Astronomy ESFRI & Research Infrastructure Cluster

Ref. Ares(2017)2719119 - 30/05/2017

ASTERICS - H2020 - 653477

Multi-messenger alert handling design

ASTERICS GA DELIVERABLE: D5.2

Document identifier:	ASTERICS-D5.2-final.docx
Date:	**2017-05-11**
Work Package:	**WP5 – CLEOPATRA**
Lead Partner:	**JIVE**
Document Status:	**Report**
Dissemination level:	**Public**
Document Link:	www.asterics2020.eu/documents/ ASTERICS-D5.2.pdf

Abstract

High-level design for a multi-messenger alert system for the automatic reaction to astronomical transient in any detectable form with a particular focus on following up

Gravitational Wave events from LIGO/Virgo at radio wavelengths by LOFAR and the EVN.

I. COPYRIGHT NOTICE

Copyright © Members of the ASTERICS Collaboration, 2015. See www.asterics2020.eu for details of the ASTERICS project and the collaboration. ASTERICS (Astronomy ESFRI & Research Infrastructure Cluster) is a project funded by the European Commission as a Research and Innovation Actions (RIA) within the H2020 Framework Programme. ASTERICS began in May 2015 and will run for 4 years.
This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nc/3.0/ or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: "Copyright © Members of the ASTERICS Collaboration, 2015. See www.asterics2020.eu for details of the ASTERICS project and the collaboration". Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/WP	Date
From			
Author(s)	Mark Kettenis Jan David Mol Sander ter Veen Antonia Rawlinson Eric Chassande-Mottin	JIVE ASTRON ASTRON ASTRON CNRS-APC	
Reviewed by	Giuseppe Cimò	JIVE / ASTRON	
Approved by	AMST		29/05/2017

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1			
2			
3			
4			

IV. APPLICATON AREA

This document is a formal deliverable for the GA of the project, applicable to all members of the ASTERICS project, beneficiaries and third parties, as well as its collaborating projects.

V. TERMINOLOGY

A complete project glossary is provided at the following page:
http://www.asterics2020.eu/about/glossary/

DMZ	Demilitarized Zone
EVN	European VLBI Array
GCN	Gamma-ray Coordination Network
GW	Gravitational Wave
LIGO	Laser Interferometer Gravitational-Wave Observatory
LDAP	Lightweight Directory Access Protocol
LOFAR	Low Frequency Array
VTP	VOEvent Transport Protocol
XSD	XML Schema Definition

VI. PROJECT SUMMARY

ASTERICS (Astronomy ESFRI & Research Infrastructure Cluster) aims to address the cross-cutting synergies and common challenges shared by the various Astronomy ESFRI facilities (SKA, CTA, KM3Net & E-ELT). It brings together for the first time, the astronomy, astrophysics and particle astrophysics communities, in addition to other related research infrastructures. The major objectives of ASTERICS are to support and accelerate the implementation of the ESFRI telescopes, to enhance their performance beyond the current state-of-the-art, and to see them interoperate as an integrated, multi-wavelength and multi-messenger facility. An important focal point is the management, processing and scientific exploitation of the huge datasets the ESFRI facilities will generate. ASTERICS will seek solutions to these problems outside of the traditional channels by directly engaging and collaborating with industry and specialised SMEs. The various ESFRI pathfinders and precursors will present the perfect proving ground for new methodologies and prototype systems. In addition, ASTERICS will enable astronomers from across the member states to have broad access to the reduced data products of the ESFRI telescopes via a seamless interface to the Virtual Observatory framework. This will massively increase the scientific impact of the telescopes, and greatly encourage use (and re-use) of the data in new and novel ways, typically not foreseen in the original proposals. By demonstrating cross-facility synchronicity, and by harmonising various policy aspects, ASTERICS will realise a distributed and interoperable approach that ushers in a new multi-messenger era for astronomy. Through an active dissemination programme, including direct engagement with all relevant stakeholders, and via the development of citizen scientist mass participation experiments, ASTERICS has the ambition to be a flagship for the scientific, industrial and societal impact ESFRI projects can deliver.

VII. EXECUTIVE SUMMARY

This document provides a high-level design for a multi-messenger alert system for the automatic follow-up of Gravitational Wave events from LIGO/Virgo at radio wavelengths by LOFAR and the EVN. It contains the scientific motivations, requirements and implementation design for such a system. It forms deliverable D5.2 for the CLEOPATRA work package (WP5) in the ASTERICS project.

Table of Contents

1 Introduction

Astronomy has traditionally focussed on observing the electromagnetic spectrum. In many cases, scientific discoveries are made by observing different parts of this spectrum. This is referred to as multi-wavelength astronomy. Recent breakthroughs in observing high-energy particles such as neutrinos and gravitational waves have extended the range of possibilities again, giving birth to the multi-messenger astronomy. Many of these observations concern events that are transient in nature. In order to truly explore the possibilities offered by multi-messenger techniques, it is important to follow-up events detected by one instrument with observations in other parts of the electromagnetic spectrum and, possibly, types of particles. The goal of task 2 within the CLEOPATRA, Work Package 5 of the ASTERCIS project, is to provide tools that facilitate the generation and distribution of these types of events, and how to react to them for helping coordinate follow-up observations. This task also intends to design a system that handles multi-messenger alerts in the context of gravitational wave observations.

The first part of this document (section 2) describes the science cases for transient alerts for the instruments involved in the task: EGO/LIGO, LOFAR and the EVN. The science cases serve as a background context from which the requirements for the system are derived. This section also provides some relevant technical information about the instruments that is relevant for observing transients. Section 3 describes the requirements for the alert handling mechanisms, which is followed by an analysis of the requirements in section 4. The results of this analysis are then used, in section 5, to formulate design principles for generation and dissemination of events, reaction to events, and coordination of follow-up observations. An outline of the plan for a demonstrator is given in section 6, followed by our preliminary conclusions. A more comprehensive set of conclusions will be provided after the demonstrator, which is deliverable D5.13 of the ASTERICS project.

2 Science Cases & Background

2.1 Transients and Gravitational Waves

In Mid-September 2015, the two detectors of the Laser Interferometer Gravitational-wave Observatory (LIGO) started a science run in an advanced configuration achieving an unprecedented sensitivity, about three times better than previous run at 200 Hz, and 100 times better in the low frequency range around 50 Hz. During this science run, the gravitational-wave (GW) signal from a distant binary black-hole merger was observed for the first time [1]. A new era for transient astronomy began with this detection.

It is important to establish connections between conventional astronomy and this new type of observations. The energy released in gravitational waves during a compact binary merger is enormous. The fraction of the rest mass emitted for binaries of stellar-mass objects ranges from one to ten percent. For instance, the source of GW150914 emitted about $3M_{\odot}c^2 \sim 10^{54}$erg [2]. For sources that are within the LIGO/Virgo horizon, if a small fraction of that energy, say a percent, is converted into electromagnetic radiation over timescales comparable to the merger, it would produce a bright electromagnetic counterpart, similar in luminosity to the brightest observed astrophysical transients.

Also, short gamma-ray bursts have long been proposed to be connected with compact binary mergers (including at least one neutron star) [3,4,5]. If this turns out to be true, gravitational wave events from binary mergers would then be associated with a prompt gamma-ray flash and multi-wavelength afterglow for "on-axis" binaries, i.e. with the binary orbital rotation axis within 20 degrees of the line of sight, for a recent review see [6]. For "off-axis" events there are predictions of an isotropic infra-red counterpart, known as a kilonova [e.g. 7,8]. The first candidate kilonova has been observed following a short gamma-ray burst, adding strength to the association between short gamma-ray bursts and compact binary mergers [9]. Additionally, there are expected to be isotropic radio counterparts at late times [e.g. 10]. Coherent, short duration radiobursts associated compact binary mergers involving at least one neutron star may occur at the time of merger [e.g. 11] or up to a few hundreds of minutes afterwards [e.g. 12].

The above energy budget argument and the existence of well-established astrophysical scenarios that connects gravitational-wave and electromagnetic emissions provide the basic motivations to pursue joint searches for both signals. Two complementary approaches have

---

[1] B. P. Abbott et al., Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* 116: 0611
[2] B. P. Abbott et al., Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* 116: 0611
[3] Lattimer J. M., Schramm D. N., 1976, ApJ, 210, 549
[4] Eichler D., Livio M., Piran T., Schramm D. N., 1989, Nature, 340, 126
[5] Gehrels N., et al., 2005, Nature, 437, 851
[6] Berger E., 2014, ARA&A, 52, 43
[7] Li L.-X., Paczynski B., 1998, ApJ, 507, L59
[8] Barnes J., Kasen D., 2013, ApJ, 775, 18
[9] Tanvir N. R., et al., 2013, Nature, 500, 547
[10] Hotokezaka K., et al., 2016, ApJ, 831, 190
[11] Usov V. V., Katz J. I., 2000, A&A, 364, 655
[12] Zhang, B, 2014, ApJL, 780, L21

been followed so far. Gravitational-wave follow-up searches in association to a variety of relevant astrophysical events, such as gamma-ray bursts, pulsar glitches, soft-gamma repeater flares, fast radio bursts, near-by supernovae, etc. have been completed and published [13]. The reverse strategy i.e., electromagnetic follow-up of gravitational-wave events will be the focus of this here.

Since 2014, LIGO and Virgo (its European counterpart) have tight collaborative agreements with more than 80 astronomer teams around the world. The instruments used by those teams cover a wide range of wavelengths from radio to very high energies. Alerts are generated from the gravitational-wave data and communicated to those teams through a dedicated network. This will change once gravitational-wave detections become routine [14]: high-confidence detections will lead to publicly released alert.

2.2 Triggered LOFAR observations

LOFAR (the LOw Frequency ARray) is a revolutionary radio telescope, operating at 10-90 and 110-240 MHz, comprising of many antennas grouped together into stations and the full array. The LOFAR Transients Key Science Project (TKSP) is conducting a range of blind transient searches, using dedicated and commensal surveys, and also triggered observations on known transient sources. Automatically handling multi-messenger transient alerts and producing transient alerts are an essential part of the science case for the TKSP.

Currently triggered observations are conducted manually via coordination with the LOFAR Radio Observatory (RO) and a minimum response timescale of 30 minutes within office hours with prior agreement. Via this strategy, the TKSP have successfully triggered on two gravitational wave alerts produced by Advanced LIGO (Abbott et al. 2016; Broderick, Rowlinson et al. in prep) and the X-ray binary V404 Cyg (Broderick et al. in prep), on timescales of days to months after the alerts. This response timescale is sufficient for incoherent sources such as synchrotron or thermal emitters; including typical afterglows of many transient sources. However, a manual response of tens of minutes is insufficient for coherent sources such as Fast Radio Bursts and predicted early time emission from compact binary mergers (for further details see Chu et al. 2016; Kaplan, Rowlinson et al. 2016), where response timescales of seconds are required. As a software-driven telescope with no moving parts, LOFAR has the fundamental capability of being able to repoint to a new position within tens of seconds, which is likely to prove revolutionary for transient astronomy at low radio frequencies. The RO and TKSP are collaborating to bring this response timescale down to minutes by fully automating and optimizing the LOFAR systems. This effort requires the

[13] Marica Branchesi for the LVC 2016 *J. Phys.:Conf. Ser.* 718 022004
[14] Open call for partnership for the EM identification and follow-up of GW candidate events, https://dcc.ligo.org/LIGO-M1300550/public

development of a broker system that is capable of automatically filtering standard transient alerts.

Transient Buffer Boards

In addition to the ordinary scheduled observations, it is also possible to follow up transients using the Transient Buffer Boards (TBBs). The TBBs contain a ring buffer that stores the last seconds of data for the selected antennas. The default setup is 5 seconds of raw data for all antennas, but there is also the possibility to select part of the bandwidth and/or less antennas to increase the buffer size. The use of this buffer has two main benefits. There is data available from before the trigger is received. This allows for observations at the same time as the trigger. Also, the telescope can be pointed in post-processing and therefore does not need to be steered in the direction of the trigger. There is one restriction here for the high band antennas (110-250 MHz). These have an analog beamformer for the tile beam, and the transient should be within this range (~30° FWHM).

The TBBs so far have been used for two science cases where triggers were used. The first science case was the detection of radio emission from cosmic rays (Schellart et al. 2013). To do this a particle detector (LORA, Thoudam et al. 2014) has been installed at the LOFAR superterp (the central core of LOFAR). This generates triggers when cosmic rays are detected. The triggers are sent to LOFAR, and 2 ms of data around the trigger time is read out for all antennas in the LOFAR core. From the radio data, the initial particle type can be reconstructed to study the evolution of the composition of cosmic rays over an energy range of 10^{17}-$10^{17.5}$ eV (Buitink et al. 2016). This is the first real-time LOFAR trigger between two different instruments.

The second science case was the localisation of Fast Radio Bursts (FRBs, Thornton et al. 2013). FRBs are short millisecond pulses with a high dispersion measure. The origin is unknown, but they are likely extra-galactic sources. An identification of the host galaxy would help in understanding FRBs and aid in using them as probes of the intergalactic medium. The first project uses a real-time trigger on streaming LOFAR data to search for FRBs (Ter Veen, PhD Thesis, RU Nijmegen). On an FRB detection, the TBBs will be read out and the data used to image the FRB and determine its position. The second project uses the Effelsberg Radio Telescope (100m). In this case, the trigger is generated at 1.4 GHz and there is a delay in the order of minutes before the signal arrives at LOFAR frequencies. This gives enough time to use a real-time trigger to trigger the FRBs. A pilot project has been conducted using known pulsars (Houben, Master Thesis, RU Nijmegen). This is the first real-time trigger between an external facility and the full LOFAR array. The communication between Effelsberg and LOFAR made use of VOEvents.

LOFAR as a source of triggers

In addition to being a responsive telescope, the TKSP conducts blind transient surveys and intends to produce transient alerts for the wider community as soon as is feasible after the observations. The TKSP has produced the LOFAR Transients Pipeline (TraP [15]), which is capable of processing large datasets of images. TraP conducts basic quality control on each input image before extracting all the detected sources and completing source association across both time and frequency to build up multi-wavelength light curves for each source. These data are stored in a large database alongside variability parameters for each source and identifications of likely transients. TraP has been proven to be successful across a wide range of radio telescopes (e.g. LOFAR [16], ASKAP [17], MWA [18], VLA[19]) and the next step in the TraP development plans is the automated production of transient alerts for communication with the wider community.

A subgroup of the TKSP have been developing AARTFAAC (the Amsterdam-ASTRON Radio Transients Facility And Analysis Centre, e.g. [20]), a whole visible sky transient detector using the central 12 stations of LOFAR, producing 1 second images in real time that will be input directly into TraP. AARTFAAC is in its final commissioning phase and is poised to detect rare bright transients (such as the transient identified by Stewart et al. 2016). After an initial period for quality control purposes, AARTFAAC will be sending out transient alerts in near real-time. AARTFAAC also intends to automatically trigger LOFAR observations.

2.3 Automatic triggers for the European VLBI Network

During the last few years there was a revolutionary change in the study of optical transients. One of the most successful optical transient surveys is the Palomar Transient Factory (PTF). While its survey speed is lower than in other surveys (e.g. CRTS), its huge success lies in the carefully planned spectroscopic (and sometimes multi-band) follow-up programmes. This could be an important lesson for the planned future radio surveys as well. The PTF opened up the parameter space of optical transients. These new transient types include the fast and faint sub-class of classical novae, and the peculiar Ca-rich gap transients that are fast evolving sub-luminous SNe. The planned upgrade of PTF, the Zwicky Transient Facility will probe the sub-day parameter space, which is currently rather unexplored; a well-established class of transients in this region is GRB afterglows (optical). Radio data on optical transients have been very sparse in the past, especially for the new types that require quick reaction time. The PTF group has carried out a near-real-time radio transient survey with the VLA in SDSS Stripe 82, coupled with contemporaneous optical monitoring and rapid follow-up at X-ray through radio frequencies. This systematic search for transient and variable radio

[15] Swinbank et al. 2015
[16] Stewart A. J. et al., 2016, MNRAS, 456, 2321
[17] Hobbs G., et al., 2016, MNRAS, 456, 3948
[18] Rowlinson A. et al., 2016, MNRAS, 458, 3506
[19] Clarke T., et al., 2016, arXiv:1603.03080
[20] Prasad et al. 2016

sources serves as a forerunner of next-generation surveys planned for WSRT/Apertif, ASKAP, LOFAR and MeerKAT.

While only a small fraction of optical transients is expected to produce radio emission, high-energy transients are great candidates for radio follow-up in general. Sources of high-energy triggers are Fermi, Swift, Integral Agile and MAXI. The currently known rate of transient events is about 0.5/day across the whole Sky (significant optical and X-ray triggers in VO-EventNet), which means that about 0.25/day would be visible for the EVN. It is not known however, what fraction of these transients produce detectable radio emission, and how much this radio emission would be delayed. Typical required radio trigger times and event durations for known and expected events vary from minutes to days. An important point is that in the next decade there will be a huge increase in transient detection rate, including triggers directly coming from the radio survey instruments already mentioned above.

The obvious advantage of VLBI is the source localization at unprecedented precision. While for the astronomical interpretation arcsecond localization would be sufficient in most cases, VLBI data have the great potential to distinguish between flaring AGN activity or another type of near-nuclear transient. The highest resolution VLBI measurements can probe non-thermal emission brightness temperatures up to about 10^{12} K, measure tiny displacements (in the 10–100 μas regime) due to source structural changes or proper motion, and they can be very helpful to measure compact source total flux densities in fields with strong arcsecond-scale diffuse emission in the host galaxy. This allows the study of a broad range of astrophysical phenomena in the Local Universe ($d \leq 200$ Mpc, or $z \sim 0.05$). For 1 mas resolution, the corresponding linear size is ~1 pc at $z = 0.05$, therefore sub-pc structures can be probed. Within this distance beamed relativistic ejecta can be resolved within a few months (here assuming $S/N \geq 10$ and size measurement or ejecta localization at the 0.1 mas level), and even mildly relativistic phenomena can be studied on similar timescales up to a few tens of Mpc distance.

3 Requirements for alert handling

In this section, we describe the requirements for the alert handling mechanisms for LOFAR, the EVN, and a Gravitational Wave detector. Standards must be defined for:

- Event generation, the act of determining whether an interesting transient observation has been made by an instrument.
- Event dissemination, the emission of such an event by the instrument over the Internet. This should to be based on *VOEvents*, possibly with extensions to the protocol.
- Event distribution, the method with which events are distributed to and amongst interested and authorised instruments

- Event reaction, the (possibly coordinated) response to event by one or more instruments. Reacting observatories must be able to perform *joint observations* based on a single event.
- Events will need to be *filtered* on viability.

Furthermore:

- The system should be *open* to 3rd parties, ranging from small optical telescopes to space programmes.
- The system should be able to handle *millions of events per night*

3.1 Requirement for LOFAR

Receiving an alert:

Essential (must):

- Format definition of transient alerts – the facilities considered in this work are interested in triggers from are sending signals in the form of VOEvents so any system needs to use these. Transient alerts will be reacted upon via a broker that can:
  - Match event to the relevant triggering proposal ID and extract proposal ranking,
  - Extract the position/s to observe,
  - Create an observing setup likely by linking to a parameter file containing a predetermined setup,
  - Determine any dwell constraints – i.e. what is the minimum acceptable delay before observations start, what is the minimum and desired observation duration,
  - Determine if there are any observation constraints – proximity to bright sources (Sun, CasA, CygA…), or source elevation.
- The instrument must provide feedback to the broker:
  - Whether and when observations start,
  - When they are complete, and
  - When data are available for processing or retrieval.
- A *"project"* is responsible for filtering events and triggering LOFAR, both using project-specific algorithms. This project can be an ASTRON-internal or an external project. The following requirements may influence inter-instrument operability:
  - An interface for the project to define which event streams to subscribe to.
  - An interface for project-supplied event filters, which determine which events to trigger on.
  - An interface for project-supplied translators from events to LOFAR specifications.

   - The project-supplied code must currently:
      - Be open for review by ASTRON.
      - Be Python 2.7 compatible.
      - Have protocols for deployment and maintenance.
- Triggering instrument unavailability must be taken into account:
   - The LOFAR telescope is not always available (for events and/or triggers), with and without prior notice, due to collaboration/project changes, quotas, scheduled and unscheduled maintenance, infrastructure availability, and priorities in the observational schedule.
   - Resources for LOFAR projects are awarded and allocated in fixed 6-month *cycles*, except for testing and development.
- Observation "shadowing", in which LOFAR follows another instrument's observational response, or another instrument follows LOFAR. An example of a system that will shadow observations is MeerLICHT and MeerKAT but there are other existing examples.
- The following parameters will need to be configurable (by the project) in order to shadow another instrument:
   - Which instrument(s) to follow,
   - Where these instruments can be reached (event-stream subscription),
   - What credentials (if any) are required for inter-instrument authentication on behalf of the project.

Important (should):

- The broker receiving and filtering transient events set-up can:
   - Determine the observing setup.
   - Communicate with telescope, e.g. to allow the broker some flexibility in observing setup if some stations / capabilities are not available.
- The triggering instrument is responsible for, and will:
   - Define which projects are allowed to trigger.
   - Define quotas for each project.
   - Define a decision tree for interrupting the telescope, to allow automatic decisions about the resources a requested observation is allowed to use, with respect to duration: start time, stations, post-processing priority, disk space, compute cycles, archiving priority, archive space, etc.

Optional (could):

- The transient alerts should contain the event classification (e.g. "gamma-ray burst"), if known, to allow consumers to filter specific galactic events.
- Communication via an electronic "handshake" with other facilities to conduct a coordinated response. E.g. LOFAR and the LOFAR international stations work together to cover different regions of the error box produced by gravitational wave events. Another use case is to have coordinated observations of a trigger where the

same region is covered by different facilities to obtain commensal multi-wavelength observations (e.g. LOFAR and APERTIF), but that each facility will only observe if the other facility is available.

- The transient alerts should allow globally unique identification, to allow instruments to match inbound events and coordinated response requests.
- The following parameters need to be defined with respect to a coordinated response when an instrument is triggered:
    - Which instruments to (possibly) do joint observations with.
    - Where these instruments can be reached for negotiation (hostname, port), and which negotiation protocol (version) they use.
    - What credentials (if any) are required for inter-instrument authentication on behalf of the project,
    - Which subset(s) of instruments are critical for the response.
- Support for advertising response windows for different antenna frequency ranges.
- Support for advertising response windows for different sub-instruments.
- Support for advertising minimal response times.

As an alert producer

Important (should):

- Create a public and open interface to broadcast events, possibly filtered by significance or event classification.
- Have a list of clients whom to inform first, to minimise latency.

Optional (could):

- Support for advertising event detection latency window(s), to allow triggering instruments to assess the possibility of a fast-enough response.
- Support for advertising event classification latency window(s).
- Support for advertising instrument availability to generate events, to allow triggering instruments to predict whether triggering can occur.

3.2 Requirements for the EVN

The EVN Programme Comittee (PC) has adopted a policy for automatic triggering based on events. A PI can submit a proposal that sets forward the criteria for observing based on external events. After the proposal has been approved it will be active for a limited time. This means that the proposal can replace other e-VLBI observations if a trigger event that meets the criteria specified in the proposal and if the running observation is ranked lower than the automatic trigger proposal.

Events will have to carry enough information to decide:

1. whether observing the source associated with the event is feasible
2. whether the event meets the criteria specified in an active EVN proposal
3. to create an observing schedule

A minimal set of parameters is:

- position
- position accuracy
- brightness
- wavelength
- classification
- time and duration of observation

Since the field of view of a VLBI array is relatively small, we need a position that is accurate enough to give us a reasonable chance that the source is within the field of view of the array. If the position accuracy is too low, the decision-making process should simply drop the event.

Since VLBI can only observe sources with a relatively high surface brightness, the brightness of an object plays an important role in the decision whether it is feasible to observe a source associated with an event.

Time of the observation, classification of the event, wavelength and the observed brightness will allow us to estimate the brightness of the source at the time and wavelength the array will be observing.

3.3 Requirements from GW detectors

GW alerts are produced by a data analysis framework able to perform transient searches and event characterization with low latency, typically minutes to hours (see Section 5.1 for more details about the general software architecture and workflow used to produce GW alerts).

GW alerts follow the VOEvent standard and are transmitted over a dedicated GCN-type communication network to observing groups that have signed a Memorandum of Understanding with LIGO and Virgo. The adaptation of GCN system to the GW domain is called GCN/LVC.

A GW alert includes basic information about the event including its time, its statistical significance measured by the False Alarm Rate (in Hz), and the type of search it originates from. The GW alert also includes a probability skymap giving the posterior probability distribution for the location of the source once it is available.

Sky maps are encoded according to the Healpix format and they are stored in FITS files. The accuracy of the source location reconstruction depends on the signal strength, its frequency contents and the number of detectors in the network. For typical binary neutron-star mergers observed by the two LIGO detectors, the 90 % confidence region corresponds to a solid angle of about 500 square degrees [21]. Adding the Virgo detector to the network results in a localization improvement by a factor of 2. Such a large sky area is generally difficult to cover even for wide-field instruments, and requires making successive observations that tiles the error region. A directional distance estimate is now coupled with the probability skymap [22]. This additional information allows to set a priority on which tiles to observe first.

It is possible to filter the alerts according to their properties, for instance select alerts of a given type (e.g., compact binary mergers) or significance (e.g., highly significant events only).

Three types of GCN/LVC Notices are communicated: ***preliminary, initial*** and ***update***.

A ***preliminary*** notice is issued after basic sanity checks and approval by operators on sites and on-call EM follow-up advocates. Here are the standard contents of the VOEvent associated to a preliminary notice.

```
TITLE:              GCN/LVC NOTICE
NOTICE_DATE:        XXXXXX
NOTICE_TYPE:        TEST LVC Preliminary
TRIGGER_NUM:        XXXXXX
TRIGGER_DATE:       XXXXXX
TRIGGER_TIME:       XXXXXX
GROUP_TYPE:         X                   [Analysis group: CBC or Burst]
SEARCH_TYPE:        X                   [Type of search e.g., Allsky]
PIPELINE_TYPE:      X                   [Name of the pipeline]
FAR:                XXXXXX [Hz]
TRIGGER_ID:         XXX
MISC:               XXX
```

An ***initial*** notice is issued after the probability skymap is available. The VOEvent includes links to skymaps files.

```
TITLE:              GCN/LVC NOTICE
NOTICE_DATE:        XXXXXX
NOTICE_TYPE:        TEST LVC Initial Skymap
TRIGGER_NUM:        XXXXXX
TRIGGER_DATE:       XXXXXX
TRIGGER_TIME:       XXXXXX
GROUP_TYPE:         X
SEARCH_TYPE:        X
PIPELINE_TYPE:      X
FAR:                XXXXXX [Hz]
TRIGGER_ID:         XXX
MISC:               XXXXXX
```

[21] B. P. Abbott, *Living Reviews in Relativity* 19, 1 (2016)
[22] L. Singer et al, *Astrophys. J.*, 829 (2016) no.1, L15

```
SKYMAP_URL:        https://gracedb.ligo.org/XXX
SKYMAP_BASIC_URL:  https://gracedb.ligo.org/XXX
EVENT_URL:         https://gracedb.ligo.org/XXX
```

Update notices are issued when more information is available, e.g., when the skymap from full event parameter estimation is available. The notice layout is identical to that of initial notices.

Member groups of the GW-EM follow-up program are encouraged to give information about which part of the error region they have covered and to announce the counterpart candidates they detected to the other members of the program, see e.g., the collection of GCN notices related to the follow-up of GW150914 [23]. This allows the deeper, multi-wavelength follow-up of the most interesting candidates by sensitive and generally narrow-field instruments that are not able to cover the entire error region.

4 Analysis of the requirements

There is still a significant gap between the information that instruments provide in the events they generate and the information instruments need to be able to observe. In many cases, further processing of events will be necessary to decide if an event is interesting and whether it makes sense to try to observe the event. Such processing will depend a lot on the science case and will be hard to do the necessary analysis in generic code. It is almost inevitable that science teams have to write specialized computer code to do the filtering in a way that meets their science goals.

The VOEvent standard seems to be mature and there seems to be consensus among various astronomy and astro-particle detection instruments that it is the way to go for future systems. VTP, the VOEvent Transport Protocol as recently been ratified as an official IVOA standard. Software to parse and transport VOEvent messages is available, most of it as Python modules.

One area of concern is that the current VTP standard will not be suitable for distributing events at very high rates (millions per night). The current VTP standard requires subscribers to a VOEvent stream to acknowledge reception of a packet. This means a VOEvent broker needs to keep track of which events have been acknowledged on a per-subscriber basis. The LSST is considering to extend VTP with a mode where acknowledgement is not necessary. This protocol would be used to send the complete stream to science teams closely attached with the LSST, which would then filter those events and rebroadcast those events passing the filter using the standard VTP mode.

[23] https://gcn.gsfc.nasa.gov/other/GW150914.gcn3

Another feature under consideration for a future version of VTP is the addition of a replay facility. This replay facility would allow subscribers to request replay of events from the recent past. This would make it possible for clients to go down for short periods without running the risk of missing an important event. Implementation of such a replay facility would only make sense for streams that are expected to produce events at a relatively low rate (i.e. a handful events per hour).

5 Design for an alert handling system

The analysis so far suggests that there is no one-size-fits-all approach. Some additional processing will be required, both for reacting to events, and for generating events for further dissemination. For proposal-driven instruments it is also necessary to have further filtering based on requirements set forward by program committees/time allocation committees. Therefore, a flexible and modular approach seems to be appropriate for the design. Such a design would be best implemented in a high-level programming language. Given that most VOEvent and VTP-related software libraries are written in Python, a well-known language in the astronomy community, it seems to be the most obvious choice for this high-level programming language.

As a first step towards implementing the design we propose to offer Python modules for filtering events based on whether an event can be observed by a participating facility or not. Such a module is currently not available and would be a great help to both observatories and science teams writing filtering pipelines.

We will then use these modules to implement a demonstrator (ASTERICS deliverable D5.13), using the LOFAR responsive telescope offering and the EVN’s automatic triggering scheme to implement an alert handling system for GW wave detections. Design principles for such an alert handling system are described in the following sub-sections on event generation, dissemination, reaction and coordinated response.

5.1 Event Generation

The gravitational wave data (time series, sampled at 16 kHz; typically, 30 GB/day) from all detectors are transferred through the Internet to computing centers. A panel of transient search algorithms are applied and generates triggers with a few second to minute latency. Figure 1 presents the overall workflow and the main building blocks of this infrastructure.

Two separate approaches are considered for searching gravitational-wave transients. On the one hand searches based on matched filtering techniques target specifically the gravitational waveform predicted for compact binary mergers. On the other hand, unmodelled, generic transient searches based on excess power methods that don't explicitly rely on a prerequisite waveform model.

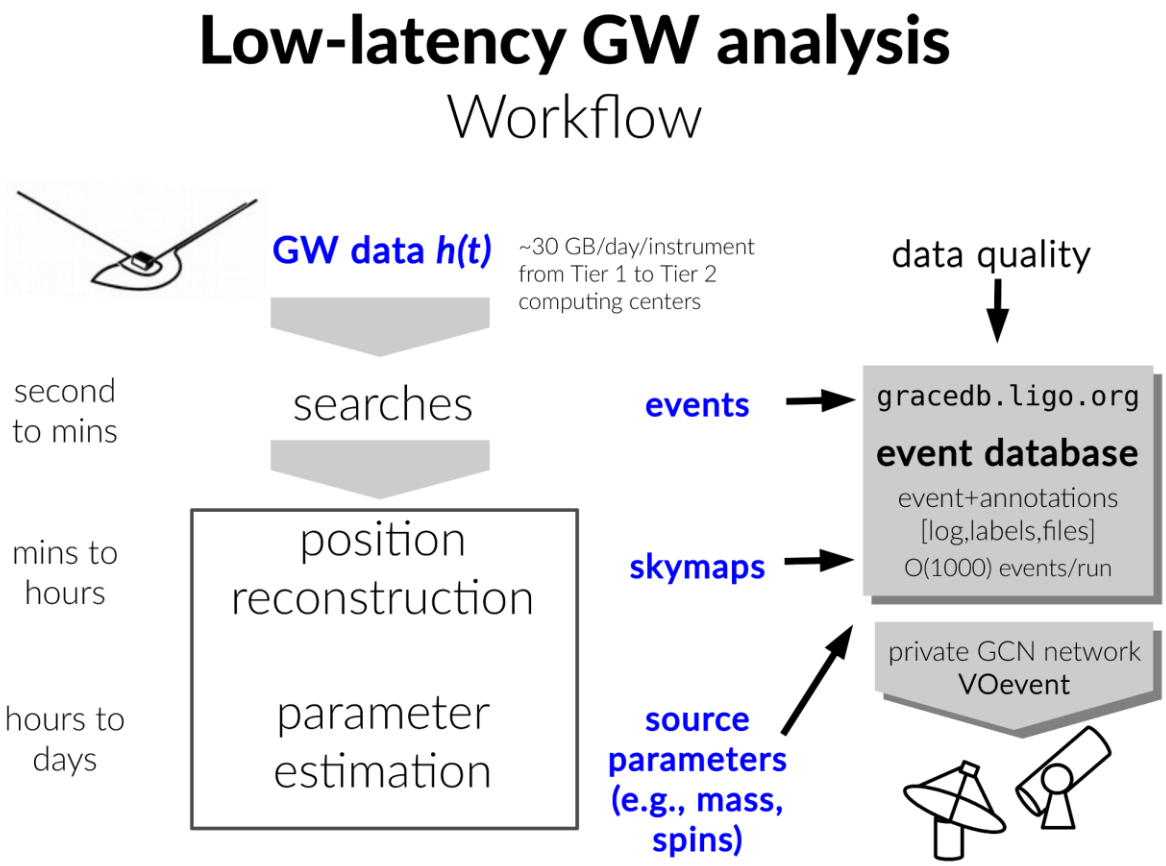

Figure 1: Schematic workflow for the generation of alerts from gravitational-wave data.

These searches provide basic information about the trigger, such as the time of arrival at each detector, the signal-to-noise ration and a preliminary estimate of the statistical significance. Depending on the type of search, the description of the triggers includes a preliminary characterization of the astrophysical source (e.g., rough estimate of the masses and spins of the compact binary mergers), or a skymap giving a first indication of the likely location of the source.

Triggers from all searches are immediately uploaded to the `gracedb.ligo.org` database. This automatically launches a series of processes. Data quality information is gathered and the probability that the trigger could be due to instrument noise is evaluated. Bayesian samplers are initiated in order to obtain the full parameter estimation for events connected to compact binary in spirals. The Bayesian estimation takes several days to deliver a final result.

Events with false-alarm rate better than 1 per month are considered as candidate for producing alert. After examination and validation by operators an initial alert is sent over a private GCN network using protocols: through socket or by email, and for the latter option the information can be given in one of the following forms: formatted text, XML VOEvent, link to GraceDB.

5.2 Event Dissemination

As follow-up events will often have similar publication restrictions as the original event, event dissemination will have to happen over both private and public channels. Such follow-up events will reference the original events through the <Citations> element in the VOEvent message. When the referenced event is private, it is left to the clients to make sure they are subscribed to the originating event stream. In this context, the replay facility that is being discussed within the IVOA could be useful. That way, if the client is not monitoring the event originating stream, it could ask for a replay if the referenced event was in the recent past. Perhaps replay of specific events (by specifying a <EventIVORN> element) could be implemented in that replay facility. This assures that the original event is only disseminated to those that should have access.

For public events, there are no restrictions and follow-up events can be published with full details. A reference to the original event should still be included. It may be beneficial to send out the original event as well, to handle the case where a client was not subscribed to the VOEvent stream where that event originated.

5.3 Event Distribution

The distribution mechanisms as set forward in the VTP standard seem to be adequate for the needs of the demonstrator as event rates for the current science cases are expected to be low. We expect to be using the Comet VOEvent broker for event distribution as there is some experience with that software among the members of the project and no know limitations that affect the goals of the project at this point.

One thing that is lacking in the current VTP standard is a mechanism for authorization. Digital signatures are being suggested as a possible way to verify the authenticity of VOEvents and could be implemented through a straightforward extension of the VOEvent standard. But authenticity of events currently isn't a major concern among the science cases.

There is a concern however about publicly distributing events. Currently LIGO only distributes events to groups who signed an MOU. This MOU also governs the distribution of follow-up events. While LIGO is expected to start distributing events publicly in the near future, this will probably only be done for high-confidence events. Since follow-up events for lower confidence events will be desirable, a mechanism to distribute events privately will need to be in place. In current VTP implementations this is usually implemented to restrict access to VOEvent brokers to specific clients by filtering on IP addresses. While not perfect this seems to provide a workable compromise. We don't rule out investigating the use of signed messages to implement authentication as suggested by the VTP standard authors in an IVOA note[24].

5.4 Event Reaction

Reaction to an event can be broadly divided into two cases. The first is the case of a singular response where a single facility responds to an event. In this case, the facility can decide autonomously if/when to observe the event. The more complicated case is the coordinated response, where it is desirable for multiple facilities to observe the same event more or less simultaneously in time.

5.4.1 Singular response

LOFAR specification flow

The LOFAR system will be extended with alert-handling functionality, called the *Responsive Telescope*. In this section, we describe the proposed design.

In LOFAR, specifications provided by the telescope operators include a set of jobs encoded in XML, created using off-line tools (e.g. “XMLGenerator”). These jobs include:

- *Observations:* the real-time reduction and recording of antenna data,
- *Pipelines:* the further non-real-time reduction of antenna data,
- *Ingests:* the transfer of reduced antenna data to remote tape clusters (“Long-term archive”, or LTA), from which the astronomer can retrieve data.

The full flow of specifications through LOFAR is shown below:

[24] http://www.ivoa.net/documents/Notes/VOEventTransport/

© Members of the ASTERICS collaboration

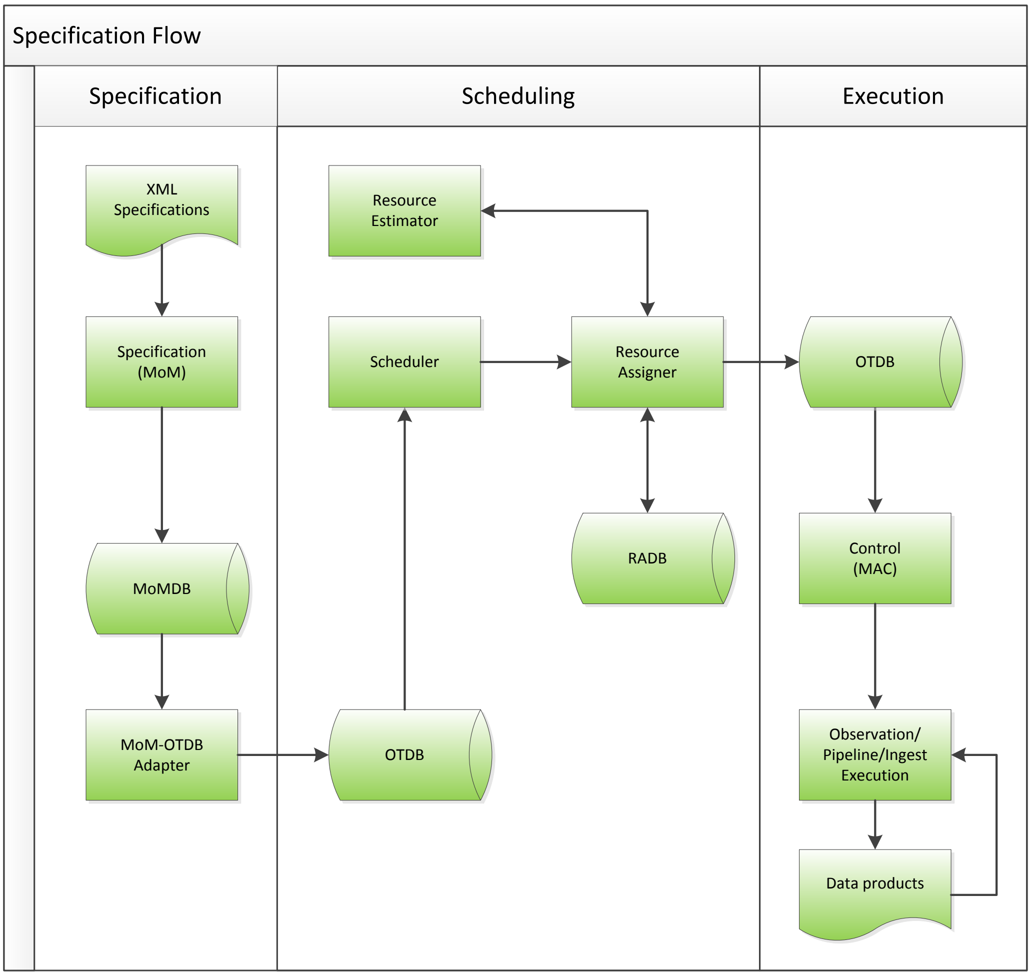

Figure 2: Specification flow in LOFAR

The XML specifications are entered in the specification subsystem ("MoM"), and transferred to the on-line database ("OTDB") on operator request, typically when it becomes part of the short-term schedule.

The observations are scheduled through human interaction with offline tools, as part of their specification (e.g. "XMLGenerator") or when adjusting the short-term schedule ("Scheduler").

The "Resource Assigner" finalises the specification by assigning disk resources and file locations, and marks the job as SCHEDULED, which implies it is ready to run. Observations are scheduled immediately, to run at their specified start time. Pipelines and ingests are scheduled when their input data was successfully created by their preceding observations or pipelines. Both pipelines and ingests use a batch-scheduling system, in which the job queue can be reordered by the operators.

All jobs are executed through the "Control" subsystem, producing data products on disk or on remote tapes. Once jobs have run, they send historical information such as quality measures and status updates back to the specification subsystem ("MoM") for annotation.

The specification subsystem and the LTA catalogue and data-retrieval services are reachable from the Internet. All other subsystems exclusively operate within the LOFAR domain.

Required functionality for a Responsive Telescope

The Responsive Telescope is the collection and implementation of services within LOFAR required to act upon external events and triggers. In LOFAR, telescope time is awarded to *projects*, representing accepted proposals for observation. Some of these projects will be allowed to interrupt the running LOFAR schedule, within the following restrictions:

- Each project has a *quota on the number of interruptions* in the observation cycle (6 months), on top of existing quotas for system resources allotted per project.
- The LOFAR Program Committee, in charge of judging and accepting the proposals for each cycle, determine the *priority ranking between projects*. Projects can only interrupt observations belonging to projects with a lower priority ranking. We deem it key to have a simple policy to create a clear understanding for both the projects responsible for the trigger and the interrupted observation on why the trigger was rejected or why an observation was interrupted.

Once accepted, a project can trigger LOFAR in two ways:

1. The *projects can supply triggers* through an Internet-facing interface. This approach provides dynamic control of when and how to trigger LOFAR, but requires active software on the astronomer's servers to produce the triggers. We opt to implement an encrypted and authenticated ReST web interface, only allowing trigger submissions from selected projects.
2. The *projects can supply event subscriptions, filters, and translators*, which allows LOFAR to receive VO events and generate triggers without further external input.

To trigger effectively and with low latency, LOFAR furthermore requires:

- The *full automation of the specification flow*, removing the need for human interaction and off-line tools
- An *optimisation for latency*, in the control flow, but also more specifically:
    - For pipelines and ingests, by allowing *automated queue manipulation* in their respective batch-scheduling systems based on project priority.
    - For the ingest, by *pinning the ingested data* on the staging disks, avoiding the need to transfer to and from tape if the project requests the data immediately.
- *Flexibility in the specifications to allow for ranges of resources to be specified*, such as a minimum/maximum duration, start time, and antennas of each type. This is needed

to increase the chance of success of the requested observation without the need to externally expose or interactively negotiate the LOFAR schedule.

Finally, LOFAR has further operational demands in supporting subsystems:

- *Monitoring of triggers* arriving and their status, with interfaces for the telescope operators as well as for each project.
- *Notifications of triggers* arriving and their status, notifying the telescope operators and the relevant project.
- *Reporting of statistics* with regards to events and triggers, both accepted and rejected, per requested time period.
- The availability of a dummy LOFAR interface to *allow projects to test their algorithms* and implementations.

To support these capabilities, we will add several components, shown in the following figure:

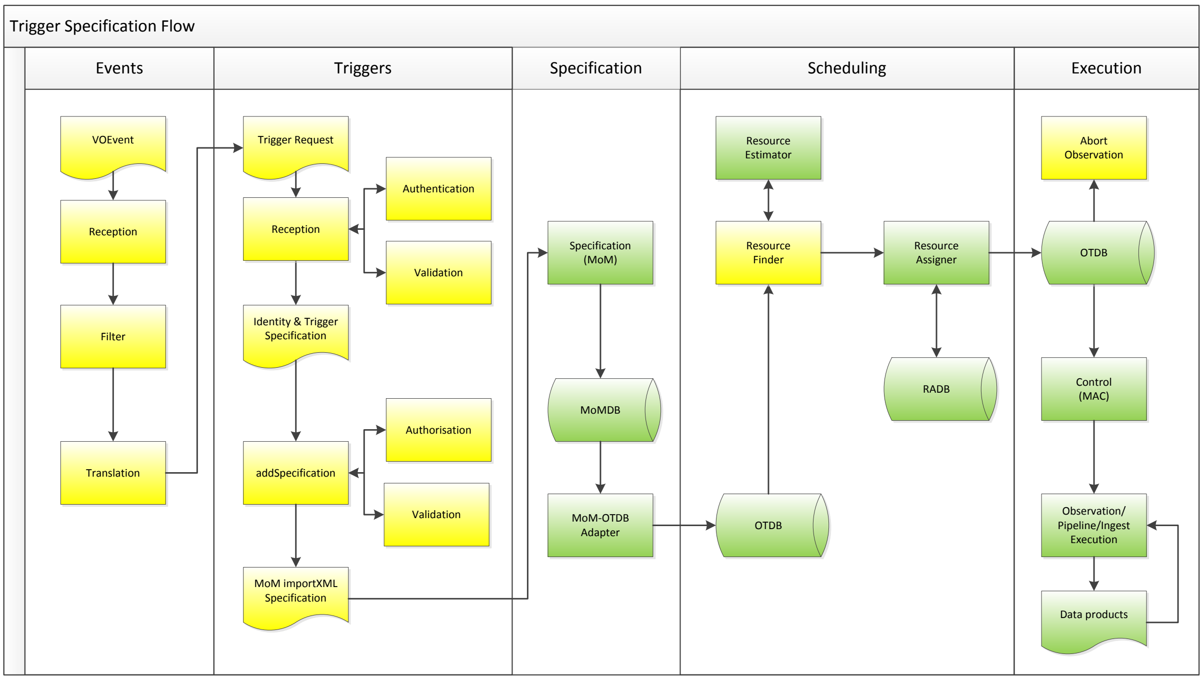

Figure 3: Responsive Telescope Component Diagram. The green boxes are existing LOFAR components, the yellow boxes are to be added.

The new components to be added are described below, along with their responsibilities:

- *Events::Reception,* a service subscribed to all external VOEvent brokers as specified by the projects that are allowed to trigger LOFAR.
- *Events::Filter*, a service running project-provided code to accept and annotate the event, or to reject it.
- *Events::Translation*, which translates annotated VOEvents into LOFAR Trigger specifications. Note that this involves converting the parameters specified by other

(non-astronomical) instruments into LOFAR terminology (a "LOFAR Trigger Specification").

- *Triggers::Reception*, a service listening to triggers on a web server in the LOFAR DMZ. Each login is authenticated against the LOFAR LDAP server, and each posted specification is validated against an XSD. Trigger-specific information is injected into the specification, resulting in a "LOFAR Specification".
- *Triggers::addSpecification*, checks whether the project is authorised to trigger LOFAR, and again validates its input against an XSD. When both succeed, the specification is translated into a "MoM specification", and send to MoM.
- *Scheduling::Resource Finder* replaces the off-line Scheduler tool, and searches for resources within the specified constraints with respect to system resources, such as the minimum/maximum duration, start time, antennas, etc.

Triggering Latency Analysis

The response time of LOFAR is determined by the time between a trigger or event is received, and the start of data recording. Although LOFAR has no moving parts in its station hardware, the specification and control chains are not optimised for latency. LOFAR currently has the following (maximum) latencies:

Subsystem	Latency
Event Handling	<not implemented>
Trigger Handling	<not implemented>
Specification	~40s
Scheduling	<human interaction>
Control	~60s
Execution	~51s
Total (except Scheduling)	*~151s*

For the Event Handling subsystem, we expect the delay to be dominated by the trigger filtering and translation as provided by the project. Additional delays to communicate between LOFAR services are in the order of milliseconds. The same holds for the Trigger Handling services. The additional delay required for automatic Scheduling is yet unknown, as it depends highly on database optimisations and the required search algorithms. All in all, we expect that a latency in the order of ~**151s** is already sufficient to test our specification flow and to cover several scientific use cases.

Theoretical Minimal Latency

The minimal delay possible in LOFAR is determined by the hardware. For this, note that there are two observing modes: the live recording and reduction of antenna data, and the Transient Buffer Boards (TBB). The TBB buffer the signals of all antennas for ~5s, and can be frozen on request. The TBB thus allows a theoretical response time of **-5s**, although its contents depend on the current antenna configuration. Since each antenna is recorded individually in the TBB, the data obtained still has a very wide Field of View. Full TBB support in LOFAR is part of the Responsive Telescope project. For regular observations, the minimal delay is determined by setting up the stations and correlator in the right mode:

Subsystem	Initialisation phase	Latency
Station	Clock switch	24s
	Bitmode switch	1s
	Splitter	1s
	RCU mode	1s
	Beam setup	4s
	Total (with clock switch)	*31s*
	Total (without clock switch)	*7s*
Correlator/Beamformer	MPI startup	4s
	Initialisation	6s
	Allocation	8s
	Total	*18s*

The above tables show that if the station is in the right clock (which is expected), a reaction time of **18s** should be possible. If a clock switch is required at the stations, the latency increases to **31s**. Note that getting to the lowest latencies requires an increasing amount of software effort, thus the actual minimal latency obtained will depend on demand.

EVN automated triggers

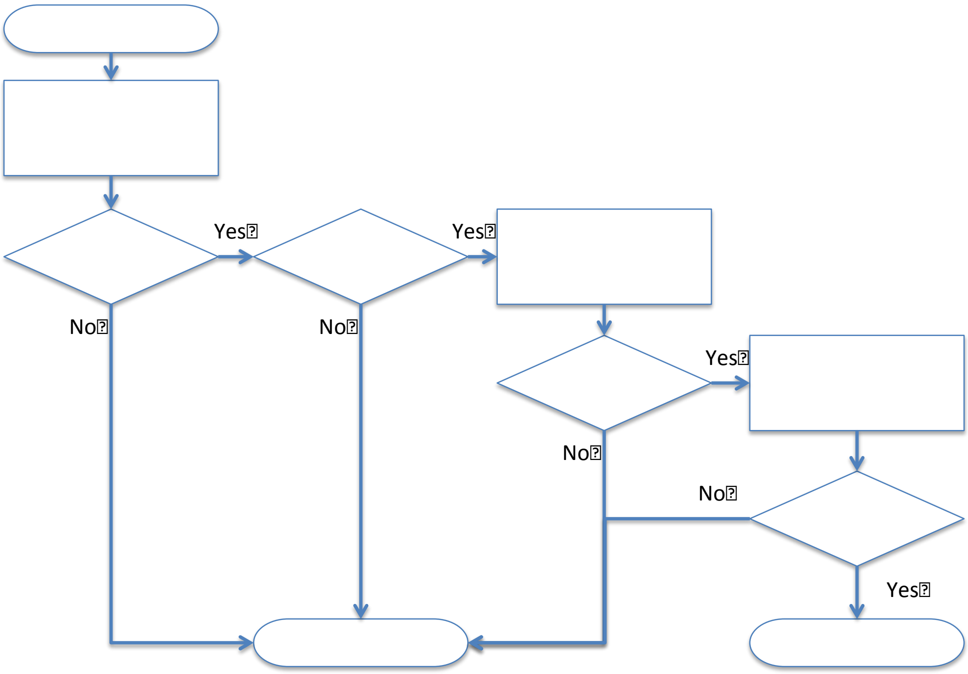

Figure 4: Flowchart for EVN automated triggers

The design of the system for automated triggers at the EVN also assumes science-driven filtering of events has been performed and it concentrates on the necessary steps to verify whether an observation triggered by a VOEvent should replace an observation currently running on the array. This can only happen on days when the EVN is observing in so-called e-VLBI mode, where data from the individual telescopes within the array are transferred over the internet and correlated in real-time at JIVE.

Triggers are only accepted for active proposals that have been approved by the EVN Programme Committee (EVN-PC). Active proposals will be stored in a database. For each e-VLBI session the (human) EVN scheduler will include all active proposals and their relative rankings in the so-called block schedule that list the planned observations for that session. This allows the system to decide if an automatic trigger should be able to override a planned observation.

There are restrictions on the observation mode of an automated trigger. The observation mode has to be identical to the planned observation mode. In practice this means that only the sources and the time intervals for which those sources are observed can be changed. The observed band, bandwidth, etc. cannot be changed.

Generating the new schedule poses a major challenge. Suitable calibration sources have to be found and scheduled for long enough that calibration of the final data set is possible. At the same time, the target source will have to be observed for long enough that enough sensitivity is reached to detect and image the source. While the scheduling is supposed to be done by the software under control of the observatory (JIVE in this case) it is largely science-driven. A possible strategy is to allow the Principle Investigator (PI) of the proposal that sends the trigger to include information about the observation strategy (calibrators, time intervals) in the trigger. This can easily be done by adding that information to the VOEvent being forwarded to the system from the science-based filter.

5.4.2 Coordinated response

So far, we have described an uncoordinated response from a single facility. A truly coordinated response to an event seems to be unfeasible. Differences in time-allocation policies, scheduling systems and response timescales make it very unlikely that a complete handshake protocol will reach an optimal agreement to observe. Instead we anticipate that instruments will send VOEvent utility messages that indicate that they (intend to) observe that include a reference to the original VOEvent that is being followed-up. The event filtering procedure for other instruments can then take this information into account and try to observe (quasi-)simultaneously with that instrument.

6 Demonstrator

The design described here will be concluded with a demonstration of a response to a GW event using LOFAR and the EVN. Since the rate of high-confidence GW events is currently of the order of a few events per year, we will almost certainly have to base this demonstration on either a replay of a historic event or a simulated event. The most realistic scenario would be to let this event trigger a follow-up at LOFAR, as LOFAR has a wide field of view. This follow-up observation would be able to narrow down the position of the event to a scale where VLBI observations make sense. Localization of the event could then trigger an EVN observation during the next window of opportunity.

7 Conclusions

The science cases described in section 2 show that many opportunities will arise once a working multi-messenger alert handling system is in place. However, the requirements in section 3 show that there is still a major gap between the events provided by an instrument and the information that other instruments require in order to follow-up. This gap needs to be filled by the science teams of these facilities, but they currently lack some of the (software) tools for bridging this gap. Providing the necessary software tools will be the next goal within task 2 of the CLEAOPATRA Work Package in ASTERICS. As a consequence of the aforementioned gap, the design presented here is still fairly high-level since it is still too early to make a formal specification of the minimal content of the VOEvent packets. These details will be decided when we are further into the implementation phase. At the end of that phase, we aim to provide a full specification and publish it in the form of an IVOA Note to assist other groups working on a VOEvent response system.