

Optimizing multitelescope observations of gravitational-wave counterparts

Michael W. Coughlin¹,^{1*} Sarah Antier²,² David Corre,³ Khalid Alqassimi,⁴ Shreya Anand,¹ Nelson Christensen,^{5,6} David A. Coulter,⁷ Ryan J. Foley,⁷ Nidhal Guessoum,⁴ Timothy M. Mikulski,⁶ Mouza Al Mualla,⁴ Draco Reed⁷ and Duo Tao¹

¹*Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*

²*APC, UMR 7164, 10 rue Alice Domon et Léonie Duquet, F-75205 Paris, France*

³*LAL, Univ Paris-Sud, CNRS/IN2P3, Bâtiment 200, 91440 Orsay, France*

⁴*American University of Sharjah, Physics Department, PO Box 26666, Sharjah, UAE*

⁵*Artemis, Université Côte d'Azur, Observatoire Côte d'Azur, CNRS, CS 34229, F-06304 Nice Cedex 4, France*

⁶*Carleton College, Northfield, MN 55057, USA*

⁷*Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA*

Accepted 2019 August 30. Received 2019 August 30; in original form 2019 July 17

ABSTRACT

The ever-increasing sensitivity of the network of gravitational-wave detectors has resulted in the accelerated rate of detections from compact binary coalescence systems in the third observing run of Advanced LIGO and Advanced Virgo. Not only has the event rate increased, but also the distances to which phenomena can be detected, leading to a rise in the required sky volume coverage to search for counterparts. Additionally, the improvement of the detectors has resulted in the discovery of more compact binary mergers involving neutron stars, revitalizing dedicated follow-up campaigns. While significant effort has been made by the community to optimize single telescope observations, using both synoptic and galaxy-targeting methods, less effort has been paid to coordinated observations in a network. This is becoming crucial, as the advent of gravitational-wave astronomy has garnered interest around the globe, resulting in abundant networks of telescopes available to search for counterparts. In this paper, we extend some of the techniques developed for single telescopes to a telescope network. We describe simple modifications to these algorithms and demonstrate them on existing network examples. These algorithms are implemented in the open-source software `gwemopt`, used by some follow-up teams, for ease of use by the broader community.

Key words: gravitational waves – telescopes.

1 INTRODUCTION

The science enabled by gravitational-wave astronomy is rapidly increasing, as the sensitivity of the network of gravitational-wave detectors with Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015) continues to improve. The third observing run (O3), which began in April 2019, has already yielded the detection of many binary black hole systems (Chatterjee et al. 2019a,b; Ghosh et al. 2019; Singer et al. 2019a,b; Shawhan et al. 2019) and a few with at least one neutron star (Chatterjee et al. 2019b; Singer et al. 2019b). This builds on the success of the Advanced LIGO and Advanced Virgo first and second observing

runs, which led to 10 binary black hole detections (Abbott et al. 2018a) and the detection of one binary neutron star (BNS) merger GW170817 (Abbott et al. 2017a). The BNS detection was unique in many ways, including the observation of the electromagnetic signature of the ejected matter. This includes (1) isotropic emission in the visible and near-infrared of the dynamical ejecta following the coalescence called a ‘kilonova’ (KN) counterpart, AT2017gfo (Chornock et al. 2017; Coulter et al. 2017; Cowperthwaite et al. 2017; Drout et al. 2017; Evans et al. 2017; Haggard et al. 2017; Hallinan et al. 2017; Kasliwal et al. 2017; Kilpatrick et al. 2017; Margutti et al. 2017; McCully et al. 2017; Nicholl et al. 2017; Pian et al. 2017; Shappee et al. 2017; Smartt et al. 2017; Utsumi et al. 2017); (2) the beaming emission of the relativistic ejecta with the short gamma-ray burst (SGRB), GRB170817A (Abbott et al. 2017c; Goldstein et al. 2017); and (3) the multiwavelength afterglow due

* E-mail: mcoughli@caltech.edu

to interaction of the jet with the interstellar environment (Alexander et al. 2017; Troja et al. 2017; Lyman et al. 2018; Mooley et al. 2018). This event yielded a variety of results, including a measurement of the expansion rate of the universe (Abbott et al. 2017b; Hotokezaka et al. 2019; Coughlin et al. 2019a), limits on the equation of state (EOS) of neutron stars (Bauswein, Baumgarte & Janka 2013; Abbott et al. 2017a; Bauswein et al. 2017; Coughlin et al. 2018b, 2019c,e; Radice et al. 2018), and the likely formation of heavy elements (Just et al. 2015; Wu et al. 2016; Abbott et al. 2017d; Kilpatrick et al. 2017; Roberts et al. 2017; Rosswog et al. 2017; Kasliwal et al. 2019a).

The detection of the optical counterpart AT2017gfo (Coulter et al. 2017) at a distance of 40 Mpc was helped by a three-detector gravitational-wave detection, constraining the final localization to $\approx 16 \text{ deg}^2$ on the sky (Abbott et al. 2019). With a sky localization of this size, there are a limited number of galaxies within the sensitivity volume of the gravitational-wave detectors, enabling a straightforward search method of observing each galaxy for new objects. Similarly, synoptic survey strategies by larger field-of-view telescopes (e.g. $\text{FOV} > 1 \text{ deg}^2$) were made easier by a localization of this size. While this was remarkably good luck for the astronomical community, by far most of the gravitational-wave events before and since were one and two-detector observations, yielding much larger localization regions. Furthermore, with the two LIGO detectors having more than twice the sensitivity of Virgo during O3, and with different antenna pattern distributions over the three detectors, the two-detector observations will be the most likely case for any BNS merger candidate in O3. This will affect the 1–50 BNS detections expected during O3, and will continue to be important into O4, when the number of expected detections varies between 4 and 80 per year (Abbott et al. 2018b). Note that the angle-averaged BNS range is already at 140 Mpc for LIGO Livingston (and about 120 Mpc for LIGO Hanford), whereas available catalogues such as GLADE (Dályá et al. 2018) are only complete below ~ 100 Mpc (although nearly complete at ~ 150 Mpc). Galaxy targeted follow-ups are significantly more limited in the case of binary black hole signals, which have generated interest for both potential gamma-ray (see e.g. Connaughton et al. 2016; Veres et al. 2019) and optical (see e.g. Smartt et al. 2017) counterpart searches. For example, the first BNS detection candidate of the O3, LIGO/Virgo S190425z, was a single detector event with an initial sky localization from BAYESTAR (Singer & Price 2016) spanning $\sim 10\,000 \text{ deg}^2$ at 155 ± 45 Mpc (Singer et al. 2019b) and an updated LAL Inference (Veitch et al. 2015) skymap that reduced the localization region to $\sim 7500 \text{ deg}^2$ (Singer et al. 2019c). There were more than 50 000 galaxies inside in the 90 per cent volume for this source (Cook et al. 2019).

These large localizations have motivated many synoptic survey systems to search for optical counterparts. These include the Zwicky Transient Facility (ZTF; Bellm et al. 2018; Masci et al. 2018; Dekany et al., in preparation; Graham et al. 2019), Palomar Gattini-IR (Moore & Kasliwal 2019; De et al., in preparation), the Dark Energy Camera (DECam; Flaugher et al. 2015), the Gravitational-wave Optical Transient Observer (GOTO; O’Brien 2018), the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS; Kaiser et al. 2010), the All-Sky Automated Survey for Supernovae (ASASSN; Shappee et al. 2014), the Asteroid Terrestrial-impact Last Alert System (ATLAS; Tonry et al. 2018), the Rapid Action Telescope for Transient Objects (TAROT; Klotz et al. 2008), and the MASTER global robotic network (Lipunov et al. 2010) amongst many others.

While the use of synoptic systems was already typical during the first and second observing runs, there is a growing trend of telescope

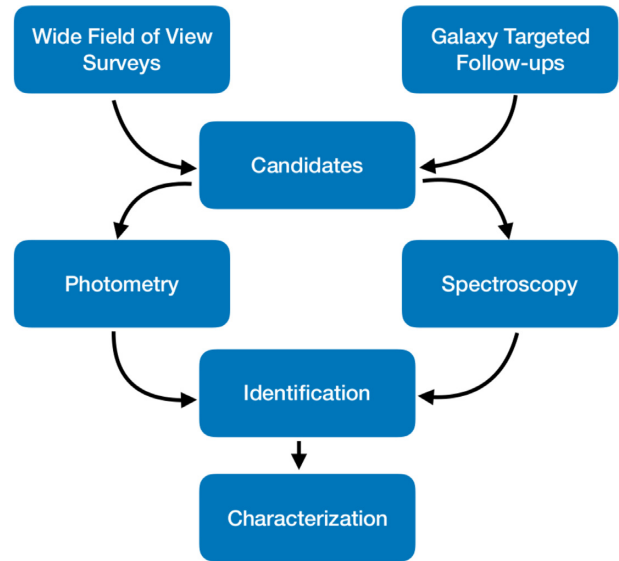


Figure 1. Flowchart of gravitational-wave electromagnetic counterpart follow-up strategy.

‘networks,’ some of which are built around these synoptic systems. These networks use various facilities to perform rapid follow-up and classification of objects (see Fig. 1). For example, ZTF, Palomar Gattini-IR, DECam, and the GROWTH-India telescope¹ (Bhalerao et al. in preparation) are scheduled by the Global Relay of Observatories Watching Transients Happen (GROWTH²) network (Coughlin et al. 2019b) [in addition to predominantly galaxy-targeted follow-up systems such as the Kitt Peak EMCCD Demonstrator (KPED) on the Kitt Peak 84 inch telescope, Coughlin et al. 2019d]. In addition, the Global Rapid Advanced Network Devoted to the Multi-messenger Addicts (GRANDMA) uses small to medium-sized telescopes spread over the entire globe, comprised of over 20 classical and robotic facilities (Antier et al., in preparation). These networks are useful for a few reasons. Due to the considerable size of the sky localizations, it is advantageous for each search to utilize telescopes capable of covering both hemispheres. In addition, coordinated observations can save precious target of opportunity time on large-aperture systems. Once the region has been imaged and candidates are identified, having worldwide coverage allows for continuous follow-up of candidates. This coverage enables identification and characterization of potential counterparts at high cadence and with multiband photometry. Having detections as early as possible is important for understanding the source mechanisms (Arcavi 2018).

There has been much recent interest in optimizing the methods to schedule observations, given the use of significant telescope time to search and follow-up electromagnetic counterparts. The gravitational-wave counterpart search effort, with the gamma-ray burst and neutrino counterpart searches closely related (Singer et al. 2013; Coughlin et al. 2019c), is unique in the community, given the significant search regions requiring coverage. Recently, an open-source codebase named gwemopt³ (Gravitational Wave – ElectroMagnetic OPTimization, Coughlin et al. 2018a) was developed, deriving concepts from the community on how to optimize

¹<https://sites.google.com/view/growthindia/>

²<http://growth.caltech.edu/>

³<https://github.com/mcoughlin/gwemopt>

optical follow-up of gravitational-wave skymaps. This includes information about how the telescopes should tile the sky, allocate available telescope time to the chosen tiles, and schedule the telescope time. We have developed generic algorithms to handle these tasks that would be useful for a wide variety of telescope set-ups; this includes telescope placement on the Earth, as well as their instrument configurations, including field of view (FOV), filters, typical exposure times, and limiting magnitudes.

While single telescope optimization remains important, it is clear that methods extending some of these methods are required for network level optimization. In this paper, we will introduce two basic extensions to the single telescope model of follow-up, which we call ‘iterative’ and ‘overlapping.’ The idea is to make straightforward extensions to the single telescope scheduling models, which have generally been shown to be robust and successful (Andreoni et al. 2019b; Coughlin et al. 2019b,c), including during the 190425z follow-up performed by GROWTH and GRANDMA (Blazek et al. 2019; De et al. 2019; Coughlin et al. 2019b; Kasliwal et al. 2019b). We note that the techniques are generic enough to be used with different scheduling algorithms, some of which we will describe below.

2 THE ITERATIVE ALGORITHM

We will briefly review the state of the art in single-telescope scheduling most relevant for a multitelescope network. One can broadly break up the process of scheduling into three categories: (1) How the telescopes should tile the sky, (2) How the telescope array should allocate time to each tile, and (3) How to schedule that time between telescopes. Ideally, all three of these would be done at once, as of course, the ability to schedule a tile should inform how much time is possible to allocate to it. In practice, `gwemopt` simply removes any tiles that are not observable during the time requested, ‘good enough’ to create sensible schedules, but still suboptimal relative to a schedule that optimizes all three simultaneously.

Despite its simplicity, this approach allows for the creation of tiles the size of the telescope’s FO with minimal overlap covering the whole sky. This is typically done using the ‘hierarchical’ and ‘greedy’ methods (Ghosh et al. 2017; Coughlin et al. 2018a); the idea of both of these tiling schemes is to decrement to zero the probability in the map enclosed in any already placed tile. In other words, each tile placed leads to a change in the skymap that is being tiled, where the locations in the map covered by that tile is set to zero. This prevents, for example, the possibility of double counting the probability contributed by a particular sky location when multiple tiles cover the same location. For some instruments, such as ZTF, the tiles are pre-determined to simplify difference imaging. ZTF in particular has both a ‘primary’ and ‘secondary’ grid, where the two grids are designed to fill in the ~ 15 per cent of the FOV that is not imaged due to gaps in between the individual CCDs. Part of the ‘job’ of the scheduling software is to optimize the use (or not) of overlapping tiles like those of ZTF. In general, taking images in the secondary grid has not been a priority for ZTF, and therefore, there are not references for all of the fields in this grid; for this reason, it will be useful to have methods to fill-in these regions with other systems.

The most important metric for any tile is the integrated spatial probability of a gravitational-wave source lying within it. This is computed by using the gravitational-wave skymaps, which report either the 2D probability $L_{GW}(\alpha, \delta)$, in right ascension α and declination δ , or 3D probability, which includes probability distributions for the luminosity distance D as a function of sky

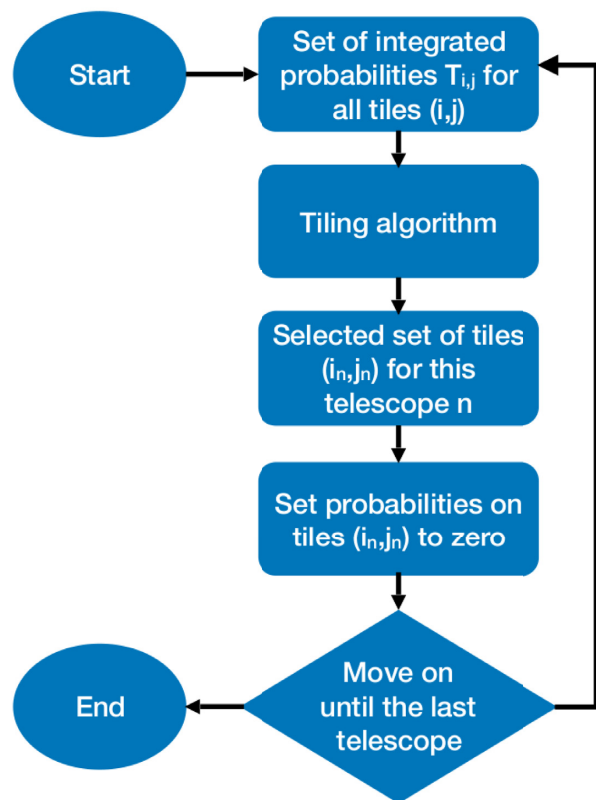


Figure 2. Flowchart of the ‘iterative’ algorithm presented in the text.

location (for explanation, see the LIGO-Virgo user guide⁴). The integrated probability in a tile is computed as a double integral over right ascension and declination:

$$T_{ij} = \int_{\alpha_i}^{\alpha_i + \Delta\alpha} \int_{\delta_i}^{\delta_i + \Delta\delta} L_{GW}(\alpha, \delta) d\Omega. \quad (1)$$

In general, a fiducial target integrated probability, usually around 90 per cent, is used to determine the number of tiles to consider for imaging and scheduling.

This brings us to our first innovation in the tiling method, where we will use the decrementing scheme from the ‘hierarchical’ and ‘greedy’ methods to expand to a multitelescope network.⁵ The first telescope in the network is scheduled as usual, yielding an optimal set of tiles for that telescope. After the first telescope is scheduled, the gravitational-wave skymap is decremented with all of the pixels covered by the first telescope’s observations set to zero. Following that, the tiling for the second telescope is computed with this modified map, and the process continues. At the end, this yields a map covered by tiles in the telescope network with minimal overlap. This algorithm is summarized in Fig. 2.

The first telescope scheduled should likely be the ‘best’ telescope in the network, i.e. the most sensitive and/or most reliable transient finder. Another option for the first telescope is the one that simply has the most localization probability observable by that site at the time of the trigger, and this option has been added as a flag to `gwemopt`. Different operators may, of course, determine what

⁴<https://emfollow.docs.ligo.org/userguide/>

⁵ ‘Hierarchical’ determines the locations of the tiles one at a time, while ‘greedy’ optimizes all locations simultaneously.

‘best’ means for them, but it is important as this telescope is likely, depending on placement on the Earth, going to be tiling the regions of highest likelihood in the skymap.

There are many considerations for what constitutes best here. Objectively, ‘etendue’ is a reasonable metric, which is the product of the aperture area of the primary mirror times the FOV covered, and so the units are degrees squared times metres squared. The Large Synoptic Survey Telescope (Ivezic et al. 2019) design is famous for optimizing around this quantity. But experience has shown that observatories whose observations strongly constrain how recently an object appeared, such as from ZTF, ATLAS, and Pan-STARRS, are incredibly important for limiting the number of objects that require follow-up. For example, while ZTF has been reporting ~ 20 objects per event, DECam follow-up has yielded an order of magnitude or more due to its lack of recent limits (Andreoni et al. 2019b; Goldstein et al. 2019). For this reason, it is not necessarily obvious that etendue is the deciding proxy for ‘best.’ For example, the telescope ordering can vary depending on a variety of metrics: (1) the significance of the event, (2) the nature of the alert, (3) the size of the gravitational-wave sky localization, (4) the distribution of the telescopes around the globe that will be used for follow-up, (5) the available filters in a given system, and (6) the delay between the trigger time and the start of observations.

In order to compare the original and iterative methods, we highlight a few examples of existing telescope networks. First of all, we include a sample ‘GROWTH’ network (Coughlin et al. 2019b), which includes ZTF, DECam, and GROWTH-India tiles. We also include the Pan-STARRS and ATLAS pair, currently scheduled and analysed by the same teams out of the University of Hawaii and Queen’s University Belfast. Finally, we include a sample of the ‘GRANDMA’ network (Antier et al., in preparation), which includes the IRIS 50 cm telescope,⁶ the Observatorio Astrofísico de Javalambre (OAJ) 80 cm,⁷ and TAROT-TCA located at Calern Observatory.⁸ The diversity of the tested networks gives a first indication of the benefits of iteration; we employ (1) a pair of telescopes with similar ~ 1 deg² FOVs in the same region with OAJ and TAROT, (2) a pair of telescopes with very different FOV with Pan-STARRS and ATLAS, located at the same observatory, and (3) a pair of telescopes with ZTF and DECam, with different FOV located at different latitudes.

We show a comparison between the original and iterative methods with S190425z sky localization area in Fig. 3. While S190425z is taken as a single example, the results can change qualitatively between, for example, skymaps with large and small sky areas, although this method is appropriate for both. To guide the reader’s eye for interpreting the plots in Fig. 3, we note that the cumulative area and probability covered are the same in both the left-hand (original) and the right-hand (iterative) panels. The black lines in each plot display the integrated probability based on the sum of the contributions of the telescopes in the network. Significant differences between the black lines and the lines from the individual telescopes indicate that the individual telescopes are imaging different areas of the sky (i.e. accumulating probability and sky coverage complementary to one another). On the left-hand side, there is usually one telescope (with the largest FOV) that tracks most closely to the overall line, while on the right-hand side, there is clear separation. The improvement is most clear at early

times, when telescopes can image different portions of the sky localization that still have significant probability. As expected, the iterative method covers both more S190425z sky localization and larger cumulative probability than independent scheduling of the individual telescopes. We show a direct comparison of the cumulative area and probability in Fig. 4.

For the GRANDMA network, the skymap coverage of S190425z is nearly doubled using the ‘iterative’ tiling from 660 to 1060 deg², with a total cumulative probability improving from ~ 0.20 to ~ 0.26 . The effect is also visible for the Pan-STARRS and ATLAS pair, even with the different FOVs. The ‘iterative’ tiling increased the total coverage by ~ 1000 deg². Indeed, the integrated probability and cumulative sky area covered is ~ 0.51 and ~ 4840 deg² with the original scheduling, while the integrated probability and cumulative sky area covered is ~ 0.55 and ~ 5735 deg² with the iterative scheduling (see Fig. 3). Note that identifying the overlapped tiles can also support multiband observation of the kilonova. Finally, the DECam-ZTF pair is naturally distributed between the north and south.

However, in the case of S190425z, the highest probability pixels were located in the sky observable by both sites. The method helps again to distribute the skymap observation more efficiently than the original tiling case. The margin of improvement of the method presented here depends strongly on the gravitational-wave localization area, but already in these three cases we see improvement. In addition, the contributions of the smallest FOV telescopes in each network is increased since the iterative method assigns yet unexplored sky coverage to any telescope in the system.

The extension of this method to the galaxy-targeting approach is straightforward. The idea of the galaxy-targeted method is to use catalogues such as GLADE (Dályá et al. 2018) or the Census of the Local Universe (CLU) catalogue (Cook et al. 2019) that list the galaxies within the gravitational-wave localization volume; the reduction in area requiring coverage makes it possible to use small FOV instruments in the counterpart searches effectively. A variety of metrics exist for determining which galaxies should be imaged (see Arcavi et al. 2017 for an example). They usually include a proxy for the location of the galaxy within the gravitational-wave localization volume and the galaxy’s mass or star formation rate. They sometimes also account for the specific telescope’s sensitivity to a given transient at the galaxy’s distance to avoid pointing at objects likely to be too far away to be detectable. In any case, for galaxy targeting, the idea is to simply remove the galaxies that have otherwise been scheduled by previous telescopes. Similar to the synoptic case, after the first telescope is scheduled, the weight associated with a given galaxy that is scheduled is set to zero. From then on, future telescopes in the network will no longer attempt to schedule that galaxy given its weight, and instead will schedule others. As an example, we use this method to schedule 11 telescopes of the GRANDMA network, each with a FOV < 1 deg², on the S190426c sky localization, a neutron star-black hole candidate with a 90 per cent credible region covering 1260 deg² at a distance of 375 ± 108 Mpc (Chatterjee et al. 2019b). Fig. 5 shows that the improvement in both the number of galaxies and the metric. More concretely, the total number of galaxies imaged improves from 1303 to 1929. For the galaxy targeted portion, the total cumulative adopted metric covered improves from 85 per cent to 99 per cent, where we have only included galaxies inside of the 90 per cent contour when computing this number.

One of the major downsides to this schema overall is the possibility of weather-related failures, and the loss in opportunity of imaging high likelihood regions of the sky. This leads to the

⁶<http://iris.lam.fr/>

⁷<https://oajweb.cefca.es/telescopes/jast-t80>

⁸<http://tarot.obs-hp.fr/>




Figure 3. Optimized coverage of S190425z. The top row shows the optimization for the ‘GROWTH’ network, which includes ZTF, DECam, and GROWTH-India tiles. The left shows the tiles drawn using the original scheduling algorithm, while the right is the same for the iterative method discussed in the text. The middle row shows the same for the Pan-STARRS and ATLAS pair. The bottom row shows the same for some telescopes of the ‘GRANDMA’ network, with IRIS, OAJ, and TAROT-Calern. The dashed lines indicate the sky area covered, while the solid line indicates the integrated probability covered of the 2D skymap.




Figure 4. Optimization comparison for the GROWTH, Pan-STARRS and ATLAS, and GRANDMA networks. The dashed lines indicate the sky area covered, while the solid line indicates the integrated probability covered of the 2D skymap. The black lines correspond to the original scheduling algorithm while the blue lines correspond to the iterative method. The yellow band is the approximate ‘daytime’ for the network, where no observations are taken for this particular skymap.




Figure 5. Optimized coverage of S190426c using the GRANDMA ‘galaxy targeting’ Network. On the left is the original algorithm where the telescopes are scheduled separately, and on the right, where they are scheduled iteratively. The total cumulative metric covered improves from 85 percent to 99 percent, and the total number of galaxies imaged improves from 1303 to 1929.

idea of ‘golden tiles,’ which are regions of the sky that are not decremented at each step. Because the inner portion of the sky localization region (say the inner 50 percent) tends to be much smaller than the outer portion of the sky localization region (say the outer 90 percent) (Singer et al. 2014), the opportunity cost of imaging the higher likelihood region is lower than the outer region. For this reason, we implemented a user-definable inner percentage of the skymap that is not incremented at each step. Based on the dimension of the telescope network, and the distribution of the telescopes over the globe, one idea would be to create coverage in groups of telescopes at similar longitudes. In each group, the optimization of the scheduling for systems at different latitudes would be performed with the method presented here. For example, if scheduled in four groups, it would allow for a maximum of a 6 h gap of observation of the same target in the sky. It seems most appropriate to have some redundancy with the golden tiles, but

with different filters, to create a balance between the possibility of kilonova detection, which requires maximum coverage, and time evolution of the kilonova in multiple filters. The percentage of the golden tiles will then be a combination of time allocation for each telescope to follow-up gravitational alerts and the size of the sky localization area. In this sense, the threshold can be based on percentage of the sky that can be covered.

Once the preferred set of tiles has been scheduled, how to allocate exposure times is the next question for the scheduling algorithm to address. We will not discuss this issue much herein, as in general this decision is very telescope and observer dependent, and has been discussed extensively elsewhere (see Coughlin et al. 2018a for a summary). In general, one popular proposed technique is to allocate exposure times proportional to the probability enclosed in a tile (see e.g. Coughlin & Stubbs 2016 or Ghosh et al. 2017). It is also possible, with the information in three-dimensional sky

localization, to image deep enough in each field to reach a fixed absolute magnitude. In general, this would mean taking longer exposures in the centre of the localization and shorter exposures at the edge of the skymap. In practice, this makes it much more complicated to compare objects found in the search. The pre-merger limits, especially those derived from a survey that uses a fixed exposure time, will generally be similar across the localization. Changing exposure times within a search will lead to significant differences in how objects are determined for follow-up. Given the desired redundancy described previously, it may also be appropriate to choose a uniform maximum magnitude based on the size and distance estimate of the localization, and then adjust the exposure time for each telescope to achieve that specific depth. For example, for a given localization, if ZTF could use 90 s exposures to cover the desired area, achieving a depth of ~ 21 , then GROWTH-India might also adjust its exposure time to achieve that same depth.

Once time has been allocated to each telescope’s tile, the next task is to schedule the observations. In most cases, there is insufficient time during the night to perform all observations, and therefore the job of the scheduling software is to optimize the subset actually taken. In general, the optimization schemes employed in *gwemopt* (Coughlin et al. 2018a) weight each tile by a combination of its integrated probability (or metric in case of the galaxy targeting) and its current and future visibility, so as to schedule as many tiles as possible while maximizing the probability covered. As part of this, the scheduling software must also account for practical constraints such as altitude and moon-sky brightness limits. There are also more telescope-specific issues such as slew rate, hour angle constraints, and camera readout times. One of the most important aspects to the scheduling is the cadence at which a telescope returns to a particular field or galaxy, and the filters that are chosen when imaging. For instance, ZTF has mainly been imaging in *g-r-g* band exposure blocks to ensure a measurement of both colour and potential change in luminosity in a single band (Coughlin et al. 2019c). The observations are separated in time not only to measure a luminosity change but also to reject asteroids and other transient objects. The scheduling algorithms used to address these issues are examined thoroughly in Coughlin et al. (2018a), and we refer the reader there for further information.

3 THE OVERLAPPING ALGORITHM

While it is convenient for each telescope to be its own follow-up resource, given the sky localizations involved and the likelihood of having tiles set at certain longitudes, it is useful to potentially image a location on the sky with multiple systems. The problem is that independent scheduling will lead to the highest probability region being imaged around the same time (at least in the case where the sites are at similar locations). In other cases, especially when the localization area is very large or the number of telescopes in the network small, it makes sense to temporally separate observations of the same field by different telescopes. For example, a patch imaged by DECam might be visible 6 h later with ZTF, so it is not necessary to wait until the second night of DECam observations to get a second epoch if ZTF also images that location.

To enable this, we once again modify the existing scheduling algorithms in a simple way. Similar to the ‘iterative’ method, we allow the first telescope to schedule its observations as before. For the next telescope, we specify a minimum difference in time (or time delay) between observations in a given field. Algorithmically, we impose this observability constraint on a given field, just like is done for moon proximity or airmass, such that it will not be




Figure 6. Cumulative histogram of the difference between telescope observations of the same patch of the sky for S190425z. We plot the original algorithm in solid and overlapping algorithm in dashed. The inset shows the original histogram. The reader should note the lack of observations within the same hour in the ‘overlapping’ traces, shown in dashed for the various networks.

scheduled within the specified time delay of a previous observation. In the example that follows, we choose one hour, more than sufficient to differentiate between real transients and asteroids, as well as potentially measuring a change in luminosity. This is appropriate, for example, for a kilonova, which is expected to show a rapid evolution in magnitude (Metzger 2017); GW170817 faded $\Delta r \sim 1$ mag per day over the first 3 d post-explosion. If the goal is to simply measure a change in luminosity over these time-scales, as opposed to discard asteroids as potential candidates, something like a 6 h delay may be more appropriate.

The difference between the normal scheduling and the ‘overlapping’ scheduling is shown in Fig. 6. We show both the histograms and the cumulative version. We would like the reader to note the lack of observations within the first hour in the ‘overlapping’ schedule case, as is expected. We also note that the number of observations in overlapping fields rapidly ‘catches up’ after that first hour, since the software will optimize around scheduling the highest probability fields once the 1 h constraint has been lifted. In this way, the algorithm is successful at not only preventing overlapping observations of the same field within the time-frame specified, significantly limiting the false positives that arise from asteroids, but also does not prevent that field from being observed by that system at a later time because, for example, that particular field had set.

4 CONCLUSION

In this paper, we have described straightforward strategies for a multitelescope network. We have shown how the introduction of two new tiling and scheduling schemes make it possible to use existing single-telescope strategies in a network capacity. This work builds upon examples of optimization in this regard done ‘by hand’ previously (Waratkar et al. 2019). Algorithms of this type open up the possibility of coordinated observations between telescopes for gravitational-wave follow-up. This also brings to bear a variety of the scheduling techniques that have been developed. For example, the slew-optimized scheduling algorithms that have recently developed (Rana, Anand & Bose 2019) are likely more important in a network where the tiles scheduled are likely more spatially separated. It also opens the possibility of prioritizing

different schedules with different telescopes. For example, one telescope might use an airmass-based optimization to maximize the sensitivity given a fixed exposure time, at the cost of not imaging the highest probability tiles as early as it might otherwise. This might be more palatable when another system is using a basic greedy algorithm to image the highest probability tiles as early as possible.

Longer term, we wish to design optimizations that will vary the number of fields accounting for scheduling constraints, instead of separating the steps of choosing fields, allocating time to them, and then scheduling. We also want to define the metrics generically enough that having more than one instrument is like having a more sensitive version of a single instrument, which is difficult, given the constraints that the telescopes are scattered around the Earth, with varying FOVs and sensitivities, etc. It should also be the case that the choices a human would intuitively make, such as using different telescopes to cover disparate parts of the localization when they are significantly separated on the sky to save slew time, are naturally accounted for in these metrics. It should be the case that the use of multiple telescopes correctly should be able to reduce slew time by minimizing the size of their patches. In addition, not relying on an ordered list of telescopes should result in even better sky coverage. There is also the open question of how a network should be optimally used, when accounting for the available time for target of opportunity observations. For example, in our analysis, we have assumed that taking complete control of each of these systems for the night following the event is appropriate. In practice, the most desirable systems should not/cannot be used on all events, given that their time is limited. Determining criteria for their optimal use should be a focus of future work.

As time goes on, the detection of kilonovae should become more frequent; this will include searching short gamma-ray burst counterparts for kilonova signatures (Guessoum, Zitouni & Mochkovitch 2018; Ascenzi et al. 2019). Thus, the future perspective will be to adopt scheduling designed for studying the physical mechanisms at stake, not only for detecting the kilonova (see e.g. Andreoni et al. 2019a). Then, instead of focusing on such things as ensuring we image an object twice in a given amount of time, we can instead prioritize metrics that are kilonova science targeted. This is where the ‘overlapping’ scheme may become important, as using multiple telescopes can turn a simple detection into a discovery by measuring rapid changes in colour and/or luminosity that may be difficult on a single system. For example, it might be useful to do a second or third round of imaging instead of exploring more of the probability volume; this may lower the odds of detecting a kilonova, but would increase the science output if it is present. The set of filters to use can be optimized depending on whether the adopted strategy prioritizes colour or luminosity variations. In addition, there are science cases for the detection as early as possible, and even for ‘non-detections’ in the early photometry (Arcavi 2018).

In conclusion, in the open source software *gwemopt* (Coughlin et al. 2018a), we have implemented a first optimization of a network level follow-up of gravitational-wave alerts, showing the substantial gains that can potentially be made by coordinated scheduling of existing telescope networks.

ACKNOWLEDGEMENTS

M. W. Coughlin is supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. Sarah Antier is supported by the CNES Postdoctoral Fellowship at Laboratoire Astroparticule et Cosmologie. David Corre is supported by

a CNRS Postdoctoral Fellowship at Laboratoire de l’Accélérateur Linéaire. N. Christensen and T. Mikulski acknowledge support from the National Science Foundation with grant number PHY-1806990, and T. Mikulski also acknowledges support from the Towsley fund at Carleton College. The UCSC (D. Coulter, R. Foley, D. Reed) team is supported in part by NASA grant NNG17PX03C, NSF grants AST-1518052 and AST-1911206, the Gordon & Betty Moore Foundation, the Heising-Simons Foundation, and by a fellowship from the David and Lucile Packard Foundation to R.J.F.

REFERENCES

- Aasi et al., 2015, *Class. Quantum Gravity*, 32, 074001
 Abbott B. P. et al., 2017a, *Phys. Rev. Lett.*, 119, 161101
 Abbott B. P. et al., 2017b, *Nature*, 551, 85
 Abbott B. P. et al., 2017c, *ApJ*, 848, L13
 Abbott B. P. et al., 2017d, *ApJ*, 850, L39
 Abbott B. P. et al., 2018a, *Phys. Rev. X*, 9, 031040, preprint (arXiv:1811.12907)
 Abbott B. P. et al., 2018b, *Living Rev. Relat.*, 21, 3
 Abbott B. P. et al., 2019, *Phys. Rev. X*, 9, 011001
 Acernese et al., 2015, *Class. Quantum Gravity*, 32, 024001
 Alexander K. D. et al., 2017, *ApJ*, 848, L21
 Andreoni I. et al., 2019a, *PASP*, 131, 068004
 Andreoni I. et al., 2019b, *ApJ*, 881, L16
 Arcavi I., 2018, *ApJ*, 855, L23
 Arcavi I. et al., 2017, *ApJ*, 848, L33
 Ascenzi S. et al., 2019, *MNRAS*, 486, 672
 Bauswein A., Baumgarte T. W., Janka H.-T., 2013, *Phys. Rev. Lett.*, 111, 131101
 Bauswein et al., 2017, *ApJ*, 850, L34
 Bellm E. C. et al., 2018, *PASP*, 131, 018002
 Blazek M. et al., 2019, GRB Coordinates Network, Circular Service, No. 24227, #1 (2019/April-0), 24227
 Chatterjee et al., 2019a, GRB Coordinates Network, Circular Service, No. 24141, #1 (2019/April-0), 24141
 Chatterjee et al., 2019b, GRB Coordinates Network, Circular Service, No. 24237, #1 (2019/April-0), 24237
 Chornock et al., 2017, *ApJ*, 848, L19
 Connaughton V. et al., 2016, *ApJ*, 826, L6
 Cook D. O. et al., 2019, *ApJ*, 880, 7
 Cook D. O. et al., 2019, GRB Coordinates Network, Circular Service, No. 24232, #1 (2019/April-0), 24232
 Coughlin M., Stubbs C., 2016, *Experimental Astronomy*, Springer, Netherlands, p. 1
 Coughlin M. W. et al., 2018a, *MNRAS*, 478, 692
 Coughlin M. W. et al., 2018b, *MNRAS*, 480, 3871
 Coughlin M. W., Dietrich T., Heinzel J., Khetan N., Antier S., Christensen N., Coulter D. A., Foley R. J., 2019a, preprint (arXiv:1908.00889)
 Coughlin M. W. et al., 2019b, preprint (arXiv:1907.12645)
 Coughlin M. W. et al., 2019c, *PASP*, 131, 048001
 Coughlin M. W. et al., 2019d, *MNRAS*, 485, 1412
 Coughlin M. W., Dietrich T., Margalit B., Metzger B. D., 2019e, *MNRAS*, 489, L91
 Coulter D. A. et al., 2017, *Science*, 358, 1556
 Cowperthwaite P. S. et al., 2017, *ApJ*, 848, L17
 D’Áya G. et al., 2018, *MNRAS*, 479, 2374
 De K. et al., 2019, GRB Coordinates Network, Circular Service, No. 24187, #1 (2019/April-0), 24187
 Drout M. R. et al., 2017, *Science*, 358, 1570
 Evans P. A. et al., 2017, *Science*, 358, 1565
 Flaugher B. et al., 2015, *AJ*, 150, 150
 Ghosh S., Chatterjee D., Kaplan D. L., Brady P. R., Sistine A. V., 2017, *PASP*, 129, 114503
 Ghosh et al., 2019, GRB Coordinates Network, Circular Service, No. 24377, #1 (2019/May-0), 24377
 Goldstein A. et al., 2017, *ApJ*, 848, L14

- Goldstein D. A. et al., 2019, *ApJ*, 881, L7
- Graham M. J. et al., 2019, *PASP*, 131, 078001
- Guessoum N., Zitouni H., Mochkovitch R., 2018, *A&A*, 620, A131
- Haggard D., Nynka M., Ruan J. J., Kalogera V., Cenko S. B., Evans P., Kennea J. A., 2017, *ApJ*, 848, L25
- Hallinan G. et al., 2017, *Science*, 358, 1579
- Hotokezaka K., Nakar E., Gottlieb O., Nissanke S., Masuda K., Hallinan G., Mooley K. P., Deller A., 2019, *Nature Astron.*
- Ivezic Z., Tyson J. A., Allsmann R., Andrew J., Angel R., 2019, *ApJ*, 873, 111
- Just O., Bauswein A., Pulpillo R. A., Goriely S., Janka H.-T., 2015, *MNRAS*, 448, 541
- Kaiser N. et al., 2010, Proc. SPIE, 7733, 77330E
- Kasliwal et al., 2017, *Science*, 358, 1559
- Kasliwal M. M. et al., 2019a, *MNRAS*, preprint ([arXiv:1812.08708](https://arxiv.org/abs/1812.08708))
- Kasliwal, et al., 2019b, GCN, 24191, 1
- Kilpatrick et al., 2017, *Science*, 358, 1583
- Klotz A., Boër M., Eysseric J., Damerdjy Y., Laas-Bourez M., Pollas C., Vachier F., 2008, *PASP*, 120, 1298
- Lipunov V. et al., 2010, *Adv. Astron.*, 2010, 349171
- Lyman J. D. et al., 2018, *Nature Astron.*, 2, 751
- Margutti R. et al., 2017, *ApJ*, 848, L20
- Masci F. J. et al., 2018, *PASP*, 131, 018003
- McCully C. et al., 2017, *ApJ*, 848, L32
- Metzger B. D., 2017, *Living Rev. Relat.*, 20, 3
- Mooley K. P. et al., 2018, *Nature*, 561, 355
- Moore A. M., Kasliwal M. M., 2019, *Nature Astron.*, 3, 109
- Nicholl et al., 2017, *ApJ*, 848, L18
- O'Brien P., 2018, 42nd COSPAR Scientific Assembly, Pasadena, CA., p. E1.15
- Pian E. et al., 2017, *Nature*, 551, 67
- Radice D., Perego A., Zappa F., Bernuzzi S., 2018, *ApJ*, 852, L29
- Rana J., Anand S., Bose S., 2019, *ApJ*, 876, 104
- Roberts L. F. et al., 2017, *MNRAS*, 464, 3907
- Rosswog S., Feindt U., Korobkin O., Wu M.-R., Sollerman J., Goobar A., Martínez-Pinedo G. 2017, *Class. Quantum Gravity*, 34, 104001
- Shappee B. J. et al., 2014, *ApJ*, 788, 48
- Shappee B. J. et al., 2017, *Science*, 358, 1574
- Shawhan et al., 2019, GRB Coordinates Network, Circular Service, No. 24098, #1 (2019/April-0), 24098
- Singer L. P., Price L. R., 2016, *Phys. Rev. D*, 93, 024013
- Singer et al., 2013, *ApJ*, 776, L34
- Singer L. P. et al., 2014, *ApJ*, 795, 105
- Singer L. P. et al., 2019a, GRB Coordinates Network, Circular Service, No. 24069, #1 (2019/April-0), 24069
- Singer et al., 2019b, GRB Coordinates Network, Circular Service, No. 24168, #1 (2019/April-0), 24168
- Singer et al., 2019c, GRB Coordinates Network, Circular Service, No. 24228, #1 (2019/May-0), 24228
- Smartt et al., 2017, *Nature*, 551, 75 EP
- Tonry J. L. et al., 2018, *PASP*, 130, 064505
- Troja E. et al., 2017, *Nature*, 551, 71
- Utsumi Y. et al., 2017, *PASJ*, 69, 101
- Veitch J. et al., 2015, *Phys. Rev. D*, 91, 042003
- Veres P., dal Canton T., Burns E., Goldstein A., Littenberg T. B., Christensen N., Preece D. R., 2019, *ApJ*, 882, 53
- Waratkar G., Kumar H., Bhalerao V., Karambelkar V., Anupama G. C., Stanzin J., 2019, GRB Coordinates Network, Circular Service, No. 24316, #1 (2019/May-0), 24316
- Wu M.-R., Fernández R., Martínez-Pinedo G., Metzger B. D., 2016, *MNRAS*, 463, 2323

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.