

Standards and Systems for Transient Response

John Swinbank*

Astronomical Institute “Anton Pannekoek”, University of Amsterdam

E-mail: swinbank@transientskp.org

Across a variety of different wavelength regimes, new facilities are being planned or are under development which will discover transients at an unprecedented, and potentially overwhelming, rate. An important challenge will be to address this deluge of data. How can we classify the detections, decide what is important, trigger appropriate follow-up observations—and do so quickly, and, where possible, with minimal human intervention? Here we discuss some of the approaches the LOFAR Transients Key Science Project has been investigating to these issues, with an emphasis on the importance of adopting common standards.

Science and Technology of Long Baseline Real-Time Interferometry:

The 8th International e-VLBI Workshop

June 22 - 26 2009

Madrid, Spain

*Speaker.

1. Introduction

LOFAR, the LOw Frequency ARray¹, is a next generation “software” radio telescope, currently under construction in the Netherlands and across Europe. When it comes online in late 2009, it will provide an unprecedented view of the low frequency radio sky, combining a large field of view with high sensitivity. The Transients Key Project [6] (TKP) will exploit this unique new capability to perform large-scale monitoring of the sky for transient phenomena—the so-called “Radio Sky Monitor” [7] (RSM). When operating, the RSM is expected to detect several new transient events per day. Each one of these events will require a response of some sort. At a minimum, they must be identified, classified, filed in the TKP’s light-curve database, and targeted for further monitoring. In many cases, however, the event will merit further investigation; perhaps the source will merit closer investigation with LOFAR itself, or perhaps a notification will be sent to the wider astronomical community requesting multi-wavelength follow-up of the most interesting or significant events.

Of course, LOFAR is not just a wide-field monitoring instrument. The TKP also anticipates responding to RSM detections as well as alerts from other facilities, perhaps reconfiguring the LOFAR system to best observe the event, or even saving the contents of LOFAR’s Transient Buffer Boards (TBBs), which will buffer all data received by each of the thousands of antennae which constitute the array—but only for a few seconds. Hence, for technical as well as scientific reasons, timely response to triggers is essential.

The huge volume of data to be searched for transients, the potentially large number of events, the multitudinous decisions which must be taken for each event, and the necessity that they should be taken rapidly, all mean that it is impractical for astronomers to manually process and respond to the incoming data. Therefore, the TKP team have been developing a software processing pipeline so that LOFAR itself can respond to new transients autonomously [12].

Of course, LOFAR is by no means the only facility which identifies and distributes notifications of transient events. Indeed, the number and sophistication of such facilities is increasing—as well, of course, as the number of events they produce. Even newer and more sophisticated radio-wavelength machines than LOFAR are already on the drawing board, while, at other wavelengths, facilities such as Pan-STARRS² and the LSST³ (Large Synoptic Survey Telescope) promise an ever-increasing deluge of transients, reaching thousands per day within the next decade.

Not only would it be wasteful for each of these next-generation “transient machines” to develop their own means of sending alerts of new events, it would also be overwhelmingly complex for systems receiving and processing such an inhomogeneous mass of input data. It is therefore vital that the astronomical community as a whole adopts common standards for the representation and transmission of transient notifications.

This document will give a brief overview of some of the internal systems developed by LOFAR’s TKP for detecting and reacting to transients. It will then describe some of the standards and systems which the TKP plans to adopt in sending and receiving alerts.

¹<http://www.astron.nl/radio-observatory/astronomers/lofar-astronomers>

²<http://pan-starrs.ifa.hawaii.edu/>

³<http://www.lsst.org/>

2. The LOFAR Transients Pipeline

The LOFAR Transients Pipeline is designed to monitor a large fraction of the sky for low frequency radio transients; the aim is both to detect new events and to build up a long-term lightcurve database of known transients.

An overview of the various stages in the LOFAR Transients Pipeline is shown in Figure 1. Image cubes—a set of simultaneous images at multiple frequency bands—will be generated by LOFAR’s Standard Imaging Pipeline at logarithmically-spaced intervals (1, 2, 5, ... 10000 seconds), and are fed directly to the Transients Pipeline. A parallelized source finding system is run on each plane of the image cube, and the results fed to a relational database system⁴. The database compares the new results with a catalogue of the area currently under investigation; it identifies new sources, and, if some known sources have not been detected, arranges for a maximum flux at their position to be calculated.

The lightcurve database produced in this way is one of the major output products of the pipeline. It is gradually migrated to a long-term storage system designed for data mining, and will be made available to the community at large. However, equally important is the response system. This is designed to automatically respond both to updates to the pipeline’s database *and* to external triggers, evaluating them, classifying them, and, if appropriate, arranging for a response to be made.

A number of possible responses might be appropriate, depending on the nature of the source and its behaviour; some possibilities were discussed in Section 1. Further, the understanding both

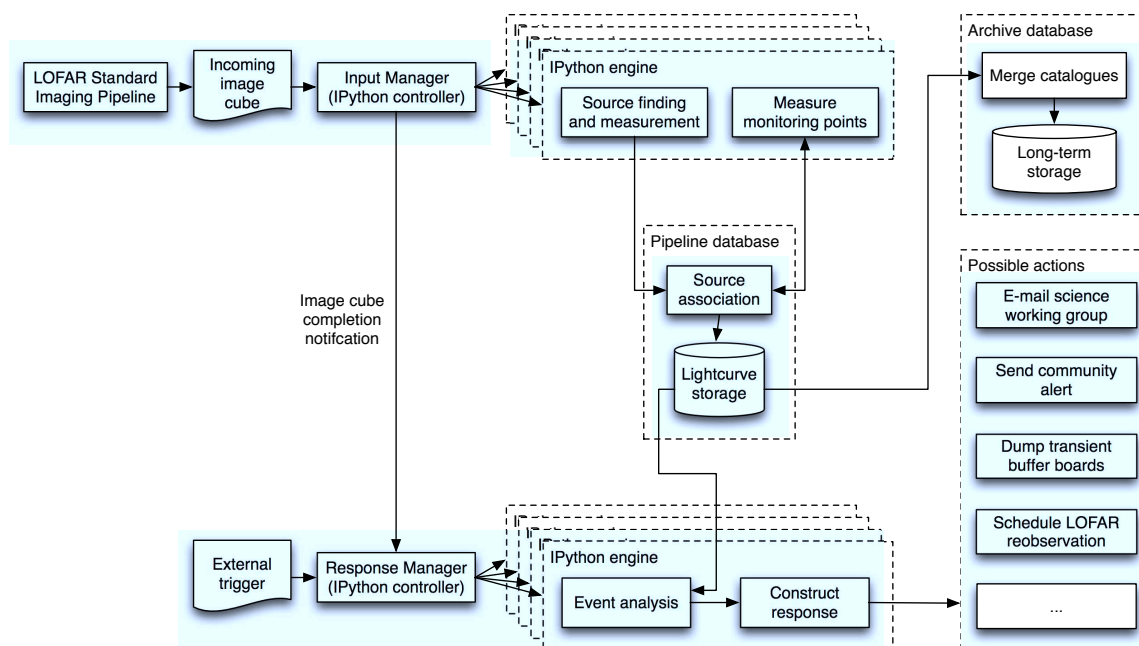


Figure 1: An overview of the LOFAR Transients Pipeline.

⁴A comparison of the performance of MySQL (<http://www.mysql.com/>) and MonetDB (<http://www.monetdb.cwi.nl/>) in the pipeline environment is currently underway.

of events and appropriate responses is expected to improve with experience. Therefore, rather than hard-coded behaviour, a dynamic, plugin-based system has been developed. Plugins will be written by astronomers in Python; they will provide support for recognising and responding to events, including a custom-developed machine-learning system [3].

3. External event notifications

Since LOFAR will be managed by a centralised control system, it is relatively easy for the Transients Pipeline to arrange for internal responses to be carried out. However, communicating with other facilities provides a more significant challenge. Although existing high-speed notification systems exist, they tend to be limited in their scope (e.g. the Gamma Ray Burst Coordinates Network [1]) and/or require significant human effort to process (e.g. The Astronomer's Telegram [9]); they are therefore unsuitable for the general-purpose, automatic system required by LOFAR. Therefore, a standardised, machine readable way of describing transient events is required. Such a system, *VOEvent*, is currently under development by the International Virtual Observatory Alliance's⁵ VOEvent Working Group^{6,7}.

3.1 VOEvent

VOEvent [11] is an XML [2] schema [5] for describing transient celestial events. *XML* provides a language suitable for the representation and sharing of structured data; the *schema* defines how that language should be used in the context of representing a transient event. The VOEvent standard is designed to make it possible to represent all the relevant information that a receiver might need to decide on an appropriate response to an event; although transmission of a VOEvent packet might imply that appropriate follow-up is requested, it is, of course, up to the receivers to decide what is important to them and configure their systems appropriately.

It is beyond the scope of this document to provide a detailed breakdown of the contents and capabilities of a VOEvent packet. Briefly, however, the means are available to specify *who* is making the report, *what* they observed, *how* they observed it, *where* and *when* it was observed—both whereabouts on the sky, and where the observatory was located—and *why* they regard it as interesting. The packet can also *cite* other VOEvents to retract, supersede or follow-up on them. All of this information is designed to be processed by machines, but the sender is encouraged to also provide human-readable analogues of the data.

Devising a machine-readable way of referring to a diverse range of events is challenging. The VOEvent system therefore makes use of *Unified Content Descriptors* or UCDs. These are defined by another IVOA recommendation [4], and provide a controlled vocabulary for describing astronomical data quantities; therefore, the quantity under discussion can be unambiguously understood by the receiving computer. Extensive use is made of other IVOA recommendations within VOEvent, notably for specifying positions and times according to accepted standards [8].

⁵<http://www.ivoa.net/>

⁶<http://www.voevent.org/>

⁷Note that the LOFAR project hopes to make extensive use of the VOEvent system, but has not, so far, been directly involved in its development.

3.2 VOEvent transport

A VOEvent XML document provides all the necessary information to perform informed follow-up observations to a transient event. However, it is also necessary to communicate that information to other facilities before they can act upon it. There is no single standard for VOEvent transport; indeed, the user is encouraged to adopt whatever system most conveniently fits their needs. For example, a VOEvent document could be delivered by e-mail, or by direct network connection, or by making it available on a web page. Further, projects can re-use existing specialist transport systems but in place with existing partners.

However, systems are being put in place to help. A simple network transport mechanism is under discussion by the VOEvent Working Group⁸. A number of sites are now exchanging VOEvent packets in real-time over a “global VOEvent backbone” based on this system. It is currently possible, with a relatively short piece of code, to connect to this backbone and observe real-time public VOEvent traffic.

Development is also taking place on other transport mechanisms. Notably, the VOEventNet Project⁹ provides XMPP [10] streams of VOEvents from various sources on its website.

3.3 VOEvent brokers

Section 3.2 described how it is possible to receive VOEvents directly from partner facilities or by connecting to a global backbone. Once an event has been received, however, it is still necessary to process it to decide if it requires action. A raw backbone connection makes this task complex: a large number of events of inhomogeneous types will be received, and each one requires parsing by the end user.

To make the task of digesting VOEvents more tractable, *broker* systems of varying degrees of complexity and sophistication can mediate between the end-user and the backbone. A broker may monitor the backbone, searching and sorting incoming events into more digestible streams based on user defined criteria; they might also provide customised notifications, e.g. by e-mail, of specific event types.

SkyAlert¹⁰ [13] is a sophisticated broker currently under development. In addition to the capabilities described above, it provides facilities for annotating each event with content provided by other databases and directly by users, so that a comprehensive dossier may be built up to help drive informed follow-up.

3.4 LOFAR inputs

VOEvents will be the primary method of sending trigger information to LOFAR. Where possible, facilities which have triggering arrangements with LOFAR will deliver information by VOEvent. In addition, the LOFAR system will monitor public VOEvent traffic, injecting it into the standard processing pipeline as described in Section 2; it will then be evaluated for response in parallel with LOFAR’s own detections.

⁸More detail at http://www.estar.org.uk/wiki/index.php/ESTAR_Event_Broker.

⁹<http://www.voeventnet.org/>

¹⁰<http://www.skyalert.org/>

4. Conclusions

The next generation of large-scale survey instruments will require sophisticated, inter-operable software systems to obtain the best scientific results. The planning and development of these systems is well underway at facilities like LOFAR, which are adopting the common standards laid down by the IVOA. Those standards continue to evolve: there has never been a better time for the e-VLBI community to become involved in the development of the infrastructure to support the future of transient astronomy.

References

- [1] S.D. Barthelmy et al., *BACODINE, the Real-Time BATSE Gamma-Ray Burst Coordinates Distribution Network*, *Ap&SS*, **231**, 235–238, 1995.
- [2] T. Bray et al., *Extensible Markup Language (XML) 1.0*, fifth edition, World Wide Web Consortium, 2008, <http://www.w3.org/TR/2008/REC-xml-20081126/>.
- [3] T. Coenen, *Automatic LOFAR Transient Classification*, Master's thesis, University of Amsterdam, 2008.
- [4] S. Derriere et al., *An IVOA Standard for Unified Content Descriptors*, version 1.1, International Virtual Observatory Alliance, 2005, <http://www.ivoa.net/Documents/REC/UCD/UCD-20050812.html>.
- [5] D.C. Fallside and P. Walmsley, *XML Schema Part 0: Primer*, second edition, World Wide Web Consortium, 2004, <http://www.w3.org/TR/2004/REC-xmlschema-0-20041028/>.
- [6] R.P. Fender and the LOFAR Transients Key Project, *The LOFAR Transients Key Project* in proceedings of *The VI Microquasar Workshop: Microquasars and Beyond*, 2006, PoS(MQW6)104.
- [7] R.P. Fender, *LOFAR Transients and the Radio Sky Monitor* in proceedings of *Bursts, Pulses and Flickering: wide-field monitoring of the dynamic radio sky*, 2007, PoS(Dynamic2007)030.
- [8] A.H. Rots, *Space-Time Coordinate Metadata for the Virtual Observatory*, version 1.33, International Virtual Observatory Alliance, 2007, <http://www.ivoa.net/Documents/REC/DM/STC-20071030.html>.
- [9] R.E. Rutledge, *The Astronomer's Telegram: A Web-based Short-Notice Publication System for the Professional Astronomical Community*, *PASP*, **110**, 754–756, 1998.
- [10] P. Saint-Andre and the Jabber Software Foundation, *Extensible Messaging and Presence Protocol (XMPP): Core*, RFC 3920, The Internet Society, 2004, <http://www.ietf.org/rfc/rfc3920.txt>.
- [11] R. Seaman et al., *Sky Event Reporting Metadata (VOEvent)*, version 1.11, International Virtual Observatory Alliance, 2006, <http://www.ivoa.net/Documents/PR/VOE/VOEvent-20060810.html>.
- [12] J. Swinbank, *A transient detection and monitoring pipeline for LOFAR*, in proceedings of *Bursts, Pulses and Flickering: wide-field monitoring of the dynamic radio sky*, 2007, PoS(Dynamic2007)044.
- [13] R.D. Williams et al., *Skylert: Real-time Astronomy for You and Your Robots*, to appear in proceedings of *ADASS 2008*, [astro-ph/0906.2186v1].