

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

²*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 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*

ABSTRACT

The simultaneous detection of electromagnetic (EM) and gravitational waves from the coalescence of two neutron stars (GW/GRB170817A) has ushered in a new era of “multi-messenger” astronomy, with EM detections spanning from gamma to radio. This great opportunity for new scientific investigations raises the issue of how the available multi-messenger 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, identifying the relevant challenges that have appeared during the latest observing run. 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.

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 twin Advanced LIGO (Aasi et al 2015) detectors, yielded the detection of a total of ten binary black hole (BBH) mergers and one binary neutron star (BNS) coalescence (Abbott et al. 2018b). Most recently, the improved sensitivity of the instruments during the third observing run (O3) has resulted in 54 detections to date - many of which are BNS or neutron star-black hole (NSBH) merger candidates (updated information can be found on the Gravitational-wave Candidate Event Database, or GraceDB¹).

Due to the association of BNS and NSBH mergers with potentially detectable EM counterparts (Metzger & Berger 2012; Chu et al. 2016), substantial efforts have been invested into optimizing follow-up observa-

tions of such candidates (e.g., Coughlin et al. 2019; 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), as well as optical and NIR transients (“kilonovae”, or KNe) powered by the decay of r-process nuclei that are ejected relatively isotropically (e.g., Li & Paczynski 1998; Nakar & Piran 2011; Metzger & Berger 2012; Piran et al. 2013; Goldstein et al. 2017; Guessoum et al. 2018).

The culmination of these follow-up efforts came to fruition on the 17th of August, 2017, unveiling the new era of multi-messenger astronomy with the detection of GW170817 along with short gamma-ray burst GRB 170817A (Abbott et al. 2017a), which were soon accompanied by the discovery of transient counterpart AT 2017gfo in NGC 4993 (D \sim 40 Mpc) (Abbott et al. 2017b,c,d; Cowperthwaite et al. 2017) and further successful broadband observations of the event. 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, ini-

¹ <https://gracedb.ligo.org/>




Figure 1. Plots of coverage for S190426c without (a) and with (b) 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. Breaking the night up into two blocks and using the SuperScheduler, 14 previously failed attempts at observation were successfully rescheduled (shown in blue), as opposed to 0 that were rescheduled without the use of the algorithm. The rectangle widths (representing exposure time) have been scaled by a factor of 50 for visualization.

tially localized to $\sim 31 \text{ deg}^2$ at the 90% credibility level and with luminosity distance $40 \pm 8 \text{ Mpc}$ (Singer 2017; Abbott et al. 2019). 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 kilonovae (e.g., Chornock et al. 2017; Coulter et al. 2017; Cowperthwaite et al. 2017; Pian et al. 2017; Smartt et al. 2017; Watson et al. 2019; Kasliwal et al. 2019) and the expansion rate of the Universe (Abbott et al. 2017; Hotokezaka et al. 2019; Coughlin et al. 2019a), 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; Coughlin et al. 2019b, 2018a, 2019c; Annala et al. 2018; Most et al. 2018; Radice et al. 2018; Abbott et al. 2018a; Lai et al. 2019).

During O3, the localization areas (at 90% credible interval) for BNS and NSBH mergers have consistently been in the thousands of square degrees; these values are in stark contrast with the aforementioned localization area of GW170817, meaning that most future detections are likely to pose challenges in obtaining significant coverage of the skymap. It is thus important to optimize our methods in performing follow-ups to GW triggers, which will greatly increase the odds of detecting an EM counterpart.

A codebase named `gwemopt`² (Gravitational-Wave ElectroMagnetic OPTimization) was hence developed

(Coughlin et al. 2018b), 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. 2019d), 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 multi-telescope observations will improve our ability to cover areas in the localization that may not be accessible to one given telescope (e.g., the localization could cover 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-

² <https://github.com/mcoughlin/gwemopt>

and single-telescope observations. In Section 2, we discuss the novel ability for `gwemopt` to take into account previously completed observations when re-scheduling, and in Section 3, we describe two features that drastically improve multi-epoch coverage of events.

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 re-scheduling 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 multi-telescope 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 probabilities associated with 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 re-observation. 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.

As shown on the left of Figure 1, the normal scheduling algorithm does not recognize that there are tiles that have already been observed (in green) when scheduling the following rounds. As a result, there are not as many previously unobserved tiles scheduled for observation (shown in blue). Conversely, the results using the SuperScheduler algorithm on the right of Figure 1 show that prioritizing unobserved tiles led to a higher number of blue tiles in the successive rounds. Evidently, incor-

porating 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. 2019) 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; 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 color change (only second to the $g-z$ pair) over the days following the detection (Andreoni et al. 2019).

It is important to promptly process images during the transient-filtering stage so 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, R 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 minutes (Bellm et al. 2019). It is justified, then, to place emphasis on scheduling at least two epochs during block scheduling³.

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

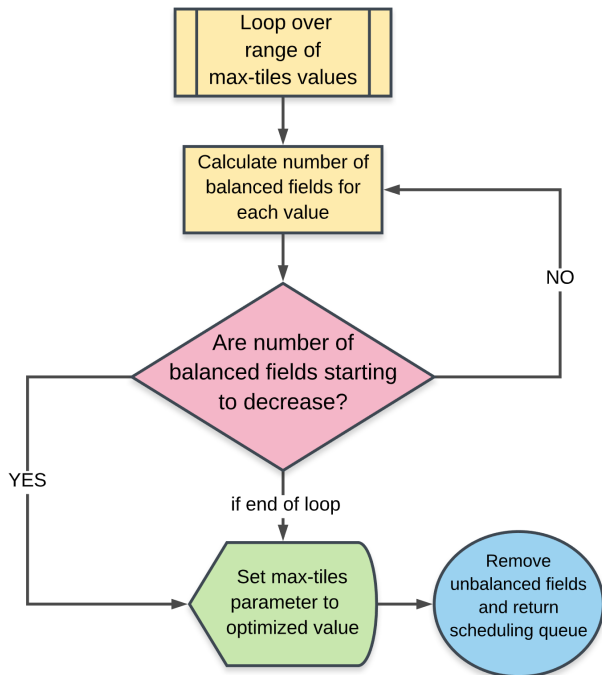


Figure 2. Step-by-step representation of the max-tiles optimization process. A “balanced field” is defined as a field that has all requested epochs scheduled.

3.1. Max-tiles optimization

Our max-tiles optimization algorithm⁴ works around the filter balancing problem by optimizing the “max-tiles” parameter, 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. This process can be visualized using the flowchart in Figure 2.

3.2. Slicing in right ascension

Although optimizing the maximum number of tiles can help to increase the amount 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 time. The block scheduling algorithm does not discriminate between continuous and disjointed localizations, and due to this limitation, has difficulty in scheduling both epochs within the appropriate observability windows.

We have hence implemented a feature to “slice” the skymap in right ascension⁵, 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 Figure 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% for the event shown in the top row, and from 2.1% to 24.9% for that shown in the bottom row.

4. CONCLUSION

In this work, we have continued to optimize the search for GW counterparts through improvements of scheduling pipelines that rely on multi-telescope 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

⁴ `--doBalanceExposure`

⁵ `--doRASlices`

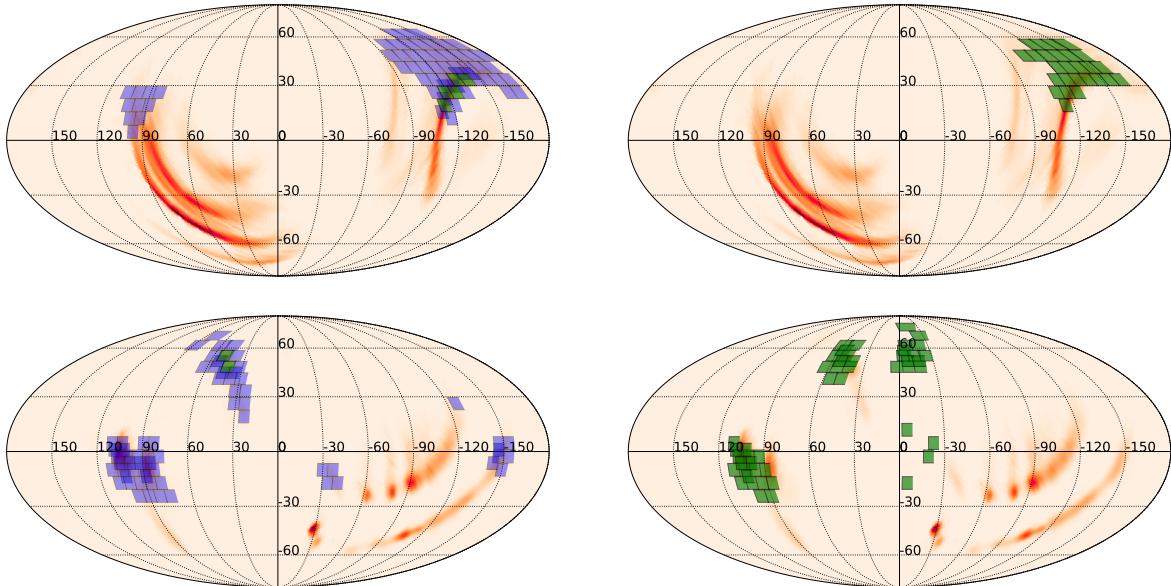


Figure 3. Skymap coverage with ZTF before and after the use of the appropriate filter balancing features discussed in Section 3. The top row displays the results for S190425z, without (on the left) and with (on the right) the use of max-tiles optimization (Section 3.1). The bottom row displays coverage for S191213g; 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 blue have not. It is evident that the number of balanced fields increases significantly when the new filter balancing features are put to use.

scheduling filter blocks with optimized slicing of the skymap. All of these improvements are important in addressing previous challenges associated with synoptic searches of large and multi-lobed localizations, and work to make future EM follow-up an overall smoother and more optimally automated process.

The dynamic scheduling and filter balancing features were implemented in `gwemopt`, an open-source scheduling software, but also contribute to the larger mission of 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 perform successful multi-wavelength follow-up of GW candidates. The ToO marshal (Coughlin et al. 2019)⁶ and the ICARE (*Interface and Communication for Addicts of the Rapid follow-up in multi-messenger Era*) pipeline (Antier et al. 2019) are the main drivers in coordinating this 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 enabling further progress during this exciting new era of GW astronomy.

M. Almualla and K. Alqassimi thank LIGO Laboratory at the California Institute of Technology for hosting the visit that led to this publication. M. W. Coughlin is supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. S. Anand acknowledges support from the GROWTH project funded by the National Science Foundation under Grant No 1545949. Nidhal Guessoum acknowledges a research grant from the Mohammed Bin Rashid Space Centre (UAE), which supported this work.

REFERENCES

Aasi et al. 2015, *Classical and Quantum Gravity*, 32, 074001

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018a, *Phys. Rev. Lett.*, 121, 161101. <https://link.aps.org/doi/10.1103/PhysRevLett.121.161101>

⁶ <https://github.com/growth-astro/growth-too-marshal>

[//link.aps.org/doi/10.1103/PhysRevLett.121.161101](https://link.aps.org/doi/10.1103/PhysRevLett.121.161101)

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *Nature*, 551, 85
- Abbott, B. P., et al. 2018b, arXiv:1811.12907
- Abbott et al. 2017a, *The Astrophysical Journal Letters*, 848, L13.
<http://stacks.iop.org/2041-8205/848/i=2/a=L13>
- . 2017b, *Astrophysical Journal Letters*, 848, L12
- . 2017c, *Nature*, 551, 85.
<http://dx.doi.org/10.1038/nature24471>
- . 2017d, *The Astrophysical Journal Letters*, 850, L39.
<http://stacks.iop.org/2041-8205/850/i=2/a=L39>
- . 2019, arXiv e-prints, arXiv:1304.0670v9
- Acernese et al. 2015, *Classical and Quantum Gravity*, 32, 024001
- Andreoni, I., et al. 2019, *Publ. Astron. Soc. Pac.*, 131, 068004
- Andreoni, I., Goldstein, D. A., Kasliwal, M. M., et al. 2020, *ApJ*, 890, 131
- Annala, E., Gorda, T., Kurkela, A., & Vuorinen, A. 2018, *Phys. Rev. Lett.*, 120, 172703. <https://link.aps.org/doi/10.1103/PhysRevLett.120.172703>
- Antier, S., Agayeva, S., Aivazyan, V., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 492, 3904.
<https://doi.org/10.1093/mnras/stz3142>
- Arcavi et al. 2017, *Nature*, 551, 64 EP .
<http://dx.doi.org/10.1038/nature24291>
- Bauswein et al. 2017, *The Astrophysical Journal Letters*, 850, L34.
<http://stacks.iop.org/2041-8205/850/i=2/a=L34>
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 131, 018002.
<https://doi.org/10.1088%2F1538-3873%2Faaecbe>
- Bellm, E. C., Ford, E. B., Tohuavohu, A., et al. 2019, *Scheduling Discovery in the 2020s*, , arXiv:1907.07817
- Chornock et al. 2017, *The Astrophysical Journal Letters*, 848, L19.
<http://stacks.iop.org/2041-8205/848/i=2/a=L19>
- Chu, Q., Howell, E. J., Rowlinson, A., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 459, 121–139.
<http://dx.doi.org/10.1093/mnras/stw576>
- Coughlin, M. W., Dietrich, T., Heinzel, J., et al. 2019a, arXiv:1908.00889
- Coughlin, M. W., Dietrich, T., Margalit, B., & Metzger, B. D. 2019b, *Monthly Notices of the Royal Astronomical Society: Letters*, 489, L91.
<https://doi.org/10.1093/mnrasl/slz133>
- Coughlin, M. W., Dietrich, T., Doctor, Z., et al. 2018a, *Monthly Notices of the Royal Astronomical Society*, 480, 3871. <http://dx.doi.org/10.1093/mnras/sty2174>
- Coughlin, M. W., Tao, D., Chan, M. L., et al. 2018b, *Monthly Notices of the Royal Astronomical Society*, 478, 692. <http://dx.doi.org/10.1093/mnras/sty1066>
- Coughlin, M. W., Dietrich, T., Antier, S., et al. 2019c, *Monthly Notices of the Royal Astronomical Society*, 492, 863. <https://doi.org/10.1093/mnras/stz3457>
- Coughlin, M. W., Antier, S., Corre, D., et al. 2019d, *Monthly Notices of the Royal Astronomical Society*, 489, 5775–5783. <https://doi.org/10.1093/mnras/stz2485>
- Coughlin, M. W., Ahumada, T., Cenko, S. B., et al. 2019e, *Publications of the Astronomical Society of the Pacific*, 131, 048001
- Coughlin et al. 2019, *The Astrophysical Journal*, 885, L19.
<https://doi.org/10.3847%2F2041-8213%2F190101>
- Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, *Science*, 358, 1556
- Cowperthwaite, P. S., Berger, E., Villar, V. A., et al. 2017, *ApJL*, 848, L17
- Goldstein, A., Veres, P., Burns, E., et al. 2017, *The Astrophysical Journal*, 848, L14
- Goldstein, D. A., Andreoni, I., Nugent, P. E., et al. 2019, *The Astrophysical Journal*, 881, L7.
<https://doi.org/10.3847%2F2041-8213%2F190101>
- Gomez, S., Hosseinzadeh, G., Cowperthwaite, P. S., et al. 2019, *The Astrophysical Journal*, 884, L55.
<https://doi.org/10.3847%2F2041-8213%2F190101>
- Guessoum, N., Zitouni, H., & Mochkovitch, R. 2018, *Astronomy & Astrophysics*, 620, A131.
<http://dx.doi.org/10.1051/0004-6361/201832940>
- Hotokezaka, K., Nakar, E., Gottlieb, O., et al. 2019, *Nature Astronomy*, doi:10.1038/s41550-019-0820-1.
<https://doi.org/10.1038/s41550-019-0820-1>
- Kasliwal, M. M., Kasen, D., Lau, R. M., et al. 2019, *Monthly Notices of the Royal Astronomical Society: Letters*, <http://oup.prod.sis.lan/mnrasl/advance-article-pdf/doi/10.1093/mnrasl/slz007/27503647/slz007.pdf>.
<https://doi.org/10.1093/mnrasl/slz007>
- Kessler, R. et al. 2015, *The Astronomical Journal*, 150, 172
- Lai, X., Zhou, E., & Xu, R. 2019, *The European Physical Journal A*, 55, 60.
<https://doi.org/10.1140/epja/i2019-12720-8>
- Li, L.-X., & Paczynski, B. 1998, *The Astrophysical Journal Letters*, 507, L59.
<http://stacks.iop.org/1538-4357/507/i=1/a=L59>
- Margalit, B., & Metzger, B. 2017, *The Astrophysical Journal Letters*, 850, doi:10.3847/2041-8213/aa991c
- Metzger, B. D., & Berger, E. 2012, *Astrophys. J.*, 746, 48
- Most, E. R., Weih, L. R., Rezzolla, L., & Schaffner-Bielich, J. 2018, *Phys. Rev. Lett.*, 120, 261103. <https://link.aps.org/doi/10.1103/PhysRevLett.120.261103>

- Nakar, E., & Piran, T. 2011, *Nature*, 478, 82.
<https://doi.org/10.1038/nature10365>
- Pian et al. 2017, *Nature*, 551, 67 EP .
<http://dx.doi.org/10.1038/nature24298>
- Piran, T., Nakar, E., & Rosswog, S. 2013, *Monthly Notices of the Royal Astronomical Society*, 430, 2121.
<http://dx.doi.org/10.1093/mnras/stt037>
- Radice, D., Perego, A., Zappa, F., & Bernuzzi, S. 2018, *The Astrophysical Journal Letters*, 852, L29.
<http://stacks.iop.org/2041-8205/852/i=2/a=L29>
- Singer, L. 2017, *GCN*, 21513
- Smartt et al. 2017, *Nature*, 551, 75 EP .
<http://dx.doi.org/10.1038/nature24303>
- Watson, D., Hansen, C. J., Selsing, J., et al. 2019, *Nature*, 574, 497