Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **495,** 4366–4371 (2020)
Advance Access publication 2020 June 1

doi:10.1093/mnras/staa1498

Dynamic scheduling: target of opportunity observations of gravitational wave events

Mouza Almualla,[1]★ Michael W. Coughlin ,[2,3]★ Shreya Anand,[3] Khalid Alqassimi,[1] Nidhal Guessoum[1] and Leo P. Singer[4,5]

[1]*Physics Department, American University of Sharjah, PO Box 26666, Sharjah, UAE*
[2]*School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA*
[3]*Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*
[4]*Astrophysics Science Division, NASA Goddard Space Flight Center, MC 661, Greenbelt, MD 20771, USA*
[5]*Joint Space-Science Institute, University of Maryland, College Park, MD 20742, USA*

Accepted 2020 May 27. Received 2020 May 27; in original form 2020 March 21

ABSTRACT

The simultaneous detection of electromagnetic and gravitational waves from the coalescence of two neutron stars (GW170817 and GRB170817A) has ushered in a new era of 'multimessenger' astronomy, with electromagnetic detections spanning from gamma to radio. This great opportunity for new scientific investigations raises the issue of how the available multimessenger tools can best be integrated to constitute a powerful method to study the transient Universe in particular. To facilitate the classification of possible optical counterparts to gravitational wave events, it is important to optimize the scheduling of observations and the filtering of transients, both key elements of the follow-up process. In this work, we describe the existing workflow whereby telescope networks such as GRANDMA and GROWTH are currently scheduled; we then present modifications we have developed for the scheduling process specifically, so as to face the relevant challenges that have appeared during the latest observing run of Advanced LIGO and Advanced Virgo. We address issues with scheduling more than one epoch for multiple fields within a skymap, especially for large and disjointed localizations. This is done in two ways: by optimizing the maximum number of fields that can be scheduled and by splitting up the lobes within the skymap by right ascension to be scheduled individually. In addition, we implement the ability to take previously observed fields into consideration when rescheduling. We show the improvements that these modifications produce in making the search for optical counterparts more efficient, and we point to areas needing further improvement.

Key words: gravitational waves – telescopes; methods: observational.

1 INTRODUCTION

The first and second observing runs of the global network of gravitational wave (GW) interferometers, comprising the Advanced Virgo (Acernese et al 2015) and the twin Advanced LIGO (Aasi et al 2015) detectors, yielded the detection of a total of 10 binary black hole mergers and 1 binary neutron star (BNS) coalescence (Abbott et al. 2019). Most recently, the improved sensitivity of the instruments during the third observing run (O3) has resulted in 56 GW candidates – many of which have been classified as BNS or neutron star–black hole (NSBH) mergers (information about these candidates can be found in the Gravitational-wave Candidate Event Database, or GraceDB[1] (Abbott et al. 2020a)).

Due to the association of BNS and NSBH mergers with potentially detectable electromagnetic counterparts (Metzger & Berger 2012; Nakar 2019), substantial efforts have been invested into optimizing follow-up observations of such candidates (e.g. Coughlin et al. 2019e; Goldstein et al. 2019; Gomez et al. 2019; Andreoni et al. 2020). These counterparts may come in the form of short gamma-ray bursts (sGRBs) accompanied by optical and near-infrared transients ('kilonovae', or KNe) powered by the decay of *r*-process nuclei that are synthesized in the merger ejecta, as well as prolonged radio emission resulting from the interaction of the sub-relativistic ejecta with the surrounding medium (e.g. Li & Paczynski 1998; Nakar & Piran 2011; Metzger & Berger 2012; Piran, Nakar & Rosswog

★ E-mail: mouzaalmuallaa@gmail.com (MA); mcoughli@caltech.edu (MWC)

[1]https://gracedb.ligo.org/

2013; Tanaka 2016; Goldstein et al. 2017; Savchenko et al. 2017; Guessoum, Zitouni & Mochkovitch 2018; Metzger 2019).

The culmination of these follow-up efforts came to fruition on the 2017 August 17, unveiling the new era of multimessenger astronomy with the detection of GW170817 along with the sGRB GRB170817A (Abbott et al. 2017d), and multiple independent discoveries of the optical transient counterpart AT 2017gfo in NGC 4993 ($D \sim 40$ Mpc) by various teams (see Abbott et al. 2017c and references therein), the first of which was announced by Coulter et al. (2017). The three Advanced LIGO and Virgo instruments had detected a signal that was determined to have likely originated from a BNS coalescence; the source was well constrained, initially localized to $\sim$31 deg^2 at the 90 per cent credibility level and with luminosity distance 40 ± 8 Mpc (LIGO Scientific Collaboration & Virgo Collaboration 2017; Abbott et al. 2017a). The unprecedented nature of these detections has since led to such scientific gains as the ability to probe into the workings of *r*-process nucleosynthesis in KNe (e.g. Chornock et al. 2017; Coulter et al. 2017; Cowperthwaite et al. 2017; Pian et al. 2017; Smartt et al. 2017; Kasliwal et al. 2019a; Watson et al. 2019) and the expansion rate of the Universe (Abbott et al. 2017b; Hotokezaka et al. 2019; Coughlin et al. 2020), as well as constrain properties of neutron stars such as mass, radius, and tidal deformability in novel ways (e.g. Bauswein et al. 2017; Margalit & Metzger 2017; Abbott et al. 2018; Annala et al. 2018; Coughlin et al. 2018b, 2019c, d; Most et al. 2018; Radice et al. 2018; Lai, Zhou & Xu 2019).

The majority of candidate BNS and NSBH mergers during O3 had localization areas leaning towards the thousands of deg^2, with the exception of probable NSBH merger S190814bv (which has an updated 90 per cent credible region of 23 deg^2 (LIGO Scientific Collaboration & Virgo Collaboration 2019b); these numbers are in stark contrast with the aforementioned localization area of GW170817. The sky localization is expected to improve in subsequent observing runs, thanks to the addition of other detectors to the GW network (Abbott et al. 2020a); however, there will still likely be events with localization areas of the order of hundreds of deg^2, which is much larger than the field of view (FOV) of most electromagnetic facilities, and so there will continue to be challenges in obtaining significant coverage of the skymap in the future. It is thus important to optimize our methods in performing follow-ups to GW triggers, which will greatly increase the odds of detecting an electromagnetic counterpart.

A codebase named `gwemopt`[2] (Gravitational-Wave ElectroMagnetic OPTimization) was hence developed (Coughlin et al. 2018a), aimed at optimizing the scheduling of Target of Opportunity (ToO) telescope observations immediately after a GW detection. This code breaks the process down into three parts: tiling, time allocation, and scheduling. During the tiling step, it takes the HEALPix GW skymap and splits it up into 'tiles' according to the FOV characteristics of the given telescope. It then goes on to allocate time to the tiles that are available for observation, which is dependent on the algorithm that is utilized for the plan. `gwemopt` finally proceeds to schedule these observations, taking into account factors such as the probability associated with the tiles, slew time, and observability. One way to further optimize the follow-up process is through the implementation of network-level telescope observations during scheduling (this is discussed in depth, for example, in Coughlin et al. 2019b), in which various telescopes around the world work together to achieve maximum coverage of the localization area for a given event. This is an especially relevant issue in the case of ToO observations, as multitelescope observations will improve our ability to cover areas in the localization that may not be accessible to one given telescope (e.g. the localization area could extend into different hemispheres); in addition, this will allow different telescopes to coordinate in imaging the same patch of the sky in different filters and perform independent visits separated in time.

In this paper, we delineate the new additions to `gwemopt` that build upon these ideas and expand on the currently available features. These features will facilitate the scheduling process in the case of both multi- and single-telescope observations. In Section 2, we discuss the novel ability for `gwemopt` to take into account previously completed observations when rescheduling, and in Section 3, we describe two features that drastically improve multi-epoch coverage of events. In Section 4, we conclude by discussing the role that these features play in the broader context of large telescope networks.

2 THE SUPERSCHEDULER ALGORITHM

Although various factors such as observability and telescope location are taken into account during the scheduling process, light pollution, bad weather conditions, and unanticipated telescope-related failures may often lead to unsuccessful attempts at observation. When scheduling or rescheduling these observations, `gwemopt` does not have any information as to whether a given tile has already been observed or not. This limitation poses some problems, since there is a possibility that the `gwemopt` pipeline will schedule tiles that were already observed rather than prioritizing unobserved tiles and increasing coverage of the localization.

This is an especially important point to consider in the case of multitelescope observations, as there should be a way to schedule different telescopes and take previous observation rounds into account. The SuperScheduler can do this by going through a given number of iterations of the scheduling process, with each iteration corresponding to an observation round. The algorithm is able to take previous rounds into account when rescheduling by reading in information about which tiles have or have not been observed; it then sets the 2D spatial probability of the GW skymap enclosed in the observed tiles to 0 before the next round is scheduled.

This step improves the efficiency of the scheduling process since `gwemopt` no longer redundantly schedules the same tiles for reobservation. The algorithm can work for multiple telescopes in each round, and the telescopes can also be changed between different iterations. In cases where observations in more than one filter are scheduled, the SuperScheduler also takes the filter in which the field was observed into account. So, if a given field has only been observed in the *g* band, for example, it can still schedule a second exposure in the *r* band the next time around rather than completely ignoring the field.

In order to test the capabilities of the SuperScheduler algorithm, we performed a simulation using the BNS merger candidate S190426c (LIGO Scientific Collaboration & Virgo Collaboration 2019a) in which a 50 per cent failure rate (exaggerated for purposes of demonstration) was assumed for the attempted observations. The performance between two different cases was compared; in the first case, the whole night was scheduled normally, and in the second case, the night was broken down into two blocks, and failed observations in the first block were taken into account when scheduling the second block. As shown on the left-hand side of Fig. 1, the normal scheduling algorithm does not take into account whether certain tiles have had successful (shown in green) or unsuccessful (shown in red) observation attempts throughout the scheduling process. As a result, there are no previously unobserved

[2] https://github.com/mcoughlin/gwemopt

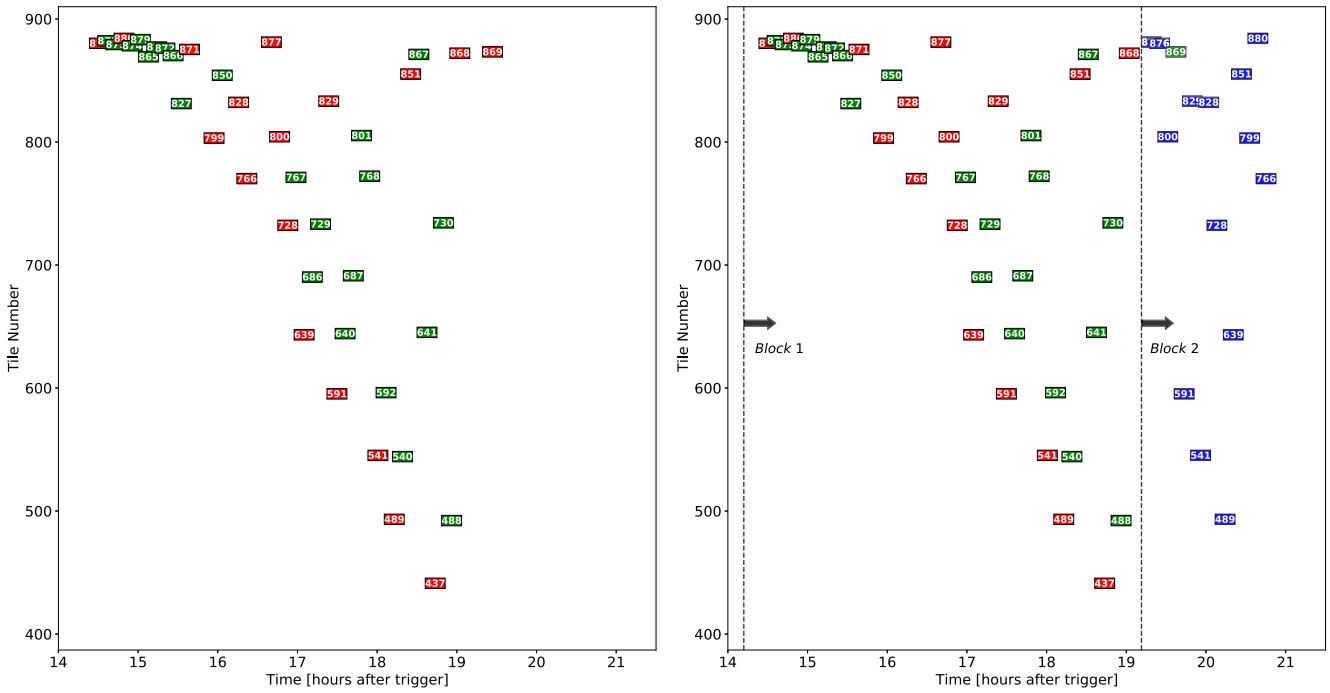

Figure 1. Plots of coverage for S190426c (LIGO Scientific Collaboration & Virgo Collaboration 2019a) without (left) and with (right) the use of the SuperScheduler algorithm. Red indicates that the corresponding field could not be observed during that respective round, and green indicates that the observation was successful. For simplicity, all fields were scheduled in the same filter for both of the cases shown. Breaking the night up into two blocks and using the SuperScheduler, 14 previously failed attempts at observation were successfully rescheduled (shown in blue). The rectangle widths (representing exposure time) have been scaled by a factor of 3 for visualization purposes, and the respective tile number is labelled at the centre of each rectangle.

tiles scheduled for observation (which would be shown in blue). Conversely, the results using the SuperScheduler algorithm on the right-hand side of Fig. 1 show that, breaking the night down into two blocks, prioritizing tiles that were not successfully observed in the first block (as indicated in the figure) led to almost all of the failed observations being rescheduled in the second block. We also note that many more fields are scheduled past the starting point of the second block (around 19 h after trigger) when using the SuperScheduler; this is because the algorithm allows the scheduler to revisit the fields that failed to be observed in the first round. This ability is not available when using the normal scheduling algorithm, and most of the localization had already been covered by that point in time, thus leading to very few additional fields scheduled past that point in comparison. Evidently, incorporating information about previous observations leads to more efficient scheduling that optimizes coverage over the course of multiple observation rounds.

3 FILTER BALANCING

If observations in multiple filters are required, `gwemopt` has the ability to implement a block-completion algorithm during the scheduling process. This means that it schedules observations in only the first filter (i.e. the first block), and then if there is time left, schedules a second pass in the next filter, and so on. This strategy minimizes the number of filter changes, which is especially advantageous since changing filters compromises observation time; the Zwicky Transient Facility (ZTF), for example, takes ∼100 s to change filters with slew time taken into account (Bellm et al. 2018).

The implementation of the block-completion algorithm may, however, lead to some challenges in scheduling observations in all requested filters for a given field. Since observations are scheduled in the second filter only after the first filter block has been completed, there will likely be a disproportionately larger number of observations in just the first filter. This issue is pertinent to the case of ToO follow-up to GW events, as strategies for the discovery of KN counterparts (Andreoni et al. 2019a) require observations in all requested filters to be satisfied (hence the term 'filter balancing'). This is because the characteristic rapid fading and reddening of KNe, as was seen with GW170817, can be used to identify candidates by acquiring images in at least two different filters (Arcavi et al. 2017; Pian et al. 2017; Smartt et al. 2017). The *g*–*i* pair in particular has been shown to be most suitable in achieving this task since more KNe are expected to be detected in the *i* filter relative to the others, and the combination also displays the largest colour change (only second to the *g*–*z* pair) over the days following the detection (Andreoni et al. 2019a).

It is important to promptly process images during the transient-filtering stage so that we can narrow down the hundreds of thousands of sources of variability to a select few candidates; high-performance image subtraction pipelines have been developed for this purpose (e.g. Kessler et al. 2015; Goldstein et al. 2019). In order to rule out moving objects such as near-Earth asteroids, the candidate must have a minimum of two detections separated by at least 30 min (Bellm et al. 2019). It is justified, then, to place emphasis on scheduling at least two epochs during block scheduling.[3]

[3]The `--doBalanceExposure` and `--doRASlices` command-line options in `gwemopt` seek to ensure this.

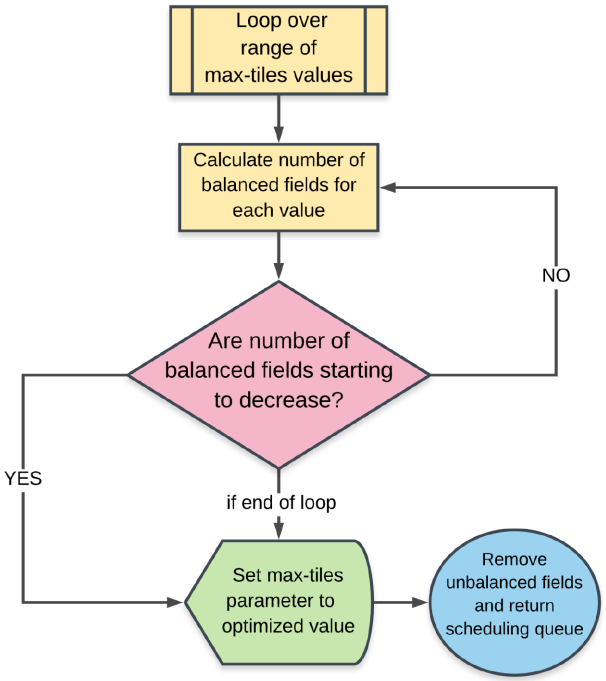

Figure 2. Step-by-step representation of the max-tiles optimization process. A 'balanced field' is defined as a field that has all requested epochs scheduled.

3.1 Max-tiles optimization

Our max-tiles optimization algorithm works around the filter balancing problem by optimizing the 'max-tiles' parameter,[4] which sets an upper limit on the number of fields that are scheduled (e.g. a max-tiles value of 15 means that a maximum of 15 fields can be scheduled). It optimizes this parameter such that the number of fields with observations in all requested filters (i.e. 'balanced' fields) is maximized, iterating through a reasonable range of max-tiles values and calculating the number of balanced fields each time. If the optimization parameter starts decreasing at any point (indicating that we have reached the point where there are too many fields to ensure all required exposures are scheduled), it exits from the loop and the max-tiles parameter is now set for the rest of the scheduling process. Any scheduled fields that do not have all of the requested observations are removed before finalizing the scheduling queue. Generally, there is no limit on how many filters (repeated or otherwise) can be specified when optimizing the max-tiles value during scheduling; however, due to difficulty in revisiting the same tile multiple times without compromising coverage, it is usually optimal to limit it to around two to three epochs maximum, although this number can vary depending on the size of the localization. This process can be visualized using the flowchart in Fig. 2.

3.2 Slicing in right ascension

Although optimizing the maximum number of tiles can help to increase the number of balanced fields, this method only proves to be effective with certain skymaps. More specifically, in cases where the skymap contains multiple disjointed 'lobes' in the probability distribution, it is still a challenge to schedule a reasonable number of balanced fields; this is because the separation in right ascension between the different lobes leads to each lobe having its own rising and setting times. The block scheduling algorithm does not discriminate between continuous and disjointed localizations, and due to this limitation, it has difficulty in scheduling both epochs within the appropriate observability windows.

We have hence implemented a feature to 'slice' the skymap in right ascension,[5] giving the scheduler the ability to distinguish between the different lobes and schedule them separately rather than treating the skymap as a whole. After slicing, the scheduler optimizes for the best order that each slice should be scheduled based on the location of the telescope. The block scheduling algorithm is still used for each slice, thus minimizing the number of filter changes; however, there are additional filter changes incorporated for the transition between each slice, which is necessary to keep up with the lobes' rising and setting times.

The results of these two features are shown in Fig. 3 for ZTF, with the left and right columns displaying the before and after skymaps. The top row displays the results for a skymap that is primarily concentrated in one area in the Northern hemisphere (most of the southern lobe is not accessible), meaning that simply using the max-tiles option is sufficient. The bottom row, in turn, shows results for a skymap in which it would be useful to use both the right ascension slicing and the max-tiles option. The number of green fields (fields with all requested exposures) increases drastically in both cases, demonstrating that these two new options are effective in solving the filter balancing problem when used appropriately. More quantitatively, the cumulative probability covered (only taking into consideration tiles that have had all requested epochs scheduled) increases from 5.7 to 11.5 per cent for the event shown in the top row, and from 2.1 to 24.9 per cent for that shown in the bottom row.

4 CONCLUSION

In this work, we have optimized the search for GW counterparts through improvements of scheduling pipelines that rely on multitelescope networks. We have presented the different features that we have implemented in the pursuit of making the scheduling of ToO observations more flexible and efficient, including taking previous/ongoing observations into account, and scheduling filter blocks with optimized slicing of the skymap. All of these improvements are important in addressing previous challenges associated with synoptic searches of counterparts in large and multilobed localizations, and work to make future electromagnetic follow-up an overall smoother and more optimally automated process.

The dynamic scheduling and filter balancing features were implemented in `gwemopt`, which is the software used to perform scheduling for both the Global Relay of Observatories Watching Transients Happen (GROWTH) and the Global Rapid Advanced Network Devoted to the Multi-messenger Addicts (GRANDMA) projects. These networks span across multiple continents, comprising tens of observatories working in a joint effort to successfully obtain multiwavelength observations of GW candidates. The dedicated follow-up of BNS and NSBH merger candidates undertaken by the GROWTH (e.g. Andreoni et al. 2019b, 2020, Coughlin et al. 2019e; Goldstein et al. 2019) and GRANDMA (e.g. Antier et al. 2019, 2020) networks throughout O3 led to the realization early on that more scheduling features would need to be implemented in order to facilitate this process, prompting

[4]The corresponding command-line option is `--doBalanceExposure`

[5]The corresponding command-line option is `--doRASlices`

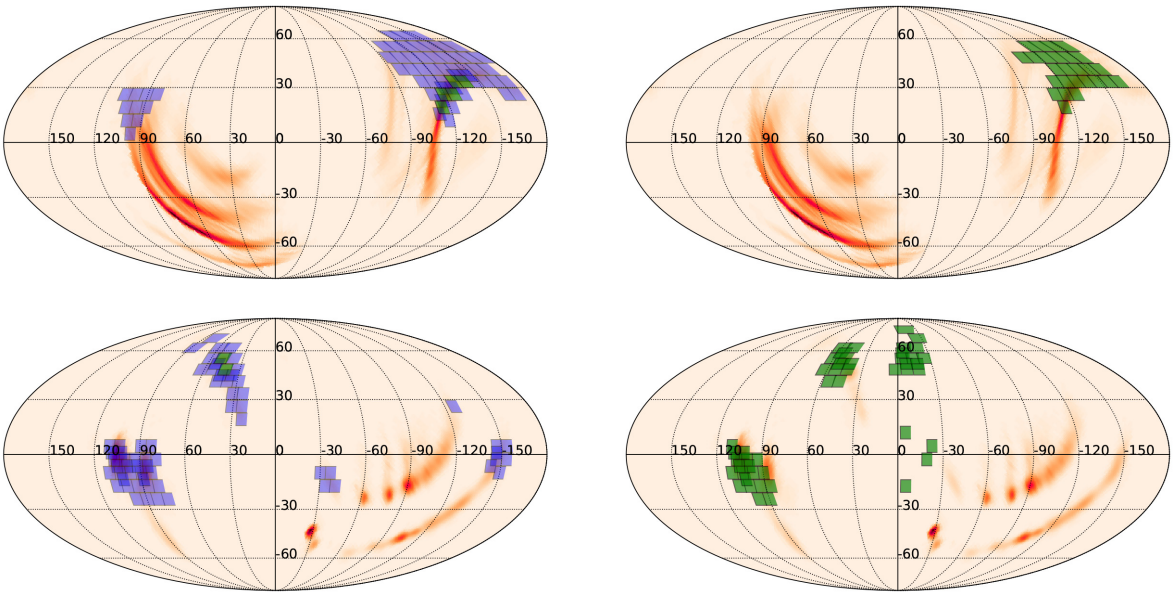

Figure 3. Skymap coverage with ZTF before and after the use of the appropriate filter balancing features discussed in Section 3. The top row displays the results for GW190425 (Abbott et al. 2020b), without (on the left) and with (on the right) the use of max-tiles optimization (Section 3.1). The bottom row displays coverage for S191213g (LIGO Scientific Collaboration & Virgo Collaboration 2019c); in this case, we compare the results when not using any of the filter balancing features (on the left), versus when both the max-tiles optimization and right ascension slicing (Section 3.2) are used (on the right). Fields represented in green have had all requested observations scheduled, while those in violet have not. It is evident that the number of balanced fields increases significantly when the new filter balancing features are put to use.

the subsequent development and implementation of the features described in Sections 2 and 3 throughout the rest of the observing run. The ToO marshal[6] (Coughlin et al. 2019a; Kasliwal et al. 2019b) and the ICARE (*Interface and Communication for Addicts of the Rapid follow-up in multi-messenger Era*) pipeline are the main drivers in coordinating the entire follow-up process for the GROWTH and GRANDMA networks, respectively, and are able to do so by combining the tiling, scheduling, and vetting processes into one cohesive platform. Optimizing all of the elements that lead up to the eventual classification of candidate counterparts is vital to an ultimately productive attempt at follow-up, and key to enable further progress during this exciting new era of GW astronomy.

ACKNOWLEDGEMENTS

MA and KA thank LIGO Laboratory at the California Institute of Technology for hosting the visit that led to this publication. MWC was supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. SA acknowledges support from the GROWTH project funded by the National Science Foundation under grant no. 1545949. NG acknowledges a research grant from the Mohammed Bin Rashid Space Centre (United Arab Emirates), which supported this work.

REFERENCES

Aasi J. 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, L12
Abbott B. P. et al., 2017d, ApJ, 848, L13
Abbott B. P. et al., 2018, Phys. Rev. Lett., 121, 161101
Abbott B. P. et al., 2019, Phys. Rev. X, 9, 031040
Abbott B. P. et al., 2020a, preprint (arXiv:1304.0670v10)
Abbott B. P. et al., 2020b, ApJ, 892, L3
Acernese F. et al., 2015, Class. Quantum Gravity, 32, 024001
Andreoni I. et al., 2019a, PASP, 131, 068004
Andreoni I. et al., 2019b, ApJ, 881, L16
Andreoni I. et al., 2020, ApJ, 890, 131
Annala E., Gorda T., Kurkela A., Vuorinen A., 2018, Phys. Rev. Lett., 120, 172703
Antier S. et al., 2019, MNRAS, 492, 3904
Antier S. et al., 2020, preprint (arXiv:2004.04277)
Arcavi I. et al., 2017, Nature, 551, 64 EP
Bauswein A. et al., 2017, ApJ, 850, L34
Bellm E. C. et al., 2018, PASP, 131, 018002
Bellm E. C. et al., 2019, BAAS, 51, 125
Chornock R. et al., 2017, ApJ, 848, L19
Coughlin M. W. et al., 2018a, MNRAS, 478, 692
Coughlin M. W. et al., 2018b, MNRAS, 480, 3871
Coughlin M. W. et al., 2019a, PASP, 131, 048001
Coughlin M. W. et al., 2019b, MNRAS, 489, 5775
Coughlin M. W., Dietrich T., Margalit B., Metzger B. D., 2019c, MNRAS, 489, L91
Coughlin M. W. et al., 2019d, MNRAS, 492, 863
Coughlin M. W. et al., 2019e, ApJ, 885, L19
Coughlin M. W. et al., 2020, Phys. Rev. Res., 2, 022006
Coulter D. A. et al., 2017, Science, 358, 1556
Cowperthwaite P. S., Berger E., Villar V. A. et al., 2017, ApJ, 848, L17
Goldstein A. et al., 2017, ApJ, 848, L14
Goldstein D. A. et al., 2019, ApJ, 881, L7
Gomez S. et al., 2019, ApJ, 884, L55
Guessoum N., Zitouni H., Mochkovitch R., 2018, A&A, 620, A131
Hotokezaka K., Nakar E., Gottlieb O., Nissanke S., Masuda K., Hallinan G., Mooley K. P., Deller A. T., 2019, Nat. Astron., 3, 940
Kasliwal M. M. et al., 2019a, MNRAS
Kasliwal M. M. et al., 2019b, PASP, 131, 038003

[6] https://github.com/growth-astro/growth-too-marshal

Kessler R. et al., 2015, AJ, 150, 172
Lai X., Zhou E., Xu R., 2019, Eur. Phys. J. A, 55, 60
Li L.-X., Paczynski B., 1998, ApJ, 507, L59
LIGO Scientific Collaboration, Virgo Collaboration, 2017, GCN Circ. 21513
LIGO Scientific Collaboration, Virgo Collaboration, 2019a, GCN Circ. 24237
LIGO Scientific Collaboration, Virgo Collaboration, 2019b, GCN Circ. 25333
LIGO Scientific Collaboration, Virgo Collaboration, 2019c, GCN Circ. 26417
Margalit B., Metzger B., 2017, ApJ, 850, L19
Metzger B. D., 2019, Living Rev. Relativ., 23, 1
Metzger B. D., Berger E., 2012, ApJ, 746, 48
Most E. R., Weih L. R., Rezzolla L., Schaffner-Bielich J., 2018, Phys. Rev. Lett., 120, 261103
Nakar E., 2019, preprint (arXiv:1912.05659)
Nakar E., Piran T., 2011, Nature, 478, 82
Pian E. et al., 2017, Nature, 551, 67
Piran T., Nakar E., Rosswog S., 2013, MNRAS, 430, 2121
Radice D., Perego A., Zappa F., Bernuzzi S., 2018, ApJ, 852, L29
Savchenko V. et al., 2017, ApJ, 848, L15
Smartt S. J. et al., 2017, Nature, 551, 75
Tanaka M., 2016, Adv. Astron., 2016, 1
Watson D. et al., 2019, Nature, 574, 497

This paper has been typeset from a TEX/LATEX file prepared by the author.