



Low-latency gravitational wave alert products and their performance at the time of the fourth LIGO-Virgo-KAGRA observing run

Sushant Sharma Chaudhary^{a,1}, Andrew Toivonen^{b,1}, Gaurav Waratkar^c, Geoffrey Mo^{d,e}, Deep Chatterjee^{d,e}, Sarah Antier^f, Patrick Brockill^g, Michael W. Coughlin^b, Reed Essick^{h,i,j}, Shaon Ghosh^k, Soichiro Morisaki^l, Pratyusava Baral^g, Amanda Baylor^g, Naresh Adhikari^g, Patrick Brady^g, Gareth Cabourn Davies^m, Tito Dal Cantonⁿ, Marco Cavaglia^a, Jolien Creighton^o, Sunil Choudhary^{p,q}, Yu-Kuang Chu^g, Patrick Clearwater^{p,q}, Luke Davis^{p,q}, Thomas Dent^r, Marco Drago^s, Becca Ewing^{t,u}, Patrick Godwin^v, Weichangfeng Guo^{p,q}, Chad Hanna^{u,w,x}, Rachael Huxford^{t,u}, Ian Harry^y, Erik Katsavounidis^{d,e}, Manoj Kovalam^{p,q}, Alvin K.Y. Li^v, Ryan Magee^v, Ethan Marx^{d,e}, Duncan Meacher^g, Cody Messick^g, Xan Morice-Atkinsonⁿ, Alexander Pace^u, Roberto De Pietri^{z,aa}, Brandon Piotrkowski^g, Soumen Roy^{bb,cc}, Surabhi Sachdev^{g,dd}, Leo P. Singer^{ee,ff}, Divya Singh^{h,u}, Marek Szczepanczyk^{gg}, Daniel Tang^{p,q}, Max Trevor^{hh}, Leo Tsukada^{tu}, Verónica Villa-Ortega^f, Linqing Wen^{p,q}, and Daniel Wysocki^g

Edited by Neta Bahcall, Princeton University, Princeton, NJ; received September 25, 2023; accepted March 16, 2024

Multimessenger searches for binary neutron star (BNS) and neutron star-black hole (NSBH) mergers are currently one of the most exciting areas of astronomy. The search for joint electromagnetic and neutrino counterparts to gravitational wave (GW)s has resumed with ALIGO's, AdVirgo's and KAGRA's fourth observing run (O4). To support this effort, public semiautomated data products are sent in near real-time and include localization and source properties to guide complementary observations. In preparation for O4, we have conducted a study using a simulated population of compact binaries and a mock data challenge (MDC) in the form of a real-time replay to optimize and profile the software infrastructure and scientific deliverables. End-to-end performance was tested, including data ingestion, running online search pipelines, performing annotations, and issuing alerts to the astrophysics community. We present an overview of the low-latency infrastructure and the performance of the data products that are now being released during O4 based on the MDC. We report the expected median latency for the preliminary alert of full bandwidth searches (29.5 s) and show consistency and accuracy of released data products using the MDC. We report the expected median latency for triggers from early warning searches (-3.1 s), which are new in O4 and target neutron star mergers during inspiral phase. This paper provides a performance overview for LIGO-Virgo-KAGRA (LVK) low-latency alert infrastructure and data products using the MDC and serves as a useful reference for the interpretation of O4 detections.

gravitational waves | multimessenger astronomy | compact binary mergers

As of May 24 2023,* aLIGO's, AdVirgo's and KAGRA's fourth observing run (O4) is underway, following a series of observing runs, which have reported the detection of the first binary black hole (BBH) in ALIGO's first observing run (O1) (1), the detection of the first binary neutron star (BNS) merger (2) and associated electromagnetic counterparts AT2017gfo (3–5) and GRB170817A (6–8) in the ALIGO's and AdVirgo's second observing run (O2), and neutron star-black hole (NSBH) (9) in the ALIGO's and AdVirgo's third observing run (O3). Focusing on neutron star (NS) mergers, there are a variety of science cases for their multimessenger counterpart searches and detections, including measurements of the NS equation of state (EoS) (10–20), the Hubble constant (19, 21–24), and r -process nucleosynthesis (4, 25–31).

The LIGO-Virgo-KAGRA (LVK)'s real-time alert infrastructure depends on several components. Broadly this includes low-latency data calibration and transfer, running modeled and unmodeled online searches, and maintaining the state of events in GRAvitational-wave candidate event database (GraceDB) following the discovery. In addition to GraceDB,[†] which serves as both the database and as an internal and external web view, the alert infrastructure includes igwn-alert,[‡] an internal messaging

Significance

We present the low-latency alert infrastructure performance and data products based on an end-to-end MDC designed to prepare for O4. We aim to provide a useful reference for the broader astronomy community following up on GW alerts, and aid in interpreting the data products made public by the LVK. The MDC covers a dense sample of compact binary coalescence (CBC) events similar to those we hope to observe during O4, demonstrating up-to-date results for the public data products. We strive to maximize the efficacy of multimessenger follow-up searches by helping astronomers utilize our data products to their full advantage. The description of the alert infrastructure may be used as a reference for astronomical experiments targeting similar timescales.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: sscwrk@mst.edu or toivo032@umn.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2316474121/-/DCSupplemental>.

Published April 23, 2024.

*<https://observing.docs.ligo.org/plan>.

[†]<https://gracedb.ligo.org/>.

[‡]<https://igwn-alert.readthedocs.io>.




Fig. 1. Task flow of low-latency alert infrastructure. The process begins with search pipeline trigger(s) on a candidate in the GW datastream, which are passed through data quality checks and compiled into a superevent. If the preferred event from the superevent passed the significant FAR cut, a preliminary alert is sent out to the public.

system to communicate the state of events, and GWCelery^{§,¶} a task queue, to cluster, annotate, and orchestrate the events, as well as publish public alerts[#] for the community to subscribe to. Fig. 1 shows the task flow of the low-latency alert infrastructure (LLAI) for candidate events. The current LLAI is a significantly upgraded version of the infrastructure used earlier, described in ref. 32, and used in the more recently reported early-warning system, reported in ref. 33. The primary changes compared to previous observing run are the transition from a XMPP based pubsub system internally to Kafka-based messaging provided by the SCiMMA broker, removal of timeouts at various places of synchronization and instead relying on labels on GraceDB to keep track of the state of the superevent, and reconfiguring the snapshotting configuration for Redis database to strike a balance between fault tolerance and avoid filling disk quota. Aside from this, upgrades to software versions of the dependencies and running computationally expensive resources on specific pool of modern hardware contributed to an improvement compared to previous observing run. The LVK Alert User Guide^{||} constitutes a living document where information and updates about this system are regularly communicated to the broader community. Further discussion of GraceDB, igwn-alert, and GWCelery is provided in *SI Appendix*.

To prepare for O4 and demonstrate performance across a variety of software and alert system improvements, we carried out a mock data challenge (MDC). This MDC constitutes a testing environment for the LLAI to prepare for O4, producing repeated sets of 40 d of data from O3, with associated simulations of CBCs to stress test the system. While the rates of simulated events (see Section E for a description of the dataset) were much higher than that expected for O4, this high rate was designed to test the various components of searches, the alert system and

the scientific deliverables before heading into O4; these include, for example, tests of the detection efficiency of online real-time low-latency searches, the rapid estimation of the binary system properties, and their associated sky localizations.

In this paper, we describe the details of the alert system for O4 and its performance based on this MDC. In addition, we provide an overview of the detection performance of real-time searches, along with the consistency and accuracy of alert data products. Sections A–D provide an overview of the LLAI and the scientific data products reported, while Section E describes the properties of the MDC, including the motivations for the choices made. Section 1 reports the properties of the LLAI as of the beginning of O4, as measured by the MDC, and Section 2 describes the conclusions and prospects for future development.

A. Searches. Low-latency GW searches consist of two categories: “modeled” CBC (34) and “unmodeled” Burst (35) searches. Modeled CBC searches target BNS, NSBH, or BBH; unmodeled searches look for signals with generic morphologies from a wide variety of astrophysical sources like core-collapse of massive stars, magnetar star-quakes, and other sources, in addition to compact binary mergers (35, 36). For the purpose of this MDC analysis, we focus on CBC searches, but also report latencies of injections found by Burst pipelines. CBC searches can be categorized as early warning, referring to premerger searches (37, 38), or full bandwidth, referring to postmerger, based on how the search truncates their templates. Each search produces candidate GW triggers and assigns them ranking statistic values and false alarm rates (FARs); the FAR for a trigger in a given search pipeline is defined as the expected rate of triggers due to detector noise, in that pipeline, with equal or higher ranking. Each search pipeline has different and independent methods of generating and ranking triggers and estimating the noise background and thus the false alarm rate (FAR); for details see refs. 39–45. In addition, the probability of astrophysical origin, p_{astro} , for a trigger is calculated for CBC searches, which is briefly introduced in Section C.

[§] <https://git.ligo.org/emfollow/gwcelery>.

[¶] <https://rtd.igwn.org/projects/gwcelery/en/latest/>.

[#] For example, those hosted by SCiMMA and NASA.

^{||} <https://emfollow.docs.ligo.org/userguide/>.

The MDC data were analyzed by four CBC searches GStreamer LIGO scientific collaboration algorithm library (GstLAL) (39, 46, 47), multi-band template analysis (MBTA) (40), Python based compact binary coalescence search software (PyCBC) (44, 48), summed parallel infinite impulse response (SPIIR) (38, 49), two burst searches coherent waveburst (cWB) (50–52), and Omicron LALInferenceBurst (oLIB) (53), and an external coincidences search rapid, on-source VEvent coincident monitor (RAVEN) (45, 54, 55). A brief description of these pipelines is given in [SI Appendix](#).

B. Selection of the Public GW Event Candidate. When a potential GW signal appears in the detector, the low-latency search pipelines analyze the signal and produce event candidates. Each pipeline can report multiple candidates for a single GW signal. The event candidates, reported within a specific time window (1 s around coalescence time for CBC searches, and 1 s around trigger time for burst searches), are collected and grouped as a *superevent*. In the collection, one GW event candidate is identified as the preferred event [defined as the event with the highest network signal-to-noise ratio (SNR) for CBC pipelines whereas lowest FAR for burst pipelines], and its properties and data products are prepared for release to the public, which include the merger time, FAR, sky localization, and classification (Section C).

Alerts are released publicly when a FAR passes the public alert threshold, currently $\text{FAR} \leq 1.6 \times 10^{-4}$ Hz (14 per day). An alert is labeled as significant when a CBC alert passes a FAR threshold of $\text{FAR} \leq 3.9 \times 10^{-7}$ Hz (one per mo) or when an unmodeled burst alert passes a FAR threshold of $\text{FAR} \leq 3.2 \times 10^{-8}$ Hz (one per year). Since multiple CBC and burst searches run in low-latency, to account for the trials factor from these different searches with statistically independent false alarms, events from CBC searches are labeled significant when a FAR passes a threshold of $\text{FAR} \leq 7.7 \times 10^{-8}$ Hz (one per 5 mo), whereas events from burst target searches require a $\text{FAR} \leq 7.9 \times 10^{-9}$ Hz (one per 4 y). In the MDC, however, CBC trials factor of 6 was used due to which CBC events were labeled significant when a FAR passes a threshold of $\text{FAR} \leq 6.4 \times 10^{-8}$ Hz (one per 6 mo). Considering the trials factor accounting for all searches, the public alert threshold is $\text{FAR} \leq 2.3 \times 10^{-5}$ Hz (two per day). Alerts that meet the public alert threshold but not the significant threshold are labeled low significance. RAVEN only uses the significant FAR thresholds when assessing its joint FAR for publication, with additional trials factors to compensate for listening to multiple GW pipelines. The LVK Alert User Guide** should be referenced for up-to-date information on trials factors during observing runs.

C. Alert Contents. Public alerts are sent in order to inform the greater astronomical community of GW events, enabling multimessenger follow-up of these events. These alerts are distributed both via GCN and the Scalable Cyberinfrastructure to support multi-messenger astrophysics†† (SCiMMA) project in order to reach maximum consumers through two broad bases of subscribers. The alerts come in two types: notices, which are machine-readable and come in a variety of formats, and GCN circulars, which are human-readable.

There are five types of notices that may be sent out for a candidate event: Early Warning, Preliminary, Initial, Update,

and Retraction. Early Warning Notices arise from dedicated premerger search pipelines, potentially enabling the release of alerts seconds before merger (33). A first Preliminary Notice is sent out when an event candidate of a superevent exceeds the public FAR threshold. Following a timeout, the preferred event is determined and a second Preliminary Notice is then issued (even if the preferred event candidate remains unchanged). Both Early Warning and Preliminary Notices are sent out if the candidate passes automatic data quality checks (56). These data quality checks are carried out by the Data Quality Report framework‡‡, and include checks for terrestrial noise and stationarity of the data, among others. In certain cases, such as when manual data quality checks yield suspicions on the astrophysical nature of the candidate, a Retraction Notice may be sent. If, however, the Early Warning or Preliminary Notice passes human vetting, then an Initial Notice is sent out accompanied by a GCN Circular to announce the detection. The final type of notice, an Update Notice, is used to send out improved estimates of alert contents based on parameter estimation when they become available. Included in each alert is an estimate of the event's probability of astrophysical origin, or p_{astro} . Both Initial and Update Notices are accompanied by human-readable GCN Circulars. This is broken up into four categories that sum to 1 by definition: P(BNS), P(NSBH), P(BBH), and P(Terrestrial), where the mass boundary between NS and Black hole (BH) is set at $3 M_{\odot}$. The specifics of individual pipeline p_{astro} calculations can be found in [SI Appendix, Text](#). If the superevent is coincident with a GCN candidate, the various data products concerning the joint candidate are included, such as the time delay, joint FAR,§§ and combined sky map if applicable.

C.1. Sky localization. One of the key data products to enable multimessenger follow-up is the rapid inference of the sky localization from GW observations. This sky localization consists of the posterior probability distribution of the source location in the sky. The sky localization, mapped either over a 2D map of right ascension and declination, or a 3D volume which also includes a distance estimate, is known as a “sky map.” Sky localization (and parameter estimation more generally) is conducted in multiple stages once a candidate is identified.

For CBC sources, BAYESian TriAngulation and rapid localization (BAYESTAR), a rapid sky localization algorithm (57), is used to generate sky maps, and may be updated by Bilby (Section D), a python-based parameter estimation pipeline that uses stochastic sampling methods (58, 59). Sky maps from BAYESTAR are released with Preliminary Notices and sky maps from Bilby are released in Update Notices. Additionally, cWB also generates localizations for burst event candidates (60).

The sky map is stored as flexible image transport system (FITS) file using the hierarchical equal area isoLatitude pixelization (HEALPix) (61) framework in the multi-order coverage (MOC) representation (62); flattened versions at a fixed HEALPix grid size are also available for superevents. MOC sky maps use adaptive division of the HEALPix grid, focusing areas of highest resolution on regions of highest probability with minimal information loss. Sky maps are made available both through the distributed alert as well as uploaded on GraceDB, where they are available for direct download. If there is a coincidence with a GRB candidate that has a sky localization, we compute the overlap integral with the GW sky map information. This weighted sky map is then included in the alert. The technique of combining sky maps from

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

†† <https://scimma.org/>.

‡‡ <https://docs.ligo.org/detchar/data-quality-report/>.

§§ <https://ligo-raven.readthedocs.io>.

two independent datasets is laid out for the first time in ref. 54, under the signal hypothesis of the Bayesian framework, and is presented in an accessible manner in ref. 63.

C.2. EM-bright. EM-bright⁴⁴ is a pipeline designed to assess whether a GW candidate is capable of producing an electromagnetic counterpart (64). A rapid assessment of EM-bright properties, HasNS and HasRemnant, is essential to trigger target of opportunity (ToO) follow-up by ground and space-based observatories. In this regard, HasNS and HasRemnant quantities are reported as a part of the automated and update discovery notices. The HasNS is the probability of the binary having a NS component, while HasRemnant is the probability of the merger leaving remnant matter postmerger in the form of dynamical or tidal ejecta.

The exact nature of EM emission from the merger is complex and depends on several factors like the properties of the ejecta, the NS EoS, and the BH mass and spin. Detailed analyses are required to assess the accurate properties of EM counterparts (see, for example, ref. 65 for a review). These are, however, impractical in a real-time setting. Aside from theoretical uncertainties, measurement uncertainties predominantly affect the assessment of EM-brightness in real-time. Note that the only real-time data products available from match filter CBC searches are the template parameters that maximize the likelihood of detection, and these values are internal. Bayesian parameter estimation from computationally cheap waveform models may be available in $\sim$ hours, as discussed later, but it is not available in the seconds after a trigger is registered. Hence, inference from the template parameters and real-time detection statistics is the feasible solution.

To this end, ref. 64 showed the application of supervised machine learning trained on a feature space involving the template parameters and detection statistics to make this inference. Training is done using large-scale simulation campaigns where the ground truth and the recovery of search pipelines are registered. The ground truth is labeled based on its intrinsic source-frame mass as having a NS component, or both mass and spin components as leaving remnant matter behind postmerger, based on a phenomenological fit to numerical relativity simulations by ref. 66. The NS EoS plays a crucial role in the labeling as stiffer EoS favor tidal disruption, and therefore prefer larger ejecta masses. While in ref. 64 a single, stiff NS EoS was used on conservative grounds, here we extend the analysis to multiple EoSs, and reweight the score based on Bayes factors computed against GW170817 (67) tidal deformability measurements for several NS EoSs presented in ref. 68. The score presented is therefore marginalized over several EoSs.

In addition to HasNS and HasRemnant, a new quantity HasMassGap, the probability that at least one component of the binary merger is in the lower mass-gap region, i.e., source-frame mass between $3 M_{\odot}$ to $5 M_{\odot}$ is computed. The technique used in computing HasMassGap is similar to the original EM-bright quantities, except the labeling is different and does not involve the knowledge of the NS EoS. The values reported for HasNS and HasRemnant use a nearest-neighbor classifier algorithm, while that used for HasMassGap use a random-forest classifier algorithm.

Similar to the sky maps, these quantities are updated from online parameter estimation samples, which are made publicly available $\sim$ hours after discovery. The parameter estimation samples allow for these quantities to be computed directly.

D. Low-latency Parameter Estimation. CBC signal candidates labeled as significant (B) are further investigated via automated Bayesian parameter estimation analysis with the Bilby library. It employs the nested sampling technique implemented in the Dynesty library (69) to explore the full parameter space of masses and spins, producing accurate inference results immune to biases included in the point estimates of masses and spins from search pipelines. It also takes into account uncertainties in detector calibration and marginalizes the posterior probability distribution over them.

To accelerate the analysis, we employ the reduced order quadrature (ROQ) technique (70–72), which approximates gravitational waveform with ROQ basis elements to reduce the computational cost of likelihood evaluations. The ROQ basis elements employed in the automated parameter estimation of O4 are presented in ref. 73.

For BNS candidates, the analysis assumes that dimensionless spins have norms less than 0.05 and are aligned with the orbital angular momentum, and employs the IMRPhenomD waveform approximant (74, 75) to recover the observed signals. With the acceleration technique, the sampling completes typically in less than 10 min. The actual time from detection to upload of results from Bilby is a few tens of minutes since this analysis starts around 5 min after signal detection, and preparing input data and postprocessing results take several minutes. Currently, the output of this analysis is not automatically made public but manually sent after it passes human vetting. Hence it is sent at the earliest when an Initial Notice is sent, and the actual latency of the update is higher than the latency of upload of the results.

In addition to this automated analysis, more costly manual analyses incorporating general spin configurations and tidal deformation of colliding objects may follow, depending on the significance of the signal. For candidates with higher masses, the automated analysis takes into account general spin configurations, and employs IMRPhenomXPHM (76) if its ROQ basis elements are available in the target mass range, and IMRPhenomPv2 (77) otherwise. This analysis typically takes hours to complete. If this analysis completes after an Initial Notice is sent, its output is released as an Update Notice. The UV-optical radiation from a kilonova is expected to fade away within ~ 48 h (5), so parameter estimation updates within $\sim$ hour are sufficient for follow-up purposes.

E. Mock Data Challenge. To create a background for the MDC, we consider the stretch of data taken between Jan 05, 2020 to Feb 14, 2020, by the LVK instruments during O3. A total of 5×10^4 simulated CBC waveforms with mass and spin distributions mentioned in Table 1 are injected into the O3 data with an interval of ~ 1 min between injections. The optimal network SNR is greater than 4 for all the injections. This is done to prevent “hopeless” injections, which are improbable to be detected in reality. The IMRPhenomPv2_NRTidalv2 waveform approximant is used to consider matter effects in case of NS components of the injections. For BHs, the same waveform is used with the tidal parameters set to zero. In order to label a component as a NS, the SLy (78) NS EoS is used, which allows for a maximum mass of $\sim 2.05 M_{\odot}$. Hence, in this scheme, the tidal deformability of component masses above this limit is set to zero consistent with being a BH. In particular, all injections above the SLy maximum mass are assumed to be BHs, and the appropriate relative rate is used for the same. These injected signals are primarily recovered by CBC pipelines, and occasionally by Burst pipelines. In the MDC exercise, we have

⁴⁴<https://git.ligo.org/emfollow/em-properties/em-bright>.

Table 1. Distribution of intrinsic properties (component masses m and spins a) of binary systems in the injection sample

Binary type	Object	$m/M_{\odot}$ (min/max)	m distribution	Max a	a distribution
<i>Compact object properties</i>					
BNS	Primary	1.0 to 2.05	Uniform	0.4	Uniform & isotropic
	Secondary	1.0 to 2.05	Uniform	0.4	Uniform & isotropic
Nsbh	Primary	1.0 to 60.0	M^{-1}	0.998	Uniform & isotropic
	Secondary	1.0 to 2.05	Uniform	0.4	Uniform & isotropic
Bbh	Primary	2.05 to 100	$M^{-2.35}$	0.998	Uniform & isotropic
	Secondary	2.05 to 100	M^1	0.998	Uniform & isotropic

The spin distributions are uniform in magnitude and isotropic in orientation, as seen in the last column.

focused most of our analysis on the output and data products of the CBC pipelines. We also note that the injection rate density used in the study is artificially high and not representative of the true discovery rate in O4. We expect $\mathcal{O}(10^2)$ CBC detections during the full duration of O4 (79), compared to $\mathcal{O}(10^3)$ of detections across the 40 d MDC cycle. Therefore, quantities like p_{astro} which rely on the background distribution, may not be the true representation as compared to a realistic signal density. The CBC injection set consists of 40.9% BNS, 35.8% NSBH, and 23.3% BBH injections.

The events are distributed uniformly in comoving volume assuming flat Λ CDM cosmology with $H_0 = 67.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_m = 0.3$ based on Planck 2018 results mentioned in Table 1 of ref. 80. The BNS systems are distributed up to a maximum redshift of $z = 0.15$, the neutron-star black-hole systems up to $z = 0.25$, and the BBHs up to $z = 1.9$. The simulated strain is projected on to the detector geometries, shifted in time to the time of experiment and streamed as 1 s segments for the search and annotation pipelines to analyze in real-time (Section A). The triggers and their annotations were reported in GraceDB for post processing studies.

This exercise is repeated in several cycles for benchmarking analysis and will continue internally during the observing run to continuously track improvements in the alert infrastructure and

provide avenues for pipelines to test their changes. The numbers reported here are those from a single cycle of 5×10^4 injections where the status of most analyses was close to their O4 configurations.

1. Results

A. Latency Measures. Due to the desire for timely follow-up by the multimessenger community, a key feature of the LLAI is dissemination of results as quickly as possible. The goal for the LLAI system is to send alerts for events within 30 s of merger time; this number sets the timescale for comparison below. Here, we perform a systematic study of the alert latency for three of the key pieces of the pipeline [a fourth, the data calibration, construction, and transfer between sites, which takes ~ 5 to 10 s, is not captured here, as well as latency from the ingestion and redistribution by general coordinates network (GCN) or scalable CyberInfrastructure for multi-messenger astrophysics (SCiMMA)]. Latency comes primarily from these three components: i) the search pipelines, ii) the event orchestrator GWCelery, and iii) GraceDB. We note that technical issues during the MDC may also cause some high-latency outliers, so the results presented are conservative when excluding the time needed for transfer and construction of the strain data.




Fig. 2. Left: Histogram of latencies for the sending of the GCN preliminary alert. We compare the MDC latencies to the median latency measured during O4a. The O4a measurement includes data calibration, construction, and transfer time, while the MDC latencies do not. Right: Histogram of latencies for Early Warning alerts. $t_{GCNPRELIM}$ corresponds to the time the GCN preliminary alert is sent, t_0 corresponds to the preferred event merger time, and t_{event} corresponds to the time of event creation.

To calculate event latencies, we compare the time an event is created and appears on GraceDB to that of the known merger time. For GW event candidates during an observing run, the merger time is defined as the time the signal peak reaches Earth's center. For CBC pipelines, we find a median (90%) latency of 12.3 s (41.4 s); for Burst pipelines, we find a median (90%) latency of 72.3 s (671.3 s). We then compare this number to the creation of a superevent; we find a median (90%) latency of 9.4 s (18.1 s). The median superevent latency is lower than the event latency simply due to the fact that the superevent may be created upon the first trigger, and that a superevent often consists of multiple events. We also make the same measurement for Early Warning alerts, shown on the right of Fig. 2; we find a median (90%) latency of -3.1 s (2.9 s). Considering the joint candidates, we find that it takes a median (90%) latency of 32.9 s (44.4 s) to find a coincidence with a GRB injection and a median (90%) latency of 35.3 s (48.4 s) to trigger a RAVEN alert.

Once the event(s) have been created, there is a request for human vetting of the alert, called the Advocate Request; we find a median (90%) latency of 12.7 s (40.1 s) to notify the advocate. To measure the latency of event communication to the community, we also measure the latency for sending of the GCN preliminary alert, which occurs for superevents that pass automated data quality checks; we find a median (90%) latency of 29.5 s (171.8 s). We show this statistic for GCN preliminary alerts on the *Left* of Fig. 2. This latency reported specifically measures that time until the GCN preliminary label is applied. We also compare this to the measured median GCN latency during O4a, the first half of O4, which, including data calibration, construction, and transfer time, is 29.7 s. The agreement between the GCN latency during the MDC and O4a shows our latency has not increased, and has likely slightly decreased as the O4a measure includes data calibration, construction, and transfer time, while the MDC measurement does not. Table 2 shows a compilation of these latency statistics for comparison. The number of candidate events within a superevent was not shown to have a noticeable effect on the latency of that event and its corresponding preliminary alert.

B. Probability of Astrophysical Origin. Each CBC pipeline uploads its own estimate of p_{astro} , as described in Section C, and is inherited by the superevent if the event is the preferred event. By matching the MDC's injected parameters to the recovered p_{astro} values from pipelines, we can test the accuracy of p_{astro} . Fig. 3 shows the recovered $P(\text{source})$ for true sources (e.g., $P(\text{BNS})$ for true injected BNS events), for superevents which pass the public alert FAR threshold. In matching these injections, we

Table 2. Measured latencies for a number of steps in the pipeline

Latency measure	Description	50% (s)	90% (s)
Superevents	$t_{\text{superevent}} - t_0$	9.4	18.1
CBC events	$t_{\text{event}} - t_0$	12.3	41.4
Burst events	$t_{\text{event}} - t_0$	72.3	671.3
Early warning events	$t_{\text{event}} - t_0$	-3.1	2.9
GW advocate request	$t_{\text{ADV_REQ}} - t_0$	12.7	40.1
GCN preliminary sent	$t_{\text{GCN_PRELIM}} - t_0$	29.5	171.8
Coincidence with GRB found	$t_{\text{EM_COINC}} - t_0$	32.9	44.4
RAVEN alert triggered	$t_{\text{RAVEN_ALERT}} - t_0$	35.3	48.4

t_0 corresponds to the event merger time reported by the pipeline, while $t_{\text{superevent}}$ and t_{event} correspond to the time of superevent or event creation. For the case of superevent latencies, t_0 is determined by the preferred event.




Fig. 3. *Top:* Histogram of “preferred” recovered $P(\text{source})$ for true sources, for superevents which pass the significant public alert threshold. The possible source classes are BNS, NSBH, and BBH. This excludes early warning events, which were not fully functional during the time of this analysis. *Bottom:* Cumulative density of the same data. We also include the distribution of correctly recovered BNS or NSBH, which checks for contamination between the two classes due to misrecovered secondary masses or effects from varying the NS/BH mass boundary. As this is a cumulative histogram, the fraction of events above a certain $P(\text{source})$ corresponds to the TPR. We see that the majority of events with a $P(\text{source})$ greater than 0.5 correctly recover the injection source type.

place the cut between NS and BH at $3 M_{\odot}$. “True” terrestrial events correspond to superevents which were not temporally matched to injections, indicating that they arose from detector noise. Injected BBHs are typically recovered confidently, with the vast majority resulting in $P(\text{BBH}) > 0.5$. Over 90% of BBHs are recovered with $P(\text{BBH}) > 0.9$. The $P(\text{BNS})$ and $P(\text{NSBH})$ distributions are less confident than the $P(\text{BBH})$ distribution; about 10% of true BNSs and NSBHs are recovered with $P(\text{BNS})$ or $P(\text{NSBH}) < 0.1$. As a check against contamination across $P(\text{BNS})$ and $P(\text{NSBH})$ due to errors in recovered masses or variation in the definitions of the mass border between NS and BH (i.e., the Tolman–Oppenheimer–Volkoff mass), we can also look at the distribution of the sum of $P(\text{BNS})$ and $P(\text{NSBH})$ for injected events where $m_2 \leq 3 M_{\odot}$. This is shown in the *Bottom* panel of Fig. 3. Compared to $P(\text{BNS})$ or $P(\text{NSBH})$ alone, $P(\text{BNS}) + P(\text{NSBH})$ performs slightly better, with $\sim 75\%$ of true BNS or NSBHs receiving $P(\text{BNS}) + P(\text{NSBH}) > 0.9$, compared to 60% and 65% for NSBH and BNS respectively. If instead we use a threshold $P(\text{source})$ of 0.5, we find a TPR of $\sim 98\%$ for BBH, $\sim 90\%$ for BNS, and $\sim 76\%$ for NSBH injections.




Fig. 4. *Top Left:* A *P-P* plot showing the BAYESTAR sky map statistics for the preferred event. The credible intervals shown in gray are based on the total number of events. *Bottom Left:* Cumulative histograms of BAYESTAR searched area for all events (blue), compared to two-interferometer (Orange), three-interferometer (Green), and Hanford (H1) and Livingston (L1) (Red) events. We see that the three-interferometer events produced smaller searched areas than the two-interferometer events by almost an order of magnitude as discussed in Section C. *Top Right:* A *P-P* plot showing the performance of BAYESTAR (blue) and Bilby (orange) generated sky maps for BNS events. The credible intervals shown in gray based on the total number of such preferred events where both BAYESTAR and Bilby sky maps are available. *Bottom Right:* Cumulative histograms showing searched area statistics for BAYESTAR and Bilby sky maps. We observe that Bilby sky maps give a lower searched area and tend to be more precise than their BAYESTAR counterparts as discussed in Section C.

C. Sky Localization. In order to evaluate localization performance, in the following, we focus on three metrics: i) localization area, ii) retrieved median distance of the source, and iii) searched area. Localization area is the area (measured in deg²) which encloses a given probability contour in the sky map; in this paper, we use 90% as the total cumulative probability threshold. Retrieved distance refers to the median of the distance distribution along the line of sight of the injected sky position. The searched area is the smallest 2D area, starting with the regions of highest probability, that contains the true location of the source; it represents a measurement of the sky area that a telescope with a small FOV relative to the sky map size would need to cover before imaging the true location. We refer the reader to ref. 57 for more details on these parameters, and use the `ligo.skymap`^{##} package to compute all metrics.

When evaluating sky map performance, we consider the preferred events for superevents that fall under the significant FAR threshold before trials factor and were detected by more than one interferometer. We exclude single detector triggers as most resulted from injections that occurred during a portion of O3 replay data where one or more detectors was not in science mode, and they may have sky localizations on the order of the entire sky. For the *Left* half of Fig. 4, we also exclude a slice of parameter space for injected NSs of mass $\leq 2M_{\odot}$ with spins ≥ 0.05 , as these events may not have a match within the pipeline template banks.

Further, in the *Top Left* panel of Fig. 4, we show the accuracy of BAYESTAR sky maps through a *P-P* plot. *P-P* plots of this format show the fraction of injections found within a given credible interval across all levels of credible intervals. The three gray lozenges around the diagonal show the three different levels of confidence (1–3 σ) for the combined BAYESTAR

^{##} <https://git.ligo.org/leo-singer/ligo.skymap>.

map sample. We find that the BAYESTAR sky maps fall just outside the credible intervals for higher credible intervals. This tells us BAYESTAR slightly overstates the precision of its sky localizations. We also show the performance of different pipelines that detected the preferred event in the given sample. See *SI Appendix* for a discussion of sky map performance for each individual pipeline. In the *Bottom Left* panel of Fig. 4, we show the cumulative trend of the searched area statistics from the combined sample. We find a median searched area of 100 to 200 deg². We then compare the two interferometer events with the three interferometer events, to show that the latter produced smaller searched areas. We also include a line for the two-interferometer case where Hanford (H1) and Livingston (L1) specifically observed the event, as those are the interferometers currently being used during O4. The LVK Alert User Guide (<https://emfollow.docs.ligo.org/userguide/>), provides up-to-date information about the interferometers currently in use. In this case, we see that there is marginal difference between the two-interferometer line, and the H1, L1 line.

Further, we probed the accuracy of the BAYESTAR sky maps compared with Bilby sky maps for BNS events. This comparison *P-P* plot and searched area histogram can be seen in the right half of Fig. 4 for the preferred event of all BNS superevents for which both sky maps were produced. In this figure, we include the events excluded in the *Left* half of Fig. 4 to demonstrate Bilby's performance even without the cuts. A plot with those cuts applied can be found in the appendix. From the *Top Right* panel of Fig. 4, we see that the BAYESTAR sky maps tend to sag below Bilby's implying that their precision was overstated as compared to Bilby. There is a trade-off between the latency of BAYESTAR sky maps available with the Preliminary GCN alert, and the improved accuracy of the Bilby sky maps that are available with the completion of parameter estimation. The cumulative searched area plot in the *Bottom Right* panel of Fig. 4 shows that typically Bilby sky maps have lower searched area by factor of 2 or more.

D. EM-Bright. We show the performance of the EM-bright classifiers (64) across all four CBC pipelines via their receiver operating characteristic (ROC) curves. For this, the NSBH boundary is chosen according to SLy EoS but the probabilities are EoS marginalized, meaning we include some uncertainty in the NS EoS in our classifications. The markers indicate three different representative thresholds (a score above which events are considered to be positively classified as the source type in question) along the ROC curves for each pipeline. Since EM-bright classifiers are trained only using GstLAL injection-recovery data (64), evaluating performance across various CBC pipelines is crucial.

In Fig. 5, we see that the HasRemnant quantity for all four pipelines has greater than a 95% TPR for a 5% False positive rate (FPR). In the *Middle* panel, we see (GstLAL, PyCBC, MBTA) perform consistently at ~97% TPR at ~3% FPR for HasNS classifier. The SPIIR pipeline purity is slightly lower compared to the other pipelines, ~94% TPR at the same misclassification fraction. One possible mitigation technique is to use more training data from the pipeline in the training process. In the last panel, for HasMassGap, we see (GstLAL, SPIIR, MBTA) perform with ~80% TPR at ~20% FPR while PyCBC performs slightly below. We expect to enhance the performance of these classifiers further in the near future by retraining the classifiers using the O3 MDC data and considering all pipelines.




Fig. 5. The ROC curves for the different EM-Bright classifiers are shown here for MDC events. The *Top*, *Middle*, and *Bottom* panels refer to HasRemnant, HasNS, and HasMassGap quantities respectively. The markers denote different representative thresholds along the curve.

2. Conclusion

In this paper, we present the performance of the low-latency alert infrastructure and associated data products based on the O3 MDC. A large simulation campaign of compact binaries, i.e., BNS, NSBH, and BBH are injected into a stretch of real data from O3. The data is taken through the entire end-to-end alert infrastructure starting from the search pipeline, data products computation, and alert generation. We demonstrate that for full bandwidth searches automated preliminary alerts, excluding time for data transfer and construction, are delivered with a median

latency of $\lesssim 30$ s, which is an improvement since O3 (81). We show that low-mass BNS injections are successfully detected by early warning searches. Annotations and alert delivery is achieved for a significant fraction of such signals before merger time with a median alert latency of ~ -3 s. It is to be noted however that alert delivery before merger time does not happen for all early warning events. In addition, through the use of this MDC dataset, we demonstrate that the data products, produced in the same workflow as planned for O4, are statistically consistent with simulated values.

The *Pastr*o values giving probability of an astrophysical BBH, BNS, or NSBH were found to correctly classify the source for the majority of events. For a threshold $P(\text{source})$ of 0.5, we find a TPR of $\sim 98\%$ for BBH, $\sim 90\%$ for BNS, and $\sim 76\%$ for NSBH injections.

The distribution of injected sky positions is found to be well recovered by the sky maps produced by both BAYESTAR and Bilby, as evidenced by Fig. 4. BAYESTAR provides low-latency sky maps that slightly overstate the precision, while Bilby provides improved accuracy upon completion of parameter estimation. The median searched area is found to be 100 to 200 deg^2 , with slight variations between the method and pipelines. We observed that Bilby sky maps have better precision, and typically gave a smaller searched area.

The EM-Bright values corresponding to the probabilities of HasNS and HasRemnant have a TPR of above $\sim 95\%$ at $\sim 5\%$ FPR across GstLAL, PyCBC, MBTA pipelines. The SPIIR pipeline is performing similarly for HasRemnant but its performance is slightly lower for HasNS. HasMassGap, on the other hand, has a TPR of above $\sim 80\%$ at $\sim 20\%$ FPR.

This paper presents the low-latency data products for O4 and their expected performance. Additional data products that expand on the current EM-bright products for determining the likelihood of a kilonova are being developed that include predictions of mass ejecta for BNS and NSBH events, as well as peak magnitudes for corresponding kilonovae. We hope to make these data products public in the future.

Data, Materials, and Software Availability. The entire mock data challenge consists of a number of cycles each with 40 d of replayed GW strain data with frequent injections. A single trigger alone on an injection contains a large set of data products and outputs, so for logistical reasons it would be best to only release very specific and reduced portions of our data. Software and data products required to reproduce figures will be provided upon reasonable request to the authors.

ACKNOWLEDGMENTS. Software and data products required to reproduce figures will be provided upon reasonable request to the authors. We thank Varun Bhalerao for review of this paper. We thank the MBTA team for their contributions and for allowing the use of their pipeline data. We are grateful for computational resources provided by LIGO Laboratory and are supported by NSF Grants No. PHY-0757058 and No. PHY0823459. This material is based upon work supported by NSF's LIGO Laboratory, which is a major facility fully funded by the NSF. This work used Expanse at the San Diego Supercomputer Cluster

through allocation AST200029—"Towards a complete catalog of variable sources to support efficient searches for compact binary mergers and their products" from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by NSF grants #2138259, #2138286, #2138307, #2137603, and #2138296. A.T. and M.W.C. acknowledge support from the NSF with grant numbers PHY-2010970 and OAC-2117997. S.S.C. and M.C. acknowledge support from the NSF with grant number PHY-2011334 and PHY-2308693. M.C. would also acknowledge support from NSF PHY-2219212. D.C. would like to acknowledge support from the NSF grants OAC-2117997 and PHY-1764464. E.K., E.M., G.M. would like to acknowledge support from NSF Grant PHY-1764464. S.M. acknowledges support from JSPS Grant-in-Aid for Transformative Research Areas(A) No. 23H04891 and No. 23H04893. S.A. thanks the MITI CNRS for their financial support. S.G. acknowledges NSF PHY-2110576. N.A., P.B., A.B., P.B., Y.K.C., C.M., and J.C. acknowledge NSF PHY-2207728. T.D. and V.V.O. have received financial support from Xunta de Galicia (CIGUS Network of research centers), by European Union ERDF and by the "María de Maeztu" Units of Excellence program CEX2020-001035-M and the Spanish Research State Agency, and are supported by the research grant PID2020-118635GB-I00 from the Spanish Ministerio de Ciencia e Innovación.

Author affiliations: ^aInstitute of Multi-messenger Astrophysics and Cosmology, Missouri University of Science and Technology, Rolla, MO 65409; ^bSchool of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455; ^cDepartment of Physics, IIT Bombay, Mumbai 400076, India; ^dMIT Kavli Institute for Astrophysics, Massachusetts Institute of Technology, Cambridge, MA 02139; ^eMIT Laser Interferometer Gravitational-Wave Observatory Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139; ^fArtemis, Observatoire de la Côte d'Azur, Université Côte d'Azur, Nice 06304, France; ^gLeonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201; ^hCanadian Institute for Theoretical Astrophysics, University of Toronto, Toronto, ON M5S 3H8, Canada; ⁱDepartment of Physics, University of Toronto, Toronto, ON M5S 1A7, Canada; ^jDavid A. Dunlap Department of Astronomy, University of Toronto, Toronto, ON M5S 3H4, Canada; ^kDepartment of Physics and Astronomy, Montclair State University, NJ 07043; ^lInstitute for Cosmic Ray Research, The University of Tokyo, Chiba 277-8582, Japan; ^mUniversity of Portsmouth, Portsmouth PO1 3FX, United Kingdom; ⁿUniversité Paris-Saclay, CNRS/IN2P3, IJCLab, Orsay 91405, France; ^oUniversity of Wisconsin, Milwaukee, WI 53211; ^pAustralian Research Council Centre of Excellence for Gravitational Wave Discovery, Hawthorn VIC 3122, Australia; ^qDepartment of Physics, University of Western Australia, Crawley WA 6009, Australia; ^rInstituto Galego de Física de Altas Enerxías, Universidade de Santiago de Compostela, 15705 Santiago de Compostela, Spain; ^sUniversità di Roma La Sapienza and INFN, Sezione di Roma, Roma I-00133, Italy; ^tDepartment of Physics, The Pennsylvania State University, University Park, PA 16802; ^uInstitute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802; ^vLaser Interferometer Gravitational-Wave Observatory (LIGO) Laboratory, California Institute of Technology, Pasadena, CA 91125; ^wDepartment of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802; ^xInstitute for Computational and Data Sciences, The Pennsylvania State University, University Park, PA 16802; ^yUniversity of Portsmouth, Portsmouth PO1 3FX, United Kingdom; ^zDipartimento di Scienze Matematiche, Fisiche e Informatiche, Università di Parma, Parma I-43124, Italy; ^{aa}Istituto Nazionale di Fisica Nucleare, Sezione di Milano Bicocca, Gruppo Collegato di Parma, Parma I-43124, Italy; ^{bb}Nikhef, Amsterdam 1098 XG, The Netherlands; ^{cc}Institute for Gravitational and Subatomic Physics, Utrecht University, Utrecht 3584 CC, The Netherlands; ^{dd}School of Physics, Georgia Institute of Technology, Atlanta, GA 30332; ^{ee}Astrophysics Science Division, NASA Goddard Space Flight Center, Code 661, Greenbelt, MD 20771; ^{ff}Joint Space-Science Institute, University of Maryland, College Park, MD 20742; ^{gg}Department of Physics, University of Florida, Gainesville, FL 32611-8440; ^{hh}Department of Physics, University of Maryland, College Park, MD 20742

Author contributions: S.S.C., A.T., G.W., G.M., D.C., S.A., P.B., M.W.C., R.E., S.G., S.M., P.B., A.B., N.A., P.B., G.C.D., T.D.C., M.C., J.C., S.C., Y.-K.C., P.C., L.D., T.D., M.D., B.E., P.G., W.G., C.H., R.H., I.H., E.K., M.K., A.K.Y.L., R.M., E.M., D.M., C.M., X.M.-A., A.P., R.D.P., B.P., S.R., S.S., L.P.S., D.S., D.T., M.T., L.T., V.V.-O., L.W., and D.W. designed research; S.S.C., A.T., G.W., G.M., D.C., S.A., P.B., G.C.D., T.D.C., M.C., J.C., S.C., Y.-K.C., P.C., L.D., T.D., M.D., B.E., P.G., W.G., C.H., R.H., I.H., E.K., M.K., A.K.Y.L., R.M., E.M., D.M., C.M., X.M.-A., A.P., R.D.P., B.P., S.R., S.S., L.P.S., D.S., M.S., D.T., M.T., L.T., V.V.-O., L.W., and D.W. performed research; S.S.C., A.T., G.W., G.M., D.C., S.A., M.W.C., S.G., and S.M. contributed new reagents/analytic tools; S.S.C., A.T., G.W., G.M., D.C., S.A., M.W.C., S.G., S.M., P.B., A.B., N.A., and Y.-K.C. analyzed data; and S.S.C., A.T., G.W., G.M., D.C., S.A., M.W.C., S.G., S.M., and N.A. wrote the paper.

1. B. P. Abbott *et al.*, Localization and broadband follow-up of the gravitational-wave transient GW150914. *Astrophys. J.* **826**, L13 (2016).
2. B. P. Abbott *et al.*, Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.* **119**, 161101 (2017).
3. D. A. Coulter *et al.*, Swope supernova survey 2017a (sss17a), the optical counterpart to a gravitational wave source. *Science* **358**, 1556–1558 (2017).
4. S. J. Smartt *et al.*, A kilonova as the electromagnetic counterpart to a gravitational-wave source. *Nature* **551**, 75 EP (2017).
5. B. P. Abbott *et al.*, Multi-messenger observations of a binary neutron star merger*. *Astrophys. J. Lett.* **848**, L12 (2017).
6. A. Goldstein *et al.*, An ordinary short gamma-ray burst with extraordinary implications: Fermi-GBM detection of GRB 170817a. *Astrophys. J.* **848**, L14 (2017).
7. V. Savchenko *et al.*, INTEGRAL detection of the first prompt gamma-ray signal coincident with the gravitational-wave event GW170817. *Astrophys. J.* **848**, L15 (2017).
8. B. P. Abbott *et al.*, Gravitational waves and gamma-rays from a binary neutron star merger: Gw170817 and grb 170817a. *Astrophys. J. Lett.* **848**, L13 (2017).

9. R. Abbott *et al.*, Observation of gravitational waves from two neutron star-black hole coalescences. *Astrophys. J. Lett.* **915**, L5 (2021).
10. A. Bauswein, O. Just, H. T. Janka, N. Stergioulas, Neutron-star radius constraints from GW170817 and future detections. *Astrophys. J. Lett.* **850**, L34 (2017).
11. B. Margalit, B. Metzger, Constraining the maximum mass of neutron stars from multi-messenger observations of gw170817. *Astrophys. J. Lett.* **850** (2017).
12. M. W. Coughlin, T. Dietrich, B. Margalit, B. D. Metzger, Multimessenger Bayesian parameter inference of a binary neutron star merger. *Mon. Notices R. Astron. Soc.: Lett.* **489**, L91–L96 (2019).
13. M. W. Coughlin *et al.*, Constraints on the neutron star equation of state from at2017gfo using radiative transfer simulations. *Mon. Notic. R. Astronom. Soc.* **480**, 3871–3878 (2018).
14. M. W. Coughlin *et al.*, Implications of the search for optical counterparts during the first six months of the Advanced LIGO's and Advanced Virgo's third observing run: Possible limits on the ejecta mass and binary properties. *Mon. Notic. R. Astron. Soc.* **492**, 863–876 (2019).
15. E. Annala, T. Gorda, A. Kurkela, A. Vuorinen, Gravitational-wave constraints on the neutron-star-matter equation of state. *Phys. Rev. Lett.* **120**, 172703 (2018).
16. E. R. Most, L. R. Weih, L. Rezzolla, J. Schaffner-Bielich, New constraints on radii and tidal deformabilities of neutron stars from gw170817. *Phys. Rev. Lett.* **120**, 261103 (2018).
17. D. Radice, A. Perego, F. Zappa, S. Bernuzzi, Gw170817: Joint constraint on the neutron star equation of state from multimessenger observations. *Astrophys. J. Lett.* **852**, L29 (2018).
18. X. Lai, E. Zhou, R. Xu, Strangeons constitute bulk strong matter: Test using gw 170817. *Euro. Phys. J. A* **55**, 60 (2019).
19. T. Dietrich *et al.*, New constraints on the supranuclear equation of state and the hubble constant from nuclear physics - multi-messenger astronomy. *Science* **370**, 1450–1453 (2020).
20. S. Huth *et al.*, Constraining neutron-star matter with microscopic and macroscopic collisions. *Nature* **606**, 276–280 (2022).
21. M. W. Coughlin *et al.*, Standardizing kilonovae and their use as standard candles to measure the Hubble constant. *Phys. Rev. Res.* **2**, 022006 (2020).
22. M. W. Coughlin *et al.*, Measuring the Hubble constant with a sample of Kilonovae. *Nat. Commun.* **11**, 4129 (2020).
23. B. P. Abbott *et al.*, A gravitational-wave standard siren measurement of the Hubble constant. *Nat. Astron.* **551**, 88–88 (2017).
24. K. Hotokezaka *et al.*, A Hubble constant measurement from superluminal motion of the jet in GW170817. *Nat. Astron.* **3**, 940–944 (2019).
25. R. Chornock *et al.*, The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo gw170817. IV. Detection of near-infrared signatures of r-process nucleosynthesis with gemini-south. *Astrophys. J. Lett.* **848**, L19 (2017).
26. D. A. Coulter *et al.*, Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source. *Science* **358**, 1556–1558 (2017).
27. P. S. Cowperthwaite *et al.*, The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. II. UV, optical, and near-infrared light curves and comparison to Kilonova models. *Astrophys. J. Lett.* **848**, L17 (2017).
28. E. Pian *et al.*, Spectroscopic identification of r-process nucleosynthesis in a double neutron-star merger. *Nature* **551**, 67 EP (2017).
29. S. Rosswog *et al.*, Detectability of compact binary merger macronovae. *Class. Quant. Grav.* **34**, 104001 (2017).
30. D. Watson *et al.*, Identification of strontium in the merger of two neutron stars. *Nature* **574**, 497–500 (2019).
31. M. M. Kasliwal *et al.*, Spitzer mid-infrared detections of neutron star merger GW170817 suggests synthesis of the heaviest elements. *Mon. Notic. R. Astron. Soc. Lett.*, (2019).
32. L. V. Low-latency gravitational-wave alerts for multimessenger astronomy during the second advanced LIGO and virgo observing run. *Astrophys. J.* **875**, 161 (2019).
33. R. Magee *et al.*, First demonstration of early warning gravitational-wave alerts. *Astrophys. J. Lett.* **910**, L21 (2021).
34. R. Abbott, Gwtc-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run (2021).
35. R. Abbott, All-sky search for short gravitational-wave bursts in the third advanced LIGO and advanced virgo run. *Phys. Rev. D*, 104 (2021).
36. R. Abbott *et al.*, All-sky search for long-duration gravitational-wave bursts in the third advanced LIGO and advanced virgo run. *Phys. Rev. D*, 104 (2021).
37. S. Sachdev *et al.*, An early-warning system for electromagnetic follow-up of gravitational-wave events. *Astrophys. J. Lett.* **905**, L25 (2020).
38. M. Kovalam *et al.*, Early warnings of binary neutron star coalescence using the SPIIR search. *Astrophys. J. Lett.* **927**, L9 (2022).
39. C. Messick *et al.*, Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data. *Phys. Rev. D* **95**, 042001 (2017).
40. F. Aubin *et al.*, The MBTA pipeline for detecting compact binary coalescences in the third LIGO-Virgo observing run. *Class. Quant. Grav.* **38**, 095004 (2021).
41. S. Hooper *et al.*, Summed parallel infinite impulse response filters for low-latency detection of chirping gravitational waves. *Phys. Rev. D* **86**, 024012 (2012).
42. J. Luan, S. Hooper, L. Wen, Y. Chen, Towards low-latency real-time detection of gravitational waves from compact binary coalescences in the era of advanced detectors. *Phys. Rev. D* **85**, 102002 (2012).
43. S. A. Usman *et al.*, The PyCBC search for gravitational waves from compact binary coalescence. *Class. Quant. Gravit.* **33**, 215004 (2016).
44. T. Dal Canton *et al.*, Real-time search for compact binary mergers in advanced LIGO and Virgo's third observing run using PyCBC live. *Astrophys. J.* **923**, 254 (2021).
45. B. J. Piotrzkowski, "Searching for Gravitational Wave Associations with High-Energy Astrophysical Transients," PhD thesis, University of Wisconsin-Milwaukee (2022).
46. L. Tsukada *et al.*, Improved ranking statistics of the GstLAL inspiral search for compact binary coalescences. *Phys. Rev. D* **108**, 043004 (2023).
47. B. Ewing *et al.*, Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA's fourth observing run (2023).
48. A. H. Nitz, T. Dal Canton, D. Davis, S. Reyes, Rapid detection of gravitational waves from compact binary mergers with PyCBC Live. *Phys. Rev. D* **98**, 024050 (2018).
49. Q. Chu *et al.*, The spool online coherent pipeline to search for gravitational waves from compact binary coalescences (2021).
50. S. Klimenko, I. Yakushin, A. Mercer, G. Mitselmakher, A coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.* **25**, 114029 (2008).
51. S. Klimenko *et al.*, Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Phys. Rev. D* **93**, 042004 (2016).
52. M. Drago *et al.*, Coherent waveburst, a pipeline for unmodeled gravitational-wave data analysis. *SoftwareX* **14**, 100678 (2021).
53. R. Lynch, S. Vitale, R. Essick, E. Katsavounidis, F. Robinet, Information-theoretic approach to the gravitational-wave burst detection problem. *Phys. Rev. D* **95** (2017).
54. A. L. Urban, "Monsters in the Dark: High Energy Signatures of Black Hole Formation with Multimessenger Astronomy," PhD thesis, University of Wisconsin-Milwaukee (2016).
55. M. A. Cho, "Low-Latency Searches for Gravitational Waves and Their Electromagnetic Counterparts with Advanced LIGO and Virgo," PhD thesis, University of Maryland (2019).
56. N. Arnaud, LIGO and Virgo detector characterization and data quality: Contributions to the O3 run and preparation for O4. *Nucl. Instrum. Meth. A* **1048**, 167945 (2023).
57. L. P. Singer, L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients. *Phys. Rev. D* **93**, 024013 (2016).
58. G. Ashton *et al.*, Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy. *Astrophys. J. Supp.* **241**, 27 (2019).
59. I. M. Romero-Shaw *et al.*, Bayesian inference for compact binary coalescences with bilby: Validation and application to the first LIGO-Virgo gravitational-wave transient catalogue. *Mon. Not. Roy. Astron. Soc.* **499**, 3295–3319 (2020).
60. S. Klimenko *et al.*, Localization of gravitational wave sources with networks of advanced detectors. *Phys. Rev. D* **83**, 102001 (2011).
61. K. M. Görski *et al.*, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *Astrophys. J.* **622**, 759–771 (2005).
62. P. Fernique *et al.*, MOC - HEALPix Multi-Order Coverage map Version 1.0 (IVOA Recommendation 02 June 2014) (2014).
63. G. Ashton *et al.*, Coincident detection significance in multimessenger astronomy. *Astrophys. J.* **860**, 6 (2018).
64. D. Chatterjee *et al.*, A machine learning-based source property inference for compact binary mergers. *Astrophys. J.* **896**, 54 (2020).
65. M. Shibata, K. Hotokezaka, Merger and mass ejection of neutron star binaries. *Annu. Rev. Nuclear Particle Sci.* **69**, 41–64 (2019).
66. F. Foucart, T. Hinderer, S. Nissanke, Remnant baryon mass in neutron star-black hole mergers: Predictions for binary neutron star mimickers and rapidly spinning black holes. *Phys. Rev. D* **98**, 081501 (2018).
67. B. P. Abbott, GW170817: Measurements of neutron star radii and equation of state. *Phys. Rev. Lett.* **121** (2018).
68. S. Ghosh *et al.*, Rapid model comparison of equations of state from gravitational wave observation of binary neutron star coalescences. *Phys. Rev. D* **104** (2021).
69. J. S. Speagle, dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences. *Mon. Not. Roy. Astron. Soc.* **493**, 3132–3158 (2020).
70. P. Canizares *et al.*, Accelerated gravitational-wave parameter estimation with reduced order modeling. *Phys. Rev. Lett.* **114**, 071104 (2015).
71. R. Smith *et al.*, Fast and accurate inference on gravitational waves from precessing compact binaries. *Phys. Rev. D* **94**, 044031 (2016).
72. S. Morisaki, V. Raymond, Rapid parameter estimation of gravitational waves from binary neutron star coalescence using focused reduced order quadrature. *Phys. Rev. D* **102**, 104020 (2020).
73. S. Morisaki *et al.*, Rapid localization and inference on compact binary coalescences with the Advanced LIGO-Virgo-KAGRA gravitational-wave detector network. *Phys. Rev. D* **108**, 123040 (2023).
74. S. Husa *et al.*, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D* **93**, 044006 (2016).
75. S. Khan *et al.*, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D* **93**, 044007 (2016).
76. G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of processing binary black holes. *Phys. Rev. D* **103**, 104056 (2021).
77. M. Hannam *et al.*, Simple model of complete precessing black-hole-binary gravitational waveforms. *Phys. Rev. Lett.* **113**, 151101 (2014).
78. E. Chabanat, P. Bonche, P. Haensel, J. Meyer, R. Schaeffer, A Skyrme parametrization from subnuclear to neutron star densities Part II. Nuclei far from stabilities. *Nucl. Phys. A* **635**, 231–256 (1998).
79. R. W. Kiendrebeogo *et al.*, Updated observing scenarios and multimessenger implications for the international gravitational-wave networks O4 and O5. *Astrophys. J.* **958**, 158 (2023).
80. N. Aghanian *et al.*, Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.* **641**, A6 (2020). [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
81. R. Abbott *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and virgo during the second part of the third observing run. *Phys. Rev. X* **13**, 041039 (2021).