Draft version March 9, 2018
Typeset using LaTeX **twocolumn** style in AASTeX62

Observational implications of lowering the LIGO-Virgo alert threshold

Ryan Lynch,[1] Michael Coughlin,[2] Salvatore Vitale,[1] Christopher W. Stubbs,[3,4] and Erik Katsavounidis[1]

[1] *Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*
[2] *Division of Physics, Math, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*
[3] *Department of Physics, Harvard University, Cambridge, MA 02138, USA*
[4] *Department of Astronomy, Harvard University, Cambridge MA 02138, USA*

ABSTRACT

The recent detection of the binary-neutron-star merger associated with GW170817 by both LIGO-Virgo and the network of electromagnetic-spectrum observing facilities around the world has made the multi-messenger detection of gravitational-wave events a reality. These joint detections allow us to probe gravitational-wave sources in greater detail and provide us with the possibility of confidently establishing events that would not have been detected in gravitational-wave data alone. In this paper, we explore the prospects of using the electromagnetic follow-up of low-significance gravitational-wave transient event candidates to increase the sample of confident detections with electromagnetic counterparts. We find that the gravitational-wave alert threshold change that would roughly double the number of detectable astrophysical events would increase the false-alarm rate by 5 orders of magnitude from 1 per 100 years to 1000 per year. We quantify the expected purity of low-significance candidate alerts issued by LIGO-Virgo as a function of the alert threshold. Our analysis suggests that increasing the number of gravitational-wave detections via the electromagnetic follow-up observations of low-significance LIGO-Virgo events will incur significant human and opportunity costs in all-sky surveys, galaxy-targeted imaging, and large-aperture spectroscopy in the near future. Nevertheless, increases in the rate of detected gravitational-wave events may partially mitigate these costs as the Advanced LIGO-Virgo network reaches its design sensitivity.

1. INTRODUCTION

The August 2017 detection of GW170817 was an event of many firsts. Not only was it the first binary-neutron-star merger detected (Abbott et al. 2017e) by the LIGO-Virgo detector network (Abbott et al. 2015; Acernese et al. 2015; Abbott et al. 2016b), it was the first gravitational-wave (GW) event confidently detected by both ground-based GW detectors and electromagnetic (EM) observatories (Abbott et al. 2017f). While the detection of GW170817 could be confidently established by GW-detector data alone, the joint EM detection enabled a vast array of rich physical insights, such as the association of short gamma-ray bursts with binary-neutron-star mergers (Abbott et al. 2017d), a new procedure for constraining the value of the Hubble parameter H_0 (Abbott et al. 2017a), and evidence of heavy-element nucleosynthesis (Abbott et al. 2017c). The high signal-to-noise ratio of GW170817 and the clarity with which it could be distinguished in both GW and EM data aided the wealth of scientific information extracted by studying it. However, given an astrophysical population of such sources, we expect that quieter GW candidates might contribute a level of scientific richness to both the GW and EM communities as well. For example, if we were able to double the number of joint GW-EM detections by performing searches for low-significance candidates, we could possibly decrease the uncertainty in GW-based measurements of H_0 by up to $\frac{1}{\sqrt{2}}$ (Chen et al. 2017a).

In this paper, we examine the extent to which searches for low-significance GW transients can augment the total ensemble of GW detections. For the purpose of establishing a baseline, let us assume that the minimum false-alarm rate (FAR) at which GWs can be confidently detected by LIGO-Virgo alone is 1 per 100 years. For a 1-year observation run with 50% coincident detector up-time, this FAR threshold roughly corresponds to a p-value significance of 3σ. This FAR also corresponds to the nominal LIGO-Virgo alert threshold proposed for issuing open public alerts in the third Advanced LIGO-Virgo observing run (Public LIGO document 2018). In effect, we will define any GW event

ryan.lynch@ligo.org

with a FAR of greater than 1 per 100 years to be a low-significance event. Under this assumption, we cannot claim low-significance LIGO-Virgo events as confident detections unless they are jointly detected by EM observations, since this joint GW-EM detection would likely have a p-value significance of greater than 3σ. Thus, the EM follow-up of low-significance GW candidates enables the detection of sources that would otherwise be undetectable. In a sense, this method is the complement to the scenario where GW events with extremely small localization volumes enable the discovery of faint EM counterparts (Chen & Holz 2016).

However, there are potential opportunity costs that EM observers must weigh when considering how many GW candidates they can reasonably follow-up. By definition, following-up GW candidates at a higher FAR threshold means a greater number of false-alarm contaminations. Furthermore, low-significance candidates are inherently faint in GW detectors, implying that they will not be as well-localized for EM observations. The combination of these two factors along with finite observational resources suggests that we should try to quantify the cost-reward payoff of GW astronomy so that we do not wrongly prioritize it over other astronomical ventures. This analysis is complementary to work optimizing EM follow-up using tiling, time allocation, and scheduling methods (Ghosh, Shaon et al. 2016; Coughlin & Stubbs 2016; Salafia et al. 2017; Chan et al. 2017; Rana et al. 2017).

2. SOURCE AND BACKGROUND RATES IN GW DETECTORS

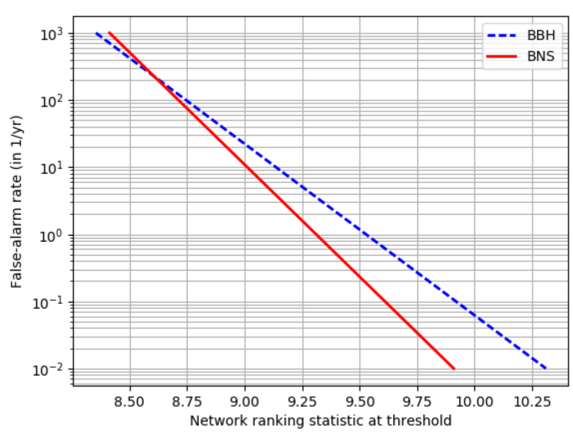

Figure 1. This log-linear plot shows the steep dependence of false-alarm rate versus the network ranking statistic (ρ) detection threshold for binary-black-hole (BBH) and binary-neutron-star (BNS) template fitting (Nitz et al. 2017). A modest $\Delta\rho = 1$ in the network ranking statistic increases the false-alarm rate by 2-3 orders of magnitude.

All search algorithms for transient GW events follow the same basic hypothesis: the signatures of GW events in every GW detector should be morphologically identical (once projection effects are taken into account) and time coincident, while detector noise need not be. The noise in each GW detector is a superposition of a Gaussian bulk and non-Gaussian noise transients. With low probability, this noise can mimic the appearance of GW events, which forms a background for the various search algorithms. One such search algorithm is PyCBC (Usman et al. 2016; Nitz et al. 2018), which uses a bank of compact-binary-coalescence templates to rank GW detection candidates according to a network ranking statistic, ρ, that combines the candidate signal-to-noise ratio (SNR) with signal consistency tests. Before we can proceed with our analysis of the prospects for low-significance GW searches, we must model how both the rate of GW search backgrounds and GW event populations behave as a function of ρ.

The bulk of the background distribution of searches for GW transients falls off steeply as a function of ρ, meaning the FAR changes by orders-of-magnitudes over narrow ranges of ρ. As a result, the LIGO-Virgo instruments have a very sharp ranking statistic cutoff above which GW events can be confidently detected and below which they cannot, with little room for ambiguity. To quantify this more precisely, we explore the background for the two LIGO detectors: one in Hanford, Washington, USA (H) and the other in Livingston, Louisiana, USA (L). The HL background rate (i.e., the FAR) decays roughly exponentially as a function of ρ for FARs between 1 per 100 years and 1000 per year. This behavior is observed both for searches for compact-binaries (Abbott et al. 2012; Nitz et al. 2017) and searches for short-duration unmodeled GW events (Abbott et al. 2017b). The steepness of this exponential falloff is determined by how easy it is for background events to mimic GWs in a given search. Thus, searches for binary-neutron-star (BNS) events have a steeper falloff than searches for binary-black-hole (BBH) events because a known time-frequency evolution is observed over longer durations for BNS events than for BBH events. Likewise, searches for short-duration unmodeled GW events have less-steep exponential falloff than for either BNS or BBH events because the former's time-frequency evolution is inherently unknown and thus less constrained.

For this paper, we focus only on searches for BBH and BNS events, since both of these source-types have already been detected by LIGO-Virgo. These detections have allowed for their rates to be observationally established. As mentioned, BNS events have already been jointly detected by LIGO-Virgo and EM observers (Abbott et al. 2017f), making them the most anticipated targets for low-significance efforts. We do not present the results for short-duration unmodeled events since we do not have any direct measurements of their rate. However, the relative results, obtained by normalizing out these unknown rates, are of similar magnitude to those for both BNS and BBH sources, resembling the results for BBH sources more closely.

We perform the exponential fit using the FAR versus ρ relationship reported in (Nitz et al. 2017), which represents the PyCBC search background when analyzing the HL data during the first Advanced LIGO-Virgo observing run. The results of this fit are shown in Fig. 1. We assume that the slopes of these fits are representative of the BBH and BNS HL searches in current and future observing runs. This assumption is based upon empirical results: we observe similar fits in published results for both Advanced (Nitz et al. 2017) and initial (Abbott et al. 2012) LIGO-Virgo observing runs. As we will soon see, the potential impact of searches for low-significance GW events is reduced as the slope of the searches' background distributions becomes steeper. Thus, any improvements that may make the GW searches' backgrounds less-heavily tailed (and hence steeper), such as reducing the non-Gaussian contribution to the GW detectors' noise or adding a third detector like Virgo to reduce the coincidence rate of non-Gaussian noise transients, may further reduce the case for low-significance GW science. While we expect these exponential fits to be representative in future observing scenarios, we can characterize the uncertainties associated with them in several ways. First, the difference in the results for the BBH and BNS searches demonstrate the uncertainties in the slope of these fits, with the BNS results representing a steeper fit and the BBH results representing a less-steep fit. Second, we allow the vertical normalization of these fits (at constant slope) to vary by up to an order of magnitude to exemplify the uncertainties in the overall rate of background events.

We must likewise find a model to describe the rate of GW events versus ρ. Assuming that GW events are distributed uniformly in volume, the *cumulative* rate of GW events exceeding an SNR threshold should roughly scale as SNR^{-3} for Advanced-era GW detectors probing the low-redshift universe (Schutz 2011; Chen & Holz 2014; Vitale 2016) (although this scaling relation will break down for third-generation GW detectors that probe higher redshifts (Vitale 2016)). By construction, we expect that the network ranking statistic $\rho \sim \mathrm{SNR}$ for real GW events (Nitz et al. 2017). As the Advanced GW detectors improve in sensitivity, the overall rate of GW events being observed will increase accordingly. We can correctly normalize these rates for an observing epoch by multiplying LIGO-Virgo's empirically motivated time-volume rate estimates (Abbott et al. 2016a, 2017e) with sensitive-volume estimates for each observing epoch. We estimate the cosmologically-corrected sensitive volume for each epoch at an SNR threshold of 8 using the online distance calculator provided by (Chen et al. 2017b). For BBH searches, we use the average total mass associated with both uniform-in-log and power-law (with a power of -2.35) mass distributions (Abbott et al. 2016a). The differences in the results for these two distributions are negligible, thus we only quote the power-law results in this paper. For BNS, we use the median total mass estimates for GW170817 (Abbott et al. 2017e).

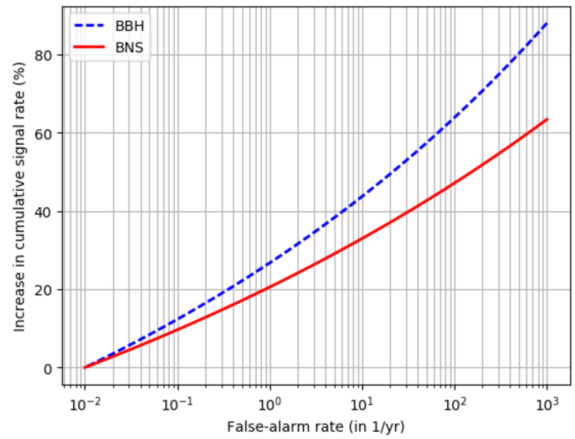

Figure 2. The relative increase in the number of GW events (in %) expected above each FAR threshold. Changing the false-alarm-rate threshold by 5 orders of magnitude increases the number of GW events by less than a factor of 2.

By combining the background and GW event models, we compute the expected rate of GW events at each FAR threshold. In Fig. 2, we plot the relative increase in GW events expected above each FAR threshold as compared to the baseline FAR threshold of 1 per 100 years. The most notable result is that increasing the FAR threshold by 5 orders of magnitudes from 1 per 100 years to 1000 per year only increases the number of detectable GW events by about 60% for BNS and 90% for BBH. In other words, because the GW event rate scales as ρ^{-3}, we need to change the ρ threshold by a factor of $\left(\frac{1}{2}\right)^{\frac{1}{3}} \sim 0.8$ to gain a factor of 2 in the number of detectable GW events. However, changing ρ from ~ 10 (corresponding to a FAR threshold of 1 per 100 years) by a factor of 0.8 increases the FAR contamination by an astonishing 5 orders of magnitude (see Fig. 1). We emphasize that these curves represent a best-case scenario in which the EM observations confidently detect every low-significance GW event. Only the events conclusively detected in the EM would be labeled as confident detections. Thus, the actual increase in the total number of confident GW detections may be significantly lower that the best-case-scenario factor of 2.

Because the uncertainty in the GW event rate is a constant normalization factor, it does not factor into Fig. 2. The only uncertainties that affect this plot are therefore related to the background fit. We manually vary the normalization of the total background rate by up to an order of magnitude, however this only results in negligible uncertainties in the expected rate of detectable GW events. The effect of the uncertainty regarding the slope of our exponential background fits is illustrated by comparing the results for the steeper BNS background to those for the less-steep BBH background. The relative increase in the rate of detectable GW events is greater for BBH searches than for BNS searches because the range of ρ spanned at these FARs is greater for BBH searches (see Fig. 1). The results of both of these searches are of similar magnitude across all FARs. Thus, we would only expect low-significance efforts to have a more sizable effect on the rate of detectable GW events if the backgrounds for any search were to become drastically less Gaussian and more heavily-tailed.

3. PURITY OF LOW-SIGNIFICANCE ALERTS

Although it is interesting to explore the GW event rate associated with each FAR threshold, we must take into account that any increase may come at a cost to observers. Issuing GW follow-up alerts to EM astronomers at higher FARs inherently means that these alerts will have higher contamination rates. In Fig. 3, we plot the expected probability of a GW alert being a real GW event at each FAR using the predicted sensitivities for three GW observing epochs (Abbott et al. 2016b): the second Advanced LIGO-Virgo observing run (O2), the third Advanced LIGO-Virgo observing run (O3), and the eventual design sensitivity Advanced LIGO-Virgo observing runs. We depict the dominant uncertainties associated with the GW event rate as the shaded regions, while plotting the results associated with the median published rates as lines. We give the explicit values of some of these probabilities (corresponding to the median published rates) at several ad hoc FAR thresholds in Table 1. At low FAR thresholds (like our baseline of 1 per 100 years), we have very high alert purity, meaning there is a great likelihood of success to offset any costs incurred to EM observers. At higher FAR thresholds, the purity is strongly dependent upon the expected event rate, i.e., the sensitivity of the detectors. In O2, LIGO-Virgo observed relatively low GW event rates, meaning high-FAR GW alerts had relatively low

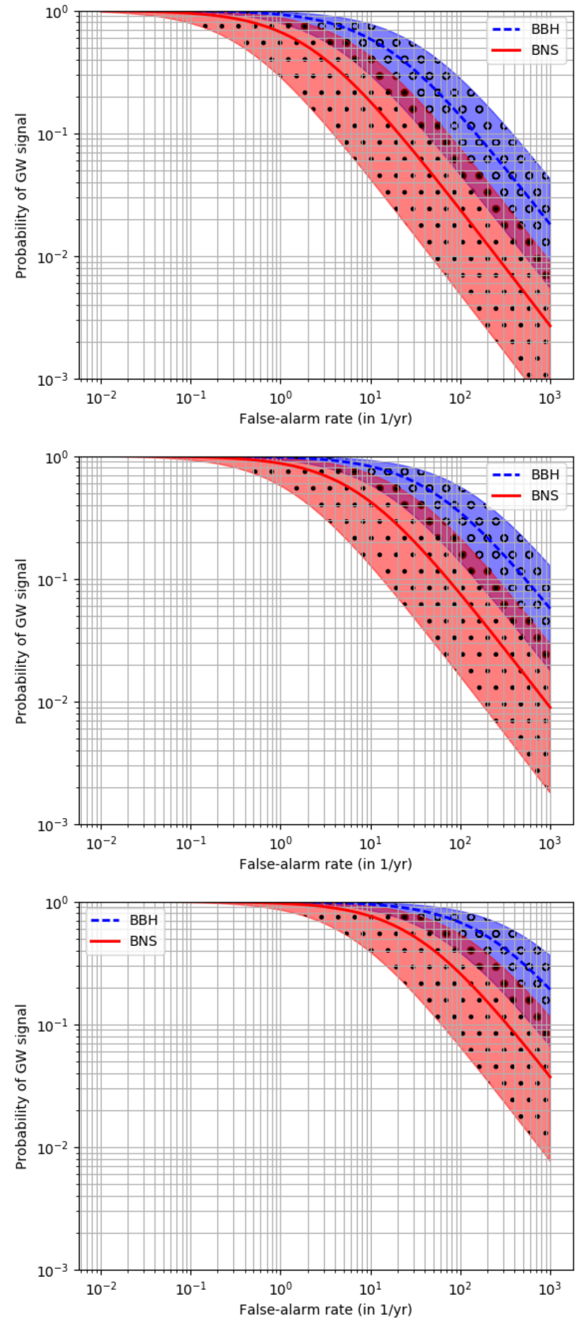

Figure 3. The probability of a GW candidate being a GW event (rather than a false alarm) at each FAR threshold for three observing epochs: O2 (top), O3 (middle), and design sensitivity (bottom). The solid lines correspond to the median published rates for the average BBH and BNS systems (Abbott et al. 2016a, 2017e), and the shaded regions correspond to the 90% confidence regions for those same rates.

probabilities of being events. However, in more sensitive observing epochs, such as when Advanced LIGO-Virgo reaches design sensitivity, the expected event rate

Table 1. The probability (in %) of GW candidates being GW events (P_{GW}), as depicted in Fig. 3, at several ad hoc FAR thresholds for various LIGO-Virgo observing epochs. We also give the corresponding fractional increase of GW events (FI_{GW}) for these same thresholds. These values correspond to the median published rates for the average BBH and BNS systems (Abbott et al. 2016a, 2017e). Note that the errors on the probabilities corresponding to rate uncertainties can be large (see Fig. 3). For O2, the probability of a candidate being a GW could degrade by more than an order of magnitude for lower alert thresholds while still not doubling the number of GW events. This degradation is less severe as LIGO-Virgo improves and reaches O3 and design levels of sensitivity.

Epoch	1 per 100 years	1 per year	1 per month	1 per week	1 per day
P_{GW} O2 BBH	99.9	92.6	54.3	22.9	4.5
P_{GW} O2 BNS	99.4	66.6	15.6	4.4	0.7
P_{GW} O3 BBH	99.9	97.6	79.6	49.4	13.5
P_{GW} O3 BNS	99.8	86.9	38.2	13.2	2.3
P_{GW} Design BBH	99.9	99.4	93.8	79.3	37.9
P_{GW} Design BNS	99.9	96.6	72.6	39.4	9.2
FI_{GW} BBH	1.0	1.27	1.45	1.58	1.77
FI_{GW} BNS	1.0	1.21	1.34	1.43	1.56

is large enough that even high-FAR GW alerts can have a high purity. For example, assuming the median event rates (Abbott et al. 2016a, 2017e), we would have needed a FAR threshold of 2 per year for BNS alerts in O2 to have a 50% chance of being real events, while at design sensitivity we could instead have a FAR threshold of 30 per year. We also note that because the lower limit on the rates of unmodeled GW sources has yet to be empirically established, there is a possibility that the purity of their alerts remains close to 0 across all FARs even when design sensitivity is reached, hence their absence from this analysis.

4. OBSERVATIONAL IMPLICATIONS FOR EM FOLLOW-UP CAMPAIGNS

In Section 3 we described how raising the FAR threshold increases the contamination rate of LIGO-Virgo alerts. There is however an additional cost associated with low-significance GW events that may directly impact their EM follow-up: events are more poorly localized by the GW detectors as their SNR decreases. A simple Fisher matrix analysis (Cutler & Flanagan 1994) shows that the uncertainty in GW distance estimates, σ_D, roughly scales as $\sigma_D \propto \mathrm{SNR}^{-2}$ (Fairhurst 2017). More in-depth calculations can be used to show that the angular area uncertainty in GW localization estimates, σ_A, roughly scales as $\sigma_A \propto \mathrm{SNR}^{-2}$ (Wen & Chen 2010), which agrees with the findings of Monte Carlo studies (Berry et al. 2015). Thus, applying an additional factor of distance squared ($\propto \mathrm{SNR}^{-2}$) to convert the angular area uncertainties to proper area uncertainties, we expect the total uncertainty in GW localization volume, σ_V, roughly scales as $\sigma_V \propto \mathrm{SNR}^{-6}$, which again agrees with the findings of Monte Carlo studies (Del Pozzo et al. 2018). Here, we assume that the network ranking statistic $\rho \sim \mathrm{SNR}$ for all low-significance GW alerts.

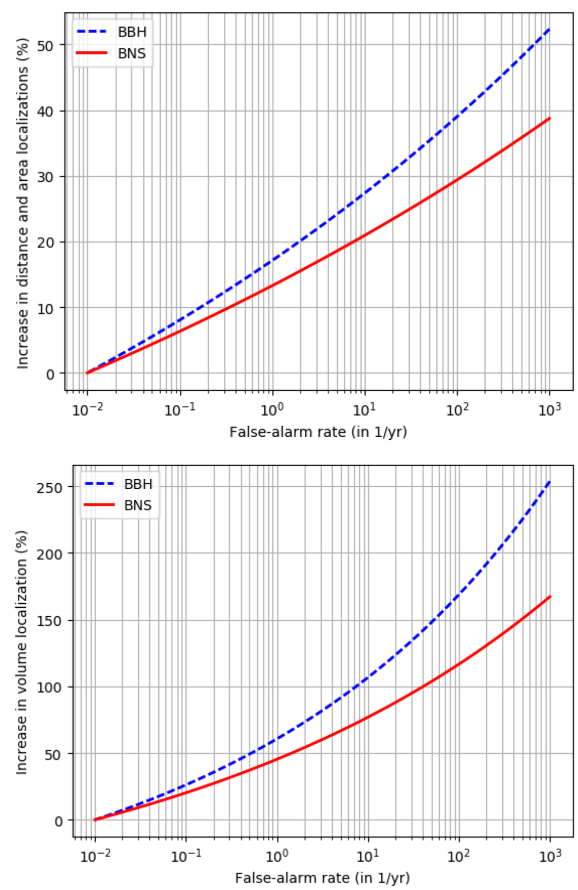

Figure 4. The relative increase (in %) in distance and angular area localizations (top) and volume localization (bottom) expected for a threshold event at each FAR. Changing the false-alarm-rate threshold by 5 orders of magnitude increases the distance and angular area localizations by less than a factor of 2 and the volume localization by less than a factor of 4.

In Fig. 4, we plot the relative increase in the GW distance, angular area, and volume localizations for threshold events, as compared to our baseline FAR of 1 per 100 years. Similarly to the results for GW event rates, we find that the relative increase in localization is a slowly-varying function of FAR. Increasing the FAR threshold by 5 orders-of-magnitude from 1 per 100 years to 1000 per year degrades the distance and angular area localizations by less than a factor of 2 and the volume localization by less than a factor 4. The reason the localization of BBH events has a larger relative increase than for BNS events is that the range of ρ spanned at these FARs is greater for BBH searches. Again, the uncertainties in these results are associated with our fits of the LIGO-Virgo backgrounds. The uncertainty associated with the total background rate normalization is explored by varying this rate by up to an order of magnitude, and the corresponding localization uncertainties are negligible. The uncertainty associated with the slope of the exponential fit is illustrated by comparing the results for the steeper-background-distribution BNS search to those of the less-steep-background-distribution BBH search. The results for both of these searches are of similar magnitude across all FARs.

We will now discuss how these observational costs will realistically affect EM follow-up. Here we focus specifically on follow-up procedures, although it should be noted that a serendipitous coincident detection of GW candidates with high-energy telescopes like Fermi/GBM, INTEGRAL/SPI-ACS, Konus/WIND (Abbott et al. 2017d) may affect the significance of GW candidates (Blackburn et al. 2015). These instruments have the advantage of continually monitoring a big fraction of the high energy sky, meaning they do not need to be run in follow-up mode. Additionally, they are usually subject to backgrounds that are overall quieter than the corresponding ones in optical bands. Thus, they present a low-cost means of potentially increasing the significance of GW events in near real time (as was the case with GW170817 (Abbott et al. 2017f)).

The EM follow-up observations of GW counterparts are undertaken in stages. Transients detected by imaging systems are assessed by spatial location (either 2- or 3-dimensional), broadband spectral characteristics, and light curve temporal evolution. These assessments can be accomplished with 2-4 meter aperture telescopes. If a viable EM counterpart is detected, large-aperture (8-10 meter class) spectroscopic observations are obtained.

We consider the impact of a higher false-alarm rate and poorer localization on EM follow-up efforts in three regimes:

- wide-field surveys, for which follow-up observations amount to re-ordering the sequence in which regions of the sky are observed;
- galaxy-targeted or other narrow-field imaging programs, which search for transients consistent with GW counterparts; and
- large-aperture spectroscopic follow-up campaigns, which obtain spectra of individual sources of interest.

4.1. *Wide-Field Sky Surveys*

For optical/infrared surveys carrying out high-cadence observations of the entire sky, responding to a GW alert is simply a matter of re-ordering the sequence of observations and perhaps changing broadband filters more rapidly than would otherwise be the case. The scientific opportunity cost of re-ordering the observations can be weighed against the probability of detecting the EM counterparts of low-significance GW alerts. This evaluation will be facilitated if the LIGO-Virgo alerts include early assessments of the relative likelihood of being a real event versus a false alarm. In lieu of such information, Fig. 3 can be used as a rough proxy.

Examples of wide-field imaging systems are the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) (Morgan et al. 2012), the Asteroid Terrestrial-impact Last Alert System (ATLAS) (Tonry 2011), the Zwicky Transient Factory (ZTF) (Bellm 2014), and eventually, the Large Synoptic Survey Telescope (LSST) (Ivezic et al. 2008). During the Advanced LIGO-Virgo runs O1 and O2, observations by survey telescopes contributed significantly to the follow-up program for many of the candidates (Smartt et al. 2016a,b; Stalder et al. 2017; Smartt, S. J. et al. 2017), both in estimating the most recent time of non-detection and in observing the fields after a GW alert.

For these systems, the primary observational cost is the loss of on-sky efficiency due to the additional filter changes, which require a time overhead that could otherwise be used for observation. The acquisition of the images can therefore be accomplished with minimal opportunity cost. This prospect of augmenting the number of joint GW-EM detections with all-sky survey systems was a prime motivation for the work described in this paper. Nevertheless, if the ratio of false alarms to actual events is high, the task of discriminating (via EM observations) a handful of GW events from an orders-of-magnitude greater number of false alarms becomes costly. This discrimination requires a combination of pipeline image processing and human intervention, with the demands of both scaling as the inverse of the alert purity. As a result, referencing Fig. 3, the use of wide-field surveys to search for low-significance GW events may not become practical until LIGO-Virgo reaches design sensitivity so that the alert purity is manageable at high FARs. We also note that the search efficiency of such discrimination may degrade at these high FARs as a result of the GW events being more distant and more poorly localized.

4.2. *Targeted Imaging Observations*

For narrow-field imaging systems that either target individual galaxies or tile the GW localization region on the sky(Coulter et al. 2017; Soares-Santos et al. 2017; Valenti et al. 2017; Arcavi et al. 2017; Tanvir et al. 2017; Lipunov et al. 2017), the EM follow-up observations are often conducted in a target-of-opportunity mode where previously scheduled programs are preempted by the GW follow-up campaign. In principle, one could envision dedicating a narrow-field follow-up system entirely to GW follow-up observations. Evaluating the scientific merit of dedicated EM follow-up system was another motivation behind the work described in this paper.

For narrow-field imaging, the increased alert rate rather than the modest increases in angular localization area (relevant for tiling) or localization volume (relevant for galaxy targeting) dominates the follow-up time requirement. Thus, the observational cost scales in proportion to the FAR. By comparison, and using as an example the H_0 measurement, the best-case benefit of such low-significance searches scale as roughly the square root of the number of additional detections, though the actual maximum reduction in uncertainty might be *even lower* than this on account of the poor distance localization of these low-significance events (Chen et al. 2017a) (Chen et al. 2017a). The discrimination challenges described for wide-field surveys still apply, but with the added scientific opportunity cost of the preempted observing program. Thus, targeted imaging observations may not have a high enough scientific payoff until the purity of LIGO-Virgo alerts improves at high FARs. Even then, the limited number of additional detections (less than a factor of 2) sets the ceiling of this payoff.

4.3. *Large-Aperture Spectroscopy*

Interrupting the observing program of a large-aperture spectroscopic telescope (such as Keck (Kasliwal et al. 2016) or Gemini (Chornock et al. 2017)) to obtain a sequence of spectra for a faint transient is arguably the most costly element in low-significance EM follow-up observations. If the overwhelming majority of the low-significance GW alerts are false alarms rather than real events, the fraction of large telescope time that is wasted is again proportional to the alert FAR. Modest-resolution spectroscopy is extremely valuable for both discrimination and characterization of EM counterparts, but at the same time large aperture telescopes are typically the most over-subscribed resource in the arsenal of follow-up tools. Given that we expect low-significance follow-up efforts to offer a relatively limited number of additional detections, we conclude that following-up low-significance GW candidates is likely to be a poor use of large aperture telescope time for the foreseeable future.

5. CONCLUSIONS

The background in gravitational-wave searches for binary systems with LIGO-Virgo is a steeply falling function of the detection statistic. As a result, a reduction on the GW alert threshold by, say, a factor of 0.8 (so that the accessible number of GW sources is doubled) would inflict an increase of *five orders of magnitude* in the false-alarm contamination This increase raises the false-alarm rate from the nominal value of 1 per 100 years to a value of 1000 per year. The lower significance also implies that these augmented events will be further away and will have larger distance, area, and volume localizations. If all of these additional GW events were confidently detected in the EM despite these contamination, distance, and localization costs, the corresponding increase in the number of detections that may be used in an ensemble measurement would be of order 2. As a result, even in the best-case scenario, we expect low-significance detections to offer limited scientific payoff for GW-based measurements.

This limited scientific payoff lies in stark contrast to the costs of low-significance EM follow-up campaigns. Even for wide-field surveys where the opportunity cost of conducting the observations is small, the task of finding the handful of GW events per 10^3 false-alarm triggers requires what may be a high investment of human capital. This is the most realistic scenario for near-future LIGO-Virgo observing epochs. However, as the LIGO-Virgo sensitivity continues to improve, the increased GW event rate results in a higher purity of low-significance GW candidates, driving the ratio of true events to false-alarms higher. Thus, EM follow-up searches for low-significance GW events may become more intriguing when the Advanced LIGO-Virgo detectors reach their design sensitivities.

6. ACKNOWLEDGMENTS

The authors acknowledge the support of the National Science Foundation and the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. MC is supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. We would also like to thank Hsin-Yu Chen, Kwan Yeung Ng, Yiwen Huang, Robert Eisenstein, Satya Mohapatra, Steve Drasco, Peter Shawhan, Alex Nitz, Tito Dal Canton, Christopher Berry, Thomas Dent, Reed Essick, and the LIGO-Virgo EM follow-up working group for useful comments and discussion. CWS is grateful to Harvard University and to the US Department of Energy (through grant DE-SC000788) for their support of this effort. This is LIGO document number P1800042.

REFERENCES

Abbott et al., B. 2012, Phys. Rev. D, 85, 082002

—. 2015, Class. Quant. Grav., 32, 074001

—. 2016a, Phys. Rev. X, 6, 041015

—. 2016b, Living Reviews in Relativity, 19, 1

—. 2017a, Nature, 551, 85

—. 2017b, Physical Review D, 95, 042003

—. 2017c, Astrophys. J., 850, L39

—. 2017d, Astrophys. J. Lett., 848, L13

—. 2017e, Phys. Rev. Lett., 119, 161101

—. 2017f, Astrophys. J. Lett., 848, L12

Acernese, F., et al. 2015, Class. Quant. Grav., 32, 024001

Arcavi, I., Hosseinzadeh, G., Howell, D. A., et al. 2017, Nature, 551, 64

Bellm, E. 2014, in The Third Hot-wiring the Transient Universe Workshop, ed. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 27–33

Berry, C. P. L., et al. 2015, Astrophys. J., 804, 114

Blackburn, L., Briggs, M. S., Camp, J., et al. 2015, Astrophys. J. Suppl., 217, 8

Chan, M. L., Hu, Y.-M., Messenger, C., Hendry, M., & Heng, I. S. 2017, The Astrophysical Journal, 834, 84

Chen, H.-Y., Fishbach, M., & Holz, D. E. 2017a, ArXiv e-prints, arXiv:1712.06531

Chen, H.-Y., & Holz, D. 2014, in APS Meeting Abstracts, C15.001

Chen, H.-Y., & Holz, D. E. 2016, ArXiv e-prints, arXiv:1612.01471

Chen, H.-Y., Holz, D. E., Miller, J., et al. 2017b, ArXiv e-prints, arXiv:1709.08079

Chornock, R., et al. 2017, Astrophys. J., 848, L19

Coughlin, M., & Stubbs, C. 2016, Experimental Astronomy, 1

Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, Science, 358, 1556

Cutler, C., & Flanagan, É. E. 1994, PhRvD, 49, 2658

Del Pozzo, W., Berry, C., Ghosh, A., Haines, T., & Vecchio, A. 2018, ArXiv e-prints, arXiv:1801.08009

Fairhurst, S. 2017, ArXiv e-prints, arXiv:1712.04724

Ghosh, Shaon, Bloemen, Steven, Nelemans, Gijs, Groot, Paul J., & Price, Larry R. 2016, Astronomy and Astrophysics, 592, A82

Ivezic, Z., Tyson, J. A., Allsman, R., Andrew, J., & Angel, R. 2008, arXiv:0805.2366

Kasliwal, M. M., Cenko, S. B., Singer, L. P., et al. 2016, Astrophys. J. Lett., 824, L24

Lipunov, V. M., Gorbovskoy, E., Kornilov, V. G., et al. 2017, Astrophys. J. Lett., 850, L1

Morgan, J. S., Kaiser, N., Moreau, V., Anderson, D., & Burgett, W. 2012, Proc. SPIE Int. Soc. Opt. Eng., 8444, 0H

Nitz, A., Harry, I., Brown, D., et al. 2018, https://doi.org/10.5281/zenodo.596388

Nitz, A. H., Dent, T., Dal Canton, T., Fairhurst, S., & Brown, D. A. 2017, Astrophys. J., 849, 118

Public LIGO document. 2018, https://dcc.ligo.org/public/0089/M1200055/002/LVCTriggerPolicy_final.pdf

Rana, J., Singhal, A., Gadre, B., Bhalerao, V., & Bose, S. 2017, The Astrophysical Journal, 838, 108

Salafia, O. S., Colpi, M., Branchesi, M., et al. 2017, The Astrophysical Journal, 846, 62

Schutz, B. F. 2011, Classical and Quantum Gravity, 28, 125023

Smartt, S. J., Chambers, K., Smith, K., et al. 2016a, Mon. Not. Roy. Astron. Soc., 462, 4094

Smartt, S. J., Chambers, K. C., W., S. K., et al. 2016b, Astrophys. J. Lett., 827, L40

Smartt, S. J. et al. 2017, Nature, 551, 75 EP

Soares-Santos, M., Holz, D. E., Annis, J., et al. 2017, Astrophys. J. Lett., 848, L16

Stalder, B., Tonry, J., Smartt, S. J., et al. 2017, The Astrophysical Journal, 850, 149

Tanvir, N. R., Levan, A. J., González-Fernández, C., et al. 2017, Astrophys. J. Lett., 848, L27

Tonry, J. L. 2011, Publ. Astron. Soc. Pac., 123, 58

Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, Classical and Quantum Gravity, 33, 215004

Valenti, S., David, Sand, J., et al. 2017, Astrophys. J. Lett., 848, L24

Vitale, S. 2016, PhRvD, 94, 121501

Wen, L., & Chen, Y. 2010, Phys. Rev. D, 81, 082001