



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. Kwang³

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

Received 2020 September 25; revised 2020 October 25; accepted 2020 October 28; published 2021 March 8

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 $20.0 \text{ Hz} \leq f \leq 585.15 \text{ Hz}$ and spin-down $-2.6 \times 10^{-9} \text{ Hz s}^{-1} \leq \dot{f} \leq 2.6 \times 10^{-10} \text{ Hz s}^{-1}$. 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 $-2.6 \times 10^{-9} \text{ Hz s}^{-1} \leq \dot{f} \leq 2.6 \times 10^{-10} \text{ Hz s}^{-1}$, 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



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

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 “ \times ” 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

$$\begin{aligned} h_+(t) &= A_+ \cos \Phi(t) \\ h_\times(t) &= A_\times \sin \Phi(t), \end{aligned} \quad (2)$$

with the “+” and “ \times ” amplitudes

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

The angle between the total angular momentum of the star and the line of sight is $0 \leq \iota \leq \pi$ and $h_0 \geq 0$ 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 $\Phi(t) = \Phi(\tau_{\text{SSB}}(t))$. The gravitational wave phase as a function of τ_{SSB} is assumed to be

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

We take $\tau_{0\text{SSB}} = 1177858472.0$ (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):

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

The $\bar{\mathcal{F}}$ is the average of \mathcal{F} over segments, in general computed at different templates for every segment. The resulting $\bar{\mathcal{F}}$ 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 $\bar{\mathcal{F}}$ at the exact fine-grid template point. We indicate the recomputed statistic with a subscript “r,” as, for example, in $\bar{\mathcal{F}}_r$.

In Gaussian noise $N_{\text{seg}} \times 2\bar{\mathcal{F}}$ follows a chi-squared distribution with $4N_{\text{seg}}$ degrees of freedom, $\chi^2_{4N_{\text{seg}}}(\rho^2)$. The noncentrality parameter $\rho^2 \propto h_0^2 T_{\text{data}}/S_h$, 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 $\bar{\mathcal{F}}$ 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

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 “persistence 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, $\hat{\beta}_{S/GLIL}$. 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 $\hat{\beta}_{S/GLIL}$ 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 hexagonal edge length d :

$$d(m_{\text{sky}}) = \frac{1}{f} \frac{\sqrt{m_{\text{sky}}}}{\pi \tau_E}, \quad (6)$$

with $\tau_E \simeq 0.021$ 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:

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

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

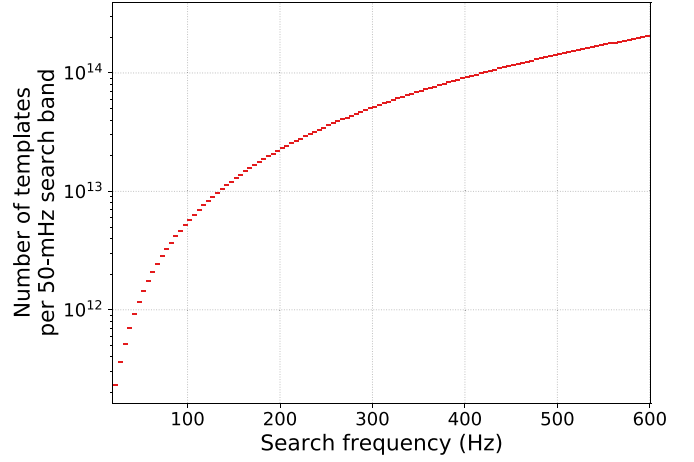


Figure 2. Number of searched templates per 50 mHz band as a function of frequency. The sky resolution increases with frequency causing the increase in the number of templates. The number of templates in frequency and spin-down is 14,970 and 8855 respectively.

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 $|\psi| \leq \pi/4$ and source position uniformly distributed on the sky (uniform in $0 \leq \alpha < 2\pi$ and in $-1 \leq \sin \delta \leq 1$). 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.

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})	$\frac{r_{\text{sky}}}{d(8.0 \times 10^{-3})}$	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	≈ 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 $\hat{\beta}_{\text{S/GLIL}}$ 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 $\hat{\beta}_{\text{S/GLIL}}$, 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.




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: $\hat{\beta}_{S/GLLR}$ and $2\bar{F}_r$ value of the loudest candidate (the candidate with the highest $\hat{\beta}_{S/GLLR}$ 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

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

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 $a = 1 \dots 4$ 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

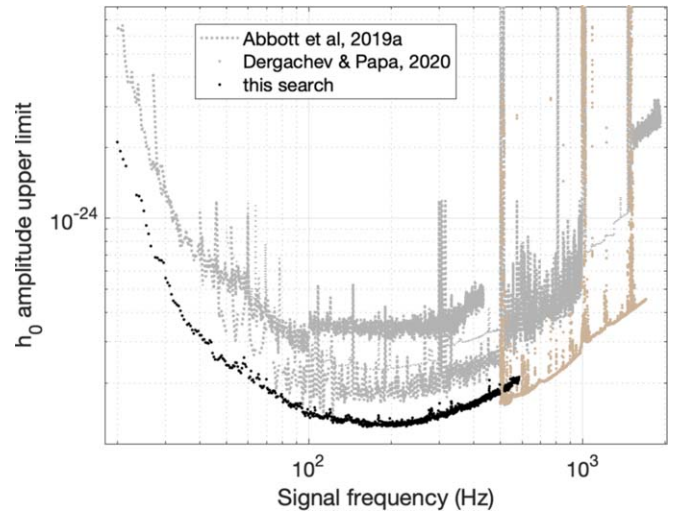


Figure 6. Smallest gravitational wave amplitude h_0 that we can exclude from the assumed population of signals (see Section 3.4). We compare our results with the latest literature: the Falcon search (Dergachev & Papa 2020a) and the LIGO results (Abbott et al. 2019) on the same data. There are multiple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




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, $-2.6 \times 10^{-9} \text{ Hz s}^{-1}$. 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 $2.6 \times 10^{-10} \text{ Hz s}^{-1}$. 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 $\approx 12 \text{ ms}$. 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.

We thank the Einstein@Home volunteers, without whose support this search could not have happened.

We acknowledge the NSF grant No. 1816904.

The follow-up searches were all performed on the ATLAS cluster at AEI Hannover. We thank Carsten Aulbert and Henning Fehrmann for their support.

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

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*
- Dröissigacker, 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