

Optical follow-up of gravitational wave triggers with DECam

K Herner¹, J Annis¹, E Berger², D Brout³, R Butler¹, H Chen⁴,
 P Cowperthwaite², H Diehl¹, Z Doctor⁴, A Drlica-Wagner¹, B Farr⁴,
 D Finley¹, J Frieman¹, D Holz³, R Kessler³, H Lin¹, J Marriner¹,
 E Nielsen¹, A Palmese^{5,1}, M Sako³, M Soares-Santos¹, F Sobreira^{6,7}
 and B Yanny¹

¹Fermi National Accelerator Laboratory, Batavia, IL

²Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA

³Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA

⁴Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL

⁵University College London, London, United Kingdom

⁶Laboratorio Interinstitucional de e-Astronomia, Rio de Janeiro, Brazil

⁷Universidade de Campinas, Campinas, Brazil

E-mail: kherner@fnal.gov

Abstract. Gravitational wave (GW) events have several possible progenitors, including black hole mergers, cosmic string cusps, supernovae, neutron star mergers, and black hole–neutron star mergers. A subset of GW events are expected to produce electromagnetic (EM) emission that, once detected, will provide complementary information about their astrophysical context. To that end, the LIGO-Virgo Collaboration has partnered with other teams to send GW candidate alerts so that searches for their EM counterparts can be pursued. One such partner is the Dark Energy Survey (DES) and Dark Energy Camera (DECam) Gravitational Waves Program (DES-GW). Situated on the 4m Blanco Telescope at the Cerro Tololo Inter-American Observatory in Chile, DECam is an ideal instrument for optical followup observations of GW triggers in the southern sky. The DES-GW program performs subtraction of new search images with respect to preexisting overlapping images to select candidate sources. Due to the short decay timescale of the expected EM counterparts and the need to quickly eliminate survey areas with no counterpart candidates, it is critical to complete the initial analysis of each night's images within 24 hours. The computational challenges in achieving this goal include maintaining robust I/O pipelines during the processing, being able to quickly acquire template images of new sky regions outside of the typical DES observing regions, and being able to rapidly provision additional batch computing resources with little advance notice. We will discuss the search area determination, imaging pipeline, general data transfer strategy, and methods to quickly increase the available amount of batch computing. We will present results from the first season of observations from September 2015 to January 2016 and conclude by presenting improvements planned for the second observing season.

1. Introduction

The detection of gravitation waves (GW) by the LIGO-Virgo Collaboration (LVC) [1] has ushered in a new era of astronomy. The first two GW events, GW150914 and GW151226 [2], resulted from binary black hole (BBH) mergers at distances of ~ 400 Mpc. The electromagnetic

(EM) signatures of BBH events are unknown. Other possible GW sources include binary neutron star (BNS) mergers and mixed (NS-BH) mergers, which have an EM signature sometime referred to as kilonova (KN). GW sources with an EM signature offer exciting scientific opportunities. For example, an analysis of the GW component can provide an independent measurement of distance to the source, while an analysis of the EM counterpart provides redshift information. Together, the two values provide an independent way to measure the rate of expansion of the Universe via the distance-redshift relation.

With these benefits in mind, the LVC has established partnerships with several collaborations around the world. Shortly after their initial analysis of a GW candidate, LVC shares the information with its partners, including estimated distance to the source, most likely type (e.g. BBH, BNS), and a map showing the probability of the source to be located at each point on the sky, an example of which is shown in Figure 1. The EM partners can then decide whether or not to search for a possible EM signature associated with candidate in a region of that sky that best suits their telescope’s coverage and maximizes the probability of detecting the source. So far, with only two detectors in the GW network, the localization of GW candidates is somewhat imprecise, meaning the LIGO probability maps typically cover hundreds of square degrees, also shown in Figure 1. LVC triggers are communicated via a private network to EM partners using the same protocol used by the gamma-ray community (GCN). EM counterpart candidates are also communicated amongst partners to enable additional followup observations.

One of the EM followup partners is the DES-GW group, consisting of members of the Dark Energy Survey (DES) and other members of the astronomical community. DES is one of the primary users of DECam, attached to the 4m Blanco Telescope at the Cerro Tololo Inter-American Observatory in Chile [3]. DECam is well-suited to the EM followup task due to its red-sensitive CCDs and its large aperture. The DES-GW program has obtained ~ 4 nights of telescope time per year. Those nights are added to the DES allocation of 105 nights per year. Once a trigger is received, our team uses an dedicated strategy code to decide whether or not to observe, how much telescope time to invest, and what sequence of observations to execute. If the decision is to observe, we then interrupt DES observations to perform the search.




Figure 1. Localization probability map received from the LVC for GW150914 (shown here in log scale). The area enclosing 90% total probability spans several hundred square degrees. Searches for EM sources are very challenging in such large areas.

2. The DES-GW Program

We can divide the DES-GW program into several phases which are realized for each trigger (Figure 2). The first phase is the immediate post-trigger assessment and determination of the overall detection probability. After that, the DES-GW group discusses its findings with DES management, and together they take a joint decision of whether or not to interrupt the main DES observational program that evening and carry out the EM search. The decision whether to perform optical followup is taken on a case-by-case basis and considers a variety of factors, including the total amount of the sky probability map that can be covered (this in turn depends on location and time of the trigger), and other observing priorities of the survey.

If the decision is to proceed with DES-GW observations, the group moves to the next phase of preparing a detailed observing plan (including a specific observational cadence) and briefing the evening’s observing team. Once the observation stage begins, automated scripts track the progress of each new search image and, once the image is transferred from CTIO, submit the corresponding image processing jobs to grid sites.

While observations and analysis are ongoing, we send updates to the network of partners via GCN. Occasionally the LVC will release an updated probability map some days after the initial trigger, in which the most probable regions for the source location may shift somewhat from the original position. In those cases we evaluate the new probability map, and modify the observational plans accordingly during successive observations.

The light from a KN fades over a matter of days. In order to identify, analyze, and report KN candidates (or other EM candidates) in time for other telescopes to perform detailed spectroscopic observations, the images from the first night of observation should be processed and analyzed within 24 hours. Meeting this deadline is a central goal of the DES-GW program, and it requires the ability to rapidly provision the required computing resources, without reliance on any single computing cluster or site.

Search program for bright gravitational wave sources




Figure 2. Overview of the DES-GW followup strategy, including communication with LIGO and other EM followup partners.

2.1. Formulation of observing plan

In general our followup observations for a trigger occur over three separate nights: the first night is as close to the initial trigger as possible, the second night is 2-3 days later (long enough to observe a decline in flux for any candidate objects), and the third night is typically 2-3 weeks after the initial trigger, by which time the light from a KN source should have faded below detection threshold. The LVC alerts contains the time of the event, the event detection algorithm, false event rate, and spatial localization map. Depending on the event detection algorithm, the maps maybe in 3 dimensions, providing distance information in each pixel. Otherwise, a global distance value is assumed. In case of merger events, the likelihood that one of the components is a neutron star (i.e. mass $< 3M_{\odot}$) may also be provided.

When an alert is received DES prepares observations in three steps: (1) production of limiting magnitude maps for the next 24 hours in time slots of about 1/2 an hour; (2) calculation of source detection probabilities for each map; (3) selection of points to observe based on the probabilities.

The calculation of a limiting magnitude proceeds from an exposure time calculator for our instrument and telescope, taking seeing, sky brightness, and exposure time as inputs. We generalize to every point on the sky that our telescope can see; this in turn depends on the time of the observation. The output is the 10σ magnitude reached in a 90 second i-band exposure, corrected for galactic extinction. This map is then multiplied by a map that is a tophat version of the telescope pointing limits: inside is 1, not possible is 0. The result is a map containing the limiting magnitude that could be achieved if we were to point the telescope there.

The calculation of the source detection probability has two components. The first models our ability to recognize a new source if we were able to detect it in the face of confusion. The second component calculates whether we could detect the source given the limiting magnitude and a source model. For the BH-BH merger events and the burst events we do not currently possess a realistic source model and we simply assume the source has an apparent magnitude of $i = 20$. For events containing a NS we use a modified Barnes and Kasen model [4] for a kilonova. We convert the distance estimate from the LIGO trigger into a distance modulus $DM = 5 \log d + 5 \log \frac{10^6}{10}$, where d is in Mpc, and the distance uncertainty into a DM uncertainty:

$$\sigma^2 = \sigma_d^2 \left(\frac{5}{d \ln 10} \right)^2 \quad (1)$$

Then the apparent magnitude of the source is $m = M + DM$ and, as gaussians can be added to form gaussians,

$$\sigma_m^2 = \sigma_M^2 + \left(\frac{5}{d \ln 10} \right)^2 \sigma^2 \quad (2)$$

The resulting gaussian is normalized $\int_0^{\infty} G(m, \sigma_m^2) dm = 1$. The probability $p(m_l)$ that the object could be seen in this pixel given a limiting magnitude of m_l is then

$$p(m) = \int_0^{m_l} G(m, \sigma_m^2) dm \quad (3)$$

If $m > m_l$, then $p(m) = \frac{1}{2} \text{erfc}(m - m_l)$. If $m < m_l$, then $p(m) = \frac{1}{2} + \frac{1}{2} \text{erf}(m - m_l)$. The $p(r)$ is weighted as uniform in volume: $\int p(r) dr = \int G(m, \sigma^2) r^2 dr$. We will need to do the same weighting in apparent magnitude space, by transforming r^2 into distance modulus

$$r = 10^{0.2(m - \bar{m}) - 5} = 10^{0.2(m - \bar{m}) + C} \quad (4)$$

$$r^2 dr = C 10^{0.6(m - \bar{m})} dm = C \exp(0.6(m - \bar{m}) \ln(10)) dm \quad (5)$$

where C is a constant. Then

$$\int_0^r Pr^2 = \int_0^r \exp\left(\frac{(m - \bar{m})^2}{s\sigma_m^2}\right) \exp(1.4(m - \bar{m})C) dm \quad (6)$$

This last equation is the probability that we will see a source given source model uncertainty and distance uncertainty for a known limiting magnitude.

One observation field is referred to as a "hex" because of the hexagonal shape of the DECam's focal plane. For a given strategy (one hex observation is one 90-second exposure in i-band for BH mergers, and a izz sequence of 90-second exposures for BNS events), we create slots of time that contain integer numbers of hexes and a duration of around a half hour. We calculate the maps and the detection probabilities for each slot. The detection probability maps are then "hexelated": the probabilities are summed inside a fixed pattern of camera pointings. For the night, the hex with the greatest probability is found and assigned to be observed in its slot. That hex is then removed from consideration (removing all probabilities for that hex at different slot times in the night), and we repeat the procedure until all slots are full. Now we have a time series of observations that maximizes the probability that we can detect the source given the reality of observing on a given night. The same time series, modified for a later time, is used during the subsequent observing nights.

2.2. Image processing software

To identify GW EM counterpart candidates within search images we use the `diffimg` [5] software, originally developed for supernova searches in DES. `diffimg` works by comparing a search image to one or more images previously taken of the same area of sky (these images are called "templates"), composing a combination of them, and subtracting the combined template from the search image. The result, known as the "difference image", should consist only of objects not present in the template images, and are thus potential candidates to be an optical counterpart of the GW trigger. The `diffimg` software, also sometimes known as the `diffimg` "pipeline", can accept as templates partially overlapping images, as well as images from DECam that were not taken on DES time (we only use such images if they have been publicly released.) Before processing each search image and its corresponding template image through the `diffimg` pipeline they all go through a pre-processing step known as single-epoch (SE) processing. The number of template images per search image is typically between 10 and 40, counting partial overlaps. On a per-CCD basis the number of overlapping template images passing quality cuts is typically less than ten.

For a typical observing cadence we expect a new image to be available for processing every three to four minutes. As soon as each new image comes in during followup observations, we run a script that calculates which DECam images overlap with the new image and can be used as templates for `diffimg` processing, checks whether SE processing is complete for that image, and then prepares a multi-stage HTCondor DAG that includes SE processing for new image and any incomplete templates in the first stage, with the `diffimg` processing of the search image in the second stage. In the first stage there is one SE job per search, and in the `diffimg` stage there is one job per CCD, for a total of 60 per search image, which run in parallel. The script also prepares a small custom tarball containing the exposure-specific `diffimg` pipeline scripts and information about what template images overlaps with the search image. The script finishes by copying the tarball to Fermilab dCache [6] for use in the corresponding grid jobs, which are then immediately submitted.

2.3. Grid processing

In order to meet the 24-hour turnaround requirements we must make extensive use of grid resources. A typical observing night will generate between 100 and 200 new images, depending

on the specifics of the observing plan and weather conditions. A typical SE processing job will take around three hours to complete, with diffing jobs running on a single CDD taking around one hour. Therefore one night’s worth of images requires between 6,000 and 12,000 CPU hours for the diffing stage, and between 300 and 600 additional hours for the SE processing stage. There could be some additional resources required if there are any template images with incomplete SE processing. It is therefore critical to use a job submission system that can access opportunistic computing resources at a wide variety of sites.

For job submission, data movement, and monitoring we use the FIFE tools developed at Fermilab, such as `jobsub` [7] for job submission, `ifdh` [8] for data movement, and FIFEMON [9] for job tracking. Using `jobsub` allows us to seamlessly submit to dedicated Fermilab computing resources, opportunistic resources on the Open Science Grid, or future dedicated DES resources such as allocation on HPC cluster or commercial clouds. Submitting to all resources at once minimizes the time it will take to complete the task, and eliminates dependency on any one resource class.

In order to be able to run jobs on multiple resource classes we must make the code base portable and have storage elements accessible at all sites. The jobs access software via a CVMFS repository [10] along with the small tarball created during DAG generation discussed above. We store the SE and diffing outputs on dedicated DES pools within the Fermilab public dCache system. Fermilab’s dCache areas are accessible to remote OSG resources and support a variety of file transfer protocols such as gridFTP and XRootD.

While we are able to take full advantage of opportunistic computing resources it is useful to know our expected processing total time if we had only opportunistic resources available (for example if dedicated Fermilab resources were unavailable on the day of a new trigger.) We can estimate the time by taking a night of DES images and passing them through the entire diffing pipe as if they were search images. We did so in early 2016 by submitting jobs for approximately 80 images, with a new image every four minutes to simulate the expected pace during real follow-up observations. During this test we submitted jobs only to opportunistic resources. We were able to quickly reach over 500 simultaneous jobs and completed over 90% of the required jobs within 10 hours of starting the test, as shown in Figure 3. With this level of resource availability we are capable of meeting of 24-hour turnaround goal using only opportunistic resources.

3. Results from first observing season

We followed up GW150914 and GW151226. We performed two analyses on GW150914 [11, 12]; one was a search for EM counterparts to the GW signal, and the other was a search for disappearing stars in the Large Magellanic Cloud, as GW events can come from a stellar core collapse with a failed supernova (which occurs in about 20% of stellar core collapses), and the stellar-rich LMC was inside the LIGO probability map. Figure 4 shows the maps for the DESGW followup observations. Neither search found any credible EM counterpart candidates. Our analysis of the second LIGO event, GW151226 [13], initially found four candidates; all were rejected as background upon further inspection. A large fraction of the diffing processing was performed on local campus resources for the second event and compared with the grid runs. These searches demonstrate the ability of the DES-GW program to quickly get on sky and perform a detailed EM counterpart searches.

4. Improvements for second observing season

The second LIGO run, known as O2, started in November 2016. We have made a number of improvements to the DES-GW processing pipeline, including more robust code, and a tool to calculate template image overlaps and run SE processing in advance of the first night of observation. We have improved the SE processing itself by developing the ability to perform on-the-fly calibration of each image. We have also worked to expand the available computing

resources by testing the diffing pipeline on commercial clouds. As part of the Fermilab HEPCloud project [14] we ran diffing processing on Amazon EC2 resources. All tests passed and the capability is ready should the DES-GW group decide that it is needed. The goal is to follow at least one BBH event. Doing so will enable us to set limits on optical emission from black hole mergers. We are also developing new tools to assist in deciding whether or not, and if so, how, to follow up on triggers. These tools help to choose a followup search strategy that takes into account both the probability area covered as well as the time it takes to execute the observational plan, by weighting the various possibilities according to how much observation




Figure 3. Count of running jobs during the DESGW OSG-only test in March 2016, grouped by OSG site. 90% of jobs completed within 10 hours of the test, with the remainder being affected by a transient database connectivity issue.




Figure 4. Sky maps for main followup of GW150914 (left), the disappearing star search in the LMC after GW150914 (middle), and the followup of GW151226 (right.) The white contours represent the 50% and 90% probability regions of the original (dashed) and updated (solid) LIGO probability maps. The yellow lines represent the fiducial DES observing region. The hexagons represent the locations of the telescope pointings during the search. In the left and right plots, the background colors indicate the DES limiting magnitude for observation, and in the center plot the background colors indicate stellar density.

time each possibility requires, how much observation time remains during the observing season, and how many more triggers the program expects to receive during the season. In this way one can select observation strategies that maximizes candidate detection probability over the entire season, not just for single events.

We have also incorporated new pieces of code for the second season which will assist in diagnostic checks of `diffimg` processing pipeline results. We calculate the background surface brightness (SB) in the coadded template image for each detection in an exposure. This SB is useful in working out problems associated with detecting events within bright galaxies. Within `diffimg`, detections are rejected and assigned flags based on the reason(s) for rejection. By gathering surface brightness and other information post-processing, we are able to assess whether these rejections are valid (and how often we may be getting them wrong.) Furthermore, the accepted detections are assigned a machine learning score corresponding to how “good” each detection is. The newly acquired background SB information can be used alongside these scores to gauge what types of backgrounds give us trouble in detection efficiency.

5. Summary

The DES-GW program is a partnership between the Dark Energy Survey and members of the astronomy community that is designed to search for electromagnetic counterparts to GW events, such as those expected from the merger of two neutron stars. DES-GW is sensitive to NS-NS mergers out to approximately 200 Mpc. 2015 and early 2016 saw a successful first observing season with optical followup of two LVC triggers. Season two started in November 2016. The DES-GW program was the most complete in terms of area covered and limiting magnitude of all EM counterpart followup programs that studied GW150914. The program has recently completed a series of improvements to the computing infrastructure for followup observation preparation and to the imaging pipeline itself. These improvements include more robust algorithms available to help inform decisions on whether to follow up on LIGO triggers. DES-GW has excellent potential for discovering an EM counterpart to a GW event, and is eagerly awaiting LIGO-Virgo triggers during the upcoming observing seasons.

Acknowledgments

This research uses services or data provided by the NOAO Science Archive. NOAO is operated by the Association of Universities for Research in Astronomy (AURA), Inc. under a cooperative agreement with the National Science Foundation. Fermilab is operated by Fermi Research Alliance, LLC under Contract number DE-AC02-07CH11359 with the United States Department of Energy.

References

- [1] B P Abbott *et al.* 2016 *Phys. Rev. Lett.* **116** 061102.
- [2] B P Abbott *et al.* 2016 *Phys. Rev. Lett.* **116** 241103.
- [3] B. Flaugher *et al.* 2015 *Astronomical J.* **150** 150.
- [4] J Barnes and D Kasen 2013 *Ap. J.* **775** 18.
- [5] R Kessler *et al.* 2015 *Astronomical J.* **150** 172.
- [6] P Fuhrmann and V Gulzow 2006 Euro-Par 2006 Parallel Processing (Springer) pp 1106-1113.
- [7] D Box 2014 *J. Phys.: Conf. Ser.* **513** 032010.
- [8] A L Lyon and M W Mengel 2014 *J. Phys.: Conf. Ser.* **513** 032068.
- [9] <https://fifemon.github.io/>
- [10] J Blomer *et al.* 2011 *J. Phys.: Conf. Ser.* **331** 042003.
- [11] M Soares-Santos *et al.* 2016 *Ap. J. Lett.* **823** L33.
- [12] J Annis *et al.* 2016 *Ap. J. Lett.* **823** L34
- [13] P Cowperthwaite *et al.* 2016 *Ap. J. Lett.* **826** L29.
- [14] G Garzoglio and O Gutsche 2015 *J. Phys.: Conf. Ser.* **664** 012001.