arXiv:2501.04598v2 [gr-qc] 30 Apr 2025

The MBTA Pipeline for Detecting Compact Binary Coalescences in the Fourth LIGO-Virgo-KAGRA Observing Run

Christopher Alléné[1], Florian Aubin[2], Inès Bentara[3], Damir Buskulic[1], Gianluca M Guidi[4,5], Vincent Juste[6], Morgan Lethuillier[3], Frédérique Marion[1], Lorenzo Mobilia[4,5], Benoît Mours[2], Amazigh Ouzriat[3], Thomas Sainrat*[2], and Viola Sordini[3]

[1] *Univ. Savoie Mont Blanc, CNRS, Laboratoire d'Annecy de Physique des Particules - IN2P3, F-74940 Annecy, France*
[2] *Université de Strasbourg, CNRS, IPHC UMR 7178, F-67000 Strasbourg, France*
[3] *Institut de Physique des 2 Infinis de Lyon (IP2I) - UMR 5822, Université de Lyon, Université Claude Bernard, CNRS, F-69622 Villeurbanne, France*
[4] *Università degli Studi di Urbino 'Carlo Bo', I-61029 Urbino, Italy*
[5] *INFN Sezione di Firenze, Sesto Fiorentino, I-50019 , Firenze, Italy*
[6] *Service de Physique Théorique, Université Libre de Bruxelles (ULB), Boulevard du Triomphe, CP225, B-1050 Brussels, Belgium*

Abstract

In this paper, we describe the Multi-Band Template Analysis (MBTA) search pipeline dedicated to the detection of compact binary coalescence (CBC) gravitational wave signals from the data obtained by the LIGO-Virgo-KAGRA collaboration (LVK) during the fourth observing run (O4), which started in May 2023. We give details on the configuration of the pipeline and its evolution compared to the third observing run (O3). We focus here on the configuration used for the offline results of the first part of the run (O4a), which are part of the GWTC-4 catalog (in preparation). We also give a brief summary of the online configuration and highlight some of the changes implemented or considered for the second part of O4 (O4b).

1 Introduction

The Multi-Band Template Analysis (MBTA) pipeline [1, 2] is one of the several pipelines used by the LIGO-Virgo-KAGRA collaboration to search for compact binary coalescences (CBC), along with GstLAL [3–6], PyCBC [7] and SPIIR [8], as well as the unmodelled search pipeline cWB [9]. It has been reliably operating since the first generation of detectors (LIGO and Virgo science runs S6 and VSR2/VSR3), and has participated in a number of detections during the past runs of the second-generation detectors [10–12]. The pipeline is operated both during the run to produce low-latency alerts ("online") and after the run to provide updated results for the catalogs ("offline") .

MBTA analyzes the data from each detector individually, using a matched-filtering technique, before making coincidences between the candidate events. Its specificity is that the matched-filtering process is split across two frequency bands, reducing the computational cost. It also implements a number of methods for rejecting noise transients (usually referred to as "glitches").

In order to improve the sensitivity of the search, and to broaden the capabilities of the pipeline, a number of upgrades have been made to MBTA with respect to [2]. In this paper, we present the configuration and methods used for the offline analysis of O4a data, also highlighting the differences with the online configuration. Section 2 describes the general features of the pipeline and the configuration for the main search, whose results are reported in the GWTC-4 catalog, while section 3 focuses on the specifics of the search for sub-solar mass objects. Section 4 presents the online early-warning search, which can produce triggers without waiting for the binary merger.

2 Description of the pipeline

In this section, we describe the pipeline as it was used for the O4a offline main search, with subsection 2.8 highlighting the configuration specific to the online search. Throughout this section, we illustrate the different parts of the pipeline using results obtained by running the offline configuration on a 40-day period from O3.

2.1 Preprocessing

MBTA analyzes the reconstructed strain data provided by the detectors after they have been properly calibrated [13, 14] and CAT1 data-quality flags, indicating severe problems in the data, have been made available [15, 16].

As in the previous run, MBTA applies several preprocessing steps before the matched-filtering process. The

*Corresponding author: thomas.sainrat@iphc.cnrs.fr

data from the detector is first resampled, usually at 4096 Hz (after applying a low-pass filter). We then apply a gating procedure, which removes stretches of bad-quality data ("glitches"). The gating is applied when a CAT1 flag is active, or when a short-timescale estimate of the detector sensitivity falls below some threshold; the data is then set to 0, using Tukey windows with 0.5-second taper time to ensure a smooth transition[1]. The local estimate of the detector sensitivity is based on the so-called binary neutron star (BNS) range, measuring the capacity of the detector to detect fiducial BNS signals, here computed at a rate of 32 Hz. It is obtained by performing a 0.25-second fast Fourier transform (FFT) of the data to estimate the noise power spectral density (PSD); in order to ensure a good accuracy, at each frequency point we take the maximum between the values in this FFT and the median of all the FFTs over the past 10 s. We then apply the gating if the range drops below an adaptive threshold set to 60% of its median value over the last 10 s. Over the O4a run, less than 0.1% of the observing time was gated by the adaptative threshold for each detector, with gated segments ~ 1.4 s-long on average. It should be noted here that short and loud astrophysical signals can trigger the gating; we therefore run an additional ungated search for this type of signals, with tighter constraints than the main search to mitigate the impact of ungated glitches.

In addition, an estimate of the PSD is produced to be used in the matched-filtering process. It is computed from the preprocessed data, by applying FFTs and taking the median of these FFTs on a certain period of time in order to get a smooth estimate; the specifics depend on the configuration and the region of the parameter space, ranging from seconds to hundreds of seconds regarding the length of the FFTs, and thousands of seconds regarding the time over which the median is estimated.

2.2 Template bank

The matched-filtering process requires simulated waveforms for the CBC signals, which are commonly called *templates*, and arranged in a template bank to cover a wide parameter space. As in previous searches, we consider this parameter space to be determined by the masses of the objects (m_1, m_2) and their dimensionless spins, assumed to be parallel to the orbital angular momentum (χ_1^z, χ_2^z) – it has been shown that this assumption does not hamper much the detection for the targeted population [17–20]. Hereafter, we will use the notations

$$m_{\rm chirp} = \frac{(m_1 m_2)^{\frac{3}{5}}}{(m_1+m_2)^{\frac{1}{5}}} \quad \text{for the chirp mass}$$

$$m_{\rm tot} = m_1 + m_2 \quad \text{for the total mass}$$

$$q = \frac{m_1}{m_2} \text{ with } m_1 > m_2 \quad \text{for the mass ratio}$$

$$\chi_{\rm eff} = \frac{m_1\chi_1^z + m_2\chi_2^z}{m_1+m_2} \quad \text{for the effective spin.}$$

Compared to O3, the parameter space has been extended to cover total masses up to $500\,{\rm M}_\odot$ in detector frame ($200\,{\rm M}_\odot$ in O3), however limiting the mass ratio to 50 (100 in O3). Spin ranges are unchanged from O3, following the astrophysical expectations for merging neutron stars ($|\chi_{1,2}^z| < 0.05$) and black holes ($|\chi_{1,2}^z| < 0.997$); the cutoff between the two is conservatively placed at $2\,{\rm M}_\odot$. We also introduce a new lower limit of 200 ms on the duration of the templates (starting at 18 Hz), which allows to remove very short templates which tend to be more susceptible to glitches. For the same reason, we remove 67 templates identified by the online O4a analysis, which have high, asymmetrical masses and high negative effective spin.

Building a template bank requires making a compromise between maximizing signal recovery and minimizing the number of templates. The technical choices and exact parameters values discussed below are the result of such process. In order to generate the bank, we first consider two regions within the parameter space, corresponding respectively to BNS and BBH with low mass ratio q, and generate "seed" banks for each of them, which are then used as a starting point to build the full bank. This allows us to tightly control the minimal match [21] in these specific regions of the parameter space of specific interest for multi-messenger astronomy, and where most of the past events have been found. A geometric placement algorithm [22, 23] is used in O4a to generate the BNS seed bank, ensuring a minimal match of 0.98. This is changed in O4b, where we instead use a stochastic algorithm [24] for technical reasons. For the low-q BBH, we use a hybrid algorithm [25][2], with a minimal match of 0.98, as most O3 events fall into that region. Finally, the bank for the full parameter space is initialized using the two seed banks and completed using the hybrid algorithm, with a minimal match of 0.965. The reference PSD used during the generation process is the "O4 high" sensitivity for the LIGO detectors, with BNS range 190 Mpc[3]. The template models used to generate the bank are the TaylorF2 approximant [26, 27] for BNS, while SEOBNRv4 [28, 29] is used for BBH; we use the implementation from the LALSimulation library [30] in both

[1] This value changed from O3 to cope with the larger dynamics in the spectrum of the LIGO Livingston data in O4.

[2] The geometrical approach relies on the knowledge of the metric, which is reliable only for low masses; for higher masses a hybrid algorithm, which relies both on the metric and a stochastic approach, gives better results.

[3] `https://dcc.ligo.org/LIGO-T2000012-v1/public`

cases. All parameters are summarized in Table 1.

MBTA needs three different template banks: the "virtual template" bank (containing around 825k templates, shown in Figure 1), covering the full frequency range, and two "real template" banks covering narrower frequency bands, with low-frequency templates (around 55k) up to 80 Hz and high-frequency templates above (around 20k), with which the matched-filtering process will actually be done. Each virtual template is then associated with a real template from each frequency band, by maximizing their match using a modified version of the PyCBC `banksim` algorithm[4]. In addition, we set apart the virtual templates merging before 80 Hz to form a "1-band" bank (around 9k templates). We also select a small (around 1k) sample of templates for the ungated search mentioned in the previous section.

All the banks described here are validated using the `banksim` algorithm, which is based on Monte-Carlo simulations. These show that more than 97% of the signals in the considered parameter space are recovered with a signal-to-noise ratio (SNR) loss lower than 3% (1.5% loss on average).

2.3 Triggers and coincidences

For each detector, MBTA applies the matched-filtering process with every real template and the results from the real templates are summed coherently by applying a time translation and a phase rotation to yield the SNR as a function of time for each virtual template. We check the maximum SNR over a typical timescale of 4 s against a threshold $\rho_{\min}$ (set to 4.4 for the main search and 10 for the ungated search) and record a trigger if this value is exceeded. At this stage, we also apply a basic χ^2 cut, where we compare the distributions of SNR in the two frequency bands to that expected from an astrophysical signal, as we describe in [1]. As in O3, we then define an initial ranking statistic (RS), given by the SNR reweighted with a χ^2 statistic comparing the SNR timeseries with the template autocorrelation, or "autoχ^2". We use the same formula and parameters given in [2]. At this stage, we also apply the data-quality reweighting described in section 2.4.

In order to identify candidate events, we look for coincident triggers between two detectors sharing the same template within a conservative time of flight window between the two detectors (for instance 15 ms between LIGO Hanford and LIGO Livingston). An O4 novelty is that we also consider some single-detector triggers on their own, as described in section 2.5.

We compute the ranking statistic for a coincidence by summing quadratically the ranking statistics of the single detector triggers, and adding a term measuring the consistency of arrival times and phases[5] across the detectors:

$$\mathrm{cRS}_{ij}^2 = \mathrm{RS}_i^2 + \mathrm{RS}_j^2 + 2\ln(P_{ij}) \quad (1)$$

where RS_i is the ranking statistic for detector i and P_{ij} depends on the arrival time and phase differences between the detectors; the latter are compared to the expected distributions assuming a uniform source population, as described in [2]. At low ranking statistics, the trigger distribution is mostly populated by noise. For further analysis, we only consider triggers with a ranking statistic higher than 7.

In order to improve the sky localization [31, 32], we also include SNR timeseries from detectors which did not produce a trigger; however, the ranking statistic is left unchanged[6]. This is for instance used during O4b when Virgo data is available.

We cluster triggers that are close in time (100 ms from any other trigger of the cluster), retaining the trigger with the highest ranking statistic as representative of the cluster.

2.4 Internal data quality assessment

During the O3 run, we introduced a new observable E_R, penalizing triggers that occur during periods showing an excess in the trigger rate, signaling poor data quality (see section 4 of [2]). This procedure was used until the end of the online O4a analysis. Although it was able to reject a number of transient noises, it showed limitations in identifying short noisy periods or glitches impacting the less dense parts of the bank (as those do not result in a significant trigger rate increase), as illustrated e.g. by the low-latency alert S230622ba [33], which had to be retracted.

We thus decided to replace it with a new technique called "SNR-Excess", still applied in post-processing as it uses the full template bank information. It takes the timeseries of the maximum SNR over the template bank[7] (for a period of -4 and $+1$ seconds around a given trigger, with a 32 Hz sampling rate) and compares it *via* a χ^2 statistic to a model based on a large population of simulated signals (injections) in O4a data covering the whole parameter space. We expect that an excess of noise may trigger many templates in the bank, while signals should only trigger matching templates, which usually allows to get a different response. As the pipeline responds differently based on detected parameters, two models are used: one for triggers with (detector-frame) total mass below 25 $M_\odot$ and another above. Figure 2 illustrates this comparison in the cases of injections and of a noise trigger. The χ^2 is renormalized to eliminate its SNR dependency. The SNR-Excess χ^2 is used to reweight the

[4] https://pycbc.org/pycbc/latest/html/banksim.html

[5] In contrast with O3, we do not include the relative amplitude term due to its minor impact.

[6] In [2], we described a method for computing a ranking statistic for three-detector triggers, but it was not used during O4.

[7] Whereas the autoχ^2 only relies on the SNR timeseries for a specific template.

	BNS seed	BBH seed	Full bank (seeds included)	Subsolar mass bank
Number of virtual templates	218182	37208	825840	2253561
Individual masses ($m_{1,2}$) [$M_\odot$]	$[1, 3]$	$[5, 500]$	$[1, 500]$	$[0.2, 10.0]$
Total mass ($m_1 + m_2$) [$M_\odot$]	$[2, 4]$	$[10, 500]$	$[1, 500]$	$[0.2, 1.0]$
Mass ratio ($m_1/m_2 \geq 1$)	$[1, 3]$	$[1, 3]$	$[1, 50]$	$[1.0, 10.0]$
Individual aligned spins ($\chi^z_{1,2}$)	$[-0.05, 0.05]$ if $m_{1,2} \leq 2$ $M_\odot$ $[-0.997, 0.997]$ else			$[-0.1, 0.1]$ if $m \leq 0.5 M_\odot$ $[-0.9, 0.9]$ else
Minimal duration from 18 Hz [s]	-	0.2		-
Minimal match	0.98		0.965	0.97
Approximants for bank generation	TaylorF2	SEOBNRv4_ROM		TaylorF2
Approximants for analysis	SpinTaylorT4	SEOBNRv4_opt		SpinTaylorT4
PSD for bank generation	L1 O4 high-sensitivity (BNS range $\sim$ 190 Mpc)			
f_0 [Hz] for bank generation	25	18		45
f_0 [Hz] for analysis	24	20		45
f_{max} [Hz] for bank generation	2048	1024		1000
f_{max} [Hz] for analysis	2048			1000
f_c [Hz]	80			120

Table 1: Parameters of the template banks used for the O4a MBTA main search and sub-solar mass search. f_0 represents the starting frequency for the templates, while f_{max} is the maximum frequency. f_c corresponds to the frequency cutoff between the two frequency bands. Also note that the waveform approximants used for bank generation are all in the frequency domain, while the ones for the analysis are in the time domain. Differences for the O4b bank are highlighted in Section 2.8.

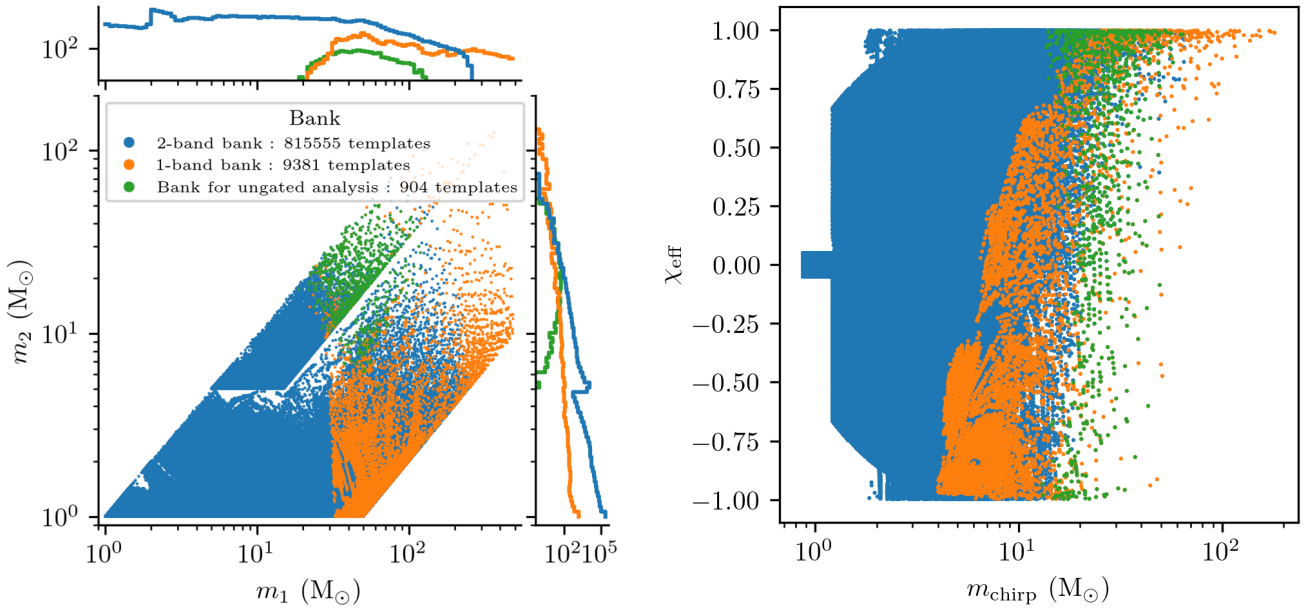

Figure 1: Distribution of the O4a bank templates in the (m_1, m_2) and $(m_{\rm chirp}, \chi_{\rm eff})$ parameter space. The peculiar form of the bank in the (m_1, m_2) space is explained by the use of seeds. The apparent holes naturally appear when building the bank, and are not a concern as events in these areas are still recovered with other templates. The three colors represent the different sub-banks defined in Section 2.2.

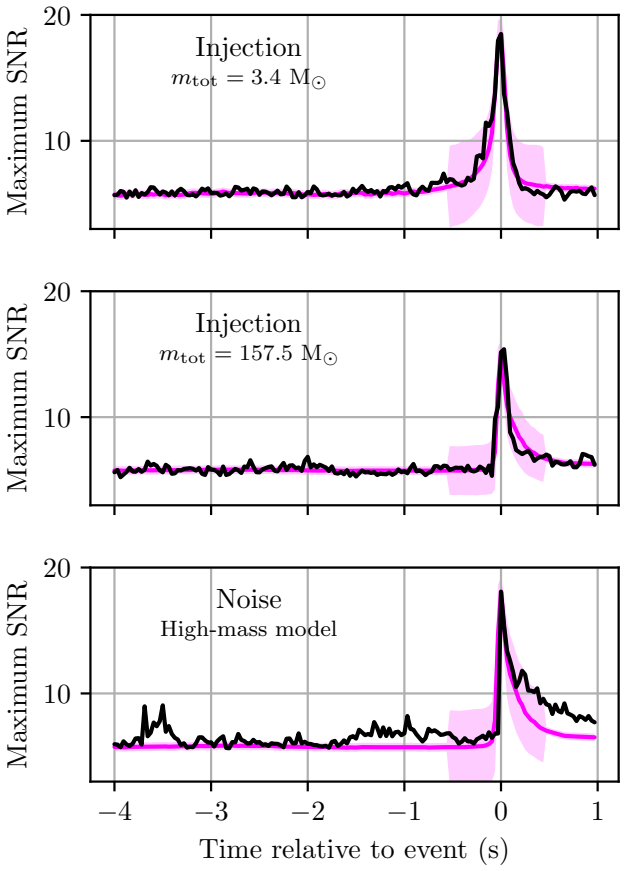

Figure 2: Illustration of the SNR-Excess process. In each case, the black curve corresponds to the timeseries of the SNR maximized over the template bank, while the magenta one is the model (the shaded region represents the uncertainty). The top two figures correspond to well behaved injected signals matching the two different models, with $\chi^2 < 1$; the bottom one corresponds to a glitch within a noisy period, with $\chi^2 \simeq 7$ reducing the ranking statistic by $\sim 30\%$.

single detector ranking statistic (RS), which is redefined as

$$\mathrm{RS} = \begin{cases} \mathrm{RS} & \text{if } \chi^2_{\text{SNR-Excess}} < 1 \\ \mathrm{RS} \times \left(\frac{1}{\chi^2_{\text{SNR-Excess}}}\right)^{0.2} & \text{else.} \end{cases} \tag{2}$$

Figure 3 displays the ranking statistic distribution when using the SNR-Excess or the E_R.

In addition, we reject part of the coincidences for which in either one of the detectors the ranking statistic was decreased by more than 15% after applying the autoχ^2 and the SNR-Excess. This cut is only applied to the noisiest part of the parameter space ($m_{\rm tot} > 106\,{\rm M}_\odot$ for $\chi_{\rm eff} > 0.33$, $m_{\rm tot} > 45\,{\rm M}_\odot$ for $-0.33 < \chi_{\rm eff} < 0.33$ and $m_{\rm tot} > 34\,{\rm M}_\odot$ for $\chi_{\rm eff} < -0.33$), and not applied to coincidences having an SNR above 9 in both LIGO detectors.

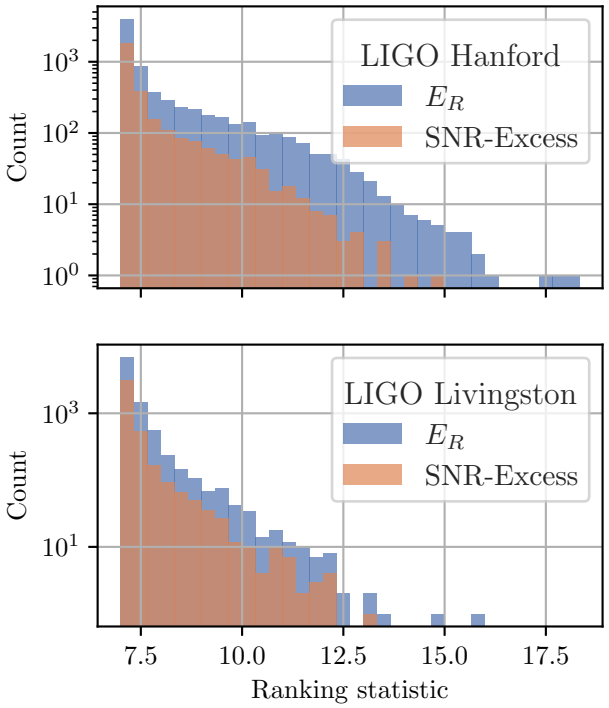

Figure 3: Distribution of the ranking statistic of single-detector triggers from the 40-day O3 rerun when applying the SNR-Excess (orange) and E_R (blue) SNR reweighting.

2.5 Single-detector events

A new addition for the O4 run is to estimate the significance of triggers from a single detector, which other CBC pipelines have also implemented [4, 34, 35]. This is motivated by the fact that about 14%[8] of the O3 run time had only one detector online, which may contain a significant number of events occurring during that time; we also expect to recover some events where two detectors are online but there is substantial signal in only one of them, which is susceptible to happen with one of the LIGO detectors and the less sensitive Virgo, representing an additional $\sim 25\%$ of the run time. In O4, we only considered single-detector triggers from LIGO detectors, whether Virgo was operating or not.

In order to compute the background for estimating the significance (see section 2.6), we cannot use the usual fake coincidence technique [2] as there is only one detector. Instead, we use the single-detector triggers distribution during double-detector time, removing coincident events to reduce contamination from real events. This distribution is fitted with an exponential function in order to extrapolate to high ranking statistic values. The validity of this extrapolation was verified by designing a different way to do the fake coincidences, based on the multi-band feature of MBTA; these investigations are de-

[8] Based on the CBC_CAT2 timelines from GWOSC

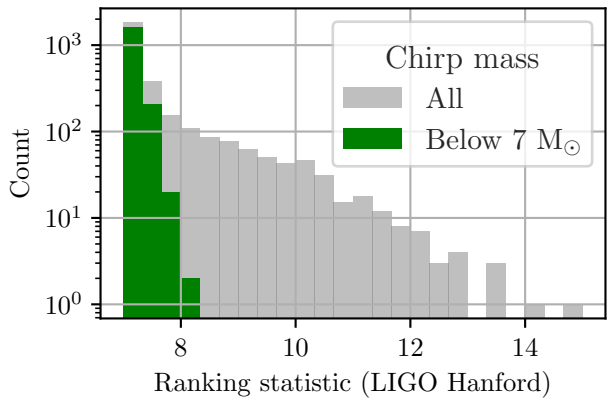

Figure 4: Ranking statistic distribution with and without the single-detector chirp mass threshold at $7\,M_\odot$.

scribed in chapter 6 of [36].

For the O4a search, we chose to only consider single-detector events with a chirp mass below 7 $M_\odot$, as illustrated by Figure 4. These events cover the BNS and most of the NSBH parameter space, which are the main targets for electromagnetic follow-up; they match longer templates, making them easier to distinguish from noise.

2.6 Probability of being astrophysical and source classification

A first metric of the significance of a trigger is the probability that it is of astrophysical origin, $p_{\rm astro} = 1 - p_{\rm terr}$ ($p_{\rm terr}$ representing the probability of the trigger being of terrestrial origin). To compute $p_{\rm astro}$, we use a method similar to the one developed for O3 [37, 38]. We define $p_{\rm astro}$ as

$$p_{\rm astro}(\mathrm{cRS}) = \frac{n_f(\mathrm{cRS})}{n_f(\mathrm{cRS}) + n_b(\mathrm{cRS})} \qquad (3)$$

with n_f and n_b[9] representing the densities over cRS (defined in equations 1 and 2) of the expected number of foreground (astrophysical) and background (noise) events. To take into account the variations of the foreground and background over the parameter space, this equation is evaluated over a number of bins that are defined over the chirp mass (switching to total mass when $m_{\rm chirp} > 7\,M_\odot$), the mass ratio (switching to effective spin when $m_{\rm chirp} > 7\,M_\odot$) and the possible detector combinations. The mass ratio and effective spin axes are divided into 3 bins each, while the chirp mass axis is divided in 43 bins, and the total mass in 40 bins; this gives 213 $(m_{\rm chirp/tot}, q/\chi_{\rm eff})$ bins containing at least one template.

The background $n_b(\mathrm{cRS})$ is measured on a reference dataset using a similar method to O3 [2]. Its logarithm is fitted with third-order polynomials; it is also rescaled by a "clustering factor" accounting for the reduction of the number of triggers during the clustering step.

The foreground is estimated by assuming the same population model as the offline injections (adjusted to the number of detections) used by the LVK to evaluate the search sensitivity for the GWTC-4 catalog. The foreground model is built by running the pipeline on these injections, which provides the distribution of the detected masses, taking into account the discreteness of the template bank. Figure 5 displays the chirp mass distribution of the population model in the source frame and as detected by the pipeline. In order to obtain the foreground as a function of the cRS, we also assume that the foreground cumulative distribution is proportional to $\mathrm{SNR}^{-3} \approx \mathrm{cRS}^{-3}$.

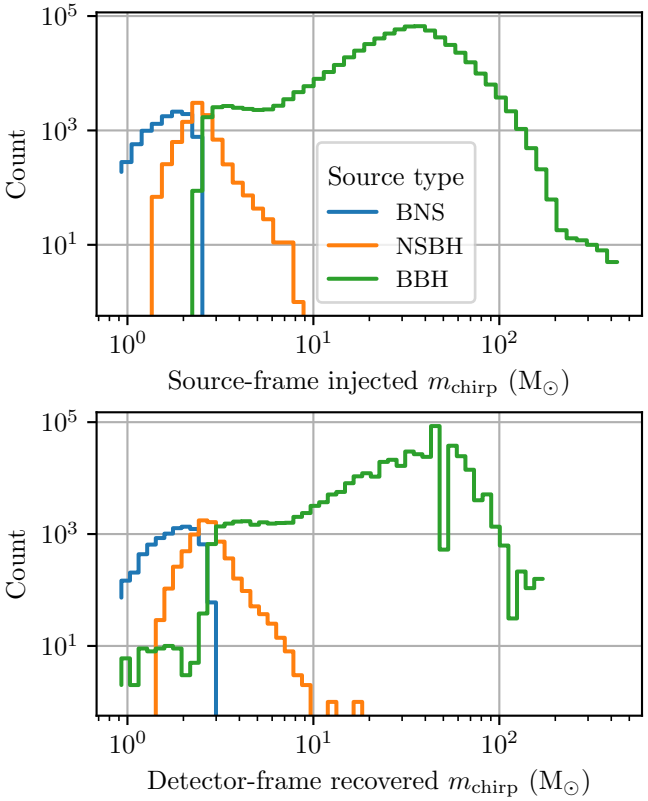

Figure 5: Distribution of the chirp mass (with logarithmic bins) for the various source types used for building the $p_{\rm astro}$ model. The top part shows the distribution of the injections (in the source frame), while the bottom part shows the events recorded by the pipeline (with a ranking statistic above 8 and within ± 0.1 s of an injection).

This process yields a relationship between $p_{\rm astro}$ and the ranking statistic for each bin, which can be parametrized. This parametrization is then used to obtain the $p_{\rm astro}$ value of a trigger, both for the online and offline analyses. The left plot of Figure 6 displays this parametrization for

[9] Related respectively to the Λ_1 and Λ_0 rates in [37].

the different bins.

In addition, we split $p_{\rm astro}$ into three source-type probabilities, such that $p_{\rm astro} = p_{\rm BNS} + p_{\rm NSBH} + p_{\rm BBH}$. These are obtained by multiplying $p_{\rm astro}$ by the relative rate of the corresponding source type recovered within each bin. In the assumed astrophysical population, the threshold between neutron star and black hole is placed at $3\,\mathrm{M}_\odot$. Figure 7 shows the confusion matrix between the MBTA classification and the true source type of the corresponding injections.

2.7 False alarm rate

In O3, the false alarm rate (FAR) was computed using a ranking statistic similar to the one described in section 2.3. In O4, in order to take into account the non-uniformity of sources across the parameter space (introduced by astrophysical priors), we use the $p_{\rm astro}$ described in the previous section as our final ranking statistic. We thus build a global FAR($p_{\rm astro}$), by estimating the number of background events above any given $p_{\rm astro}$ value. This has the additional benefit of ensuring consistency between the FAR and $p_{\rm astro}$ values.

The first step is done independently for each ($m_{\rm chirp/tot}$, $q/\chi_{\rm eff}$, detector combination) bin, which we will denote by the (i, d) indices (with i representing the intrinsic parameters and d the detector combination, HL, H-Lon, L-Hon, H and L[10]). We aim to build a $\mathrm{FAR}_{(i,d)}(p_{\rm astro})$ parametrization, which is done by combining the $p_{\mathrm{astro},(i,d)}(\mathrm{cRS})$ and $n_{b,(i,d)}(\mathrm{cRS})$ parametrizations obtained in the previous section. The $p_{\mathrm{astro},(i,d)}(\mathrm{cRS})$ parametrization is inverted to get a $\mathrm{cRS}_{(i,d)}(p_{\rm astro})$; the $n_{b,(i,d)}(\mathrm{cRS})$ distribution corresponds to the one used to compute the "background" term in $p_{\rm astro}$. By combining the two, we get a one-to-one relationship between $p_{\rm astro}$ and the local FAR for a given bin. Figure 6 illustrates this process.

We then combine the results for different detector combinations to create a single parametrization for each ($m_{\rm chirp/tot}, q/\chi_{\rm eff}$) bin, which is obtained with

$$\mathrm{FAR}_i(p_{\rm astro}) = \sum_{\substack{d\in\{\mathrm{HL,H,L,}\\ \mathrm{H-Lon,L-Hon}\}}} w_d\ \mathrm{FAR}_{(i,d)}(p_{\rm astro}) \quad (4)$$

where $w_d = T_d/T_{\rm total}$ represents the fraction of the observing time for this detector combination over the total time $T_{\rm total}$ (with $T_{\rm total} = T_{\rm HL} + T_{\rm H} + T_{\rm L}$ and $T_{\rm H-Lon} = T_{\rm L-Hon} = T_{\rm HL}$)

Finally, we combine all the bins together to form a single FAR($p_{\rm astro}$) mapping, allowing to obtain a consistent relationship between FAR and $p_{\rm astro}$, by simply computing

$$\mathrm{FAR}(p_{\rm astro}) = \sum_{i\in\mathcal{O}} \mathrm{FAR}_i(p_{\rm astro}) \quad (5)$$

where $\mathcal{O}$ represents the set of 213 ($m_{\rm chirp/tot}, q/\chi_{\rm eff}$) bins.

For the offline run, these models are built for each chunk of 2 months. Figure 8 illustrates the resulting FAR($p_{\rm astro}$) model. Figure 9 shows that, in the noise-dominated regime, the number of events at a given FAR threshold behaves according to the expectation. For the online run, we instead use a different model for each possible source type, as described in Appendix A.

2.8 Online operations

The main search has been running online since the beginning of O4. The configuration during O4a had a few differences with the offline analysis:

- The 67 templates mentioned in subsection 2.2 were not removed.
- The SNR-Excess reranking and cut discussed in section 2.4 were not yet commissioned, and thus the older "E_R" method was used.
- The $p_{\rm astro}$ model was the one used for the O3b offline analysis (see [37]), adjusting the event rates to match the increased sensitivity.
- As detailed in Appendix A, we used separate FAR($p_{\rm astro}$) models depending on the source type, in order to improve the chance of detecting BNS and NSBH sources.

For the O4b online configuration, we integrate the following changes:

- A new template bank was generated using noise curves measured during O4a; we also restrict the maximum mass ratio within the bank to less than 20 (instead of 50).
- The SNR-Excess reranking is part of the online search, without the cut introduced at the end of section 2.4.
- The $p_{\rm astro}$ model is the same as the one for the O4a offline analysis.
- We still use a different FAR($p_{\rm astro}$) model per source type.
- We also filter data from Virgo, which joined the second part of the run. Given the lower sensitivity compared to the two LIGO detectors, the data is not used to estimate the significance, but to improve sky localisation.

In Table 2, we provide a summary of the performance of the various online searches during the O4a run.

[10] H and L represent LIGO Hanford and Livingston respectively; the "-on" means that the SNR in the second detector was too low to make a coincidence. During O4b, we do not differentiate triggers with or without Virgo.

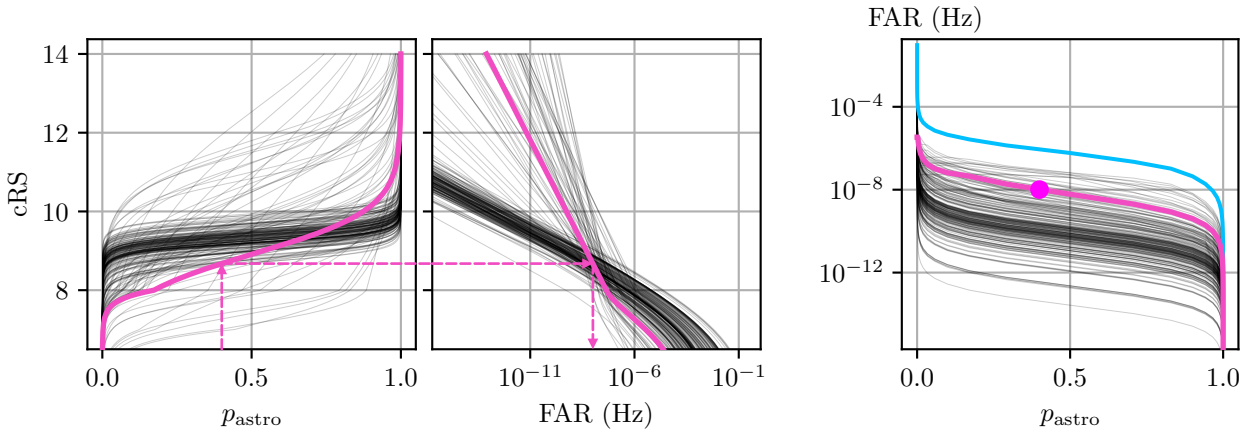

Figure 6: Representation of the process for generating the $\mathrm{FAR}(p_{\mathrm{astro}})$. Each of the black lines represents a different ($m_{\mathrm{chirp/tot}}$, q/χ_{eff}, HL) bin; the magenta line is an example of a bin with high-mass components. The first plot shows the (inverted) $p_{\mathrm{astro}(i,\mathrm{HL})}(\mathrm{cRS})$ relation, while the second one shows the (inverted) $\mathrm{FAR}_{(i,\mathrm{HL})}(\mathrm{cRS})$; a given value of p_{astro} yields a FAR value as shown by the arrows. The plot on the right shows the resulting $\mathrm{FAR}_{(i,\mathrm{HL})}(p_{\mathrm{astro}})$, with the dot representing the point determined from the arrows; the blue curve is the sum of all the individual curves, yielding $\mathrm{FAR}_{\mathrm{HL}}(p_{\mathrm{astro}})$.

			Early-warning		
Search	Main	SSM	42 Hz	50 Hz	58 Hz
Latency 50% [5%;95%] quantiles (s)	21 [19; 25]	47 [42; 55]	−6 [−15; 7]	6 [−3; 11]	9 [2; 19]
Low-significance events	36	~ 50	4	4	11
Low-significance events seen as significant by another pipeline	18	/		0	
Significant events	30	/		0	
Retractions	1	/		0	

Table 2: Summary of the results of the various online searches during the O4a run, based on the events uploaded to the gravitational-wave events candidate database (GraceDB). An event is labelled as "low-significance" if its FAR is below 4.63×10^{-6} Hz and "significant" if it is below 7.72×10^{-8} Hz, corresponding respectively to 2/day and 1/month, adjusted with a trial factor of 5 to account for the various pipelines (these thresholds are the ones used for public alerts, as detailed in `https://emfollow.docs.ligo.org/userguide/analysis/index.html`). As a reference, all pipelines produced a total of 1610 low-significance alerts, 81 validated significant alerts (including 8 single-detector events with $m_{\mathrm{chirp}} > 7\,\mathrm{M}_\odot$, not targeted by MBTA) and 11 retractions. The latency is given relative to the merger time, and represents the time when the event is successfully registered in GraceDB. This includes around 10 s of latency to receive the data from the detectors. The first number corresponds to the median latency, while the numbers inside brackets are the 5% and 95% quantiles. We do not include event counts for the SSM search as no results were made public during O4a. The SSM search is described in Section 3, and the early-warning searches in Section 4.

2.9 Performance illustration

As mentioned at the beginning of the section, we ran the O4a offline configuration of the pipeline on a 40-day period from O3. In Table 3, we list the most significant events recovered by this test run, which we compare to the events from GWTC-3 [38]. Out of the 9 GWTC-3 events from this period, 6 are recovered significantly ($p_{\mathrm{astro}} > 0.5$). Among the 3 remaining events, two of them were not very significant in GWTC-3, while the third one is effectively a single-detector trigger in the O4a configuration (as Virgo data is not analyzed), with a chirp mass above the threshold of $7\,\mathrm{M}_\odot$. We also identify an additional candidate with a p_{astro} slightly above 0.5.

3 Sub-solar mass search

An interesting target for gravitational-wave astronomy are compact objects with a mass possibly below 1 solar mass (sub-solar mass or SSM objects). There is no widely accepted astrophysical channel that predicts the existence of compact objects in this mass range, but several possibilities have been proposed. These include primor-

O3 rerun				GWTC-3			MBTA GWTC-3	
GPS time	RS	IFAR (y)	p_{astro}	Identifier	IFAR (y)	p_{astro}	IFAR (y)	p_{astro}
1264316116.4	16.3	9.06e+07	1.00	GW200129_065458	1e+05	0.99		
1264693411.6	10.8	6.11e+06	1.00	GW200202_154313	1e+05	0.99		
1263097407.7	10.9	3.32e+06	1.00	GW200115_042309	1e+05	0.99	182	0.99
1264213229.9	9.1	4.76	0.96	GW200128_022011	233	0.99	0.303	0.98
1265202095.9	9.9	2.64	0.93	GW200208_130117	3.23e+03	0.99	2.17	0.99
1265273710.2	9.0	0.341	0.70	GW200209_085452	21.7	0.97	0.0833	0.97
1264333383.1	7.8	0.24	0.61					
1265235995.9	7.9	0.0291	0.14	GW200208_222617	0.208	0.70	0.00238	0.02
1265361793.0	8.6	0.00355	0.01	GW200210_092254	0.833	0.54		
1262879936.1	13.6	N/A	N/A	GW200112_155838	1e+05	0.99		

Table 3: Most significant events retrieved by the search ($p_{\mathrm{astro}} > 0.5$) in the partial re-analysis of O3 data. We are able to recover more events than the pipeline contributed to GWTC-3, and with increased IFAR values for MBTA events already in GWTC-3. The improvement depends on the region of the parameter space where the event is recovered. There are 3 GWTC-3 events from the period which were not retrieved : GW200112_155838, which was only observed by LIGO Livingston and Virgo, and thus seen as a single detector event (with m_{chirp} around 27) in the O4 configuration, and GW200208_222617 and GW200210_092254, with respective FAR values of 4.8 per year and 1.2 per year in the catalog, which were not seen as significant in the O3 MBTA offline analysis either. We also retrieve an additional event which was not included in GWTC-3 as it was below the $p_{\mathrm{astro}} > 0.5$ threshold; it was however reported in [39] as GW200129_114245, with a p_{astro} of 0.53 and IFAR of 0.04 y.

Figure 7: Confusion matrix for the source classification (according to the largest p_{source} value). Among the NSBH misclassified as BBH, 73.6% still have $p_{\mathrm{NSBH}} > 0.1$.

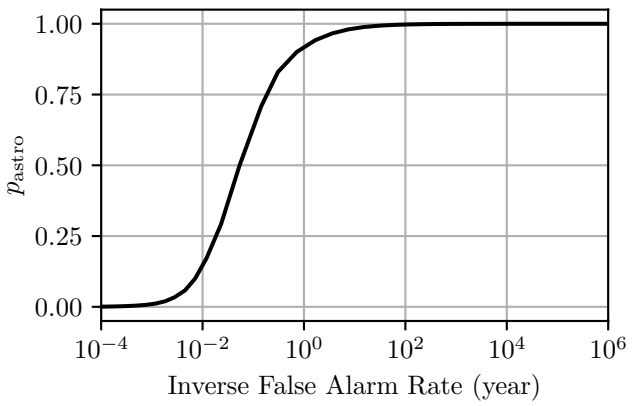

Figure 8: Relation between p_{astro} and the inverse false alarm rate (IFAR), derived from the FAR(p_{astro}) parametrization, for the first two months of the O4a offline analysis.

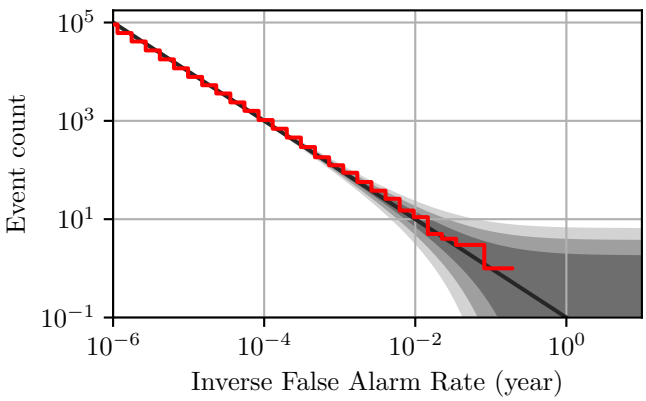

Figure 9: Cumulative distribution of the IFAR of the triggers from the partial O3 rerun. The black line represents the expected amount of noise triggers for 36.6 days of observation, with error bars as the shaded regions. Triggers matching events from GWTC-3 were removed.

dial black holes or dark black holes [40]. The recent study of supernova remnant HESS J1731-347 [41], hinting to the existence of a low-mass neutron star, is an additional motivation for this search. A discovery would point to a new formation channel, and possibly new physics. An absence of detection improves constraints on dark matter models involving such objects.

During O3, MBTA ran an offline search for CBC events comprising an SSM component (hereafter "SSM search" for simplicity), along with the PyCBC and GstLAL pipelines [40, 42]. No events were observed at the time. During the O4 run, we also run an online SSM search, separate from the main search.

3.1 Template bank

Similarly to the main search, the SSM search needs three template banks: the virtual template bank and the two real template banks corresponding to the high-frequency and low-frequency bands. The three banks are generated using a geometric algorithm; the parameters can be found in Table 1. The bank contains about 2.2 million virtual templates, which represents an 18% increase compared to the O3 offline SSM bank (due to evolution of the sensitivity curve). The distribution of the templates across the mass dimensions of the parameter space is displayed in Figure 10.

3.2 Processing

Most of the pipeline operation is similar to the main search; three important differences exist:

- The templates start at 45 Hz in order to reduce the computational cost of the search. This results in a 7% loss in SNR compared to starting at 24 Hz like the main search
- The SNR-Excess uses a different max SNR time-series model, appropriate for the SSM template bank.
- Significance estimate: As there is no existing astrophysical prior for an SSM population, we do not compute p_{astro} for these events, and compute the FAR using the same method as in O3 [2]; however, we now also consider single-detector triggers, in a similar fashion to the main search.

3.3 Online

During the O4a run, we operated the SSM search online for the first time, to be able to respond quickly to a candidate, if any, and to pave the way towards sending public alerts in the future. The configuration used was close to the offline search, with the following differences:

- Single-detector triggers were not considered.
- As in the online main search, the "E_R" method was used instead of SNR-Excess.

Candidates with FAR values below 2/hour were submitted to the GraceDB database. However, public alerts were not enabled in the low-latency system.

4 Early-warning search

The detection of multi-messenger events is one of the important goals of the LIGO-Virgo-KAGRA collaboration: such events, while rare, enable a tremendous amount of science. Only one such event (GW170817 [43]) has been observed up until now, but proved extremely useful to the scientific community, providing insights on the behaviour of extreme matter [44] and putting tight constraints on the speed of gravity [45].

A promising course of action to increase the probability of observing such an event (in particular in channels requiring the instruments to be pointed in the correct direction at the event time) are the so-called "early-warning" alerts. If the GW signal from the source is strong enough, it is possible to detect it before the merger actually takes place, allowing electromagnetic observatories to quickly react and hopefully observe the immediate aftermath of the merger. This is especially true for gamma-ray bursts, which are typically expected a few seconds after the merger. However, implementing such alerts is not trivial, as a large part of the SNR comes from the last few seconds of the coalescence, and requires important optimizations in order to achieve a significantly lower latency.

The early-warning searches [46] focus on BNS signals, and their principle is identical to the main search, except that the templates are cut at a frequency f_{max} much lower than the 2048 Hz for the main search. Table 4 gives the ranges of achievable time gain for each search for a

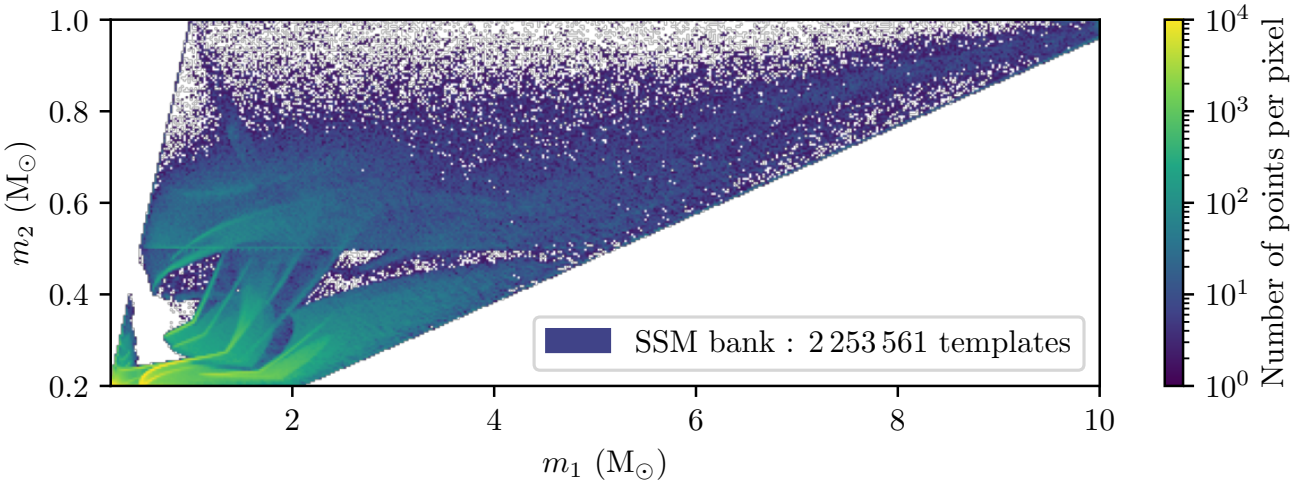

Figure 10: Virtual template bank for the O4a sub-solar mass search in the (m_1, m_2) plane.

f_{max}[Hz]	34	42	50	58
f_0[Hz]	24 (search) 15 (bank generation)			
Fraction of recovered SNR	19%	31%	44%	52%
Time to merger [s]	[68,15]	[39,8]	[25,5]	[17,4]
Individual masses [$M_\odot$]	$m_1, m_2 \in [1, 2.5]$			
Individual spins	$s_1^z, s_2^z \in [-0.05, 0.05]$			
Minimal match	0.98			
Number of templates	1562	2364	4990	5045

Table 4: Parameters for the early-warning searches. f_0 and f_{max} are defined as in Table 1. As for the main search, only aligned spins are taken into account. The SNR fractions are given relative to the 24-2048 Hz band used by the main search, for the LIGO Hanford detector. The times from f_{max} to merger are the maximum and minimum over the explored parameter space.

reference system. Choosing the upper frequency cutoff requires a compromise between the strength of the signals that can be detected (the lower the maximum frequency is, the stronger the signal needs to be) and the latency with which they are reported. We choose to run several instances of the early-warning search with different maximum frequencies in order to maximize the chances of reporting an interesting event; the chosen upper frequency limits for filtering are 42, 50 and 58 Hz, each step leading to roughly a factor 2 increase in collected SNR^2. A 34 Hz search was also implemented, but it was difficult to build a reliable background model due to low-frequency artifacts, leading to this particular search being turned off early during the run. Early-warning searches are performed in a single frequency band for simplicity and due to the small amount of templates. Each of the early-warning searches has a specific template bank, generated using the same geometric algorithm as the BNS seed from the main bank; the parameters are described in Table 4. They only cover the BNS parameter space: this represents both the most interesting signals as they are more likely to have an electromagnetic counterpart, and the ones which are the most likely to be detected early as they are longer.

The early-warning searches have other differences with the main search:

- The frequency upper limit for computing the gating range is set to 40 Hz (in order to focus on glitches in the frequency range specific to early-warning, such as scattered light [16]), and uses a 0.3125 s tapering window in order to optimize the latency.
- It does not use the SNR Excess feature (instead using the "E_R" method as for the online main search).
- We consider only HL coincident events, without filtering Virgo data.

- The computation of the FAR and of $p_{\rm astro}$ is slightly different from the main search. As all templates correspond to BNS signals, we assume that $p_{\rm astro} = p_{\rm BNS}$; we compute $p_{\rm astro}$ over a single bin covering the entire parameter space. The foreground rate is rescaled based on an estimate of the percentage of SNR recovered by the early-warning search compared to the main search (given in Table 4).

No significant early-warning events were observed by MBTA during O4a. During the O4b run, we reported two moderately significant candidates (S240429an and S240623dg), which were later retracted as no corresponding events were found by the main search.

5 Computing resources usage

Running expensive computations is a growing concern among scientific collaborations as they strive to reduce the environmental impact of their activities. In the spirit of addressing this issue, we report on the computing resources used by MBTA.

The main benefit of the two-band architecture is a large reduction of the computational cost of the filtering. Table 5 details the number of CPUs needed to run an instance of the different search types. As a comparison, the hardware used for all MBTA online operations (including the main search, extra instances of the search for test purposes, monitoring and development servers) represents around 2% of the computing resources allocated to online analyses at the main LVK low-latency computing centers (at CIT and Virgo), or 3.5% of the CPU-based resources allocated to CBC search pipelines.

A rough estimate of the carbon footprint based on [47] shows that around 22.3 tons of CO_2 were emitted during the online O4a run (including all instances of the pipeline). A single offline run of the pipeline emits around 600 kg of CO_2 (corresponding to 1.89 M HS06.h[11]) for the main search and 1.5 tons for the SSM search (note that the large difference compared to the online is mostly due to the lower-carbon electrical mix in France, where the offline run is done, compared to Italy, where the online search is hosted).

6 Conclusion

We have described many improvements to the MBTA pipeline, including the new SNR-Excess noise rejection method, a tighter relationship between the FAR and $p_{\rm astro}$ (which is now computed in low-latency) and the inclusion of single-detector triggers. Two new online searches were also added, with the sub-solar mass being run in low-latency for the first time, and the early-warning search. The pipeline was able to contribute to many of low-latency alerts during the O4 run, achieving a latency below 30 seconds. Better glitch management led to a smaller number of retracted alerts compared to O3. Further improvements will continue to be implemented while analyzing new datasets offline and be deployed online in the future.

[11] HS06 or HEP-SPEC06 is a benchmarking procedure commonly used in high-energy physics.

Acknowlegements

We express our gratitude to Gijs Nelemans and Jérôme Novak for their work in carrying out the internal LVK review of the MBTA pipeline. We thank our LVK collaborators from the CBC and low-latency groups for constructive comments. We also thank Dimitri Estevez for commenting on an early version of the paper. We thank Reed Essick (and the entire LVK Rates and Populations group) for developing the foreground model and injection sets used within our work. This analysis exploits the resources of the computing facility at the EGO-Virgo site, and of the Computing Center of the Institut National de Physique Nucléaire et Physique des Particules (CC-IN2P3 / CNRS). This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

Appendix A Multiple FAR($p_{\rm astro}$) models

In section 2.7, we detailed how we obtain a global FAR($p_{\rm astro}$) parametrization for the offline run. For the online configuration, we choose a slightly different approach in order to increase the chances of detecting a multi-messenger event. Instead of summing the FAR$_i(p_{\rm astro})$ over all bins and therefore all source types, we build four different

Search	Main	SSM	Early-warning (all)	Development and validation
Core count	128	160	88	269
Processor type	Xeon Gold 6242	Xeon Gold 6242	Xeon E5-2680	Mixed
Resource percentage	20%	24%	14%	42%

Table 5: Computational resources used by the various MBTA searches during O4a. For each search type, there are two instances, one running on data and the other on data with injections for monitoring purposes.

models corresponding to different source types: BNS, NSBH, low-mass BBH (BBH-low) and high-mass BBH (BBH-high), with a split at $m_{\rm chirp} > 7\mathrm{M}_\odot$ for the BBH motivated by the threshold for single-detector triggers. This is achieved by modifying equation 5 as:

$$\mathrm{FAR}_s(p_{\rm astro}) = w_s \sum_{i \in \mathcal{O}_s} \mathrm{FAR}_i(p_{\rm astro}) \qquad (6)$$

where $\mathcal{O}_s$ is the set of bins for which the source s is the dominant one (i.e. represents the largest number of injections found in this bin when building the $p_{\rm astro}$ model). When associating a FAR to a candidate, we use the model corresponding to its dominant source type. Effectively, this divides the parameter space into four different "searches". We thus need to apply a trial factor w_s to each of the models. We choose here to share the background budget equally between BNS, NSBH and BBH, giving a trial factor of 1/3 for each (for the BBH, this is further split with BBH-low representing 10% and BBH-high the remaining 90%, determined empirically with the $m_{\rm chirp}$ distribution of the population model). This gives values of $w_{\rm BNS} = w_{\rm NSBH} = 0.333$, $w_{\rm BBH-low} = 0.033$ and $w_{\rm BBH-high} = 0.3$. By placing the BNS and NSBH on an equal footing with the BBH, we thus encourage their detection. Note that this process only affects the FAR value; the $p_{\rm astro}$ is unchanged. Figure 11 shows these different models: for a same $p_{\rm astro}$ value, the BNS and NSBH will be assigned a higher IFAR value (meaning the event is given more significance) than the BBH.

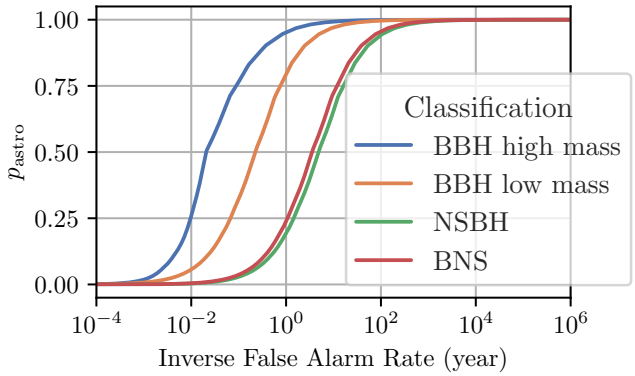

Figure 11: $\mathrm{FAR}(p_{\rm astro})$ parametrizations for the online analysis, using one model per source type.

References

[1] T. Adams, D. Buskulic, V. Germain, et al. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Classical and Quantum Gravity*, 33(17):175012, aug 2016. doi:10.1088/0264-9381/33/17/175012.

[2] F. Aubin, F. Brighenti, R. Chierici, et al. The MBTA pipeline for detecting compact binary coalescences in the third LIGO–virgo observing run. *Classical and Quantum Gravity*, 38(9):095004, apr 2021. doi:10.1088/1361-6382/abe913.

[3] L. Tsukada, P. Joshi, S. Adhicary, et al. Improved ranking statistics of the gstlal inspiral search for compact binary coalescences, 2023.

[4] B. Ewing, R. Huxford, D. Singh, et al. Performance of the low-latency gstlal inspiral search towards ligo, virgo, and kagra's fourth observing run, 2023.

[5] S. Sakon, L. Tsukada, H. Fong, et al. Template bank for compact binary mergers in the fourth observing run of advanced ligo, advanced virgo, and kagra. *Phys. Rev. D*, 109:044066, Feb 2024. doi:10.1103/PhysRevD.109.044066.

[6] C. Hanna, J. Kennington, W. Niu, et al. Template bank for sub solar mass compact binary mergers in the fourth observing run of advanced ligo, advanced virgo, and kagra, 2024.

[7] T. D. Canton, A. H. Nitz, B. Gadre, et al. Real-time search for compact binary mergers in advanced ligo and virgo's third observing run using pycbc live. *The Astrophysical Journal*, 923(2):254, dec 2021. doi:10.3847/1538-4357/ac2f9a.

[8] Q. Chu, M. Kovalam, L. Wen, et al. Spiir online coherent pipeline to search for gravitational waves from compact binary coalescences. *Phys. Rev. D*, 105:024023, Jan 2022. doi:10.1103/PhysRevD.105.024023.

[9] S. Klimenko, G. Vedovato, M. Drago, et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Phys. Rev. D*, 93:042004, Feb 2016. doi:10.1103/PhysRevD.93.042004.

[10] LIGO Scientific Collaboration, J. Aasi, B. P. Abbott, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, April 2015. doi:10.1088/0264-9381/32/7/074001.

[11] F. Acernese, M. Agathos, K. Agatsuma, et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, January 2015. doi:10.1088/0264-9381/32/2/024001.

[12] Y. Aso, Y. Michimura, K. Somiya, et al. Interferometer design of the KAGRA gravitational wave detector. *Physical Review D*, 88(4):043007, August 2013. doi:10.1103/PhysRevD.88.043007.

[13] A. D. Viets, M. Wade, A. L. Urban, et al. Reconstructing the calibrated strain signal in the advanced ligo detectors. *Classical and Quantum Gravity*, 35(9):095015, April 2018. ISSN 1361-6382. doi:10.1088/1361-6382/aab658.

[14] F. Acernese and others (The Virgo Collaboration). Calibration of advanced virgo and reconstruction of the detector strain h(t) during the observing run o3. *Classical and Quantum Gravity*, 39(4):045006, January 2022. ISSN 1361-6382. doi:10.1088/1361-6382/ac3c8e.

[15] S. Soni, B. K. Berger, D. Davis, et al. Ligo detector characterization in the first half of the fourth observing run, 2024.

[16] F. Acernese, M. Agathos, A. Ain, et al. Virgo detector characterization and data quality: results from the o3 run. *Classical and Quantum Gravity*, 40(18):185006, aug 2023. doi:10.1088/1361-6382/acd92d.

[17] B. Abbott et al. Search of s3 ligo data for gravitational wave signals from spinning black hole and neutron star binary inspirals. *Phys. Rev. D*, 78:042002, Aug 2008. doi:10.1103/PhysRevD.78.042002.

[18] C. Van Den Broeck, D. A. Brown, T. Cokelaer, et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Phys. Rev. D*, 80:024009, 2009. doi:10.1103/PhysRevD.80.024009.

[19] T. D. Canton, A. P. Lundgren, and A. B. Nielsen. Impact of precession on aligned-spin searches for neutron-star–black-hole binaries. *Physical Review D*, 91(6), mar 2015. doi:10.1103/physrevd.91.062010.

[20] B. P. Abbott et al. Gw150914: First results from the search for binary black hole coalescence with advanced ligo. *Phys. Rev. D*, 93:122003, Jun 2016. doi:10.1103/PhysRevD.93.122003.

[21] B. J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Physical Review D*, 53(12):6749–6761, June 1996. ISSN 1089-4918. doi:10.1103/physrevd.53.6749.

[22] D. A. Brown, I. Harry, A. Lundgren, et al. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Phys. Rev. D*, 86:084017, 2012. doi:10.1103/PhysRevD.86.084017.

[23] I. W. Harry, A. H. Nitz, D. A. Brown, et al. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. *Phys. Rev. D*, 89(2):024010, 2014. doi:10.1103/PhysRevD.89.024010.

[24] I. W. Harry, B. Allen, and B. S. Sathyaprakash. A Stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D*, 80:104014, 2009. doi:10.1103/PhysRevD.80.104014.

[25] S. Roy, A. S. Sengupta, and P. Ajith. Effectual template banks for upcoming compact binary searches in advanced-ligo and virgo data. *Physical Review D*, 99(2), January 2019. ISSN 2470-0029. doi:10.1103/physrevd.99.024048.

[26] A. Buonanno, B. R. Iyer, E. Ochsner, et al. Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80:084043, Oct 2009. doi:10.1103/PhysRevD.80.084043.

[27] A. Buonanno, Y. Chen, and M. Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Phys. Rev. D*, 67:104025, May 2003. doi:10.1103/PhysRevD.67.104025.

[28] A. Bohé, L. Shao, A. Taracchini, et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Physical Review D*, 95(4), feb 2017. doi:10.1103/physrevd.95.044028.

[29] C. Devine, Z. B. Etienne, and S. T. McWilliams. Optimizing spinning time-domain gravitational waveforms for advanced LIGO data analysis. *Classical and Quantum Gravity*, 33(12):125025, may 2016. doi:10.1088/0264-9381/33/12/125025.

[30] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. free software (GPL), 2018. doi:10.7935/GT1W-FZ16.

[31] Y. Fujii, T. Adams, F. Marion, et al. Fast localization of coalescing binaries with a heterogeneous network of advanced gravitational wave de-

tectors. *Astropart. Phys.*, 113:1–5, 2019. doi: 10.1016/j.astropartphys.2019.04.008.

[32] LIGO Scientific Collaboration Virgo Collaboration KAGRA Collaboration. Ligo-virgo-kagra public alerts user guide.

[33] Ligo Scientific Collaboration, VIRGO Collaboration, and Kagra Collaboration. LIGO/Virgo/KAGRA S230622ba: Retraction of GW compact binary merger candidate. *GRB Coordinates Network*, 34065:1, June 2023.

[34] S. Sachdev, S. Caudill, H. Fong, et al. The gstlal search analysis methods for compact binary mergers in advanced ligo's second and advanced virgo's first observing runs, 2019.

[35] G. S. Cabourn Davies and I. W. Harry. Establishing significance of gravitational-wave signals from a single observatory in the pycbc offline search. *Classical and Quantum Gravity*, 39(21):215012, September 2022. ISSN 1361-6382. doi:10.1088/1361-6382/ac8862.

[36] V. Juste. *Search for compact binary coalescence using a single gravitational waves detector with the MBTA data analysis pipeline*. Ph.D. thesis, Université de Strasbourg, 2023.

[37] N. Andres, M. Assiduo, F. Aubin, et al. Assessing the compact-binary merger candidates reported by the mbta pipeline in the ligo–virgo o3 run: probability of astrophysical origin, classification, and associated uncertainties. *Classical and Quantum Gravity*, 39(5):055002, feb 2022. doi:10.1088/1361-6382/ac482a.

[38] R. Abbott et al. Gwtc-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run. *Physical Review X*, 13(4), December 2023. ISSN 2160-3308. doi: 10.1103/physrevx.13.041039.

[39] A. H. Nitz, S. Kumar, Y.-F. Wang, et al. 4-ogc: Catalog of gravitational waves from compact binary mergers. *The Astrophysical Journal*, 946(2):59, March 2023. ISSN 1538-4357. doi: 10.3847/1538-4357/aca591.

[40] R. Abbott et al. Search for subsolar-mass binaries in the first half of advanced ligo's and advanced virgo's third observing run. *Phys. Rev. Lett.*, 129:061104, Aug 2022. doi:10.1103/PhysRevLett.129.061104.

[41] V. Doroshenko, V. Suleimanov, G. Pühlhofer, et al. A strangely light neutron star within a supernova remnant. *Nature Astronomy*, 6(12):1444–1451, Dec 2022. ISSN 2397-3366. doi:10.1038/s41550-022-01800-1.

[42] R. Abbott et al. Search for subsolar-mass black hole binaries in the second part of advanced ligo's and advanced virgo's third observing run. *Monthly Notices of the Royal Astronomical Society*, page stad588, 02 2023. ISSN 0035-8711. doi:10.1093/mnras/stad588.

[43] B. Abbott, R. Abbott, T. Abbott, et al. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119(16), October 2017. ISSN 1079-7114. doi: 10.1103/physrevlett.119.161101.

[44] D. Watson, C. J. Hansen, J. Selsing, et al. Identification of strontium in the merger of two neutron stars. *Nature*, 574(7779):497–500, October 2019. ISSN 1476-4687. doi:10.1038/s41586-019-1676-3.

[45] B. P. Abbott et al. Gravitational waves and gamma-rays from a binary neutron star merger: Gw170817 and grb 170817a. *The Astrophysical Journal Letters*, 848(2):L13, oct 2017. doi:10.3847/2041-8213/aa920c.

[46] C. Alléné. *Étude des pré-déclenchements lors de la recherche à faible latence de coalescences de binaires compactes par le logiciel MBTA*. Ph.D. thesis, Université Savoie Mont Blanc, 2024.

[47] L. Lannelongue, J. Grealey, and M. Inouye. Green algorithms: Quantifying the carbon footprint of computation. *Advanced Science*, 8(12):2100707, 2021. doi:https://doi.org/10.1002/advs.202100707.