

Einstein@Home search for continuous gravitational waves from Cassiopeia A

Sylvia J. Zhu,^{1,2,*} Maria Alessandra Papa,^{1,2,4,†} Heinz-Bernd Eggenstein,^{2,3} Reinhard Prix,^{2,3} Karl Wette,^{2,3} Bruce Allen,^{2,4,3} Oliver Bock,^{2,3} David Keitel,^{2,3,5} Badri Krishnan,^{2,3} Bernd Machenschalk,^{2,3} Miroslav Shaltev,^{2,3} and Xavier Siemens⁴

¹Max-Planck-Institut für Gravitationsphysik, am Mühlenberg 1, 14476 Potsdam-Golm, Germany

²Max-Planck-Institut für Gravitationsphysik, Callinstraße 38, 30167 Hannover, Germany

³Leibniz Universität Hannover, Welfengarten 1, 30167 Hannover, Germany

⁴University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA

⁵Universitat de les Illes Balears, IAC3—IEEC, E-07122 Palma de Mallorca, Spain

(Received 30 August 2016; published 28 October 2016)

We report the results of a directed search for continuous gravitational-wave emission in a broad frequency range (between 50 and 1000 Hz) from the central compact object of the supernova remnant Cassiopeia A (Cas A). The data come from the sixth science run of LIGO, and the search is performed on the volunteer distributed computing network Einstein@Home. We find no significant signal candidate and set the most constraining upper limits to date on the gravitational-wave emission from Cas A, which beat the indirect age-based upper limit across the entire search range. At 170 Hz (the most sensitive frequency range), we set 90% confidence upper limits on the gravitational-wave amplitude h_0 of $\sim 2.9 \times 10^{-25}$, roughly twice as constraining as the upper limits from previous searches on Cas A. The upper limits can also be expressed as constraints on the ellipticity of Cas A; with a few reasonable assumptions, we show that at gravitational-wave frequencies greater than 300 Hz we can exclude an ellipticity of $\gtrsim 10^{-5}$.

DOI: [10.1103/PhysRevD.94.082008](https://doi.org/10.1103/PhysRevD.94.082008)

I. INTRODUCTION

Isolated neutron stars with nonaxisymmetric asymmetries are thought to be one of the best sources for continuous gravitational-wave emission. We report the results of a directed search for continuous gravitational-wave emission from the central compact object of the supernova remnant Cassiopeia A (Cas A) with the Laser Interferometer Gravitational-Wave Observatory (LIGO). Directed searches, in which the source and therefore the sky position are specified, are generally more sensitive than all-sky surveys. The reason is that typically fewer templates are needed for directed searches than for all-sky surveys; this results in a smaller trials factor and hence in a smaller weakest detectable signal at fixed detection confidence.

At an age of a few hundreds of years, Cas A is one of the youngest known supernova remnants [1]. Its young age means that any asymmetries in the central compact object that were produced at birth are likely still present. Based on x-ray observations, the central compact object is most likely a neutron star with a low surface magnetic field

strength [2]. No pulsed electromagnetic emission has been observed from the central object, so its spin parameters are unknown.

Assuming the central object is a neutron star, its asymmetries would be expected to continuously produce slowly evolving and nearly monochromatic gravitational waves (e.g., Ref. [3]). We perform a search for this gravitational-wave emission from Cas A using data from the sixth LIGO science run with the volunteer distributed computer network Einstein@Home [4].

For the remainder of this paper, when we refer to Cas A, we are referring to the central compact object.

II. SEARCH

A. Data used in this search

The two LIGO interferometers are located in the United States in Hanford, Washington, and Livingston, Louisiana, a separation distance of 3000 km [5]. The last science run of initial LIGO, S6, took place between July 2009 and October 2010 [6]. For this analysis, we only use data taken between February 6, 2010 (GPS time 949461068 s), and October 20, 2010 (GPS time 971629632 s), selected for the best sensitivity [7].

Unlike what we did for previous Einstein@Home searches [8,9], we do not perform any upfront line cleaning to remove known artifacts.

*sylvia.zhu@aei.mpg.de

†maria.alessandra.papa@aei.mpg.de

Published by the American Physical Society under the terms of the *Creative Commons Attribution 3.0 License*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

TABLE I. The search parameters (rounded to the first decimal point) are listed. t_{ref} is the reference time at which the values of f and \dot{f} are defined. γ_1 and γ_2 are the refinement factors for the \dot{f} and \ddot{f} grids, respectively, during the incoherent summation stage.

| | |
|-----------------------------------|-----------------------|
| T_{coh} (h) | 140 |
| t_{ref} (GPS s) | 960541454.5 |
| N_{seg} | 44 |
| δf (Hz) | 5.4×10^{-7} |
| $\delta \dot{f}$ (Hz s $^{-1}$) | 8.2×10^{-12} |
| $\delta \ddot{f}$ (Hz s $^{-2}$) | 1.9×10^{-18} |
| γ_1 | 90 |
| γ_2 | 60 |

B. Search setup

We perform a semicoherent search and rank the results according to the line-robust statistic \hat{O}_{SGL} [10], in a manner similar to Refs. [8,9]. The basic ingredient is the averaged \mathcal{F} statistic [11,12], $\overline{\mathcal{F}}$, computed using the Global Correlation Transform (GCT) method [13,14]. In a stack-slide search, the time series data are partitioned into $i = 1 \dots N_{\text{seg}}$ segments of length T_{coh} each. The data from every segment are match filtered against a set of signal templates, each specified by a set of parameters (the signal frequency f , the first-order spindown \dot{f} , the second-order spindown \ddot{f} , and the sky position) to produce values of the detection statistic \mathcal{F}_i for each segment (the coherent step). These \mathcal{F}_i values are then combined to produce an average value of the statistic across the N_{seg} segments, $2\overline{\mathcal{F}}$, which is the core statistic that we use in these analyses (the incoherent step). The values of the signal template parameters f , \dot{f} , and \ddot{f} are given by a predetermined grid. The f grid spacing (i.e., the separation between two adjacent values of f in the search) is kept the same for the coherent and incoherent steps, while the spacings for the \dot{f} and \ddot{f} grids for the incoherent summing are finer by factors of 90 and 60, respectively. The search parameters are summarized in Table I and were derived using the optimization scheme described in Ref. [15] assuming a run duration of 6 months on Einstein@Home.

C. Detection and ranking statistics

The $2\overline{\mathcal{F}}$ statistic gives a measure of the likelihood that the data resemble a signal vs Gaussian noise; therefore, signals are expected to have high values of $2\overline{\mathcal{F}}$. However, line disturbances in the data can also result in high values of $2\overline{\mathcal{F}}$. The line-robust statistic, \hat{O}_{SGL} , was designed to address this by testing the signal hypothesis against a composite noise model comprising a combination of Gaussian noise and a single-detector spectral line. The \hat{O}_{SGL} parameters are tuned as described in Ref. [10] using

simulations so that the detection efficiency of \hat{O}_{SGL} performs as well as $2\overline{\mathcal{F}}$ in Gaussian noise and better in the presence of lines. For this search, the value of c^* (related to the tuning parameter in the choice of prior; see Ref. [10]) is set to be 34.8, which corresponds to a Gaussian false-alarm probability of 10^{-9} .

The \hat{O}_{SGL} distribution even in Gaussian noise is not known analytically. Therefore, although we use the \hat{O}_{SGL} toplists to find the best signal candidates, we still use $2\overline{\mathcal{F}}$ as the detection statistic for ascertaining a candidate's significance. For a stack-slide search with N_{seg} segments, the $N_{\text{seg}} \times 2\overline{\mathcal{F}}$ distribution in Gaussian noise follows a chi-squared distribution with $4N_{\text{seg}}$ degrees of freedom [16].

D. Parameter space

Since the spin parameters of Cas A are unknown, our search encompasses a large range of possible gravitational-wave frequencies f , namely, from 50 to 1000 Hz. For a given value of f , the \dot{f} and \ddot{f} ranges are given by the following specifications,

$$\dot{f} \in [-f/\tau_{\text{NS}}, 0] \quad (1)$$

$$\ddot{f} \in [0, 2f/\tau_{\text{NS}}^2], \quad (2)$$

where τ_{NS} is the fiducial age of the neutron star, taken to be 300 years. As discussed in Ref. [3], this choice of τ_{NS} is on the young end of the age estimates, which yields a larger search parameter space than other, less conservative choices. The searched parameter space at each value of f is a rectangle in the $\dot{f} - \ddot{f}$ plane, and the search volume increases quadratically with f (Fig. 1).

Compared to Refs. [3,17,18], the largest magnitude of the first-order spindown parameter is the same, corresponding to a conservative assumption (in the sense that it allows for the broadest range of first-order spindown values) on the average braking index at a fixed age of the object. The range of the second-order spindown is constructed differently here than in Refs. [3,18] in that it does not depend on \dot{f} . The highest searched value of \ddot{f} is $n\dot{f}_{\text{max}}^2/f$, with $n = 2$ being the instantaneous braking index. The searches mentioned above took this as the *lower* boundary of the \ddot{f} range and set the upper boundary at $n = 7$. Our choice does not search such a broad range of \ddot{f} values and is driven by ease of setup of the search. Observational data on braking indexes support this choice [19].

We estimate that searching over third-order spindown values is not necessary. We do this by counting how many templates are needed to cover the third-order spindown range. The third-order spindown template extent $\Delta\ddot{f}$ in a semicoherent search with mismatch m is

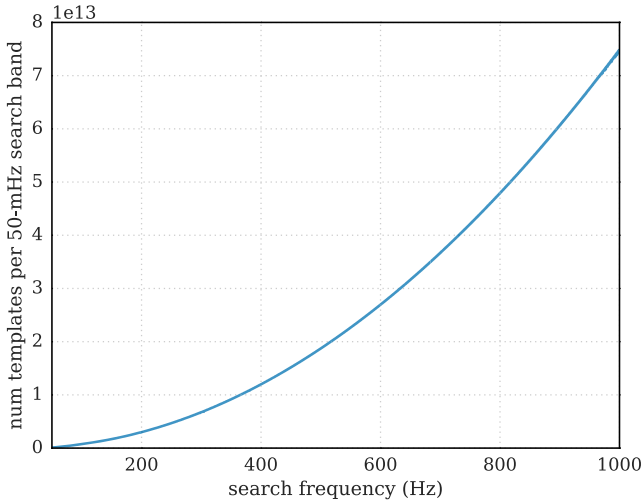


FIG. 1. For this search, the number of templates per 50 mHz search band increases quadratically with f . At each value of f , the \dot{f} search range is $[-f/\tau_{\text{NS}}, 0]$, and the \ddot{f} search range is $[0, 2f/\tau_{\text{NS}}^2]$. In total, 4.99×10^{17} templates are included in this search.

$$\Delta\ddot{f} = \frac{1}{\gamma_3} \frac{2520\sqrt{m}}{\pi T_{\text{coh}}^4}. \quad (3)$$

For this search, we set $m = 0.2$ and $\gamma_3 \approx 3.89 \times 10^5$ [20]. The template extent of Eq. (3) is $\sqrt{mg^{33}}$ where g^{ij} is the inverse of the phase metric [21]. The third-order spindown range, consistent with the choices of Eq. (2), is $6f/\tau_{\text{NS}}^3$. With $\tau_{\text{NS}} = 300$ years, we find that we do not need more than a single template to cover the third-order spindown range; therefore, we do not need to add a third-order spindown dimension.

The location of Cas A is known to within $\sim 1''$, which is smaller than the sky resolution of our search. Hence, we only search a single sky position (right ascension = 23h 23m 28s, declination = $58^\circ 58' 43''$).

E. Distribution of the computational load

The search runs on volunteer computers in the Einstein@Home network and is split into 9.2 million work units (WUs), with each WU designed to run for about 6 h on a modern PC. A single WU encompasses a 50 mHz range in f and the entire range of \dot{f} at the start value of f , along a single slice out of the \ddot{f} range. The results from WUs that search over the same 50 mHz range are combined into a single band, and these multiple WUs together cover the entire \dot{f} range at that value of f . Each WU searches through approximately 5×10^{10} templates and returns two lists of results corresponding to the 3000 templates with the highest values of the $2\mathcal{F}$ and \hat{O}_{SGL} statistics (described in Sec. II C), called the toplist. The total number of templates included in this search is 4.99×10^{17} .

F. Semiautomatic identification of disturbances

When the noise is purely Gaussian, the $2\overline{\mathcal{F}}$ distribution is well modelled, and the significances of signal candidates can be determined in a straightforward manner. However, disturbances generate deviations from the expected distribution. In order to meaningfully use the same statistical analysis on all of the candidates, the disturbed 50 mHz bands must be excluded from the search. Previous searches [8,9] relied on a visual inspection of the full data set in order to identify the disturbed bands, which is a very time consuming endeavor. Here, we introduce a semiautomatic method that greatly reduces the number of bands that need to be visually inspected.

We use two indices to identify bands that cannot automatically be classified as undisturbed: 1) the density of toplist candidates in that band and 2) their average $2\overline{\mathcal{F}}$. We classify as undisturbed those bands of which the maximum density and average $2\overline{\mathcal{F}}$ are well within the bulk distribution of the values for these quantities in the neighbouring frequency bands and mark the remainder as potentially disturbed and in need of visual inspection.

The size of the toplist and the frequency grid spacing are fixed. Therefore, when a disturbance is present in a 50 mHz band, the toplist within that band disproportionately include templates in the parameter space near the disturbance. We look for evidence of disturbances in the 50 mHz bands using a method that mimics and replaces the visual inspection used in previous searches [8]: for a given band, we calculate the density of candidates in a 10×10 grid in $f - \dot{f}$ space and take the maximum density as an indicator of how disturbed the band is likely to be. Since disturbances also manifest as deviations in the $2\overline{\mathcal{F}}$ distribution, we use the mean of $2\overline{\mathcal{F}}$ as an additional indicator of how much a band is disturbed. A visual representation of these concepts is shown in Fig. 2.

Because the search volume increases with f , both the mean of $2\overline{\mathcal{F}}$ and the candidate density vary with frequency. To account for this effect, we compare the observed maximum density and mean $2\overline{\mathcal{F}}$ from each band with the distribution of maximum density and mean $2\overline{\mathcal{F}}$ values in sets of 200 contiguous 50 mHz bands (10 Hz). These constitute our reference distributions.

Since the majority of the bands are undisturbed, the reference distributions are composed of a well-defined bulk (from the undisturbed bands) with tails (disturbed bands), as illustrated in Fig. 3. We define the ‘‘bulk’’ of each distribution by eye and then mark the bands that fall outside of this bulk on either side as being potentially disturbed; we generally expect disturbed bands to be in the upper ends of the distributions (that is, to have particularly large values of maximum density and mean $2\overline{\mathcal{F}}$) but also include bands in the lower ends so as not to miss any unexpected disturbed behavior. We proceed with a full visual inspection only of this potentially disturbed subset. Figure 3 shows the

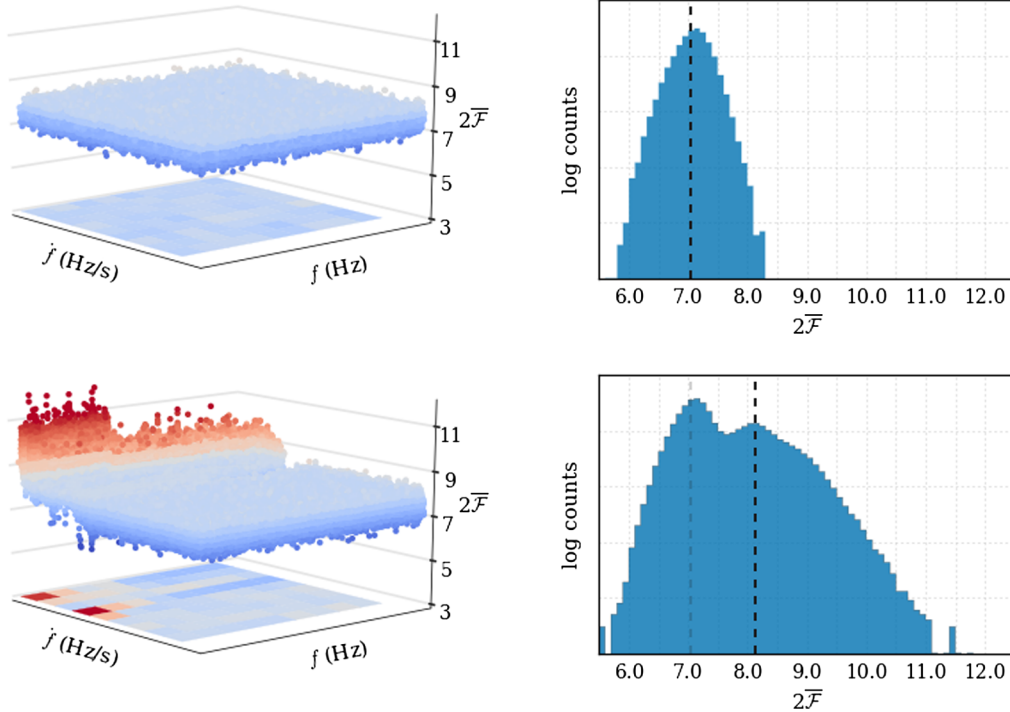


FIG. 2. The maximum density (left) and the mean $2\overline{\mathcal{F}}$ value (right) for the candidates are the two metrics we use to identify potentially disturbed bands. *Left*: undisturbed bands (an example in the top panels) have a very uniform density of candidates in $f - \dot{f}$, while disturbed bands (an example in the bottom panels) present marked overdensities. The $2\overline{\mathcal{F}}$ values in the $f - \dot{f}$ plane are shown in the 3D plot, while the candidate density is shown in the 2D projections. The maximum density in a disturbed band tends to be much higher (here, more red) than the maximum density in an undisturbed band. *Right*: The $2\overline{\mathcal{F}}$ distribution in an undisturbed band (top) and in a disturbed band (bottom).

reference distributions for the bands between 90 and 100 Hz. These are typical examples and illustrate how the “by eye” definition of the bulk of the distributions is not subtle. When selecting the bulk, we err on being conservative: when in doubt, we label bands as being potentially disturbed, as these will be reinspected later.

If a signal were present, it would not be excluded because of the automated procedure. On the one hand, if it were so weak that the band would not be marked as

disturbed, the band would automatically be included in the analysis. On the other hand, if it were strong enough that the automated procedure marked its band as being potentially disturbed, then it would be visually inspected by a human who would recognize the signature of a signal and not discard the band.

This method still requires human input in two steps: first, to define the bulk of the reference distributions and, second, to inspect the subset of potentially disturbed bands.

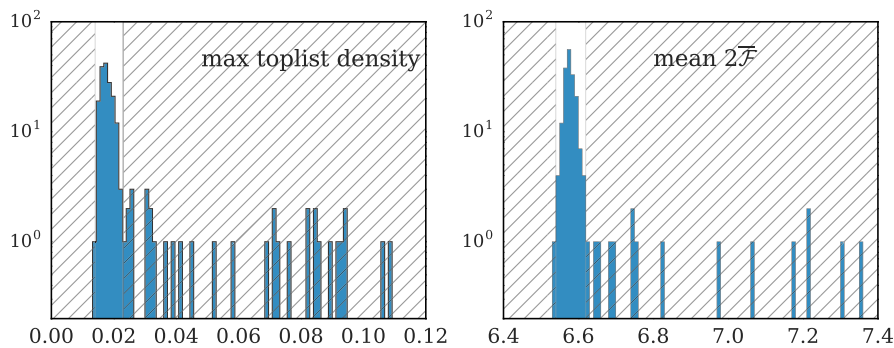


FIG. 3. The distributions of maximum toplist density (left) and mean toplist $2\overline{\mathcal{F}}$ (right) are shown for a sample 10 Hz frequency range. Both distributions consist of an undisturbed body with a disturbed tail (hatched). All 50 mHz bands that fall within the hatched areas are marked as potentially disturbed.

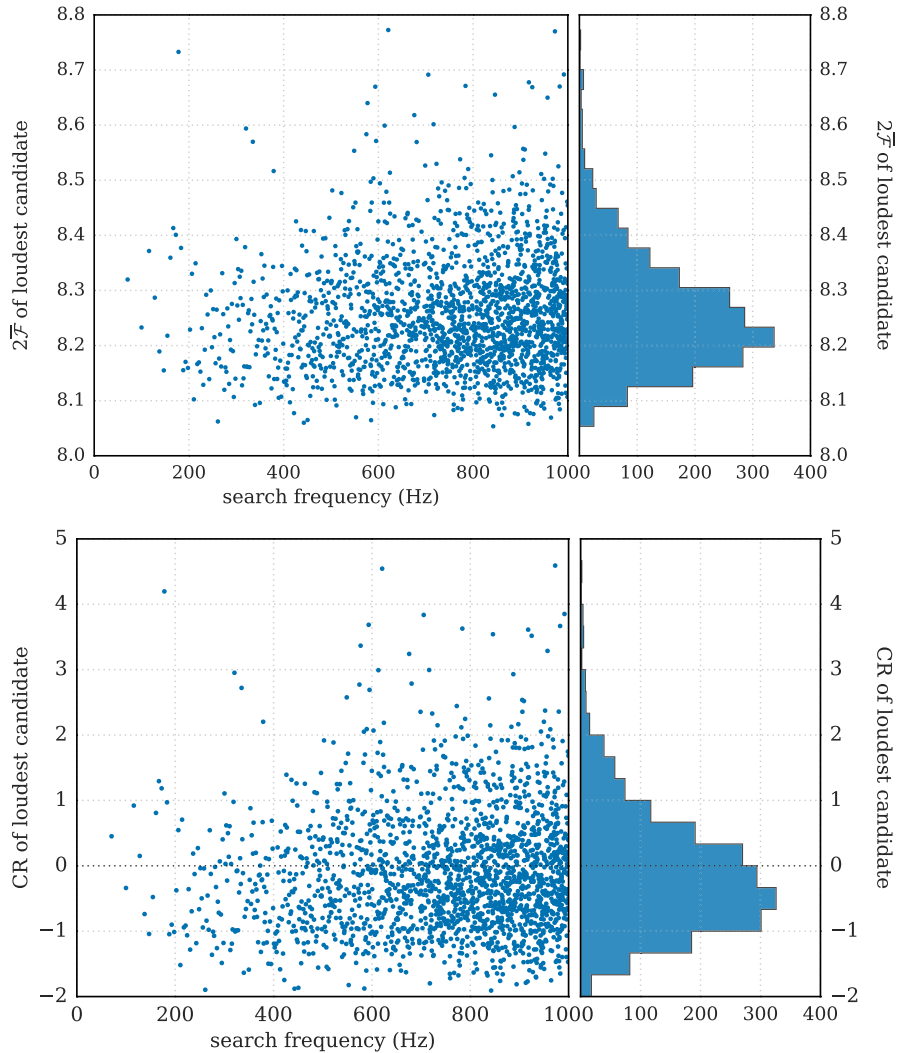


FIG. 4. For each of the 2000 partitions, we determine the $2\bar{\mathcal{F}}$ of the loudest candidates (top) as well as their CR values (bottom), where CR is defined in Eq. (4).

However, the “calibration” work necessary for determining the bulk of the reference distributions only requires the inspection of two distributions every 10 Hz rather than multiple distributions every 50 mHz. Furthermore, the bands that do not pass the undisturbed-classification criteria and require visual inspection are only 15% of the total set. Overall, this procedure still cuts down the required time from multiple days with multiple people to a few hours by a single person.

This procedure requires minimal tuning and relies only on the assumption that the reference distributions are predominantly undisturbed. This has so far been our experience on all the LIGO data sets that we have inspected. We are confident that this method can be applied to other sets of gravitational-wave data.

When we compare this method against a full visual inspection of a few search frequency ranges (50 to 100 Hz, 450 to 500 Hz, and 950 to 1000 Hz), it identifies $\sim 95\%$ of

the disturbed bands and misses only the most marginal disturbances. After we apply this method to the entire frequency range, we exclude a total of 1991 50 mHz bands as being disturbed ($\sim 10\%$); these are listed in Table S2.

G. Analysis of undisturbed bands

The $2\bar{\mathcal{F}}$ distribution in Gaussian noise only depends on the number of effectively independent templates searched (N). However, the grid spacings are chosen to maximize signal recovery, so the N templates are not fully independent. The observed distribution is instead described by an effective number of templates $N_{\text{eff}} < N$. The value of N_{eff} is obtained by fitting the distribution of the loudest candidates (i.e., the highest values of $2\bar{\mathcal{F}}$).

We divide the entire set of 50 mHz bands across our search frequency range into 2000 partitions of approximately

equal parameter space volume, which results in $\sim 2 \times 10^{14}$ templates per partition. In order to create these partitions, we calculate an exact partitioning of the total search volume and divide the full range of 50 mHz bands so that the number of templates in each partition best matches the number of templates in the exact partitioning. Since the number of templates in a band grows with frequency, the frequency width spanned by each partition decreases with increasing frequency. This can be clearly seen in the left panels of Fig. 4. In this way, since each partition contains roughly the same number of search templates, the expected loudest candidate in each partition is the same and is drawn from the same underlying distribution, defined by N_{eff} . For this search, we find that $N_{\text{eff}}/N \approx 0.65$.

Figure 4 shows both the $2\bar{\mathcal{F}}$ (top) and the critical ratio (CR) (bottom) for the loudest candidates. We define CR as

$$CR := \frac{2\bar{\mathcal{F}}_{\text{meas}} - 2\bar{\mathcal{F}}_{\text{exp}}}{\sigma_{2\bar{\mathcal{F}}}}, \quad (4)$$

where $2\bar{\mathcal{F}}_{\text{meas}}$ is the measured value of the loudest, $2\bar{\mathcal{F}}_{\text{exp}}$ is the expected value of the loudest, and $\sigma_{2\bar{\mathcal{F}}}$ is the expected standard deviation for the loudest over a partition. The loudest candidate over the entire search is in the 620.85 Hz band and has a $2\bar{\mathcal{F}}$ value of 8.77; this is also the most significant candidate, with a CR of 4.56. However, if we consider the entire searched parameter space rather than just the partition at 620.85 Hz, the CR value of the most significant candidate drops to <0 ; i.e., the expected loudest is actually higher than the loudest that we observe. This tells us that our search has not revealed any gravitational-wave signal from Cas A in the targeted waveform parameter space, as even the template that most resembles a signal has a statistical significance that is well within the expectations due to random chance.

We convert the CR values of the loudest candidates to p-values to represent the chance probability of finding a partition-loudest candidate as significant as or more significant than what was measured in the search. The results are plotted in Fig. 5, along with the expected distribution of p-values in Gaussian noise. There is a small systematic deviation from the expected distribution which arises from a subtle difference between the $\hat{\mathcal{O}}_{\text{SGL}}$ and $2\bar{\mathcal{F}}$ toplists and is not due to any physical effect.

III. UPPER LIMITS

We find no candidates with $CR > 5$ and no excess in the p-value distribution. Therefore, we set frequentist 90% upper limits on the continuous gravitational-wave strain $h_0^{90\%}$ in our search range using the process described in previous works [8,9], which we summarize below.

The $h_0^{90\%}$ in a partition is the gravitational-wave amplitude at which 90% of a population of signals with

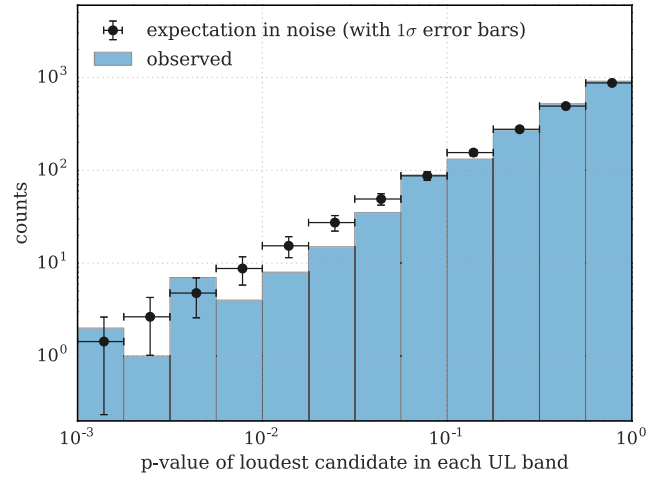


FIG. 5. The p-values for the loudest candidate in each upper limit band is plotted in the blue histogram, and the expectation in Gaussian noise is shown in the black scatter points. We do not find any excess in our search. The small systematic deviation in our data from the expected is caused by a subtle difference in the $\hat{\mathcal{O}}_{\text{SGL}}$ and $2\bar{\mathcal{F}}$ rankings.

parameters within the partition would produce a more significant candidate than the most significant candidate measured by the search in that partition. We determine $h_0^{90\%}$ by injecting signals at fixed amplitudes bracketing the $h_0^{90\%}$ level, then running the search on these injections and counting how many injections were recovered (i.e., how many produced a candidate more significant than the loudest measured by the actual search). Because this injection-and-recovery procedure is time consuming, we perform it on only a subset of 20 representative partitions—uniformly distributed in frequency in the search range—rather than the full set of 2000 and use these results to derive the upper limits in all the other partitions.

For each of the 20 injection partitions, we fit a sigmoid to the detection efficiency (the fraction of recovered injections) as a function of injection amplitude to determine both the value of $h_0^{90\%}$ and the 1- σ uncertainty on $h_0^{90\%}$. We determine the $h_{0,\text{CR}_i}^{90\%,j}$ in each of the injection partitions corresponding to different detection criteria binned by CR, with $\text{CR}_i = [0, 1, 2, 3, 4, 5]$. For each CR_i , we derive the corresponding sensitivity depths

$$\mathcal{D}_{\text{CR}_i}^{90\%,j} = \frac{\sqrt{S_h(f_j)}}{h_{0,\text{CR}_i}^{90\%,j}} \quad [1/\sqrt{\text{Hz}}]. \quad (5)$$

By design, the sensitivity depths of this search are roughly constant across the different partitions. We estimate the sensitivity depths by averaging the values across the injection partitions:

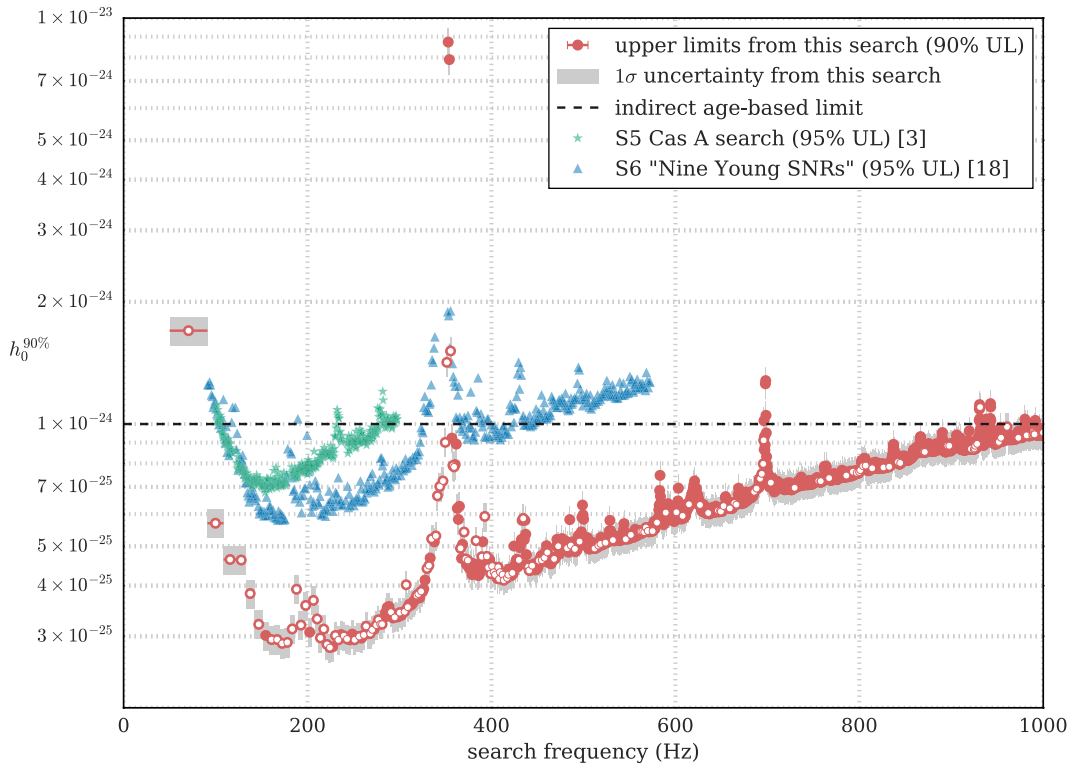


FIG. 6. 90% confidence strain amplitude upper limits in each of the 2000 partitions. The results for partitions that contain only undisturbed 50 mHz bands are plotted in the filled red circles, while the results for partitions with disturbed 50 mHz bands are plotted in the open red circles. We also plot the 95% confidence upper limits from two previous searches on Cas A in green and blue. Our upper limits beat the so-called indirect age-based limit [17] across the band.

$$\mathcal{D}_{\text{CR}_i}^{90\%} = \frac{1}{20} \sum_{j=1}^{20} \mathcal{D}_{\text{CR}_i}^{90\%,j}. \quad (6)$$

For each of the remaining partitions, at frequencies around f_k , we derive the upper limit as

$$h_0^{90\%}(f_k) = \frac{\sqrt{S_h(f_k)}}{\mathcal{D}_{\text{CR}_i(f_k)}^{90\%}}, \quad (7)$$

where $\text{CR}_i(f_k)$ is the significance bin of the loudest candidate of the partition at f_k and $S_h(f_k)$ is the power spectral density of the data. $\mathcal{D}_{\text{CR}_i}^{90\%} \approx 70 \text{ Hz}^{-1/2}$ for this search.

Our upper limits are plotted in Fig. 6 in red with $1\text{-}\sigma$ uncertainties in gray and provided in tabular form as Supplemental Material [22]. The uncertainties in $h_0^{90\%}$ that we report here are propagated from the statistical uncertainties in fitting the recovery. The partitions containing disturbed bands (which were not included in the analysis) are marked with open circles.

The upper limit value near 170 Hz, where the detectors are the most sensitive, is 2.9×10^{-25} . This value is roughly two times lower than the previous most constraining upper

limit on Cas A [18], plotted in blue, which also used S6 data. Our upper limits are also more than twice as constraining as an earlier Cas A search, plotted in green [3], which ran on S5 data.¹ Our upper limits beat the so-called indirect age-based limit [17] across the vast majority of the frequency range.

IV. CONCLUSIONS

The upper limits on the gravitational-wave strain from Cas A translate into constraints on the shape of Cas A. As described in Ref. [24], a neutron star's mass distribution can be described by the ellipticity ϵ , where

$$\epsilon = \frac{|I_{xx} - I_{yy}|}{I_{zz}} \quad (8)$$

and I_{zz} is the principal moment of inertia of the star around its rotational axis. If a neutron star at a distance D and

¹However, we note that the other two searches produced 95% upper limits rather than 90% upper limits; the latter is the standard for the broad surveys by Einstein@Home [8,9,23]. The ratio between the 90% and the 95% confidence upper limits is ~ 1.1 .

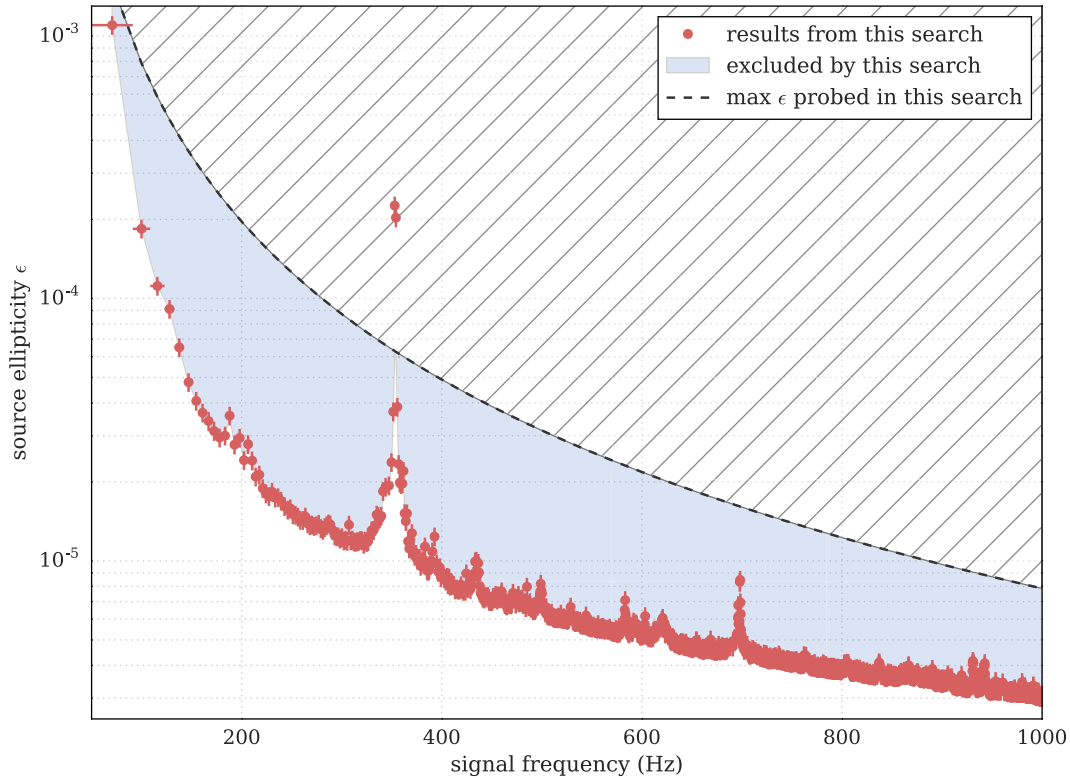


FIG. 7. With a few reasonable assumptions, we can convert the upper limits on the gravitational-wave strain to upper limits on the ellipticity of Cas A. The shaded area denotes the source ellipticities (as a function of signal frequency) that are excluded by this search; ellipticities in this region would have produced signals that this search would have detected. The dashed line marks the spindown ellipticity probed by this search and is set by our choice of \dot{f} search range.

spinning at a frequency $f/2$ has a nonaxisymmetric distortion ϵ , then it will produce a continuous gravitational wave with a frequency f and amplitude h_0 . These quantities are related to each other as follows:

$$\epsilon = \frac{h_0 D}{f^2} \frac{c^4}{4\pi^2 I_{zz} G}. \quad (9)$$

Equation (9) shows how we can reexpress the constraints on the gravitational-wave amplitude as constraints on the ellipticity. We take the distance to Cas A to be 3.4 kpc [1] and I_{zz} to be 10^{38} kg m².

These constraints on source ellipticity are shown in Fig. 7. For instance, if Cas A is emitting gravitational waves at around 200 Hz (and therefore spinning at a frequency of 100 Hz), its ellipticity should be less than a few times 10^{-5} , since we would have been able to detect gravitational waves produced by larger ellipticities.

The maximum ellipticity is the ellipticity necessary to sustain emission at the spindown limit, i.e., when all of the lost rotational energy is radiated as gravitational waves. This spindown ellipticity is

$$\epsilon^{sd} = \sqrt{\frac{5c^5}{32\pi^4 G} \frac{x|\dot{f}|}{If^5}} \quad \text{with } x = 1, \quad (10)$$

where \dot{f} is twice the spin-frequency derivative.

The highest spindown ellipticity for an object emitting gravitational waves at a frequency f that our search could have detected can be computed from Eq. (10) by setting $\dot{f} = f/300$ yr. For an isolated system, if \dot{f} is twice the spin-frequency derivative, larger ellipticities would violate energy conservation. For this reason, we only highlight the region between the ellipticity upper limit curve and the spindown ellipticity curve as excluded by the search. However, we note that systems in general could have ellipticities larger than the spindown ellipticity if the gravitational wave \dot{f} (the apparent \dot{f}) differs from the intrinsic one due to, for example, radial motion.

ACKNOWLEDGMENTS

The authors thank the Einstein@Home volunteers who have supported this work by donating compute cycles of their machines, and acknowledge the support from the Max Planck Society and the Leibniz Universität Hannover. All

the postprocessing computational work for this search was carried out on the ATLAS supercomputing cluster at the Max-Planck-Institut für Gravitationsphysik/Leibniz Universität Hannover. Maria Alessandra Papa, Bruce Allen, and Xavier Siemens gratefully acknowledge the

support from NSF PHY Grant No. 1104902. We also thank the Continuous Wave Group of the LIGO Scientific Collaboration for useful discussions. This document has been assigned LIGO Laboratory document number LIGO-P1600212.

-
- [1] R. A. Fesen, M. C. Hammell, J. Morse, R. A. Chevalier, K. J. Borkowski, M. A. Dopita, C. L. Gerardy, S. S. Lawrence, J. C. Raymond, and S. van den Bergh, The expansion asymmetry and age of the Cassiopeia A supernova remnant, *Astrophys. J.* **645**, 283 (2006).
- [2] W. C. G. Ho and C. O. Heinke, A neutron star with a carbon atmosphere in the Cassiopeia A supernova remnant, *Nature* **462**, 71 (2009).
- [3] J. Abadie *et al.* (LIGO Scientific Collaboration), First search for gravitational waves from the youngest known neutron star, *Astrophys. J.* **722**, 1504 (2010).
- [4] <https://www.einsteinathome.org>.
- [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration), LIGO: The Laser Interferometer Gravitational-Wave Observatory, *Rep. Prog. Phys.* **72** (2009).
- [6] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Sensitivity achieved by the LIGO and Virgo gravitational wave detectors during LIGO's sixth and Virgo's second and third science runs, [arXiv:1203.2674](https://arxiv.org/abs/1203.2674).
- [7] M. Shaltev, Optimizing the StackSlide setup and data selection for continuous-gravitational-wave searches in realistic detector data, *Phys. Rev. D* **93**, 044058 (2016).
- [8] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaborations), Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data running on the Einstein@Home volunteer distributed computing project, [arXiv:1606.09619](https://arxiv.org/abs/1606.09619) [*Phys. Rev. D* (to be published)].
- [9] A. Singh, M. A. Papa, H.-B. Eggenstein, S. Zhu, H. Pletsch, B. Allen, O. Bock, B. Maschenchalk, R. Prix, and X. Siemens, Results of an all-sky high-frequency Einstein@Home search for continuous gravitational waves in LIGO's fifth science run, *Phys. Rev. D* **94**, 064061 (2016).
- [10] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, Search for continuous gravitational waves: Improving robustness versus instrument artifacts, *Phys. Rev. D* **89**, 064023 (2014).
- [11] P. Jaranowski, A. Królak, and B. F. Schutz, Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection, *Phys. Rev. D* **58**, 063001 (1998).
- [12] C. Cutler and B. F. Schutz, Generalized \mathcal{F} -statistic: Multiple detectors and multiple gravitational wave pulsars, *Phys. Rev. D* **72**, 063006 (2005).
- [13] H. J. Pletsch and B. Allen, Exploiting Large-Scale Correlations to Detect Continuous Gravitational Waves, *Phys. Rev. Lett.* **103**, 181102 (2009).
- [14] H. J. Pletsch, Parameter-space metric of semicoherent searches for continuous gravitational waves, *Phys. Rev. D* **82**, 042002 (2010).
- [15] R. Prix and M. Shaltev, Search for continuous gravitational waves: Optimal StackSlide method at fixed computing cost, *Phys. Rev. D* **85**, 084010 (2012).
- [16] K. Wette, Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars, *Phys. Rev. D* **85**, 042003 (2012).
- [17] K. Wette *et al.*, Searching for gravitational waves from Cassiopeia A with LIGO, *Classical Quantum Gravity* **25**, 235011 (2008).
- [18] J. Aasi *et al.* (LIGO Scientific Collaboration), Searches for continuous gravitational waves from nine young supernova remnants, *Astrophys. J.* **813**, 1 (2015).
- [19] O. Hamil, J. R. Stone, M. Urbanec, and G. Urbancová, Braking index of isolated pulsars, *Phys. Rev. D* **91**, 063007 (2015).
- [20] M. Shaltev, Coherent follow-up of continuous gravitational-wave candidates: minimal required observation time, *J. Phys. Conf. Ser.* **363**, 012043 (2012).
- [21] K. Wette, Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars, *Phys. Rev. D* **92**, 082003 (2015).
- [22] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevD.94.082008> for the table of upper limit values and the list of 50-mHz bands that were excluded from this search.
- [23] M. A. Papa *et al.*, Hierarchical follow-up of sub-threshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data, [arXiv:1608.08928](https://arxiv.org/abs/1608.08928).
- [24] J. Ming, B. Krishnan, M. A. Papa, C. Aulbert, and H. Fehrmann, Optimal directed searches for continuous gravitational waves, *Phys. Rev. D* **93**, 064011 (2016).