



Publication Year	2016
Acceptance in OA	2020-05-20T14:56:38Z
Title	XMM-Newton Slew Survey Observations of the Gravitational Wave Event GW150914
Authors	Troja, E., Read, A. M., Tiengo, A., SALVATERRA, Ruben
Publisher's version (DOI)	10.3847/2041-8205/822/1/L8
Handle	http://hdl.handle.net/20.500.12386/25009
Journal	THE ASTROPHYSICAL JOURNAL
Volume	822



XMM-NEWTON SLEW SURVEY OBSERVATIONS OF THE GRAVITATIONAL WAVE EVENT GW150914

E. TROJA^{1,2}, A. M. READ³, A. TIENGO^{4,5}, AND R. SALVATERRA⁵

¹ NASA Goddard Space Flight Center, 8800 Greenbelt Rd, Greenbelt, MD 20771, USA

² Department of Physics and Astronomy, University of Maryland, College Park, MD 20742-4111, USA

³ Department of Physics and Astronomy, Leicester University, Leicester LE1 7RH, UK

⁴ Istituto Universitario di Studi Superiori, piazza della Vittoria 15, I-27100 Pavia, Italy

⁵ Istituto di Astrofisica Spaziale e Fisica Cosmica Milano, INAF, via E. Bassini 15, I-20133 Milano, Italy

Received 2016 March 21; accepted 2016 March 29; published 2016 April 22

ABSTRACT

The detection of the first gravitational wave (GW) transient GW150914 prompted an extensive campaign of follow-up observations at all wavelengths. Although no dedicated *XMM-Newton* observations have been performed, the satellite passed through the GW150914 error region during normal operations. Here we report the analysis of the data taken during these satellite slews performed two hours and two weeks after the GW event. Our data cover 1.1 and 4.8 deg² of the final GW localization region. No X-ray counterpart to GW150914 is found down to a sensitivity of 6×10^{-13} erg cm⁻² s⁻¹ in the 0.2–2 keV band. Nevertheless, these observations show the great potential of *XMM-Newton* slew observations for searching for the electromagnetic counterparts of GW events. A series of adjacent slews performed in response to a GW trigger would take $\lesssim 1.5$ days to cover most of the typical GW credible region. We discuss this scenario and its prospects for detecting the X-ray counterpart of future GW detections.

Key words: gamma-ray burst: general – gravitational waves – X-rays: general

1. INTRODUCTION

On 2015 September 14 at 09:50:45 UTC, the advanced Laser Interferometer Gravitational-Wave Observatory (LIGO; Aasi et al. 2015) triggered and detected its first gravitational wave (GW) event, dubbed GW150914 (Abbott et al. 2016b). This was a revolutionary discovery, which confirmed the predictions of general relativity (Abbott et al. 2016f) and opened a new window into the study of our universe. The refined LIGO analysis yields a robust detection (Abbott et al. 2016c), and shows that the observed waveform is consistent with the merger of a binary black hole (BBH) system with relatively “heavy” masses of $\approx 36 M_{\odot}$ and $\approx 29 M_{\odot}$, thus providing the first observational evidence that these astrophysical systems exist (Belczynski et al. 2010; Abbott et al. 2016a).

Despite its crude localization (≈ 590 deg²), GW150914 enjoyed an extensive observing campaign across the electromagnetic spectrum (Abbott et al. 2016e). Although no visible counterpart was expected from a BBH merger, Connaughton et al. (2016) reported the presence of a weak, short duration (~ 1 s) gamma-ray excess, starting 0.4 s after the GW trigger and coming from a sky area consistent with the GW localization. This event was not detected during the more sensitive *INTEGRAL* observations (Savchenko et al. 2016), thus its nature, as well as its association with GW150914, remain uncertain. Nonetheless, several new scenarios were proposed in order to explore the possible electromagnetic signatures of a BBH merger, including short gamma-ray bursts and their accompanying afterglows (e.g., Murase et al. 2016; Perna et al. 2016). In this respect, X-ray observations of GW events represent a promising route toward the discovery of their electromagnetic counterparts. The X-ray window offers several advantages when compared to other energy bands: (a) most ($\sim 80\%$) short gamma-ray bursts (GRBs) have an X-ray afterglow detection (D’Avanzo et al. 2014), compared to only 30% detected in the optical band (Kann et al. 2011), showing that X-ray observations are far more efficient in detecting these

events; (b) space-based X-ray observations are not subject to atmospheric constraints, which often hamper or degrade ground-based searches; and (c) the number of X-ray candidates expected within the LIGO localizations is significantly lower than in deep optical surveys.

The major obstacle for X-ray observatories is to rapidly observe the large GW error regions with an adequate sensitivity. In the case of GW1509014, the entire LIGO error region was scanned by MAXI with a shallow sensitivity of $\approx 10^{-9}$ erg cm⁻² s⁻¹ in the 2–20 keV energy band (Abbott et al. 2016e). Deeper Target of Opportunity (ToO) observations of the field were performed by the *Swift* X-ray Telescope, which covered an area of 5 deg² with a sensitivity of $6 \times 10^{-13} - 6 \times 10^{-12}$ erg cm⁻² s⁻¹ (Evans et al. 2016). In this work, we report the serendipitous *XMM-Newton* (Jansen et al. 2001) slew survey observations of the LIGO localization. In Section 2 we describe our observations, in Section 3 we present our results, and in Section 4 we propose a novel observational strategy in order to maximize the impact of *XMM-Newton* in the era of GW astronomy.

2. OBSERVATIONS

The *XMM-Newton* Slew Survey is based solely on data from the European Photon Imaging Camera (EPIC) pn camera (Strüder et al. 2001). The in-orbit slew speed of 90° hr⁻¹ results in an exposure time between 1–11 s. The soft band (0.2–2 keV) sensitivity limit of *XMM-Newton* slews is 6×10^{-13} erg cm⁻² s⁻¹, close to that of the *ROSAT* All-Sky Survey (Voges et al. 1999). In the hard (2–12 keV) band the slew data goes significantly deeper (4×10^{-12} erg cm⁻² s⁻¹) than previous surveys. For details on the construction and characteristics of the *XMM-Newton* slew survey catalog, see Saxton et al. (2008).

At the time of the GW150914 event, *XMM-Newton* was performing a pointed observation toward the direction of R.A. = 18^h23^m14^s.15, decl. = $-01^{\circ}27'38''$.3, far from the LIGO localization region of GW150914. Shortly afterwards,

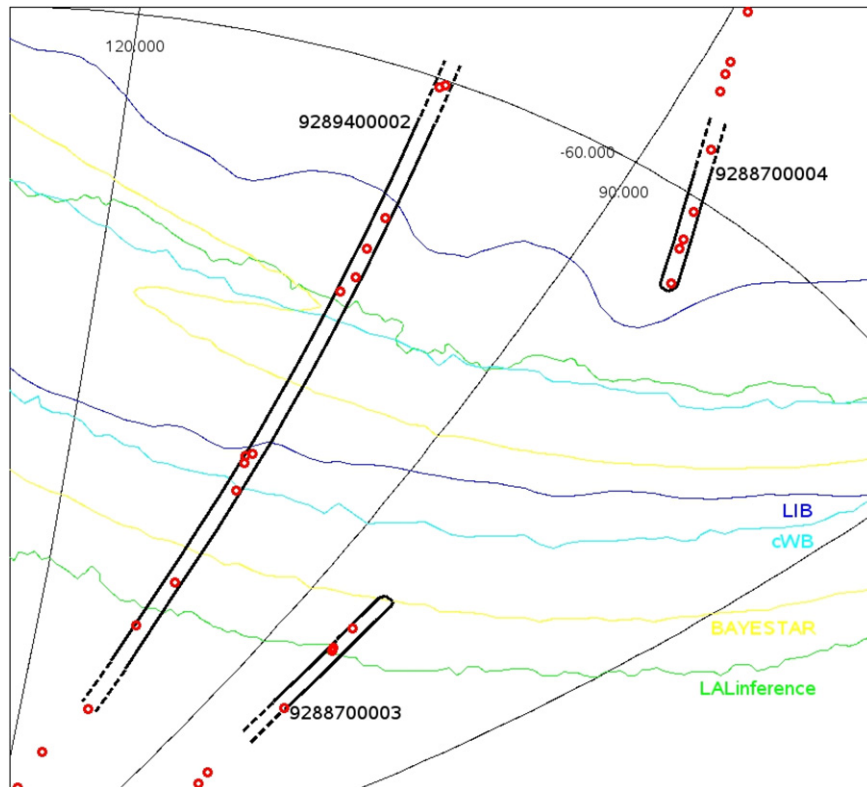


Figure 1. A map of the *XMM-Newton* slews and slew sources from the two weeks after the GW150914 event. Slew sources are shown as the small circles. Slews are shown as the thin black strips (9288700003 extends to the south, 9288700004 to the north, and 9289400002 in both directions). The LIGO 90% confidence regions from the four different pipelines—LALInference, BAYESTAR, cWB (sky), and LIB—are shown. No other slews in the two week period after the GW150914 event lie anywhere close to the LIGO localization regions.

Table 1
Log of *XMM-Newton* Slew Observations

OBSID	Start Time (UT)	$T-T_{\text{GW}}$	LIGO Coverage	
			All (deg ²)	LALInf (deg ²)
9288700003	2015 Sep 14 11:55:03	2.1 hr	1.1	1.1
9288700004	2015 Sep 15 01:34:53	0.66 day
9289400002	2015 Sep 28 23:28:39	14.6 day	6.5	4.8

the satellite slewed and by chance crossed the LIGO 90% error region. Many slews were subsequently made across the sky as the spacecraft moved between its scheduled pointing positions.

Here we analyze the *XMM-Newton* slews made in the two weeks directly after the GW150914 event. Six slews were made during that time, but only three intercept or come close to the LIGO localization maps, as listed in Table 1. The data reduction and analysis used are in many ways similar to the standard slew analysis (Saxton et al. 2008). However, a number of improvements to *XMM-Newton* slew data analysis have been made, and were used in the present analysis, and these will be described in a future paper (A. M. Read et al. 2016, in preparation).

3. RESULTS

A detailed comparison of the slew coverage with the localization regions from the four different LIGO pipelines—LALInference, BAYESTAR, cWB (sky), and LIB—is shown in Figure 1. Since GW150914 is a compact binary coalescence (CBC) event, the LALInference map is considered to be the

most accurate, authoritative, and final localization for this event (Abbott et al. 2016e). We therefore prioritize the LALInference map, but still consider all four maps together as a wider total region. We calculate that slew 9288700003 covers ≈ 1.1 deg² of the LALInf (and of the total) region, and slew 9289400002 covers ≈ 4.8 deg² of the LALInf region (and ≈ 6.5 deg² of the total region). Slew 9288700004 marginally overlaps with the LIGO localization and only covers an area outside the 90% confidence region. These values are tabulated in Table 1.

Source detection was performed using a semi-standard `ebxdetect (local) + espinemap + ebxdetect (map) + emldetect` method, and performed on a single image containing the single events (`pattern = 0`) in the 0.2–0.5 keV band, plus single and double events (`pattern = 0–4`) in the 0.5–12 keV band. Source detection is performed separately in three energy bands: a total band (0.2–12 keV), a hard band (2–12 keV), and a soft band (0.2–2 keV). The source position and background maps are computed separately in each band, and this may result in small differences between the total band and the soft+hard band. The sources detected in the three bands were then combined to produce an initial catalog. Twelve sources were detected in the slew areas coincident with the (total) LIGO region, and X-ray parameters for these are given in Table 2. The sources are also indicated (together with sources detected outside of the LIGO region) in Figure 1. Each detection was assigned a quality flag, from 1 (likely real) to 4 (likely spurious). The detection likelihood (DET_ML) was computed by the `emldetect` task and is defined as $\text{DET_ML} = -\ln P$, where P is the probability the detection is spurious due to a Poissonian fluctuation (Watson et al. 2009).

Table 2
XMM-Newton Slew Sources Detected within the Total LIGO Localization Region

Name	R.A. (deg)	Decl. (deg)	Source Counts			Detection Likelihood			Quality	Catalog ID
			0.2–12 keV	2–12 keV	0.2–2 keV	0.2–12 keV	2–12 keV	0.2–2 keV		
OBSID: 928870003										
XSS J052742.8-763133	81.92836	−76.52586	3.9	...	3.9	10.5	...	10.4	2	1SWXRT J052742.4-763128
XSS J052914.5-762522	82.31060	−76.42279	15.2	4.4	9.9	48.4	10.2	25.2	1	1RXS J052918.2-762514
XSS J052952.5-753757	82.46876	−75.63253	3.6	8.8	3	WISE J052952.46-753805.9
OBSID: 928940002										
XSS J064838.0-641619	102.15847	−64.27199	3.7	...	3.8	9.7	...	10.3	1	1RXS J064837.2-641624
XSS J064903.4-661054	102.26428	−66.18164	5.3	...	5.5	13.0	...	15.1	2	WISE J064903.95-661045.2
XSS J064943.8-651636	102.43274	−65.27666	3.2	9.0	3	WISE J064943.68-651638.7
XSS J065106.2-664320	102.77593	−66.72232	3.8	11.0	4	...
XSS J065226.0-732221	103.10855	−73.37250	3.9	10.1	4	...
XSS J065327.3-720907	103.36414	−72.15190	3.7	9.5	2	WISE J065327.99-720903.2
XSS J065436.5-722926	103.65240	−72.49045	3.8	15.8	1	1RXS J065433.3-722928
XSS J065520.8-721745	103.83708	−72.29580	5.5	8.3	3	WISE J065521.78-721738.0
XSS J065720.5-764311	104.33549	−76.71979	3.5	10.1	2	ATPMN J065720.9-764309

Saxton et al. (2008) quote the astrometric uncertainty of slew sources to be about $8''$ (68% confidence error radius), and Warwick et al. (2012) quote a 90% error circle radius for their slew sample of $10''$. Searches for multiwavelength counterparts made use of the facilities at Vizier, Simbad, and NED. Specific cross-correlations were made with *WISE* and with 2MASS (as in Warwick et al. 2012), and in the case of *ROSAT*, extra allowances were made to account for the uncertainty in the *ROSAT* positions. The results are reported in Table 2 (last column). Apart from XSS J065106.2-664320 and XSS J065226.0-732221, the remaining 10 sources all have *WISE* counterparts within $10''$. The chance of a random positional coincidence with a *WISE* source is only 20%. Among them, four correspond to cataloged X-ray sources. One source, XSS J064838.0-641619, is coincident with the nearby elliptical galaxy NGC 2305 at a distance of 48 Mpc. However, it is detected at a flux level consistent with the *ROSAT* observations. Based on Table 2, we conclude that no new X-ray source was detected in our observation.

4. STRATEGY

XMM-Newton mainly operates in pointing mode, with a minimum exposure of 5 ks per observation. This time-constraint makes it highly impractical to rapidly tile large areas of the sky or to follow-up a large number of candidates. Slew observations are executed only between two pointed observations. Indeed, the 90% error region of GW150914 was partially covered by *XMM-Newton* slews only by chance. Nevertheless, the observations presented in the previous section covered 6 deg^2 (8 deg^2) of the LALInf (total) credible region with a sensitivity comparable to the *Swift*/XRT follow-up (Evans et al. 2016). Here we explore the possibility of performing *XMM-Newton* Target of Opportunity Slew Surveys in response to future GW triggers.

Given the $14'$ radius EPIC-pn field of view and the slew speed of 90° hr^{-1} , the 590 deg^2 area of the GW150914 error region could in principle be covered in about 14 hr. However, overhead times could significantly increase the total time required to perform such a large area survey. Their impact can be evaluated on the basis of a special test for a large area slew

survey that was performed at a reduced speed in 2006 September (revolution 1242). This slow slew survey covered a 140 deg^2 sky region (a rectangle of $40^\circ \times 3^\circ.5$; see Figure 2) within a single satellite orbit (~ 1.5 days of scientific observing time) with 16 partially overlapping slews at a reduced speed of 30° hr^{-1} . During this test, the time overhead between two adjacent slews was ~ 3 ks. A similar survey at the normal slew speed of 90° hr^{-1} would take only ~ 20 hr and an area larger than 40% of the GW150914 error region could be covered within a single *XMM-Newton* orbit. Without requiring an overlap of the slews, the area surveyed in an orbit would increase to more than 70% of the GW150914 error region. Furthermore, in future runs the localization accuracy of GW transients is expected to improve thanks to the increased sensitivity, larger bandwidth and the addition of other detectors to the GW network. A 50% reduction in the error region is already expected for the O2 run (2016-2017), with 14% of the detections being localized within 20 deg^2 (Abbott et al. 2016d), an area which could be covered by slews in ≈ 30 ks. By 2019 this fraction is expected to increase up to $\approx 30\%$ (Abbott et al. 2016d).

The actual observing efficiency also depends on the sky position of the GW trigger, which determines the slew length, its direction, and the sky regions that cannot be observed due to the satellite visibility constraints. About 1/6 of the sky is accessible by *XMM-Newton* at any given time, and some slewing directions may require more complex satellite operations and produce a worse astrometric accuracy, possibly causing an imperfect coverage of the region to be surveyed. Nonetheless, even if the whole error region cannot be covered due to visibility and slewing constraints, the proposed strategy will be an efficient way to search a significant fraction of a large sky area for an X-ray counterpart.

The typical *XMM-Newton* response time to a ToO trigger is $\lesssim 10$ –12 hr, and in the best case, it can be as small as 4 to 5 hr (e.g., Scharrel 2015). As we discuss later on, a very fast turnaround would be preferable in some exceptional cases (e.g., a joint GRB-GW detection or a well-localized GW source), and a slower reaction time is instead preferred to search for off-axis afterglow emission (Granot et al. 2002; van Eerten &




Figure 2. Image of the Galactic Center obtained by slewing *XMM-Newton* during revolution 1242. The image covers a $40 \times 3.5 \text{ deg}^2$ area, comparable to LIGO localizations. This test was performed at a reduced slew speed of 30° hr^{-1} , which is not required in our case. We propose that slew observations of future GW triggers should be performed at a standard slew speed of 90° hr^{-1} . This would allow us to cover a 20 deg^2 area in $\approx 30 \text{ ks}$.

MacFadyen 2011). In order to maximize the scientific return, the data should be analyzed in “real time” to identify and localize possible candidates with sufficient accuracy to enable further follow-up campaigns. At the moment a fast automatic analysis is performed on slew data, but it requires orbital data that become available with a delay of 1–2 weeks. Without the orbital data, the slew source positions within a few arcmin could be obtained directly from raw slew data, relying on the nominal slewing direction and the relative position of known bright X-ray sources. This would make it possible to reduce the processing time to few hours.

In Figure 3 we estimate the detectability of short GRB afterglows during an *XMM-Newton* slew observation. First, we considered the known sample of *Swift* short GRBs with measured redshift (e.g., D’Avanzo et al. 2014). The left panel shows their X-ray fluxes distribution at $t = 1$ day and at a distance $D = 200 \text{ Mpc}$, which is the horizon distance for NS–NS mergers at the LIGO design sensitivity. Most of them (>60%) lie above the *XMM-Newton* slew survey detection threshold, $f_X = 1.2 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ (0.2–12 keV). However, one has to consider that the observed sample is shaped by complex observational biases, which tend to favor the detection and localization of the brightest events. We therefore simulated a sample of GRB afterglows following standard prescriptions for a spherical fireball (Granot & Sari 2002), and post-jet-break scalings from Sari et al. (1999). The explosion properties of short GRBs are not well constrained, and we made an educated guess about the values and distributions of afterglow parameters. We assumed an electron spectral index $p \approx 2.3$,

and simulated three populations of explosions with the following parameters.

- a. A total energy release E uniformly distributed between 10^{48} and 10^{51} erg, a medium density n between 0.001 and 1 cm^{-3} , an electron energy fraction ϵ_e between 10^{-3} and 0.1, a magnetic energy fraction ϵ_B between 10^{-4} and 10^{-2} , and an opening angle θ_j between 1° and 40° .
- b. A total energy release with an average $\langle E \rangle = 10^{49}$ erg and Gaussian dispersion $\sigma_E = 1$ dex, n between 0.001 and 1 cm^{-3} , $\langle \epsilon_e \rangle = 0.1$ with Gaussian dispersion $\sigma_{\epsilon_e} = 0.5$ dex, ϵ_B between 10^{-4} and 10^{-2} , and $\langle \theta_j \rangle = 10^\circ$ with Gaussian dispersion $\sigma_\theta = 0.5$ dex.
- c. $\langle E \rangle = 10^{50}$ erg with Gaussian dispersion $\sigma_E = 1$ dex, $\langle n \rangle = 0.1$ with $\sigma_n = 1$ dex, $\langle \epsilon_e \rangle = 0.1$ with $\sigma_{\epsilon_e} = 0.5$ dex, $\langle \epsilon_B \rangle = 0.1$ with $\sigma_B = 0.5$ dex, and $\langle \theta_j \rangle = 10^\circ$ with $\sigma_\theta = 0.5$ dex.

For each set of parameters, we simulated 340,000 afterglows and calculated their X-ray fluxes in the 0.2–12 keV at $t = 1$ day. We assumed an intrinsic absorption N_H , $i = 10^{21} \text{ cm}^{-2}$, and a flat galactic column density N_H , $g = 3 \times 10^{20} \text{ cm}^{-2}$. The results are shown in Figure 3 (middle panel), where we report the cumulative flux distribution. For the most optimistic case (c) the fraction of visible X-ray afterglows is nearly 80%, and for the less optimistic scenario (a) is still 25%. These estimates should be taken with a grain of salt, as they simplify a complex and poorly constrained phenomenon. The predicted fluxes, calculated for on-axis observers, are significantly fainter if the jet is seen far off-axis

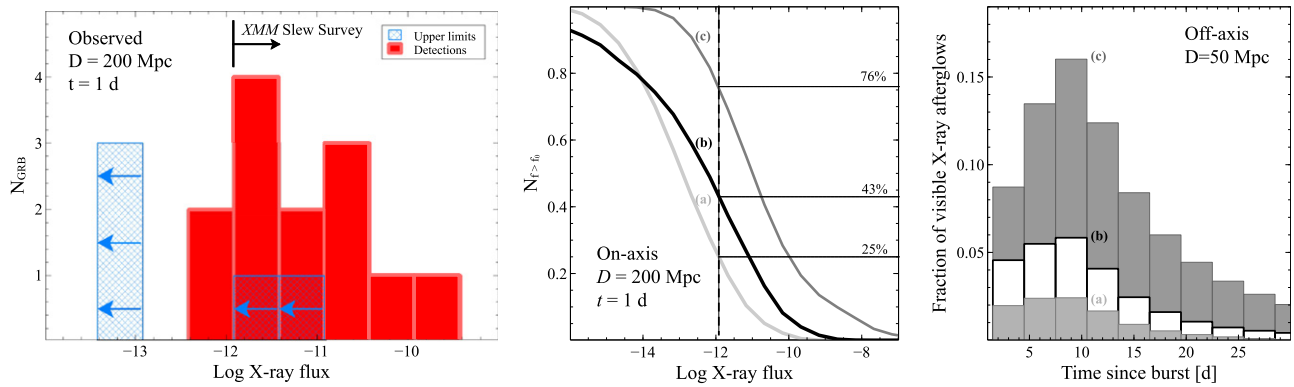


Figure 3. Left panel: X-ray flux distribution for short GRB afterglows with known distance. Fluxes were calculated in the 0.2–12 keV energy band at a time $t = 1$ day after the burst onset. Values were then redshifted at a common distance $D = 200$ Mpc. Middle panel: cumulative flux distribution for simulated on-axis X-ray afterglows at $D = 200$ Mpc, and at a time $t = 1$ day post-burst. The input parameters used for the simulations are described in the text. The vertical dashed line indicates the *XMM-Newton* slew sensitivity in the 0.2–12 keV energy band. Depending on the details of the GRB explosion the fraction of visible afterglows ranges between $\approx 25\%$ and $\approx 75\%$. Right panel: fraction of visible off-axis X-ray afterglows as a function of time. We assumed a distance of $D = 50$ Mpc, and the same input parameters of on-axis simulations. The peak ranges between $\approx 3\%$ and $\approx 15\%$ at $t \approx 8$ days post-burst.

($\theta_{\text{obs}} > \theta_j$). Viewing angle effects could therefore affect our calculations, decreasing the number of visible X-ray afterglows by a factor ~ 3 to over 100, depending on the (unknown) beaming distribution of short GRB jets. This may be alleviated by the fact that the source inclination impacts not only the afterglow flux, but also the strength of the emitted gravitational radiation. Face-on binaries, i.e., with their rotation axis pointing toward us, are more likely to be detected by a factor of 3.4 (Schutz 2011). It is plausible that the GRB jet will form along the same axis, thus suggesting that on-axis short GRBs should also be stronger GW sources. A joint GW-GRB detection would give us a high level of confidence to search for an on-axis X-ray afterglow. However, even in this case, a standard follow-up strategy could not be effective. Given its large field of view, the *Fermi* Gamma-ray Burst Monitor is most likely to detect such a joint event, thus providing only a coarse localization (e.g., Connaughton et al. 2016) comparable to the LIGO one.

In the absence of a GRB trigger, the most likely counterpart would be an orphan afterglow, whose emission is much fainter, and therefore detectable within a much smaller volume. By adopting the simple analytical model of Granot et al. (2002) with the same explosion parameters listed above, and folding in the probability distribution of Schutz (2011) for the binary orientation, we derive in Figure 3 (right panel) the prospects for detecting such off-axis emission with *XMM-Newton*. For a distance of 50 Mpc, the fraction of detectable events ranges between 3% and 15%, much smaller than for an on-axis explosion. However, Figure 3 also shows that the optimal reaction time is slow, peaking between 7 and 15 days after the burst.

Finally, we consider that the past decade of afterglow studies has unveiled the presence of new emission components in addition to the standard forward shock emission described in Granot & Sari (2002). Of particular interest is the emission from a stable millisecond magnetar, which could power a bright and nearly isotropic X-ray transient. Zhang (2013) estimated a bright and persistent X-ray emission from a magnetar-driven relativistic wind lasting for several hours after the burst onset. Metzger & Piro (2014) instead considered the emission from the remnant pulsar wind nebula, and estimated a peak X-ray luminosity of 10^{43} – 10^{44} erg s $^{-1}$ on a timescale of $\lesssim 1$ day. At a distance of 200 Mpc, this corresponds to a flux of 2×10^{-12} – 2×10^{-11} erg cm $^{-2}$ s $^{-1}$, above the *XMM-Newton*

slew survey sensitivity. In case of a particularly well-localized GW event (≈ 20 deg 2), we argue that a rapid observation of the GW field would still be extremely valuable in order to search for these possible X-ray counterparts. Within this field we expect to detect on average 10 X-ray candidates with high significance, which can be easily targeted for subsequent follow-up with *Swift* or optical/nIR ground-based facilities.

5. SUMMARY

We report *XMM-Newton* observations of the field of GW150914, the first LIGO detection. Although no X-ray counterpart was found, these serendipitous observations show the great potential of *XMM-Newton* slews for searching for electromagnetic counterparts of GWs. A single slew of only 7 minutes covered 4.8 deg 2 of the LIGO region down to a sensitivity of 6×10^{-13} erg cm $^{-2}$ s $^{-1}$ (0.2–2 keV). An observing strategy consisting of a series of adjacent slews (Figure 2) could survey a large fraction of future LIGO localizations within a single *XMM-Newton* orbit. In order to maximize the chances of success without excessively impacting *XMM-Newton* operations, we suggest three possible responses to future GW triggers: (a) a rapid ($\lesssim 1$ day) high-priority ToO of the GW region in case of a simultaneous GRB trigger with poor localization; (b) a slower response (≈ 7 days) ToO in case of a rare, nearby ($\lesssim 50$ Mpc) GW transient; and (c) a rapid ToO in case of a well-localized (≈ 20 deg 2) GW event. The latter case could be assigned a medium-priority status, i.e., be implemented only during working hours. On the basis of the experience of the *XMM-Newton* Slew Survey, we expect about 0.5 candidates per square degree, much less than the sources expected from deep optical surveys. Multi-wavelength follow-up observations may be required in order to characterize and identify the selected sources. For a canonical fireball spectral index $\beta_{\text{OX}} \gtrsim 0.6$, the optical counterparts of *XMM-Newton* slew candidates are expected to have $r \lesssim 22$ mag, which can be reached with small aperture (~ 1.5 m) telescopes in reasonably short ($\lesssim 10$ minute) exposures.

We thank Richard Saxton and the anonymous referee for useful and constructive comments that helped improve the manuscript. This work is based on observations obtained with *XMM-Newton*, an ESA science mission with instruments and

contributions directly funded by ESA Member States and NASA. This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France, and of the NASA/IPAC Extragalactic Database (NED), which is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

- Aasi, J., Abbott, B. P., Abbott, T., et al. 2015, *CQGra*, **32**, 074001
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *ApJL*, **818**, L22
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvL*, **116**, 061102
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, arXiv:1602.03843
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016d, *LRR*, **19**, 1
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016e, *ApJL*, submitted (arXiv:1602.08492)
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016f, arXiv:1602.03841
- Belczynski, K., Dominik, M., Bulik, T., et al. 2010, *ApJL*, **715**, L138
- Connaughton, V., Burns, E., Goldstein, A., et al. 2016, arXiv:1602.03920
- D'Avanzo, P., Salvaterra, R., Bernardini, M. G., et al. 2014, *MNRAS*, **442**, 2342
- Evans, P. A., Kennea, J. A., Barthelmy, S. D., et al. 2016, *MNRAS*, in press (arXiv:1602.03868)
- Granot, J., Panaitescu, A., Kumar, P., & Woosley, S. E. 2002, *ApJL*, **570**, L61
- Granot, J., & Sari, R. 2002, *ApJ*, **568**, 820
- Jansen, F., Lumb, D., Altieri, B., et al. 2001, *A&A*, **365**, L1
- Kann, D. A., Klose, S., Zhang, B., et al. 2011, *ApJ*, **734**, 96
- Metzger, B. D., & Piro, A. L. 2014, *MNRAS*, **439**, 3916
- Murase, K., Kashiyama, K., Meszaros, P., Shoemaker, I., & Senno, N. 2016, *ApJL*, in press (arXiv:1602.06938)
- Perna, R., Lazzati, D., & Giacomazzo, B. 2016, *ApJL*, **821**, L18
- Sari, R., Piran, T., & Halpern, J. P. 1999, *ApJL*, **519**, L17
- Savchenko, V., Ferrigno, C., Mereghetti, S., et al. 2016, *ApJL*, **820**, L36
- Saxton, R. D., Read, A. M., Esquej, P., et al. 2008, *A&A*, **480**, 611
- Schartel, N. 2015, *GCN*, **18154**, 1
- Schutz, B. F. 2011, *CQGra*, **28**, 125023
- Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, **365**, L18
- van Eerten, H. J., & MacFadyen, A. I. 2011, *ApJL*, **733**, L37
- Voges, W., Aschenbach, B., Boller, T., et al. 1999, *A&A*, **349**, 389
- Warwick, R. S., Saxton, R. D., & Read, A. M. 2012, *A&A*, **548**, A99
- Watson, M. G., Schröder, A. C., Fyfe, D., et al. 2009, *A&A*, **493**, 339
- Zhang, B. 2013, *ApJL*, **763**, L22