The Astrophysical Journal Letters, 975:L19 (12pp), 2024 November 1 https://doi.org/10.3847/2041-8213/ad87ce

Swiftly Chasing Gravitational Waves across the Sky in Real Time

Aaron Tohuvavohu[1,2,3], Jamie A. Kennea[4], Christopher J. Roberts[5], James DeLaunay[4], Samuele Ronchini[4,6], S. Bradley Cenko[7,8], Becca Ewing[9], Ryan Magee[10], Cody Messick[11], Surabhi Sachdev[12], and Leo P. Singer[7]
[1] David A. Dunlap Department of Astronomy and Astrophysics, University of Toronto, Toronto M5S 3H7, Canada
[2] Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto M5S 3H7, Canada
[3] Cahill Center for Astronomy and Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA
[4] Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA
[5] Astrophysics Projects Division, NASA Goddard Space Flight Center, Mail Code 440, Greenbelt, MD 20771, USA
[6] Institute for Gravitation & the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA
[7] Astrophysics Science Division, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA
[8] Joint Space-Science Institute, University of Maryland, College Park, MD 20742, USA
[9] Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA
[10] LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA
[11] University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA
[12] School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA
Received 2024 October 7; revised 2024 October 14; accepted 2024 October 15; published 2024 October 30

Abstract

We introduce a new capability of the Neil Gehrels Swift Observatory, dubbed "continuous commanding," that achieves 10 s latency response time on orbit to unscheduled target-of-opportunity requests received on the ground. We show that this will allow Swift to respond to premerger (early-warning) gravitational-wave (GW) detections, rapidly slewing the Burst Alert Telescope (BAT) across the sky to place the GW origin in the BAT field of view at or before merger time. This will dramatically increase the GW/gamma-ray burst (GRB) codetection rate and enable prompt arcminute localization of a neutron star merger. We simulate the full Swift response to a GW early-warning alert, including input sky maps produced at different early-warning times, a complete model of the Swift attitude control system, and a full accounting of the latency between the GW detectors and the spacecraft. 60 s of early warning can double the rate of a prompt GRB detection with arcminute localization, and 140 s guarantees observation anywhere on the unocculted sky, even with localization areas $\gg$1000 deg^2. While 140 s is beyond current GW detector sensitivities, 30–70 s is achievable today. We show that the detection yield is now limited by the latency of LIGO/Virgo cyberinfrastructure and motivate a focus on its reduction. Continuous commanding has been integrated as a general capability of Swift, significantly increasing its versatility in response to the growing demands of time-domain astrophysics. We demonstrate this potential on an externally triggered fast radio burst (FRB), slewing 81° across the sky, and collecting X-ray and UV photons from the source position <150 s after the trigger was received from the Canadian Hydrogen Intensity Mapping Experiment, thereby setting the earliest and deepest such constraints on high-energy activity from nonrepeating FRBs. The Swift Team invites the community to consider and propose novel scientific applications of ultra-low-latency UV, X-ray, and gamma-ray observations.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gamma-ray bursts (629); Space telescopes (1547)

1. Introduction

The merger of two compact objects provides a rich laboratory in which to explore the fundamental physics of spacetime. If at least one of these objects is baryonic, the potential yield expands enormously, with accretion physics, ultrarelativistic jet launching, heavy-element nucleosynthesis, perturbed supranuclear matter, and more all emerging rapidly following the merger (E. Burns 2020). Advanced gravitational-wave (GW) detectors have begun to detect these binary coalescences in their late inspiral and merger phases, providing merger times, source classification, and sky localizations (R. Abbott et al. 2023). Mining the rich postmerger physics requires electromagnetic facilities to rapidly locate and study these phenomena before the quickly evolving processes shut off or fade below detectability. In one case, this has been partially successful, with the GW detection of a binary neutron star (BNS) merger (GW170817; B. P. Abbott et al. 2017a, 2017b), a gamma-ray burst (GRB) detected 1.7 s later (A. Goldstein et al. 2017; V. Savchenko et al. 2017), and the subsequent search for and discovery of the associated kilonova at 11 hr postmerger in the optical and then UV and IR (V. A. Villar et al. 2017, and references therein), followed by detection of a synchrotron afterglow at 16 days after the merger in the X-ray (D. Haggard et al. 2017; R. Margutti et al. 2017; E. Troja et al. 2017) and radio (K. D. Alexander et al. 2017; G. Hallinan et al. 2017). These observations, and follow-up over the years since merger, have yielded a tremendous amount of information on the production of heavy *r*-process elements and the chemical enrichment of the Universe, the physics and structure of relativistic jets, and the equation of state of matter at supranuclear densities (R. Margutti & R. Chornock 2021).

However, much was missed. Focused observations with humanity's most sensitive facilities did not begin until ~1 day after the merger (B. P. Abbott et al. 2017b). Despite tremendous efforts, significant uncertainty remains in the composition of the ejecta that drove the kilonova and the

mechanism driving the early emission due to model degeneracy. Some of these degeneracies can only be broken with early observations, during the kilonova's rise within a few hours postmerger (I. Arcavi 2018). Additionally, several hypothesized emission components emerge at significantly earlier times. These include fast thermal UV emission driven by the beta decay of free neutrons after ∼880 s (S. R. Kulkarni 2005; B. D. Metzger et al. 2015; C. Dean et al. 2021), UV and X-ray emission from a magnetar remnant driving fast winds and shock breakout (G.-L. Wu et al. 2022; L. Combi & D. M. Siegel 2023), very early TeV emission from the forward shock in the jet (B. Banerjee et al. 2023), an optical flash on a timescale of tens of seconds from synchrotron emission driven by the reverse shock (G. Oganesyan et al. 2023), and more. It is clear that there is compelling physics that presents itself only briefly, within the first seconds to hours postmerger. Current approaches to counterpart discovery and characterization mostly focus on the thermal (kilonova) emission (which rises on timescales of hours) and involve time-consuming wide-field optical searches and transient characterization (K. Herner et al. 2020; T. Ahumada et al. 2024). Even in the case of a well-localized GW event, the typical time to optical counterpart discovery and follow-up with more powerful telescopes is hours to days. This approach is not suitable for accessing the early-time physics.

The ideal scenario would be to locate the BNS promptly, to a precision sufficient for immediate follow-up with our most powerful focused telescopes, which typically have field of view (FOV) ~arcminutes across. This would bypass the slow wide-field search stage and allow access to the early-time physics with sensitive observations across the electromagnetic spectrum. To do this, we should use the earliest known emission. The gamma rays are produced within seconds of merger, and existing instruments can determine gamma-ray source positions to the required precision. However, the combination of instantaneous FOV, sensitivity, and localization precision required is challenging to achieve. While all-sky gamma-ray instruments do exist (e.g., Fermi/GBM, INTEGRAL/SPI-ACS), their localization precision is comparable to or worse than the GW localizations. In contrast, arcminute-scale localizations for gamma rays can be determined, but the various existent techniques to achieve this necessarily limit the FOV. Simultaneously, high sensitivity is required due to the large fraction of BNS mergers where the GRB jet is viewed off-axis, resulting in fainter gamma-ray emission.

The Burst Alert Telescope (BAT; S. D. Barthelmy et al. 2005) on the Neil Gehrels Swift Observatory (N. Gehrels et al. 2004) is the most sensitive GRB detector in current operation and capable of regular 1′–3′ precision positions. A. Tohuvavohu et al. (2020) and J. DeLaunay & A. Tohuvavohu (2022) introduced new data products and analysis techniques for Swift/BAT that significantly (∼4×) increase the detection rate and sensitive volume for faint GRBs. These capabilities have led to multiple searches, including real-time joint searches with the International Gravitational Wave Observatory Network (IGWN), with results providing the deepest upper limits to date on gamma-ray emission at the time of GWs (C. Fletcher et al. 2024; G. Raman et al. 2024; S. Ronchini et al. 2024). However, Swift/BAT is a directional instrument, providing superior sensitivity and localization precision over "only" ∼2.2 sr, or ∼17% of the sky. This limits the current detection yield at least as much as the sensitivity does, as BNS mergers are known to be rare in the local Universe. If only we knew where to point.

In parallel, the increasing sensitivity of the IGWN (particularly at low frequencies) has raised a unique opportunity to detect and roughly localize BNS mergers seen in GW data before merger (K. Cannon et al. 2012). These early-warning searches work in the same manner as traditional matched-filter searches on the GW data but instead use templates truncated at several different final frequencies during the inspiral, equivalent to different times before merger. However, the current sensitivities of the GW detectors allow confident detection no more than 70 s before merger and with localization areas of $>\mathcal{O}(10^3)$ deg^2. Not many instruments are able to take advantage of such an opportunity. This would require sensitivity to the prompt emission produced in BNS mergers, a response time of tens of seconds across the sky, and an instantaneous FOV of >1000 deg^2. Here we investigate the capabilities of Swift/BAT for this purpose and show that the spacecraft and instrument are, in principle, capable of successfully utilizing an early-warning detection to ensure counterpart discovery.

Swift was designed to solve the mystery of short GRBs and their connection to neutron star mergers, to quote the eponymous song:

> Swiftly swirling, gravity twirling
> Neutron stars about to collide.
> Off in a galaxy so far away,
> Catastrophic interplay,
> Roller coaster gamma-ray ride.
> Super bright explosion then,
> Never to be seen again,
> How are we supposed to know?
> How about a telescope rotation [Swift is the satellite that swings]
> Swiftly onto the location [Onto those brightly bursting things]
> Of its panchromatic afterglow? [To grab the multiwavelength answer of what makes them glow!]
> –Padi Boyd, Swift Song, 1999

Swift has been slewing immediately and autonomously within seconds of emission time since launch, triggered by a GRB detected on board. However, its capacity for a rapid response to external triggers has long been hampered by operational and command and control limitations. Here we complete the final prophesied evolution of Swift, swinging it rapidly across the sky to chase down neutron stars before they merge.

In Section 2, we discuss the expected gamma-ray emission and the relevant timescales necessary for observation. In Section 3, we describe the detection of GWs premerger and discuss some challenges inherent to using these alerts. In Section 4, we formulate the problem of early-warning response and describe some constraints of Swift that shape the solution space. In Section 5, we describe a rapid trajectory optimization algorithm to quickly solve the optimal response to an early-warning GW alert. In Section 6, we present results after applying this strategy to a large sample of early-warning alerts. In Section 7, we explore practical latencies that can reduce the amount of time Swift has to respond and demonstrate a new Swift capability for <10 s target-of-opportunity (TOO) response. In Section 8, we conclude with some suggestions for LIGO/Virgo/Kagra necessary to enable this science and a broader invitation to the community to use the new capabilities developed for Swift.

2. Emission Targets and Timescales

BNS mergers are known to produce prompt GRB emission shortly after merger. In the case of GW170817, this signal began ∼1.7 s after merger time with a duration of ∼2.5 s. Because the jet launch delay timescale itself is expected to be very short, the physical origin of the delay in the emission seen likely depends on propagation of the jet to the gamma-ray production site (P. Beniamini et al. 2020). This will in turn vary with the jet Lorentz factor and the circum-merger environment. The prompt emergence of the signal and its short duration necessitate observation at, not after, merger time. However, there can also be longer-lived gamma-ray emission.

Up to ∼10% of the Swift-detected short GRBs (which are thought to result from compact object mergers; P. D'Avanzo 2015) are found to have extended emission, lasting ∼10 −100 s after the prompt phase (Y. Kaneko et al. 2015). The origin is unknown but is often interpreted as continued energy injection from a short-lived (magnetar) postmerger remnant (B. P. Gompertz et al. 2014). In this case, the shutoff slope of the extended emission can be connected to either the magnetic-braking-driven spin-down or the GW-driven spin-down of the remnant (H.-J. Lü et al. 2020). The combination of GW and gamma-ray observations can thereby uniquely constrain the nature of the postmerger remnant. We note that Swift/BAT with its imaging capabilities is especially sensitive to weak, long emission. This longer-lived emission raises the possibility to detect and localize the merger site even if BAT arrives on target up to ∼10−100 s after merger.

The recent observations of 170817-like kilonova signatures embedded in the afterglows of long GRBs 211211A (J. C. Rastinejad et al. 2022; E. Troja et al. 2022) and 230307A (Y.-H. Yang 2024; A. J. Levan et al. 2024) have allowed their BNS merger origin to be established (as opposed to the death of massive stars, thought to be the origin of most long-duration GRBs; O. Luongo & M. Muccino 2021). These bursts have durations of ∼50 and 200 s, respectively. Codetection with GWs and prompt localization and follow-up will be crucial to probe how to power such long emission given the energy budget available in BNS mergers. In this case, it may be possible to detect and localize the merger site even if BAT arrives at the target ∼50−100 s after the merger.

Additionally, there are a plethora of models that predict "precursor" emission before merger time (e.g., B. M. S. Hansen & M. Lyutikov 2001). These involve some interaction of the two neutron stars before merger, either of their magnetospheres (B. D. Metzger & C. Zivancev 2016; E. R. Most & A. A. Philippov 2022) or tidal strain (D. Tsang 2013), triggering various emission mechanisms. The total energy budget available to drive such processes increases rapidly as the neutron stars approach merger, so the emission is generally expected no more than 10 s before the merger time (M. Lyutikov 2024). Detection of precursor emission provides one of the only ways to directly probe the material/nuclear properties of the neutron stars (D. Neill et al. 2021) or their magnetic fields (A. M. Beloborodov 2021) before the messy merger. The maximum achievable luminosity of these phenomena is substantially lower than the prompt GRB emission, challenging the capabilities of current gamma-ray detectors.[13] Swift/BAT is the most sensitive GRB detector currently operating and would have the best chance at detection, assuming arrival on target ∼10 s before merger.

[13] Radio would be a better bet (M. Lyutikov 2024).

3. Early-warning GW Detections

A BNS system produces GW emission throughout its lifetime, slowly shedding angular momentum and inspiraling toward each other while increasing their orbital frequency. The earliest possible detection therefore depends on the low-frequency sensitivity of the detector. The time to merger t_c as a function of GW frequency f is

$$t_c = \frac{5}{256}\frac{c^5}{(G\pi f)^{8/3}} \cdot \frac{1}{\mathcal{M}^{5/3}}, \quad (1)$$

where $\mathcal{M}$ is the chirp mass of the system (P. C. Peters & J. Mathews 1963). The GW signal enters the operating band of the current ground-based detectors at 10–30 Hz, depending on signal amplitude.

R. Magee et al. (2021) first demonstrated the practical capability to detect BNS signal candidates and distributed alerts before merger on a replay of Observing Run 3 (O3) data. Alongside the search itself, they also highlighted some cyberinfrastructure challenges associated with data transfer, calibration, and alert generation to meet the strict latency demands. In O4, all four IGWN real-time compact binary coalescence (CBC) pipelines (gstLAL, pyCBC, MBTA, and SPIIR) now operate early-warning searches alongside their full bandwidth searches (T. Adams et al. 2016; A. H. Nitz et al. 2020; S. Sachdev et al. 2020; M. Kovalam et al. 2022). These pipelines utilize a small number (four to six) of different discrete truncation frequencies (29–58 Hz) corresponding to different premerger times (60–9 s, assuming a 1.4–1.4 $M_\odot$ BNS system). These truncated templates are used in a matched-filter search, allowing signals to be detected in the late inspiral stage as they accumulate signal in the detectors. Depending on the distance, sky position, and source masses of the BNS, a confident detection (network SNR > 10) can be achieved up to 70 s before merger in the case of a GW170817-like signal (40 Mpc) or up to 50 s before merger in the case of a GW190425-like signal (200 Mpc).

Here we use the Early Warning Gravitational Wave Data Release (S. Sachdev et al. 2020) to study the characteristics of early-warning detections and localizations and the optimal strategy for early-warning response with Swift. This data release comprises ∼2 million injected signals with known parameters representative of the expected BNS population demographics in the local Universe (to $z = 0.2$). Injections that are detected are tagged with their gstLAL-recovered SNRs, and BAYESTAR (L. P. Singer & L. R. Price 2016) produced 3D sky maps at truncation frequencies of 29, 32, 38, 49, 56, and 1024 Hz, corresponding to roughly 60, 45, 30, 15, 10, and 0 s before merger (for a fiducial 1.4–1.4 $M_\odot$ system). S. Sachdev et al. (2020) present the localization area evolving as a function of truncation frequency, showing that the localization precision (characterized by 90% enclosed area) evolves with time and accumulated SNR, slowly improving before rapidly converging in the last few seconds before merger. However, GW localizations can be multimodal, with even relatively small sky areas split across opposite hemispheres. The maximum probability position can therefore shift back and forth across the whole sky before converging, as shown in Figure 1.

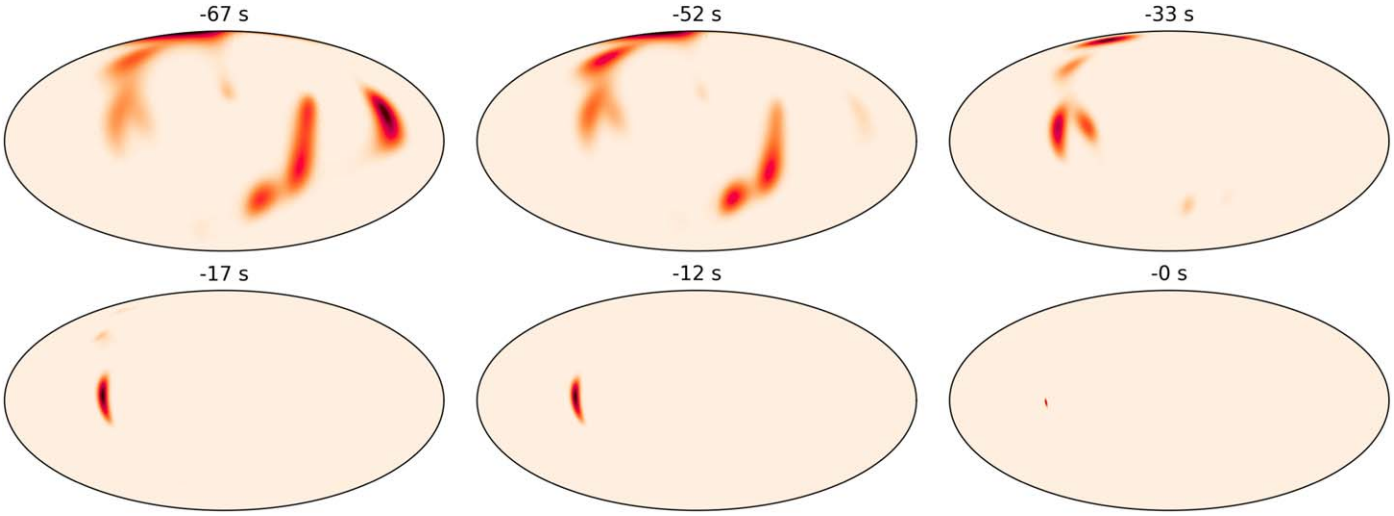

Figure 1. GW localization sky maps at different truncation frequencies (premerger times) for a 1.4–1.2 $M_{\odot}$ merger at 50 Mpc. This example clearly shows the highest-probability mode shifting between three locations in the sky, separated by ∼180°. In this case, a slew decision made on information from the earliest sky map could unknowingly result in looking at the wrong part of the sky.

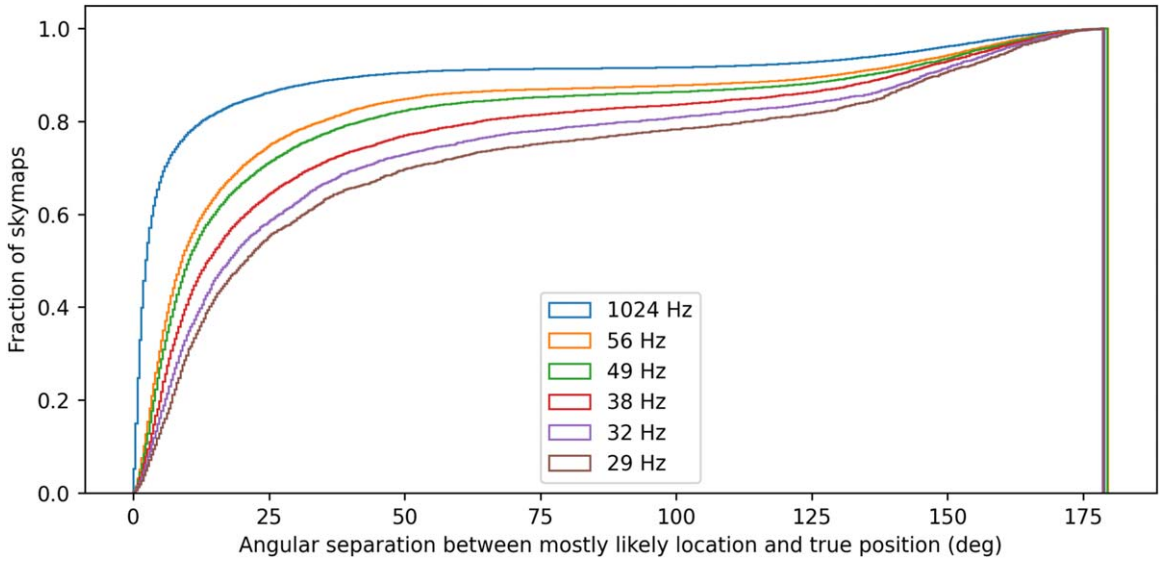

Figure 2. Cumulative distribution of the angular separation between the highest-probability pixel in a sky map and the true location for different truncation frequencies. An observer with a telescope of a given FOV radius can easily read off when to make a slew decision, depending on their risk posture and response time to the target. The BAT FOV is oblong, but its semiminor axis is 40°, thus guaranteeing coverage of the true position >60% of the time, if pointed successfully at the highest-probability location from the 29 Hz (∼–60 s) sky map.

Since this use case requires instantaneous coverage at or before merger time by a telescope with a finite FOV, a perhaps more relevant concern is the positional stability of the localization with time. In Figure 2, we compute the angular separation between the most probable sky pixel in each of the ∼50,000 early-warning sky maps and the true (injected) position of the merger as a function of truncation frequency. We find that the true location is within 40° of the most probable sky map pixel >60% of the time at 29 Hz and increasing for all higher truncation frequencies. In contrast, conditional on there being a BNS merger somewhere on the sky, it will land in the BAT FOV (semiminor axis of 40°) ∼17% of the time if we do nothing. To first order, this suggests that it is almost always beneficial to use the early-warning information, modulo constraints. We will explore this more thoroughly in later sections.

We note that the Early Warning Data Release was produced assuming the LIGO and Virgo detectors at design sensitivity, which has not yet been achieved as of O4. This mostly affects the absolute rate of expected detections and their distances, although changes to the low-frequency behavior of the detectors can and will shape the evolution of the sky localization with time. Additionally, while these results pertain to detection prospects and parameter estimations achievable with realistic data, they do not include search pipeline or other cyberinfrastructure latency. We will return to this important issue in Section 7.

4. Swift Observing Scenarios

We are interested in the optimal Swift response to the following question. If a BNS early-warning alert is received at T0-X s until merger with sky map Y, and Swift is currently observing target Z, what is the action that maximizes the

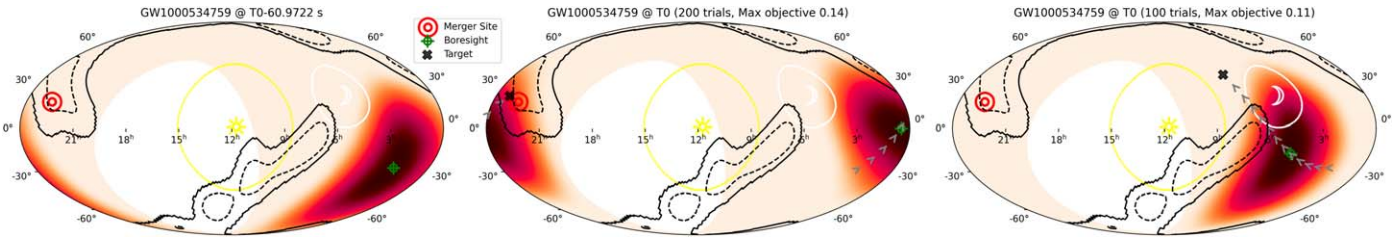

Figure 3. A GW early-warning alert at 29 Hz. Plotted are the GW sky map 50% and 90% containment contours and BAT FOV (heat map) at the exact truncation frequency time relative to merger T0. Left shows the orientation at the time the alert is produced, with little coverage of the GW sky map. Center and right show the optimal trajectory found by the sampler (after 200 and 100 trials, respectively) and the BAT FOV at merger time if that trajectory is executed. The portion of the sky occulted by the Earth at the time is in white, and the Sun and Moon boresight constraints are marked. The bull's-eye marks the true location of the merger, not known to the trajectory optimizer. The black cross is the commanded target of the slew, and the green reticle is the Swift boresight.

chance of the true BNS location falling inside the BAT FOV at merger time?

The response to this question can take the following forms: do nothing (high confidence that the GW origin is in the current BAT FOV), wait for improved information (expect more information in low latency that can substantially inform response), or immediately slew along a trajectory that will maximize the BAT/GW coverage at T0. Some specific characteristics and constraints of the Swift spacecraft and BAT shape the optimal approach.

Once Swift begins to slew, it cannot be interrupted and redirected to another location until the slew is complete. This has significant consequences in decision-making under uncertainty, with the added expectation that updated information may (or may not) arrive. Given the average duration of Swift slews until settling, this lockout period can be >100 s. Since the GW sky maps evolve with time and can contain multiple modes spanning a large fraction of the sky, a plausible scenario is that one is received, a slew is initiated toward the most probable area, and then an updated sky map arrives showing a refined localization on the other side of the sky. Ideally, we would like to avoid that.

The BAT cannot trigger while slewing but can collect data for analysis on the ground. The BAT onboard triggers cannot run while the spacecraft is slewing and are only activated once the "settled" flag is received from the attitude control system (ACS). This means that a typical onboard trigger at $\sim T_{\rm GRB} + 6$ s and associated alerts are not feasible during these times. However, data can be collected during slews and promptly sent to the ground for rapid analysis (A. Tohuvavohu et al. 2020), and many GRBs are regularly detected during slews in this way (A. Tohuvavohu et al. 2021a, 2021b; J. DeLaunay et al. 2022, 2024). Indeed, over longer exposure times, analyses performed on data taken during slews are actually more sensitive than inertially pointed data, as the slew breaks degeneracies in the projected mask pattern, thereby reducing systematic noise in the reconstructed sky images (A. J. Copete 2012).

The BAT coded FOV has significant rotational asymmetry. The coded FOV is the sky area inside which BAT can provide arcminute localizations for detected bursts. BAT is capable of detecting bursts outside of the coded FOV with a reduced sensitivity but is unable to provide localizations to arcminute precision for these bursts. This coded FOV is $\sim 120^\circ \times 80^\circ$ at its maximal dimensions (but it is not rectangular; see, e.g., Figure 3), thereby implying a strong roll-angle dependence on the sky coverage for a given boresight pointing. Only a subset of roll angles are available for a given target boresight at a given time, constrained by spacecraft capabilities. Since GW localization regions are also highly asymmetric, this can significantly affect feasible coverage.

We do not have direct control over the precise Swift slew trajectory. The autonomy on board Swift and its ACS allows significant operational flexibility and substantially reduces risk. However, it also constrains some capability. When given target coordinates, Swift determines its own slew path that the ACS deems optimal and safe, which will depend on the instantaneous orientation of observing constraints (such as the Sun, Moon, Earth limb, and others). Some theoretically possible slew trajectories are therefore unavailable for use, and any specific trajectory desired needs to be reverse engineered from the target coordinates that would generate it at a given time from a fixed origin. Additionally, the velocity profile (of critical importance!) depends sensitively on the target position and roll angle.

5. Making the Right Choice: Trajectory Optimization

From the above considerations, it is clear that the trajectory Swift takes as it slews to a new position will strongly impact the rate of prompt BNS detections. Here we examine the optimum slew trajectory that maximizes the fraction of the GW localization posterior contained inside the BAT coded FOV at merger time. The Swift ACS was designed with significant redundancy, with six reaction wheels arranged for mutually orthogonal three-axis control providing torque for rapid slews across the sky. In 2022 January, one of these reaction wheels failed, and Swift now operates in a five-wheel mode, reducing the maximum torque available. In this mode, the maximum slew rate is ~ 0.8 deg s^{-1}, but this maximum rate and maximum acceleration are not always achievable and depend strongly on the orientation of constraints on the sky and the roll-angle change during the slew. Additionally, the Swift boresight must avoid transgressing within 27° of the Earth limb, 45° from the Sun, and 22° from the Moon for instrument safety during slews. As such, the minimum slew distance between two points can often be substantially longer than the great circle distance. We note that the Earth limb moves $\sim 3.^\circ 8$ minute^{-1} from the perspective of Swift, but slew safety constraints are evaluated by the ACS at the beginning of slews, not the end. Improper trajectory evaluation can therefore lead to Swift getting trapped between two constraints. These and other similar vagaries necessitate careful modeling when attempting to optimize the trajectory. In what follows, we use a complete model of the Swift ACS in its five-wheel mode to achieve this.

In the general case, it is not sufficient to simply solve for the Swift attitude at merger time that maximizes the GW sky map containment in the BAT FOV and then slew toward this. In

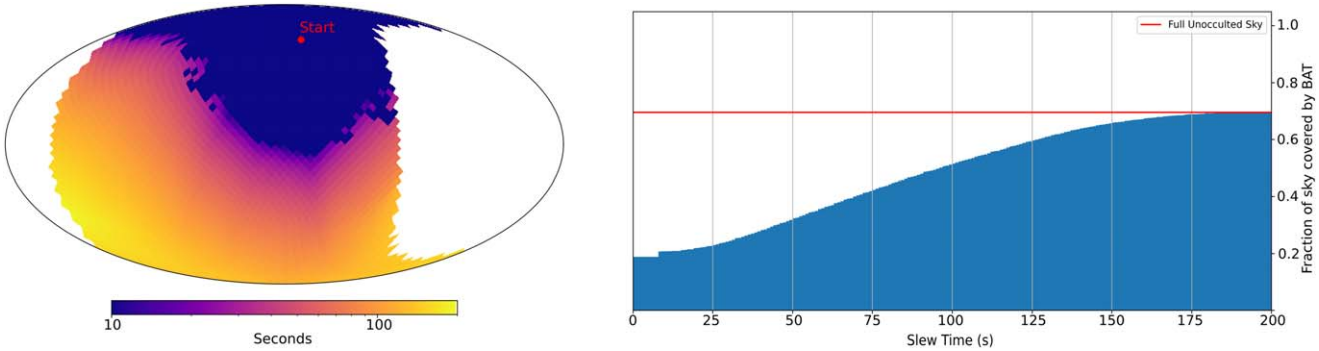

Figure 4. Left: minimum slew time required to put any location on the sky inside the BAT FOV; white sky areas are occulted by the Earth. Right: the cumulative fraction of sky accessible to the BAT FOV as a function of slew time. If given a target, Swift can put almost any position on the unocculted sky into the BAT FOV in <175 s.

most cases, the strictly optimal attitude is not achievable by merger time, and the attitude as a function of time along that trajectory will often not have optimal coverage at merger time. That is, the path to the optimal attitude at merger is not always the optimal trajectory. This, combined with the fact that we cannot directly shape the Swift trajectory, means that fully exploring the space of allowable Swift trajectories is required. In Figure 4, we show an example of the minimum time required to put any location on the unocculted sky into the BAT FOV by exhaustively searching over all possible trajectories, a computationally expensive task ill-suited for real-time operations.

The response to early-warning GW alerts sets strong requirements on latency. We have allocated 1 s of latency between alert receipt and target uplink to the spacecraft. This includes processing the alert, determining the response strategy, and finding the optimal trajectory and the target that will achieve it. Meeting this latency requirement necessitates determining the optimum trajectory very rapidly. We formulate an objective function `slewmaxcorr`, which takes as hyperparameters the slew target R.A., decl., roll, and fixed parameters of $t_{\rm start}$, time until merger, and the GW sky map. It then computes the full time-resolved slew trajectory beginning at $t_{\rm start}$ and finds the instantaneous Swift attitude along this path at merger time. The inner product between the BAT FOV coverage at this time and the GW sky map (masked for occultation) is then calculated and returned. Hereafter, we refer to this quantity as the objective value, and its maximization identifies the optimal Swift trajectory. We maximize `slewmaxcorr`, sampling over the space of allowed target coordinates (constrained samples are pruned). `slewmaxcorr` is optimized for speed but still requires $\sim$10 ms to evaluate per trial. This implies that we need to find an optimized solution in $\leqslant$100 trials in order to meet our total latency requirements.

The combination of operational constraints, nonlinear slew dynamics, irregular GW sky maps, and nonuniform sensitivity across the FOV leads to a highly nonconvex and multimodal optimization problem where traditional gradient-based methods can easily get trapped in local optima. An example is shown in Figure 3. We employ a global optimization strategy with probabilistic sampling and trial pruning. We utilize the tree-structured Parzen estimator (TPE; J. Bergstra et al. 2011), which works by constructing probabilistic models to efficiently explore the search space, allowing it to focus on areas with higher potential for improvement while still exploring broadly. This probabilistic approach enables TPE to effectively guide the optimization even in complex, nonconvex spaces. We see this performance in Figure 5, where we compare TPE to a random (independent) sampler and a covariance matrix adaptation evolution strategy (CMA-ES) sampler on a wide range of input conditions on trajectory optimization. On average, TPE reaches better objective values after 100 trials, although all samplers show substantial variance. This makes TPE the best choice when considering both the speed of convergence and the quality of the solutions found within a fixed number of trials. Convergence is typically achieved within 100–150 trials, although pathological cases do exist. For the specific case depicted in Figure 3, a slow convergence means that the merger is not successfully observed when triggering a slew after 100 trials. For this case, the commanded final roll angle has greater relative "importance" compared to the R.A., decl. due to its strong impact on the maximum slew velocity and the starting position relative to the multimodal sky map. We note that additional sampler tuning, transfer learning from subsamples of the simulation set, and further optimization of the objective function are likely to improve performance, particularly for these edge cases.

In the following section, we utilize the TPE trajectory optimizer, which is allowed to maximize the number of trials performed under a time-out of 1 s, to match the performance requirements of the real-time early-warning response system.

6. Simulations of Swift Response to GW Early-warning Alerts and Projected Yields

With a response strategy in hand, we aim to faithfully reproduce the actions that Swift will take upon receiving an early-warning alert and study the expected yield. From the early-warning data release, 2737 (4131) [7775] injections have detections at 29 (32) [38] Hz. For each detection, we create an instantaneous GW alert and ingest it as we would in real time. For each early-warning alert, we project the sky maps onto the sky as seen by Swift at the time of the alert. In order to preserve the relationship between time, sky location, SNR, and localization precision, we retain the true time of the signal and calculate the actual orbital track, the time-resolved orientation and size of observing constraints and occultations, and the as-flown observing schedule of Swift. We compute the effective sensitivity of BAT to the true position of the BNS (known from the injection) and find that $\sim$17% of all mergers will serendipitously occur inside the BAT FOV, consistent with expectations from the size of the FOV (2.2 sr).[14] We also find

[14] Note that this implies that the Swift observing schedule is indeed completely uncorrelated with the IGWN antenna pattern.

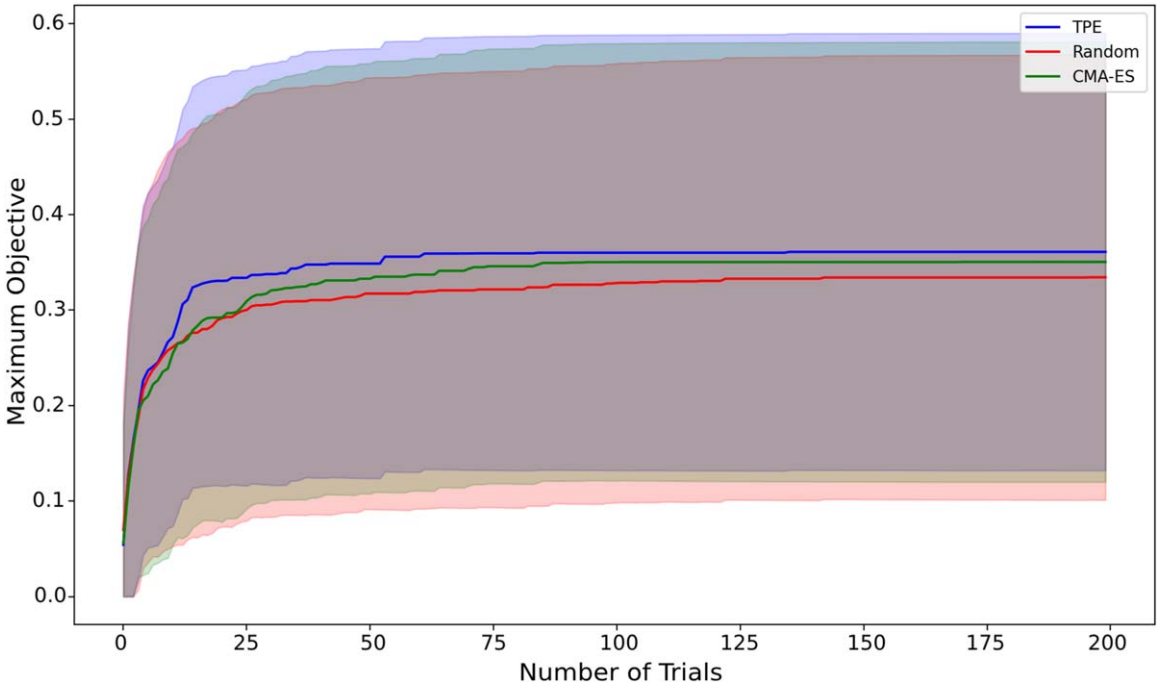

Figure 5. Comparison of best objective values achieved by the TPE, random, and CMA-ES samplers as a function of trials. The shaded areas represent the standard deviation across >1000 runs on different sky maps and starting conditions. TPE shows slightly superior overall performance and consistency, achieving the best results on average, although all samplers show significant performance variability across trajectory conditions.

that 32% will be occulted by the Earth at merger time, consistent with the fraction of sky blocked by Earth at Swift's orbital altitude. We then apply the trajectory optimizer described in Section 4 and obtain a target R.A., decl., roll to send to the spacecraft. We propagate the slew resulting from this target command and find the position and orientation of the BAT FOV on the sky at merger time for a range of slew start latencies from 0 to 60 s. As such, this latency fully encapsulates all in principle reducible latencies, including IGWN pipeline latencies and Swift command latencies, as will be discussed further in Section 7, but assumes fixed detector sensitivity and Swift ACS. We again compute the effective sensitivity of BAT to the true position of the BNS at merger time and record these results. See Figures 6 and 7 for representative samples from this simulation set.

We define our success as the true location of the BNS landing inside the BAT coded FOV, where arcminute-precision positions can be found for detected GRBs. The BAT coded FOV does not have uniform sensitivity, and we define the edge as the 10% coding fraction, although brighter bursts are regularly localized at lower coding. The BAT imaging effective area is ∼linear in coding fraction. The results for a response with 29 Hz sky maps are shown in Figure 8 (left), where the reverse cumulative fraction of the BAT coding fraction at the true BNS location is shown as a function of time between slew start and merger.

Using sky maps produced at 29 Hz, we expect an early-warning response to saturate our success rate at ∼50% of all mergers (a 3× relative improvement), even at zero latency and with an optimal trajectory. This is because 30% of the time, the true position will be blocked by the Earth, and an uncorrelated ∼35% (30%) of the time, we will look in the wrong direction, given the localization precision at 29 (32) Hz and the size of our FOV (see Figure 2). We note that the Earth cannot be made smaller, nor Swift raised in its orbit, so the only way to raise this success ceiling is to improve the GW detector sensitivity such that sky localizations are better constrained at low truncation frequency. When utilizing higher truncation frequency (less warning) sky maps, the precision improves, but the trade is for less warning time. We show that using the earliest sky map available is the correct choice given the Swift slew speed in the right panel of Figure 8.

Our results show that with zero latency between the GW signal arriving at 29 Hz and Swift beginning to slew, we can more than double the rate of mergers falling inside the BAT FOV and of prompt arcminute localizations. But there is no such thing as a zero-latency system. In Figure 8, we see the strong impact of latency on the success rate, and in the following section, we investigate the real-world latency of the various infrastructure between the wiggling LIGO mirrors and the Swift spacecraft ACS.

7. Latency Budget

Thus far, we have focused on irreducible, fixed latencies originating from the evolution of the GW signal given the current IGWN sensitivity, the size of the BAT FOV, and the ACS of Swift. However, there are significant nonintrinsic additional latency overheads associated with both.

7.1. IGWN

We have so far assumed instantaneous access to the information contained in the GW data at a given time/truncation frequency, ignoring the associated latencies of recording the data at the interferometers, calibration, data transfer to cluster, search pipeline filtering, parameter estimation (including sky location), and alert distribution. S. S. Chaudhary et al. (2024) characterize the latency of the GW alert products via a mock data challenge but do not include latency originating before the data are available for searches. However, they do note that the total

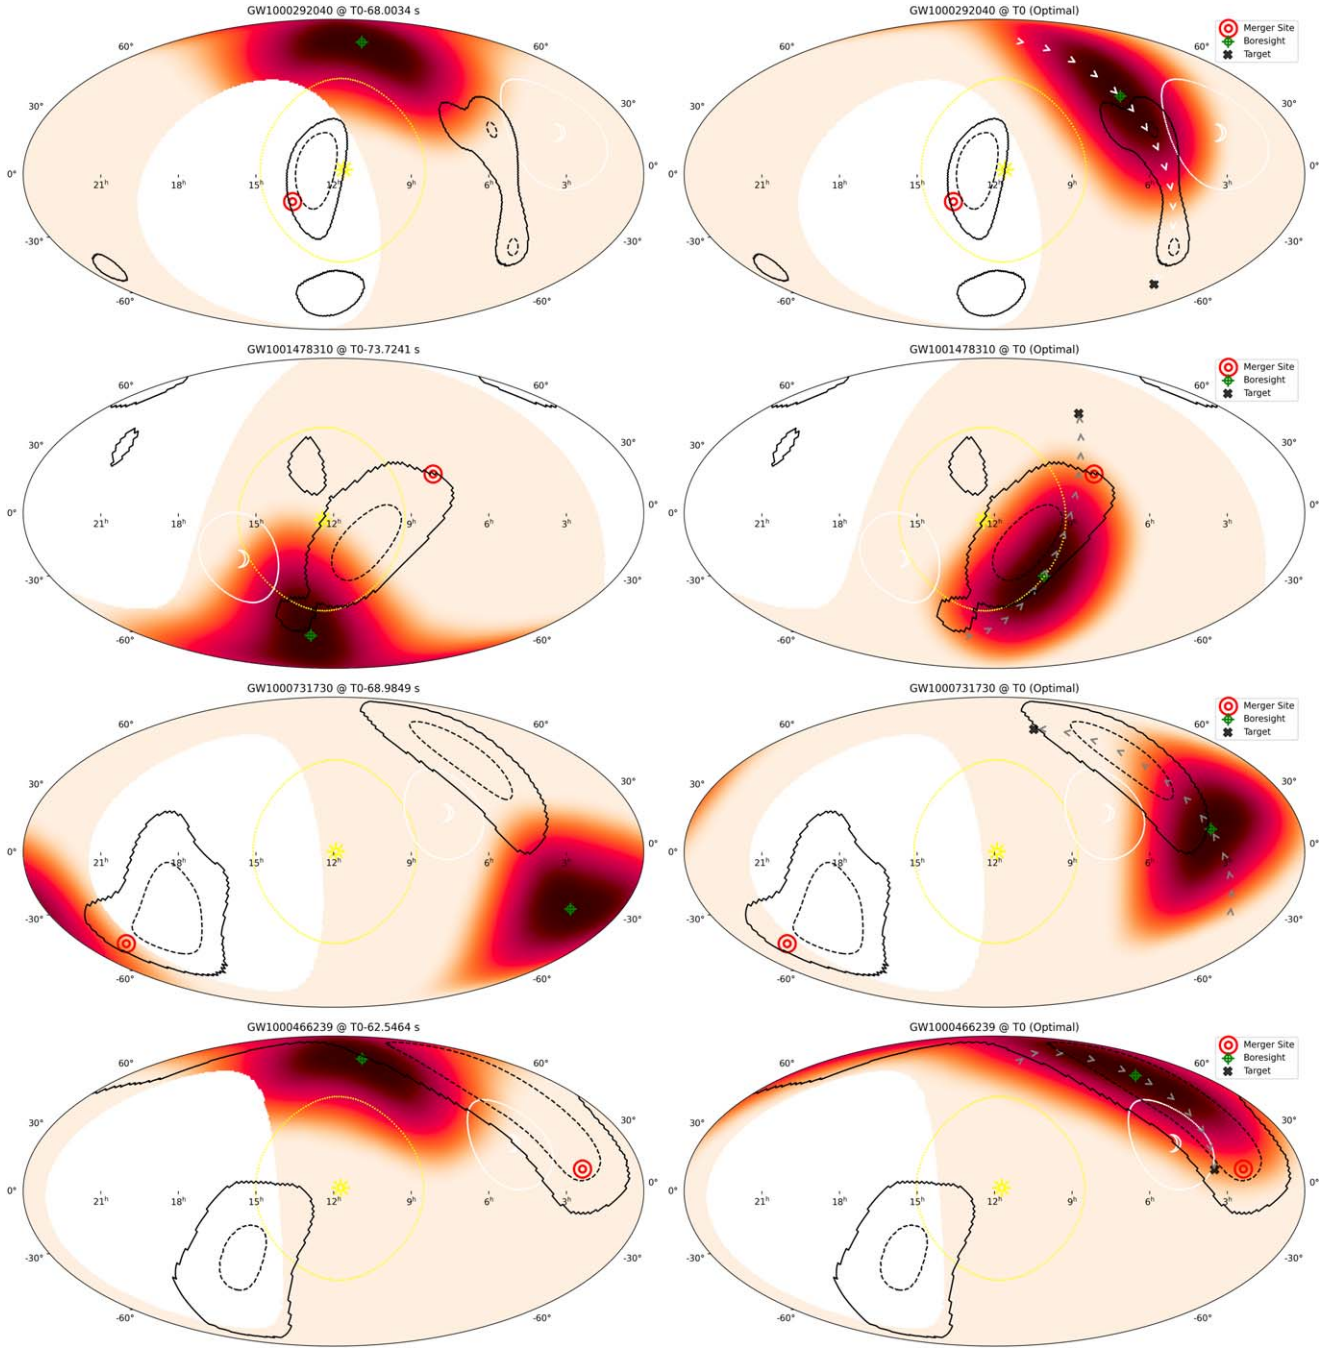

Figure 6. A sample of GW early-warning sky maps at a truncation frequency of 29 Hz (∼−60 s before merger). The left column shows BAT coverage at the alert time, and the right panel shows the optimal trajectory found and BAT coverage at merger time if the trajectory is executed. Note a tragic (rare) occurrence in the third row of slewing away from the correct position, not yet known. But note also marked successes in the second and fourth panels.

measured latency from real data, characterized as GW event time (merger) to alert production, has a median of 29.7 s in O4a. This is comprised of ∼5–10 s latency in data acquisition, calibration, and transfer to the computing cluster; ∼5–10 s for search pipelines to find an event and trigger; ∼1 s for sky map creation; and another ∼7–10 s for parameter estimation, annotation, and alert orchestration.

We repeat this end-to-end latency measurement for all of O4 (to date), measuring the latency directly as the difference between the merger time and received alert time. We purposefully screen out the early-warning alerts (all thus far retracted) so as to reliably measure the full-stack latency from GW signal in the detector to alert. This latency can then be safely added on top of any CBC analysis, including early warning. Figure 9 shows the latency distribution for both O4a and O4b. We measure a median latency of ∼34 s, higher than presented in S. S. Chaudhary et al. (2024). This latency overhead is comparable to the early warning achievable at current detector sensitivity. We highlight that the current pipeline latency mostly obviates the effectiveness of the early-warning alerts at providing actionable information before merger.

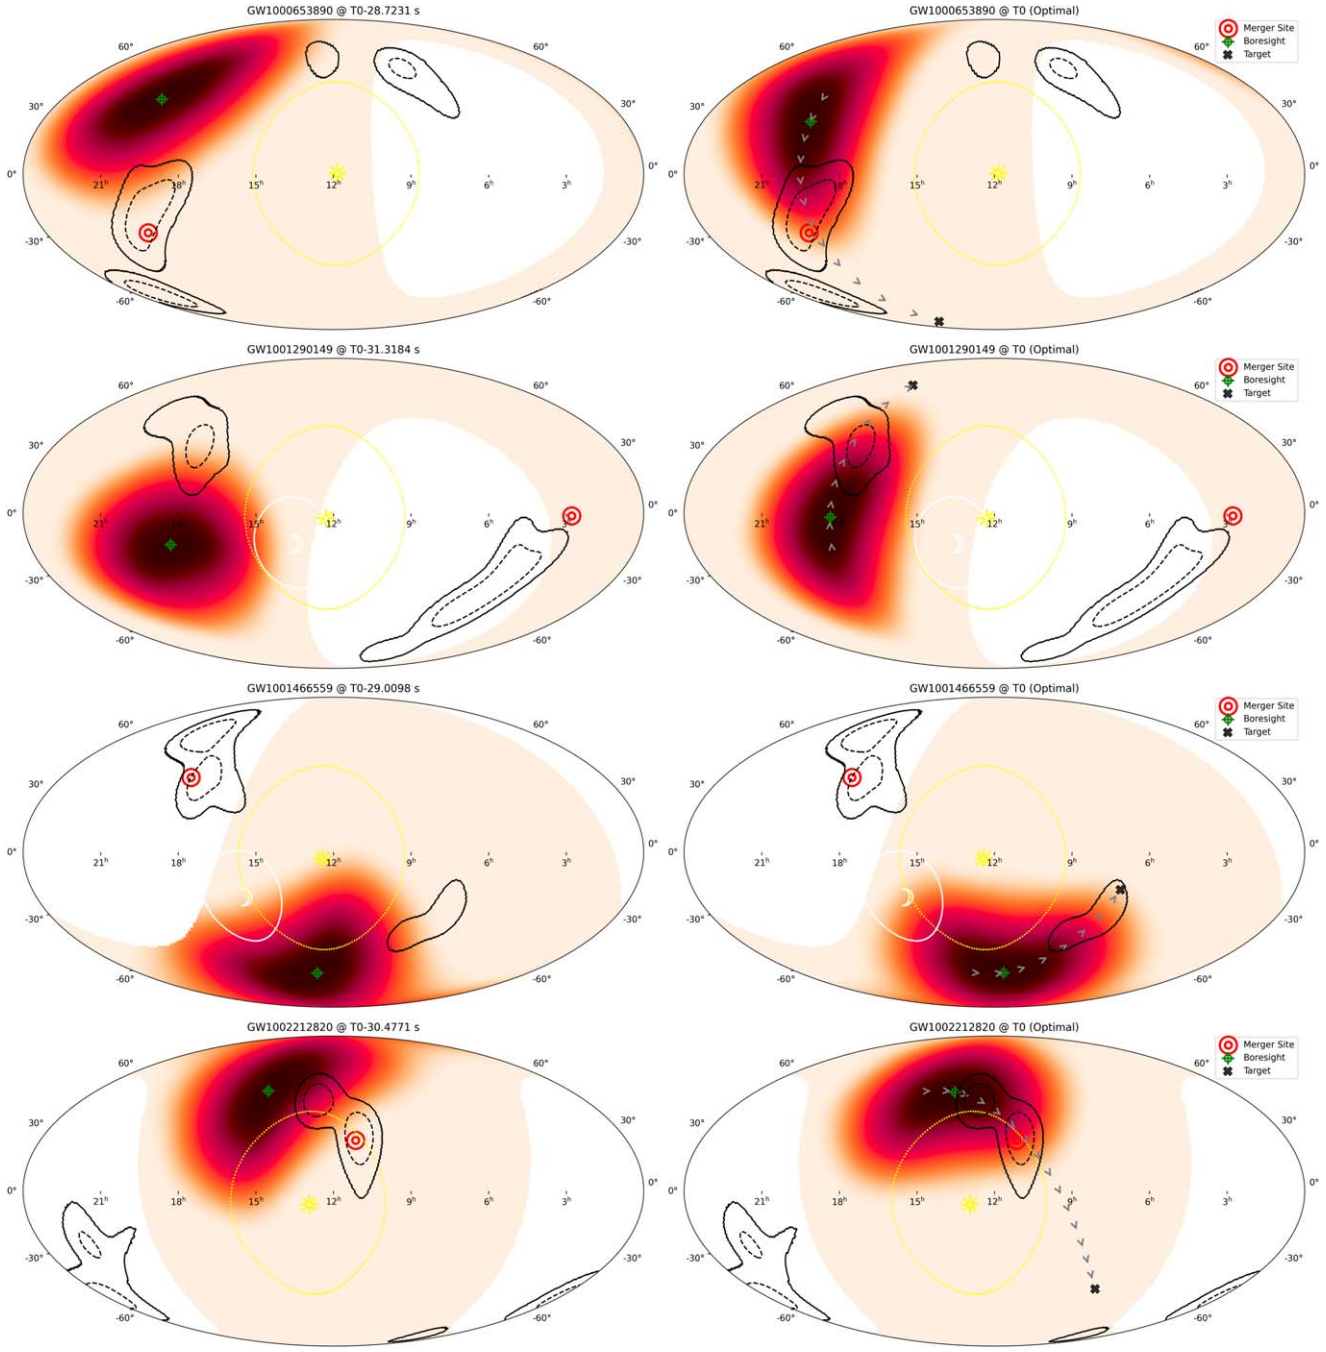

Figure 7. Same as Figure 6 but for a truncation frequency of 38 Hz. Note the successes in the first and fourth panels.

7.2. *Swift*

On the follow-up side, we have considered only the slew latencies of the Swift ACS, assuming instant slew start. We have so far ignored the ground processing of the alert, decision latency, and telecommand latency to the spacecraft. In A. Tohuvavohu et al. (2020), we presented a novel pipeline that reduced the on-demand command latency of Swift from 4 hr to 14 minutes, measured as the difference between alert receipt on the ground and command receipt on orbit. The majority of this 14 minute latency floor stemmed from an immutable characteristic of on-demand scheduling of the Tracking Data and Relay Satellite System (TDRSS), which does not permit new contacts scheduled autonomously to begin less than 10 minutes from request time to ensure adequate time for service configuration changes between different Space Network (SN) users (C. J. Roberts et al. 2021).

However, once the Swift Mission Operations Center (MOC) and the satellite are connected, the command latency through TDRSS should be very low. We can break this into a few fundamental components, tracing the journey of a command bit from the Swift MOC to Swift on orbit.

1. Terrestrial network delays, MOC to SN ground system.
2. Propagation delay through SN ground terminal.
3. Time of flight from SN ground terminal to TDRS satellite (space-to-ground link).

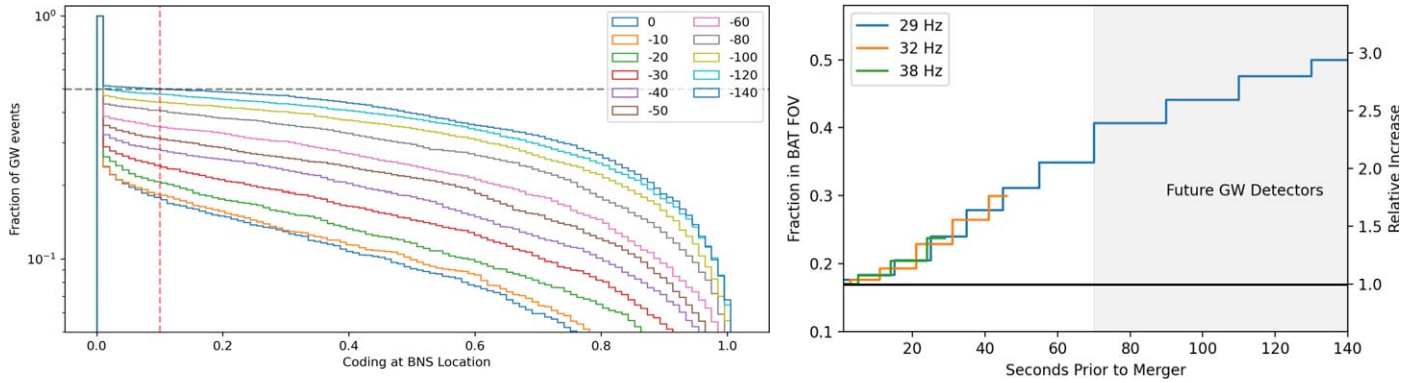

Figure 8. The reverse cumulative distribution of the BAT normalized sensitivity to the true BNS location using 29 Hz sky maps at various response latencies. We mark the saturation condition for 29 Hz at 50% of the mergers and the "edge" of the BAT FOV at 0.1. One can easily read off that an early-warning time of −60 s (pink) doubles the fraction of mergers inside the FOV. Note that the relative gain increases with coding, meaning that early-warning response not only maximizes the fraction of mergers in the FOV but also independently increases the average sensitivity to those already in the FOV. On the right, we present the fraction of mergers inside the BAT FOV as a function of warning time for sky maps derived at different truncation frequencies. We can easily see that, despite the improved localization precision at higher frequencies, acting on the earliest sky map is always better than waiting given Swift's finite slew speed. We also show capabilities assuming an early-warning time of >70 s, achievable only by upgraded GW detectors. For these cases, we assume the same sky maps as produced at 29 Hz by Advanced LIGO/Virgo.

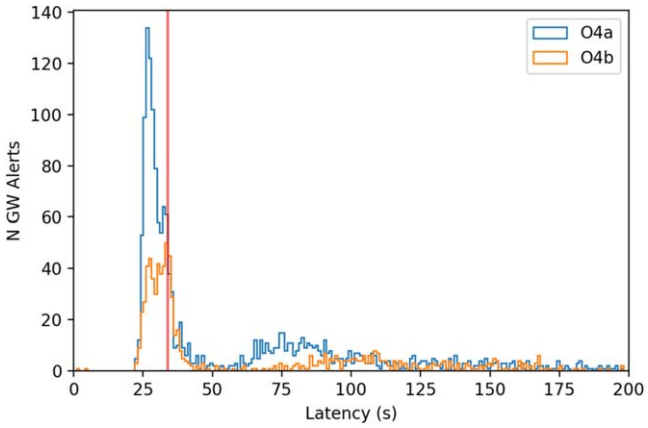

Figure 9. The latency of GW alerts during O4, measured as the difference between alert receipt and merger time. The red line denotes the median latency across all alerts in O4. Note that these latencies must be added on top of early-warning negative latencies, leading most "early-warning" alerts to be distributed after merger.

4. Propagation delay through the TDRS satellite.
5. Time of flight from TDRS satellite to Swift (space-to-space link).
6. Swift receiver waiting for the entire command telemetry frame at our data transfer rate.

The sum of these latencies is conservatively less than 2 s! However, achieving this would require Swift to already be in a scheduled contact at the unknown future time of a neutron star merger.

Motivated by the unique scientific opportunity presented by early-warning GW detection and the high instantaneous availability fraction (⩾85%) of TDRSS Multiple Access Forward service demonstrated by A. Tohuvavohu et al. (2020) and C. J. Roberts et al. (2021), the SN has enabled the Swift mission to preschedule semicontinuous contact service when available. After several test periods, this service began ongoing operation in spring 2023, now achieving an average availability of ∼80%. We call this capability "continuous commanding." Effectively, we are now always in touch with Swift.[15]

After the spacecraft has received the slew request, it determines if the requested slew is safe to begin immediately. This takes <1 s. If so, it notifies the instruments that a slew is about to begin and will wait up to 2 s for UVOT to reply with an acknowledgment before beginning to slew. As such, we should expect an absolute minimum of 3–5 s latency from command sent on the ground until slew start.

7.3. Swift Demonstration

In order to characterize the true performance of this new system and exemplify broader application to time-domain science, we demonstrate the Swift response to an external fast radio burst (FRB) trigger received from the Canadian Hydrogen Intensity Mapping Experiment. This FRB arrived on Earth at T0 = 2023-09-05T15:18:46.48 UTC (arrival time dedispersed to infinite frequency), and the VOEvent reporting this detection was received at the Swift MOC 46 s later, at $T_{\rm alert}$ = 15:19:33. Here we provide a subsequent timeline for Swift repointing:

$T_{\rm alert}$ + 31 s	TOO command sent from MOC to *Swift*
$T_{\rm alert}$ + 40 s	Slew start, 81° to go.
$T_{\rm alert}$ + 88 s	Target position enters BAT FOV.
$T_{\rm alert}$ + 149 s	XRT and UVOT begin recording photons from target position.

Note the 9 s delay between sending the command and Swift beginning to slew! Compare this to the previous best

[15] It would have been much easier if our space telescope were in geosynchronous orbit.

performance of on-demand response, at tens of minutes. In this case, the 31 s delay between alert received and command sent originates from a suboptimal ingestion and safety check procedure in the Swift ground system. Using the new procedure described in previous sections reduces the latency from receipt to command being sent to 1 s. Continuous commanding has since been integrated into general Swift operations, significantly increasing observatory versatility and efficiency. It has been used to send hundreds of TOO commands to date. The median delay between command sent and slew start is presently 10 s, and the minimum achieved is 6 s. We continue to investigate ways to further decrease this latency.

8. Conclusions

Exceptional scientific yield is possible with early-time (<1 minute) access to the postmerger environment. Prompt GRB detection and arcminute localization is presently the most promising avenue to achieve this, in the absence of decihertz GW detectors. Swift's superior sensitivity and precision positions are the best suited among current instruments but are balanced with an FOV that covers only one-sixth of the sky, significantly suppressing detection rates given the low rate of BNS mergers in the local Universe. Undaunted, we investigate the question of whether Swift can point at the location of a BNS before merger by utilizing early-warning GW detections.

We apply our Swift ACS model and trajectory optimizer to a large set of realistic simulations of early-warning GW detections and sky maps. Swift is indeed capable of exploiting early-warning GW alerts, slewing rapidly across the sky to put the GW location in the BAT FOV by merger time. With 60 s of early warning, this will more than double the rate of BNS mergers with prompt arcminute positions. While the localization precision increases with later warning time, for Swift and with the current IGWN sensitivity, responding to the earliest alert is the optimal strategy. Given the finite Swift slew speed, higher-latency alerts result in a significant reduction in the expected yield.

The IGWN is now sufficiently sensitive that the practical latencies between GW arrival in the interferometers and information reaching the Swift spacecraft now limit capabilities. This is a good thing, because in most cases cyberinfrastructure is easier to improve than GW interferometers or spacecraft already in orbit. However, we need to do it!

We demonstrate a new Swift operational capability whereby we can begin slewing the spacecraft within 10 s of an external TOO being received on the ground. With this innovation, the prompt arcminute detection rate is now fundamentally limited by the Swift slew speed (which cannot be improved) and the IGWN alert latency and software infrastructure (which can). However, the IGWN alert latency is presently ∼30 s, mostly obviating the practical utility of early-warning alerts.

Several authors have outlined prospects for improved early-warning capability with upgraded sensitivity of future detectors (A. H. Nitz et al. 2020; R. Magee & S. Borhanian 2022), as well as novel statistical approaches (C. Chatterjee & L. Wen 2023; E. Marx et al. 2024). We point out that in the current IGWN implementation, the search pipelines themselves only contribute ∼one-third of the total internal latency of the system, and coequal focus must be placed on the latency of data transfer, calibration (or not), and alert orchestration. We recommend a significant focus on reduction of IGWN internal latencies to allow for this extremely compelling scientific opportunity.

In addition to GW early-warning response with BAT, "continuous commanding" opens new opportunities for science with Swift. Upcoming surveys like CHORD (K. Vanderlinde et al. 2019), the Argus Array (H. Corbett et al. 2022; N. M. Law et al. 2022), LAST (E. O. Ofek et al. 2023), and ULTRASAT (Y. Shvartzvald et al. 2024) will survey the sky at very high cadence and distribute transient alerts at high rates and low latency. We anticipate significant demand from transients discovered by these surveys. This system is already being utilized by a community program for characterization of orphan supernovae discovered by the Zwicky Transient Facility, and the Swift Team invites the community to consider and propose novel scientific applications of ultra-low-latency UV, X-ray, and gamma-ray observations.

Acknowledgments

We thank the Swift Flight Operations Team for their invaluable assistance bringing continuous commanding online and for consistent safe stewardship of Swift. We further thank Eric Siskind, Jamie Rollins, Alex Nitz, Chad Hanna, Reed Essick, Chris Matzner, Maria Drout, and many others for enlightening discussions.

Facilities: Swift, LIGO, EGO:Virgo.

Software: ligo.skymap, swift-tools, optuna (T. Akiba et al. 2019), astropy (Astropy Collaboration et al. 2013; A. M. Price-Whelan et al. 2018).

ORCID iDs

Aaron Tohuvavohu https://orcid.org/0000-0002-2810-8764
Jamie A. Kennea https://orcid.org/0000-0002-6745-4790
Christopher J. Roberts https://orcid.org/0000-0002-8866-7891
James DeLaunay https://orcid.org/0000-0001-5229-1995
Samuele Ronchini https://orcid.org/0000-0003-0020-687X
S. Bradley Cenko https://orcid.org/0000-0003-1673-970X
Becca Ewing https://orcid.org/0000-0001-9178-5744
Ryan Magee https://orcid.org/0000-0001-9769-531X
Cody Messick https://orcid.org/0000-0002-8230-3309
Surabhi Sachdev https://orcid.org/0000-0002-0525-2317
Leo P. Singer https://orcid.org/0000-0001-9898-5597

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, PhRvL, 119, 161101
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, ApJL, 848, L12
Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, PhRvX, 13, 041039
Adams, T., Buskulic, D., Germain, V., et al. 2016, CQGra, 33, 175012
Ahumada, T., Anand, S., Coughlin, M. W., et al. 2024, arXiv:2405.12403
Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, arXiv:1907.10902
Alexander, K. D., Berger, E., Fong, W., et al. 2017, ApJL, 848, L21
Arcavi, I. 2018, ApJL, 855, L23
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Banerjee, B., Oganesyan, G., Branchesi, M., et al. 2023, A&A, 678, A126
Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, SSRv, 120, 143
Beloborodov, A. M. 2021, ApJ, 921, 92
Beniamini, P., Duran, R. B., Petropoulou, M., & Giannios, D. 2020, ApJL, 895, L33
Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B. 2011, in Advances in Neural Information Processing Systems, 24, ed. J. Shawe-Taylor et al. (New York: Curran Associates, Inc.), L33
Burns, E. 2020, LRR, 23, 4
Cannon, K., Cariou, R., Chapman, A., et al. 2012, ApJ, 748, 136
Chatterjee, C., & Wen, L. 2023, ApJ, 959, 76

Chaudhary, S. S., Toivonen, A., Waratkar, G., et al. 2024, PNAS, 121, e2316474121
Combi, L., & Siegel, D. M. 2023, PhRvL, 131, 231402
Copete, A. J. 2012, PhD thesis, Harvard Univ., Massachusetts
Corbett, H., Vasquez Soto, A., Machia, L., et al. 2022, Proc. SPIE, 12189, 1218910
D'Avanzo, P. 2015, JHEAp, 7, 73
Dean, C., Fernández, R., & Metzger, B. D. 2021, ApJ, 921, 161
DeLaunay, J., & Tohuvavohu, A. 2022, ApJ, 941, 169
DeLaunay, J., Tohuvavohu, A., Kennea, J. A., & Raman, G. 2022, GCN, 32023, 1
DeLaunay, J., Tohuvavohu, A., Kennea, J. A., et al. 2024, GCN, 36672, 1
Fletcher, C., Wood, J., Hamburg, R., et al. 2024, ApJ, 964, 149
Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, ApJ, 611, 1005
Goldstein, A., Veres, P., Burns, E., et al. 2017, ApJL, 848, L14
Gompertz, B. P., O'Brien, P. T., & Wynn, G. A. 2014, MNRAS, 438, 240
Haggard, D., Nynka, M., Ruan, J. J., et al. 2017, ApJL, 848, L25
Hallinan, G., Corsi, A., Mooley, K. P., et al. 2017, Sci, 358, 1579
Hansen, B. M. S., & Lyutikov, M. 2001, MNRAS, 322, 695
Herner, K., Annis, J., Brout, D., et al. 2020, A&C, 33, 100425
Kaneko, Y., Bostancı, Z. F., Göğüş, E., & Lin, L. 2015, MNRAS, 452, 824
Kovalam, M., Kaium Patwary, M. A., Sreekumar, A. K., et al. 2022, ApJL, 927, L9
Kulkarni, S. R. 2005, arXiv:astro-ph/0510256
Law, N. M., Corbett, H., Galliher, N. W., et al. 2022, PASP, 134, 035003
Levan, A. J., Gompertz, B. P., Salafia, O. S., et al. 2024, Natur, 626, 737
Lü, H.-J., Yuan, Y., Lan, L., et al. 2020, ApJL, 898, L6
Luongo, O., & Muccino, M. 2021, Galax, 9, 77
Lyutikov, M. 2024, arXiv:2402.16504
Magee, R., & Borhanian, S. 2022, ApJ, 935, 139
Magee, R., Chatterjee, D., Singer, L. P., et al. 2021, ApJL, 910, L21
Margutti, R., Berger, E., Fong, W., et al. 2017, ApJL, 848, L20
Margutti, R., & Chornock, R. 2021, ARA&A, 59, 155
Marx, E., Benoit, W., Gunny, A., et al. 2024, arXiv:2403.18661
Metzger, B. D., Bauswein, A., Goriely, S., & Kasen, D. 2015, MNRAS, 446, 1115
Metzger, B. D., & Zivancev, C. 2016, MNRAS, 461, 4435
Most, E. R., & Philippov, A. A. 2022, MNRAS, 515, 2710
Neill, D., Newton, W. G., & Tsang, D. 2021, MNRAS, 504, 1129
Nitz, A. H., Schäfer, M., & Dal Canton, T. 2020, ApJL, 902, L29
Ofek, E. O., Ben-Ami, S., Polishook, D., et al. 2023, PASP, 135, 065001
Oganesyan, G., Karpov, S., Salafia, O. S., et al. 2023, NatAs, 7, 843
Peters, P. C., & Mathews, J. 1963, PhRv, 131, 435
Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, AJ, 156, 123
Raman, G., Ronchini, S., Delaunay, J., et al. 2024, arXiv:2407.12867
Rastinejad, J. C., Gompertz, B. P., Levan, A. J., et al. 2022, Natur, 612, 223
Roberts, C. J., Burke, J. C., Benson, M. J., et al. 2021, JAIS, 18, 333
Ronchini, S., Bala, S., Wood, J., et al. 2024, ApJL, 970, L20
Sachdev, S., Magee, R., Hanna, C., et al. 2020, ApJL, 905, L25
Savchenko, V., Ferrigno, C., Kuulkers, E., et al. 2017, ApJL, 848, L15
Shvartzvald, Y., Waxman, E., Gal-Yam, A., et al. 2024, ApJ, 964, 74
Singer, L. P., & Price, L. R. 2016, PhRvD, 93, 024013
Tohuvavohu, A., DeLaunay, J., Kennea, J. A., & Raman, G. 2021a, GCN, 30393, 1
Tohuvavohu, A., DeLaunay, J., Kennea, J. A., & Raman, G. 2021b, GCN, 30325, 1
Tohuvavohu, A., Kennea, J. A., DeLaunay, J., et al. 2020, ApJ, 900, 35
Troja, E., Fryer, C. L., O'Connor, B., et al. 2022, Natur, 612, 228
Troja, E., Piro, L., van Eerten, H., et al. 2017, Natur, 551, 71
Tsang, D. 2013, ApJ, 777, 103
Vanderlinde, K., Liu, A., Gaensler, B., et al. 2019, in Canadian Long Range Plan for Astronomy and Astrophysics White Papers, 2020, 28
Villar, V. A., Guillochon, J., Berger, E., et al. 2017, ApJL, 851, L21
Wu, G.-L., Yu, Y.-W., & Li, S.-Z. 2022, Univ, 8, 633
Yang, Y.-H., et al. 2024, Natur, 626, 742