



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Adaptive clustering procedure for continuous gravitational wave searches

Avneet Singh, Maria Alessandra Papa, Heinz-Bernd Eggenstein, and Sinéad Walsh

Phys. Rev. D **96**, 082003 — Published 16 October 2017

DOI: [10.1103/PhysRevD.96.082003](https://doi.org/10.1103/PhysRevD.96.082003)

An adaptive clustering procedure for continuous gravitational wave searches

Avneet Singh^{1,2,3,†}, Maria Alessandra Papa^{1,2,4,‡}, Heinz-Bernd Eggenstein^{1,2}, Sinéad Walsh^{1,2,4}

¹Max-Planck-Institut für Gravitationsphysik, Albert Einstein Institute, am Mühlenberg 1, 14476, Potsdam-Golm

²Max-Planck-Institut für Gravitationsphysik, Albert Einstein Institute, Callinstraße 38, 30167, Hannover

³Leibniz Universität Hannover, Welfengarten 1, 30167, Hannover

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

In hierarchical searches for continuous gravitational waves, clustering of candidates is an important post-processing step because it reduces the number of noise candidates that are followed-up at successive stages [1][7][12]. Previous clustering procedures bundled together nearby candidates ascribing them to the same root cause (be it a signal or a disturbance), based on a predefined cluster volume. In this paper, we present a procedure that adapts the cluster volume to the data itself and checks for consistency of such volume with what is expected from a signal. This significantly improves the noise rejection capabilities at fixed detection threshold, and at fixed computing resources for the follow-up stages, this results in an overall more sensitive search. This new procedure was employed in the first Einstein@Home search on data from the first science run of the advanced LIGO detectors (O1) [11].

I. INTRODUCTION

In searches for continuous gravitational wave (CW) signals (e.g. [1–6, 12, 13, 15]), like in many other gravitational wave searches, the detection statistic can be triggered both by signals and by noise disturbances. Furthermore, when the signal or disturbance is strong, it typically does not trigger only a single template waveform but also many nearby ones.

‘Clustering’ is the procedure through which we assess elevated detection statistic template points close enough to each other in parameter space that might arise from the same root cause, i.e. signal or noise disturbance. The reason for doing this is that the clustering properties help discriminate candidates due to signals from the candidates due to disturbances, and in certain cases (e.g. loud disturbances), bundle together large numbers of candidates together which one does not need to assess separately. In case of hierarchical sub-threshold searches (e.g. [7, 12]), clustering is performed on the candidates from the first stage. This significantly reduces the number of candidates for subsequent follow-up at fixed threshold on the detection statistic. Hence, at fixed computing budget for the follow-up stages, clustering allows to lower the threshold and increase the sensitivity of the search.

In previous searches using a clustering procedure, the cluster volume was constant, based on the average clustering properties of signals [7, 12]. In this paper, we present a clustering method that is adaptive, i.e. it adapts the clustering size in each dimension to the local distribution of candidates in parameter space, and then it requires consistency in clustering among the different dimensions. We have named it AdCl procedure (*Adaptive Clustering Procedure*).

As the name suggests, the AdCl procedure adapts its parameters to the data. If the data were pure Gaussian noise, all this sophistication would not be necessary. Hence, in order to illustrate the AdCl under realistic and relevant condi-

tions, throughout this paper we use small (50 mHz) frequency-domain snippets of data from the first Advanced LIGO observing run (O1).

The paper is organised as follows. In section II, we introduce the fundamental idea behind the adaptive clustering procedure; in section III, we detail how it functions and introduce the various parameters that characterise it. In section IV, we present and compare the performance of this procedure against the clustering procedure used in previous searches. The last section summarises the main findings and discusses prospects.

II. CLUSTERING OF CANDIDATES

A typical all-sky CW search covers the entire sky, a large frequency range and a certain range of spin-down values. In this parameter space, grids are set up and a detection statistic is computed at each grid point.

We indicate a generic grid point with $\lambda_i \equiv (f_i, \dot{f}_i, \alpha_i, \delta_i)$, with $i = 1 \dots N$, and the detection statistic calculated at that grid point with Γ_i . Here, α_i, δ_i are the equatorial sky coordinates of the signal template, while f_i and \dot{f}_i denote the frequency and the first-order spin-down respectively. The result of the search are the ensemble of $\kappa_i \equiv (\lambda_i, \Gamma_i)$. We concentrate on the subset of these results that are interesting, i.e. where the detection statistic values are elevated above some predefined threshold (Γ_{\perp}). Let’s assume that there are M such results. We will refer to these as the candidates.

Operationally, the clustering procedure is an iterative process and it was first introduced in [7]: we begin with the highest detection statistic value in our results, corresponding to, say, candidate $\kappa_{i(1)}$, where “1” identifies the first iteration of the clustering procedure (i.e. the first cluster). The candidate $\kappa_{i(1)}$ is also called the *seed* for the first cluster. We then find elevated detection statistic values “nearby” λ_i , and we associate them with $\kappa_{i(1)}$. These set of points will form the first cluster, and they – along with the seed $\kappa_{i(1)}$ – will be referred to as the

[†]avneet.singh@aei.mpg.de; [‡]maria.alessandra.papa@aei.mpg.de

occupants of the cluster. We proceed to remove these occupants associated with $\kappa_{i(1)}$ from the original set of candidates. In the next iteration, we consider the candidate with the highest detection statistic value among the remaining candidates, now $\kappa_{i(2)}$, i.e. the seed for the second cluster. We again find elevated detection statistic values nearby $\kappa_{i(2)}$ and associate them with it. The occupants of the second cluster are again removed from the set of remaining candidates. This process is repeated with $\kappa_{i(3)}$, $\kappa_{i(4)}$, $\kappa_{i(5)}$ and so on. The process ends when we have no more seeds left above a certain predefined detection statistic threshold (Γ_S).

The core of the AdCl procedure lies in identifying an over-density of candidates in frequency and spin-down around each seed, in determining its extension, and in checking whether that set of candidates also presents an over-density in sky around its seed. These features are trademark signatures of signals (e.g. see Fig.1 top-panel), and they are not shared by most noise fluctuations/disturbances (e.g. see Fig.1 bottom-panel). We note that previous clustering algorithms did not require such coincident over-densities between frequency-spin-down and the sky, and hence, lead to a higher number of false alarms.

Furthermore, the AdCl procedure dynamically defines the clustering neighbourhood based on the data itself. In contrast, previous procedures derived a static clustering neighbourhood around the seed based on average clustering properties of the signals independently of the data. Thus, the AdCl procedure enables us to bundle together any over-density that extends over large volumes of parameter space as a single follow-up candidate, and hence, the number of candidates to follow-up from highly populated parameter space regions decreases significantly.

III. THE CLUSTER SIZE

A. A measure of distance in frequency and spin-down space (F-space)

The clustering is applied to a set of candidates χ_1 whose detection statistic value is above a certain threshold Γ_L ; Let's assume that there are M such candidates:

$$\chi_1 := \{\kappa_\ell\} \quad | \quad \Gamma_\ell \geq \Gamma_L, \quad (1)$$

where $1 \leq \ell \leq M$.

In general, at each iteration i , the clustering procedure defines a new cluster, and it does this by operating on a set of candidates χ_i . We indicate the seed for the i -th cluster with $\kappa_{\ell(i)}$, with $\ell(i)$ being the index that corresponds to the candidate with the loudest detection statistic value among the candidates in χ_i . We constrain the cluster seed to exceed a fixed threshold Γ_S , which in general is larger than Γ_L . The clustering procedure stops at iteration $N_c + 1$ when there are no more candidates with detection statistic values above Γ_S in χ_{N_c+1} , i.e. when $\Gamma_{\ell(N_c+1)} < \Gamma_S$.

At each iteration i , we define as Euclidean distance $\mathcal{R}_{i,k}^F$ in frequency and spin-down space (F-space) between the cluster

seed $\kappa_{\ell(i)}$ and every other candidate κ_k in χ_i :

$$\mathcal{R}_{i,k}^F := \sqrt{\left[\frac{f_k - f_{\ell(i)}}{\delta f}\right]^2 + \left[\frac{\dot{f}_k - \dot{f}_{\ell(i)}}{\delta \dot{f}}\right]^2} \quad \forall \quad \kappa_k \in \chi_i, \quad (2)$$

where, δf and $\delta \dot{f}$ are the frequency and spin-down grid spacings used in the search. Note that at fixed $\mathcal{R}_{i,k}^F$, (2) is an ellipse in F-space centered at $(f_{\ell(i)}, \dot{f}_{\ell(i)})$ and with axes of half-length $\delta f \times \mathcal{R}_{i,k}^F$ and $\delta \dot{f} \times \mathcal{R}_{i,k}^F$.

B. Distribution of distances in F-space

We define the cluster radius for the i -th cluster based on the distribution of the distances $\mathcal{R}_{i,k}^F$ in F-space. In order to derive such a distribution, we must bin the distances $\mathcal{R}_{i,k}^F$ appropriately.

The binning in F-space naturally takes the form of concentric elliptical annuli (f^r, \dot{f}^r) at distances $\mathcal{B}_{i,r}^F$ from the seed. The index r denotes the different bins. The edges (f^1, \dot{f}^1) of the first bin are defined by the equation

$$\sqrt{\left[\frac{f^1 - f_{\ell(i)}}{\mathcal{B}_{i,1}^F \delta f}\right]^2 + \left[\frac{\dot{f}^1 - \dot{f}_{\ell(i)}}{\mathcal{B}_{i,1}^F \delta \dot{f}}\right]^2} = 1. \quad (3)$$

The successive bins are defined by the recursive relation

$$[\mathcal{B}_{i,r+1}^F]^2 - [\mathcal{B}_{i,r}^F]^2 = [\mathcal{B}_{i,1}^F]^2 \quad \text{for all bins } r = 1, 2, 3 \dots, \quad (4)$$

which requires that the area of the annuli is constant and equal to $\pi[\mathcal{B}_{i,1}^F]^2$ (see Fig.2). Note that each annulus encloses an equal number of parameter space points. The relation (4) can be explicitly solved to yield

$$\mathcal{B}_{i,r}^F = \sqrt{r} \mathcal{B}_{i,1}^F \quad \text{for all bins } r = 1, 2, 3 \dots, \quad (5)$$

The value of $\mathcal{B}_{i,1}^F$ is chosen as

$$\mathcal{B}_{i,1}^F = \frac{1}{N_F} \mathcal{R}_{i,\max}^F \quad \text{with} \quad \mathcal{R}_{i,\max}^F = \max_k \{\mathcal{R}_{i,k}^F\}. \quad (6)$$

Comparing (6) with (5) and setting $\mathcal{R}_{i,\max}^F = \mathcal{B}_{i,r}^F$, we find that N_F^2 is equal to the total number of r bins. N_F is determined as follows: the candidate count $n_{i,r}^F$ in the various r bins is determined for a test value of N_F , say N_{F_t} . If the condition,

$$n_{i,1}^F(N_{F_t}) \geq C_F \langle n_{i,r}^F(N_{F_t}) \rangle \quad \text{where } r = 1, 2, 3 \dots, \quad (7)$$

is not satisfied, we iteratively decrease N_{F_t} by one ($N_{F_t} \rightarrow N_{F_t} - 1$) until (7) is verified, and set $N_F = N_{F_t}$. In (7), the angled brackets indicate the average over the r bins, and $C_F > 1$. Note that N_{F_t} should be large enough such that $\mathcal{B}_{i,1}^F$ is comparable with the signal containment region in F-space [12]. Further, C_F encodes the over-density requirement, and for low amplitude signals, this requirement is very lax: $C_F \gtrsim 1$, which means that the procedure picks the finest binning for which we at least do not have an under-density around the seed.




FIG. 1. (*top panel*) Distribution of \mathcal{F} -statistic values ($2\overline{\mathcal{F}}$) in the parameter space near a fake signal in noise (from LIGO O1 data). Note the elevated $2\overline{\mathcal{F}}$ values in the neighborhood of the injection. The elevated $2\overline{\mathcal{F}}$ values are clearly coincident in frequency-spindown and the sky. The location of the injection is marked with ‘+’.

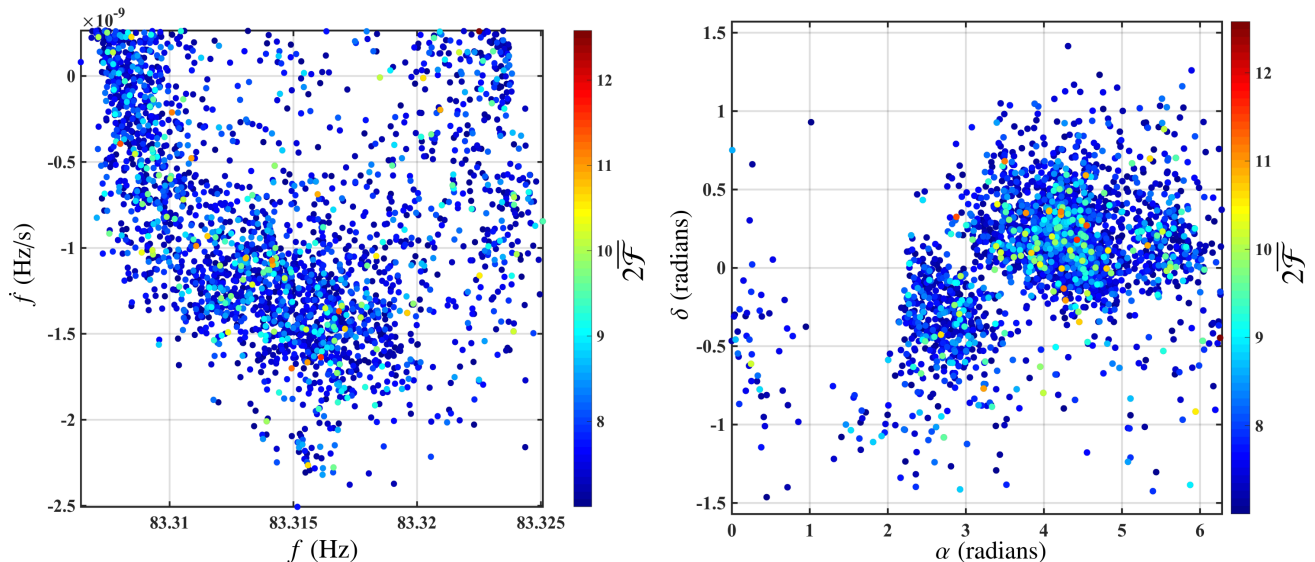


FIG. 1. (*bottom panel*) Distribution of \mathcal{F} -statistic values ($2\overline{\mathcal{F}}$) in the parameter space in the vicinity of a typical noise disturbance in LIGO O1 data. In contrast with Fig.1 (top-panel), the elevated $2\overline{\mathcal{F}}$ values due to the disturbance are not coincident between frequency-spindown and the sky.

We note that in a sub-threshold search, the clustering procedures are in principle sensitive to the parameter Γ_L : the overdensities of signal candidates due to a weak signal will only be observable down to certain detection statistic values, below which the density of noise candidates will be high enough that the over-density due to the signal candidates will not be appreciable. The threshold Γ_L could, in principle, be optimally placed at the level just above when this effect begins to take place. However, this is difficult to determine. By setting $C_F \gtrsim 1$, we appreciate the smallest over-density possible, and hence, ease the dependency of the procedure on Γ_L .

If for some i -th cluster, no resolution (no N_F value) can be found that meets the requirement of (7), then only the seed $\kappa_{\ell(i)}$ is removed from χ_i and the resulting set of candidates defines χ_{i+1} . The i -th cluster, ϕ_i , is classified as a *single-occupant-cluster*.

In Fig.3, we compare the distribution of $\mathcal{R}_{i,k}^F$ values from searches ran on noise data (blue curve), and on noise data plus a CW signal (red curve). The red distribution presents a clear maximum near the seed $\kappa_{\ell(i)}$, i.e. there is an evident over-density of candidates near the seed. We want to estimate the extent of this over-density, and cluster the candidates that form




FIG. 2. Example of annular binning in F-space, defined by (3)–(5). The values of the parameters are: $\delta f = 8.3 \times 10^{-7}$ Hz, $\delta \dot{f} = 1.3 \times 10^{-13}$ Hz/s. An ad-hoc value for $\mathcal{B}_{i,1}^F$ of 1.2×10^3 ($N_F = 50$) is taken. The area within each annuli is constant, hence the annuli get thinner with distance from the center. Note that the figure shows only the first 5 annuli for clarity; in total, there are $N_F^2 = 2500$ such annuli.




FIG. 3. Distribution of $\mathcal{R}_{i,k}^F$ for a noise-only data-set (red) and for a data-set also containing a fake signal (blue). Note the flatness of the distribution for data without any signal. The grid spacing $\mathcal{B}_{i,1}^F$ in F-space is defined by $N_F = 50$.

this over-density together.

C. Cluster size in F-space

For every i -th cluster, we use the distribution of distances $\mathcal{R}_{i,k}^F$ in order to determine the cluster radius in F-space: the cluster radius \mathcal{R}_i^{F*} is the value of the distance defined in (2) at which we have the first local minimum of $n_{i,r}^F$. Since the $\mathcal{R}_{i,k}^F$ histogram has typically many fluctuations, in order to estimate more accurately the position of the first minimum of the underlying distribution, we determine its shape with a fitting procedure which smoothens out the random fluctuations.

We fit the data $n_{i,r}^F$ in two stages. In the first stage, the data




FIG. 4. Fit to the distribution of $\mathcal{R}_{i,k}^F$ for a noise-only data-set (top panel) and for a data-set also containing a fake signal (bottom panel). The grid spacing $\mathcal{B}_{i,1}^F$ is defined by $N_F = 50$. We can see that the fitting procedure contours the shape of the distribution while ignoring small scale fluctuations, and returns a good measure of the over-density.

is separately fitted with two functions G (a superposition of Gaussians) and S (a superposition of sinusoids):

$$G(x) = \sum_{l=1}^{m_1} G_l(x); \quad S(x) = \sum_{l=1}^{m_2} S_l(x). \quad (8)$$

This step is implemented using a compiled MATLAB executable (using the package *fit*), which provides support for $m_1, m_2 \in [1, 8]$. For each fit, we choose the highest value of m_1 and m_2 that is able to fit the data within the standard tolerances defined by the program. The fitted curves G and S are summed and re-normalised, and the output is then fit again with a Gaussian function, yielding g_i^F . This second fit smoothens out the small scale fluctuations and leaves us with a clear view of the over-densities in F-space.

Finally, we can identify the local minimum of g_i^F closest to the origin and take that as the radius \mathcal{R}_i^{F*} in F-space of the i -th cluster. In Fig.4, we show an example of the fitting procedure on purely noise data (top panel) and in data containing a fake signal (bottom panel).

D. Hill parameters in F-space and further constraints

The distribution of candidates in parameter space is very diverse, depending on the nature of the noise in the data. Be-

cause of this, even an adaptive procedure, such as the one described above, may still generate clusters that spuriously assemble together candidates that are actually independent. In order to counter this, instead of setting more stringent criteria, for example a higher threshold C_F , it is more effective to produce a first estimate of the cluster based on liberal parameters, and then scrutinize its topological properties in detail, and further accept, discard or modify the cluster based on these. The topological properties that we consider for a cluster i are the so-called ‘‘hill parameters’’¹ *prominence* P_i , *dominance* D_i and *goodness* G_i :

$$P_i := \frac{\mathcal{R}_i^{F*}}{\max_k \{\mathcal{R}_{i,k}^F\}}, \quad (9)$$

$$D_i := \frac{g_i^F(\mathcal{B}_{i,1}^F) - g_i^F(\mathcal{R}_i^{F*})}{g_i^F(\mathcal{B}_{i,1}^F)}, \quad (10)$$

$$G_i := \frac{|\ln_{i,1}^F - g_i^F(\mathcal{B}_{i,1}^F)|}{\ln_{i,1}^F + g_i^F(\mathcal{B}_{i,1}^F)}. \quad (11)$$

The cluster candidates from the set χ_i are further inspected to check if:

$$P_i \leq P_{th}, \quad D_i \geq D_{th}, \quad G_i \leq G_{th}. \quad (12)$$

These thresholds (P_{th} , D_{th} , G_{th}) on the hill parameters restrict the topology of clusters: P_{th} restricts the fraction of the available parameter space that the cluster occupies; D_{th} bounds the contrast between the density of candidates near the seed and at the cluster edge; G_{th} specifies the minimum agreement between the fitted curve and the observed density near the seed. The nature and the values of the constraints in (12) is such that they exclude clusters that extend too far in the F-space, and at the same time, show very little contrast with respect to the local noise background; thus we shortlist the kind of clusters that we typically expect from signals.

When a cluster in F-space fails to meet any of the criteria given by (12), we shortlist candidates $n_{i,1}^F$ from the distribution that fall within the first bin $\mathcal{B}_{i,1}^F$ around the seed and discard all other candidates from the iteration. This is equivalent to resetting $\mathcal{R}_i^{F*} = \mathcal{B}_{i,1}^F$. This choice is justified because the failing of the hill parameters means that the shortlisted cluster is not topologically consistent with what we require from a cluster of that extent. However, the initial over-density still remains near the seed and it might be due to a low amplitude signal. In this regard, we do not discard the whole cluster. On the other hand, if the criteria in (12) are met, we shortlist all the candidates, including the seed, that fall within our estimated cluster radius \mathcal{R}_i^{F*} , and discard all other candidates outside the cluster radius.

The candidates clustered in F-space constitute the χ_i^{sky} set and their clustering properties in the sky are considered further.

E. A measure of distance in the sky (S-space)

We now want to determine whether the shortlisted candidates in χ_i^{sky} show any over-density in sky around the seed. If any over-density is found, the candidates constituting this over-density will form the final i -th cluster.

As in F-space, for each candidate $\kappa_k \in \chi_i^{sky}$, we introduce a distance in the sky, $\mathcal{R}_{i,k}^S$, to the seed of the i -th cluster under consideration:

$$\mathcal{R}_{i,k}^S := \sqrt{[x_k - x_{\ell(i)}]^2 + [y_k - y_{\ell(i)}]^2} \quad \forall \kappa_k \in \chi_i^{sky}. \quad (13)$$

This definition is justified when the search grids are uniform on some plane (x, y) , for example the ecliptic plane (e.g. see [12]) or the equatorial plane (e.g. see [4]). The transformation equations between the sky coordinates (α, δ) and (x, y) for a uniform grid on ecliptic plane (λ, β) , are:

$$\begin{cases} x = \cos \lambda \cos \beta \\ y = \sin \lambda \cos \beta, \end{cases} \quad (14)$$

with

$$\begin{cases} \lambda = \tan^{-1} \left[\frac{\sin \alpha \cos \varphi + \tan \delta \sin \varphi}{\cos \alpha} \right] \\ \beta = \sin^{-1} [\sin \delta \cos \varphi - \sin \alpha \cos \delta \sin \varphi]. \end{cases} \quad (15)$$

In the expressions above, $\varphi = 23.4^\circ$ is the angle of obliquity of the ecliptic with respect to the celestial equatorial plane². The ecliptic plane represents the S-space after this transformation.

F. Distribution of distances in S-space

The binning of the $\mathcal{R}_{i,k}^S$ values is performed in a similar fashion as previously done in F-space. The edges of the bins, labeled by r , of the i -th cluster, satisfy the following relation:

$$[\mathcal{B}_{i,r+1}^S]^2 - [\mathcal{B}_{i,r}^S]^2 = [\mathcal{B}_{i,1}^S]^2 \quad \text{for all bins } r = 1, 2, 3 \dots \quad (16)$$

This recursive relation describes concentric circular annuli in the (x, y) plane enclosing equal areas; the annuli naturally get thinner as we move away from the seed, as shown in Fig.2. The first bin is a circle and its area is proportional to $[\mathcal{B}_{i,1}^S]^2$.

$\mathcal{B}_{i,1}^S$ is chosen based on the clustering properties of signals. Precisely, it will depend on the 99% containment region of the search [12]. This region defines a neighbourhood around a cluster seed originating from a signal, within which the true signal parameters are contained with 99% confidence. If we indicate with d_{sky} the width of the search pixels in the (x, y) plane [4, 12], and with $N^{99\%}$ the diameter of the 99% containment region expressed in number of pixels, then we can express $\mathcal{B}_{i,1}^S$ as

$$\mathcal{B}_{i,1}^S := \frac{N^{99\%} + N_S}{2} d_{sky}, \quad (17)$$

¹We adopt the notion of hill parameters from the concept of ‘topographic prominence’ used in topography/geography, e.g. see [8].

²Note that in (15), λ must be translated to its correct quadrant by adding 180° or subtracting 180° .




FIG. 5. Distribution of $\mathcal{R}_{i,k}^S$ values for a noise-only data-set (red) and for a data-set also containing a fake signal (blue): $\mathcal{B}_{i,1}^S$ for the noise-only case is defined by $N_S = 0$, while for signal it is set to $N_S = 6$.

where, N_S is a parameter that has to be tuned as shown in section IV. Further, (17) says that the first bin in the sky, i.e. the circle with radius $\mathcal{B}_{i,1}^S$, contains all sky pixels within the 99% containment region, plus (or minus) a tuning term N_S . We continue binning recursively according to (16) until the width of the bin becomes smaller than a sky-grid pixel; all candidates lying more distant than that point are ignored.

We find that in disturbed data that contains a large number of noise outliers, a single value of N_S independent of the loudness Γ_i of the cluster seed under consideration, makes this clustering procedure very slow. The reason is that very large values of Γ_i are often associated with many candidates clustered in F-space (highly populated χ_i^F) that are distributed almost isotropically in the sky. In this situation, if the resolution in the sky ($\mathcal{B}_{i,1}^S$) is high, the sky-clustering step eliminates one candidate at the time as a single-occupant-cluster, and this is very inefficient. The solution is to decrease the resolution (increase $\mathcal{B}_{i,1}^S$ by increasing N_S) with Γ_i .

In Fig. 5, we show the re-normalised distribution of $\mathcal{R}_{i,k}^S$ for a fake signal and near-Gaussian noise.

G. Cluster in S-space

In order to estimate the cluster radius in S-space, we check for over-densities by analysing the distribution of $\mathcal{R}_{i,k}^S$.

If the first bin is the most highly populated (i.e. $n_{i,1}^S = \max_r \{n_{i,r}^S\}$), all the candidates contained within a distance \mathcal{R}_i^{S*} are clustered together:

$$\mathcal{R}_i^{S*} = \min_r \left\{ \mathcal{B}_{i,r}^S : \frac{n_{i,r}^S - n_{i,r+1}^S}{n_{i,r}^S} > C_S \right\}. \quad (18)$$

\mathcal{R}_i^{S*} is the smallest distance at which we have a relative drop in the density of candidates above a certain threshold C_S . All candidates within \mathcal{R}_i^{S*} constitute, together with the seed, the final i -th cluster, ϕ_i . The set of candidates considered for the next clustering iteration is $\chi_{i+1} = \chi_i - \phi_i$.

The value of C_S is chosen based on the localisation properties of signals and leaning on the conservative side, i.e. to-

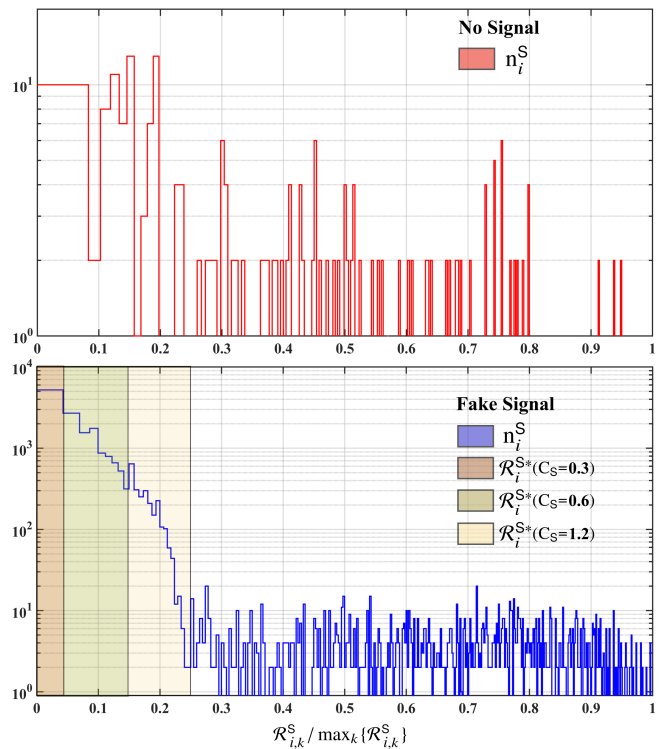


FIG. 6. Clustering properties in the sky for a data-set containing a signal (lower panel) and a noise data-set (upper panel). The values of N_S are 6 and 0, respectively for the lower and upper panels. Note the variation in \mathcal{R}_i^{S*} for given values of C_S . Note that the first bin in the pure noise case (top-panel) is not the most highly populated, and hence failed the clustering criteria ($n_{i,1}^S \neq \max_r \{n_{i,r}^S\}$), irrespective of the value of C_S .

ward lower values of C_S . For instance, in the bottom panel of Fig. 6, we see that $\mathcal{R}_i^{S*}(1.2) > \mathcal{R}_i^{S*}(0.6) > \mathcal{R}_i^{S*}(0.3)$. Indeed, the lower value of C_S clusters less candidates, but the candidates excluded at this iteration will likely form their own separate cluster at the next iteration. If this second set of points were due a signal, with a lower C_S , they would be associated to the correct seed.

If the first bin is not the most highly populated, the final cluster ϕ_i will contain only the seed $\kappa_{\ell(i)}$. All the other candidates remain un-clustered, and available for association with another cluster in the set $\chi_{i+1} = \chi_i - \kappa_{\ell(i)}$.

This recursive procedure continues until there are no more candidate seeds, i.e. no more candidates with detection statistic value above the threshold Γ_S . In Fig. 8, 9 in appendix A, we show a snapshot of the procedure for the first iteration on data χ_1 for a fake signal and LIGO O1 noise.

IV. PERFORMANCE

We characterise the performance of the AdCl procedure and compare it with the old clustering procedure, used in [12]. We show how the tuning parameters were chosen in an actual search [11], with parameters given in table I.

Quantity	Value
T_{obs}	4 months
T_{coh}	210 hours
N_{seg}	12
δf	8.3×10^{-7} Hz
$\delta \dot{f}$	1.3×10^{-13} Hz/s
$d_{\text{sky}} (f = 100\text{Hz})$	20 arcmin

TABLE I. The clustering procedure is applied to the output from this all-sky search.

The two clustering procedures are compared at the same value of seed threshold Γ_S , and with the other parameters optimally tuned.

A. Clustering parameters

We will consider two different data inputs to the clustering procedure, one suitable for a high-significance search (loud signals), and the other for a sub-threshold search (weak signals). In the former search, the detection statistic is $2\overline{\mathcal{F}}$, the corresponding Γ_S and Γ_L thresholds are 12.0 and 10.5 respectively, and the value of N_S for the i -th cluster is:

$$N_S(2\overline{\mathcal{F}}_i) = \begin{cases} 0 & \text{if } 2\overline{\mathcal{F}}_i < 18 \\ 2\overline{\mathcal{F}}_i - 18 & \text{if } 18 \leq 2\overline{\mathcal{F}}_i \leq 48 \\ 30 & \text{if } 2\overline{\mathcal{F}}_i > 48. \end{cases} \quad (19)$$

Alternatively, for the second search, the detection statistic is the line- and the transient line-robust statistic $\hat{\beta}_{S/\text{GLIL}}$ [9, 11], the corresponding Γ_S and Γ_L thresholds are 5.5 and 4, respectively, and N_S for the i -th cluster is:

$$N_S(\hat{\beta}_{S/\text{GLIL}}^i) = \begin{cases} 0 & \text{if } \hat{\beta}_{S/\text{GLIL}}^i < 15 \\ \hat{\beta}_{S/\text{GLIL}}^i - 15 & \text{if } 15 \leq \hat{\beta}_{S/\text{GLIL}}^i \leq 35 \\ 31 & \text{if } \hat{\beta}_{S/\text{GLIL}}^i > 35. \end{cases} \quad (20)$$

This is the set-up appropriate for a search like [11].

The reason why we consider searches with different detection statistics is historical: at the time when we started characterising the AdCl procedure, we were planning to use it for a high-significance search on quiet bands, as done in [4]. In this case, the simplest detection statistic to use is $2\overline{\mathcal{F}}$, and all the false alarm and detection efficiency studies were performed with this statistic. It was only later that we realised that the quality of the data in the low-frequency range was such that a high-significance search was not possible: we would have many candidates above threshold, and we would have to carry out a large scale follow-up. Due to these complications, the use of the $\hat{\beta}_{S/\text{GLIL}}$ was necessary. In the absence of large disturbances, the empirical relationship between the two detection statistics is $2\overline{\mathcal{F}} \equiv 0.419 \hat{\beta}_{S/\text{GLIL}} + 10.855$.

The other parameters are chosen as described in the previous sections and they are equal for both types of searches, and

their values are:

$$\begin{cases} N_F \in [25, 50], C_F = 1.2 \\ C_S = 0.25 \\ P_{\text{th}} = 0.25, D_{\text{th}} = 0.05, G_{\text{th}} = 0.1. \end{cases} \quad (21)$$

On the other hand, the old clustering uses a fixed cluster size corresponding to the 99% containment regions in the various dimensions. In case of the high-threshold $2\overline{\mathcal{F}}$ search:

$$2\overline{\mathcal{F}} \equiv \begin{cases} \Delta f = 1.15 \times 10^{-4} \text{ Hz}, \\ \Delta \dot{f} = 5.6 \times 10^{-11} \text{ Hz/s}, \\ \Delta^{\text{sky}} = 6 \times 6 \text{ sky-pixels}, \end{cases} \quad (22)$$

while for the sub-threshold $\hat{\beta}_{S/\text{GLIL}}$ search:

$$\hat{\beta}_{S/\text{GLIL}} \equiv \begin{cases} \Delta f = 1.85 \times 10^{-4} \text{ Hz}, \\ \Delta \dot{f} = 8.5 \times 10^{-11} \text{ Hz/s}, \\ \Delta^{\text{sky}} = 9 \times 9 \text{ sky-pixels}. \end{cases} \quad (23)$$

B. Safety

Naturally, the clustering procedure needs to be safe, i.e. it should not discard real signals. Thus, we choose the clustering parameters to yield the lowest false alarm rate for a very low false dismissal rate. We now show the detection efficiencies for the clustering parameters outlined in (19), (20), (21).

We estimate the detection efficiency by performing Monte-Carlo simulations of gravitational wave signals in real data taken from the LIGO O1 run. By using the real LIGO data instead of fake Gaussian noise, we derive a realistic benchmark of the performance. In a nutshell, the fake signals are added to the real data, the search is run, and the clustering procedure is applied.

The population of signals have parameters reasonably uniformly distributed in frequency, spin-down and sky-position, and with amplitudes that yield the detection statistic values shown in Fig.7.

The detection efficiency \mathcal{E} is defined as the ratio of the number of candidates from signals recovered by the clustering procedure with the total number of signals with detection statistic value above Γ_S . For a signal to be recovered by the clustering procedure, we require that the signal parameters lie within the 99% containment region of the seed parameters (we remind the reader that the detection statistic value of the seed must also exceed Γ_S). This means that if there were a follow-up stage on the cluster seeds, the true signal parameters would lie within the searched region, and if there were no follow-ups, the signal parameters would lie within the quoted parameter uncertainties.

C. Noise Rejection

We estimate the false alarm rate by applying the clustering procedure to the same search output data as described in the previous section, just without fake signals.




FIG. 7. Distribution of the values of the detection statistics of the sub-threshold and high-significance signals added to the data to characterize the performance of the clustering procedure. Note that the high-significance signal population is not purely high-significance; it also contains a few signals at low values of the detection statistic (less than 10% below $2\bar{\mathcal{F}} = 14.0 \equiv \hat{\beta}_{S/GLL} = 7.5$). Meanwhile, the sub-threshold search may be considered as purely sub-threshold (none above $\hat{\beta}_{S/GLL} = 10.5 \equiv 2\bar{\mathcal{F}} = 15.3$).

		AdCl Procedure	Old Procedure
High-significance $2\bar{\mathcal{F}}$ search	NR	65.9%	$\leq 40.0\%$
	\mathcal{E}	97.6%	95.1%
Sub-threshold $\hat{\beta}_{S/GLL}$ search	NR	90.5%	$\leq 74.1\%$
	\mathcal{E}	95.5%	$> 95.0\%$

TABLE II. Comparison of the noise rejection (NR) and the detection efficiencies (\mathcal{E}) of high-significance and sub-threshold searches between the new and old clustering procedures.

The input to the clustering procedure are \mathcal{N}_{in} candidates, with detection statistic values greater than Γ_S . At the output of the clustering procedure, we have \mathcal{N}_{out} candidates. We define the noise reduction factor NR as:

$$NR := 1 - \frac{\mathcal{N}_{out}}{\mathcal{N}_{in}} \quad \text{on noise.} \quad (24)$$

Naturally, $0 \leq NR \leq 1$, and higher values of NR denote lower number of noise candidates after the clustering procedure.

D. Results

The performance results for the AdCl and the old clustering procedures are shown in table II. For a high-significance search, the detection efficiency, exceeding 95%, is high for both the procedures, but the new clustering has a noise rejection which is significantly higher (nearly 66% versus 40%) than the one achieved by the previous method.

In a sub-threshold search, we set a low enough threshold on the detection statistic of the seed (Γ_S) such that we expect a large number of candidates to exceed this limit, just due to random noise. The underlying idea behind this is that with successive follow-up stages, one is able to weed out the noise and identify a signal that, at the first stage of the hierarchy, was hidden by a multitude of false alarms. In this regime, the clustering procedure operates in an environment of the most uniformly and densely populated candidates. The signal signature used by the clustering procedure are local over-densities around the cluster seed, coincident in F-space and S-space. But the cluster seed is, at every iteration i , the loudest candidate in the set χ_i , and when the signal is weak, i.e. its amplitude is comparable to the amplitude of many of the candidates, it might not be picked as a seed. For this reason, the detection efficiency is lower for a sub-threshold search with respect to a high-threshold search. In order to compare the performance of AdCl procedure with the old procedure, we fix the detection efficiency at $> 95\%$ by lowering Γ_L to 3.4 for the old procedure (keeping $\Gamma_L = 4.0$ for AdCl procedure). In this case, the AdCl procedure improves the noise rejection (NR) by 22% over the old procedure.

The results of table II refer to signal-frequency bands where the data is fairly uniformly distributed in parameter space, i.e. there are no extended regions of the parameter space that host enhanced values of the detection statistic values, as in the case of the top panel of Fig. 1. Moreover, the AdCl procedure performs very well in disturbed conditions, and this is important because the disturbed regions typically yield a lot of spurious candidates.

In noisy regions, the new clustering procedure has a NR of 98.9%, compared to $\leq 91.1\%$ for the old procedure in a $2\bar{\mathcal{F}}$ search. We expect similar results for noisy data in a $\hat{\beta}_{S/GLL}$ search. The NR values in the disturbed bands are higher than those in quiet bands because each cluster comprises more candidates above Γ_S in noisy bands than in the quiet bands. This is expected merely due to higher density of disturbances. The new clustering procedure has a higher NR than the old method because it adapts the cluster size to the local over-density and can get as big (or small) as it needs, in order to accommodate the features in the data.

A rigorous quantitative assessment of the detection efficiency in disturbed bands is hard to make because the results would depend not only on the location of the fake signals in parameter space but also their numbers with respect to the disturbances. In such scenarios, there is no unbiased way to pick the fake signal population. However, based on the fact that for a cluster to be identified, we only require a seed above threshold and concurrent clustering around that seed in both F-space and in S-space, we do not expect the presence of more can-

didates due to disturbances (which generally do not cluster in the parameter space) to interfere too much with the identification of the signal clusters. On the contrary, the old procedure does not require a local over-density around the seed and it might happen that a signal candidate gets associated with a higher random fluctuation; this cluster may not satisfy the over-density criteria in the AdCl procedure which may have led to a wrong estimation of the follow-up region. Thus, by requiring the seed to be centered at a local over-density, the new procedure avoids this type of occurrence. This might slightly favour the detection efficiency of the AdCl procedure with respect to the old one.

V. CONCLUSIONS

The clustering procedure that we propose in this paper is more effective at reducing the number of candidates to be considered in follow-up stages while achieving comparable, if not better, detection efficiency with respect to the procedure used in previous searches. Since we operate at fixed computing budget, the number of candidates that a given follow-up stage can search, is fixed. Hence, a higher noise rejection means a lower detection threshold. In a search like the Einstein@Home O1 low-frequency search [11], the new clustering has allowed us to lower the $\beta_{S/GLL}$ threshold. In disturbed bands, the noise rejection is even higher.

There are two main reasons for the observed improvements. The first reason is that the AdCl procedure is more demanding than the old one, i.e. a cluster has to display a more pronounced over-density of candidates compared to nearby noise. The second reason is that, since the cluster size is estimated on the data itself, the clustering algorithm adapts itself to it and is capable of bundle together a large number of candidates arising from extended regions of parameter space.

Another advantage of the AdCl procedure compared to the old one is that, by relying on local over-densities of candidates, the false alarm rate does not increase with a decreasing value of the threshold Γ_L , which is significantly necessary for low-significance searches.

However, the AdCl procedure may well go through many iterations before discarding a single candidate as a single occupancy cluster and restoring the rest of the candidates for future consideration. This, especially in noisy bands, can make

it rather slow. The variable sky-binning depending on the seed amplitude is a way to ease this issue, and quite certainly, further use will inspire other ways to make the procedure faster in all noise conditions.

Currently, the tuning parameters (N_F , C_F , N_S , C_S) and the hill parameters (P_{th} , D_{th} , G_{th}) are chosen to represent the approximate topology of the clusters that we expect from signals. These approximate values are chosen upon visual inspection of the fake signals injected in LIGO data at many values of the signal amplitude h_0 . To improve the estimates on these parameters, one would require to perform a much larger number of Monte-Carlo simulations (in λ and h_0), and then estimate the cluster properties. This remains a difficult task due to limited computational resources and very large parameter space of the tuning and hill parameters. The modeling of clusters arising from non-Gaussian noise (such as instrumental artifacts) is even more difficult, especially for unknown sources of disturbances. In principle, this modeling could however help better discern between signals and noise.

There are certainly other possible ways to perform adaptive clustering. One of the methods is to employ machine learning. Besides that, one could also perform more complex parameter space correlation studies of the detection statistic values, similar to the studies done for cosmic microwave background (CMB) surveys [14] and large scale structure (LSS) surveys [10]. However, such complex analysis methods require much cleaner data, and they are certainly an overkill for the current data-sets.

ACKNOWLEDGMENTS

This procedure was used in [11], and we thank Sergey Klimenko and Evan Goetz for the review of the application of this new clustering procedure to the results of that search. 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. This document has LIGO DCC number P1700123.

-
- [1] J Aasi *et al.* *Directed search for continuous gravitational waves from the Galactic center.* **Phys. Rev. Lett.**, 88(10):102002, 2013.
 - [2] J Aasi *et al.* *Implementation of an F-statistic all-sky search for continuous gravitational waves in Virgo VSR1 data.* **Class. Quant. Grav.**, 31(16):165014, 2014.
 - [3] J Aasi *et al.* (LIGO Scientific Collaboration). *Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data.* **Phys. Rev. D**, 87(8):042001, 2013.
 - [4] B Abbot *et al.* (LIGO Scientific 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(10):102002, 2016.
 - [5] B Abbott *et al.* (LIGO Scientific Collaboration). *Einstein@Home search for periodic gravitational waves in LIGO S4 data.* **Phys. Rev. D**, 79(2):022001, 2009.
 - [6] B P Abbott *et al.* (LIGO Scientific Collaboration). *Einstein@Home search for periodic gravitational waves in early S5 LIGO data.* **Phys. Rev. D**, 80(4):042003, 2009.
 - [7] B Behnke, M A Papa, and R Prix. *Postprocessing methods used in the search for continuous gravitational-wave signals from*

the Galactic Center. **Phys. Rev. D**, 91(6):064007, 2015.

- [8] A Helman. *The Finest Peaks: Prominence and other Mountain Measures*. (E-Book), 2012. ISBN 978-1-41223-664-5.
- [9] D Keitel. *Robust semicoherent searches for continuous gravitational waves with noise and signal models including hours to days long transients*. **Phys. Rev. D**, 93(8):084024, 2016.
- [10] F Leclercq. *Bayesian large-scale structure inference and cosmic web analysis*. **arXiv**, 1605.08420(-):237, 2015.
- [11] LIGO Scientific Collaboration and Virgo Collaboration. *First low-frequency all-sky search for continuous gravitational waves in advanced LIGO data*. Submitted to **Phys. Rev. D** [v1](#), 2017.
- [12] 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(12):122006, 2016.

- [13] A Singh *et al.* *Results of an all-sky high-frequency Einstein@Home search for continuous gravitational waves in LIGO's fifth science run*. **Phys. Rev. D**, 94(6):064061, 2016.
- [14] M Tristram and K Ganga. *Data analysis methods for the cosmic microwave background*. **Rept. Prog. Phys.**, 70(6):899, 2007.
- [15] S Zhu *et al.* *An Einstein@Home search for continuous gravitational waves from Cassiopeia A*. **Phys. Rev. D**, 94(8):082008, 2016.

VI. APPENDIX

A. The first cluster

We now illustrate the different phases of the first iteration of the clustering procedure on two small snippets of data from the LIGO O1 run with and without a fake signal (Fig.8, Fig.9).

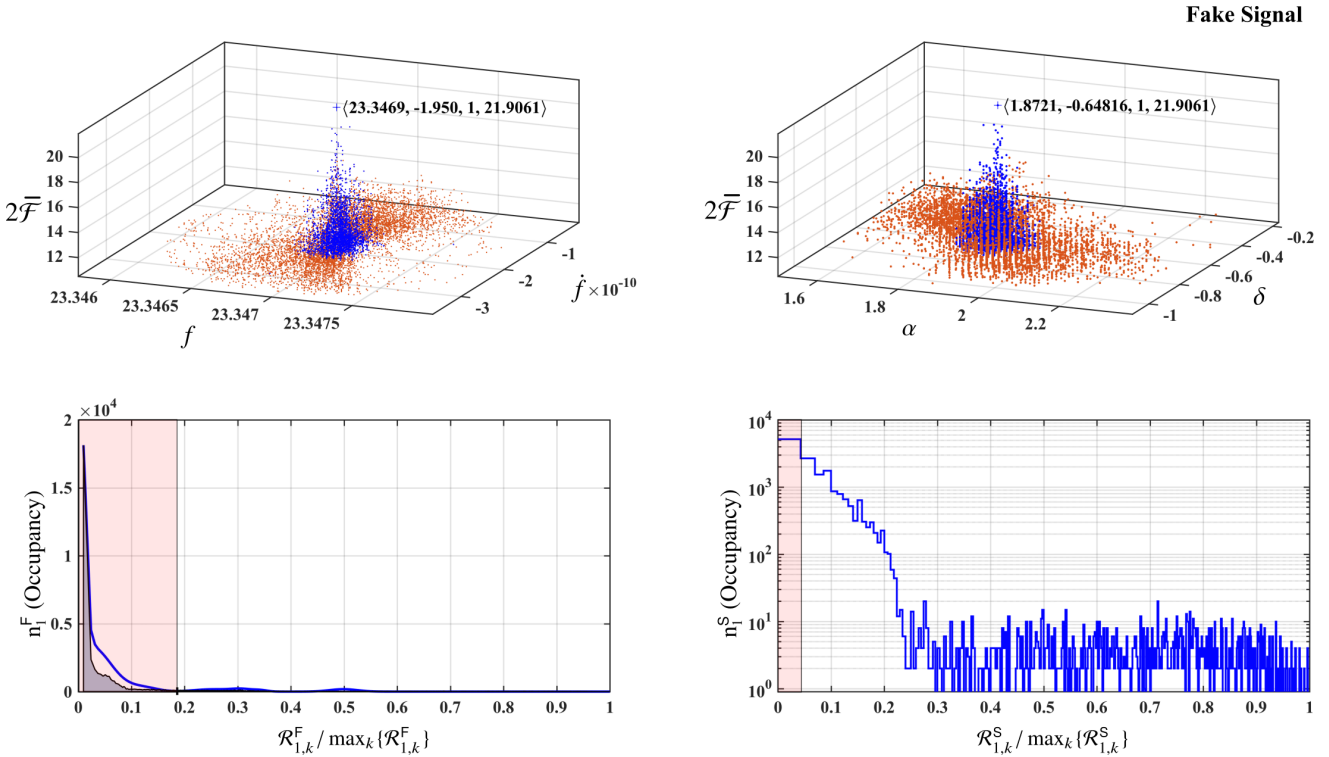


FIG. 8. (*Signal case*) The orange points are the candidates in χ_1^{sky} ; the blue points are the subset of these that form the final cluster ϕ_1 . The corresponding distributions for $\mathcal{R}_{i,k}^F$ and $\mathcal{R}_{i,k}^S$ are shown in the second row of plots. The shaded regions extend up to $\mathcal{R}_i^{F^*}$ (left plot) and $\mathcal{R}_i^{S^*}$ (right plot). The seed is marked with a '+'. The numbers in the brackets by the '+' denote: $\langle f, \dot{f}, i, \Gamma_i \rangle$ in F-space and $\langle \alpha, \delta, i, \Gamma_i \rangle$ in S-space, where $\Gamma \equiv 2\bar{\mathcal{F}}$.

Noise data

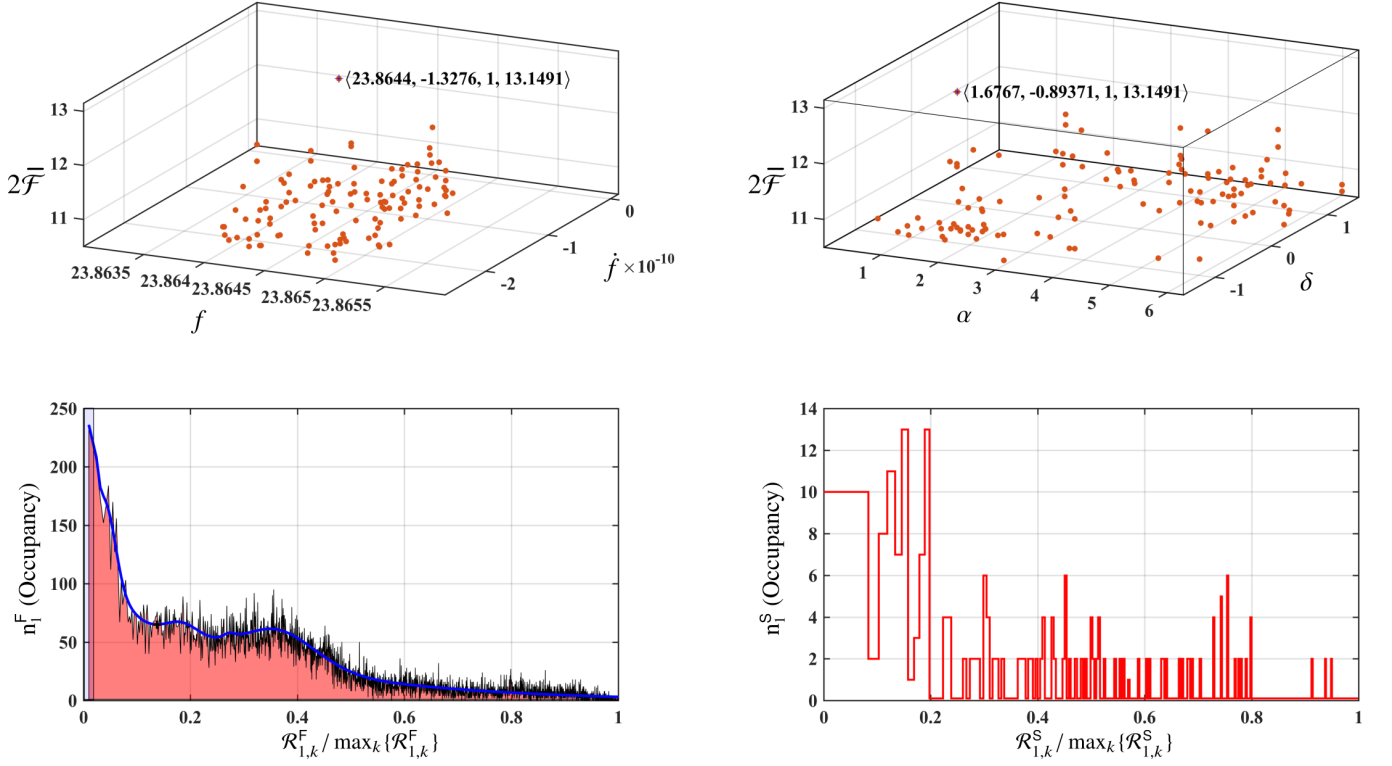


FIG. 9. (*Noise only*) The orange points are the candidates in χ_1^{sky} . The corresponding distributions for $\mathcal{R}_{i,k}^F$ and $\mathcal{R}_{i,k}^S$ are shown in the second row of plots. The shaded regions extend up to \mathcal{R}_i^{F*} (left plot) and \mathcal{R}_i^{S*} (right plot). Note that \mathcal{R}_i^{F*} fails the hill parameters test and is reset to $\mathcal{B}_{i,1}^F$. The distribution in S-space satisfies $n_{i,1}^S \neq \max_r \{n_{i,r}^S\}$ (i.e. no over-density is sky near the seed), so ϕ_1 is a single-occupant-cluster. The seed is marked with a '+'. The numbers in the brackets by the '+' denote: $\langle f, \hat{f}, i, \Gamma_i \rangle$ in F-space and $\langle \alpha, \delta, i, \Gamma_i \rangle$ in S-space, where $\Gamma \equiv 2\bar{\mathcal{F}}$.