

A new veto for continuous gravitational wave searches

Sylvia J. Zhu^{1,2,a}, Maria Alessandra Papa^{1,2,3,b}, Sinéad Walsh³

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

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

³ *University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

(Dated: July 18, 2017)

We present a new veto procedure to distinguish between continuous gravitational wave (CW) signals and the detector artifacts that can mimic their behavior. The veto procedure exploits the fact that a long-lasting coherent disturbance is less likely than a real signal to exhibit a Doppler modulation of astrophysical origin. Therefore, in the presence of an outlier from a search, we perform a multi-step search around the frequency of the outlier with the Doppler modulation turned off (DM-off), and compare these results with the results from the original (DM-on) search. If the results from the DM-off search are more significant than those from the DM-on search, the outlier is most likely due to an artifact rather than a signal. We tune the veto procedure so that it has a very low false dismissal rate. With this veto, we are able to identify as coherent disturbances $>99.9\%$ of the 6349 candidates from the recent all-sky low-frequency Einstein@Home search on the data from the Advanced LIGO O1 observing run [1]. We present the details of each identified disturbance in the Appendix.

I. INTRODUCTION

In searches for continuous gravitational waves (CWs) from rotating neutron stars with asymmetries, detector artifacts can partly mimic the behavior of astrophysical signals and yield high values of the detection statistic, so that these artifacts look more like signals than Gaussian noise. In recent years, detection statistics that are more robust to lines (the line-robust statistic [2]) and transient disturbances (the line-and-transient-robust statistic [3]) have been developed. However, in spite of the fact that these exhibit higher detection efficiencies in the presence of the disturbances that they designed to be robust against, searches that have used such statistics can still suffer from a large number of loud candidates caused by detector artifacts [1].

A standard CW signal is approximately monochromatic with a phase evolution that is characterized by a frequency f_0 and its time-derivatives, e.g., \dot{f}_0 . The frequency f at which a detector on Earth receives the signal is Doppler shifted over time with respect to f_0 as the Earth rotates and orbits around the Sun.

Conceptually, when searching for a signal from a given sky location, the data analysis techniques correct the data for the Doppler shift and “reconcentrate” all the signal power in a narrow frequency range — optimally, a single frequency bin — at f_0 .

Coherent detector artifacts of terrestrial origin do not exhibit this Doppler modulation but can still match a

signal wave-shape well enough to yield a significant value of the detection statistic in a search for astrophysical signals. However, in the results of a search for waveforms without Doppler modulation, the significance of an artifact should increase with respect to the results of an astrophysical search whereas the significance of an astrophysical signal should decrease. We can use this difference in behaviors to set up a veto.

The standard CW waveforms are derived from astrophysical modeling of neutron stars and their general form is well constrained. In contrast, detector artifacts can take many forms. Some artifacts have known causes, and if their existence is known ahead of time, data from the affected frequency ranges can be and often are replaced with Gaussian noise (or “cleaned”; e.g., [1, 4, 5]). However, many artifacts only become apparent after a search has been performed; this is especially true when new data are analyzed.

It is impossible to quantify the performance of a veto that tests for detector disturbances other than by using data containing the coherent disturbances whose detrimental impact we want to mitigate. In its current form, this veto was developed for the Einstein@Home all-sky low frequency search of the data from the first observing run of Advanced LIGO (O1) [1]. The veto characterization and exploration were performed with this search in mind, so many specific aspects of it (e.g., grid spacings, frequency band) reflect the [1] search setup. However, the general principles are not search dependent, and this veto can easily be modified to accommodate other search setups.

This paper is organized as follows: Section II describes the veto procedure; Section III demonstrates its effectiveness in a real-world application; and Section IV summa-

^aemail: sylvia.zhu@aei.mpg.de

^bemail: maria.alessandra.papa@aei.mpg.de

rizes and discusses results and prospects.

II. VETO CHARACTERIZATION

A. Detection statistic

This method can be used to test candidates from any continuous wave search. However, we concentrate here on Einstein@Home-like searches, which have typically included the deepest investigations [6].

Einstein@Home searches for continuous waves typically search over 10^{17} different templates and return only the top candidate per million. These results must then be analyzed with automated methods that rely on the data being reasonably well behaved, so all frequency bands that contain visible large-scale patterns are set aside and not analyzed upfront. Hence, the remaining data are fairly “clean.” The line-and-transient-robust statistic [2, 3] is used to select the candidates that are most likely to be signals. An estimate of their significance against Gaussian noise fluctuations is performed using the standard average $2\bar{\mathcal{F}}$ statistic (where the average is taken over the different segments of the semi-coherent search) [7, 8].

In this paper, we concentrate on applying the veto to candidates that have survived a whole hierarchy of follow-ups, the last stage of which is a fully coherent search over the entire observation time [1, 6]. In this case, the reference detection statistic is $2\mathcal{F}$ rather than the average $2\bar{\mathcal{F}}$. In stationary Gaussian noise, the $2\mathcal{F}$ statistic follows a central χ^2 distribution with four degrees of freedom (χ_4^2). If a signal is present, the observed distribution is a noncentral χ^2 distribution with four degrees of freedom and a noncentrality parameter ρ^2 that depends on the strength of the signal ($\chi_4^2(\rho^2)$) [7].

B. The DM-off waveform

The heart of the DM-off veto is the comparison between the values of the detection statistic with an astrophysical waveform template bank and the values of the detection statistic from waveforms without Doppler modulation. The astrophysical waveforms not only present a frequency modulation due to the Doppler shift but also an amplitude modulation. This is due to the fact that interferometric detectors have non-uniform antenna sensitivity patterns across the sky and hence that the detector response to the same signal changes in time as the Earth rotates. The amplitude modulation depends on waveform parameters that are not explicitly searched for (hence called the nuisance parameters), because the \mathcal{F} -statistic is analytically maximised over them. We keep the amplitude modulation in the DM-off template waveforms for two reasons: The detection statistic remains a χ^2 distribution with 4 degrees of freedom, which allows for a straightforward comparison with the results of the

astrophysical search; and, the amplitude of the disturbance waveform is allowed to vary and potentially match a larger variety of disturbance types than a fixed amplitude modulation would allow.

How similar are the DM-on and DM-off template waveforms? To answer this question, we simulate 300 signals for an observation time and with a duty factor similar to that of the Advanced LIGO O1 observing run. The signal frequencies lie in the following ranges: 20-25 Hz, 50-55 Hz, and 90-95 Hz; the sky positions of the source are random in the sky; the cosine of the inclination angle is uniformly distributed; and the data has no noise. We search each signal using a fully coherent \mathcal{F} -statistic search with a search setup like that of the last follow-up stage of [1]. We use two template banks: For the DM-on search we use the standard astrophysical template waveforms, and for the DM-off search (our veto search) we use DM-off template waveforms. For the DM-on search, we use a single template that is perfectly matched to the signal. For the DM-off search, we use the template bank described in Section II D. From each search we consider the highest detection statistic value: $2\mathcal{F}_{\text{DM-on}} = \rho_{\text{DM-on}}^2$ and $2\mathcal{F}_{\text{DM-off}} = \rho_{\text{DM-off}}^2$. We then compute the mismatch μ :

$$\mu = \frac{\rho_{\text{DM-on}}^2 - \rho_{\text{DM-off}}^2}{\rho_{\text{DM-on}}^2}. \quad (1)$$

Fig. 1 shows this mismatch at the three frequency ranges. As expected, higher mismatches are found for signals that experience larger Doppler shifts. This occurs at the ecliptic equator (where the relative motion between the Earth and the source is greatest) and at higher frequencies (as the difference between the Doppler-shifted frequency and the source frequency is proportional to the source frequency). We can see that these waveform families are very different: the average mismatch is very high ($>98\%$) for even the low frequency signals at 20 Hz.

We expect the mismatch to increase for search setups with longer ($>$ several months) observation times because the Doppler signature is even stronger for such signals. O1 had a particularly short duration, so DM-off waveforms corresponding to future LIGO observing runs will, in general, have even larger mismatches with the astrophysical signals.

C. Proof of principle

We illustrate the concept of this veto by comparing the $2\mathcal{F}_{\text{DM-on}}$ and the $2\mathcal{F}_{\text{DM-off}}$ of two of the ten candidates that survived the penultimate stage of a search for continuous wave signals on data from the last Initial LIGO run (S6) [6]. We consider candidate 1 and 6 of Table II of [6]. After the last stage, candidate 1 did not exceed the $2\bar{\mathcal{F}}$ detection threshold and further studies found an association with a comb of lines. In contrast, the detection statistic value for candidate 6 did exceed the detection

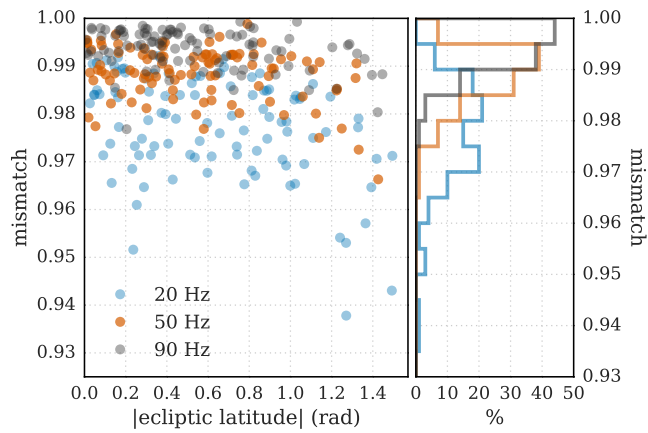


FIG. 1. The mismatch gives a measure of the dissimilarity between the DM-off waveforms and an astrophysical signal. The left panel here shows the mismatch as a function of the latitude of the source, where zero is at the equator and $\pi/2$ is at the poles. The right panel has the histograms for those mismatches. There is a dependence on the latitude and on frequency of the signal, as one would expect; sources at lower ecliptic latitudes and with higher frequencies have larger mismatches.

threshold; however, candidate 6 was due to a fake signal that was present in the data for validation purposes.

Figure 2 shows the original search results and the DM-off search results for these two candidates. The comparison is striking: the $2\bar{\mathcal{F}}$ statistic associated with candidate 1 in a DM-off search is a factor of 1.6 larger than the values from the original search, clearly showing that a non-astrophysical waveform matches the data better than the astrophysical one. The opposite happens for the candidate associated with a fake signal: when the data is searched with non-astrophysical waveforms, the significance of the candidate disappears and becomes comparable with the neighboring noise floor.

D. Implementation

In the previous section, we have always referred to the multi-detector $2\mathcal{F}$ (or average $2\bar{\mathcal{F}}$) statistic. This statistic coherently combines the data from all available detectors in order to determine whether the data are more likely to contain Gaussian noise or a waveform. However, high values of the detection statistic can result when the data contain some coherent disturbance — even only in one detector — that looks more like a signal than Gaussian noise. The multi-detector DM-off search finds the DM-off waveform that best matches the data consistently across the detectors; if a coherent disturbance occurs only in one of them, it also attempts to take into account the data from the other detector and may yield a detection statistic value that is not as significant as the single detector value and also not significant enough to exceed a veto threshold. However, if the data contain a signal,

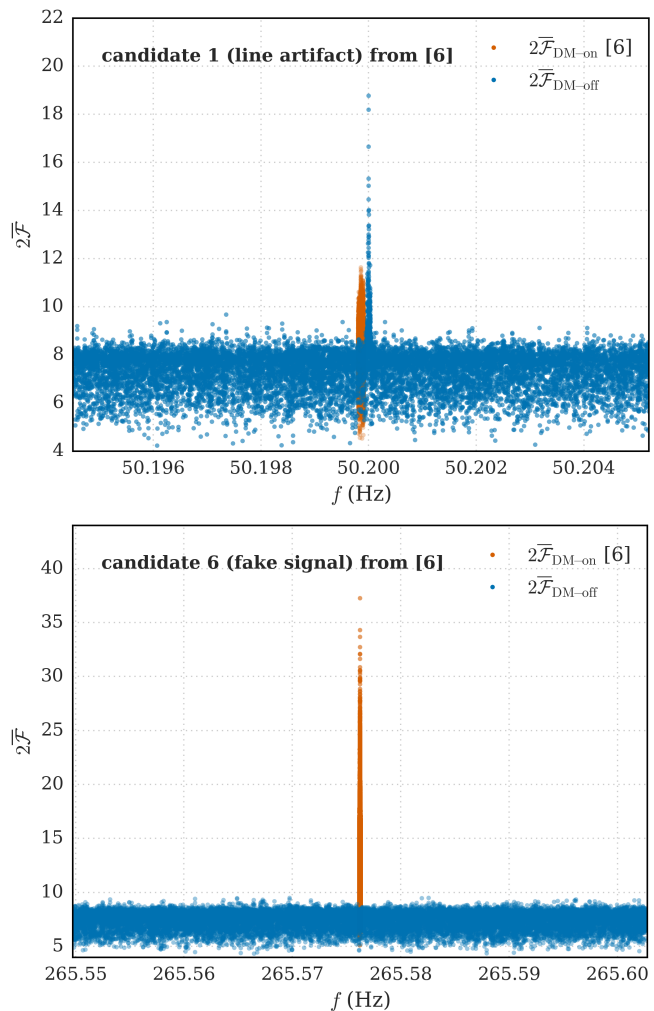


FIG. 2. Comparison between the original search results (orange) and the DM-off search results (blue) for candidates 1 and 6 of [6]. Candidate 1 (top) did not pass the detection threshold set at $2\bar{\mathcal{F}} = 15$ in [6]; candidate 6 (bottom) did, but was associated with a fake continuous wave signal in the data. After a DM-off search, the line that caused candidate 1 becomes more prominent while the signal from candidate 6 disappears into the noise. We use these two candidates to illustrate that the DM-off veto concept works.

both the multi-detector and single-detector $2\mathcal{F}_{\text{DM-on}}$ values will be larger than the corresponding $2\mathcal{F}_{\text{DM-off}}$ ones. Therefore, we perform the $2\mathcal{F}_{\text{DM-on}} - 2\mathcal{F}_{\text{DM-off}}$ comparison not only between the multi-detector statistics, but also between the single-detector ones. For a candidate to pass the veto, we require that all the DM-off $2\mathcal{F}$ statistics — both multi-detector and single-detector — successfully pass the thresholds.

The aim of the DM-off search is to check whether it is possible to produce a more significant detection statistic value from a non-astrophysical template bank using the same data (a very small frequency portion of the entire data set, on the order of mHz) that produced a significant candidate in the astrophysical-signal search. If this

is the case, then the original candidate is discarded. The starting point is the DM-on search candidate, defined by its parameters $(\alpha^*, \delta^*, f^*, \dot{f}^*)^{\text{DM-on}}$. Consider a DM-on search that explored the whole sky, a broad frequency band, and a spindown-range $(\dot{f}_{\min}, \dot{f}_{\max})$ with typically $|\dot{f}_{\max}| \ll |\dot{f}_{\min}|$ ($\dot{f}_{\min} < 0$). The DM-off search has to include the DM-off waveforms with the highest overlap with the candidate waveform $(\alpha^*, \delta^*, f^*, \dot{f}^*)^{\text{DM-on}}$. A conservative choice (meaning a choice that will include more waveforms than necessary, but will not exclude any) is to take the following:

- a frequency range Δf around f^* equal to $\Delta f_{(\text{spindown})} + \Delta f_{(\text{Doppler})}$, where $\Delta f_{(\text{spindown})} = \dot{f}^* T_{\text{obs}}$ is the maximum spindown frequency shift over the entire observation time and $\Delta f_{(\text{Doppler})}$ is the Doppler shift frequency during the observation time from a source at the sky position (α^*, δ^*) of the candidate;
- a broader spindown range than that of the original search: $(\dot{f}_{\min}, |\dot{f}_{\min}|)$, where $\dot{f}_{\min} < 0$;
- the whole sky.

There is a natural scale for template bank spacings of f and \dot{f} for a DM-off search over an observation time T_{obs} : $\frac{1}{T_{\text{obs}}}$ and $\frac{1}{T_{\text{obs}}^2}$, respectively. These are the smallest differences in f and \dot{f} that would be resolvable in a sinusoidal signal search over an observation duration T_{obs} . With the generous search ranges given above, such resolutions can result in high computing costs. On the other hand, the resolution in the sky is only moderate; it is determined by the rate of change across the sky of the detectors' antenna patterns (amplitude modulation functions; see Eqs. (12) and (13) of [7]). These have a periodicity of 12 hrs and, for the purpose of waveform overlap, can be considered constant over 30 minute intervals. This means that $O(100)$ points in the sky suffice to provide adequate sky resolution for the DM-off waveforms.

In a standard search, the grid spacings must be sufficiently fine in order to minimize the possibility of missing a signal; that is, the goal of a DM-on search is to find signal candidates. In contrast, the goal of a DM-off search is to *remove* noise-candidates, so coarser grids merely reduce the veto's power to identify artifacts. We can overcome this limitation by applying a finer DM-off grid in a later step, after we have eliminated most of the candidates, and in this way save computing cycles.

We have designed the DM-off veto in 3 steps. The first two steps use a search grid in f and \dot{f} that is 10×10 times coarser than the DM-on search grid. In all the DM-off steps we use a sky grid that comprises 45 sky points isotropically distributed on the celestial sphere.

1. **Step 1: Coarse grid, multi-detector.** A multi-detector DM-off search around each candidate is run using coarse f and \dot{f} grids. The largest value of $2\mathcal{F}_{\text{DM-off}}$ is considered and for a candidate to

pass to the next stage its detection statistic value must stay below a predetermined threshold.

2. **Step 2: Coarse grid, single-detector.** Since instrumental artifacts are not expected to be coherent across detectors, the DM-off search is then run on the data from the detectors separately on the surviving candidates from Step 1. A candidate's $2\mathcal{F}_{\text{DM-off}}$ values must be lower than the predetermined thresholds in *all* the detectors separately in order to not be vetoed; that is, if a candidate looks like an artifact in any detector, then it is discarded.
3. **Step 3: Fine grid, multi-detector and single-detector.** In this stage, the f and \dot{f} grid spacings are the same as the grids from the original search, and typically this means that f and \dot{f} are over-resolved. Assuming that the location in parameter space of the local maximum of $2\mathcal{F}_{\text{DM-off}}$ for the candidate has been identified in the previous stages, this stage is a refinement in its immediate neighborhood. 4000×4000 fine-grid points in frequency and spindown are explored, with both a multi-detector search and single-detector searches. The center point for each of these searches is the maximum that was recorded in the previous steps. If a candidate's $2\mathcal{F}_{\text{DM-off}}$ value is above the predetermined threshold in any of these three fine grid searches, then the candidate is discarded.

E. Thresholds and veto safety

We concentrate now on the thresholds of the veto in the case of candidates stemming from a fully coherent search like the last stage (FU3) of [1]; i.e., a 2-detector search on approximately 4 months of data. These studies would need to be repeated for candidates from a different search, but the general concept and the effectiveness of the procedure are well illustrated even in a particular case, like this one.

To determine the veto thresholds, we consider a population of 1500 fake signals in artificial Gaussian noise. We perform both a DM-on search and a DM-off search and compare the respective detection statistic values. We set the thresholds so that the veto is safe; i.e., no signal is discarded. We ensure that this happens at all the steps listed in Section II D.

Most of the fake signals we search have frequencies around 20 Hz, but we include 400 fake signals at 50 Hz and 90 Hz as well to check for consistency across frequencies. This reflects the population of FU3 candidates from [1], which are more abundant at lower frequencies. The $2\mathcal{F}_{\text{DM-on}}$ values of our fake signals vary between a few (consistent with noise) to over ten thousand; i.e., we sample the expected behavior for a wide range of signal strengths. Realistically, we expect any CW signal we detect to be on the lower end of this range, as otherwise they would have already been found by a previous search.

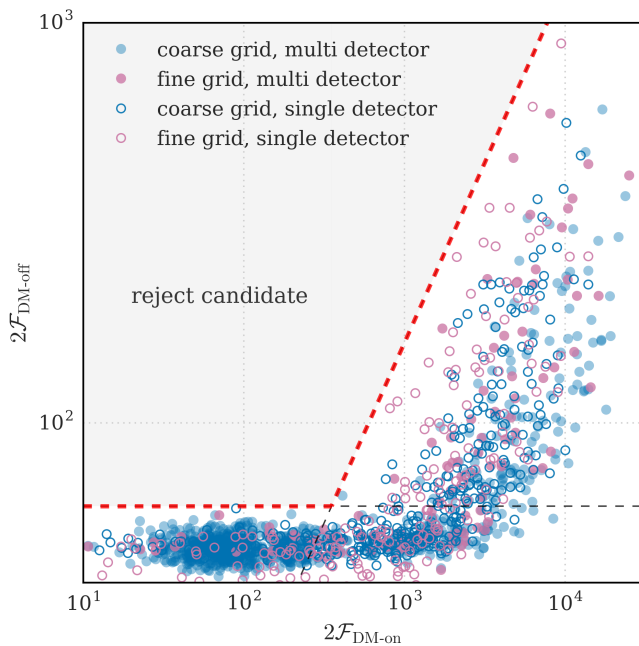


FIG. 3. Data containing fake noise plus 1500 simulated signals are searched with both the DM-off and the DM-on searches. The highest detection statistics are recorded as $2\mathcal{F}_{\text{DM-off}}$ and $2\mathcal{F}_{\text{DM-on}}$, respectively. For signals with $2\mathcal{F}_{\text{DM-on}} \lesssim 400$, the $2\mathcal{F}_{\text{DM-off}}$ is not distinguishable from the noise. The loudest value of $2\mathcal{F}_{\text{DM-off}}$ found for this group is 61.25. For louder $2\mathcal{F}_{\text{DM-on}}$, the $2\mathcal{F}_{\text{DM-off}}$ values are significantly above the noise and increase with $2\mathcal{F}_{\text{DM-on}}$. The veto threshold is indicated by the red dashed line, and was chosen so that none of the fake signals would have been rejected hence yielding a false dismissal rate of astrophysical signals smaller than one in 1500. Candidates that fall in the region above the red dashed line do not behave like signals, and therefore can be rejected.

The DM-off search parameter space and grids were described in the previous Section.

Figure 3 shows the $2\mathcal{F}_{\text{DM-off}}$ vs $2\mathcal{F}_{\text{DM-on}}$ values for the 1500 fake signals. For values of $2\mathcal{F}_{\text{DM-on}} \lesssim 400$, the $2\mathcal{F}_{\text{DM-off}}$ value is always less than 62 (the largest observed $2\mathcal{F}_{\text{DM-off}}$ is 61.25) and independent of the $2\mathcal{F}_{\text{DM-on}}$. In fact, the $2\mathcal{F}_{\text{DM-off}}$ values for these weaker signals are consistent with noise: When the Doppler modulation is turned off, the detection statistic value is lower than or comparable to the highest expected value over as many independent trials as there are templates in the bank, simply due to noise fluctuations.

For each injection, we determine the $2\mathcal{F}_{\text{DM-on}}$ value by running an FU3 astrophysical search on the simulated data in a neighborhood of the signal parameters and taking the loudest resulting detection statistic value.

Our simulation includes all of the steps in the DM-off veto described in Section IID. Since the results from the different DM-off searches do not differ significantly from each other, we use the same thresholds for all steps (Figure 3).

For candidates with $2\mathcal{F}_{\text{DM-on}} \leq 400$, we set a flat

threshold at 62; that is, we reject any candidate at any step with $2\mathcal{F}_{\text{DM-off}} \geq 62$. For candidates with $2\mathcal{F}_{\text{DM-on}} > 400$, we set a threshold based on their $2\mathcal{F}_{\text{DM-on}}$ values such that none of our injections would have been rejected. The threshold, then, is defined as the following:

$$2\mathcal{F}_{\text{DM-off}}^{\text{thr}} = \begin{cases} 62 & 2\mathcal{F}_{\text{DM-on}} \leq 400, \\ 10^{0.9 \log(2\mathcal{F}_{\text{DM-on}} - 0.5)} & 2\mathcal{F}_{\text{DM-on}} > 400. \end{cases} \quad (2)$$

We reject a candidate if $2\mathcal{F}_{\text{DM-off}} > 2\mathcal{F}_{\text{DM-off}}^{\text{thr}}$. This threshold is indicated in Figure 3 by the dashed red line. Based on our simulations, this threshold yields a false dismissal rate of less than one in 1500, or $< 0.07\%$.

III. EFFECTIVENESS OF THIS VETO WITH A WORKED EXAMPLE

In order to illustrate the effectiveness of this veto, we now show its noise rejection performance on real detector data. We cannot resort to synthetic data because only real data has the type of coherent disturbances that make this veto necessary in the first place, and these populations of disturbances are not characterized. Therefore, in this Section, we illustrate the application of the veto on the results of a low frequency search for continuous signals on LIGO data [1].

Compared to previous runs, the low frequency range of the data in this first Advanced LIGO run was plagued by many coherent disturbances. In spite of all the “cleaning” efforts aimed at identifying disturbed spectral regions up front, a large excess of detection statistic values above the pre-set threshold was found. A hierarchical procedure consisting of a clustering algorithm [9] and three follow-up searches reduced the number of candidates from over 15 million to several thousand. In the end, 6349 signal candidates survived the final fully coherent multi-detector follow-up stage (FU3). With the veto that we present in this paper, we are able to discard the overwhelming majority of these (6345 out of 6349).

The 6349 candidates are not distributed uniformly across the 20-100 Hz frequency range, but are instead clustered in 57 frequency bands (see left panel of Fig.9 in [1]). These 6349 candidates do not include 201 candidates stemming from a fake signal present in the data for pipeline-validation purposes, with frequencies at 52.8 Hz.

A. Veto application

After **Step 1**, 90% of the initial 6349 candidates are rejected as being consistent with instrumental artifacts. 653 candidates from 22 frequency ranges pass Step 1 (and 8 from the fake signal, all from the same frequency band).

Figure 4 shows an example of a candidate that does not survive Step 1. Based on the DM-off search results, we understand that this particular candidate, like all the others in its frequency region, is caused by a stationary

Stage	num surviving candidates	num remaining frequency bands
DM-on Stage 3	6349 (201)	57 (1)
DM-off Step 1	653 (8)	22 (1)
DM-off Step 2	101 (5)	10 (1)
DM-off Step 3	4 (1)	4 (1)

TABLE I. The results presented in [1] use the DM-off veto to test the 6349 candidates that survived multiple stages of an astrophysical search. In order to illustrate the performance of the veto, we present the details of that application. The numbers in parentheses denote the number of candidates or frequency bands attributed to a fake signal at 52.8 Hz. After applying the three steps of the DM-off veto, only four candidates survive (as well as a fifth candidate due to the fake signal at 52.8 Hz).

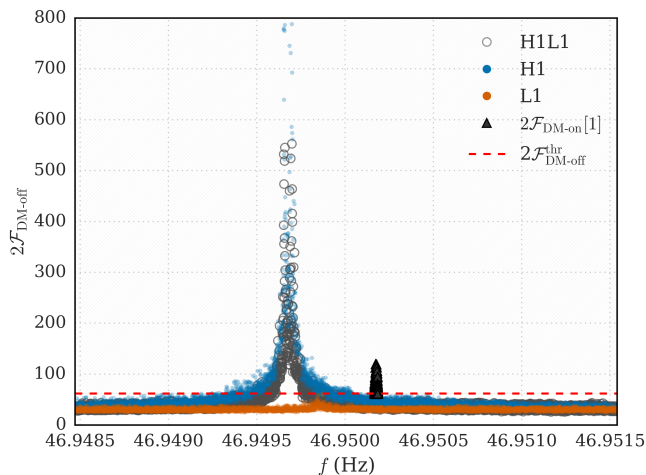


FIG. 4. An example of a candidate that does not survive Step 1. For this candidate, $2\mathcal{F}_{\text{DM-off}}$ in H1 exceeds the threshold (the horizontal red dashed line) and so this candidate is discarded. This particular candidate’s elevated $2\mathcal{F}_{\text{DM-on}}$ value is caused by a stationary line in H1.

line in H1. The candidate has a value of $2\mathcal{F}_{\text{DM-on}} = 80$ at the end of FU3; based on our simulations, signals that would produce $2\mathcal{F}_{\text{DM-on}}$ values such as this should have $2\mathcal{F}_{\text{DM-off}} < 62$; in contrast, for this candidate $2\mathcal{F}_{\text{DM-off}} \approx 560$ in H1, and so it is discarded. We note that the DM-off result for L1 is well below the threshold.

After **Step 2**, 85% of the 653 candidates are rejected, and 101 candidates from 10 frequency ranges (plus 5 candidates from the same frequency range due to the fake signal) pass to the next and final step. Figure 5 shows an example of a candidate that survives Step 1 but not Step 2. For this candidate, the H1 $2\mathcal{F}_{\text{DM-off}}$ value exceeds the threshold, but the L1 and the combined H1-L1 $2\mathcal{F}_{\text{DM-off}}$ values do not. In this particular example, the original $2\mathcal{F}_{\text{DM-on}}$ value is 79, while the highest $2\mathcal{F}_{\text{DM-off}}$ value, obtained by using only the H1 data (blue dots), is around 75 and comes from a template with $\dot{f} = 1.4 \times 10^{-10}$ Hz/s. We conclude that this candidate’s elevated $2\mathcal{F}_{\text{DM-on}}$ value

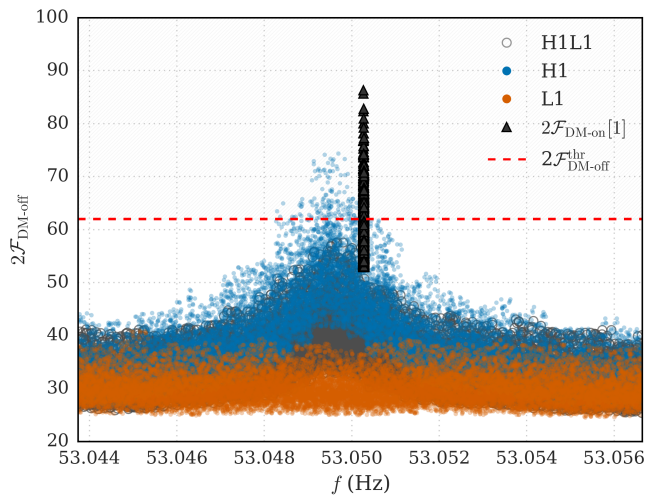


FIG. 5. An example of a candidate that does not survive Step 2. This candidate comes from a wandering line in H1.

is due to a wandering line in H1. Candidates such as this one illustrate the importance of requiring that a candidate’s single detector and multi-detector DM-off statistics must all pass the thresholds, and of not restricting the waveform model to stationary lines (i.e., $\dot{f} \neq 0$).

Only four candidates pass **Step 3** (Figure 6), plus one from the fake signal. These four candidates come from different frequency ranges.

Although this veto is optimized for computational efficiency, its cost is not negligible. Its application to the 6349 candidates used ≈ 6000 CPU-cores \times 4 hours on nodes of the Atlas cluster [10].

IV. CONCLUSION

The basic idea behind the DM-off veto is very simple: How likely is it that a candidate is due to a continuous *astrophysical* signal at a certain frequency compared to a continuous *terrestrial* signal in the same frequency region? As we have shown, we can simply “turn off” the Doppler modulation of the astrophysical signal and construct a family of coherent waveforms that has small overlaps with the astrophysical signals and models well the coherent disturbances in the data.

Based on this, we have developed a simple but highly effective veto to identify coherent disturbances among candidates of coherent continuous wave searches. The veto thresholds must be tuned depending on the coherent search set-up with simulations of the veto performance on astrophysical signals, and we show how to do this with a particular example in Section II E.

When this veto is applied to the 6349 candidates from the last and fully coherent stage of a hierarchical search [1], it was able to identify the overwhelming majority of these ($> 99.9\%$) as being due to noise artifacts. The results of the search, including the veto, are presented

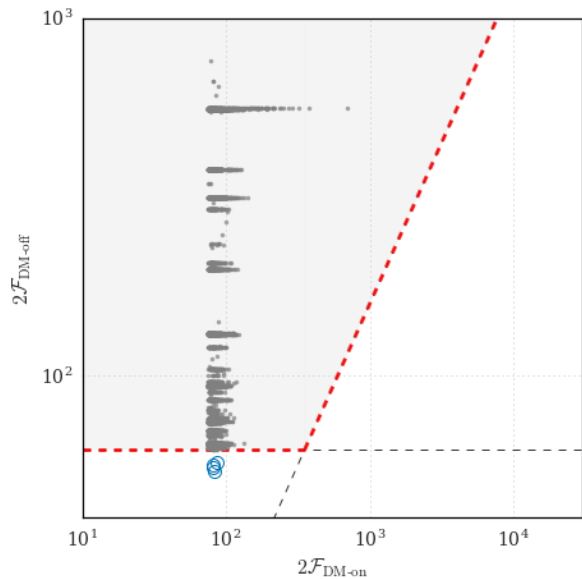


FIG. 6. $2\mathcal{F}_{\text{DM-off}}$ versus $2\mathcal{F}_{\text{DM-on}}$ for the 6349 candidates of the example that we illustrate in Section III. The dashed line shows the veto threshold of Eq. 2, and the pink and blue points are described in Figure 3. The region above the dashed line is the candidate-rejection region. Four signal candidates (blue circles) survive the DM-off veto, as well as fifth candidate that is due to a fake signal. The remaining 6345 candidates (gray dots) are rejected as being caused by detector artifacts.

in [1]. In this paper, we show examples from this veto application, in order to illustrate the method.

The veto is implemented in a hierarchy of three steps: if a candidate is discarded at step 1 (or 2) it is not subject to step 2 and (or) 3. This reduces the computational burden of this veto, which, however, still requires an average of ≈ 4 CPU-core hours per candidate.

After applying this veto, we not only discard spurious candidates in the astrophysical searches but also know more about the weak coherent artifacts in the data. We know where they lie in frequency, what detector is affected, and whether the line is stationary or wandering (and, through the $f_{\text{DM-off}}$, we have a measure of how much it moves). For instance, half of the 57 frequency ranges that produced the 6349 signal candidates contain stationary lines. In the Appendix, we provide a Table that characterizes the artifacts, based on this veto.

It may be possible in the future to use a modified ver-

sion of this veto in order to identify stationary lines before an astrophysical search, and to treat these regions differently or even preemptively exclude them from the analysis.

Not all hierarchical searches for continuous gravitational wave signals have a final fully coherent stage; see, for instance, [6]. In [6], no candidates survived Stage 4, which was a 22 segment (280 hrs coherent observation time) semi-coherent search. In this context, in a forthcoming paper, we will explore at what stage the semi-coherent version of this veto might be more effective than the standard hierarchy of semi-coherent follow-ups. In Section II C we ran the DM-off searches using the search setup of the first stage of [6], so we in fact have used a semi-coherent version of the DM-off veto to illustrate its concept. We are therefore confident that a more rigorous generalization will be possible.

In this paper we consider continuous signals from isolated neutron stars. However, the DM-off veto can be easily generalized to test candidates from other continuous sources, including neutron stars in binary systems.

ACKNOWLEDGMENTS

The DM-off veto procedure was used in [1], and we thank Sergey Klimenko and Evan Goetz for the review of the application of this new veto to the results of that search. We also thank Andrew Melatos, David Keitel, Greg Ashton and Grant Meadors for useful comments. M A Papa and S Walsh gratefully acknowledge the support from NSF PHY Grant 1104902. All computational work for this search was carried out on the ATLAS super-computing cluster at the Max-Planck-Institut für Gravitationsphysik, Hannover and Leibniz Universität Hannover. The authors thank to the LIGO Scientific Collaboration for access to the data and gratefully acknowledge the support of the United States National Science Foundation (NSF) for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. This document has LIGO (<https://dcc.ligo.org>) DCC number P1700114.

-
- [1] LIGO and Virgo Collaborations, “First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data,” (2017), submitted to PRD, arXiv:1707.02669 [gr-qc].
- [2] D. Keitel *et al.*, “Search for continuous gravitational waves: Improving robustness versus instrument arti-

facts,” Phys. Rev. D **89**, 060423 (2014).

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

- [4] A. Singh *et al.*, “Results from an all-sky high-frequency Einstein@Home search for continuous gravitational waves in the LIGO 5th Science Run,” *Phys. Rev. D* **94**, 064061 (2016).
- [5] B. P. Abbott *et al.*, “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**, 102002 (2016).
- [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 data,” *Phys. Rev. D* **94**, 122006 (2016).
- [7] 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).
- [8] C. Cutler and B. F. Schutz, “Generalized \mathcal{F} -statistic: Multiple detectors and multiple gravitational wave pulsars,” *Phys. Rev. D* **72**, 063006 (2005).
- [9] A. Singh *et al.*, “An adaptive clustering procedure for continuous gravitational wave searches,” (2017), submitted to PRD, arXiv:1707.02676 [gr-qc].
- [10] https://www.aei.mpg.de/24838/02_Computing_and_ATLAS.

1. Table of artifacts

H1 f (Hz)	H1 \dot{f} (Hz/s)	L1 f (Hz)	L1 \dot{f} (Hz/s)	comments
20.404	2.172×10^{-11}	—	—	
—	—	21.459	-7.036×10^{-13}	
26.176	-1.259×10^{-10}	—	—	broad feature
26.309	1.051×10^{-11}	—	—	loud stationary line
—	—	26.343	1.600×10^{-10}	
26.525	1.144×10^{-11}	—	—	
—	—	27.487	-7.263×10^{-11}	
27.591	4.150×10^{-10}	—	—	
—	—	27.842	-1.638×10^{-12}	
—	—	27.895	-6.890×10^{-11}	broad feature
28.947	2.099×10^{-12}	—	—	
31.141	8.244×10^{-11}	—	—	
—	—	31.379	2.099×10^{-12}	
31.402	-2.686×10^{-11}	—	—	
—	—	31.512	6.469×10^{-11}	
—	—	31.763	-7.036×10^{-13}	
33.332	-8.177×10^{-12}	33.333	-5.374×10^{-12}	very loud in L1
34.826	2.099×10^{-12}	—	—	
35.762	2.306×10^{-13}	35.763	2.306×10^{-13}	known cause (Acromag binary output chassis)
37.289	-3.506×10^{-12}	—	—	
37.310	1.165×10^{-12}	—	—	
—	—	39.763	-7.036×10^{-13}	
42.847	2.749×10^{-10}	—	—	
—	—	42.918	3.967×10^{-12}	
—	—	43.684	-7.036×10^{-13}	
44.703	-5.374×10^{-12}	—	—	loud stationary line
44.888	1.646×10^{-10}	—	—	
44.999	2.306×10^{-13}	—	—	
45.347	-9.330×10^{-10}	—	—	broad feature
46.928	2.524×10^{-10}	—	—	
46.950	-2.572×10^{-12}	—	—	
—	—	47.683	-3.506×10^{-12}	
48.969	1.319×10^{-10}	—	—	broad feature
50.250	-9.111×10^{-12}	—	—	
51.009	2.095×10^{-10}	—	—	
52.617	2.078×10^{-11}	—	—	
52.807	5.178×10^{-10}	52.807	2.898×10^{-10}	
53.050	1.497×10^{-10}	—	—	
53.385	-1.564×10^{-9}	—	—	broad feature
54.724	1.172×10^{-9}	—	—	broad feature
57.130	1.768×10^{-10}	—	—	
59.523	-1.098×10^{-11}	—	—	
59.604	-7.036×10^{-13}	—	—	
66.666	-1.191×10^{-11}	—	—	
66.757	3.506×10^{-12}	—	—	
66.877	5.836×10^{-12}	—	—	
74.505	-2.572×10^{-12}	—	—	loud stationary line
75.033	3.272×10^{-10}	—	—	broad feature
—	—	83.316	6.770×10^{-12}	
83.445	2.306×10^{-13}	83.447	-1.638×10^{-12}	
85.830	-2.572×10^{-12}	—	—	
89.406	-2.572×10^{-12}	—	—	loud stationary line
90.040	9.178×10^{-11}	—	—	broad, multi-peaked feature
96.348	-1.051×10^{-9}	—	—	no obvious artifact

TABLE II. Using **Step 2** of the DM-off search, we identify 53 detector artifacts in the Advanced LIGO detectors at frequencies between 20 and 100 Hz. For each artifact, we report the f and \dot{f} values that are associated with the highest $2\mathcal{F}_{\text{DM-off}}$ values in that frequency band if $2\mathcal{F}_{\text{DM-off}} > 62$. Most of these artifacts are weak stationary ($\dot{f} \lesssim 10^{-13}$ Hz/s) or wandering lines, but some are particularly loud or broad. We find 41 detector artifacts in H1 and 17 in L1. A few frequency regions contain artifacts in both detectors. One of the artifacts (at 35.76 Hz) has a known cause that was independently discovered during an observing run.