Optimizing the choice of analysis method for all-sky searches for continuous gravitational waves with Einstein@Home

Sinéad Walsh, Karl Wette, Maria Alessandra Papa, and Reinhard Prix

Optimising the choice of analysis method for all-sky searches for continuous gravitational waves with Einstein@Home

Sinéad Walsh,[1, 2, 3] Karl Wette,[4, 2, 3] Maria Alessandra Papa,[2, 1, 3] and Reinhard Prix[2, 3]

[1] *Department of Physics, University of Wisconsin, Milwaukee, WI 53201, USA*
[2] *Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-30167 Hannover, Germany*
[3] *Institut für Gravitationsphysik, Leibniz Universität Hannover, D-30167 Hannover, Germany*
[4] *ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav) and Centre for Gravitational Physics, Research School of Physics and Engineering, The Australian National University, Acton ACT 2601, Australia*

Rapidly rotating neutron stars are promising sources of continuous gravitational wave radiation for the LIGO and Virgo interferometers. The majority of neutron stars in our galaxy have not been identified with electromagnetic observations. All-sky searches for isolated neutron stars offer the potential to detect gravitational waves from these unidentified sources. The parameter space of these blind all-sky searches, which also cover a large range of frequencies and frequency derivatives, presents a significant computational challenge. Various methods have been designed to perform these searches within the limits of available computing resources. Recently, a search method called Weave has been proposed to achieve template placement with a minimal number of templates. We employ a mock data challenge to assess the ability of this method to recover signals in an all-sky search for unknown neutron stars, and compare its search sensitivity with that of the global correlation transform method (GCT), which has been used for all-sky searches with the Einstein@Home volunteer computing project for a number of years. We find that the Weave method is 14% more sensitive than GCT for an all-sky search on Einstein@Home, with a sensitivity depth of 57.9 ± 0.6 $\frac{1}{\sqrt{\text{Hz}}}$ at 90% detection efficiency, compared to $50.8^{+0.7}_{-1.1}$ $\frac{1}{\sqrt{\text{Hz}}}$ for GCT. This corresponds to a 50% increase in the volume of sky where we are sensitive with the Weave search. We also find that the Weave search recovers candidates closer to the true signal position. In the all-sky search studied here the improvement in candidate localisation would lead to a factor of 70 reduction in the computing cost required to follow up the same number of candidates. We assess the feasability of deploying the search on Einstein@Home, and find that Weave requires significantly more memory than is typically available on a volunteer computer. We conclude that, while GCT remains the best choice for deployment on Einstein@Home due to its lower memory requirements, Weave presents significant advantages for the subsequent hierarchical follow-up searches of interesting candidates.

I. INTRODUCTION

Continuous gravitational waves (CWs) from isolated neutron stars (NSs) are a potential source of detectable gravitational waves. CW radiation is emitted by rotating NSs with non-axisymmetric deformations. The signal is expected to be relatively stable over many years. While the amplitude of CW signals is expected to be small, the continuous nature of the signal allows us to integrate the signal over large time spans of data to distinguish it from noise.

Broad-band all-sky searches cover the whole sky over a broad range of frequency and frequency derivative in order to detect CW radiation from unknown NSs. All-sky searches in the Laser Interferometer Gravitational-Wave Observatory (LIGO) [1, 2] and Virgo [3, 4] data have so far not resulted in detection. Instead upper limits have been placed on the amplitude of CWs from isolated NSs [5–9]. The advanced detectors, which began operation in 2015, will eventually have a sensitivity to these weak signals over an order of magnitude better than that of the previous generation, with the largest gains at frequencies below 100 Hz.

The most sensitive search for CW signals is performed with a fully coherent integration over a large timespan, months to years, of data. The computational power re-
quired for the integration increases rapidly with the observation time of the data. When searching for CW signals over a broad frequency and spindown range, and over the whole sky, a fully coherent search quickly becomes computationally unfeasible [10, 11]. This is the motivation for semi-coherent search methods. The data is split into shorter segments which are analysed coherently. Then, for each point in the parameter space $(f_0, \dot{f}, \alpha, \delta)$, the coherent results in each segment are combined incoherently to determine the detection statistic over the entire set of data. For limited available computing power, these semi-coherent search methods achieve higher sensitivity than could be achieved with a fully coherent search [11].

Einstein@Home is a volunteer computing project where members of the public donate their idle computing power to the search for long lasting astrophysical signal [14]. The donated computing power allows for broader and more sensitive searches than previously possible [5, 7]. The Einstein@Home GW search that we consider here is described further in [5].

Some of the first semi-coherent searches used the Hough transform method [12, 13], which sums weighted binary counts in each segment, depending on whether the normalised power in that segment exceeds a certain threshold. The parameter space used to sum the coherent results across segments is refined in both sky

and spindown. This method was applied to all-sky searches with the Einstein@Home computing project [14] on LIGO S5 data [15]. The global correlation transform method (GCT) [16] was later developed and claimed to be more sensitive at fixed computing cost, without the need for sky refinement in the incoherent part of the search. This search method was used in the Einstein@Home all-sky searches in LIGO S6 and O1 data [5, 7]. The GCT method is able to use arbitrarily long coherent segment lengths, however, its sensitivity cannot be reliably estimated analytically, instead extensive injection-and-recovery studies are needed to design an all-sky search using this method [6].

The purpose of the study presented in this article is to assess a recently-developed semi-coherent search method to detect CWs from isolated neutron stars. A principal claim of this method is that it achieves template placement with a minimal number of templates required to achieve a desired maximum mismatch on the semi-coherent and coherent search grid in all-sky searches [17–20]. This would make it a milestone tool, achieving the highest search sensitivity possible at fixed computing cost, with increases in sensitivity only possible if the code can be made faster. An additional, and significant, benefit of having optimal template placement is that we can predict the sensitivity without the need for extensive injections-and-recovery simulations. The implementation of this semi-coherent search method is described in detail in [21], and is hereafter referred to as Weave. The implementation is freely available as part of the LALSuite [22] gravitational-wave data analysis library.

We want to examine the detection efficiency of this search method in an all-sky search for isolated neutron stars with the Einstein@Home project. This type of search is currently performed with the GCT method. With the GCT search method, the optimal search setup (i.e. the optimal combination of search parameters) at fixed computational cost is determined by running Monte-Carlo injection-and-recovery studies of many search setups [6].

With the Weave search method, it should be possible to semi-analytically determine optimal search setups, without the need for extensive Monte-Carlo injection-and-recovery studies. This would greatly simplify the process of setting up an all-sky search. This is because with Weave we can reliably predict the distribution of relative SNR-loss (in the following referred to as *mismatch*) of the Weave template bank for any given search-setup (see [20, 21] for details, and [23] for an implemenation of the mismatch predictor). This is a key input required for accurate sensitivity *estimation* (as described in [24]), as well as for the semi-analytic sensitivity *optimization* method developed in [23, 25]. This method has since been extended to take advantage of empirical non-linear mismatch predictions (as given in [20]), and work is ongoing to apply this framework to the Weave search method. The results of this will be reported in future work.

The first objective of this paper is to compare the achievable detection efficiency of the Weave search method with that of existing semi-coherent methods, especially the GCT search method. The second objective of this paper is to determine if the Weave search method can be applied within the technical limitations of searches with the Einstein@Home project.

This comparison is first made using the same data as in the mock data challenge (MDC), which was used to compare five all-sky search methods for the detection of gravitational waves from unknown neutron stars [26]. The comparison is repeated in simulated detector data (Gaussian noise) with more fake signals, to allow for a more precise comparison with the GCT search. This comparison is made for the standard CW signal model described in Sections II and IV.

An overview of the search method is presented in Section III. The MDC is described briefly in Section IV; and the choice of optimal setup for the Weave search is described in Section V. Section VI describes how the methods are compared and the results of the comparison.

II. THE SIGNAL

Gravitational waves (GWs) emitted from non-axisymmetric NSs are typically described by a signal which remains relatively stable over years of observation [27]. The strain amplitude of the GW is proportional to the ellipticity, ε, defined as

$$\varepsilon = \frac{|I_{xx} - I_{yy}|}{I_{zz}}, \tag{1}$$

where I_{zz} is the principal moment of inertia of the star, and I_{xx} and I_{yy} are the moments of inertia about the axes. The strain amplitude of the GW at the detector, assuming a rigidly rotating triaxial body, is then given by

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

where f is the frequency of the GW, G is the Gravitational constant, c is the speed of light, and d is the distance to the NS. For a star steadily rotating around a principal axis of inertia, the frequency of the GW is at twice the rotational frequency of the NS. The frequency evolves over time as energy is lost due to various dissipation mechanisms, including GW emission. The first time derivative of the frequency, $\dot{f}$, is referred to as spindown.

The signal model is described by eight parameters, four phase evolution parameters $\lambda = (f_0, \dot{f}, \alpha, \delta)$ and four amplitude parameters $(h_0, \iota, \psi, \phi_0)$, where ι is the inclination angle between the line of sight to the NS and its rotation axis, ψ is the polarisation angle and ϕ_0 is the initial phase of the signal at a reference time τ_0.

The parameter f_0 is the GW frequency at τ_0. The GW frequency in the Solar System Barycenter (SSB) at time

τ is given by

$$f_{\mathrm{SSB}}(\tau) = f_0 + \dot{f}(\tau - \tau_0), \qquad (3)$$

assuming the second time derivative of the frequency, $\ddot{f}$, is negligible. The transformation between the time t at which a certain wavefront arrives at a detector and the time τ at which the same wavefront arrives at the SSB is

$$\tau(t) = t + \frac{\mathbf{r}(t) \cdot \mathbf{n}}{c} + \Delta_{E\odot} - \Delta_{S\odot} \qquad (4)$$

where $\mathbf{r}(t)$ is the position vector of the detector in the SSB frame, $\mathbf{n}$ is the unit-vector pointing towards the sky location of the source, and $\Delta_{E\odot}$ and $\Delta_{S\odot}$ are the relativistic Einstein and Shapiro time delays respectively. Ignoring the relativistic corrections, the instantaneous frequency $f(t)$ of a continuous GW signal as observed at the detector is related to the frequency $f(\tau)$ at the SSB by the relation:

$$f(t) = f_{\mathrm{SSB}}(\tau) + f_{\mathrm{SSB}}(\tau)(\frac{\mathbf{v}(\tau) \cdot \mathbf{n}}{c}) \qquad (5)$$

where $\mathbf{v}(\tau)$ is the detector velocity with respect to the SSB frame. The frequency measured in the detector, $f(t)$, is shifted due to Doppler effect from the motion of the Earth around the Sun and by the rotation of the Earth.

In a blind all-sky search there is the potential for the detection of signals produced by different emission mechanisms (e.g. r-modes [28]). The ability of the all-sky search methods to recover such signals is not examined in this study. Here we assume the signal follows the model described above.

III. THE WEAVE METHOD

The search begins with many months of data being split into segments of a few days each. For the coherent analysis, the multi-detector $\mathcal{F}$-statistic [27, 29] is computed for each segment and for each parameter space point on the coherent grid. The coherent grid is determined by the Weave method, which is founded on the use of the reduced supersky metric, introduced in Section B IV of [17]. The reduced supersky metric is a close approximation to the phase metric, and expressed in a convenient coordinate system with respect to which the metric is constant and numerically well-conditioned.

An average $2\mathcal{F}$-statistic, $2\overline{\mathcal{F}}$, is calculated by summing $2\mathcal{F}$ values for a single template across the segments. (A template waveform is defined by the four phase evolution parameters $\lambda = (f_0, \dot{f}, \alpha, \delta)$.) A second grid in parameter space is set up, the so-called semi-coherent grid. This grid is finer than the coherent grid in spindown and sky

[19]. The detection statistic at every point in the semi-coherent grid is estimated by summing detection statistic values at the closest points on the coherent grid. Weave uses a metric based method described in [19] to determine the closest point on the coherent grid.

IV. THE MOCK DATA CHALLENGE

In [26], a MDC was employed to empirically compare the performance of different all-sky search methods. This was done by simulating the detector response to CW signals in data from the S6 LIGO science run [30], with fake signals at frequencies ranging from 40 to 1550 Hz. Each of the search pipelines then performed a search over the data to assess its ability to recover the signal.

The sensitivity is measured in terms of the smallest signal amplitude detectable with a given efficiency $h_{\mathrm{smallest\text{-}detectable}}$. The smallest detectable amplitude scales linearly with the noise $\sqrt{S_h}$ so the ratio of the smallest detectable amplitude to the $\sqrt{S_h}$ is roughly constant for a given search across the different frequencies. The smaller the ratio, the higher the sensitivity of the search. For this reason the notion of sensitivity depth $\mathcal{D}$ was introduced [31]: $\mathcal{D} := \frac{\sqrt{S_h}}{h_{\mathrm{smallest\text{-}detectable}}}$. The higher the sensitivity depth, the more sensitive a search is. When comparing the Weave and GCT searches we quote here their sensitivity depths at 90% detection efficiency.

Note that in [26] the signal strength is normalised by S_h defined as the harmonic mean over both detectors of the harmonic sum power spectral density (PSD) of the data, at the frequency of the signal. Instead here we define S_h as the harmonic mean over both detectors of the harmonic mean PSD of the data, at the frequency of the signal. One can approximately[1] convert the former into sensitivity depth by multiplying by $\sqrt{N_{\mathrm{det}}} \approx 1.4$, as 2 detectors were used. A discussion of why the harmonic mean over per-detector PSDs should be used in the calculation of S_h is included in [24].

In the MDC of [26], different search configurations were assumed for the Einstein@Home search with GCT in different frequency ranges, namely from 40 to 500 Hz, 500 to 1000 Hz and 1000 to 2000 Hz. Here we perform the same MDC with Weave only in the frequency range of 40 to 500 Hz, which is particularly interesting because it covers the most sensitive range of the detectors.

In Section VI D 2 we assess the performance of the Weave search method with this MDC data and compare it to the sensitivity of the searches in Figure 3 of [26].

We are most interested in comparing the Weave search method with the GCT method which is used by the

---

[1] assuming roughly equal number of SFTs from each detector

Einstein@Home search in the MDC. However, the difference in sensitivity between the Weave search and the GCT search is less than the uncertainty on the efficiency measurement in the MDC (This is shown in Section VI D 2, Figure 4). Therefore, we generate a more extensive data set so we can make a more precise comparison between the Weave and GCT search methods.

The fake signals in this new set are created in almost the same way as the fake signals used in the MDC, described in Section IV A of [26]. Some differences are, we only generate injections in the frequency range of 40 to 500 Hz, and we do not include injections with positive $\dot{f}$ or with non-zero $\ddot{f}$, both of which were found to have no effect on the measured efficiency.

In the MDC, the h_0 of fake signals was obtained by randomly drawing the SNR from a uniform distribution. However, it was shown in [32] that, at fixed signal strength h_0, a uniform SNR distribution implies an unphysical distribution of ι which is biased towards higher absolute values of $\cos \iota$. The surplus of higher absolute values of $\cos \iota$ results in an overestimate of the CW search sensitivity by $\approx 30\%$ [32] at fixed signal strength. When generating new fake signals, we instead obtain the h_0 by randomly drawing from a uniform distribution of normalised signal strength.

The range of signal strength was chosen to bracket the 90% detection efficiency for both GCT and Weave, because search sensitivities are typically compared at high detection efficiencies [26]. These are comparable to the upper limits values expected from the searches, in case of a null result.

For the MDC the fake signals were injected into 15 months of LIGO S6 data. The GCT search was performed using only the last nine months of this data set. Here, we do not use LIGO S6 data. Instead, we generate nine months of Gaussian noise data with the same duty cycle as the LIGO S6 data. For this comparison between Weave and GCT we are interested in the sensitivity in well behaved data. By generating fake data we avoid being affected by unclassified detector disturbances in S6 data.

V. CHOOSING THE SEARCH SETUP

The search parameters and selection criteria to use in an Einstein@Home all-sky search depend on the parameter space to be covered and the computational budget. The GCT search set-up is the same as the one examined in the MDC [26]. The Weave search is determined through an optimisation process, aimed at achieving the maximum sensitivity depth at fixed computational cost over the same search parameter space as for the GCT search. In particular we considered an all-sky search covering 40 to 500 Hz and 2.9e-9 Hz/s, and 6 months run-

time on Einstein@Home [2]. We also set a threshold on the recovered detection statistic such that we would expect a fixed number of candidates above threshold from false alarms in random noise in an all-sky search. This threshold and number of candidates is established, in Sections VI A and VI D 1, to be such that the time taken to followup candidates from the Weave or GCT search would be equivalent.

At a given coherent segment length, the Weave search takes two input parameters, the maximum metric mismatch for the coherent search grid, $\tilde{\mu}_{\max}$, and for the semicoherent search grid, $\hat{\mu}_{\max}$. These will determine the search grids. We find the optimal search setup by running Monte-Carlo injection-and-recovery studies of many search setups, as done for the GCT method in [6]. We then repeat for different coherent segment lengths. We compare the detection efficiencies at a threshold corresponding to a fixed false alarm rate, and choose the most sensitive setup.

Einstein@Home searches are designed to run for a certain number of months on the Einstein@Home computing project. Therefore, we need to be able to predict the total runtime of different search setups on Einstein@Home. The runtime model for the Weave search method is described in Section IV C of [21], and its implementation is documented in [23].

The runtime prediction relies on having an estimate of the total number of search templates. We also need to know the number of templates to determine the false alarm rate for different thresholds on the detection statistic.

A. Determining the number of templates

Knowing the total number of templates searched is necessary for two reasons: 1) to determine the runtime of a search setup and 2) to predict the false alarm rate in Gaussian noise at different $2\overline{\mathcal{F}}$ thresholds, as described in [5].

A model to estimate the number of templates in a Weave search was not available when the study in this paper began. Instead, the number of templates was obtained by running the Weave search code over a small region of parameter space and counting the number of templates generated. The number of templates on the full search volume was then determined by scaling arguments to be 7×10^{16} templates on the coherent grid and 1×10^{18} semi-coherent templates. The GCT all-sky search would have 1×10^{15} templates on the coherent grid, and 2×10^{17} semi-coherent templates.

———

[2] This assumes the search is run continuously on 8000 'fast' CPUs. In reality, there are more Einstein@Home CPUs, many slower than the CPUs assumed for the runtime estimate. At the time of the MDC [26], 8000 'fast' CPUs was determined to be an reasonable estimate for the speed of an Einstein@Home search.

A model to estimate the number of templates in a Weave search has since been developed, and is described in Section IV B of [21]. The model has been implemented in a script which is publicly available, and given in [21]. The number of templates estimated with this script is expected to be accurate within a factor of a few. Such a deviation in the number of templates could result in a factor of a few difference in the estimated runtime for the search. This uncertainty is too large for a search on Einstein@Home; a search which is designed to run for six months in Einstein@Home should not run for a year or more.

This script estimates that there are 6.1×10^{16} templates on the coherent grid and 2.8×10^{18} semi-coherent templates. This results in an approximate runtime for the all-sky search of 13.5 months on Einstein@Home. This is twice the runtime estimated by running the Weave search over a small region of parameter space and extrapolating from the number of templates counted.

The number of templates required to cover any parameter space includes some padding of the parameter-space boundaries. When the total template count is extrapolated, the padding is also extrapolated, but the padding doesn't extrapolate in the same way as the bulk template count. When the script from [21] estimates the number of templates for the entire search over the whole sky and frequency range, it does not assume the same padding which would apply when the parameter space is split into smaller regions. Einstein@Home searches are split into smaller regions of parameter space, which are searched separately. To estimate the number of templates in an Einstein@Home search, the script which estimates the number of templates should be run separately for each sub-region in parameter space. When the number of templates are estimated in this way, there are 2.3×10^{16} templates on the coherent grid and 1×10^{18} semi-coherent templates. This results in an approximate runtime for the all-sky search of 4.9 months on Einstein@Home.

This estimate of 4.9 months is reasonably similar to the runtime of 6 months estimated by calculating the number of templates from scaling arguments, and is significantly smaller than the runtime estimate of 13.5 months obtained by estimating the number of search templates for the entire search as if it was a single search. However, for this search spanning 40 to 500 Hz, running this script in series for each 0.05 Hz band took more than a week. Therefore, the script to estimate the number of templates from [21] is not suitable for an Einstein@Home search. However, it is suitable to estimate the number of templates for all-sky searches which are not split into many sub-regions in parameter space.

B. The optimal search setup

The best Weave setup found for an all-sky search from 40 to 500 Hz, over nine months of data, which runs for six months on Einstein@Home, is given in Table I. This setup is found using Monte-Carlo injection-and-recovery studies as described in Section V. The GCT search parameters used in the MDC are shown for reference in Table II. The number of semi-coherent templates in each dimension of parameter space are shown for a nominal region of parameters space in Table III.

In Figure 1 we plot the mismatch between the loudest recovered $2\mathcal{F}$ and the maximum possible $2\mathcal{F}$ i.e the $2\mathcal{F}$ we would recover with a search template exactly at the signal parameters. This mismatch indicates how much of the signal SNR we are losing due to the offset between the true signal parameters and the parameters of the closest search template. The Weave search has a mean mismatch of 0.49. The GCT search has a mean mismatch of 0.59.

T_{coh} (h)	60
$N_{segments}$	100
$\tilde{\mu}_{max}$	0.2
$\hat{\mu}_{max}$	20

TABLE I. Search parameters for the Weave MDC search, as determined by Monte-Carlo injection-and-recovery studies. T_{coh} is the coherent segment length in hours. $\tilde{\mu}_{max}$ is the maximum coherent mismatch. The frequency spacing on the coherent grid for this setup is 1.61×10^{-6} Hz. $\hat{\mu}_{max}$ is the maximum semicoherent mismatch.

T_{coh} (h)	60
$N_{segments}$	90
δf (Hz)	3.61×10^{-6}
$\delta \dot{f}$ (Hz/s)	1.16×10^{-10}
sky factor	0.01
$\dot{f}$ refine	230

TABLE II. Einstein@Home GCT MDC search parameters. The sky grid resolution is determined by the sky factor as shown in Equation 8 of [26] The spindown resolution used on the fine grid, for the semi-coherent part of the search, is given by $\delta \dot{f}$ divided by the $\dot{f}$-refine value.

$N_{templates}$	Weave	GCT
0.05 Hz	31,055	13,850
2.9×10^{-9} Hz/s	1,204	5,750
all-sky at 44Hz	67,122	8,090

TABLE III. The number of semi-coherent templates in each dimension of parameter space for a nominal region of parameter space: 0.05 Hz, 2.9×10^{-9} Hz/s and the whole sky at 44 Hz for Weave and 50 Hz for GCT. The GCT sky grid templates are generated for every 10 Hz, so searches over frequencies from 40 to 50 Hz all use the 50 Hz sky grid. The number of sky grid templates in a Weave search scales with frequency.
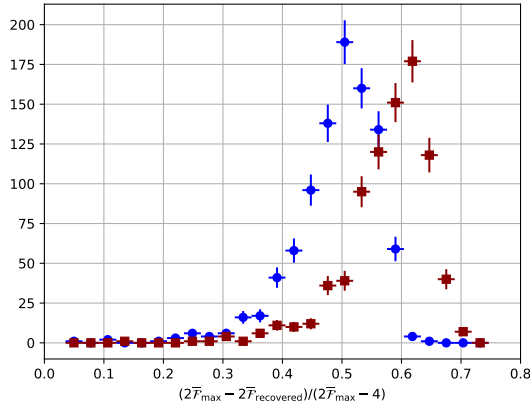
FIG. 1. Histogram of the mismatch between the maximum $2\overline{\mathcal{F}}$, the $2\overline{\mathcal{F}}$ of a search template at the exact signal parameters, and the loudest recovered $2\overline{\mathcal{F}}$ from a GCT search (red) and a Weave search (blue).

VI. COMPARISON OF METHODS

A. Effective number of search templates

The distribution of the average $2\mathcal{F}$ detection statistic, $2\overline{\mathcal{F}}$, in Gaussian noise is expected to follow a chi-squared distribution with $4N_{seg}$ degrees of freedom. From this we can estimate the number of false alarms above a $2\overline{\mathcal{F}}$ value for a search with N independent templates [33].

In Section V A we calculate the number of templates, $N_{\text{templates}}$, in the full all-sky search. Assuming these templates are independent, a $2\overline{\mathcal{F}}$ threshold of 6.1 should result in 40 million false alarms in the all-sky search. However, if the search templates are not fully independent the effective number of independent templates is smaller.

In this section we perform many searches over a small region of parameter space in Gaussian noise. For Weave the search range is 0.05 Hz, 1e-10 Hz/s, and a sky patch with 0.03 radians radius. For GCT the search range is 0.05 Hz, 3e-10 Hz/s, and a sky patch with a radius of 0.03 when projected onto the ecliptic plane. The Weave search has $\approx 3e + 8$ templates and the GCT search has $\approx 5e + 8$ templates.

We compare the distribution of loudest recovered $2\overline{\mathcal{F}}$, $2\overline{\mathcal{F}}_{\text{max}}$, with what we expect if the search templates are independent. By fitting the expected $2\overline{\mathcal{F}}_{\text{max}}$ distribution to the measured distribution, we can determine the effective number of templates in the search. [3]

The results are shown for the GCT search and for the Weave search in Figure 2. For the GCT

---

[3] The shape of the $2\overline{\mathcal{F}}_{\text{max}}$ distribution over correlated templates does not always agree with the $2\overline{\mathcal{F}}_{\text{max}}$ distribution over N_{eff} independent templates, see Appendix C of [24]

search we measure an effective number of templates $N_{eff} = 0.43 N_{\text{templates}}$. For the Weave search $N_{eff} = 0.69 N_{\text{templates}}$.

In an all-sky search with GCT, with $N_{eff} = 0.43 N_{\text{templates}}$, we would expect 16 million false alarms in Gaussian noise with $2\overline{\mathcal{F}} > 6.17$, the threshold used for the MDC. In an all-sky search with Weave, with $N_{eff} = 0.69 N_{\text{templates}}$, a threshold of $2\overline{\mathcal{F}} > 6.155$ would produce the same number of false alarms in Gaussian noise.
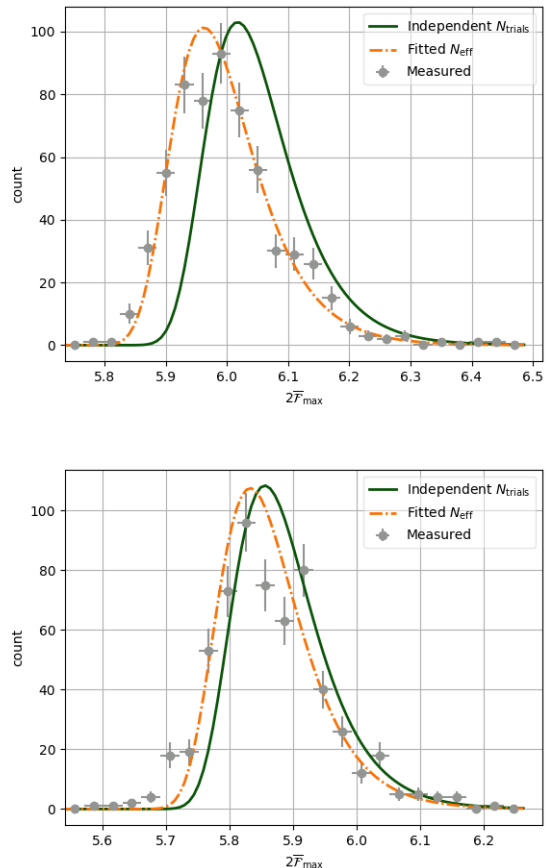

FIG. 2. The distribution of recovered $2\overline{\mathcal{F}}_{\text{max}}$ values in 600 repeated runs over Gaussian noise for GCT (top) and Weave (bottom). If the search templates were independent the measured $2\overline{\mathcal{F}}_{\text{max}}$ (grey) would agree with expected $2\overline{\mathcal{F}}_{\text{max}}$ for $N_{\text{templates}}$ (green). The effective number of templates is obtained with a fit to the measured $2\overline{\mathcal{F}}_{\text{max}}$, the results of the fit are shown in orange.

B. Memory usage

When a search is run on the Einstein@Home computing project, the parameter space to search is split into cells, and each volunteer computer searches a cell corresponding to a specific work-unit (WU) of the global

task. The cells are chosen such that a single WU will run for approximatly 8 hours on a volunteer computer. The loudest candidates recovered from each cell are then returned to the Einstein@Home servers.

While many modern computers have more than 10 GB of memory, Einstein@Home includes older computers with less memory. The searches run single-threaded, so typically one task is executed per available core, meaning the memory we can use is limited to what is available per CPU core. Therefore, to be able to run on Einstein@Home volunteer computers, we limit the memory a WU can use to 1 GB. Normally, the parameter space of an all-sky search is split into 0.05 Hz bands, and then the number of sky grid points is scaled to result in a search which runs for 8 hours. The GCT searches consume a managable amount of memory with such WUs. The Weave search given in Table I, when split up to run for 8 hours over 0.05 Hz, reaches a memory consumption of 2-7 GB, depending on how you split the parameter space.

There are a number of ways to reduce the memory consumption for a WU. As an example, it is possible to stage, in sequence, multiple searches within a WU such that the total runtime of the WU remains close to 8 hours, but at any given time the memory consumption is smaller than if all the work was staged at the same time. This will results in more output files as each search staged will output a list of the most significant candidates, but this is something which can be handled by Einstein@Home. Alternatively, the Weave executable has an option to internally split the frequency and spindown parameter space of a WU into subsets, to limit the memory consumption.

With both of the options for reducing the memory, there are limitations to how small you can make the parameter space. If you reduce the frequency range so there are fewer than $\approx 10k$ frequency bins in a single search, the runtime will increase as the resampling method for calculating the $\mathcal{F}$ becomes less efficient. The search runtime is also found to increase for small spindown search ranges and for small sky patch sizes.

As a concrete example we attempt to determine the WU configuration which would work for an Einstein@Home search with the Weave search given in Table I. However, we find that if the frequency band is less than 0.05 Hz the runtime increases, for a spindown range less than 2.9e-9 Hz/s the runtime increases, and for more than 500 sky patches at 150 Hz the runtime increases. At this limiting search configuration covering 0.05 Hz, 2.9e-9 Hz/s and 500 sky patches at 150 Hz, the memory consumption is 7 GB. If we decrease the WU parameter space to 0.04 Hz, 2.9e-9 Hz/s split internally in 10 partitions, and 500 sky patches at 150 Hz, the memory usage is 1 GB at a cost of a 50% increase in the total search runtime.

A semi-analytical model for predicting the memory usage is described in [21]. This can be used to design WUs within the memory limit constraints given by specific hardware. One of the parameters needed by the mem-

ory model is the maximum size of the Weave code cache, and this cannot easily be predicted. Instead, the user must run the Weave code in a special mode, which is faster than the normal mode, which measures the size of the cache without computing any $2\mathcal{F}$ values. This must be done for each WU configuration considered. This is a step in the deployment of a Weave search on Einstein@Home which is not necessary for a GCT search. On the other hand, memory footprints of a few GBs should pose no problem if using Weave on a dedicated super computing cluster.

C. The line-robust statistic

The most recent all-sky search on Einstein@Home [7] has used the $\hat{\beta}_{\mathrm{S/GLtL}}$ statistic as well as the $2\overline{\mathcal{F}}$ statistic. This is a line- and transient-robust statistic that tests the signal hypothesis against a noise model which, in addition to Gaussian noise, also includes single-detector continuous or transient spectral lines [34, 35].

Including the $\hat{\beta}_{\mathrm{S/GLtL}}$ statistic increases the runtime and memory usage of a search, as the additional statistics must be calculated for each search template. For the GCT and Weave searches considered here, including $\hat{\beta}_{\mathrm{S/GLtL}}$ increases the runtime by $\approx 40\%$ and $\approx 250\%$, respectively. This increase may be different for other setups. The additional computing cost for Weave would require a repeat of the optimisation procedure in Section V with a runtime model which includes the extra statistic, which would ultimately lead to a lower efficiency than the current setup at fixed computing cost.

For the setup considered here, the memory increase is negligible for GCT, which remains well below 1GB, whereas it increases by a factor of ≈ 3 for Weave. This memory increase can potentially be managed as described in Section VI B, at the expense of a higher runtime.

The increase in runtime and memory is due to the fact that the extra statistics are calculated for every point in parameter space. However, this statistic is only beneficial (compared to the $2\overline{\mathcal{F}}$ statistic) in regions of parameter space where the data is disturbed by detector artifacts, which was 10% of the frequency bands in the S6 all-sky search [5]. By calculating the $\hat{\beta}_{\mathrm{S/GLtL}}$ statistic only on the disturbed frequency bands it may be possible to apply the optimal Weave search setup within the computing budget while retaining the benefits of the $\hat{\beta}_{\mathrm{S/GLtL}}$ statistic.

Due to the memory limitations of Einstein@Home (Section VI B) it would not be possible to run this post-processing stage on Einstein@Home, instead it would be run on a computing cluster. Since the original search is designed to run for 6 months over 8000 CPU cores, if as little as 5% of the frequency bands are disturbed the search would take a month to run continuously on 5000 cluster CPU cores. If 10-20% of the frequency bands are

disturbed, rerunning with the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic becomes prohibitively expensive.

When we say a frequency band is disturbed, we mean that the list of the candidates with the largest $2\overline{\mathcal{F}}$ statistic in that frequency band is dominated or saturated by candidates due to a detector artefact. When we use the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic in the search, a separate list of the candidates with the largest $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic is output. The advantage of including the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic is when a frequency band which appears disturbed in the $2\overline{\mathcal{F}}$-statistic list is not disturbed in the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic list, which means the candidates from the detector artefact have been suppressed and there is still the potential to recover a signal in this frequency band.

Even if 10% of frequency bands are disturbed, only a fraction of these will be recovered by using the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic. If this fraction is as low as 20% (2% of the total), we would not consider the exclusion of the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic to be a noteworthy disadvantage of the Weave search. However, it is realistic to expect that 50% of disturbed bands would be recovered by using the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic. For the comparison in this paper we consider a search using only the $2\overline{\mathcal{F}}$ statistic.

D. Results of the comparison

1. Signal parameter estimation

In an all-sky search, above-threshold candidates are passed through hierarchical refinement stages designed to exclude candidates from noise while retaining signal candidates [6]. In the MDC, the parameter space around the candidates in the first refinement stage is used to determine which candidates from the initial stage would result in a detection.

The distance between the signal parameters and the parameters recovered in the MDC, in frequency, spindown and sky position, is shown in Figure 3. Here the distance in sky is represented by dR, the angular separation between two sky positions in radians. The dR scales approximately proportional to the frequency of the signal. Here we see that the Weave candidates have a smaller uncertainty on the signal parameters.

For the GCT search the first refinement stage searches $\delta f \pm 1.9{\times}10^{-4}\,\mathrm{Hz}$, $\delta\dot f \pm 3.46{\times}10^{-11}\,\mathrm{Hz/s}$ and a sky patch of 1.2 initial-search sky grid bins around the selected candidate. After the GCT search $\sim 90\%$ of signals have the loudest candidate within this region around the signal parameters. We find that after the Weave search $\sim 98\%$ of the loudest candidates recovered are within this region around the signal parameters.

After the Weave search, 90% of signals have the loudest candidate within $\delta f \pm 3.9{\times}10^{-5}\,\mathrm{Hz}$, $\delta\dot f \pm 8.2{\times}10^{-12}\,\mathrm{Hz/s}$ and a sky patch with a radius 0.3 the radius of the GCT 90% confidence region. The computing power used to follow-up these candidates scales roughly linearly with
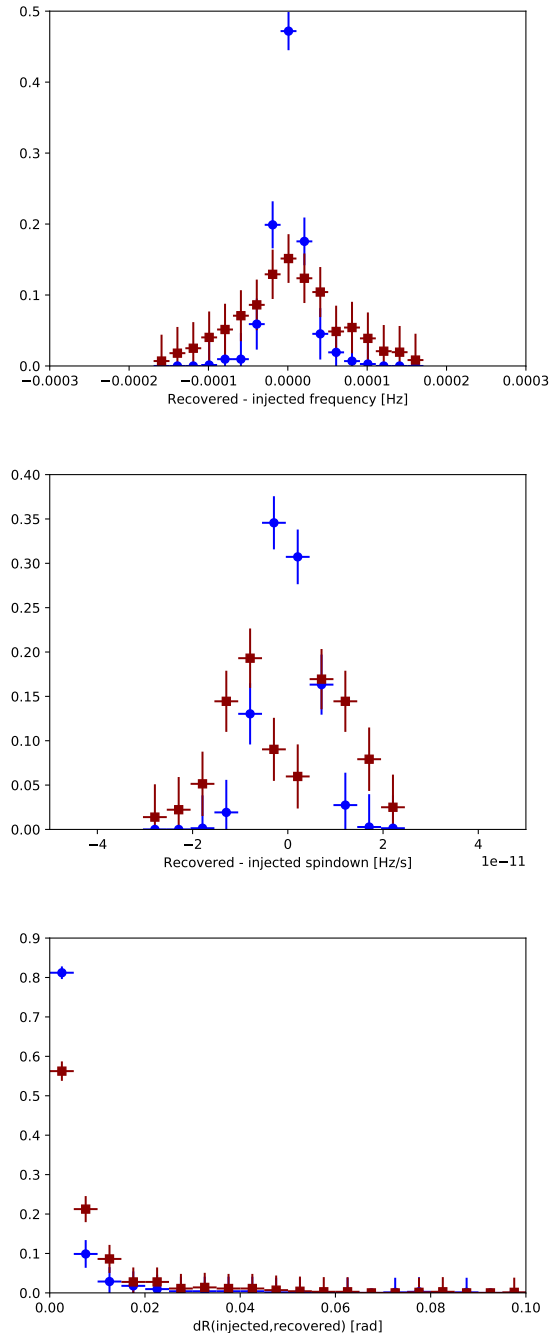

FIG. 3. The distance between the signal and the recovered candidate, in frequency, spindown and sky position, when the candidate with the highest SNR is chosen for Weave (blue) and GCT (red). The search can recover more than one candidate per injection. For strong signals, there are in fact many detection candidates around the signals true parameter values. We study the distribution for the distances of candidates from the true signal parameter values for the detection candidates with the highest $2\overline{\mathcal{F}}$. The dR distribution extends beyond the limits shown in the plot with the max dR value at 0.27 for Weave and 0.38 for GCT.

δf, $\delta \dot{f}$ and the square of the sky patch radius. Per candidate, the first Weave refinement stage would be ~ 200 times less computationally expensive than the GCT first refinement stage. This means that for the same computational cost, we can reduce the $2\overline{\mathcal{F}}$ threshold to follow-up 200 times more candidates.

Alternatively, we can increase the size of the first refinement stage search so that more signal is recovered by the refinement. We find that after the Weave search 97% of signals have the loudest candidate within $\delta f \pm 6.3 \times 10^{-5}$ Hz, $\delta \dot{f} \pm 1 \times 10^{-11}$ Hz/s and a sky patch with a radius 0.4 the radius of the GCT 90% confidence region. Per candidate, the first Weave refinement stage is ~ 70 times less computationally expensive than the GCT first refinement stage. This means for the same computational cost, we can reduce the $2\overline{\mathcal{F}}$ threshold to follow-up 70 times more candidates.

In an all-sky search with Weave, with $N_{eff} = 0.69 N_{\text{templates}}$, a threshold of $2\overline{\mathcal{F}} > 6.155$ would produce 16 million false alarms in Gaussian noise. With a ~ 70 times faster follow-up search per candidate, we can afford to followup 1120 million candidates, corresponding to a $2\overline{\mathcal{F}}$ threshold of 5.91. Decreasing $2\overline{\mathcal{F}}_{\text{threshold}}$ from 6.155 to 5.91 opens the possibility of detecting a GW strain which is $\approx 6\%$ smaller, as the minimum detectable h_0^2 is proportional to the non-centrality parameter, $2\overline{\mathcal{F}} - 4$. The $2\overline{\mathcal{F}}$ thresholds for the comparison in this paper are shown in Table IV.

Search	$2\overline{\mathcal{F}}_{\text{threshold}}$
GCT	6.17
Weave	5.91

TABLE IV. A signal candidate must have $2\overline{\mathcal{F}} > 2\overline{\mathcal{F}}_{\text{threshold}}$ to be passed to the first refinement stage. The GCT threshold is the same as used for the MDC [26]. The Weave threshold was derived in Sections VI A and VI D 1.

2. MDC detection efficiency

We are primarily concerned with the ability of the searches to recover CW signals. The detection efficiency is the fraction of signals which are considered detected, and it is the benchmark that we will use to compare pipeline performance. The detection efficiency is measured as a function of the inverse of signal strength, h_0, normalised to the noise of the detector $\sqrt{S_h}$.

Unlike [26], we only consider injections which do not overlap with known detector artefacts. For this comparison we are only concerned with the ability of the Weave search method to recover signal in well behaved data. The detection efficiency, measured for MDC injections which do not overlap with a known noise line, is shown in Figure 4. This can be compared directly with Figure 3 in [26].

In Figure 4, Weave appears to be more sensitive than GCT. However, the uncertainty region is too large to quantify the improvement. In the next section we compare the detection efficiency with the new set of fake signals described in Section IV.

The dependence of the detection efficiency on signal strength is obtained with a sigmoidal fit to the MDC results. A non-linear least squares fit is performed using the *scipy.optimize.curve_fit* package in python, in the same way as done for the MDC.
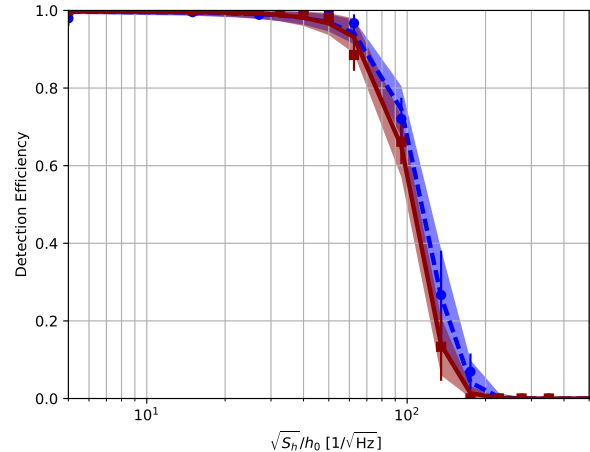

FIG. 4. Measured detection efficiency of Weave (blue dashed line) and GCT (red solid line) for the MDC injections. The curves and error bands are obtained by fitting sigmoids to the data, see Section VI D 2. Note that in [26] the signal strength is normalised by a different definition of S_h, where the harmonic sum over per-detector PSDs is used instead of the harmonic mean. One can approximately convert the MDC results into sensitivity depth at any efficiency by multiplying by $\sqrt{N_{\text{det}}} \approx 1.4$, as 2 detectors were used.

3. Detection efficiency in Gaussian noise

In this section we compare the detection efficiency of the Weave and GCT search methods using signals in Gaussian noise. The signal parameters are described in Section IV. The main difference with the MDC signals is that these signal are chosen to be randomly uniform in $\sqrt{S_h}/h_0$ ($\frac{1}{\sqrt{\text{Hz}}}$), so that at any fixed $\sqrt{S_h}/h_0$ ($\frac{1}{\sqrt{\text{Hz}}}$) the distribution of $\cos \iota$ is uniform. This results in a lower apparent efficiency than measured in Figure 4, where the signals were chosen to be randomly uniform in SNR, so that at a fixed $\sqrt{S_h}/h_0$ ($\frac{1}{\sqrt{\text{Hz}}}$) there are more signals at higher $|\cos \iota|$.

While the distribution of detection efficiency in Figure 4 appears to have a sigmoid shape, we find that when we increase the precision of our efficiency measurement, the distribution is not well described by a sigmoid. This has also been discussed in [24]. To estimate the sensi-

tivity depth at 90% detection efficiency we perform a linear fit between the surrounding data points, again, by performing a non-linear least squares fit using the *scipy.optimize.curve_fit* package in python. The uncertainty band is obtained with a linear fit to the binomial uncertainty on the measured efficiency. This straight line is not expected to desribe the measured efficiency across a broad range of sensitivity depths.

At 90% detection efficiency, Weave is 14% more sensitive than GCT with a sensitivity depth of 57.9 ± 0.6 $\frac{1}{\sqrt{\mathrm{Hz}}}$ at 90% detection efficiency, compared to $50.8^{+0.7}_{-1.1}$ $\frac{1}{\sqrt{\mathrm{Hz}}}$ for GCT. If we assume CW sources are isotropically distributed, this corresponds to a 50% increase in the volume of sky where we are sensitive.
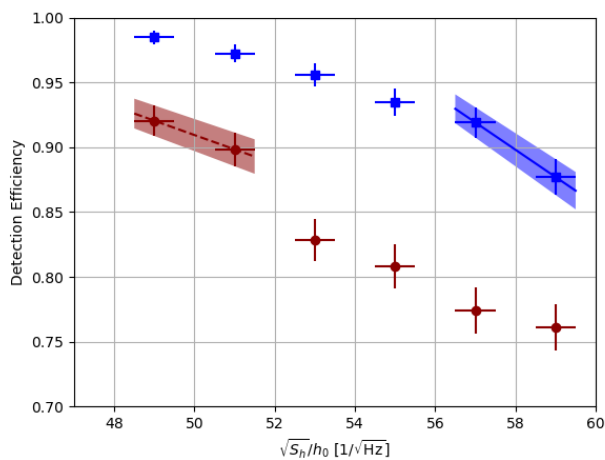

FIG. 5. Measured detection efficiency for Weave (blue dashed line) and GCT (red solid line) for injections in Gaussian noise which are uniform in $|\cos \iota|$ at fixed $\sqrt{S_h}/h_0$ ($\frac{1}{\sqrt{\mathrm{Hz}}}$).

VII. HIERARCHICAL FOLLOW-UP OF INTERESING CANDIDATES

As mentioned in Section VI D 1, above-threshold candidates from an all-sky search are passed through hierarchical refinement stages designed to exclude candidates from noise while retaining signal candidates [6, 7]. At each refinement stage, a search is performed in the parameter space around the candidate with a higher coherent segment length and/or a finer search grid than the previous stage.

At each refinement stage, the uncertainty on the signal parameters from the previous stage determines the region in parameter space which needs to be searched. In Section VI D 1 we show that the Weave search recovers candidates closer to the true signal parameters than the GCT search. With the smaller parameter uncertaintly, we would have more computing power to invest in using a larger segment length and finer grids at each refinement stage.

We consider the possibility of using the Weave search from Table I, to follow-up above threshold candidates from an all-sky search with the GCT search from Table II. Such a follow-up could search, for example, $\pm 2.5 \times 10^{-4} Hz, \pm 4 \times 10^{-11} Hz/s$ and a sky patch with a radius of 2.0 initial-search sky grid bins around each candidate. This is larger than the 90% confidence region mentioned in VI D 1.

The runtime for the Weave search over this parameter space is less than 10 seconds, and the memory is 50 Mb. For a computing cluster with 5000 CPUs, it would take 9 hours to follow-up the 16 million candidates from GCT with this Weave search. This small investment would result in candidates recovered much closer to the true signal parameters.

Each stage of the hierarchical refinement procedure requires the user to determine the search setup for that stage. This is done using the same extensitve Monte Carlo simulations used to determine the search setup for the initial all-sky search.

VIII. CONCLUSION

We have considered the state-of-the art Weave search method in its current implementation for Einstein@Home all-sky searches for continuous waves from isolated NS. We compare this new method with the GCT search method currently used in Einstein@Home searches. The comparison is initially made by performing the MDC which has previously been used to compare the methods used for all-sky searches for isolated NS in LIGO and Virgo data [26]. However, for a more precise comparison we generate a larger set of signals in Gaussian noise.

We compare the detection efficiency of an Einstein@Home all-sky search using the GCT search method and the new Weave search method. We find that the Weave search achieves 14% higher sensitivity depth than the GCT search. This increase translates to a 50% increase in the volume of parameter space within which we are sensitive to a CW signal.

The Weave search designed here, for 9 months of LIGO S6 data, is also expected to be optimal for an all-sky search over the 9 months of LIGO O2 data, though in O2 there is a 1 month gap which may affect the optimisation procedure.

While we measure an improvement in sensitivity, the Weave method has disadvantages when considered for a search on Einstein@Home. As explained in Section VI B, the Weave search requires more memory than a GCT search. This memory needs to be well constrained before deploying the search on volunteer computers, some of which have limited memory capabilities. With the Weave search considered here, limiting the memory to something which can be deployed on Einstein@Home results in a 50% increase in the search runtime. This memory limitation could potentially be mitigated by adapting

the Weave software for multi-threading.

Furthermore (Section VI C), including the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic in an all-sky search with Weave would increase the runtime and memory of the search by a factor of 2 to 3. If less than $\approx 5\%$ of frequency bands are disturbed, the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic could be applied to recover disturbed bands, however, if 10-20% of the frequency bands contain disturbances, rerunning with the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic becomes prohibitively expensive.

A major benefit of the Weave method is the potential to predict the search sensitivity, which should make it possible to semi-analytically determine optimal search setups. Work is ongoing to apply the optimization framework to the Weave method, which would greatly simplify the process of setting up an Einstein@Home search. The Weave memory model could be incorporated into the optimization framework to find a search setup with low memory requirements. In addition, having an semi-analytic optimization framework would allow us to consider a broader range of search options, as we would not be limited by the time and computing cost of extensive injections-and-recovery studies.

We find that Weave is promising for the hierarchical follow-up of candidates from an all-sky search. These candidates tend to come from regions of parameter space which have been identified as not being disturbed by detector artefacts, so the search can be run without using the $\hat{\beta}_{\mathrm{S/GLtL}}$-statistic. These follow-ups tend to run on clusters, so can handle larger memory footprints. However, since these searches cover a small region of param-

eter space around each candidate, the Weave memory requirement is small. Since the signal parameter recovery is more accurate with Weave, Section VI D 1, each follow-up stage would cover a smaller parameter space than if the candidates were from a GCT search. This allows for more sensitive follow-up searches at each stage than would be possible with a GCT search. Finally, we need to perform injection-and-recovery searches to determine the search setup to use for each stage. When the semi-analytic optimization framework can be applied to determine the best Weave search setup without injection-and-recovery studies the design of hierarchical follow-up searches will be greatly simplified.

We conclude that the GCT method remains the best search method for deployment on the Einstein@Home computing project, due to its lower memory requirements, while Weave presents significant advantages for the subsequent hierarchical follow-up searches of interesting candidates.

[1] The LIGO Scientific Collaboration, "LIGO: the Laser Interferometer Gravitational-Wave Observatory", *Rept. Prog. Phys* 72 (2009) 076901

[2] The LIGO Scientific Collaboration, "Advanced LIGO", *Class. Quant. Grav.* 32 (2015) 074001

[3] The Virgo Collaboration, "Status of Virgo", *Class. Quant. Grav.* 25 (2008) 114045

[4] The Virgo Collaboration, "Advanced Virgo: a 2nd generation interferometric gravitational wave detector", *Class. Quant. Grav.* 32 (2015) 024001

[5] The LIGO Scientific Collaboration and the Virgo Collaboration, "Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data running on the Einstein@Home volunteer distributed computing project", *Phys. Rev. D* 94 (2016) 102002

[6] 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", *Phys. Rev. D* 94 (2016) 122006

[7] The LIGO Scientific Collaboration and the Virgo Collaboration, "First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data", *Phys. Rev. D* 96 (2017) 122004

[8] The LIGO Scientific Collaboration and the Virgo Collaboration, "All-sky search for periodic gravitational waves

in the O1 LIGO data", *Phys. Rev. D* 96 (2017) 062002

[9] The LIGO Scientific Collaboration and the Virgo Collaboration, "Full Band All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data", *Phys. Rev. D* 97 (2018) 102003

[10] S. Frasca, P. Astone and C. Palomba, "Evaluation of sensitivity and computing power for the Virgo hierarchical search for periodic sources", *Class. Quantum Grav.* 22 (2005) S1013

[11] P. R. Brady and T. Creighton, "Searching for periodic sources with LIGO. II. Hierarchical searches", *Phys. Rev. D* 61 (2000) 082001

[12] B. Krishnan et al, "The Hough transform search for continuous gravitational waves", *Phys. Rev. D* 70 (2004) 082001

[13] B. Krishnan, "Wide parameter search for isolated pulsars using the Hough transform", *Class. Quant. Grav* 55 (2005) S1265-S1276

[14] einsteinathome.org

[15] The LIGO Scientific Collaboration and the Virgo Collaboration, "Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data", *Phys. Rev. D* 87 (2013) 042001

[16] H. J. Pletsch and B. Allen, "Exploiting Large-Scale Correlations to Detect Continuous Gravitational Waves",

Phys. Rev. Lett. 103 (2009) 181102

[17] K. Wette and R. Prix, "Flat parameter-space metric for all-sky searches for gravitational-wave pulsars", *Phys. Rev. D* 88 (2013) 123005

[18] K. Wette, "Lattice template placement for coherent all-sky searches for graviational-wave pulsars", *Phys. Rev. D* 90 (2014) 122010

[19] K. Wette, "Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars", *Phys. Rev. D* 92 (2015) 082003

[20] K. Wette, "Empirically extending the range of validity of parameter-space metrics for all-sky searches for gravitational-wave pulsars", *Phys. Rev. D* 94 (2016) 122002

[21] K. Wette, S. Walsh, R. Prix and M. A. Papa, "Implementing a semicoherent search for continuous gravitational waves using optimally-constructed template banks", *Phys. Rev. D* 97 (2018) 123016

[22] The LIGO Scientific Collaboration, "LIGO Algorithm Library - LALSuite", *free software (GPL)* (2018)

[23] K. Wette et al, "OctApps: a library of Octave functions for continuous gravitational-wave data analysis", *JOSS* 3(26) (2018) 707

[24] C. Dreissigacker, R. Prix, and K. Wette, "Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches", *Phys. Rev. D* 98 (2018) 084058

[25] R. Prix and M. Shaltev, "Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost", *Phys. Rev. D* 85 (2012) 084010

[26] S. Walsh et al., "A comparison of methods for the detection of graviataional waves from unknown neutron stars", *Phys. Rev. D* 94 (2016) 124010

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

[28] B. Owen, "How to adapt broad-band gravitational-wave searches for r-modes", *Phys. Rev. D* 82 (2010) 104002

[29] C. Cutler and B. F. Schutz, "The generalized F-statistic: multiple detectors and multiple GW pulsars", *Phys. Rev. D* 72 (2005) 063006

[30] The LIGO Scientific Collaboration and the Virgo Collaboration, "Characterization of the LIGO detectors during their sixth science run", *Class. Quantum Grav.* 32 (2015) 115012

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

[32] K. Wette, "Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars", *Phys. Rev. D* 85 (2012) 042003

[33] The LIGO Scientific Collaboration, "First search for gravitational waves from the youngest known neutron star", *ApJ* 722 (2010)

[34] D. Keitel, R. Prix, M. A. Papa, P. Leaci, M. Siddiqi, "Search for continuous gravitational waves: Improving robustness versus instrumental artifacts", *Phys. Rev. D* 89 (2014) 064023

[35] D. Keitel, "Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients", *Phys. Rev. D* 93 (2016) 084024