






Research Paper

A VOEvent-based automatic trigger system for the Murchison Widefield Array

P. J. Hancock¹ , G. E. Anderson¹, A. Williams², M. Sokolowski¹, S. E. Tremblay^{1,3}, A. Rowlinson^{4,5}, B. Crosse¹, B. W. Meyers^{1,3,6} , C. R. Lynch^{1,7}, A. Zic⁷, A. P. Beardsley⁸ , D. Emrich² , T. M. O. Franzen¹, L. Horsley², M. Johnston-Hollitt¹, D. L. Kaplan⁹, D. Kenney¹, M. F. Morales¹⁰, D. Pallot¹¹, K. Steele², S. J. Tingay¹, C. M. Trott^{1,3}, M. Walker², R. B. Wayth^{1,3}  and C. Wu¹¹

¹International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia, ²Curtin Institute of Radio Astronomy, Curtin University, GPO Box U1987, Perth WA 6845, Australia, ³ARC Centre of Excellence for All-sky Astrophysics (CAASTRO), ⁴ASTRON, the Netherlands Institute for Radio Astronomy, Postbus 2, NL-7990 AA Dwingeloo, the Netherlands, ⁵Anton Pannekoek Institute for Astronomy, University of Amsterdam, Postbus 94249, NL-1090 GE Amsterdam, the Netherlands, ⁶CSIRO Astronomy and Space Science, PO Box 76, Epping, NSW 1710, Australia, ⁷School of Physics, Sydney Institute for Astronomy, The University of Sydney, NSW 2006, Australia, ⁸School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA, ⁹Department of Physics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA, ¹⁰Department of Physics, University of Washington, Seattle, WA 98195, USA, ¹¹International Centre for Radio Astronomy Research, University of Western Australia, Crawley 6009, Australia

Abstract

The Murchison Widefield Array (MWA) is an electronically steered low-frequency (<300 MHz) radio interferometer, with a ‘slew’ time less than 8 s. Low-frequency (~100 MHz) radio telescopes are ideally suited for rapid response follow-up of transients due to their large field of view, the inverted spectrum of coherent emission, and the fact that the dispersion delay between a 1 GHz and 100 MHz pulse is on the order of 1–10 min for dispersion measures of 100–2000 pc/cm³. The MWA has previously been used to provide fast follow-up for transient events including gamma-ray bursts (GRBs), fast radio bursts (FRBs), and gravitational waves, using systems that respond to gamma-ray coordinates network packet-based notifications. We describe a system for automatically triggering MWA observations of such events, based on Virtual Observatory Event standard triggers, which is more flexible, capable, and accurate than previous systems. The system can respond to external multi-messenger triggers, which makes it well-suited to searching for prompt coherent radio emission from GRBs, the study of FRBs and gravitational waves, single pulse studies of pulsars, and rapid follow-up of high-energy superflares from flare stars. The new triggering system has the capability to trigger observations in both the regular correlator mode (limited to ≥0.5 s integrations) and using the Voltage Capture System (VCS, 0.1 ms integration) of the MWA and represents a new mode of operation for the MWA. The upgraded standard correlator triggering capability has been in use since MWA observing semester 2018B (July–Dec 2018), and the VCS and buffered mode triggers will become available for observing in a future semester.

Keywords: GRBs – FRBs – instrumentation – software – transients

(Received 20 November 2018; revised 16 September 2019; accepted 30 September 2019)

1. Introduction

The steady improvement in radio astronomy technology in recent decades has allowed for the deep study of the physics associated with transient astronomical events (Fender & Bell 2011), either outbursting or explosive incoherent radio sources such as supernova, gamma-ray bursts (GRBs) (Frail et al. 1997), X-ray binaries, and tidal disruption events (Rees 1988) or coherent sources such as fast radio bursts (FRBs, Thornton et al. 2013) and pulsars (Fender et al. 2015b). Radio emission traces

relativistic ejecta and unusual emission mechanisms, the observation of which allows astronomers to directly probe total energy budgets, magnetic fields, and the properties and structure of the interstellar and intergalactic media. However, the most extreme physics takes place at the very start of the transient event, such as the supernova shock ‘breakout’ serendipitously detected in the X-ray band by Soderberg et al. (2008), or the optical flash associated with the reverse shock emission from a GRB as seen by Galama et al. (1999) and Vestrand et al. (2014). In the case of FRBs, the transient event consists entirely of a single burst, with no afterglow yet detected (Williams & Berger 2016). In order to capture such short-lived associated emission, it is necessary for a variety of telescopes covering the entire electromagnetic spectrum to be capable of automatic and rapid follow-up of newly discovered astronomical transients.

Rapid response follow-up of transients has previously been primarily conducted in the GHz regime, with telescopes such as

Author for correspondence: P. J. Hancock, E-mail: Paul.Hancock@Curtin.edu.au

Cite this article: Hancock PJ, Anderson GE, Williams A, Sokolowski M, Tremblay SE, Rowlinson A, Crosse B, Meyers BW, Lynch CR and Zic A, Beardsley AP, Emrich D, Franzen TMO, Horsley L, Johnston-Hollitt M, Kaplan DL, Kenney D, Morales MF, Pallot D, Steele K, Tingay SJ, Trott CM, Walker M, Wayth RB and Wu C. (2019) A VOEvent-based automatic trigger system for the Murchison Widefield Array. *Publications of the Astronomical Society of Australia* 36, e046, 1–11. <https://doi.org/10.1017/pasa.2019.40>

a 12-m dish based at the CSIRO Parkes Observatory, Australia (1.4 GHz; Bannister *et al.* 2012), the 26-m dish located at the Mount Pleasant Radio Observatory, Australia, (2.3 GHz; Palaniswamy *et al.* 2014), and the Arcminute Microkelvin Imager (AMI) based at the Mullard Radio Astronomy Observatory (MRAO), UK (14–18 GHz; Staley *et al.* 2013; Anderson *et al.* 2014; Fender *et al.* 2015b; Anderson *et al.* 2018a). The Australia Telescope Compact Array has also recently been equipped with rapid response systems (e.g. Anderson *et al.* 2018c), further expanding the rapid response frequency coverage in the GHz range from 1–20 GHz.^a

While the first rapid response radio experiments were conducted at 151 MHz with the Cambridge Low Frequency Synthesis Telescope based at the MRAO (Green *et al.* 1995; Dessenne *et al.* 1996), such programs in the MHz domain are only now resurfacing with the construction of the new generation of low-frequency radio telescopes in preparation for the Square Kilometre Array (SKA). For example, the first station of the Long Wavelength Array (LWA1, 10–88 MHz; Taylor *et al.* 2012; Ellingson *et al.* 2013) has commissioned two rapid response triggering modes with a 2-min response time (Yancey *et al.* 2015). In October 2017, the High Band Array (120–168 MHz) of the Low-Frequency Array (LOFAR) enabled a rapid response system capable of triggering within 3–5 min on transient alerts (referred to as the LOFAR Responsive Telescope^b). Meanwhile, the MWA, which operates in the 80–300 MHz frequency range (Tingay *et al.* 2013; Wayth *et al.* 2018), has been running a functional, yet somewhat limited rapid response mode since December 2014 (Kaplan *et al.* 2015). There are also all-sky low-frequency radio experiments that have been (or are capable of being) on-sky at the time of GRBs and gravitational wave events, including the LWA1 Prototype All Sky Imager (LWA1-PASI; Obenberger *et al.* 2014), the Owens Vally Radio Observatory Long Wavelength Array (OVRO-LWA, 27–84 MHz; Anderson *et al.* 2018b; Callister *et al.* 2019), and the LOFAR Low Band Array (10–90 MHz) Amsterdam ASTRON Radio Transient Facility and Analysis Centre (AARTFAAC; Prasad *et al.* 2014, 2016).

The MHz frequency range offers a number of benefits over the GHz regime due to a combination of intrinsic emission properties, propagation effects, and detector properties. Non-thermal coherent emission typically has an inverted spectrum making such sources brighter at MHz frequencies. The arrival time of pulsed signals is delayed with decreasing frequency due to dispersion caused by the ionised intergalactic and interstellar media (Taylor & Cordes 1993), which means that MHz observations can be reliably triggered by gamma-ray, X-ray, or even GHz observations and still be on target before the signal arrives. Radio interferometers in the MHz regime naturally have wide fields of view, are electronically steered, and often have large fractional bandwidths. The MHz frequency range is the ideal observing band to search for prompt emission associated with transients.

1.1. A rapid response system for the MWA

Since 2015, the MWA has been capable of receiving socket-based alerts from the Gamma-ray Coordinates Network (GCN^c). Custom software (based on that used for GRB triggering by Bannister *et al.* 2012) was built to analyse incoming GRB alerts

^ahttps://www.narrabri.atnf.csiro.au/observing/users_guide/html/chunked/ch02s05.html

^b<https://www.astron.nl/radio-observatory/lofar-system-capabilities/responsive-telescope/responsive-telescope>

^cgcn.gsfc.nasa.gov/gcn3_circulars.html

from both the Niel Gehrels *Swift* Observatory (hereafter referred to as *swift*; Gehrels *et al.* 2004) Burst Alert Telescope (BAT; Barthelmy *et al.* 2005) and the *Fermi* Gamma-ray Burst Monitor (GBM; Meegan *et al.* 2009) and then automatically schedule 30 min of observations at the source position in the standard MWA observing mode. Due to the limited amount of information available in the socket-based GCN alerts, this system occasionally triggered on events that were not GRBs. The *Fermi*-GBM and *Swift*-BAT can generate multiple alerts for the same burst, with updated positional information arriving at later times. The typical 1 sigma error radius indicated in a *Fermi*-GBM trigger is 5–15^d plus systematic uncertainties. At 150 MHz, the MWA has a field of view which is 15° radius at half power, which means that the *Fermi*-GBM positions would fall within this region approximately 2/3 of the time. The MWA's socket-based rapid response system was unable to incorporate updated information from subsequent alerts on the same event into the observing schedule, resulting in the final GRB positions from *Fermi*-triggered events sometimes being located near the edge of or even outside of the MWA's large field of view. Nonetheless, the socket-based alert system was successful at automatically observing GRBs. Indeed, the observation of the short-duration GRB 150424A sets the most stringent upper limits (3 Jy at 132 MHz) on prompt radio emission from any type of GRB (Kaplan *et al.* 2015).

We present the upgraded MWA rapid response system, which now responds to transient alerts transmitted via the Virtual Observatory Event standard (VOEvent; Seaman *et al.* 2011) that provides a machine readable format for the communication of astrophysical events. All of the events that are distributed via the GCN are also distributed as VOEvents.^e The GCN-distributed VOEvents contain more detailed information than that provided in the socket-based alerts including: Moon and Sun angular distance, event position in multiple coordinate systems, spacecraft location, alternative classifications, and (probably most importantly) the probability of a given alert being (for example) a genuine GRB. This increase in information makes it possible to trigger follow-up observations with more confidence, to reduce the fraction of false positives, and to update observations as new information becomes available. The VOEvent format is an XML-based format that can be easily interpreted by a variety of software such as the VOEVENT-PARSE python module (Staley *et al.* 2013; Staley 2014). We utilise the software packages provided by the '4 Pi Sky VOEvent Broker' (Staley & Fender 2016) and the Comet VOEvent client (Swinbank 2014) to parse and filter VOEvent transient alerts, enabling the automation of transient follow-up with the MWA.

The triggering system described here is able to observe in three modes. The first mode is the regular correlator setup, which is used for most science observations with the MWA, and has a time resolution of ≥ 0.5 s. The second mode is to capture voltages using the Voltage Capture System (VCS; Tremblay *et al.* 2015) with a time resolution of 0.1 ms. The third and final mode is a buffered capture mode, wherein the telescope is scheduled to capture voltages (with the VCS) to a ring buffer but only write to disk once a trigger is received.

In this paper, we first review some science motivations for a triggered observing system on the MWA (Section 2). Next, we describe the recently developed back-end service (Section 3.1) for the MWA telescope to enable fast rescheduling of the telescope

^dgcn.gsfc.nasa.gov/fermi.html

^egcn.gsfc.nasa.gov/voevent.html

and the corresponding front-end (Section 3.2) which will receive VOEvents and submit observing requests to the MWA. We then describe the VOEvent filter that is in place to respond to *Swift* and *Fermi* GRBs (Section 4) as a case study. We summarise and discuss future developments in Section 5.

2. Science motivations

This project was originally motivated by the desire to probe the very early-time low-frequency radio emission from GRBs. The overlap between short GRBs and gravitational wave events (as demonstrated by the simultaneous detection of GW170817 and GRB 170817A; Abbott et al. 2017a, 2017b), and the possibility of FRB-like signals being produced by these events (for a review of these emission mechanisms see Rowlinson & Anderson 2019), provides further motivation. The incorporation of voltage buffer triggers has made the triggering system particularly useful for the study of FRBs and intermittent pulsars. Additionally, the fact that *Swift* and other space telescopes have on-board catalogues of known flaring stars, some of which are expected to have associated radio flares means that triggered observations of M-Dwarf flares are another immediate science motivation. In lieu of a complete list of science applications, we discuss here just those that are being actively pursued by the MWA using the triggering service. Further discussion of the science applications of the MWA can be found in Bowman et al. (2013) and Beardsley et al. (2019).

2.1. GRB and gravitational wave follow-up

The study of short-duration GRBs (short GRBs) is highly topical as they are linked with binary neutron star (BNS), or neutron star (NS) - black hole (BH) binary mergers, which are the main class of gravitational wave (GW) events known to have electromagnetic counterparts (e.g. GW170817; Abbott et al. 2017a, 2017b). There are several theories that predict such mergers should produce prompt, coherent emission in the form of a pulse, perhaps similar to FRBs, whether due to magnetic braking as the magnetic fields of the NSs are synchronised to the binary rotation before the merger (Hansen & Lyutikov 2001; Lyutikov 2013), persistent or pulsating pulsar-like emission from a short-lived (<1000 s), highly magnetised, supramassive NS remnant (Totani 2013; Rowlinson et al. 2013; Metzger, Berger, & Margalit 2017), or the collapse of said supramassive NS into a BH (Falcke & Rezzolla 2014; Zhang 2014). The detection of such prompt emission associated with the existence and/or collapse of a supramassive NS would allow us to constrain the equation-of-state of nuclear matter (e.g. Lattimer 2012; Lasky et al. 2014).

Both short GRBs and GW events are therefore exciting targets for rapid response observations with MWA. The MWA has previously been involved in performing rapid response triggered observations of both *Swift*- and *Fermi*-detected GRBs (e.g. GRB 150424A; Kaplan et al. 2015), as well as in the multi-wavelength follow-up of GW events (e.g. GW170817; Abbott et al. 2017a; Andreoni et al. 2017). As the MWA is electronically steered, its rapid response mode is capable of automatically repointing the telescope within 14 s of receiving a transient alert (see Section 3.1.1). This means that the MWA could be on target in time to detect any prompt emission associated with a BNS merger. In fact, Rowlinson & Anderson (2019) made predictions for such prompt emission associated with short GRBs over a wide range of redshifts and showed that MWA is very competitive for detecting such signals at low radio frequencies, particularly at the earliest timescales (seconds to minutes post-burst).

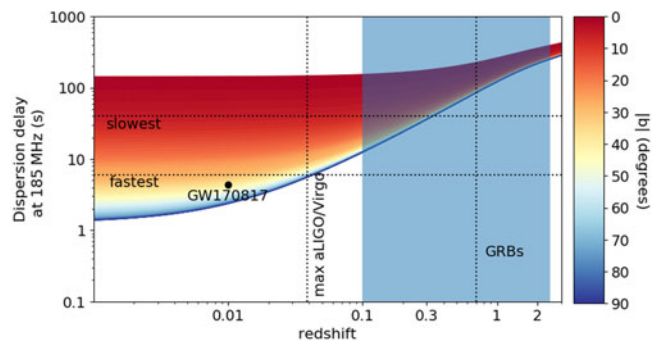


Figure 1. The dispersion delay at 185 MHz due to the intergalactic and interstellar medium as a function of redshift and absolute Galactic Latitude. The Galactic contribution is calculated from the model of Yao et al. (2017), while the intergalactic contribution is computed from the model of Inoue (2004). The redshift range of short GRBs is indicated in the shaded region, with the vertical dashed line representing the average redshift of 0.7. The expected dispersion delay for GW170817 and the horizon for detecting BNS mergers with aLIGO/Virgo during the O3 season are also indicated. The two horizontal lines indicate the fastest and slowest response times for the MWA, which are discussed in Section 3.1.1.

Another benefit of the MWA for follow-up observations is its large field of view. While *Swift* GRBs are usually localised to within a few arcminutes, *Fermi* GRBs can have positional errors of up to tens of degrees, which can be well encompassed by the MWA's large field of view (see Section 4 for further details). In the case of GW events, the triggers from aLIGO/Virgo do not report a single pointing direction, but a probability map that spans many thousands of square degrees. This large area of sky needs to be surveyed quickly in order to catch any prompt emission. Again, the MWA's wide field of view makes it possible to quickly cover these large regions, and the automated rapid response system described here allows for such a tiling of observations (see Kaplan et al. 2016).

Another consideration for both GRBs and GW events is the dispersion delay in the arrival time of any prompt emission at low frequencies, and whether MWA can be on-source in time to detect any such signals that may have been emitted at the time of the merger. Short GRBs have been detected by *Swift* at redshifts of between 0.1 and 2.5 with an average of $z = 0.7$ (Rowlinson et al. 2013). Yao, Manchester, & Wang (2017) model the Galactic contribution of dispersion measure as a function of sky position, while Inoue (2004) provides a model of dispersion measure from the intergalactic medium as a function of redshift. Combining the Galactic and extra-galactic components with the observed redshift range for GRBs, we compute that the arrival time of a prompt signal associated with the event would be dispersion-delayed 12–404 s at 185 MHz and take 4–132 s to cross the MWA's 30.72 MHz bandwidth (Taylor & Cordes 1993). However, the current sensitivity horizon for BNS events during the aLIGO/Virgo O3 run is 170 Mpc, so the expected dispersion delay for GW events can be much lower than that for GRBs at high Galactic Latitude. The expected dispersion delay at 185 MHz as a function of redshift is shown in Figure 1, with the range Galactic dispersion measure contribution shown in colour. The fastest and slowest reaction times of the MWA are indicated by horizontal dashed lines (see Section 3.1.1). We also show the range in redshift for a population of short GRBs, the expected dispersion delay for the BNS merger GW170817, as well as the horizon limit for the aLIGO/Virgo O3 observing run. Figure 1 therefore shows that the MWA will likely be on-target in time to observe prompt radio emission associated with most GRBs, however is unlikely to respond fast enough to

GW events, even if the aLIGO/Virgo alerts were instantaneous (see James *et al.* 2019a, for MWA triggering predictions on negative latency GW alerts).

An additional consideration is the effect of time sampling on the signal detection. The new ability of MWA to trigger higher temporal and spectral resolution observations using the VCS (see Section 3.1.3) will increase the sensitivity to millisecond-duration pulses by at least an order of magnitude due to such pulses no-longer being smeared out over the coarse 0.5 s sampling of the standard correlator (Rowlinson & Anderson (2019) show the MWA VCS to be the most competitive triggering instrument for probing prompt, coherent radio emission from binary mergers). However, there are also deleterious effects of multipath scattering along the line of sight that act in opposition to this sensitivity improvement and are especially potent at low frequencies since the pulse broadening time scale, τ , is strongly frequency-dependent, where typically $\tau \propto \nu^{-4}$ (e.g. Geyer *et al.* 2017; Krishnakumar, Joshi, & Manoharan 2017; Bansal *et al.* 2019; Kirsten *et al.* 2019). The voltage data will be sensitive to this effect and would allow us to resolve and place constraints on the scatter broadening of such pulses. It is not expected that the pulses will be broadened so extremely (i.e. to beyond the 0.5 s correlated observation time sampling) to entirely mitigate the effective sensitivity gained from acquiring the high time resolution time series data. Therefore, the estimated order of magnitude improved sensitivity from capturing voltage data is a reasonably optimistic scenario.

2.2. FRB observations

To date, there are no reported detections of FRBs at frequencies lower than 400 MHz, even though a number of research groups have employed different techniques across a number of telescopes to search for them (e.g. Coenen *et al.* 2014; Tingay *et al.* 2015; Rowlinson *et al.* 2016; Chawla *et al.* 2017). Detecting the low-frequency radio emission from FRBs, if it exists, would give unique insight into the emission energetics and would help to narrow down the progenitors from the large number of current candidates^f (as discussed below). Similarly, the large fractional bandwidths inherent in low-frequency observations are of interest since the spectral modulation of FRBs has a high variance within the population (Shannon *et al.* 2018). Despite this deficit, there is renewed hope at detecting FRBs at these frequencies with the Canadian Hydrogen Intensity Mapping Experiment (CHIME) detection of FRBs at the 400 MHz lower limit of their bandpass (Boyle & CHIME/FRB Collaboration 2018).

The highly dispersed nature of these signals allows for a potentially more efficient use of low-frequency radio telescope time if one can trigger based on real-time detections of FRBs at higher frequencies. In the absence of such an automated triggering service, Sokolowski *et al.* (2018) have made use of shadowing observations to have the MWA co-observe with the Australian SKA Pathfinder (ASKAP, Johnston *et al.* 2008; Hotan *et al.* 2014). In this shadowing setup, ASKAP observes the sky in a fly's-eye mode whilst recording baseband data (James *et al.* 2019b), which is then processed off-line to search for FRBs, resulting in 20 new detections (Bannister *et al.* 2017; Shannon *et al.* 2018). Simultaneously, the MWA observed an overlapping area of sky using the standard correlator mode (10 kHz/0.5 s resolution). For each FRB detection by ASKAP, the MWA data were imaged and analysed for FRB emission. The MWA did not detect any emission from

the ASKAP-detected FRBs in this mode of operation, providing insights into the spectral index of this enigmatic class of objects. Namely, the non-detections are inconsistent with the mean spectral index of $\alpha = -1.8 \pm 0.3$ that is measured for the brightest ASKAP detections. Since pulse broadening cannot explain the non-detections with the MWA, this suggests a spectral turn over at frequencies above 200 MHz and plausible mechanisms are discussed, including: intrinsic spectral behaviour of the radiation process(es), free-free absorption, or caustic/scintillation-induced amplification at higher frequencies (Sokolowski *et al.* 2018).

The current MWA-ASKAP shadowing setup has the MWA operating in the normal correlator mode, as the VCS data rate is too large to sustain for the multiple hours of observations that are carried out each day. The ability for the MWA to operate in a buffered trigger mode, combined with a real-time detection system on ASKAP, would allow the MWA data to be captured at the best time and frequency resolution and only store those that are coincident with an FRB candidate. This would increase the MWA's sensitivity to FRBs by at least an order of magnitude.

For FRBs with a DM as low as 100 pc/cm^3 , which corresponds to a distance of $\sim 130 \text{ Mpc}$ (depending on the line of sight, cf. FRB 171020; Shannon *et al.* 2018), the time of arrival difference between detection at 1 GHz and pulse arrival at 100 MHz is just 40 s. This time delay is long enough that a real-time FRB detection system on a 1 GHz telescope can detect an event, generate a trigger, and pass it to a low-frequency instrument such as the MWA, before the pulse arrives at MHz frequencies. The time delay is short enough, however, that an automated system is required both at the detection and follow-up stations.

2.3. Pulsar observations

Pulsars provide unparalleled Galactic laboratories to study astrophysical coherent emission processes. While the emission from most pulsars is extremely regular, some pulsars exhibit irregular emission, such as sporadic emission (e.g. giant pulses Meyers *et al.* 2017), or the switching of emission states (e.g. Kramer *et al.* 2006; Lorimer *et al.* 2012; Young *et al.* 2014). This sporadic emission occurs on time scales from seconds to months, and these pulsars pose a major challenge to understanding the underlying physics of the pulsar radio emission mechanism. Simultaneous multi-frequency or contemporaneous high-energy and radio observations (e.g. Oronsaye *et al.* 2015; Meyers *et al.* 2017; Abdo *et al.* 2010; Hermsen *et al.* 2013, 2018) suggest that giant pulses and state-switching behaviour are extremely broadband.

The vast majority of southern hemispheric pulsars are lacking low-frequency radio coverage, particularly for sporadic or intermittent pulsars. Low-frequency observations with the MWA are thus very promising as they can reveal emission characteristics that are substantially different to those observed at higher frequencies. In particular, sporadically emitting pulsars are inherently difficult to observe without regular monitoring and long dwell times, which is not currently feasible given the limited observing time available with the VCS (see Section 3.1.3). The newly developed voltage buffer mode has been designed to mitigate these shortcomings, especially the dwell time constraint (see Section 3.1.4 for details).

Both the MWA rapid response triggering and buffering modes allow us to once again bypass the large VCS data recording rate, which is unsustainable for observations much longer than 1 h. For example, telescopes that have regular, real-time pulsar monitoring programmes, such as the recently upgraded Molonglo

^ffrbtheorycat.org

Observatory Synthesis Telescope (UTMOST; Bailes et al. 2017; Venkatraman Krishnan et al. 2019), can trigger the MWA VCS and/or the MWA buffering mode when, for example, intermittent pulsars are active, rather than relying on serendipity and potentially wasting valuable telescope-observing time and resources. The recent detection of low-frequency emission from the intermittent pulsar J1107-5907 (Meyers et al. 2018) provides an excellent demonstration of the MWA's ability to conduct such coordinated broadband observations. As it is, simultaneous broadband observations involving multiple telescopes can provide valuable insights into the pulsar emission mechanism, such as the spectral index distribution (from both integrated profiles and single pulses, e.g. Meyers et al. 2017; Jankowski et al. 2018) and single pulse energy distributions (e.g. Burke-Spolaor et al. 2011; Meyers et al. 2018), both of which are intimately tied to the emission physics.

2.4. M-Dwarf flares

M-dwarf stars are known to produce frequent, powerful flares that are detectable across the entire electromagnetic spectrum. Simultaneous multi-wavelength observations of these flares have provided a window into the processes of plasma acceleration and heating within stellar atmospheres, which drive the flaring emission mechanisms in different wavebands (e.g. Osten et al. 2016, 2010, 2005; Fender et al. 2015b).

At low radio frequencies, stellar flares are often coherent and highly polarised in nature. Recently, the MWA has been successful in detecting faint, polarised flares from the well-known flare star UV Ceti (Lynch et al. 2017), uncovering a population of low-frequency flares two orders of magnitude fainter than single-dish detections made before the 1980s (e.g. Nelson et al. 1979; Spangler, Shawhan, & Rankin 1974; Lovell 1964; Slee 1963).

Some M-type stars exhibit gamma-ray or X-ray 'super-flares', which are bright enough to trigger the *Swift*-BAT system. Recently, AMI triggered on a *Swift* gamma-ray super-flare from the nearby binary system DG CVn, detecting a bright (~ 100 mJy), incoherent flare at 15 GHz, 6 min after the gamma-ray detection, and an additional 90 mJy flare approximately 24 h afterwards (Fender et al. 2015b). These observations at radio frequencies were also accompanied by simultaneous observations using UV and optical facilities (Osten et al. 2016). These multi-wavelength observations enabled a detailed analysis of the flare energetics, the relation of this powerful flare to lower-energy solar flares, and on the potential impact of such flares on the habitability of close-in planets around M-dwarfs (Osten et al. 2016).

It is currently unknown whether there is a low-frequency radio counterpart to the gamma and X-ray emission observed during M-dwarf super-flares. Recently, Argiroffi et al. (2019) reported strong evidence of a coronal mass ejection for HR9024, where X-ray spectroscopy of the stellar flare was used to map plasma motions during the event. One possible source of low-frequency radio emission associated with powerful flares is from Type II bursts, produced during coronal mass ejections (Webb & Howard 2012; Crosley & Osten 2018; Gopalswamy 2006). Rapid response MWA observations of M-dwarf super-flares following triggers from high-energy facilities such as *Swift* and the *Monitor of All-sky X-ray Image (MAXI)* would enable the potential detection of associated prompt, low-frequency flares, and any subsequent low-frequency emission associated with magnetospheric or coronal mass ejection-associated plasma motions.

3. MWA triggering service

The MWA rapid response triggering system is divided into two parts: a back-end and a front-end. The back-end is a web service that has been installed on an on-site server, forming part of the MWA Monitor and Control system. This back-end system accepts requests from clients via the Internet. An entirely separate front-end (which can be run externally or on-site) parses incoming VOEvents, makes decisions about when to trigger a new observation (or re-point an existing triggered observation), and calls the web service to schedule the observations. Multiple front-ends responsible for monitoring and parsing different VOEvent streams and/or transient source types can be run in parallel.

Separating the science (what VOEvents to trigger on and why) from the scheduling function lets the operations team handle the code that directly controls the telescope schedule, while allowing astronomers in the science project teams to write their own parsing code to decide which events to follow and what observation and follow-up strategies to adopt. The front-end system is able to accept pointing and frequency parameters as lists, which will then be iterated over to generate the final set of observing commands. We now describe the two services.

3.1. Back-end web service

The back-end web service includes the following functions that are called by generating an HTTP request to a particular URL with a set of parameters:

- **obslist** - When given a desired override duration, return a summary of all observations already in the schedule over that time period.
- **busy** - When given a science project ID code and a desired override duration (provided in seconds from the present), **busy** returns 'False' if that science project is authorised to override all of the observations already in the schedule over that time period.
- **triggerobs** - When supplied with a science project ID code and associated private key, **trigger_obs** will take a set of observational parameters (described later) and generate a set of observations using the standard MWA correlator. If that science project is authorised to override all the observations already in the schedule over the requested time period, that period in the schedule is cleared and the requested observation/s are added to the schedule.
- **triggervcs** - like **triggerobs**, but schedules observations with the VCS (Tremblay et al. 2015) provided there is enough free disk space on the voltage capture servers.
- **triggerbuffer** - like **triggervcs**, but for use when the VCS is currently in buffered mode (see Section 3.1.3). **triggerbuffer** does not accept observational parameters except for the observing duration, as the other parameters are already set when the telescope is put into the buffer mode. Calling **triggerbuffer** will cause the ring buffer to be drained and for VCS observations to continue as normal afterward.

When deciding if observations can be interrupted, the back-end software considers only the project codes for the existing and requested observations. The prioritisation of transient projects authorised to override active observations or another transient programme is decided by, and at the discretion of the MWA

Director,⁸ and encoded in a configuration file maintained by the operations team. The only other additional constraint is that ongoing VCS observations cannot currently be interrupted, but this may change in the future.

Several input observational parameters can be specified in order to fulfill different science requirements and to optimise the quality of the resulting data. The observational parameters can be given as lists so that observations will be scheduled that span multiple positions on the sky and multiple frequencies. These inputs can include:

- **(ra,dec) | (alt, az) | source:** A pointing direction or source name (from a limited local list of typical targets). Positions and names can be supplied as lists, and the back-end system will tile the sky accordingly.
- **avoidsun:** Whether to modify the given pointing direction/s to keep the desired target near the primary beam centre, but to minimise any contributed power from the Sun by placing it in a primary beam null (see Section 3.1.2). This option has no effect if the Sun is below the horizon.
- **freespec:** One or more frequency specifiers defining the (arbitrary) set of 24 coarse channels to record out of the coarse channels that define the MWA's 80–300 MHz frequency response. Each coarse channel is 1.28 MHz wide making a contiguous bandwidth of 30.72 MHz; however, non-contiguous channel numbers may be specified. Pointing directions and frequency specifiers are duplicated, for example, if two pointing directions are given, and three frequency specifiers, then each target direction will be observed at each of the three chosen frequency sets.
- **nobs, exptime:** The number of observations to schedule for each frequency/pointing combination, and the length of each observation. By default, 15 consecutive observations of 120 s each are scheduled because the MWA analogue beamformers do not track sidereal motion during a single observation.
- **frees, inttime:** The correlator averaging parameters to use - frequency resolution (currently 10, 20, or 40 kHz) and time resolution (currently 0.5, 1.0, or 2.0 s).
- **calibrator, calexptime:** Whether to schedule calibrator observation/s after the triggered source, and if so, what source to calibrate on, and how long the calibrator observation/s should be. The user can also let the system choose a calibrator source automatically, by setting calibrator to be 'True'. If more than one frequency specifier was given, then the calibrator will be observed at each of the given frequency sets.

3.1.1. Latency

Fast response times are extremely important for a rapid response triggering system, with lower latencies allowing astronomers to probe more exotic transient physics. As the MWA is electronically steered, the repointing of the telescope using the rapid response system is not limited by sky slew time, but rather in the automatic canceling and scheduling of observations.

The MWA observing schedule is stored in a set of database tables on a PostgreSQL^h server on-site, with start and stop times stored as the number of seconds since the Global Positioning System (GPS) epoch, referred to as 'GPS seconds'.ⁱ

⁸Including a consideration of the proposal scores assigned by the MWA Time Assignment Committee.

^hwww.postgresql.org

ⁱIS-GPS-200

The MWA schedule works on a natural cadence of 8 s: all observations must start and stop on an integer multiple of 8 GPS seconds. This means that truncating an existing observation and inserting a new observation will cause the new observation to begin on an 8-s boundary, leading to a natural latency of 0–8 s. In practice, the Monitor and Control system gives the various components of the telescope time to prepare by sending their new configuration 4 s ahead of the start of each observation. This means that a running observation cannot have its stop time changed to a value less than 4 s in the future, and a new observation cannot be scheduled to start less than 4 s in the future. This effectively shifts the natural latency to be 4–12 s. The triggering software and web back-end systems have a processing time of around 2 s. The Sun-avoidance check (Section 3.1.2, below) is not run when the Sun is below the horizon, but when it is run can add up to 10 s of latency. Choosing a calibrator automatically can take up to 8 s, and changing the correlator mode to or from the VCS mode requires a dummy observation of 8 s to be inserted into the schedule. Including all the delays and cadences, the total latency period between the arrival of the transient alert notice and the start of a triggered observation is 6–40 s when triggering either the correlator or VCS rapid response mode. For triggers of the voltage buffer, **no such latency exists** as all of the above-mentioned setup is performed as the telescope enters the buffered mode, and the buffer can hold up to 150 s of data, providing a 'negative' latency for this mode.

3.1.2. Sun avoidance

Non-Solar daytime observations with the MWA are both possible and feasible. For example, the rapid response observation of GRB 150424A occurred during the day (Kaplan *et al.* 2015). However, the main concern with daytime observations is the location of the Sun relative to the pointing direction and primary beam side-lobes. Even if the Sun is located within a 1% primary beam sidelobe, the power it contributes to the dataset creates artifacts across any resulting images which are difficult to remove. Ideally, observations can be positioned such that the Sun is in or close to a primary beam null (Figure 2), thus minimising its power contribution. Since the primary beam of the MWA changes rapidly over the 30.72 MHz bandwidth, it may not be possible to place the Sun in a primary beam null over this entire band.

It is possible to adjust the requested pointing direction such that the sensitivity to the Sun is minimised, whilst the sensitivity in the direction of interest is maximised. Therefore, if Sun avoidance is requested for alerts received during daytime, the system selects a pointing direction from the list of 197 MWA 'grid-points' (pointing directions representing exact delays for all tile dipoles) in order to minimise the contribution from the Sun. The ratio, $r = B_{\text{trg}}/B_{\text{sun}}$, of the MWA primary beam sensitivity in the directions of the target (B_{trg}) and the Sun (B_{sun}) calculated at the centre of the observing band is used as a metric. The values of the ratio r for different 'grid-points' can be as low as zero when the target is near the null of the primary beam ($B_{\text{trg}} \rightarrow 0$) and can approach infinity as the Sun approaches the primary beam null ($B_{\text{sun}} \rightarrow 0$). Depending on the relative pointing directions of the target and Sun, the ratio r can vary significantly (by even a few orders of magnitude) over the 30.72 MHz bandwidth. Therefore, a grid point that optimises r at the central frequency may be sub-optimal at either ends of the band. Further revisions can be made during the data reduction level by sacrificing the upper- and lower-most parts of the band, which can result in an overall improvement in the quality and sensitivity of the images.




Figure 2. An all sky map showing: the radio continuum from the Haslam map (background grayscale; Haslam et al. 1982), the location of bright sources from the GLEAM catalogue (blue circles; Hurley-Walker et al. 2017), calibrators (named, white circles), solar system objects, and contours of the MWA primary beam normalised to the pointing direction (green and black). During this observation, the Sun was placed into a 0.1% sidelobe, by the Sun avoidance code.




Figure 3. The dependence of the standard deviation of noise in Stokes V images on the ratio ($r = B_{trg} / B_{sun}$) of the primary beam sensitivity in the directions of the target (B_{trg}) and the Sun (B_{sun}). The noise increases noticeably at ratios $r < 1000$.

The optimal pointing is required to satisfy conditions: $r > 1000$ and $B_{trg} \geq B_{min}$, where B_{min} is the minimum acceptable beam response in the direction of the target ($B_{min} = 50\%$). This strategy is supported by Figure 3 which shows that the higher the ratio r , the better the properties of the image. These initial criteria sometimes cannot be met, in which case we relax the criteria to allow $B_{min} \geq 10\%$. This allows the software to quickly find a pointing direction with minimal response towards the Sun (i.e. maximising the ratio r), which still provides at least 10% beam sensitivity towards a target. Furthermore, the primary beam computations are one of the main contributions to the overall latency of the trigger system (see Section 3.1.1) so we opt not to decrease B_{min} in stages, but instead jump directly to the minimum acceptable limit.

In order to test the Sun avoidance procedure, we performed daytime observations in all grid-point directions above elevation 50° . The data quality was evaluated in terms of the standard deviation of the noise in Stokes V images (σ_v), since the image noise

is not side-lobe confusion limited in circular polarisation as it would be in total intensity. The 25-s Stokes V images were of the quality expected for the night sky ($\sigma_v \approx 2 \text{ mJy beam}^{-1}$, based on predictions using the 2016 MWA beam model; Sokolowski et al. 2017) for pointing directions where $r > 1000$, whilst σ_v noticeably increased in the pointing directions with $r < 1000$ (Figure 3).

3.1.3. Voltage capture mode

The MWA is capable of capturing high time and frequency resolution data using the VCS. The VCS observing mode has a substantial data rate ($\sim 28 \text{ TB/h}$), far greater than the rate at which it can be transported to the MWA archive in Perth in real-time so the data are initially stored on-site. The on-site storage currently limits VCS observations to $\sim 90 \text{ min}$, so whilst it is possible to observe with the VCS and possible to trigger this mode with the back-end mentioned in Section 3.1, additional constraints need to be placed on such observations. To ensure that future VCS observations are not disrupted by a VCS trigger, a check is made to determine the current and expected disk use on-site at the MWA. If the requested triggered observations would result in future observations failing due to insufficient disk space, then the observing request will be rejected.

Each VCS observation stores data across 32 RAIDs simultaneously (2 per VCS server) so when a `triggervcs` call is made, the software determines which RAID has the least free space and, based on that value, calculates how much time can be recorded. The triggering system then checks what VCS observations are scheduled for the next 24 h¹ and subtracts this scheduled observing time from the total recording time available. If the duration of the requested trigger observation is shorter than the remaining time, then a VCS trigger is allowed.

3.1.4. Voltage buffer mode

There are some science cases where it is necessary to have the MWA actively follow the pointing direction of other telescopes (henceforth referred to as ‘shadowing’), and the optimal observing strategy also requires the high time and frequency resolution provided by the VCS (e.g. ASKAP FRBs; see Section 2.2). Unfortunately, the recording limit of $\sim 90 \text{ min}$ is not conducive to shadowing observations. To address this, a new VCS observing mode was developed where the critically sampled ($100 \mu\text{s}/10 \text{ kHz}$ time and frequency resolution) tile voltages are stored in a ring buffer within the on-board memory of the VCS servers until a trigger is received, thereby mitigating the recording limit. In this buffer mode, the pointing direction and frequency selection is set prior to observing, and the telescope is collecting (but not always recording) data for the duration of this observation. As such, normal telescope observations cannot continue when this mode is selected, unlike normal triggered observations that interrupt and override normal scheduled observations.

While operating in this mode, the voltage streams from all 128 tiles are stored in memory for up to 150 s before being discarded on a first-in-first-out basis, so at any given time, the last 150 s of data from the telescope is buffered in memory. When a trigger is received in this mode, the VCS software begins to write the buffered data to disk, while continuing to record new data to memory. This process runs in at least real time; thus, after $\sim 2 - 3 \text{ min}$, the VCS servers have drained the memory buffers and will be operating as in a normal VCS observation, where the data

¹For the current system. This is a configurable parameter.

are written directly to disk. Voltages will continue to be written to disk until the requested trigger stop time is reached. This mode is non-interrupting, in that the telescope is technically observing in Voltage Capture Mode (see Section 3.1.3) for the scheduled duration even though potentially no data are being written to disk and should therefore be thought of as a much more efficient, pre-scheduled VCS shadowing observation. Throughout the nominated observing time, the software will automatically re-point the telescope at a given duty cycle (typically 5–15 min) so that the target position is always within the tile primary beam. It is possible to optimise these pointing directions in advance using the Sun avoidance mentioned in Section 3.1.2.

The voltage buffering functionality potentially allows for long-duration (i.e. >90 min) VCS observations where only those data containing an event are actually written to disk, thereby significantly reducing the overall data rate while providing a more versatile shadowing capability to the MWA. Additionally, the buffer mode is critical for observations where the delay between high- and low-frequency emission is not large (e.g. $\lesssim 40$ s). For instance, sporadic emission from pulsars, where the dispersive delays can be many seconds, but still not long enough to process, send, and receive triggers, re-points the MWA and begin VCS recording before the pulse arrives at ~ 150 MHz.

An early prototype of this observing mode has been tested in exactly this circumstance, when UTMOST was used to trigger an MWA VCS recording of individual pulses from the intermittent pulsar J1107–5907 (Meyers *et al.* 2018). In this case, the dispersion delay ($DM = 40.75 \text{ pc cm}^{-3}$) between the two observing bands (835 and 154 MHz for UTMOST and the MWA, respectively) is only ~ 8 s, shorter than the typical trigger latency, but easily within the voltage buffer capacity.

3.2. VOEvent handler front end

The front end code uses the ‘4 Pi Sky VOEvent Broker’ (Staley & Fender 2016) and the COMET VOEvent client (Swinbank 2014). The COMET client operates by receiving VOEvents and then making asynchronous calls to an external programme. We implement a simple queue to ensure that events are processed serially and in the order in which they are received. The script `PUSH_VOEVENT.PY` is called by COMET and pushes the received VOEvent onto a handler service (`VOEVENT_HANDLER.PY`) via a Remote Procedure Call. The VOEvents are then passed to one or more registered handler functions in plugin libraries, which process the VOEvents using the `VOEVENT-PARSE` python module (Staley *et al.* 2013; Staley 2014) and submit observing requests if required. The plugin event handlers return either True or False depending on whether an observing request was made. `VOEVENT_HANDLER.PY` will pass a VOEvent to all registered handlers stopping once a trigger request has been made or the list has been exhausted. The module `TRIGGERSERVICE.PY` provides wrapper functions that abstracts the calls to the back-end web services: `can_interrupt_now`, `obslist`, `triggerobs`, and `triggervcs`. The module `HANDLERS.PY` provides a `TriggerEvent` object that holds information about triggered events, including a cache of VOEvents associated with the event. It also provides a high-level wrapper for triggering observations and also sends email notifications relating to the trigger. Additionally, `HANDLERS.PY` provides a wrapper function for retrieving the positional information from a VOEvent. All the scripts noted thus far are provided to allow MWA science teams to write python scripts with a focus on parsing the information within the VOEvent without having to worry



Figure 4. Logical flow for receiving a *Swift* GRB alert. The outcomes are either to trigger a new observation, update the current observation (green), or to not observe (red).

about the underlying data format or the HTTP requests that are used to trigger observations.

The front-end service is available at github.com/MWATelescope/mwa_trigger under an Academic Free License (AFL 3.0), along with the VOEvent parser for the GRB follow-up described below.

4. Case study: GRB follow-up

In order to parse VOEvents of high-priority GRBs to the MWA back-end web service described in Section 3.1, we created the script `GRB_FERMI_SWIFT.PY`, which is capable of filtering GRBs detected by the *Fermi*-GBM and *Swift*-BAT instruments. *Swift*-detected events are less common than *Fermi* as the BAT instrument can only observe one-sixth of the sky when compared to the GBM, which monitors $\sim 50\%$ of the sky at a time. However, *Swift*-BAT events have a far better positional accuracy (1 arcmin–4 arcmin; Gehrels *et al.* 2004) compared to the *Fermi*-GBM (with initial positional errors that are several 10 s of degrees, with final position accuracy usually $< 10^\circ$; Connaughton *et al.* 2015). We therefore give *swift*-detected events priority over *Fermi*-detected events. We also prioritise potential short GRBs over long GRBs, due to their association with gravitational wave events, and given that the majority of the emission mechanisms are expected to escape along the jet axis, they are less likely to be absorbed by the merger ejecta, which is expected to be more concentrated along the equatorial plane (Zhang 2014). The high-level decision tree that `GRB_FERMI_SWIFT.PY` implements for parsing *Swift* and *Fermi* notices to MWA is outlined in Figures 4 and 5.

In Figures 4 and 5, the box ‘trigger new observation’ represents a request for a new observation, which does not guarantee that an observation will take place. The back-end triggering service makes the necessary checks to ensure that the requested observations are permitted to override the current observing programme before they are scheduled. We place an additional

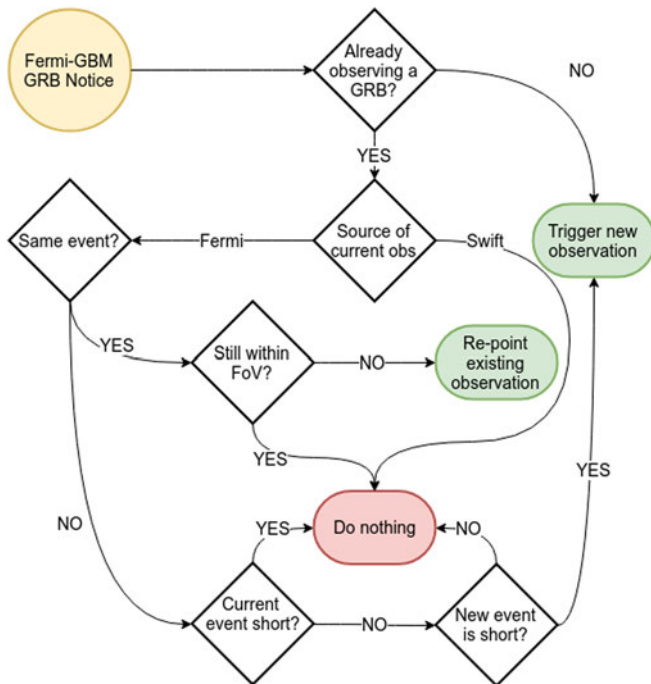


Figure 5. Logical flow for receiving a *Fermi* GRB alert. The outcomes are either to trigger a new observation, update the current observation (green), or to not observe (red).

constraint that prevents any observing requests for a GRB that is below an elevation limit of 30° . *Fermi* VOEvents also come with a MOST_LIKELY parameter, which indicates the class of object that is most likely to have generated the *Fermi* alert. We reject any alerts for which the likelihood of the event being a GRB is $<50\%$. All *Swift* events are automatically compared to the Onboard Source Catalogue (OSC). If all known transients can be rejected, the *Swift* alert will automatically define the trigger as a real GRB, again allowing for further VOEvent filtering. Note that the automatic *Swift* comparison to the OSC means that VOEvents can also be monitored for other source types, such as known X-ray binaries, flare stars and magnetars, potentially motivating other triggering programmes.

For all alerts, we make the assumption that the most recent notice will have the most up-to-date information and so we will re-point to the newest position if it is not in the field of view^k of the current observation. The exception is that *Fermi* alerts arrive in three different flavours: flight position, ground position, and final position. For a single event, *Fermi* can generate zero, one, or more of each flavour of alert and not always in the same order. We ascribe a positional reliability hierarchy to the *Fermi* alerts with the highest to lowest reliability being: final > ground > flight. Notices of the same flavour will cause a position update check, but those of lower reliability cannot update positions previously generated from a higher priority alert. For example, a ground position notice will generate a re-pointing request if the best position was generated from a previous ground position or flight position alert but not if it was generated from a final position alert.

A given GRB can also be detected by both *Swift* and *Fermi*. The *Fermi*-GBM triggers often arrive before the *Swift*-BAT triggers (e.g. GRB 130427A; Anderson et al. 2014). Given the much better positional precision of *Swift* generated events, they are prioritised over the *Fermi* alerts. We therefore make no effort to cross-identify

^kIn practice we define a minimum angular offset (default is 10 deg), and only new positions that are more distant than this will cause a new trigger.

GRBs that are detected by both telescopes. Such GRBs will always generate observing requests and will do so at the more precise *Swift* position. While we prioritise the follow-up of short GRBs, it is not possible to determine the long/short nature of a GRB at the time of detection, the classification of which usually requires human inspection of the gamma-ray data. We therefore only trigger on those *Fermi* events for which the integration time (the time in seconds it takes the gamma-ray signal to reach the triggering threshold of the instrument, which is different to the burst duration) is ≤ 2.048 s. This 2 s is based on the classification of the short GRB class (Kouveliotou et al. 1993), whilst the fractional quantity incorporates the resolution of the duration in *Fermi* notices. We use the same cut-off for *Swift* events, allowing us to re-point at any new *Swift* triggers that are more likely to be a short GRB (see Figure 4).

5. Summary and future development

We have developed a VOEvent-based rapid response system for the MWA telescope. This development was motivated by the need to perform low-latency observations of GRBs detected by the *Swift* and *Fermi* telescopes and to enable efficient use of the voltage capture system for the study of GRBs, FRBs, and intermittent pulsars. The software runs in two parts: a back-end system that is part of the MWA Monitor and Control system, which exposes multiple web interfaces, and a front-end system that can be run anywhere, responds to VOEvents, and interacts with the back-end system. The front-end system, along with example VOEvent handlers, is available at github.com/MWATelescope/mwa_trigger. Contributions, bug reports, and feature requests are encouraged.

The system is still being developed, and new features are being planned. The latency of the back-end system is currently 6–040 s. A natural 4–12 s latency exists as part of the scheduling cycle of the MWA, and this may be reduced as part of a future upgrade to the MWA. The remaining 2–28 s latency is due to: Sun-avoidance calculations, calibrator selection and scheduling, and correlator mode change operations. The Sun is not the only bright source that causes consternation: bright radio sources in the primary beam side-lobes can also cause calibration issues during both day time and night time observations. An extension of the Sun avoidance code is planned for night time observations, so as to place these difficult sources in the null of the primary beam at the central observing frequency. Currently, the back-end system inserts observations into the schedule by expunging and truncating existing observations, and there is not yet any concept of cancelling a triggered observation. In the future, we plan to have the deleted observations cached so that if a trigger is cancelled, the MWA schedule can be re-instated, reducing the impact on other observers.

This work represents a new mode of operation for the MWA. The upgraded standard correlator triggering capability has been implemented since MWA observing semester 2018B, and the VCS and buffered mode triggers will become available for observing in a future semester.

Acknowledgements. We thank the referee for their comments and suggestions, which improved the quality of this publication. We thank Natasha Hurley-Walker, Freya North-Hickey, and Jean-Pierre Macquart for their suggestions that led to the creation of Figure 1. We thank the current and previous MWA Directors, Melanie Johnston-Hollitt and Randall Wayth, for providing director’s time during the development and testing of this service. We acknowledge the Wajarri Yamatji people as the traditional owners of the site on which the MWA operates.

Funding. GEA is the recipient of an Australian Research Council Discovery Early Career Researcher Award (project number DE180100346) funded by the Australian Government. BWM acknowledges the contribution of an Australian Government Research Training Program Scholarship in supporting this research. SET and BWM acknowledge the Australian Research Council grant CE110001020 (CAASTRO).

Software. The front end software makes use of COMET (Swinbank 2014) and VOEVENT-PARSE (Staley & Fender 2016) to process the incoming VOEvents. Both the front-end and back-end software rely on the Astropy (The Astropy Collaboration et al. 2013; Astropy Collaboration et al. 2018), numpy (van der Walt, Colbert, & Varoquaux 2011), and scipy (Jones et al. 2001) python modules.

Facilities. This scientific work makes use of the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. We acknowledge the Pawsey Supercomputing Centre, which is supported by the Western Australian and Australian Governments. This work made use of data supplied by the *Swift* satellite. *Swift*, launched in November 2004, is a NASA mission in partnership with the Italian Space Agency and the UK Space Agency. *Swift* is managed by NASA Goddard. Penn State University controls science and flight operations from the Mission Operations Center in University Park, Pennsylvania. Los Alamos National Laboratory provides gamma-ray imaging analysis.

Services. This research has made use of NASA's Astrophysics Data System and the online cosmology calculator tool¹ Wright (2006).

References

- Abbott, B. P., et al. 2017a, *ApJ*, 848, L12
 Abbott, B. P., et al. 2017b, *ApJ*, 848, L13
 Abdo, A. A., et al. 2010, *ApJ*, 708, 1254
 Anderson, G. E., et al. 2014, *MNRAS*, 440, 2059
 Anderson, G. E., et al. 2018a, *MNRAS*, 473, 1512
 Anderson, G. E., et al. 2018b, GRB Coordinates Network, 23467, 1
 Anderson, M. M., et al. 2018c, *ApJ*, 864, 22
 Andreoni, I., et al. 2017, *PASA*, 34, e069
 Argiroffi, C., et al. 2019, *NA*, p. 328
 Astropy Collaboration et al. 2018, *AJ*, 156, 123
 Bailes, M., et al. 2017, *PASA*, 34, e045
 Bannister, K. W., Murphy T., Gaensler B. M., & Reynolds J. E. 2012, *ApJ*, 757, 38
 Bannister, K., et al. 2017, *ApJL*, 841, L12
 Bansal, K., Taylor, G. B., Stovall K., & Dowell J. 2019, *ApJ*, 875, 146
 Barthelmy, S. D., et al. 2005, *SSR*, 120, 143
 Beardsley, A. P., et al. 2019, Science with the Murchison Widefield Array: Phase I Results and Phase II Opportunities
 Bowman, J. D., et al. 2013, *PASA*, 30, e031
 Boyle, P. C., & CHIME/FRB Collaboration 2018, ATEL, 11901, 1
 Burke-Spolaor, S., et al. 2011, *MNRAS*, 416, 2465
 Callister, T. A., et al. 2019, *ApJ*, 877, L39
 Chawla, P., et al. 2017, *ApJ*, 844, 140
 Coenen, T., et al. 2014, *AA*, 570, A60
 Connaughton, V., et al. 2015, *AJSS*, 216, 32
 Crosley, M. K., & Osten, R. A. 2018, *ApJ*, 862, 113
 Dessenne, C. A. C., et al. 1996, *MNRAS*, 281, 977
 Ellingson, S. W., Craig, J., Dowell, J., Taylor, G. B., & Helmboldt, J. F. 2013, arXiv e-prints, p. arXiv:1307.0697
 Falcke, H., & Rezzolla, L. 2014, *A&A*, 562, A137
 Fender, R., Stewart, A., Macquart, J. P., Donnarumma, I., Murphy, T., Deller, A., Paragi, Z., & Chatterjee, S. 2015a, in Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 51
 Fender, R. P., Anderson, G. E., Osten, R., Staley, T., Rumsey, C., Grainge, K., & Saunders, R. D. E. 2015b, *MNRAS*, 446, L66
 Fender, R. P., & Bell, M. E. 2011, *BASI*, 39, 315
 Frail, D. A., Kulkarni, S. R., Nicastro, L., Feroci, M., & Taylor, G. B. 1997, *Nature*, 389, 261
 Galama, T. J., et al. 1999, *Nature*, 398, 394
 Gehrels, N., et al. 2004, *ApJ*, 611, 1005
 Geyer, M., et al. 2017, *MNRAS*, 470, 2659
 Gopalswamy, N. 2006, *AGU, Geophysical Monograph Series, Washington DC*, 165, 207
 Green, D. A., et al. 1995, *ASS*, 231, 281
 Hansen, B. M. S., & Lyutikov, M. 2001, *MNRAS*, 322, 695
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, 47, 1
 Hermsen, W., et al. 2013, *Science*, 339, 436
 Hermsen, W., et al. 2018, *MNRAS*, 480, 3655
 Hotan, A. W., et al. 2014, *PASA*, 31, e041
 Hurley-Walker, N., et al. 2017, *MNRAS*, 464, 1146
 Inoue, S. 2004, *MNRAS*, 348, 999
 James, C. W., Anderson, G. E., Wen, L., Bosveld, J., Chu, Q., Kovalam, M., Slaven-Blair, T. J., & Williams, A. 2019a, *MNRAS Letters*, slz129
 James, C. W., et al. 2019b, *PASA*, 36, e009
 Jankowski, F., van Straten, W., Keane, E. F., Bailes, M., Barr, E. D., Johnston, S., & Kerr, M. 2018, *MNRAS*, 473, 4436
 Johnston, S., et al. 2008, *EA*, 22, 151
 Jones, E., et al. 2001, SciPy: Open Source Scientific Tools for Python
 Kaplan, D. L., et al. 2015, *ApJL*, 814, 6
 Kaplan, D. L., Murphy, T., Rowlinson, A., Croft, S. D., Wayth, R. B., & Trott, C. M. 2016, *PASA*, 33, id.e050 7 pp. 33
 Kirsten, F., Bhat, N. D. R., Meyers, B. W., Macquart, J.-P., Tremblay, S. E., & Ord, S. M., 2019, *ApJ*, 874, 179
 Kouveliotou, C., Meegan, C. A., Fishman, G. J., Bhat, N. P., Briggs, M. S., Koshut, T. M., Paciesas, W. S., & Pendleton, G. N. 1993, *ApJ*, 413, L101
 Kramer, M., Lyne, A. G., O'Brien, J. T., Jordan, C. A., & Lorimer, D. R. 2006, *Science*, 312, 549
 Krishnakumar, M. A., Joshi, B. C., & Manoharan, P. K. 2017, *ApJ*, 846, 104
 Lasky, P. D., Haskell, B., Ravi, V., Howell, E. J., & Coward, D. M. 2014, *PRD*, 89, 047302
 Lattimer, J. M. 2012, *ARNPS*, 62, 485
 Lorimer, D. R., Lyne, A. G., McLaughlin, M. A., Kramer, M., Pavlov, G. G., & Chang, C. 2012, *ApJ*, 758, 141
 Lovell, B. 1964, *Nature*, 203, 1213
 Lynch, C. R., Lenc, E., Kaplan, D. L., Murphy, T., & Anderson, G. E. 2017, *ApJ*, 836, L30
 Lyutikov, M. 2013, *ApJ*, 768, 63
 Meegan, C., et al. 2009, *AJ*, 702, 791
 Metzger, B. D., Berger, E., & Margalit, B. 2017, *ApJ*, 841, 14
 Meyers, B. W., et al. 2017, *ApJ*, 851, 20
 Meyers, B. W., et al. 2018, preprint, 1811, arXiv:1811.01142
 Nelson, G. J., Robinson, R. D., Slee, O. B., Fielding, G., Page, A. A., & Walker, W. S. G. 1979, *MNRAS*, 187, 405
 Obenberger, K. S., et al. 2014, *ApJ*, 785, 27
 Oronsaye, S. I., et al. 2015, *ApJ*, 809, 51
 Osten, R. A., Hawley, S. L., Allred, J. C., Johns-Krull, C. M., & Roark, C. 2005, *ApJ*, 621, 398
 Osten, R. A., et al. 2010, *ApJ*, 721, 785
 Osten, R. A., et al. 2016, *ApJ*, 832, 174
 Palaniswamy, D., Wayth, R. B., Trott, C. M., McCallum, J. N., Tingay, S. J., & Reynolds, C. 2014, *ApJ*, 790, 63
 Prasad, P., Wijnholds, S. J., Huizinga, F., & Wijers, R. A. M. J. 2014, *A&A*, 568, A48
 Prasad, P., et al. 2016, *JAI*, 5, 1641008
 Rees, M. J. 1988, *Nature*, 333, 523
 Rowlinson, A., Anderson, G. E. 2019, arXiv e-prints, p. arXiv:1905.02509
 Rowlinson, A., O'Brien, P. T., Metzger, B. D., Tanvir, N. R., & Levan, A. J. 2013, *MNRAS*, 430, 1061

¹www.astro.ucla.edu/~wright/CosmoCalc.html

- Rowlinson, A., et al. 2016, *MNRAS*, 458
- Seaman, R., et al. 2011, *IVOA Recommendation*, 11 July 201, p. 711
- Shannon, R. M., et al. 2018, *Nature*, 562, 386
- Slee, O. B. 1963, *Nature*, 199, 991
- Soderberg, A. M., et al. 2008, *Nature*, 453, 469
- Sokolowski, M., et al. 2017, *PASA*, 34, e062
- Sokolowski, M., et al. 2018, *ApJL*, 867, L12
- Spangler, S. R., Shawhan, S. D., & Rankin, J. M. 1974, *ApJ*, 190, L129
- Staley, T. D. 2014, Astrophysics Source Code Library, record ascl:1411.003
- Staley, T. D., & Fender R. 2016, arXiv e-prints, p. [arXiv:1606.03735](https://arxiv.org/abs/1606.03735)
- Staley, T. D., et al. 2013, *MNRAS*, 428, 3114
- Swinbank, J. J. 2014, *A&C*, 7, 12
- Taylor, G. B., et al. 2012, *JAI*, 1, 1250004
- Taylor, J. H., & Cordes J. M. 1993, *ApJ*, 411, 674
- The Astropy Collaboration et al. 2013, *A&A*, 558, 9
- Thornton, D., et al. 2013, *Science*, 341, 53
- Tingay, S. J., et al. 2013, *PASA*, 30, 21
- Tingay, S. J., et al. 2015, *AJ*, 150, 1
- Totani, T. 2013, *PASJ*, 65, L12
- Tremblay, S. E., et al. 2015, *PASA*, 32, e005
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CS&E*, 13, 22
- Venkatraman Krishnan, V. et al. 2019, arXiv e-prints, p. [arXiv:1905.02415](https://arxiv.org/abs/1905.02415)
- Vestrand, W. T., et al. 2014, *Science*, 343, 38
- Wayth, R. B., et al. 2018, *PASA*, 35
- Webb, D. F., & Howard T. A. 2012, *LRSP*, 9, 3
- Williams, P. K. G., & Berger E. 2016, *ApJL*, 821, L22
- Wright, E. L. 2006, *PASP*, 118, 1711
- Yancey, C. C., et al. 2015, *ApJ*, 812, 168
- Yao, J. M., Manchester, R. N., & Wang, N. 2017, *ApJ*, 835, 29
- Young, N. J., Weltevrede, P., Stappers, B. W., Lyne, A. G., & Kramer, M. 2014, *MNRAS*, 442, 2519
- Zhang, B. 2014, *ApJ*, 780, L21