sharp bin-centered line in L1 at 787.3 Hz
120	14.1	2.8	608.065 95	−19.469	268.412	−28.434	Sharp bin-centered line in H1 at 608 Hz
121	14.1	2.8	899.256 24	−27.714	263.423	−28.511	Strong broad lines in H1
135	14.0	3.1	599.496 00	−17.127	267.287	−28.898	Strong broad lines in H1
137	14.0	3.1	771.051 17	−8.789	265.082	−22.538	(FU1)
139	13.9	3.1	587.372 28	−0.374	265.864	−29.281	(FU1)
145	13.9	3.2	864.060 26	−9.286	262.769	−25.668	Sharp bin-centered line in H1 at 864 Hz
146	13.9	3.3	575.233 99	−5.628	268.609	−27.074	Hardware injected pulsar 2
149	13.9	3.3	764.656 86	−8.411	267.769	−27.209	Sharp bin-centered line in L1 at 764.6 Hz
151	13.9	3.3	817.316 37	−23.376	267.491	−28.124	(FU1)
156	13.8	3.4	773.675 02	−4.093	265.308	−23.202	
161	13.8	3.5	618.061 67	−13.849	268.606	−25.285	Slope in H1 spectrum
166	13.8	3.5	738.022 55	−0.388	263.204	−27.112	(FU0)
169	13.7	3.6	629.864 33	−29.236	267.845	−26.116	Large broad line in L1
170	13.7	3.6	769.353 42	−26.944	267.449	−22.402	(FU1)
172	13.7	3.6	686.755 65	−11.731	264.065	−29.473	(FU0)
176	13.7	3.7	764.656 87	−8.411	267.773	−27.188	Sharp bin-centered line in L1 at 764.6 Hz
180	13.6	3.7	683.402 67	−12.421	265.532	−28.007	Slight slope in L1 (FU0)
186	13.6	3.8	799.267 03	−5.724	266.255	−28.143	(FU0)
190	13.6	3.8	824.590 28	−2.196	268.212	−27.435	(FU0)
192	13.6	3.8	645.946 31	−15.366	266.741	−26.423	
202	13.5	3.9	727.325 68	−17.438	267.008	−27.029	(FU1)
203	13.5	3.9	539.858 63	−8.429	267.373	−29.396	Near 60 Hz line
215	13.5	4.1	851.689 71	−18.266	266.383	−22.291	(FU1)
219	13.4	4.1	489.119 59	−5.482	265.242	−25.347	Nearby lines
227	13.4	4.2	694.426 37	−26.409	267.828	−27.898	(FU0)
229	13.4	4.2	581.710 75	−9.491	263.276	−28.388	(FU1)
232	13.4	4.3	713.463 88	−2.209	266.696	−26.761	Strong bin-centered line in L1 at 713.4 Hz
243	13.3	4.3	575.256 58	−18.581	265.147	−29.940	Hardware injected pulsar 2
244	13.3	4.3	583.358 70	−8.394	269.473	−29.052	Strong broad line in H1 at 583.317 (FU0)
248	13.3	4.4	763.951 14	−21.211	266.312	−24.681	Hardware injected pulsar 9
249	13.3	4.4	680.273 62	−14.639	263.404	−25.497	(FU0)
251	13.3	4.4	608.001 20	−16.846	263.648	−28.505	Sharp bin-centered line in H1 at 608 Hz
254	13.3	4.4	770.006 05	−19.542	264.245	−27.578	(FU1)
260	13.3	4.4	772.834 74	−27.344	267.338	−26.027	Sharp bin-centered line in L1 at 772.8 Hz
261	13.3	4.4	809.984 76	−27.979	263.634	−25.694	Sharp bin-centered line in L1 at 810 Hz
264	13.3	4.4	694.426 38	−26.413	267.818	−27.711	(FU0)
270	13.3	4.5	878.165 83	−14.756	267.454	−26.247	Strong bin-centered line in L1 at 878.1 Hz
282	13.2	4.5	547.675 10	−9.729	263.964	−27.205	(FU0)
287	13.2	4.6	829.860 26	−8.384	265.709	−27.241	Strong bin-centered line in L1 at 829.8 Hz
290	13.2	4.6	792.702 63	−16.328	266.224	−29.232	(FU0)
291	13.2	4.6	655.434 80	−2.403	263.896	−25.645	(FU0)
297	13.2	4.6	848.062 91	−16.696	264.046	−26.114	Strong bin-centered line in H1 at 848 Hz
299	13.2	4.6	725.527 68	−5.731	263.854	−25.707	(FU0)

(Table continued)

TABLE II. *(Continued)*

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA_{J2000} degrees	DEC_{J2000} degrees	Description
303	13.2	4.6	782.792 14	−20.134	265.043	−26.320	(FU1)
309	13.1	4.7	599.203 98	−18.747	269.644	−26.805	Big broad lines in H1 near 599.14 and 599.16 Hz
310	13.1	4.7	698.222 24	−22.943	263.921	−25.129	(FU0)
311	13.1	4.7	763.951 14	−21.208	266.313	−24.999	Hardware injected pulsar 9
313	13.1	4.7	527.127 74	−7.311	266.075	−25.556	(FU0)
316	13.1	4.8	753.975 28	0.849	266.035	−25.250	(FU0)
317	13.1	4.8	844.158 41	−9.201	262.865	−24.558	(FU1)
319	13.1	4.8	799.638 51	−3.359	265.959	−29.050	(FU1)
323	13.1	4.8	718.028 58	−22.968	266.134	−27.019	(FU1)
325	13.1	4.8	527.641 48	−23.659	268.915	−24.650	(FU0)
335	13.0	4.9	621.850 99	−11.539	266.052	−27.996	Sloping H1 spectrum
337	13.0	4.9	676.377 58	−0.769	268.788	−24.882	(FU0)
340	13.0	4.9	678.392 54	−3.871	263.974	−28.642	

TABLE III. Outliers in frequency range 900–1200 Hz that passed the detection pipeline excluding regions heavily contaminated with violin modes. Outliers marked with "line" had strong narrow-band disturbances identified near the outlier location. We have marked outliers not consistent with the target signals at one of the semicoherent $\mathcal{F}$-statistic follow-ups with "(FU0/1/2)", depending on the stage at which they did not pass the detection thresholds. Frequencies are converted to epoch GPS 1130529362.

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA_{J2000} degrees	DEC_{J2000} degrees	Description
5	25.2	−20.4	1176.697 99	−26.024	264.288	−23.567	Strong bin-centered line in L1 at 1176.6 Hz
16	20.8	−10.8	955.008 51	−26.744	269.405	−24.278	Sharp line in L1
19	19.3	−7.6	910.171 53	−8.556	264.028	−24.077	Large broad line in H1 at 910.1 Hz
33	17.5	−4.0	1176.586 14	−3.254	263.501	−25.977	Strong bin-centered line in L1 at 1176.6 Hz
39	17.0	−2.9	1120.099 15	−25.758	264.594	−27.448	Strong bin-centered line in H1 at 1120 Hz
40	16.9	−2.6	910.183 76	−22.563	268.269	−25.989	Large broad line in H1 at 910.1 Hz
58	15.5	0.1	1173.802 11	−28.189	264.045	−27.467	Strong bin-centered line in L1 at 1173.7 Hz
59	15.5	0.1	1128.383 43	−12.747	267.927	−27.526	Strong bin-centered line in L1 at 1128.3 Hz
63	15.3	0.4	906.633 79	−17.648	264.077	−23.762	Strong bin-centered line in L1 at 906.6 Hz
68	15.2	0.6	1105.157 33	−26.774	263.713	−26.330	(FU1 w. O2 data)
76	15.0	1.2	1128.393 96	−25.644	263.670	−27.537	Strong bin-centered line in L1 1128.3 Hz
77	14.9	1.3	946.923 21	−20.521	266.068	−27.668	(FU0)
81	14.8	1.4	977.672 44	−18.241	264.215	−26.981	Strong bin-centered line in L1 at 977.6 Hz
86	14.7	1.6	1130.322 68	−1.751	267.651	−29.016	(FU0)
101	14.5	2.1	983.469 51	−11.611	264.204	−23.932	Strong bin-centered line in L1 at 983.4 Hz
106	14.4	2.3	976.019 56	−26.851	265.826	−29.672	Line in H1 at 976 Hz
113	14.3	2.5	957.885 53	−2.921	266.766	−24.376	Sharp bin-centered line in L1 at 957.8 Hz
114	14.3	2.5	932.288 76	−0.777	265.303	−25.123	Sharp bin-centered line in L1 at 932.2 Hz

(Table continued)

TABLE III. *(Continued)*

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	Description
116	14.2	2.6	1083.851 97	−19.421	266.328	−24.585	(FU0)
118	14.2	2.7	1117.937 10	−29.371	266.174	−24.014	
125	14.1	2.9	1192.547 99	0.039	267.676	−26.744	(FU0)
126	14.1	2.9	1144.817 99	−12.261	265.508	−28.590	(FU0)
130	14.0	3.0	1146.010 80	−14.459	264.505	−26.227	(FU1)
133	14.0	3.1	1056.491 44	−6.807	263.080	−25.072	(FU1)
143	13.9	3.2	916.791 25	−21.067	266.033	−24.889	(FU0)
147	13.9	3.3	1055.064 00	−21.219	267.870	−25.985	(FU0)
148	13.9	3.3	1148.128 64	−9.921	268.460	−28.864	Strong bin-centered line in L1 at 1148.1 Hz
158	13.8	3.4	1193.005 46	−15.631	266.493	−25.264	(FU0)
159	13.8	3.5	911.769 58	−9.399	266.039	−25.180	(FU0)
160	13.8	3.5	1130.731 54	−3.169	263.000	−26.335	(FU0)
162	13.8	3.5	953.400 39	−10.813	266.270	−27.596	
174	13.7	3.6	1087.755 30	−28.077	269.614	−26.525	(FU2)
181	13.6	3.7	969.522 38	−8.313	265.371	−29.459	(FU0)
206	13.5	4.0	1159.915 42	−24.221	264.096	−25.279	(FU0)
208	13.5	4.0	1142.866 54	−6.054	263.246	−24.169	(FU0)
210	13.5	4.0	934.782 61	−2.273	268.362	−25.152	(FU0)
211	13.5	4.0	1080.014 73	−27.759	266.170	−28.283	Strong coincident bin-centered lines in H1 and L1 at 1080 Hz
213	13.5	4.1	1127.802 27	−19.557	267.194	−23.981	(FU1)
223	13.4	4.2	970.662 43	−29.248	266.863	−25.559	(FU0)
228	13.4	4.2	931.299 79	−10.836	266.227	−29.987	
233	13.4	4.3	1151.536 14	−8.588	265.597	−23.441	
237	13.3	4.3	1145.637 73	−5.498	268.460	−29.377	(FU0)
242	13.3	4.3	983.476 59	−7.236	265.003	−27.225	Strong bin-centered line in L1 at 983.4 Hz
245	13.3	4.3	1197.770 64	−27.963	269.015	−27.561	
253	13.3	4.4	903.290 02	−15.369	266.450	−25.044	(FU0)
258	13.3	4.4	938.884 34	−9.556	264.187	−28.599	(FU0)
265	13.3	4.4	953.400 39	−10.816	266.269	−27.442	
266	13.3	4.5	906.901 04	−20.814	263.359	−26.506	Broad line in H1 near 906.82 Hz
267	13.3	4.5	1069.138 74	−2.926	263.432	−27.102	
272	13.2	4.5	1033.967 10	−20.787	262.779	−25.972	(FU0)
273	13.2	4.5	1039.414 70	0.882	267.926	−24.697	
276	13.2	4.5	1121.718 65	−15.156	263.124	−27.755	(FU0)
279	13.2	4.5	1102.259 11	−26.261	266.982	−27.473	(FU1)
284	13.2	4.5	1055.961 11	−20.166	266.229	−28.653	(FU0)
285	13.2	4.5	1081.910 76	−29.388	264.805	−29.647	(FU0)
295	13.2	4.6	1143.156 76	−5.409	269.119	−26.704	(FU1)
300	13.2	4.6	951.012 13	−4.077	266.334	−28.538	(FU0)
305	13.2	4.7	1070.506 37	−4.409	264.544	−27.297	
307	13.1	4.7	1123.018 94	−10.841	268.163	−24.666	(FU0)
314	13.1	4.7	945.033 86	−24.222	267.786	−26.393	Bump in L1
318	13.1	4.8	989.778 30	−4.182	269.861	−25.516	Disturbed H1 spectrum
330	13.1	4.8	1176.336 62	−22.274	264.698	−27.271	(FU1)
336	13.0	4.9	985.143 27	−23.229	265.622	−26.558	Many strong nearby lines in H1
338	13.0	4.9	1090.818 73	−0.502	264.444	−29.353	(FU0)
339	13.0	4.9	1196.002 79	−2.451	267.251	−24.656	(FU0)
342	13.0	4.9	1037.605 85	−26.743	263.801	−28.995	Strong bin-centered line in L1 at 1037.5 Hz
350	13.0	4.9	1197.484 07	−23.759	269.911	−25.523	

106

TABLE IV. Outliers in frequency range 1200–1400 Hz that passed the detection pipeline excluding regions heavily contaminated with violin modes. Outliers marked with "line" had strong narrow-band disturbances identified near the outlier location. We have marked outliers not consistent with the target signals at one of the semicoherent $\mathcal{F}$-statistic follow-ups with "(FU0/1/2)", depending on the stage at which they did not pass the detection thresholds. Frequencies are converted to epoch GPS 1130529362.

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA_{J2000} degrees	DEC_{J2000} degrees	Description
1	30.5	−32.4	1220.623 44	−16.068	265.994	−24.436	Induced by hardware injection 7 (FU0)
37	17.0	−2.9	1360.092 84	−16.252	262.922	−27.103	Strong bin-centered line in H1 at 1360 Hz
42	16.8	−2.5	1276.226 72	−0.304	268.176	−23.187	Strong bin-centered line in L1 at 1276.1 Hz
55	15.8	−0.4	1202.299 27	−6.043	264.510	−24.894	Strong bin-centered line in L1 at 1202.2 Hz
56	15.7	−0.2	1376.122 53	−4.253	269.645	−26.036	Strong bin-centered line in H1 at 1376 Hz
67	15.3	0.6	1280.129 32	0.519	268.744	−23.550	Strong bin-centered line in H1 at 1280 Hz
69	15.2	0.8	1270.395 56	−12.976	268.516	−24.920	Strong bin-centered line in L1 at 1270.3 Hz
78	14.9	1.3	1328.132 75	0.159	269.294	−26.578	Strong bin-centered line in H1 at 1328 Hz
82	14.8	1.5	1202.299 85	−5.886	264.920	−25.698	Strong bin-centered line in L1 at 1202.2 Hz
91	14.7	1.8	1376.098 11	−20.913	267.292	−26.665	Strong bin-centered line in H1 at 1376 Hz (FU0)
93	14.6	1.8	1355.197 92	−22.717	265.274	−27.347	(FU1)
96	14.6	2.0	1303.930 01	−13.662	268.017	−29.082	(FU0)
107	14.3	2.4	1321.589 65	−19.634	263.445	−25.948	Strong bin-centered line in L1 at 1321.5 Hz
111	14.3	2.5	1352.118 73	−4.071	266.570	−28.811	Bin-centered line in H1 at 1352 Hz
112	14.3	2.5	1254.440 64	−15.222	265.991	−25.692	(FU1)
115	14.2	2.6	1301.622 21	−15.454	265.848	−28.031	
117	14.2	2.6	1202.303 56	−26.059	267.238	−23.480	Strong bin-centered line in L1 at 1202.2 Hz
123	14.1	2.9	1270.393 17	−17.622	268.335	−25.117	Strong bin-centered line in L1 at 1270.3 Hz
127	14.0	3.0	1386.494 65	−13.231	267.856	−23.177	(FU0)
128	14.0	3.0	1380.261 31	−19.417	263.195	−24.722	(FU0)
129	14.0	3.0	1264.076 44	−19.604	265.130	−29.366	Line in at 1264 Hz in H1
138	13.9	3.1	1249.432 52	−2.559	265.309	−23.949	(FU0)
140	13.9	3.2	1373.863 00	−7.084	263.597	−27.357	(FU0)
141	13.9	3.2	1205.724 93	−26.927	264.610	−28.241	(FU0)
142	13.9	3.2	1271.073 64	−17.608	266.148	−28.676	(FU0)
144	13.9	3.2	1366.559 54	−0.596	262.870	−26.589	(FU0)
150	13.9	3.3	1331.308 90	−21.827	263.364	−28.484	(FU1)
152	13.8	3.4	1264.103 09	−6.094	262.596	−26.501	Strong bin-centered line in H1 at 1264 Hz
163	13.8	3.5	1315.289 28	−21.894	270.072	−25.232	(FU0)
165	13.8	3.5	1380.261 31	−19.417	263.200	−24.633	(FU0)
171	13.7	3.6	1269.918 95	−21.536	267.428	−28.206	(FU0)
173	13.7	3.6	1276.990 16	−0.241	265.645	−29.501	(FU1)
175	13.7	3.7	1332.838 14	−16.049	265.781	−27.931	(FU0)
177	13.7	3.7	1372.181 44	−19.329	264.461	−27.028	(FU1)
179	13.7	3.7	1267.041 68	−8.514	268.806	−23.928	(FU0)
182	13.6	3.7	1232.095 95	−8.929	266.351	−26.068	Strong bin-centered line in H1 at 1232 Hz

(Table continued)

TABLE IV. *(Continued)*

Idx	SNR	log$_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	Description
188	13.6	3.8	1254.440 64	−15.229	265.990	−25.311	(FU1)
189	13.6	3.8	1394.570 34	−8.916	269.267	−23.752	(FU1)
191	13.6	3.8	1270.386 26	−19.122	263.603	−25.026	Strong bin-centered line in L1 at 1270.3 Hz
194	13.6	3.8	1318.615 37	−2.694	263.896	−27.405	Line in L1 at 1318.6 Hz ???
195	13.6	3.9	1262.658 32	−14.294	266.151	−27.129	(FU0)
199	13.5	3.9	1212.897 60	−9.841	262.769	−25.319	(FU1)
207	13.5	4.0	1202.298 41	−27.188	263.136	−26.481	Strong bin-centered line in L1 at 1202.2 Hz
214	13.5	4.1	1232.091 53	−18.288	268.914	−24.416	Strong bin-centered line in H1 at 1232 Hz
218	13.5	4.1	1270.386 24	−19.118	263.585	−25.347	Strong bin-centered line in L1 at 1270.3 Hz
220	13.4	4.1	1313.365 32	0.507	266.019	−23.381	(FU0)
236	13.3	4.3	1242.801 88	0.099	263.747	−28.690	(FU0)
238	13.3	4.3	1210.064 45	−0.084	264.836	−26.046	
241	13.3	4.3	1317.506 08	−14.391	267.541	−28.868	(FU0)
246	13.3	4.3	1280.917 94	−20.272	265.138	−29.732	(FU1)
247	13.3	4.4	1222.091 62	−9.369	268.642	−26.238	Strong bin-centered line in L1 at 1222 Hz
255	13.3	4.4	1253.011 07	−5.959	265.096	−25.776	
257	13.3	4.4	1323.778 88	0.764	268.106	−29.354	(FU0)
262	13.3	4.4	1253.011 07	−5.964	265.098	−25.512	(FU0)
263	13.3	4.4	1261.986 90	−6.648	265.929	−24.334	(FU0)
269	13.3	4.5	1329.921 57	−2.304	264.893	−25.113	(FU0)
271	13.2	4.5	1371.332 76	−6.367	267.733	−27.949	(FU0)
274	13.2	4.5	1299.819 00	−3.076	265.372	−26.913	(FU0)
278	13.2	4.5	1289.140 48	−18.428	265.039	−27.429	(FU0)
283	13.2	4.5	1292.033 74	−18.191	263.987	−24.625	(FU0)
286	13.2	4.6	1298.510 21	−3.954	269.268	−26.330	(FU0)
288	13.2	4.6	1321.829 47	−26.962	267.202	−26.389	
289	13.2	4.6	1399.830 41	−29.214	264.793	−29.091	
293	13.2	4.6	1247.689 62	−21.128	265.493	−23.383	Strong bin-centered line in L1 at 1247.6 Hz (FU0)
294	13.2	4.6	1225.387 54	−28.327	264.042	−25.595	(FU0)
298	13.2	4.6	1359.459 10	−2.732	266.847	−26.802	(FU0)
308	13.1	4.7	1282.331 63	−13.246	263.533	−26.441	(FU0)
322	13.1	4.8	1225.722 67	−17.528	263.821	−25.398	(FU0)
327	13.1	4.8	1249.951 09	−23.121	266.789	−27.051	(FU1)
331	13.1	4.8	1297.000 50	−1.628	267.856	−26.736	(FU0)
332	13.1	4.8	1215.840 97	−17.271	269.914	−27.022	(FU0)
333	13.0	4.9	1295.239 73	−11.871	266.940	−28.091	(FU0)
334	13.0	4.9	1306.384 13	−18.809	265.273	−22.916	(FU0)
341	13.0	4.9	1345.967 37	−26.084	263.623	−26.514	(FU1)
343	13.0	4.9	1216.271 44	−10.614	267.855	−25.711	(FU0)
344	13.0	4.9	1259.455 64	−9.417	267.383	−28.047	(FU0)
345	13.0	4.9	1235.245 32	−0.544	263.166	−27.175	(FU1)
347	13.0	4.9	1308.952 61	−3.621	264.053	−25.778	(FU0)
348	13.0	4.9	1324.989 25	−27.563	267.805	−28.168	(FU0)
349	13.0	4.9	1272.224 05	−27.694	264.174	−27.709	
351	13.0	4.9	1334.419 73	−22.024	262.929	−28.009	(FU0)

TABLE V. Outliers above 1400 Hz that passed the detection pipeline excluding regions heavily contaminated with violin modes. Outliers marked with "line" had strong narrow-band disturbances identified near the outlier location. We have marked outliers not consistent with the target signals at one of the semicoherent $\mathcal{F}$-statistic follow-ups with "(FU0/1/2)", depending on the stage at which they did not pass the detection thresholds. Frequencies are converted to epoch GPS 1130529362.

Idx	SNR	$\log_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	Description
2	27.8	−26.2	1457.987 71	−5.923	267.265	−23.236	Broad line in L1
3	27.5	−25.5	1495.877 36	−19.359	263.979	−25.411	Broad line in L1
12	21.6	−12.6	1469.414 04	−12.041	267.020	−24.399	Strong bin-centered line in L1 at 1469.3 Hz
15	20.9	−11.0	1467.552 57	−19.926	265.951	−29.741	Broad disturbance in H1 (?)
18	20.2	−9.6	1469.441 00	−1.796	269.546	−24.514	Strong bin-centered line in L1 at 1469.3 Hz
21	19.1	−7.1	1421.116 33	−6.941	263.048	−26.468	Strong bin-centered line in L1 at 1421 Hz
23	18.1	−5.1	1478.753 95	−0.269	268.825	−26.084	Broad line in H1
26	17.8	−4.4	1469.406 83	−14.406	263.178	−26.341	Strong bin-centered line in L1 at 1469.3 Hz
29	17.7	−4.3	1421.118 41	−6.437	264.313	−27.030	Strong bin-centered line in L1 at 1421 Hz
30	17.6	−4.2	1408.104 31	−25.129	266.189	−23.920	Strong bin-centered line in H1 at 1408 Hz
35	17.4	−3.6	1418.202 92	−20.386	267.609	−24.774	Strong bin-centered line in L1 at 1418.1 Hz
36	17.1	−3.1	1421.118 41	−6.436	264.313	−27.043	Strong bin-centered line in L1 at 1421 Hz
41	16.8	−2.5	1484.505 12	−6.777	269.295	−26.268	Broad line in H1
53	16.2	−1.2	1467.473 64	−1.878	268.521	−27.384	Broad line in H1
60	15.5	0.2	1478.752 26	−0.909	267.825	−27.981	Broad line in H1
62	15.4	0.3	1499.440 86	−7.311	263.883	−25.493	Nearby broad lines in H1 and L1
73	15.0	1.1	1484.756 42	−29.816	265.225	−26.457	Broad line in H1
74	15.0	1.1	1401.698 54	−22.076	264.511	−26.586	(FU1)
88	14.7	1.7	1458.958 08	−15.962	268.240	−26.171	
90	14.7	1.8	1400.733 83	−8.001	268.040	−25.138	(FU1)
92	14.6	1.8	1492.341 25	−8.216	266.388	−23.361	Broad line in L1
94	14.6	1.9	1497.841 01	−4.821	263.660	−28.526	Broad line in L1
100	14.5	2.1	1484.487 10	−0.763	267.286	−27.112	Broad lines in H1
132	14.0	3.1	1443.818 41	−9.151	268.034	−25.533	Strong bin-centered line in L1 at 1443.7 Hz
136	14.0	3.1	1497.841 02	−4.823	263.665	−28.451	Broad line in L1
154	13.8	3.4	1498.674 20	−19.913	264.048	−23.412	Broad line in H1
167	13.7	3.5	1472.023 03	−12.624	263.554	−25.278	Broad line in L1
183	13.6	3.7	1402.870 61	−20.561	269.872	−26.485	
193	13.6	3.8	1457.636 10	−4.906	263.216	−24.295	Broad line in L1
196	13.6	3.9	1454.893 91	−20.218	268.541	−25.625	
198	13.5	3.9	1442.617 15	−16.934	267.499	−25.456	
200	13.5	3.9	1462.091 24	−28.426	264.860	−27.360	Broad line in H1
201	13.5	3.9	1488.982 02	−9.992	265.988	−27.058	(FU0)
204	13.5	3.9	1443.500 58	−8.113	268.399	−26.734	(FU0)
205	13.5	4.0	1496.135 32	−8.503	264.346	−26.994	Broad line in L1
209	13.5	4.0	1499.430 21	−15.954	264.590	−25.146	Nearby broad lines in H1 and L1
216	13.5	4.1	1498.677 34	−21.961	267.824	−29.502	Broad line in H1
224	13.4	4.2	1499.428 49	−16.246	263.688	−28.354	Nearby broad lines in H1 and L1
226	13.4	4.2	1489.976 25	−29.213	269.022	−24.757	(FU0)
230	13.4	4.2	1465.368 84	−1.351	266.338	−23.042	(FU0)
231	13.4	4.3	1408.128 49	−2.961	264.812	−28.640	Strong bin-centered line in H1 at 1408 Hz
235	13.3	4.3	1482.006 82	−4.574	263.839	−26.682	Disturbed spectrum in H1
239	13.3	4.3	1408.107 86	−17.896	270.286	−26.385	Strong bin-centered line in H1 at 1408 Hz
250	13.3	4.4	1459.401 66	−29.472	269.924	−27.156	(FU0)
252	13.3	4.4	1460.139 53	−25.686	267.954	−26.966	(FU0)
256	13.3	4.4	1408.110 02	−10.656	268.386	−27.919	Strong bin-centered line in H1 at 1408 Hz
259	13.3	4.4	1482.623 19	−22.633	264.730	−27.878	Nearby broad line in H1
268	13.3	4.5	1445.848 29	−16.384	268.512	−28.667	(FU0)

(Table continued)

TABLE V. *(Continued)*

Idx	SNR	log$_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	Description
275	13.2	4.5	1474.296 69	−28.242	265.215	−26.055	Nearby broad line in H1, disturbed H1 spectrum
277	13.2	4.5	1462.208 88	−19.916	268.963	−25.106	(FU0)
296	13.2	4.6	1478.107 69	−26.661	265.631	−22.822	Bin-centered line in L1 at 1478 Hz, disturbed H1 spectrum
301	13.2	4.6	1459.945 53	−3.364	265.599	−29.328	(FU0)
302	13.2	4.6	1455.788 04	−1.216	263.403	−28.173	(FU1)
312	13.1	4.7	1472.966 26	−6.889	263.067	−25.410	(FU1)
320	13.1	4.8	1432.487 86	−23.659	267.769	−28.544	(FU0)
321	13.1	4.8	1496.769 07	−17.039	262.628	−26.112	(FU0)
324	13.1	4.8	1499.427 98	−18.013	268.803	−25.498	Nearby broad lines in H1 and L1
326	13.1	4.8	1475.089 10	−24.164	268.123	−29.057	Nearby strong line in H1, disturbed spectrum
328	13.1	4.8	1430.124 80	−23.206	266.342	−24.490	
329	13.1	4.8	1430.124 80	−23.204	266.342	−24.625	(FU0)
346	13.0	4.9	1423.419 85	−2.116	264.904	−26.065	(FU0)
352	13.0	4.9	1499.108 49	−0.374	267.745	−26.798	Broad line in L1

TABLE VI.   Outliers in 495–520 Hz and 990–1033 Hz regions heavily contaminated with violin modes. Outliers marked with "line" had strong narrow-band disturbances identified near the outlier location. We have marked outliers not consistent with the target signals at one of the semicoherent $\mathcal{F}$-statistic follow-ups with "(FU0/1/2)", depending on the stage at which they did not pass the detection thresholds. Frequencies are converted to epoch GPS 1130529362.

Idx	SNR	log$_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	
4	27.3	−25.0	508.260 88	−9.104	269.109	−29.202	Broad line in H1 at 508.222
7	23.6	−16.9	1030.758 07	−24.288	265.630	−28.137	Forest of strong lines in L1
9	23.2	−16.0	501.548 72	−22.486	264.981	−23.309	Large line in H1, violin mode region
10	22.7	−14.9	1018.714 52	−13.869	268.055	−23.245	Strong line in L1
11	22.0	−13.4	505.620 39	−6.776	264.643	−27.775	Large lines in H1, violin mode region
14	21.0	−11.2	1027.446 09	−28.797	266.018	−29.859	Forest of strong lines in L1
22	18.9	−6.7	1014.132 65	−5.347	269.223	−28.565	Forest of strong lines in L1
25	17.9	−4.7	1030.760 62	−27.591	269.921	−27.663	Forest of strong lines in L1
27	17.7	−4.4	505.633 42	−0.324	266.622	−29.454	Large lines in H1, violin mode region
28	17.7	−4.3	505.683 86	−20.164	267.657	−25.001	Large lines in H1, violin mode region
31	17.6	−4.1	1008.586 25	0.313	263.517	−27.178	Strong broad line in H1, line in L1
38	17.0	−2.9	505.721 51	−16.584	265.219	−27.670	Large line in H1, violin mode region
46	16.5	−1.9	1006.003 95	−14.306	267.233	−27.920	Strong broad lines in H1
48	16.5	−1.9	1021.203 75	−22.584	267.795	−24.931	Lines in L1
49	16.4	−1.7	509.197 31	0.039	266.911	−29.803	Violin mode region
50	16.3	−1.5	1031.088 95	−27.254	263.465	−24.793	Lines in L1
51	16.3	−1.5	506.973 95	−12.404	265.161	−27.594	Large line in H1, violin mode region
52	16.3	−1.4	1027.539 60	−15.576	269.243	−24.104	Forest of strong lines in L1
54	15.8	−0.6	509.196 26	−0.289	264.906	−26.767	Violin mode region
57	15.6	−0.2	1017.170 41	−15.632	263.374	−24.850	
61	15.4	0.3	1029.137 74	−21.237	263.826	−28.459	Forest of strong lines in L1 (FU0)
64	15.3	0.5	505.728 44	−13.179	267.146	−23.848	Large broad and narrow lines in H1, L1, violin mode region
65	15.3	0.5	1027.534 47	−15.899	264.371	−27.442	Forest of strong lines in L1
66	15.3	0.6	1014.135 50	−0.511	269.734	−28.397	Forest of strong lines in L1
71	15.1	0.9	992.021 21	−22.994	269.830	−28.298	Strong broad line in H1, lines in L1
79	14.9	1.4	503.010 53	−14.412	264.576	−29.521	Large lines in H1, violin mode region

(Table continued)

TABLE VI. *(Continued)*

Idx	SNR	log$_{10}$ (GFA)	Frequency Hz	Spindown nHz/s	RA$_{J2000}$ degrees	DEC$_{J2000}$ degrees	
83	14.8	1.5	1006.180 12	−21.694	263.313	−24.036	Strong bin-centered line in L1 at 1006.1 Hz
87	14.7	1.6	1030.760 44	−17.694	264.609	−25.125	Forest of strong lines in L1
95	14.6	1.9	1003.748 67	−4.776	265.096	−24.430	Strong broad line in H1
97	14.6	2.0	1013.039 47	−0.404	268.508	−28.026	Disturbed background in H1
98	14.5	2.0	1026.144 62	−28.021	267.878	−27.584	Broad line in L1
105	14.4	2.3	1004.025 92	−2.816	265.980	−25.168	(FU0)
108	14.3	2.4	1004.025 91	−2.814	265.980	−25.265	(FU0)
109	14.3	2.4	1003.736 14	−16.516	266.800	−28.487	Strong broad line in H1
110	14.3	2.5	1029.196 72	1.012	270.018	−27.237	Forest of lines in L1
122	14.1	2.8	993.557 83	−25.844	265.053	−24.977	(FU0)
124	14.1	2.9	511.996 12	−0.903	266.165	−25.709	Sharp bin-centered line at 512 Hz
131	14.0	3.0	1000.066 19	−13.539	266.692	−26.396	Lines in H1 and L1
134	14.0	3.1	1011.005 09	−18.583	269.680	−25.059	Strong line in L1, disturbed H1 spectrum
153	13.8	3.4	1003.752 96	−3.369	265.458	−27.670	Strong broad line in H1
155	13.8	3.4	1025.636 09	−7.339	265.091	−23.025	Forest of lines in L1
157	13.8	3.4	994.505 37	−0.149	264.616	−28.649	(FU0)
164	13.8	3.5	990.288 98	−0.769	264.972	−26.915	Disturbed H1 spectrum
168	13.7	3.5	1031.395 56	−7.501	263.845	−23.332	Forest of lines in L1
178	13.7	3.7	1010.615 05	−17.337	265.335	−24.634	Disturbed H1 spectrum, lines in L1
184	13.6	3.8	991.260 21	−23.166	266.511	−27.602	Disturbed H1 spectrum, lines
185	13.6	3.8	1006.003 96	−14.314	267.225	−27.634	Strong broad lines in H1, lines in L1
187	13.6	3.8	1031.161 38	−14.838	264.511	−23.556	Forest of lines in L1
197	13.6	3.9	509.198 95	0.414	269.633	−27.167	Large lines, violin mode region
212	13.5	4.0	1017.383 48	−16.619	268.845	−24.206	(FU0)
217	13.5	4.1	1018.829 20	−17.391	264.399	−26.832	(FU0)
221	13.4	4.1	1011.028 16	−16.411	267.760	−22.824	Disturbed H1 spectrum, lines in L1
222	13.4	4.1	1016.775 36	−0.176	267.870	−23.024	(FU0)
225	13.4	4.2	993.089 01	−25.739	269.788	−25.063	(FU1)
234	13.4	4.3	992.256 22	−22.391	265.644	−27.798	Forest of lines in H1 and L1
240	13.3	4.3	1005.021 79	−22.107	269.152	−28.822	(FU0)
280	13.2	4.5	1031.491 41	−25.294	266.685	−27.054	Forest of lines in L1
281	13.2	4.5	1013.136 66	−22.974	268.774	−25.644	Large disturbance in H1
292	13.2	4.6	997.434 74	−5.569	266.898	−26.932	Strong bin-centered line in L1 at 997.4, disturbed H1 spectrum
304	13.2	4.6	1002.788 26	−1.404	268.608	−27.280	Strong line in L1
306	13.1	4.7	1000.062 13	−15.341	264.654	−27.047	Lines in H1 and L1 (FU0)
315	13.1	4.7	991.620 69	−12.746	263.013	−27.125	Strong broad line in L1

[1] K. Riles, Recent searches for continuous gravitational waves, Mod. Phys. Lett. A **32**, 1730035 (2017).

[2] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), First low-frequency Einstein@Home all-sky search for continuous gravitational waves in advanced LIGO data, Phys. Rev. D **96**, 122004 (2017).

[3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search for periodic gravitational waves in the O1 LIGO data, Phys. Rev. D **96**, 062002 (2017).

[4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Full band all-sky search for periodic gravitational waves in the O1 LIGO data, Phys. Rev. D **97**, 102003 (2018).

[5] V. Dergachev and M. A. Papa, First Loosely Coherent All-Sky Search for Periodic Gravitational Waves in the O1 LIGO Data, arXiv:1902.05530 [Phys. Rev. Lett. (to be published)].

[6] J. Aasi *et al.* (LIGO Scientific and Virgo Collaborations), A search of the Orion spur for continuous gravitational waves

using a "loosely coherent" algorithm on data from LIGO interferometers, Phys. Rev. D **93,** 042006 (2016).

[7] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, The LIGO Open Science Center, J. Phys. Conf. Ser. **610,** 012021 (2015).

[8] LIGO Open Science Center, https://doi.org/10.7935/CA75-FM95, https://losc.ligo.org.

[9] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, Classical Quantum Gravity **32,** 074001 (2015).

[10] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW150914: The Advanced LIGO Detectors in the Era of First Discoveries, Phys. Rev. Lett. **116,** 131103 (2016).

[11] V. Dergachev, On blind searches for noise dominated signals: A loosely coherent approach, Classical Quantum Gravity **27,** 205017 (2010).

[12] V. Dergachev, Loosely coherent searches for sets of well-modeled signals, Phys. Rev. D **85,** 062003 (2012).

[13] V. Dergachev, Efficient loosely coherent searches for medium scale coherence lengths, arXiv:1807.02351 [Phys. Rev. D (to be published)].

[14] N. K. Johnson-McDaniel and B. J. Owen, Maximum elastic deformations of relativistic stars, Phys. Rev. D **88,** 044004 (2013).

[15] P. Covas *et al.*, Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO, Phys. Rev. D **97,** 082002 (2018).

[16] B. Abbott *et al.* (LIGO Scientific Collaboration), All-sky search for periodic gravitational waves in LIGO S4 data, Phys. Rev. D **77,** 022001 (2008).

[17] B. P. Abbott *et al.* (LIGO Scientific Collaboration), All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data, Phys. Rev. Lett. **102,** 111102 (2009).

[18] B. Abbott *et al.* (The LIGO and Virgo Scientific Collaboration), All-sky search for periodic gravitational waves in the full S5 data, Phys. Rev. D **85,** 022001 (2012).

[19] V. Dergachev, A novel universal statistic for computing upper limits in ill-behaved background, Phys. Rev. D **87,** 062001 (2013).

[20] A. G. Lyne, S. H. Mankelow, J. F. Bell, and R. N. Manchester, Radio pulsars in Terzan 5, Mon. Not. R. Astron. Soc. **316,** 491 (2000).

[21] M. Cadelano, S. M. Ransom, P. C. C. Freire, F. R. Ferraro, J. W. T. Hessels, B. Lanzoni, C. Pallanca, and I. H. Stairs, Discovery of three new millisecond pulsars in Terzan 5, Astrophys. J. **855,** 125 (2018).

[22] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The ATNF pulsar catalogue, Astrophys. J. **129,** 1993 (2005).

[23] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett. **116,** 061102 (2016).

[24] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevD.99.084048 for numerical values of upper limits.

[25] ATNF Pulsar Catalogue, http://www.atnf.csiro.au/research/pulsar/psrcat/.

[26] B. Behnke, M. A. Papa, and R. Prix, Postprocessing methods used in the search for continuous gravitational-wave signals from the Galactic Center, Phys. Rev. D **91,** 064007 (2015).

[27] C. Dreissigacker, R. Prix, and K. Wette, Fast and accurate sensitivity estimation for continuous-gravitational-wave searches, Phys. Rev. D **98,** 084058 (2018).

[28] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Directional Limits on Persistent Gravitational Waves from Advanced LIGO First Observing Run, Phys. Rev. Lett. **118,** 121102 (2017).

[29] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search for continous gravitational waves from isolated neutron stars using Advanced LIGO O2 data, arXiv:1903.01901.

112

Conclusion and Outlook

As the recent history of gravitational wave astronomy shows, once the first gravitational wave event is detected, a plethora of others follow, and the new field blooms. This first breakthrough has yet to be made for continuous gravitational waves. But the current efforts on the experimental side and on the analysis side will bear fruits, and the first detection of a continuous gravitational wave will happen. After that, more will come, and with the ability to observe them continuously, we will leverage an increase in the signal-to-noise ratio with every new observation run / detector, which opens up possibilities for exciting new science.

Detections of continuous waves will help to reduce the many uncertainties in the make-up of neutron stars and matter in these extreme conditions not testable in our laboratories. Continuous gravitational waves give us access to an invisible population of neutron stars, and broaden our understanding of fundamental physics.

In this thesis multiple important steps towards a detection were addressed and applied, from the initial data preparation of the calibrated gravitational wave detector data, to preparing and conducting searches for unknown isolated neutron stars in our galaxy, to crucial postprocessing methods weeding out billions of candidates.

Non-detections are informative: a null result restricts the maximum possible deformation of the sources. The ellipticities probed by the various searches presented here are physically plausible. Future more sensitive searches will constrain the ellipticity even more and could potentially rule out some models of the neutron star interior, or lead to a detection.

While the detector data becomes more and more sensitive, we are also improving our methods and searches, of which several were presented in this thesis:

The data preparation method in Chapter 2 is a significant development over a compact binary coalescence gating method [99] used to prepare the data for [12]. Being a method conceived for searches for short-lived signals, it was inadequate for continuous wave searches. The method that I developed was applied to O1, O2 and O3 data, and most of the group's continuous wave searches of Advanced LIGO data used in fact data that I had conditioned, [25, 37, 44, 68, 71, 72, 77, 91, 105, 106].

The method is easy to use and self-adapting to the data, as showcased on the results of the third observation run: With minimal tuning and within a week of the data being released, the method produced similar results to the "self-gated" data as prepared by LIGO, which took them a considerable time to prepare [100, 109].

The framework was further developed to prepare data for any search, i.e. using Short Fourier Transforms (SFTs) of arbitrary time-baseline within minimal setup time. Thanks to this improvement, the data conditioning is now being used also for

binary searches, streamlining the data preparation process significantly, across the board.

The gating method might be further improved: Although most glitches have support over a large range of frequencies, there are classes of glitches where the power is concentrated in certain regions. So far the full bandwidth is used in the whitening process, whereas a frequency-dependent identification and removal might lead to a more effective noise mitigation.

Chapter 3 presents a new clustering method, which is an important post-processing step. Large-scale Einstein@Home searches return billions of candidates but not all candidates are independent: disturbances and signals can give rise to thousands of related candidates. Clustering finds and groups together, *clusters*, candidates due to the same cause. Each cluster conveniently summarizes the properties of the entire set. The morphologies of noise and signal clusters vary, thus through clustering we can veto a high number of noise candidates. The method is crucial because it reduces the number of candidates to follow up to a manageable set (of, say, some million clusters,) without dismissing candidates due to signals.

The method is very innovative with respect to previous clustering methods, which emphasize the detection statistic value. In contrast I showed that for weak signals buried in noise the detection statistic values become less meaningful: there are millions of noise candidates with "more significant" detection statistic values. I showed that in this regime, the overdensity of candidates, compared to the background, is a more effective indicator of the presence of a signal. In the regime of strong signals my approach has no downside, since the signals featuring moderate to high detection statistic values also come with many nearby associated candidates. My density clustering provides a 30-fold smaller follow-up at fixed sensitivity, compared to previous methods, and a 10% sensitivity improvement at fixed computing cost.

The density clustering method was applied to the Einstein@Home all-sky search in O2 data (Chapter 4 [93]), the directed search for G347.3 (Chapter 8 [66]), a few directed searches in preparation, and the Einstein@Home all-sky search in O3 data ([94, in preparation]).

Perhaps the most interesting application is an ongoing one, namely to the results of all-sky searches for signals from neutron stars in binary systems. The parameter space for these signals has higher dimensions compared to the isolated signals, and where this is a problem for previous clustering methods, the performance of density clustering excels in higher dimensions.

Chapter 4 showcases how everything must come together for the most difficult source of gravitational wave astronomy: an all-sky search for continuous gravitational waves. The eight-stage hierarchical follow-up scheme examines over $350\,145$ candidates and retains only six, which are the fake continuous signals in the data stream for validation purposes. The upper limits on the gravitational wave amplitude h_0 and neutron star ellipticity ε are more stringent compared to upper limits from any other search on the same waveform space. The results of this search are used in Singh and Papa [91] to constrain the gravitational wave amplitude over a large and unexplored range of binary orbital parameters, namely objects in long-periodicity binaries. These would still appear as a signal in Stage 0 of the Einstein@Home search, albeit an isolated search,

The searches presented in this thesis rule out large regions of gravitational wave amplitude and neutron star equatorial ellipticities: the Einstein@Home O2 all-sky

search constrains isolated neutron stars within 100 pc spinning faster than 12 ms with ellipticities in the few 10^{-7} and faster than 5 ms with $\varepsilon \geqslant 10^{-7}$. In the regions of higher abundance of neutron stars, ellipticities of $\varepsilon \sim 10^{-6}$ (Galactic Center) and $\varepsilon \leqslant 7 \times 10^{-7}$ (Terzan 5) are ruled out for isolated neutron stars. The searches for known or assumed neutron stars constrained the continuous wave amplitude h_0 between 7.7×10^{-27} and 1.5×10^{-25}, translating to constraints on the ellipticities at the $\varepsilon \sim 10^{-8}$ to $\sim 10^{-4}$ level.

Most of the searches presented in this thesis are the most sensitive searches for their respective targets in their data set and parameter space. The probed ellipticities could, according to some models, be supported by the neutron star crust. There is, however, much uncertainty on the actual values of the equatorial ellipticity. According to Johnson-McDaniel and Owen [56] the neutron star crust could support ellipticities of around 10^{-5}, or in the 10^{-6} regime according to [69, 98], whereas other models find much lower ellipticities around 10^{-9} [47]. Woan et al. [104] suggest that millisecond pulsars might have a *minimum* ellipticity of $\sim 10^{-9}$. The searches presented are therefore probing interesting regions of ellipticities, and demonstrate the impact of novel techniques. It is exactly through the constant pursuit of improvement that searches will eventually enable the detection of continuous gravitational waves. The work presented in this thesis contributes to this.

The past is written, but the future is left for us to write. And we have powerful tools, openness, optimism, and the spirit of curiosity.

Jean-Luc Picard

[1] J. Abadie et al. Beating the spin-down limit on gravitational wave emission from the Vela pulsar. *Astrophys. J.*, 737:93, 2011. DOI: 10.1088/0004-637X/737/2/93.

[2] B. Abbott et al. Beating the spin-down limit on gravitational wave emission from the Crab pulsar. *Astrophys. J. Lett.*, 683:L45–L50, 2008. DOI: 10.1086/591526. [Erratum: Astrophys.J.Lett. 706, L203–L204 (2009)].

[3] B. Abbott et al. The Einstein@Home search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 79:022001, 2009. DOI: 10.1103/PhysRevD.79.022001.

[4] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016. DOI: 10.1103/PhysRevLett.116.061102.

[5] B. P. Abbott et al. Tests of general relativity with GW150914. *Phys. Rev. Lett.*, 116(22):221101, 2016. DOI: 10.1103/PhysRevLett.116.221101. [Erratum: Phys.Rev.Lett. 121, 129902 (2018)].

[6] B. P. Abbott et al. A gravitational-wave standard siren measurement of the Hubble constant. *Nature*, 551(7678):85–88, 2017. DOI: 10.1038/nature24471.

[7] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017. DOI: 10.1103/PhysRevLett.119.161101.

[8] B. P. Abbott et al. Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophys. J. Lett.*, 848(2):L12, 2017. DOI: 10.3847/2041-8213/aa91c9.

[9] B. P. Abbott et al. Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *Astrophys. J. Lett.*, 848 (2):L13, 2017. DOI: 10.3847/2041-8213/aa920c.

[10] B. P. Abbott et al. Tests of General Relativity with the Binary Black Hole Signals from the LIGO-Virgo Catalog GWTC-1. *Phys. Rev. D*, 100(10):104036, 2019. DOI: 10.1103/PhysRevD.100.104036.

[11] B. P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_\odot$. *Astrophys. J. Lett.*, 892(1):L3, 2020. DOI: 10.3847/2041-8213/ab75f5.

[12] Benjamin P. Abbott et al. First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data. *Phys. Rev. D*, 96 (12):122004, 2017. DOI: `10.1103/PhysRevD.96.122004`.

[13] R. Abbott et al. All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems. *Phys. Rev. D*, 103(6):064017, 2021. DOI: `10.1103/PhysRevD.103.064017`.

[14] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run. *arXiv e-prints*, page arXiv:2111.03606, 11 2021.

[15] R. Abbott et al. Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences. *Astrophys. J. Lett.*, 915(1):L5, 2021. DOI: `10.3847/2041-8213/ac082e`.

[16] R. Abbott et al. Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data. *Phys. Rev. D*, 106(4):042003, 2022. DOI: `10.1103/PhysRevD.106.042003`.

[17] R. Abbott et al. Searches for Gravitational Waves from Known Pulsars at Two Harmonics in the Second and Third LIGO-Virgo Observing Runs. *Astrophys. J.*, 935(1):1, 2022. DOI: `10.3847/1538-4357/ac6acf`.

[18] R. Abbott et al. Narrowband Searches for Continuous and Long-duration Transient Gravitational Waves from Known Pulsars in the LIGO-Virgo Third Observing Run. *Astrophys. J.*, 932(2):133, 2022. DOI: `10.3847/1538-4357/ac6ad0`.

[19] M. Ali Alpar and Altan Baykal. Pulsar Braking Indices, Glitches and Energy Dissipation In Neutron Stars. *Mon. Not. Roy. Astron. Soc.*, 372:489, 2006. DOI: `10.1111/j.1365-2966.2006.10893.x`.

[20] Bruce Allen. Spherical ansatz for parameter-space metrics. *Phys. Rev. D*, 100 (12):124004, 2019. DOI: `10.1103/PhysRevD.100.124004`.

[21] D. P. Anderson. BOINC: A System for Public-Resource Computing and Storage. In *Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing (GRID04)*, pages 4–10, 2004.

[22] D. P. Anderson, C. Christensen, and B. Allen. Designing a Runtime System for Volunteer Computing. In *Proceedings of the 2006 ACM/IEEE conference on Supercomputing*, pages 126–136, 2006.

[23] Georgios C. Antonopoulos, Mirko S. Rakoski, Benjamin Steltner, Stefan Kalies, Tammo Ripken, and Heiko Meyer. Experimental setup combining digital holographic microscopy (DHM) and fluorescence imaging to study gold nanoparticle mediated laser manipulation. In Gabriel Popescu and YongKeun Park, editors, *Quantitative Phase Imaging*, volume 9336 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 93360G, March 2015. DOI: `10.1117/12.2079322`.

[24] Georgios C. Antonopoulos, Benjamin Steltner, Alexander Heisterkamp, Tammo Ripken, and Heiko Meyer. Tile-Based Two-Dimensional Phase Unwrapping for Digital Holography Using a Modular Framework. *PLoS ONE*, 10(11):e0143186, November 2015. DOI: `10.1371/journal.pone.0143186`.

[25] Anjana Ashok, Banafsheh Beheshtipour, Maria Alessandra Papa, Paulo C. C. Freire, Benjamin Steltner, Bernd Machenschalk, Oliver Behnke, Bruce Allen, and Reinhard Prix. New Searches for Continuous Gravitational Waves from Seven Fast Pulsars. *Astrophys. J.*, 923(1):85, 2021. DOI: `10.3847/1538-4357/ac2582`.

[26] D. Bhattacharya and E. P. J. van den Heuvel. Formation and evolution of binary and millisecond radio pulsars. *Physics Reports*, 203(1-2):1–124, January 1991. DOI: `10.1016/0370-1573(91)90064-S`.

[27] Gerhard Bohm and Guenter Zech. *Introduction to statistics and data analysis for physicists*. DESY, Hamburg, 2014. ISBN 978-3-935702-88-1. DOI: `10.3204/DESY-BOOK/statistics`.

[28] BOINC. `http://boinc.berkeley.edu/`, 2020.

[29] Patrick R. Brady and Teviet Creighton. Searching for periodic sources with LIGO. 2. Hierarchical searches. *Phys. Rev. D*, 61:082001, 2000. DOI: `10.1103/PhysRevD.61.082001`.

[30] Patrick R. Brady, Teviet Creighton, Curt Cutler, and Bernard F. Schutz. Searching for periodic sources with LIGO. *Phys. Rev. D*, 57:2101–2116, 1998. DOI: `10.1103/PhysRevD.57.2101`.

[31] Aaron Buikema et al. Sensitivity and performance of the Advanced LIGO detectors in the third observing run. *Phys. Rev. D*, 102(6):062003, 2020. DOI: `10.1103/PhysRevD.102.062003`.

[32] Craig Cahillane and Georgia Mansell. Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four. *Galaxies*, 10(1):36, 2022. DOI: `10.3390/galaxies10010036`.

[33] ATNF Pulsar Catalogue. `http://www.atnf.csiro.au/research/pulsar/psrcat`, 2022.

[34] Jorge L. Cervantes-Cota, S. Galindo-Uribarri, and G-F. Smoot. A Brief History of Gravitational Waves. *Universe*, 2(3):22, 2016. DOI: `10.3390/universe2030022`.

[35] The ATLAS computing cluster. `https://www.aei.mpg.de/346572/atlas-computing-cluster`, 2022.

[36] P. B. Covas and Alicia M. Sintes. First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data. *Phys. Rev. Lett.*, 124(19):191102, 2020. DOI: `10.1103/PhysRevLett.124.191102`.

[37] P. B. Covas, M. A. Papa, R. Prix, and B. J. Owen. Constraints on r-modes and Mountains on Millisecond Neutron Stars in Binary Systems. *Astrophys. J. Lett.*, 929(2):L19, 2022. DOI: `10.3847/2041-8213/ac62d7`.

[38] Thayne Currie, John Debes, Timothy J. Rodigas, Adam Burrows, Yoichi Itoh, Misato Fukagawa, Scott Kenyon, Marc Kuchner, and Soko Matsumura. Direct Imaging Confirmation and Characterization of a Dust-Enshrouded Candidate Exoplanet Orbiting Fomalhaut. *Astrophys. J. Lett.*, 760:L32, 2012. DOI: `10.1088/2041-8205/760/2/L32`.

[39] Curt Cutler, Iraj Gholami, and Badri Krishnan. Improved stack-slide searches for gravitational-wave pulsars. *Phys. Rev. D*, 72:042004, 2005. DOI: `10.1103/PhysRevD.72.042004`.

[40] Vladimir Dergachev, Maria Alessandra Papa, Benjamin Steltner, and Heinz-Bernd Eggenstein. Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the galactic center. *Phys. Rev. D*, 99(8):084048, 2019. DOI: `10.1103/PhysRevD.99.084048`.

[41] G. Desvignes et al. High-precision timing of 42 millisecond pulsars with the European Pulsar Timing Array. *Mon. Not. Roy. Astron. Soc.*, 458(3):3341–3380, 2016. DOI: `10.1093/mnras/stw483`.

[42] Irina Dvorkin, Jean-Philippe Uzan, Elisabeth Vangioni, and Joseph Silk. Exploring stellar evolution with gravitational-wave observations. *Mon. Not. Roy. Astron. Soc.*, 479(1):121–129, 2018. DOI: `10.1093/mnras/sty1414`.

[43] Albert Einstein. The Foundation of the General Theory of Relativity. *Annalen Phys.*, 49(7):769–822, 1916. DOI: `10.1002/andp.19163540702`.

[44] Liudmila Fesik and Maria Alessandra Papa. First Search for r-mode Gravitational Waves from PSR J0537–6910. *Astrophys. J.*, 895(1):11, 2020. DOI: `10.3847/1538-4357/ab8193`. [Erratum: Astrophys.J. 897, 185 (2020)].

[45] Tobin T. Fricke et al. DC readout experiment in Enhanced LIGO. *Class. Quant. Grav.*, 29:065005, 2012. DOI: `10.1088/0264-9381/29/6/065005`.

[46] Raphael Galicher, Christian Marois, Ben Zuckerman, and Bruce Macintosh. Fomalhaut b: Independent Analysis of the Hubble Space Telescope Public Archive Data. *Astrophys. J.*, 769:42, 2013. DOI: `10.1088/0004-637X/769/1/42`.

[47] Fabian Gittins and Nils Andersson. Modelling neutron star mountains in relativity. *Mon. Not. Roy. Astron. Soc.*, 507(1):116–128, 2021. DOI: `10.1093/mnras/stab2048`.

[48] T. Gold. Rotating neutron stars as the origin of the pulsating radio sources. *Nature*, 218:731–732, 1968. DOI: `10.1038/218731a0`.

[49] P. Haensel, A. Y. Potekhin, and D. G. Yakovlev. *Neutron stars 1: Equation of state and structure*, volume 326. Springer, New York, USA, 2007. DOI: `10.1007/978-0-387-47301-7`.

[50] B. Haskell, L. Samuelsson, K. Glampedakis, and N. Andersson. Modelling magnetically deformed neutron stars. *Mon. Not. Roy. Astron. Soc.*, 385:531–542, 2008. DOI: 10.1111/j.1365-2966.2008.12861.x.

[51] A. Hewish, S. J. Bell, J. D. H Pilkington, P. F. Scott, and R. A. Collins. Observation of a rapidly pulsating radio source. *Nature*, 217:709–713, 1968. DOI: 10.1038/217709a0.

[52] S. Hild et al. DC-readout of a signal-recycled gravitational wave detector. *Class. Quant. Grav.*, 26:055012, 2009. DOI: 10.1088/0264-9381/26/5/055012.

[53] Wynn C. G. Ho and Craig O. Heinke. A neutron star with a carbon atmosphere in the cassiopeia a supernova remnant. *Nature*, 462(7269):71–73, 2009. DOI: 10.1038/nature08525. URL https://doi.org/10.1038/nature08525.

[54] Piotr Jaranowski and Andrzej Krolak. Gravitational-Wave Data Analysis. Formalism and Sample Applications: The Gaussian Case. *Living Rev. Rel.*, 8: 3, 2005.

[55] Piotr Jaranowski, Andrzej Krolak, and Bernard F. Schutz. Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection. *Phys. Rev.*, D58:063001, 1998. DOI: 10.1103/PhysRevD.58.063001.

[56] Nathan K. Johnson-McDaniel and Benjamin J. Owen. Maximum elastic deformations of relativistic stars. *Phys. Rev. D*, 88:044004, 2013. DOI: 10.1103/PhysRevD.88.044004.

[57] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. Roy. Astron. Soc.*, 331:203, 2002. DOI: 10.1046/j.1365-8711.2002.05180.x.

[58] Dana Jones and Ling Sun. Search for continuous gravitational waves from Fomalhaut b in the second Advanced LIGO observing run with a hidden Markov model. *Phys. Rev. D*, 103(2):023020, 2021. DOI: 10.1103/PhysRevD.103.023020.

[59] M. Kramer et al. Tests of general relativity from timing the double pulsar. *Science*, 314:97–102, 2006. DOI: 10.1126/science.1132305.

[60] James M. Lattimer. Neutron star structure and the equation of state. *Prog. Theor. Phys. Suppl.*, 186:1–8, 2010. DOI: 10.1143/PTPS.186.1.

[61] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. lscsoft.docs.ligo.org/lalsuite/, 2018.

[62] Bennett Link, Lucia M. Franco, and Richard I. Epstein. Starquake-induced magnetic field and torque evolution in neutron stars. *Astrophys. J.*, 508:838, 1998. DOI: 10.1086/306457.

[63] U. Lombardo and H. J. Schulze. Superfluidity in neutron star matter. *Lect. Notes Phys.*, 578:30–53, 2001.

[64] D. R. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*, volume 4. Cambridge University Press, Cambridge, UK, 2004.

[65] R N Manchester, G B Hobbs, A Teoh, and M Hobbs. The Australia Telescope National Facility pulsar catalogue. *Astron. J.*, 129:1993, 2005. DOI: `10.1086/428488`.

[66] J. Ming, M. A. Papa, A. Singh, H.-B. Eggenstein, S. J. Zhu, V. Dergachev, Y. Hu, R. Prix, B. Machenschalk, C. Beer, O. Behnke, and B. Allen. Results from an Einstein@Home search for continuous gravitational waves from Cassiopeia A, Vela Jr. and G347.3. *Phys. Rev.*, D100(2):024063, 2019. DOI: `10.1103/PhysRevD.100.024063`.

[67] Jing Ming, Maria Alessandra Papa, Badri Krishnan, Reinhard Prix, Christian Beer, Sylvia J. Zhu, Heinz-Bernd Eggenstein, Oliver Bock, and Bernd Machenschalk. Optimally setting up directed searches for continuous gravitational waves in Advanced LIGO O1 data. *Phys. Rev. D*, 97(2):024051, 2018. DOI: `10.1103/PhysRevD.97.024051`.

[68] Jing Ming, Maria Alessandra Papa, Heinz-Bernd Eggenstein, Bernd Machenschalk, Benjamin Steltner, Reinhard Prix, Bruce Allen, and Oliver Behnke. Results From an Einstein@Home Search for Continuous Gravitational Waves From G347.3 at Low Frequencies in LIGO O2 Data. *Astrophys. J.*, 925(1):8, 2022. DOI: `10.3847/1538-4357/ac35cb`.

[69] J. A. Morales and C. J. Horowitz. Neutron Star Crust Can Support A Large Ellipticity. *Monthly Notices of the Royal Astronomical Society*, 517(4):5610–5616, 9 2022. DOI: `10.1093/mnras/stac3058`.

[70] Jerzy Neyman and Egon Sharpe Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Phil. Trans. Roy. Soc. Lond. A*, 231(694-706):289–337, 1933. DOI: `10.1098/rsta.1933.0009`.

[71] L. Nieder, C. J. Clark, C. G. Bassa, J. Wu, A. Singh, J. Y. Donner, B. Allen, R. P. Breton, V. S. Dhillon, H. B. Eggenstein, J. W. T. Hessels, M. R. Kennedy, M. Kerr, S. Littlefair, T. R. Marsh, D. Mata Sánchez, M. A. Papa, P. S. Ray, B. Steltner, and J. P. W. Verbiest. Detection and timing of gamma-ray pulsations from the 707 Hz pulsar J0952−0607. *Astrophys. J.*, 883(1):42, 5 2019. DOI: `10.3847/1538-4357/ab357e`.

[72] L. Nieder et al. Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated Einstein@Home. *Astrophys. J. Lett.*, 902(2):L46, 2020. DOI: `10.3847/2041-8213/abbc02`.

[73] Alexander H. Nitz, Thomas Dent, Gareth S. Davies, Sumit Kumar, Collin D. Capano, Ian Harry, Simone Mozzon, Laura Nuttall, Andrew Lundgren, and Márton Tápai. 2-OGC: Open Gravitational-wave Catalog of binary mergers from analysis of public Advanced LIGO and Virgo data. *Astrophys. J.*, 891:123, 3 2020. DOI: `10.3847/1538-4357/ab733f`.

[74] Feryal Özel and Paulo Freire. Masses, Radii, and the Equation of State of Neutron Stars. *Ann. Rev. Astron. Astrophys.*, 54:401–440, 2016. DOI: `10.1146/annurev-astro-081915-023322`.

[75] F. Pacini. Energy emission from a neutron star. *Nature*, 216(5115):567–568, 1967. DOI: `10.1038/216567a0`. URL `https://doi.org/10.1038/216567a0`.

[76] M. A. Papa, Bernard F. Schutz, and A. M. Sintes. Searching for continuous gravitational wave signals: The hierarchical Hough transform algorithm. *ICTP Lect. Notes Ser.*, 3:431–442, 2001.

[77] M. Alessandra Papa, Jing Ming, Eric V. Gotthelf, Bruce Allen, Reinhard Prix, Vladimir Dergachev, Heinz-Bernd Eggenstein, Avneet Singh, and Sylvia J. Zhu. Search for Continuous Gravitational Waves from the Central Compact Objects in Supernova Remnants Cassiopeia A, Vela Jr., and G347.3–0.5. *Astrophys. J.*, 897(1):22, 2020. DOI: `10.3847/1538-4357/ab92a6`.

[78] Maria Alessandra Papa et al. Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data. *Phys. Rev. D*, 94(12):122006, 2016. DOI: `10.1103/PhysRevD.94.122006`.

[79] Holger J. Pletsch. Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection. *Phys. Rev. D*, 78:102005, 2008. DOI: `10.1103/PhysRevD.78.102005`.

[80] Holger J. Pletsch. Parameter-space metric of semicoherent searches for continuous gravitational waves. *Phys. Rev. D*, 82:042002, 2010. DOI: `10.1103/PhysRevD.82.042002`.

[81] Holger J. Pletsch and Bruce Allen. Exploiting global correlations to detect continuous gravitational waves. *Phys. Rev. Lett.*, 103:181102, 2009. DOI: `10.1103/PhysRevLett.103.181102`.

[82] Reinhard Prix. Gravitational Waves from Spinning Neutron Stars. *Astrophys. Space Sci. Lib.*, 357:651–685, 2009. DOI: `10.1007/978-3-540-76965-1_24`.

[83] Reinhard Prix and Miroslav Shaltev. Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost. *Phys. Rev. D*, 85:084010, 2012. DOI: `10.1103/PhysRevD.85.084010`.

[84] Brendan T. Reed, Alex Deibel, and C. J. Horowitz. Modeling the Galactic Neutron Star Population for Use in Continuous Gravitational-wave Searches. *Astrophys. J.*, 921(1):89, 2021. DOI: `10.3847/1538-4357/ac1c04`.

[85] Keith Riles. Searches for Continuous-Wave Gravitational Radiation, 6 2022.

[86] N. Sartore, E. Ripamonti, A. Treves, and R. Turolla. Galactic neutron stars I. Space and velocity distributions in the disk and in the halo. *Astron. Astrophys.*, 510:A23, 2010. DOI: `10.1051/0004-6361/200912222`.

[87] Bernard F. Schutz. *A First Course in General Relativity*. Cambridge Univ. Pr., Cambridge, UK, 1985.

[88] Bernard F. Schutz and M. Alessandra Papa. End-to-end algorithm for hierarchical area searches for long duration GW sources for GEO-600. In *34th Rencontres de Moriond: Gravitational Waves and Experimental Gravity*, pages 199–206, 1999.

[89] Bernard F Schutz and Franco Ricci. Gravitational Waves, Sources, and Detectors. *arXiv e-prints*, page arXiv:1005.4735, 5 2010.

[90] M. Servillat, C. O. Heinke, W. C. G. Ho, J. E. Grindlay, J. Hong, M. van den Berg, and S. Bogdanov. Neutron star atmosphere composition: the quiescent, low-mass X-ray binary in the globular cluster M28. *Monthly Notices of the Royal Astronomical Society*, 423(2):1556–1561, 06 2012. ISSN 0035-8711. DOI: `10.1111/j.1365-2966.2012.20976.x`. URL `https://doi.org/10.1111/j.1365-2966.2012.20976.x`.

[91] Avneet Singh and Maria Alessandra Papa. Opportunistic search for continuous gravitational waves from compact objects in long-period binaries. *arXiv e-prints*, page arXiv:2208.14117, 8 2022.

[92] Avneet Singh, M. Alessandra Papa, Heinz-Bernd Eggenstein, and Sinead Walsh. Adaptive clustering procedure for continuous gravitational wave searches. *Phys. Rev. D*, 96(8):082003, 2017. DOI: `10.1103/PhysRevD.96.082003`.

[93] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang. Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data. *Astrophys. J.*, 909(1):79, 2021. DOI: `10.3847/1538-4357/abc7c9`.

[94] B. Steltner, M. A. Papa, H. B. Eggenstein, R. Prix, and B. Machenschalk. Deep Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O3a Public Data. *in preparation*, 2023.

[95] Benjamin Steltner, Thorben Menne, Maria Alessandra Papa, and Heinz-Bernd Eggenstein. Density-clustering of continuous gravitational wave candidates from large surveys. *Phys. Rev. D*, 106(10):104063, 11 2022. DOI: `10.1103/PhysRevD.106.104063`.

[96] Benjamin Steltner, Maria Alessandra Papa, and Heinz-Bernd Eggenstein. Identification and removal of non-Gaussian noise transients for gravitational-wave searches. *Phys. Rev. D*, 105(2):022005, 2022. DOI: `10.1103/PhysRevD.105.022005`.

[97] Stephen R. Taylor and Davide Gerosa. Mining Gravitational-wave Catalogs To Understand Binary Stellar Evolution: A New Hierarchical Bayesian Framework. *Phys. Rev. D*, 98(8):083017, 2018. DOI: `10.1103/PhysRevD.98.083017`.

[98] Greg Ushomirsky, Curt Cutler, and Lars Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *Mon. Not. Roy. Astron. Soc.*, 319:902, 2000. DOI: `10.1046/j.1365-8711.2000.03938.x`.

[99] Samantha A. Usman et al. The PyCBC search for gravitational waves from compact binary coalescence. *Class. Quant. Grav.*, 33(21):215004, 2016. DOI: `10.1088/0264-9381/33/21/215004`.

[100] J. Wang and K. Riles. A semi-coherent directed search for continuous gravitational waves from supernova remnants in the ligo o3 data set. presentation at the April 2021 APS meeting, session Y16.00007, 2021.

[101] J. M. Weisberg, J. H. Taylor, and L. A. Fowler. Gravitational Waves From an Orbiting Pulsar. *Sci. Am.*, 245:66–74, 1981. DOI: `10.1038/scientificamerican1081-74`.

[102] Karl Wette. Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars. *Phys. Rev. D*, 92(8):082003, 2015. DOI: `10.1103/PhysRevD.92.082003`.

[103] Clifford M. Will. The Confrontation between General Relativity and Experiment. *Living Rev. Rel.*, 17:4, 2014. DOI: `10.12942/lrr-2014-4`.

[104] G. Woan, M.D. Pitkin, B. Haskell, D.I. Jones, and P.D. Lasky. Evidence for a Minimum Ellipticity in Millisecond Pulsars. *Astrophys. J. Lett.*, 863(2):L40, 2018. DOI: `10.3847/2041-8213/aad86a`.

[105] Yuanhao Zhang, Maria Alessandra Papa, Badri Krishnan, and Anna L. Watts. Search for Continuous Gravitational Waves from Scorpius X-1 in LIGO O2 Data. *Astrophys. J. Lett.*, 906(2):L14, 2021. DOI: `10.3847/2041-8213/abd256`.

[106] Sylvia J. Zhu, Masha Baryakhtar, Maria Alessandra Papa, Daichi Tsuna, Norita Kawanaka, and Heinz-Bernd Eggenstein. Characterizing the continuous gravitational-wave signal from boson clouds around Galactic isolated black holes. *Phys. Rev. D*, 102(6):063020, 2020. DOI: `10.1103/PhysRevD.102.063020`.

[107] M. Zimmermann. Gravitational Waves From Rotating and Precessing Rigid Bodies. II. General Solutions and Computationally Useful Formulas. *Phys. Rev. D*, 21:891–898, 1980. DOI: `10.1103/PhysRevD.21.891`.

[108] M. Zimmermann and E. Szedenits. Gravitational Waves from Rotating and Precessing Rigid Bodies: Simple Models and Applications to Pulsars. *Phys. Rev. D*, 20:351–355, 1979. DOI: `10.1103/PhysRevD.20.351`.

[109] J. Zweizig and K. Riles. Information on self-gating of h(t) used in o3 continuous-wave searches. Technical Document T2000384 available at `https://dcc.ligo.org/`, 2021.

Acknowledgements

I want to thank a lot of people for their help and support throughout the work on this thesis. These have been exciting years and I look forward to the future with great anticipation.

I would like to express my deepest appreciation to Maria Alessandra Papa for being my advisor, for leading such a wonderful research group and giving me the opportunity to work in this field. She has motivated me to move forward, but has also always allowed me the freedom to independently explore different approaches and take unconventional paths. The discussions with her have definitely improved this work.

I would also like to extend my deepest gratitude to the Max Planck Society, and the directors of the AEI for leading such an amazing institute full of incredible people. In particular Karsten Danzmann, who awakened my curiosity about gravitational waves with his phenomenal first-semester experimental physics lectures, and Bruce Allen for leading the data analysis subdivision, Einstein@Home and ATLAS.

Special thanks go to the Einstein@Home volunteers, without whom my work - in the truest sense of the word - would have been impossible. I would like to thank the Einstein@Home and ATLAS staff for maintaining these two amazing resources. I am particular grateful to Carsten Aulbert, who was always there to help whenever a problem occurred.

I also wish to thank all my co-authors on the published papers for their contributions, for which I am very grateful.

It has been a pleasure to work in this institute with the amazing people and nice colleagues. In particular Heinz-Bernd Eggenstein who I closely collaborated with in the tedious tasks of line investigations and visual identification of disturbed bands, and to whom I could always go to discuss things. The same goes for Maximillian Bensch with his help regarding writing GPU code, the python-lalsuite interface. I thank Reinhard Prix and Maximillian Bensch for maintaining our software repositories. I thoroughly enjoyed working with my cluster-buddy Thorben Menne, culminating in holding the group hostage for 2.5 hours presenting why density clustering is great. I thank Anjana Ashok and Lars Nieder for always being there for me, useful discussions, and exactly the right amount of silly discussions and distractions.

Several people played a decisive role that I am working in this field. I thank Thomas Hilscher for giving me the inspiration to study physics, after I realized very late that "Computational Methods in Engineering" is not the right study for me. I also thank my bachelor thesis supervisor Georgios Antonopoulos who introduced me to data analysis, and Avneet Singh for teaching the "Data Analysis Lab" and subsequently introducing me to Maria Alessandra Papa.

I also want to thank my friends for being there and supporting me over the years.

Finally, I cannot begin to express my gratitude to my family, my parents, grandma, grandpa, and my fiancée for their love and support throughout my life.

Curriculum Vitae

Personal Details

Name	Benjamin Steltner
Date of Birth	02.03.1991
Place of Birth	Hannover, Germany

Education

2018 – 2023 **PhD studies at Gottfried Wilhelm Leibniz Universität Hannover, Germany**,
Doctoral thesis at the Max Planck Institute for Gravitational Physics (Albert Einstein Institute), supervisor Prof. Dr. Papa, "All-Sky Searches for Continuous Gravitational Waves",
Affiliated member of the LIGO Scientific Collaboration (2017 - 2018),
Supervisor of the student lab "Data Analysis Lab" (2019),
Member of the local organizing committee for the hybrid conference "The Gravitational Wave Physics and Astronomy Workshop (GWPAW 2021)" (2021)

2015 – 2018 **Master studies in Physics at Gottfried Wilhelm Leibniz Universität Hannover, Germany**,
Master thesis at the Max Planck Institute for Gravitational Physics (Albert Einstein Institute, supervisor Prof. Dr. Papa, "Preparing Gravitational Wave Data for Continuous Wave Searches",
Minor subject: Computer Science

2011 – 2015 **Bachelor studies in Physics at Gottfried Wilhelm Leibniz Universität Hannover, Germany**,
Bachelor thesis at the Laser Zentrum Hannover, supervisor Prof. Dr. Heisterkamp, "Design and implementation of an algorithm for Phase Unwrapping",
Minor subject: Computer Science

2003 – 2010 **Abitur (equivalent to A level)**, *Hannah-Arendt-Gymnasium, Barsinghausen, Germany*
Focus subjects: Physics, Mathematics, Politics

1. B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang. Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data. *Astrophys. J.*, 909(1):79, 2021. DOI: `10.3847/1538-4357/abc7c9`

2. Benjamin Steltner, Maria Alessandra Papa, and Heinz-Bernd Eggenstein. Identification and removal of non-Gaussian noise transients for gravitational-wave searches. *Phys. Rev. D*, 105(2):022005, 2022. DOI: `10.1103/PhysRevD.105.022005`

3. Benjamin Steltner, Thorben Menne, Maria Alessandra Papa, and Heinz-Bernd Eggenstein. Density-clustering of continuous gravitational wave candidates from large surveys. *Phys. Rev. D*, 106(10):104063, 11 2022. DOI: `10.1103/PhysRevD.106.104063`

4. Anjana Ashok, Banafsheh Beheshtipour, Maria Alessandra Papa, Paulo C. C. Freire, Benjamin Steltner, Bernd Machenschalk, Oliver Behnke, Bruce Allen, and Reinhard Prix. New Searches for Continuous Gravitational Waves from Seven Fast Pulsars. *Astrophys. J.*, 923(1):85, 2021. DOI: `10.3847/1538-4357/ac2582`

5. Jing Ming, Maria Alessandra Papa, Heinz-Bernd Eggenstein, Bernd Machenschalk, Benjamin Steltner, Reinhard Prix, Bruce Allen, and Oliver Behnke. Results From an Einstein@Home Search for Continuous Gravitational Waves From G347.3 at Low Frequencies in LIGO O2 Data. *Astrophys. J.*, 925(1):8, 2022. DOI: `10.3847/1538-4357/ac35cb`

6. L. Nieder et al. Discovery of a Gamma-ray Black Widow Pulsar by GPU-accelerated Einstein@Home. *Astrophys. J. Lett.*, 902(2):L46, 2020. DOI: `10.3847/2041-8213/abbc02`

7. L. Nieder, C. J. Clark, C. G. Bassa, J. Wu, A. Singh, J. Y. Donner, B. Allen, R. P. Breton, V. S. Dhillon, H. B. Eggenstein, J. W. T. Hessels, M. R. Kennedy, M. Kerr, S. Littlefair, T. R. Marsh, D. Mata Sánchez, M. A. Papa, P. S. Ray, B. Steltner, and J. P. W. Verbiest. Detection and timing of gamma-ray pulsations from the 707 Hz pulsar J0952−0607. *Astrophys. J.*, 883(1):42, 5 2019. DOI: `10.3847/1538-4357/ab357e`

8. Vladimir Dergachev, Maria Alessandra Papa, Benjamin Steltner, and Heinz-Bernd Eggenstein. Loosely coherent search in LIGO O1 data for continuous gravitational waves from Terzan 5 and the galactic center. *Phys. Rev. D*, 99 (8):084048, 2019. DOI: `10.1103/PhysRevD.99.084048`

9. Georgios C. Antonopoulos, Mirko S. Rakoski, Benjamin Steltner, Stefan Kalies, Tammo Ripken, and Heiko Meyer. Experimental setup combining digital holographic microscopy (DHM) and fluorescence imaging to study gold nanoparticle mediated laser manipulation. In Gabriel Popescu and YongKeun Park, editors, *Quantitative Phase Imaging*, volume 9336 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 93360G, March 2015. DOI: `10.1117/12.2079322`

10. Georgios C. Antonopoulos, Benjamin Steltner, Alexander Heisterkamp, Tammo Ripken, and Heiko Meyer. Tile-Based Two-Dimensional Phase Unwrapping for Digital Holography Using a Modular Framework. *PLoS ONE*, 10(11):e0143186, November 2015. DOI: `10.1371/journal.pone.0143186`

The other publications authored as part of the LIGO scientific collaboration are not listed.