Detecting long-duration narrow-band gravitational wave transients associated with soft gamma repeater quasiperiodic oscillations

David Murphy, Maggie Tse, Peter Raffai, Imre Bartos, Rubab Khan, Zsuzsa Márka, Luca Matone, Keith Redwine, and Szabolcs Márka

Detecting Long-Duration Narrow-Band Gravitational Wave Transients Associated with Soft Gamma Repeater Quasi-Periodic Oscillations

David Murphy,[1, 2, a] Maggie Tse,[1, 2] Peter Raffai,[1, 3, 4] Imre Bartos,[1, 2] Rubab Khan,[5]
Zsuzsa Márka,[1, 2] Luca Matone,[1, 2] Keith Redwine,[6] and Szabolcs Márka[1, 2]

[1]*Department of Physics, Columbia University, New York, NY 10027, USA*
[2]*Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA*
[3]*Institute of Physics, Eötvös University, 1117 Budapest, Hungary*
[4]*MTA-ELTE EIRSA "Lendület" Astrophysics Research Group*
[5]*Department of Astronomy, Ohio State University, Columbus, OH 43210, USA*
[6]*Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA*

We have performed an in-depth concept study of a gravitational wave data analysis method which targets repeated long quasi-monochromatic transients (triggers) from cosmic sources. The algorithm concept can be applied to multi-trigger data sets in which the detector-source orientation and the statistical properties of the data stream change with time, and does not require the assumption that the data is Gaussian. Reconstructing or limiting the energetics of potential gravitational wave emissions associated with quasi-periodic oscillations (QPOs) observed in the X-ray lightcurve tails of soft gamma repeater flares might be an interesting endeavour of the future. Therefore we chose this in a simplified form to illustrate the flow, capabilities, and performance of the method. We investigate performance aspects of a multi-trigger based data analysis approach by using $\mathcal{O}\,(100\,\text{s})$ long stretches of mock data in coincidence with the times of observed QPOs, and by using the known sky location of the source. We analytically derive the probability density function of the background distribution and compare to the results obtained by applying the concept to simulated Gaussian noise, as well as off-source playground data collected by the 4-km Hanford detector (H1) during LIGO's fifth science run (S5). We show that the transient glitch rejection and adaptive differential energy comparison methods we apply succeed in rejecting outliers in the S5 background data. Finally, we discuss how to extend the method to a network containing multiple detectors, and as an example, tune the method to maximize sensitivity to SGR 1806−20 flare times.

PACS numbers: 95.85Sz, 95.30.Sf, 95.55.Ym

I. INTRODUCTION

Soft gamma repeaters (SGRs) are mostly galactic objects that emit short bursts of X-ray and gamma radiation at irregular intervals. At times these sources emit exceptionally energetic flares that last hundreds of seconds, with peak luminosities of up to $10^{46}\,\text{erg s}^{-1}$ [1]. The most luminous SGR flare observed to date was recorded on 27 December 2004, and originated from SGR 1806−20 [2]. Quasi-periodic oscillations (QPOs) were observed at late times in the lightcurve tail by the Rossi X-Ray Timing Explorer (RXTE) [3] and the Ramaty High Energy Solar Spectroscopic Imager (RHESSI) [4] satellites, most notably at 92.5 Hz and 626.5 Hz [1, 2]. Although the exact timing of the QPO signals is unclear since the peak flare saturated both detectors, the signals can be unambiguously observed beginning approximately $170 - 200\,\text{s}$ after the initial flare [1, 2].

The $\sim 8\,\text{s}$ period of the lightcurve tail observed during the December 2004 hyperflare [2], as well as the phase-dependence of the QPO features, suggests that the parent object might be a type of slowly rotating neutron star. In this model the phase dependence can be interpreted as localization to a specific region of the star's surface, and the observed frequencies may be related to particular seismic modes [1, 2], making QPOs an interesting probe of the astroseismology of SGRs. In particular, if QPOs are indeed related to seismic vibrations of a neutron star, then we can estimate the thickness of the star's crust, and possibly infer what type of material makes up the star's interior [5]. It has also been suggested that changes in the observed QPOs over time may be indicative of physical changes in the source, such as a decaying magnetic field [6].

The magnetar model [7] proposes that SGRs are a type of highly magnetized NS with a surface magnetic field in excess of $10^{15}\,\text{G}$. In this model QPOs can be explained by a starquake event, in which the star's magnetic field becomes twisted and then violently rearranges, exciting crustal or global seismic modes and releasing gamma ray bursts with the observed QPO structures [8, 9]. Numerical studies in which the magnetar is treated as a coupled system consisting of an elastic crust and a magnetohydrodynamic fluid core have demonstrated a number of interesting effects. These include the dynamical generation of instabilities and asymmetric quasi-equilibria in the magnetic field [10], as well as a coupling between discrete crustal modes and continuous bands of global Alfvén modes, with QPOs appearing in the late time starquake oscillations near the edges of the Alfvén spectral

[a]Electronic address: djm2131@columbia.edu

bands [11–15]. This numerical work also suggests that some QPOs may be exceptionally strong because they are close in frequency to one of the crustal modes [12, 13, 15]. The same dynamics may also drive the emission of gravitational waves (GWs) by exciting the magnetar's seismic modes; in particular, crustal f-modes [16]. Although it is argued that at realistic values of the magnetic field strength f-modes are unlikely to be sufficiently excited to produce an observable signal in second-generation gravitational wave detectors [17, 18], Zink et al. note that low-frequency modes up to 100 Hz remain unstudied, and may be sufficiently excited to produce detectable GW signals [18].

A number of searches for gravitational radiation associated with SGRs and magnetars have been conducted. In particular, a search for GWs associated with f-mode ringdowns in six magnetars targeting the frequency ranges $1 - 3\,\text{kHz}$, $100 - 200\,\text{Hz}$, and $100 - 1000\,\text{Hz}$, with time windows of $\sim 200\,\text{ms}$, found no evidence of a GW signal [19]. Other searches targeting short-duration ($\sim 0.3\,\text{s}$) GWs from SGRs [20] as well as stacked GW emissions associated with the SGR 1900+14 storm event [21] also found no evidence of GWs associated with SGR bursts at energies accessible by past interferometric gravitational wave detectors.

Our work extends a search algorithm developed by Matone and Márka [22], targeting narrow-band, long-duration GW signals associated with SGR QPOs. This algorithm is based on an excess energy method and thus does not presume any particular waveform for the emitted GW signal. In addition, the algorithm also does not require exact knowledge of the timing of the GW signal relative to the electromagnetic trigger; we choose conservative $\mathcal{O}(100\,\text{s})$ long time windows encompassing the entire flare event. The algorithm was previously applied to LIGO data recorded during the December 2004 SGR $1806-20$ hyperflare event to constrain the energetics of any associated GW signal, and set a lowest bound of $7.67 \times 10^{46}\,\text{erg}$ on the total GW emissions associated with the QPO at 92.5 Hz [23]. A similar algorithm targeting long-duration GW transients using a complementary cross-power technique is discussed in [24], and is also relevant to searches for GWs from SGRs in the LIGO data.

The LIGO network of interferometric gravitational wave detectors has completed 6 science runs to date. Science runs 5 and 6 lasted from November 4th, 2005 to September 30th, 2007, and from July 7th, 2009 to October 20th, 2010, respectively. During S5 alone ~ 150 EM flares originating from SGR $1806-20$ were observed. If SGR flares are indeed coupled to seismic vibrations of a neutron star as suggested by the magnetar model, one might expect that each burst should excite the same, characteristic frequencies of the star's crust regardless of the luminosity of the flare, leading to repeated GW emissions at the same frequencies. As a result, a comprehensive algorithm analyzing all triggers and QPO frequencies overlapping with the LIGO data set for evidence of repeated GW emissions from SGR flares might prove itself

more sensitive than an algorithm targeting a single burst or QPO frequency. Although we illustrate the concept through a 10 Hz frequency band centered on the 92.5 Hz QPO observed during the December 2004 SGR $1806-20$ hyperflare, the algorithm can be applied to any frequency window of interest, allowing the method to target each of the observed QPO frequencies.

The algorithm concept presented has been designed to efficiently process large, multi-detector data sets containing multiple SGR flares from a single source, and includes a number of data quality measures intended to account for nonstationarity and deviations from Gaussianity in the data stream of a physical detector operating for months or years. The method, like the methods mentioned above, benefits from knowledge of the arrival time of an electromagnetic counterpart to the hypothesized GW signal, as well as knowledge of the sky position of the source. In addition, we further assume that the recurrent SGR bursts contain the same QPO structures and seismic modes observed in hyperflares, and introduce a new detection criterion based on the two-sample Kolmogorov-Smirnov (K-S) test and Fisher's method, allowing us to increase the sensitivity of the algorithm relative to analyzing a single event. We also discuss the problem of variable detector orientation for each trigger, and discuss how to optimize the concept accordingly. We show that with the proper data conditioning procedures the behaviour of real data can be modeled well mathematically and approximated numerically through gaussian statistics.

The paper is organized as follows: in Sec. II we discuss the algorithm and analytically derive the probability density function (PDF) of the background distribution for Gaussian white noise. In Sec. III we discuss the implementation of the algorithm as a concept pipeline and data quality measures. Sec. IV verifies the implementation by applying the pipeline to randomly generated Gaussian noise and comparing the results to the predictions of the mathematical model. Finally, in Sec. V we present a concept study of the off-source detector background and orientation using playground data segments from LIGO's fifth science run [25].

II. MATHEMATICAL MODEL OF THE ALGORITHM

The time domain algorithm splits the data stream by bandpass filtering into three channels of a specific bandwidth motivated by X-ray observation constraints Δf: the QPO channel, which is centered on the QPO frequency, and the *up* and *down* channels, centered at $f_{\text{up}} = f_{\text{QPO}} + \Delta f$ and $f_{\text{down}} = f_{\text{QPO}} - \Delta f$, respectively. The bandwidth Δf for example can be derived from the measured full width at half maximum (FWHM) of the QPOs observed in the EM (X-ray) emissions [35]. We define the energy of a strain time series $\xi(t)$ in an inter-

val Δt to be

$$E = \int_{\Delta t} dt\, |\xi(t)|^2. \qquad (2.1)$$

We can compare the energy contained in the QPO band to the energy contained in the adjacent bands during an interval Δt by computing the excess energy

$$E^{\text{excess}} = E_{\text{QPO}} - w\left(E_{\text{up}} + E_{\text{down}}\right), \qquad (2.2)$$

where w is a locally determined averaging factor [36]. Defining the excess energy in this manner allows us to isolate the contribution of a GW signal associated with the QPO to the total energy of the data stream while simultaneously suppressing broadband instrumental noise. It is also necessary to account for temporal non-stationarity in the data stream since we perform a comprehensive analysis with data spanning several years, during which changes in the properties of the detector noise are too significant to be ignored. We accomplish this by computing the excess energy in the time interval $(t_{\text{center}}, t_{\text{center}} + \Delta t)$, the excess energy in the time interval $(t_{\text{center}} - \Delta t, t_{\text{center}})$, and then computing the difference in excess energy

$$\Delta E^{\text{excess}} = \qquad\qquad\qquad (2.3)$$
$$= \; E^{\text{excess}}_{(t_{\text{center}},\,t_{\text{center}}+\Delta t)} - E^{\text{excess}}_{(t_{\text{center}}-\Delta t,\,t_{\text{center}})}$$

resulting in a quantity which can be compared directly across different times. For on-source measurements we choose t_{center} to be the time of an observed trigger [37], and for off-source measurements t_{center} is randomly sampled from stretches of data disjoint from the region hypothesized to contain a GW signal. The value of Δt is chosen based on the timescale of the SGR flares, and is the same for both on-source and off-source measurements: since we do not, in general, have detailed information regarding the length of every SGR burst, we choose Δt to be the full duration of the longest observed lightcurve tail from a given SGR as an illustration here. For SGR 1806−20 this is $\Delta t = 300\,\text{s}$ based on the December 2004 hyperflare. This choice of Δt is intended to be conservative to ensure that we capture associated GW signals. The method for calculating ΔE^{excess} is depicted schematically in Fig. 2.

By considering each sampling point in the data stream as an independent, identically distributed normal random variable we can construct a theoretical model for the algorithm. In the following section we derive an explicit formula for the probability density function of ΔE^{excess}. These results are then used in Sec. IV to check the algorithm's performance by comparing numerical simulations to the theoretical results.

A. Sampling Distribution of ΔE^{excess}

Here we model the strain noise of the detector by band-limited Gaussian white noise for which each sampling point, ξ_i, is an independent [38] Gaussian random variable with zero mean and standard deviation σ. The energy of the strain noise in a time interval Δt is given by the discrete analogue of (2.1)

$$E = \frac{1}{F_{\text{s}}} \sum_{i=1}^{n} \xi_i{}^2, \qquad (2.4)$$

where F_{s} is the sampling frequency and $n = F_{\text{s}}\Delta t$ is the total number of sampling points. E follows a χ^2-distribution with PDF [26]

$$f_E(x) = \begin{cases} \frac{F_{\text{s}}{}^{n/2}}{\sigma^n 2^{n/2}\Gamma(n/2)} x^{n/2-1} e^{-\frac{F_{\text{s}}}{2\sigma^2}x} & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases} \qquad (2.5)$$

Assuming that the channels are uncorrelated and the signal energy is distributed isotropically in frequency space it suffices to choose $w = 1/2$ in (2.2), and we can then define a random variable ΔE^{excess} as

$$\Delta E^{\text{excess}} = E^{\text{excess}}_{\text{after flare}} - E^{\text{excess}}_{\text{before flare}} = \qquad (2.6)$$
$$= \left[E^{\text{after}}_{\text{QPO}} - \frac{1}{2}\left(E^{\text{after}}_{\text{up}} - E^{\text{after}}_{\text{down}}\right)\right] -$$
$$- \left[E^{\text{before}}_{\text{QPO}} - \frac{1}{2}\left(E^{\text{before}}_{\text{up}} - E^{\text{before}}_{\text{down}}\right)\right]$$

where each of the terms has PDF (2.5). The PDF of ΔE^{excess} can be found using an identity for characteristic functions (here denoted by φ) [27]

$$\varphi_{\sum_i a_i X_i}(s) = \prod_i \varphi_{X_i}(a_i s). \qquad (2.7)$$

The characteristic function associated with (2.5) is

$$\varphi_E(s) = \frac{1}{\left(1 - \frac{2\sigma^2}{F_s}is\right)^{n/2}}, \qquad (2.8)$$

so, using (2.7) and the inverse Fourier transform we find that

$$f_{\Delta E^{\text{excess}}}(x) = \qquad\qquad\qquad (2.9)$$
$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} ds\, \frac{\cos(sx)}{\left(1 + \frac{4\sigma^4}{F_s^2}s^2\right)^{n/2}\left(1 + \frac{\sigma^4}{F_s^2}s^2\right)^{n}}.$$

The cumulative distribution function (CDF) is

$$F_{\Delta E^{\text{excess}}}(x) = \qquad\qquad\qquad (2.10)$$
$$= \frac{1}{2\pi} \int_{-\infty}^{x} du \int_{-\infty}^{\infty} ds\, \frac{\cos(su)}{\left(1 + \frac{4\sigma^4}{F_s^2}s^2\right)^{n/2}\left(1 + \frac{\sigma^4}{F_s^2}s^2\right)^{n}}.$$

Note that the CDF can similarly be calculated for any other values of w.

III. DESCRIPTION OF THE PIPELINE

The pipeline is designed to sequentially process the LIGO data stream in the time domain and take ΔE^{excess} measurements around times associated with SGR flares. We distinguish between flare measurements and background measurements. Flare measurements are performed at the time of an observed trigger, whereas background measurements are performed at times before or after a trigger and used to characterize the detector background. Fig. 1 depicts a schematic of the algorithm.
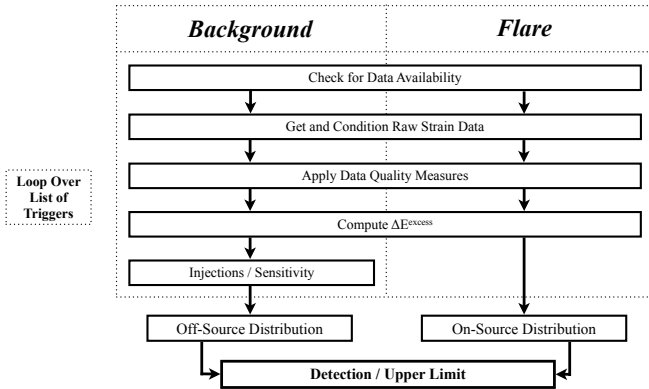

FIG. 1: A block diagram illustrating the steps in the analysis pipeline concept. The data conditioning step is discussed in Sec. III A. Data quality measures are intended to account for variability in the quality of data recorded by a physical instrument and are discussed in Sec. III B. The injection step is optional, and is used to place an upper limit in the event that the pipeline fails to make a detection.

Using a list of the start and end times of detector lock stretches that contain one or more astrophysical triggers, the pipeline sequentially processes each trigger, performing a single measurement at the time of the trigger, and a fixed number of background measurements before and after the trigger. For each measurement the data is filtered to isolate the QPO band and the two adjacent frequency bands. The algorithm then applies data quality measures to the filtered data streams as described in Sec. III B, and computes ΔE^{excess} as defined by (2.3). After iterating over the full list of triggers the resulting background and flare distributions are compared using the two-sample K-S test [28], and detection is claimed if the distributions are distinguishable with 5σ confidence. Otherwise, we inject a simulated waveform of known energy into a single pre-flare background measurement corresponding to each trigger to determine the sensitivity — defined to be the minimum injected energy such that the background and foreground distributions meet our detection criteria — and set an upper limit on the energy of any GW signal contained in the data stream.

A. Data Conditioning

The ingoing data stream is conditioned by applying a series of three zero-phase, bandpass, Butterworth filters, resulting in the upper, lower, and QPO channels. Sufficiently large buffer time intervals are included and cut off after filtering at both ends of the data stream to make sure that no transients remain in the data due to filtering. The difference in excess energy is then calculated according to (2.3). Injections are performed by generating the plus and cross components ($h_+(t)$ and $h_\times(t)$, respectively) of the simulated GW signal with a given root-sum-square strain (h_{rss})

$$h_{\text{rss}} \equiv \sqrt{\int dt \left(h_+^2(t) + h_\times^2(t) \right)} \qquad (3.1)$$

and adding it to the input data prior to the data conditioning step. We can write

$$h(t) = F_+ \cdot h_+(t) + F_\times \cdot h_\times(t) \qquad (3.2)$$

where and F_+ and $F_\times$ are antenna factors determined by the orientation of the detector relative to the source at the time of a given trigger and the choice of the GW polarization frame. The antenna factors satisfy $0 \leq F_+, F_\times \leq 1$ and encode the directional sensitivity of the detector to an incoming plus- (F_+) or cross- ($F_\times$) polarized GW [29]. The injection process is performed once per trigger, resulting in a simulated foreground distribution that can be used to determine the sensitivity of the algorithm.

The root-sum-square strain measured at the detector (h_{rss}) can be translated into a total physical energy ($E_{\text{GW}}^{\text{iso}}$) that the source would emit in GWs if the emission was isotropic [23]:

$$E_{\text{GW}}^{\text{iso}} = \frac{\pi^2 c^3}{G} d^2 f_0^2 h_{\text{rss}}^2, \qquad (3.3)$$

where d is the distance between the detector and the source, and f_0 is the central frequency of the GW signal [39]. In the event of a non-detection this relation allows us to use the measured sensitivity to constrain the energetics of GW emissions at the source.

B. Data Quality

A number of measures have been implemented to account for instrumental transients and the variability in data quality across long stretches of time. We rely primarily on three techniques: (i) the application of data quality flags, (ii) transient glitch rejection, and (iii) the use of a locally-determined weighting factor w to measure the excess energy (defined by (2.2)).

1. Data quality flags are constructed by the LSC's detector characterization working group to characterize noise contamination in the detector's data

stream, and mark the times of known environmental noise sources and detector calibration issues. The flags are divided into categories 1, 2, 3, and 4 based on severity and duration. Here we reject any data marked with a flag designated category 1 — times when the detector was not taking data in the design configuration — or category 2 — times when well-understood instrumental glitches were observed. We also avoid standard and significant line features while selecting the frequency bands, and we propose to apply lock-in filtering techniques to mitigate the effects of any potentially unavoidable line features. Detailed descriptions of the flags can be found in [30].

2. Transient glitch rejection provides a second method for reducing noise contamination. Since we expect a long-duration signal lasting hundreds of seconds we can safely reject transient glitches over much smaller time scales. This is accomplished using a method similar to the glitch veto algorithm discussed in [23]: we divide the data stream of each, filtered channel into tiles of length $\Delta\tau$, typically of order $1\,\mathrm{s}$ or less, compute distributions of the energy per tile for each of the three channels, and recursively veto any tiles with energy more than ± 4 standard deviations from the mean tile energy. We then divide each channel by a duty factor

$$DF = 1 - \frac{\mathrm{length\,(veto)}}{\mathrm{length\,(data)}} \tag{3.4}$$

to renormalize the remaining data in this example [40]. The veto is applied to the strain time series of each filtered channel independently since glitches are not generally distributed evenly in frequency space. This process is iterated until either all outliers have been removed or the fraction of data vetoed passes a cutoff value, typically chosen to be 10%, at which point the data segment is rejected as unreliable.

3. Another source of data contamination arises from the curvature of the LIGO noise floor. Naively choosing $w = 1/2$ in the definition of the excess energy statistic (2.2) results in a quantity that does not properly account for differences in the background noise across multiple frequency bands. Since the properties of the detector noise change with time we are forced to calculate an appropriate value of w for each measurement [41]. With this in mind we define a local weighting factor w by averaging the four quantities

$$w_{\pm 1,\pm 2} = \left| \frac{E_{\mathrm{QPO}}}{E_{\mathrm{up}} + E_{\mathrm{down}}} \right|_{t_w = \pm t_{w,1},\ \pm t_{w,2}} \tag{3.5}$$

calculated by summing the energies of each channel in the time interval $[t_w - \Delta t, t_w + \Delta t]$ for each of the four times $t_w = \pm t_{w,1}, \pm t_{w,2}$. This is depicted
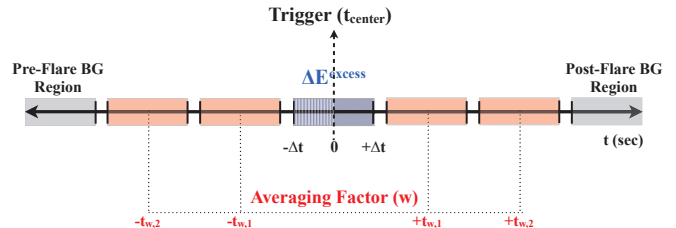
schematically for a hypothetical measurement in Fig. 2.

FIG. 2: A schematic illustrating the regions of the time-series used to perform a $\Delta E^{\mathrm{excess}}$ measurement with a locally determined weighting factor w. The measurement is centered about a time t_{center} which is either the time of an observed SGR trigger or randomly sampled from the background data. The value of w is obtained by calculating the ratio (3.5) using the data contained in the interval $[t_w - \Delta t, t_w + \Delta t]$ for each of the four times $t_w = \pm t_{w,1}, \pm t_{w,2}$. These four values are averaged to estimate w, and then $\Delta E^{\mathrm{excess}}$ is computed using this value in (2.2) and (2.3). The part of the diagram corresponding to $\Delta E^{\mathrm{excess}}$ is divided into a solid region and a dashed region to emphasize that it is the difference in excess energy between these two regions which we ultimately calculate using the locally determined w.

We choose Δt to be the same time interval used for computing excess energy measurements, and $t_{w,1}$ and $t_{w,2}$ to be the minimum adjacent times which ensure that none of the measurements overlap. With this scheme we can use the data immediately preceding and following the region used for a given $\Delta E^{\mathrm{excess}}$ measurement to determine the appropriate value of w in (2.2). Defined in this way, w accounts for differences between channels due to the shape of the noise floor, as well as changes in the noise floor over time.

IV. GAUSSIAN NOISE SIMULATIONS

In order to validate the concept and test the mathematical model of section II we simulate the result of applying the pipeline to a set of detected SGR $1806-20$ flares coinciding with the data from LIGO's fifth science run (S5). This is accomplished by configuring the pipeline to perform a realistic study but replacing the LIGO data stream with stretches of randomly generated Gaussian white noise. This is not unreasonable as the data quality cut studies indicated that the "cleaned" data is very similar to a Gaussian random set (see [25] and references therein).

A. Simulated Detector Background

We construct a background sampling distribution by calculating $\Delta E^{\mathrm{excess}}$ for simulated data consisting of 240

seconds of randomly generated Gaussian white noise with strain spectral noise density $\tilde{n} = 10^{-22}\,\mathrm{Hz}^{-1/2}$. The value of $\tilde{n}$ is chosen so that the simulated noise is similar in magnitude to the LIGO noise floor during S5. We choose $f_{\mathrm{QPO}} = 92.5\,\mathrm{Hz}$ and $\Delta f = 10\,\mathrm{Hz}$, corresponding to the frequency and FWHM of the strongest QPO signal observed in the lightcurve tail of the 27 December, 2004 SGR $1806-20$ hyperflare event, and $\Delta t = 30\,\mathrm{s}$ [23]. This value of Δt is shorter than the $\sim 300\,\mathrm{s}$ duration of the lightcurve tail observed during the SGR $1806-20$ hyperflare, nevertheless for the purposes of a concept study, is less computationally expensive while providing essentially the same insight. We also compute the theoretical PDF of the background distribution using the same parameters to compare the theoretical result (2.9) to the simulation.

In anticipation of later work utilizing LIGO data, we also construct a second background distribution incorporating a simulated veto step to demonstrate that our algorithm is robust with respect to the inclusion of data quality vetoes. This is accomplished by randomly removing 5% of the sampling points from the strain time series of each, filtered channel independently, and then applying a duty factor (3.4) to each channel. We also obtain a fitted PDF by using nonlinear least squares regression to fit the theoretical CDF (2.10) to the empirical CDF of the resulting background distribution. Fig. 3 depicts the results of these simulations.

In order to characterize the goodness of fit we compute the mean percentage error (MPE)

$$\mathrm{MPE} = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{\hat{F}(x_i) - F(x_i)}{F(x_i)}\right| \qquad (4.1)$$

where $\{x_i\}$ is the set of observations which make up the background sampling distribution, n is the number of elements of $\{x_i\}$, $\hat{F}(x)$ is the empirical CDF associated with $\{x_i\}$, and $F(x)$ is the theoretical (no veto) or fit (veto) CDF obtained from (2.10). For the no veto simulation we find

$$\mathrm{MPE}_{\mathrm{no\ veto}} = 0.0125 \qquad (4.2)$$

indicating that the theoretical PDF/CDF obtained by numerically integrating (2.9)/(2.10) with the same parameters used to generate the simulated Gaussian noise accurately describes the simulated detector background. The MPE for the veto simulation is

$$\mathrm{MPE}_{\mathrm{veto}} = 0.0196 \qquad (4.3)$$

indicating that the fit of the theoretical model to the simulated detector background with a veto step is also accurate. In addition, we observe that the fit parameters $(n_{\mathrm{veto}}, \tilde{n}_{\mathrm{veto}}) = (493285, 10^{-22}\,\mathrm{Hz}^{-1/2})$ for the veto case are within 0.4% of the theoretical values $(n_{\mathrm{no\ veto}}, \tilde{n}_{\mathrm{no\ veto}}) = (491520, 10^{-22}\,\mathrm{Hz}^{-1/2})$ for the no veto case, demonstrating that the influence of the veto step on the detector background distribution is small.
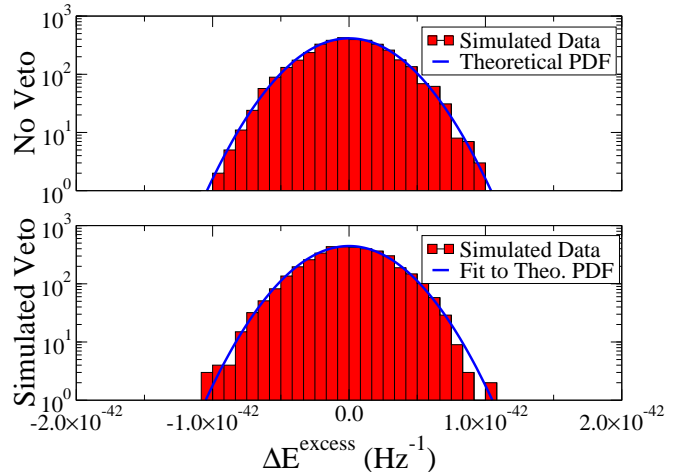

FIG. 3: Detector background with and without a simulated data quality veto step. The parameters used to generate the Gaussian noise are $n = 491520$ and $\tilde{n} = 10^{-22}\,\mathrm{Hz}^{-1/2}$, where $n = F_s\Delta t$ is the total number of sampling points used to make a $\Delta E^{\mathrm{excess}}$ measurement and $\tilde{n}$ is the strain spectral density of the noise. The upper plot shows the agreement between the detector background with no vetoes (histogram) and the theoretical PDF (solid curve) obtained by numerically integrating (2.9). The lower plot shows a fit of the theoretical PDF to the modified detector background after incorporating simulated data quality vetoes. The parameters of the fit, obtained using nonlinear least-squares regression to fit the theoretical CDF (2.10) to the empirical CDF of the resulting detector background, are $n = 493285$ and $\tilde{n} = 10^{-22}\,\mathrm{Hz}^{-1/2}$.

B. Sensitivity

The sensitivity is determined by injecting waveforms with given h_{rss} into the background described in Sec. IV A. Though in general we would expect a GW signal from a magnetar to be elliptically polarized [16], for simplicity's sake we have used sine-Gaussian waveforms

$$h(t) = A\sin\left(2\pi f_c t + \phi\right)e^{-\left(\frac{t-t_0}{\tau}\right)^2} \qquad (4.4)$$

linearly polarized along the plus-direction, and parameterized by the quality factor $Q = \sqrt{2}\pi\tau f_c \gtrsim 1000$, where A is the peak amplitude, f_c is the central frequency, τ is the $1/e$ decay time, ϕ is the phase of the sinusoid relative to the gaussian envelope, and t_0 is the time of the waveform peak. We choose $f_c = 92.5\,\mathrm{Hz}$, corresponding to the QPO at $92.5\,\mathrm{Hz}$, and choose values of t_0 and Q such that all but a negligible fraction of the injected h_{rss} is contained in the interval $[t_{\mathrm{center}}, t_{\mathrm{center}} + \Delta t]$. Plus-polarized injections with a given value of h_{rss} are constructed by first generating a waveform $h(t)$ with functional form (4.4) and $A = 1$, calculating the energy of this waveform, and then rescaling so that the total energy is given by $F_+^2 h_{\mathrm{rss}}^2$. The injection waveform satisfies $0 \le E_{h(t)} \le h_{\mathrm{rss}}^2$ and $E_{h(t)} = h_{\mathrm{rss}}^2$ in the idealized case

$F_+ = 1$.

A foreground sampling distribution is constructed by making one such injection into each of 108 simulated, Gaussian noise segments, where the value 108 is obtained by counting the number of electromagnetic SGR 1806−20 triggers observed during S5 for which enough science-quality H1-L1 coincident data is available to compute ΔE^{excess} with $\Delta t = 30\,\text{s}$ [42]. The simulated background and foreground distributions are then compared using the two-sample K-S test. This procedure is performed 100 times for each choice of h_{rss} to estimate the percentage of simulation trials in which we can claim detection with 5σ significance [43]. We iterate by interpolating the resulting injected h_{rss} versus detection percentage curve to find the sensitivity, defined to be the minimum injected h_{rss} such that we can claim detection 50% of the time. In Fig. 4 we examine the sensitivity for a Gaussian noise simulation of the LIGO S5 data set containing 108 triggers. For simplicity we ignore the detector orientation and set $F_+ = F_\times = 1$ for all triggers. As the average value for F_+ and $F_\times$ is usually around 0.6, this simplification underestimates the sensitivity by an average factor of ~ 1.67.
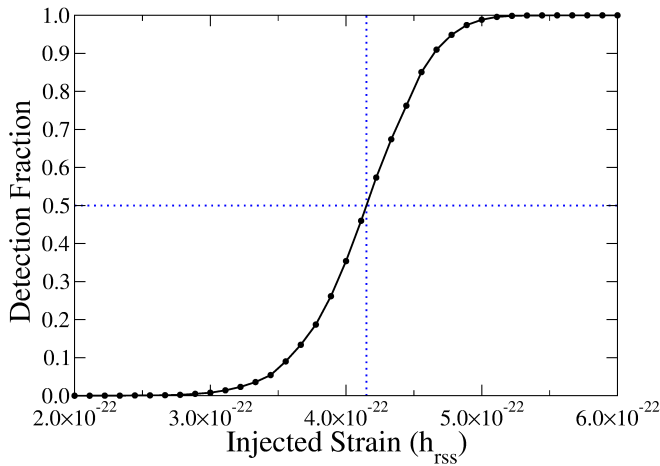

FIG. 4: A simulation of the sensitivity for Gaussian white noise with strain spectral noise density $\tilde{n} = 10^{-22}\,\text{Hz}^{-1/2}$. The sensitivity is defined to be the minimum injected h_{rss} such that the foreground and background distributions are distinguishable via the two-sample K-S test with 5σ significance 50% of the time. We consider a background sampling distribution with 20,000 samples and a trigger distribution with 108 samples, and use linearly-polarized sine-Gaussian injections. Each injection is performed 100 times assuming ideal detector-source orientation to calculate the detection percentage.

We find that the sensitivity with negligible uncertainty for the simulated data set, with trigger times corresponding to those astrophysically detected by satellites, is

$$h_{\text{rss}}^{\text{sens}} = 4.15 \times 10^{-22}\,\text{Hz}^{-1/2}. \tag{4.5}$$

Assuming a detector-source distance of $14\,\text{kpc}$ for SGR 1806−20 [32] this corresponds to an isotropic GW emission at the source (3.3) of

$$E_{\text{GW}}^{\text{iso}} = 1.10 \times 10^{46}\,\text{erg} \tag{4.6}$$

V. APPLICATION TO OFF-SOURCE S5 DATA

In Sec. IV we demonstrated that (2.9) accurately describes the background modeled through simulated, Gaussian data. Here we perform a similar analysis using the data from LIGO's fifth science run (S5) for the 4-km Hanford detector (H1) [25]. We first analyze the effects of incorporating the data quality measures described in Sec. III B to account for deviations from Gaussianity due to noise contamination in the LIGO data stream. Then we calculate the background ΔE^{excess} distribution of H1 during S5. Finally, we discuss how to extend the algorithm to perform a multi-detector study utilizing the full LIGO network, and optimize the method for SGR 1806−20 triggers observed by the H1-L1 network during S5. We again choose $f_{\text{QPO}} = 92.5\,\text{Hz}$, $\Delta t = 30\,\text{s}$, and $\Delta f = 10\,\text{Hz}$, based on observations of the SGR 1806−20 hyperflare event [23].

A. S5 Data Quality

In order to characterize the weighting factor w we calculate the energy in each of the three channels — E_{QPO}, E_{up}, and E_{down} — in the time interval $[t_{\text{center}} - \Delta t, t_{\text{center}} + \Delta t]$ and compute the sampling statistic

$$\hat{w} = \frac{E_{\text{QPO}}}{E_{\text{up}} + E_{\text{down}}}. \tag{5.1}$$

In figure 5 this measurement is performed for a series of times $\{t_{\text{center}}\}$ distributed throughout the S5 run up to April 2007.

We observe that, on average, $\hat{w} < 0.5$, consistent with the shape of the S5 LIGO noise floor [25]. In addition, the large fluctuations in $\hat{w}$ from sample to sample indicate the need to determine the weighting factor in the excess energy (2.2) locally for each measurement. In the remainder of the paper all ΔE^{excess} measurements are performed using the scheme for determining w described in Sec. III B.

We investigate the influence of the data quality measures on the shape of the H1 background ΔE^{excess} distribution by performing three trials, corresponding to three veto methodologies. In the first no data quality measures are applied. In the second we exclude data marked with category 1 or category 2 flags. Finally, in the third we incorporate data quality flags as well as a transient glitch rejection step using $10\,\text{ms}$ tiles and a 4σ veto threshold. The resulting distributions are summarized in Fig. 6.
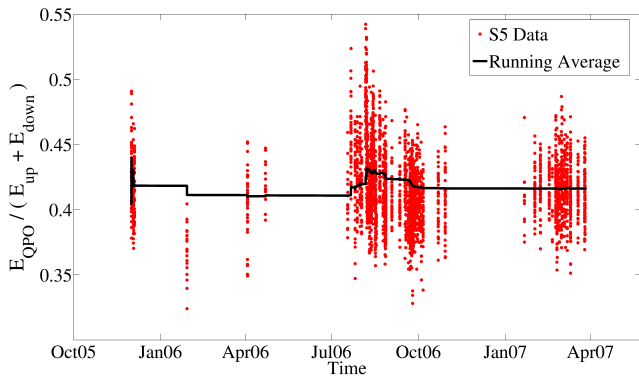
We conclude that although the inclusion of DQ flags

FIG. 5: $\hat{w}$, defined by (5.1), calculated using H1 data for a series of background times distributed throughout the S5 run up to April 2007. The running average is computed at time t by averaging all measurements occurring prior to t.
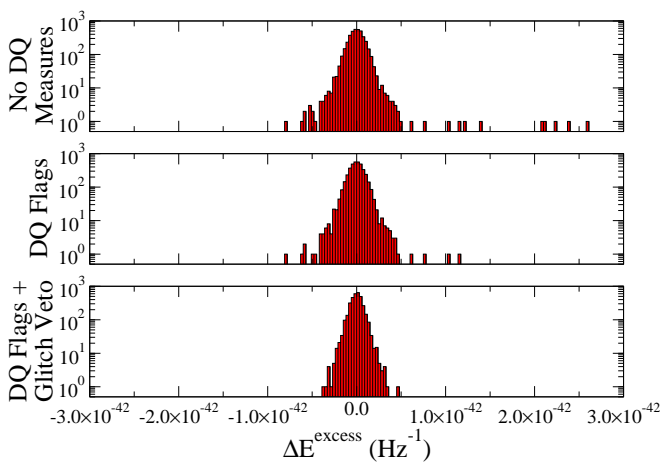

FIG. 6: Veto methodology and resulting H1 detector background using S5 data. In the top figure no data quality measures are applied. In the middle figure data marked with category 1 or category 2 DQ flags is removed. In the bottom figure category 1 and 2 flags are applied as well as a transient glitch rejection step (see Sec. III B) using 10 ms tiles and a 4σ veto threshold.

alone goes a long way towards successfully removing the loudest outliers, the glitch veto step is also necessary to fully account for deviations from Gaussianity and to establish strong agreement between the distribution derived from real data and the mathematical model/simulations based on gaussian statistics.

B. H1 Detector Background

We find, after incorporating all of the data quality procedures described in Sec. III B, that the resulting $\Delta E^{\mathrm{excess}}$ distribution is accurately described by the the-

oretical PDF (2.9) if we partition the S5 run into two, independent sub-runs: the first lasts from the beginning of S5 in November 2005 to June 2006, and the second lasts from June 2006 to the end of S5 in October 2007. The division in June 2006 corresponds to a commissioning break during which the H1 detector received a number of upgrades. As a result, the detector configurations in each of the sub-runs are considerably different, and, in the context of our analysis, should be treated as distinctly different detector systems. Fig. 7 shows the background distributions obtained for each of the two sub-runs.
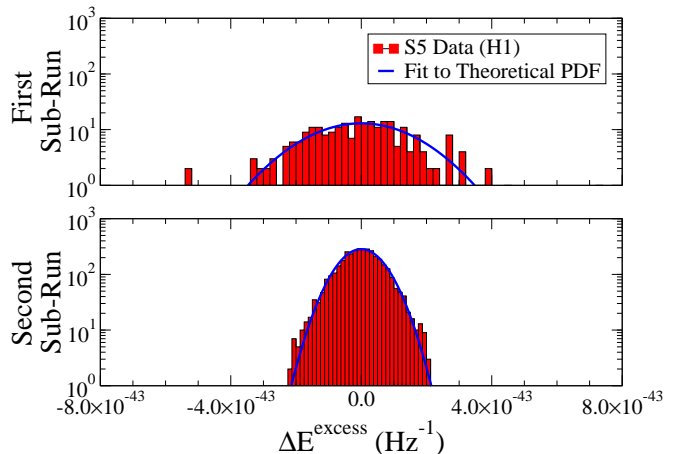

FIG. 7: H1 detector background (histogram) for each of the sub-runs November 2005 – June 2006 and June 2006 – October 2007 using category 1 and category 2 DQ flags as well as a glitch rejection step with 10 ms tiles and a 4σ veto threshold (see Sec. III B). For the production of the lower plot, only data before April 2007 was used. The solid curves are fits of the theoretical PDF (2.9). The fit parameters are obtained using nonlinear least squares regression to fit the theoretical CDF (2.10) to the empirical CDF of the background distribution.

Given the agreement between the resulting background distributions and the theoretical model (2.9), we conclude that our data quality methodology is sufficient to account for non-Gaussianity arising from noise contamination in the S5 data.

C. Extension to Multi-Detector Networks

Extending the concept to encompass of the full LIGO network requires us to consider how to pool information recorded by multiple instruments into a single analysis, as well as the orientation of each detector at the times of the observed triggers. We extend to a multi-detector network by first performing the single-detector analysis described in Sec. III for each, independent detector. This gives a set of p-values representing the significance with

which the two-sample K-S test can distinguish between the background and foreground distributions in the data stream of each detector. Assuming a network of N independent detectors, Fisher's method computes the test statistic

$$X^2 = -2 \sum_{i=1}^{N} \ln p_i \qquad (5.2)$$

which follows a χ^2-distribution with $2N$ degrees of freedom [33]. We claim detection for the detector network if X^2 is sufficiently large that we can reject the null hypothesis with 5σ significance. We choose to explore Fisher's method over a coherent method for two main reasons: (i) Fisher's method does not require us to assume any relationship between the waveforms observed in different detectors beyond the existence of excess energy in the QPO band, and (ii) a significant fraction of triggers will only be recorded by one detector because of downtime, and coherent methods are not applicable for these triggers.

Likewise, we can use Fisher's method to define the sensitivity for the detector network. We obtain a set of antenna factors for a given detector by calculating the antenna response function at the time of each trigger in the Dominant Polarization Frame (DPF) of that detector [44]. For a single detector, this is the frame in which the detector response is entirely in the plus direction ($F_\times \equiv 0$) [29]. We then get a multi-detector sensitivity result by first fixing a value of $h_{\rm rss}$, and then performing the single-detector sensitivity analysis described in Sec. IV B for each detector in the network individually, weighting injections with the appropriate set of antenna factors. We then combine the single-detector results using Fisher's method, and define the sensitivity of the network to be the minimum injected $h_{\rm rss}$ such that we can reject the null hypothesis with 5σ significance in 50% of trials. In the remainder of the paper we will consider the H1-L1 network, where L1 is the 4-km Livingston detector.

In addition, we must consider the fact that the detector network is not necessarily equally sensitive to all triggers observed from a single source since the relative positions of the detectors and source change with time. Fig. 8a shows the quantity F_+ for each trigger and for H1 and L1 separately. Since the algorithm becomes more sensitive as the number of triggers increases, but is less sensitive if the detector is poorly aligned for some of the triggers, we expect that there is an optimal subset of the full list of triggers for which the sensitivity of a multi trigger approach is maximized. We perform this optimization by repeating the sensitivity analysis using simulated, Gaussian noise described in Sec. IV B for H1 and L1 individually, weighting the injection waveforms with the appropriate antenna factors, and then combining the results using Fisher's method to get a sensitivity for the hypothetical H1-L1 network. We then find the worst-aligned member (smallest value of F_+) of the combined

list of H1 and L1 triggers, exclude this trigger for both detectors, and compute the sensitivity again. In Fig. 8b we plot the sensitivity as a function of the number of triggers, obtained by iterating this process.
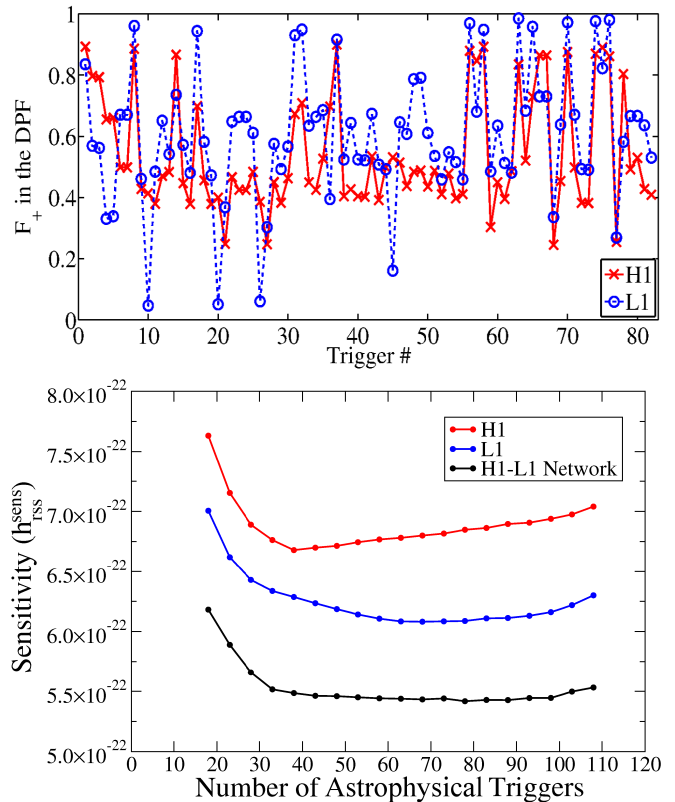

FIG. 8: The upper plot shows the F_+ values for H1 and L1 in their respective Dominant Polarization Frames (DPF-s, see [29]) calculated for the direction of the SGR 1806−20 source at the times of the astrophysical triggers observed during S5. In the lower plot we compute the sensitivity for Gaussian noise with strain spectral noise density $\tilde{n} = 10^{-22}\,{\rm Hz}^{-1/2}$ as a function of the number of triggers included in the analysis with the highest corresponding F_+ values. In each case we choose the number of these triggers such as to maximize sensitivity (i.e. minimize $h_{\rm rss}^{\rm sens}$).

We conclude that the sensitivity of the algorithm to GW emissions from SGR 1806−20 is maximized for the S5 H1-L1 data set if we restrict only to the 78 triggers for which the detectors were best aligned with the source. The same procedure can easily be repeated to tune the algorithm for other data sets or sources. Furthermore, after repeating the simulation of Sec. IV B with the detector orientation taken into account we find that the single detector sensitivities are

$$(h_{\rm rss}^{\rm sens})_{\rm H1} = 6.68 \times 10^{-22}\,{\rm Hz}^{-1/2} \qquad (5.3)$$
$$(h_{\rm rss}^{\rm sens})_{\rm L1} = 6.08 \times 10^{-22}\,{\rm Hz}^{-1/2} \qquad (5.4)$$

and the sensitivity for the H1-L1 network obtained using

Fisher's method is

$$(h_{\text{rss}}^{\text{sens}})_{\text{H1-L1}} = 5.42 \times 10^{-22}\,\text{Hz}^{-1/2}. \qquad (5.5)$$

Again assuming a detector-source distance of $14\,\text{kpc}$ for SGR $1806-20$, the corresponding isotropic GW emission energies are

$$\left(E_{\text{GW}}^{\text{iso}}\right)_{\text{H1}} = 2.84 \times 10^{46}\,\text{erg} \qquad (5.6)$$

$$\left(E_{\text{GW}}^{\text{iso}}\right)_{\text{L1}} = 2.35 \times 10^{46}\,\text{erg} \qquad (5.7)$$

$$\left(E_{\text{GW}}^{\text{iso}}\right)_{\text{H1-L1}} = 1.87 \times 10^{46}\,\text{erg} \qquad (5.8)$$

VI. CONCLUSION

We have investigated aspects of a multi-trigger data analysis concept targeting the detection of long-duration narrow-band gravitational wave transients. Potential astrophysical motivators include quasi-periodic oscillations of magnetars that are already observed in the electromagnetic spectrum following soft gamma repeater hyperflare events. This concept significantly extends a single-trigger method introduced by Matone and Márka [22] both in data handling sophistication and comprehensive use of available external information. The method is based on an excess-energy statistic targeting long-duration, narrow-band signals, and incorporates a number of data quality measures to account for nonstationarity in the data stream of a physical detector over long periods of time. The detection threshold and sensitivity are both defined in terms of the two-sample Kolmogorov-Smirnov test as a viable example.

Analytical results for the PDF and CDF of the detector background distribution enabled us to demonstrate that these expressions accurately describe the output of the pipeline applied to (i) Gaussian white noise and (ii) off-source LIGO S5 data from the 4-km Hanford detector after reasonable data quality cuts. In turn, they showed that the justified data quality cuts are efficient in removing the major contamination from the data. It is important to note that a useful side product of the method might be the ability to identify distinct running condition of the gravitational wave detectors (for a similar work in this topic, see [34]).

We also demonstrated how to combine the data streams of multiple, independent detectors into a single study on an alternative non-coherent way, and illustrated how to optimize the algorithm for a future search involving the S5 Hanford-Livingston data set. Simulating the optimal trigger list and using Gaussian white noise with strain spectral noise density $\tilde{n} = 10^{-22}\,\text{Hz}^{-1/2}$ resulted in an upper bound on the total energy emitted in isotropic gravitational radiation by SGR $1806-20$ in the 30 seconds following a flare of

$$\left(E_{\text{GW}}^{\text{iso}}\right)_{\text{H1-L1}} = 1.87 \times 10^{46}\,\text{erg} \times \left(\frac{d}{14\,\text{kpc}}\right)^2, \qquad (6.1)$$

where d is the detector-source distance.

Due to the uncertainties of currently available model predictions, throughout our study we assumed that (i) the central frequency, bandwidth, and time scale of a gravitational wave signal associated with an SGR QPO event can directly be taken from the electromagnetic observations of the QPO signal, (ii) the duration of the gravitational wave signals are the same for every QPO event of the same SGR source, and (iii) the duration of all the gravitational wave signals associated with the astrophysical triggers from an SGR source can be approximated with the duration of the longest flare event among the triggers (e.g. in case of SGR $1806-20$, it is the 27 December 2004 hyperflare event). These assumptions however might turn out to be too optimistic as theoretical model predictions become more sophisticated in the future. Our method can be applied to any chosen central frequency, and can be generalized to a wide scale of bandwidths. However, a more detailed study is necessary to see the consequences of having different durations for the gravitational wave signals emitted during the different flare events of the same SGR source. We leave this study to be part of a future work.

VII. ACKNOWLEDGMENTS

[1] A.L. Watts and T.E. Strohmayer. Detection with RHESSI of high frequency x-ray oscillations in the tail of the 2004 hyperflare from SGR $1806-20$. *Astrophys. J. Lett.*, 637:L117–20, 2006.

[2] G.L. Israel et al. Discovery of rapid x-ray oscillations in the tail of the SGR 180620 hyperflare. *Astrophys. J. Lett.*, 628:L53–6, 2005.

[3] J. H. Swank. The Rossi X-ray Timing Explorer: capabilities, achievements and aims. *Advances in Space Research*, 38:2959–2963, January 2006.

[4] R.P. Lin et al. High-energy solar spectroscopic imager (HESSI) small explorer mission for the next (2000) solar maximum. *Proc. SPIE*, 3442:2–12, 1998.

[5] A.L. Watts and T.E. Strohmayer. Neutron star oscillations and QPOs during magnetar flares. *Adv. Space Res.*, 40:1446–1452, 2007.

[6] A.M. El-Mezeini and A.I. Ibrahim. Discovery of quasi-periodic oscillations in the recurrent burst emission from SGR 1806−20. *Astrophys. J. Lett.*, 721:L121–L125, 2010.

[7] R. C. Duncan and C. Thompson. Formation of very strongly magnetized neutron stars - implications for gamma-ray bursts. *Astrophys. J. Lett.*, 392:L9–L13, June 1992.

[8] S.J. Schwartz et al. The gamma-ray giant flare from SGR 1806−20: evidence for crustal cracking via initial timescales. *Astrophys. J. Lett.*, 627:L129–32, 2005.

[9] D.M. Palmer et al. A giant gamma-ray flare from the magnetar SGR 1806−20. *Nature*, 434:1107–9, 2005.

[10] P.D. Lasky et al. Hydromagnetic instabilities in neutron stars. *Astrophys. J.*, 735:L20, 2011.

[11] Y. Levin. On the theory of magnetar QPOs. *Mon. Not. Roy. Astron. Soc.*, 377:159–167, 2007.

[12] M. van Hoven and Y. Levin. Magnetar oscillations i: strongly coupled dynamics of the crust and the core. *arXiv:1006.0348*, 2010.

[13] M. van Hoven and Y. Levin. Magnetar oscillations ii: spectral method. *arXiv:1110.2107*, 2011.

[14] H. Sotani, K. Kokkotas, and N. Stergioulas. Alfvén QPOs in magnetars. *Mon. Not. Roy. Astron. Soc.*, 385:L5, 2008.

[15] K. Glampedakis, L. Samuelsson, and N. Andersson. Elastic or magnetic? a toy model for global magnetar oscillations with implications for QPOs during flares. *Mon. Not. Roy. Astron. Soc.*, 371:L74–L77, 2006.

[16] I. Bartos, P. Brady, and S. Márka. How gravitational-wave observations can shape the gamma-ray burst paradigm. *arXiv:1212.2289*, 2012.

[17] Y. Levin and M. van Hoven. On the excitation of f-modes and torsional modes by magnetar giant flares. *Mon. Not. Roy. Astron. Soc.*, 418:659-663, 2011.

[18] B. Zink, P.D. Lasky, and K.D. Kokkotas. Are gravitational waves from giant magnetar flares observable? *Phys. Rev. D*, 85:024030, 2012.

[19] J. Abadie et al. Search for gravitational wave bursts from six magnetars. *Astrophys. J.*, 734:L35, 2011.

[20] B. Abbott et al. Search for gravitational wave bursts from soft gamma repeaters. *Phys. Rev. Lett.*, 101:211102, 2008.

[21] B.P. Abbott et al. Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm. *Astrophys.J.*, 701:L68–L74, 2009.

[22] L. Matone and S. Márka. Search algorithm for the detection of long-duration narrow-band transients in GW interferometers. *Class. Quant. Grav.*, 24:S649–S657, 2007.

[23] B. Abbott et al. Search for gravitational wave radiation associated with the pulsating tail of the SGR 1806−20 hyperflare of 27 December 2004 using LIGO. *Phys. Rev. D*, 76(6):062003, 2007.

[24] E. Thrane, S. Kandhasamy, C.D. Ott, W.G. Anderson, N.L. Christensen, et al. Long gravitational-wave transients and associated detection strategies for a network of terrestrial interferometers. *Phys. Rev. D*, 83:083004, 2011.

[25] B. Abbott et al. LIGO: the laser interferometer gravitational-wave observatory. *Rept. Prog. Phys*, 72:076901, 2009.

[26] G.B. Arfken and H.J. Weber. *Mathematical methods for physicists*. Academic Press, 6th edition, June 2005.

[27] E. Parzen. *Modern probability theory and its applications*. Wiley, 1992.

[28] N.J. Salkind. *The encyclopedia of measurement and statistics*. Sage, 2006.

[29] M. Drago. *Search for transient gravitational wave signals with unknown waveform in the LIGO-Virgo network of interferometric detectors using a fully coherent algorithm*. PhD thesis, UNIPD, 2010.

[30] B. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth LIGO science run. *Phys. Rev. D*, 80:102001, Nov 2009.

[31] J. Shao and D. Tu. *The jackknife and bootstrap*. Springer Series in Statistics. Springer Verlag, 1995.

[32] Gilad Svirski, Ehud Nakar, and Eran O. Ofek. SGR 1806−20 distance and dust properties in molecular clouds by analysis of a flare x-ray echoes. *Mon. Not. Roy. Astron. Soc.*, 415:2485, 2011.

[33] R.A. Fisher. *Statistical methods for research workers*. MacMillan, 14th edition, 1970.

[34] M. W. Coughlin (for the LIGO Scientific Collaboration and the Virgo Collaboration) Identification of long-duration noise transients in LIGO and Virgo *Class. Quant. Grav.*, 28:235008, 2011.

[35] For the case of the December 2004 SGR 1806−20 hyperflare a table summarizing the frequencies and FWHMs of the observed QPOs can be found in [23]. If assuming, based on the magnetar model, that the QPO frequencies and bandwidths are the same as the f-modes of the neutron star, these can also be the frequencies and bandwidths of the gravitational radiation [17, 18].

[36] Naively one would choose $w = 1/2$, but the noise floor of a physical instrument is neither flat, linear, or constant across all frequency bands. Instead, we use a locally determined value that is discussed in more detail in Sec. III B.

[37] We define the trigger time to be the arrival time of the observed EM burst.

[38] For a band-limited noise sample with bandwidth Δf and length Δt, the independence of the sampling points is an approximation that is only valid when $\Delta t \gg \Delta f^{-1}$.

[39] The isotropic equivalent energy of the GW emission, $E_{\mathrm{GW}}^{\mathrm{iso}}$, is measured in *erg*s, and should not be confused with the energy of a strain time series defined in (2.1) and measured in seconds.

[40] In general, when a nontrivial veto step is applied to the data stream of a physical detector, independent measurements will have different lengths of data removed by the veto based on the magnitude and duration of any glitches. The duty factor is proportional to the length of data remaining, and is necessary to renormalize the post-veto data stream so that independent measurements have comparable total energies. If the signal exhibits strong time dependence as it evolves this renormalization will introduce model/signal dependent distortions to the signal energy measurement, nevertheless it is efficient for the locally flat background.

[41] This nonstationarity can be seen in Fig. 5, which shows a plot of w calculated using LIGO data.

[42] For $\Delta t = 300$ s there are 82 such triggers.

[43] This corresponds to a K-S test p-value of $p = 2.7 \times 10^{-7}$.

[44] Since we are computing the energy of the injected wave-

form, which is polarization independent, we are free to pick a different polarization frame for each detector.