

HOT-WIRING THE TRANSIENT UNIVERSE 3

P. R. Wozniak, M. J. Graham,
A. A. Mahabal, & R. Seaman



Proceedings of the Third Hot-Wiring Transient Universe Workshop

November 2013

Santa Fe, New Mexico

Editors

P. R. Woźniak, M. J. Graham, A. A. Mahabal, and R. Seaman

Organized by Los Alamos National Laboratory

Sponsored by Los Alamos National Laboratory,
National Optical Astronomy Observatory,
Large Synoptic Survey Telescope,
and International Virtual Observatory Alliance

Foreword

Science and Engineering for Time Domain Astronomy

The third edition of the Hot-wiring the Transient Universe Workshop took place at the Eldorado Hotel in Santa Fe, NM between November 13 and 15, 2013. The meeting explored opportunities and challenges of massively parallel time domain surveys coupled with rapid coordinated multi-wavelength follow-up observations. The interdisciplinary agenda includes future and ongoing science investigations, information infrastructure for publishing observations in real time, as well as novel data science to classify events and systems to optimize follow-up campaigns. Time domain astronomy is at the fore of modern astrophysics and crosses fields from solar physics and solar system objects, through stellar variability, to explosive phenomena at galactic and cosmological distances. Recent rapid progress by instruments in space and on the ground has been toward a continuous record of the electromagnetic sky with ever increasing coverage, sensitivity, and temporal resolution. With the advent of gravitational wave and neutrino observatories we are witnessing the birth of multi-messenger astronomy.

Acknowledgements

Many individuals contributed their skills and efforts to make the HTU-III workshop successful.

The Organizing Committee developed an excellent scientific program that attracted very strong participation from the time-domain astronomical community despite serious restrictions on conference travel and the United States federal government shutdown that ended less than 4 weeks before the meeting.

The workshop could not have happened without Rob Seaman who has been promoting close interactions between science and technology oriented astronomers for as long as I remember. Rob knows all the secret ingredients necessary for a productive workshop series.

Barbara Roybal provided expert advise and excellent administrative support that made it possible for LANL to host HTU-III.

Many thanks to Rachel O'Donoghue and the Eldorado Hotel staff for putting together a great venue and highly professional site support, Brandy Putt for preparing welcome packages and helping with the registration process, and Jim Douglass for providing poster stands.

Pete Marenfeld created another awesome workshop poster, the best yet of an impressive series that make the Hotwiring meetings even more enjoyable. Iair Arcavi took the conference photo.

We thank the LDRD program at LANL for sponsoring the event. Funding for posters was provided by LSST.

Przemek Woźniak
Chairman, Organizing Committee

Workshop Organization

Przemek Woźniak	Organizing Committee Chair
Rob Seaman	Program Chair
Brandy Putt	Conference Secretary
Barbara Roybal	Conference Coordinator

Organizing Committee

Przemek Woźniak	Los Alamos National Laboratory
Tom Vestrand	Los Alamos National Laboratory
Josh Bloom	University of California, Berkeley
Matthew Graham	California Institute of Technology
Lynne Jones	University of Washington
Mansi Kasliwal	Carnegie Observatories
Ashish Mahabal	California Institute of Technology
Tom Matheson	National Optical Astronomy Observatory
Andrej Prsa	Villanova University
Rob Seaman	National Optical Astronomy Observatory
Rachel Street	Las Cumbres Observatory Global Telescope Network
John Swinbank	Low-Frequency Array for Radio Astronomy
Lucianne Walkowicz	Princeton University
Roy Williams	Laser Interferometer Gravitational Wave Observatory

CONTENTS

Foreword	iii
Acknowledgements	v
Workshop Organization	vii
Massively Parallel Time-Domain Astrophysics: Challenges & Opportunities	1
Autonomous Infrastructure for Observatory Operations	3
<i>R. Seaman</i>	
The Follow-up Crisis: Optimizing Science in an Opportunity Rich Environment	5
<i>T. Vestrand</i>	
Machine Learning for Time-Domain Discovery and Classification	7
<i>J. Richards</i>	
The Variable Sky	9
<i>S. Ridgway et al.</i>	
Hot-Wiring Flare Stars: Optical Flare Rates and Properties from Time-Domain Surveys	15
<i>A. Kowalski</i>	
Time-Domain Surveys: Transient Searches	17
Transient Alerts in LSST	19
<i>J. Kantor</i>	
The Zwicky Transient Facility	27
<i>E. Bellm</i>	
SkyMapper and Supernovae	35
<i>R. Scalzo</i>	
The Catalina Real-time Transient Survey (CRTS)	37
<i>A. Drake</i>	
Pan-STARRS Transients, Recent Results, Future Plans	39
<i>K. Chambers</i>	
Gaia – Revealing the Transient Sky	41
<i>N. Walton</i>	
Photometric Science Alerts from Gaia	43
<i>H. Campbell et al.</i>	

Time-Domain Surveys: Moving Objects and Exo-Planets	51
Small Body Populations According to NEOWISE	53
<i>A. Mainzer</i>	
The Catalina Sky Survey for Near-Earth Objects	55
<i>E. Christensen</i>	
Time-Series Photometric Surveys: Some Musings	57
<i>S. Howell</i>	
Passing NASA's Planet Quest Baton from Kepler to TESS	63
<i>J. Jenkins</i>	
The ATLAS All-Sky Survey	65
<i>L. Denneau</i>	
The Performance of MOPS in a Sensitive Search for Near-Earth Asteroids with the Dark Energy Camera	67
<i>D. Trilling & L. Allen</i>	
 Time-Domain Surveys: Beyond Optical Photometry	 69
The TeV Sky Observed by the High-Altitude Water Cherenkov Observatory (HAWC)	71
<i>B. Dingus</i>	
Hearing & Seeing the Violent Universe	73
<i>S. Nissanka</i>	
Follow-up of LIGO-Virgo Observations of Gravitational Waves	75
<i>R. Williams</i>	
The Needle in the Hundred-Square-Degree Haystack: from Fermi GRBs to LIGO Discoveries	77
<i>L. Singer</i>	
ARTS – the Apertif Radio Transient System	79
<i>J. van Leeuwen</i>	
Radio Adventures in the Time Domain	81
<i>D. Frail</i>	
The Karl G. Jansky VLA Sky Survey (VLASS): Defining a New View of the Dynamic Sky	83
<i>S. Myers</i>	
VLA Searches for Fast Radio Transients at 1 TB/hour	85
<i>C. J. Law et al.</i>	

Nuts and Bolts: Telescope Networking	93
Time to Revisit the Heterogeneous Telescope Network	95
<i>F. V. Hessman</i>	
GCN/TAN: Past, Present & Future. Serving the Transient Community's Need	103
<i>S. Barthelmy</i>	
VOEvent: Where We Are; Where We're Going	105
<i>J. Swinbank</i>	
Time Series Data Visualization in World Wide Telescope	113
<i>J. Fay</i>	
RTS2 and BB – Network Observations	115
<i>P. Kubanek</i>	
Multi-Telescope Observing: the LCOGT Network Scheduler	117
<i>E. Saunders & S. Lampoudi</i>	
Nuts and Bolts: Algorithms and Event Brokers	125
Novel Measures for Rare Transients	127
<i>A. Mahabal</i>	
The Modern Automated Astrophysics Stack	129
<i>J. Bloom</i>	
PESSTO – The Public ESO Spectroscopic Survey for Transient Objects	131
<i>M. Fraser et al.</i>	
Time Series Explorer	137
<i>J. Scargle</i>	
ANTARES: The Arizona-NOAO Temporal Analysis and Response to Events System	145
<i>T. Matheson</i>	
Bayesian Time-Series Selection of AGN Using Multi-Band Difference-Imaging in the Pan-STARRS1 Medium-Deep Survey	151
<i>S. Kumar</i>	
The South African Astro-informatics Alliance (SA3)	159
<i>S. Barway</i>	
Predicting Fundamental Stellar Parameters from Photometric Light Curves	161
<i>A. Miller</i>	
State-Based Models for Light Curve Classification	163
<i>A. Becker</i>	

Follow-up Science, Opportunities and Strategies	165
Early Time Optical Emission from Gamma-Ray Bursts	167
<i>D. Kopač, A. Gomboc, & J. Japelj</i>	
Burst of the Century? A Case Study of the Afterglow of Nearby Ultra-Bright GRB 130427A	175
<i>D. Perley</i>	
Optical Interferometry and Adaptive Optics of Bright Transients	177
<i>F. Millour et al.</i>	
Multi-Color Robotic Observations with RATIR	185
<i>N. Butler</i>	
The Robotic FLOYDS Spectrographs	187
<i>D. Sand</i>	
Dynamic Follow-up of Transient Events with the LCOGT Robotic Telescope Network	189
<i>R. Street et al.</i>	
Rapid Follow-up in iPTF and the Science it Enables	197
<i>I. Arcavi</i>	
Transient Alert Follow-up Planned for CCAT	199
<i>T. Jenness</i>	
Data Triage of Astronomical Transients: A Machine Learning Approach	205
<i>U. Rebbaapragna</i>	
Lessons Learned and into the Future	207
Toward an Intelligent Event Broker: Automated Transient Classification	209
<i>P. Woźniak</i>	
How to Really Describe the Variable Sky	211
<i>M. Graham</i>	
The Radio Transient Sky	213
<i>J. Lazio</i>	
Astrophysics in the Era of Massive Time-Domain Surveys	215
<i>G. Djorgovski</i>	

Poster Papers	217
Following up Fermi GBM Gamma-Ray Bursts	219
<i>V. Connaughton et al.</i>	
The TOROS Project	225
<i>M. C. Diaz et al.</i>	
Program	231
Participants	239
Author Index	243

Massively Parallel Time-Domain Astrophysics: Challenges & Opportunities

Autonomous Infrastructure for Observatory Operations

*Rob Seaman
National Optical Astronomy Observatory*

Abstract

This is an era of rapid change from ancient human-mediated modes of astronomical practice to a vision of ever larger time domain surveys, ever bigger "big data", to increasing numbers of robotic telescopes and astronomical automation on every mountaintop. Over the past decades, facets of a new autonomous astronomical toolkit have been prototyped and deployed in support of numerous space missions. Remote and queue observing modes have gained significant market share on the ground. Archives and data-mining are becoming ubiquitous; astroinformatic techniques and virtual observatory standards and protocols are areas of active development. Astronomers and engineers, planetary and solar scientists, and researchers from communities as diverse as particle physics and exobiology are collaborating on a vast range of "multi-messenger" science. What then is missing?

The Follow-up Crisis: Optimizing Science in a Opportunity Rich Environment

*Tom Vestrand
Los Alamos National Laboratory*

Abstract

Rapid follow-up tasking for robotic telescopes has been dominated by a one-dimensional uncoordinated response strategy developed for gamma-ray burst studies. However, this second-grade soccer approach is increasing showing its limitations even when there are only a few events per night. And it will certainly fail when faced with the denial-of-service attack generated by the nightly flood of new transients generated by massive variability surveys like LSST. We discuss approaches for optimizing the scientific return from autonomous robotic telescopes in the high event range limit and explore the potential of a coordinated telescope ecosystem employing heterogeneous telescopes.

Machine Learning for Time-Domain Discovery and Classification

*Joseph Richards
Lawrence Berkeley National Laboratory*

Abstract

To maximize the scientific returns from modern time-domain projects, sophisticated machine-learning tools must be used. Our group has been on the cutting edge of the methodological and algorithmic development for time-domain astronomical data analysis. I will describe several problems in which we have made great strides, including real-time discovery and classification of transient events, photometric supernova typing, and probabilistic classification of variable stars from long-baseline time series. I will describe our use of manifold learning for feature extraction in multi-band supernova light curves, active learning to overcome sample-selection biases, and semi-supervised learning to maximally leverage existing sets of labeled and unlabeled data. These algorithmic advances have already reaped benefits for discovery and classification in real-time surveys and hold a tremendous amount of promise moving forward.

The Variable Sky

*Stephen T. Ridgway, Thomas Matheson, Kenneth J. Mighell, and Knut A. Olsen
National Optical Astronomy Observatory*

*Steve B. Howell
NASA Ames Research Center*

1 Introduction

A top-down characterization of variability in stars and galaxies allows us to predict the rates of discovery and the total numbers of variable targets that will be detected in deep synoptic surveys. The goal is to reduce the uncertainties from more than an order of magnitude to less than or of order $2\times$. These numbers may be useful for estimating the scale of alert distribution and characterization tasks, and the scope of the demand for target-of-opportunity follow-up.

2 Synoptic Survey Alerts and the NOAO Variable Sky Project

One of the unique products of a synoptic (here understood as repeating) survey is the production of alerts on detection of variable targets. For some such targets, rapid follow-up will be desired, which will be enabled by immediate publication. The Large Synoptic Survey Telescope (LSST) will publish alerts on all targets which show variation from a fiducial measurement, and will publish them within 60 seconds [6]. For a survey devised for its high data throughput, this suggests a challenging computational task (of particular concern to the project) and a heavy demand for follow-up facilities (of special interest to our observatory). A Google search shows estimated LSST alert rates in the range 5000 to 2,000,000 per night, and while these may involve varying definitions and assumptions, the large range and the lack of documentation is a concern.

The NOAO Variable Sky project, developed by the authors, addresses the alert rates expected for synoptic surveys. This paper collects summary results for LSST for the high latitude sky. The high latitude sky will be the hunting ground for faint, rare, extragalactic sources, and for this part of the sky, contamination by galactic variable stars will be low. Details will be published elsewhere.

3 Finding the Needle in the Haystack

Rapid alert publication supports two distinct event types—those which are known or immediately identified, and those which are unexpected or not yet identified. An automated classifier, commonly called a Broker [2], can be used to filter an alert stream in order to select useful events. However, if the alert rate is high, then even a small mis-classification rate could easily obscure or delay rare discoveries. In this context, it is important to distinguish between discovery alerts and repeat observation alerts. Eventually it will be useful to dig deeper and determine the fraction of discovery alerts that actually have sufficient archival information (as non-variable sources) to support classification and thus implicitly to estimate the alert rate for which there is no such data.

With this preface we describe below the types of variable and transient sources considered, based on estimates of which source types can be expected to dominate the alert stream.

4 Variable Star Discovery Rate per Night

A bottom-up enumeration of all variable star types did not work well, and so we took a top-down approach. The Kepler survey was used to characterize the variability of the most numerous stellar spectral classes, in terms of variability probability distribution functions as a function of temperature and (for the cooler stars) of luminosity. These probability functions were then applied to a simulated star catalog generated with the Besancon Galaxy synthesis model [10], giving a probability of variability vs amplitude for every star in the catalog. Detection limits from an LSST exposure calculator [4] were used to determine the detectability at the 5σ level for every star, and the probabilities summed to predict the number of detectable variable stars. This was carried out for samples in and near the Galactic plane, and for an arc through the south celestial pole.

In Figure 1, the discovery rate for variable stars is shown on a per-night basis, integrated over the high latitude sky (taken as $|l| > 20$ degrees). The discovery rate is based on the observed statistics of variability, and how many new variables will be detected based on the length of the observing sequence.

Guided by the LSST criterion for issuing alerts, new alerts will be issued for most variable stars most times that they are observed. Thus these detections all contribute to the alert rate, also shown in Figure 1. However, as known variables, they will have a history and existing characterization, greatly reducing the data distribution load on the alert system, the characterizing burden on the Broker, and presumably the follow-up effort required for classification.

5 AGN Discovery Rate per Night

A similar approach was employed in the study of AGN variability and discovery rates. However, here the available data on variability are far less numerous and homogeneous. A luminosity function for galaxies [8] was combined with the additional assumptions that 2.5% of all galaxies have AGNs [5], and that all AGNs are variable at the ~ 10 mmag level [3] or more. A study of AGN variability [5] provided a variability probability vs amplitude function. The most poorly determined quantity is the probability distribution of variability time scales. In the example shown here, a characteristic time constant of 6 months was assumed, which determines the slope in the discovery rate, though a more complete description might allow for a range of time scales. As with stars, it is assumed that once variation has become evident, that AGN will continue to generate alerts after the discovery, hence contribute to on-going alert rate. In Figure 1, the AGN alert rate is assumed equal to the initial AGN discovery rate, since there is insufficient statistical characterization for a more elaborate model.

6 Variable QSO Discoveries per Night

The surface density of QSOs [9] and the predicted cumulative distribution of magnitude differences [7] are used to predict the discovery rate. The discovery rate for QSO's is also shown in Figure 1 in the same way as for AGNs. The increase with time of the QSO variability amplitude (described by the structure function) tends to flatten the discovery rate. The time constant is not well characterized, so the discovery rate is shown here as flat though it will begin to decline after a few years.

7 LSST Alert Rate vs. Discovery Rates

Figure 1 shows an estimated high latitude alert rate of $\simeq 10^5$ /night. This would be an ominous number if it were necessary to analyze this many new targets nightly. However, as shown in Figure 1, we predict that the number of discovered sources of the most abundant and familiar types will be more nearly $\simeq 1000$ /night. Furthermore, most of the QSOs and AGNs, when detected as variables, will already have long time series giving good colors, and most likely a history of sub- 5σ variability, and these two pieces of evidence will support a high confidence of classification. These targets aside, SNe will be among the most numerous discoveries, and these will be readily identified with a combination of history and association to galaxies.

We do not include here cataclysmic variables, owing to lack of good luminosity functions and variability probability information for the more numerous, evolved CVs. As we will discuss elsewhere, if the number density of detectable CVs is as large

as permitted by stellar population estimates, CV discoveries could compete with QSOs. However, most CVs will also be identifiable from archival colors and low-level variability.

In order to determine the discovery rate of variable sources which cannot be immediately classified from archival data, we note that for such targets it will be essential to obtain immediate forced photometry on the target location, since the observational history below the catalog brightness cutoff will be necessary, and probably sufficient, for characterizing most of the faintest targets.

Of course a residual number of alerts will be generated by targets which have been fainter than the forced photometry stacked limit prior to discovery - these will number among the few sources which are initially both anonymous and (individually at least) unforeseeable. The expected counts for cool, flaring dwarfs, AGNs and QSOs brightening from below to above the detection limit can be estimated by extension of the analysis described above, but this exercise is deferred to a future study.

8 Number of objects in the LSST transient and variable catalog

The integral of the discovery rates gives a good measure of the size of the LSST transient and variable catalog, which is predicted to trend towards a little less than 10^7 high latitude sources over the 10 years survey.

9 What about the Galactic plane?

The study produced a map of variable number densities on the sky. The expected number of variable stars in the LSST-observed plane is $\simeq 2.5$ orders of magnitude larger than in the high latitude sky. However, it is not correct to multiply the star discovery rate in Figure 1 by this factor, as LSST may spend less time in the plane than in the high latitudes. Furthermore, the analysis above assumes photon- and calibration-limited detectability of variability. In the galactic plane, crowding may set the limit significantly above this. Finally, the trigger level for variability alerts is a selectable parameter, and it is not likely to be set at a level that overwhelms feasible processing.

10 Another Synoptic Survey – GAIA

Extending this analysis to the European Space Agency astrometric mission GAIA shows the perhaps unexpected result that in spite of an enormous difference in telescope apertures, GAIA and LSST will detect essentially the same numbers of variable

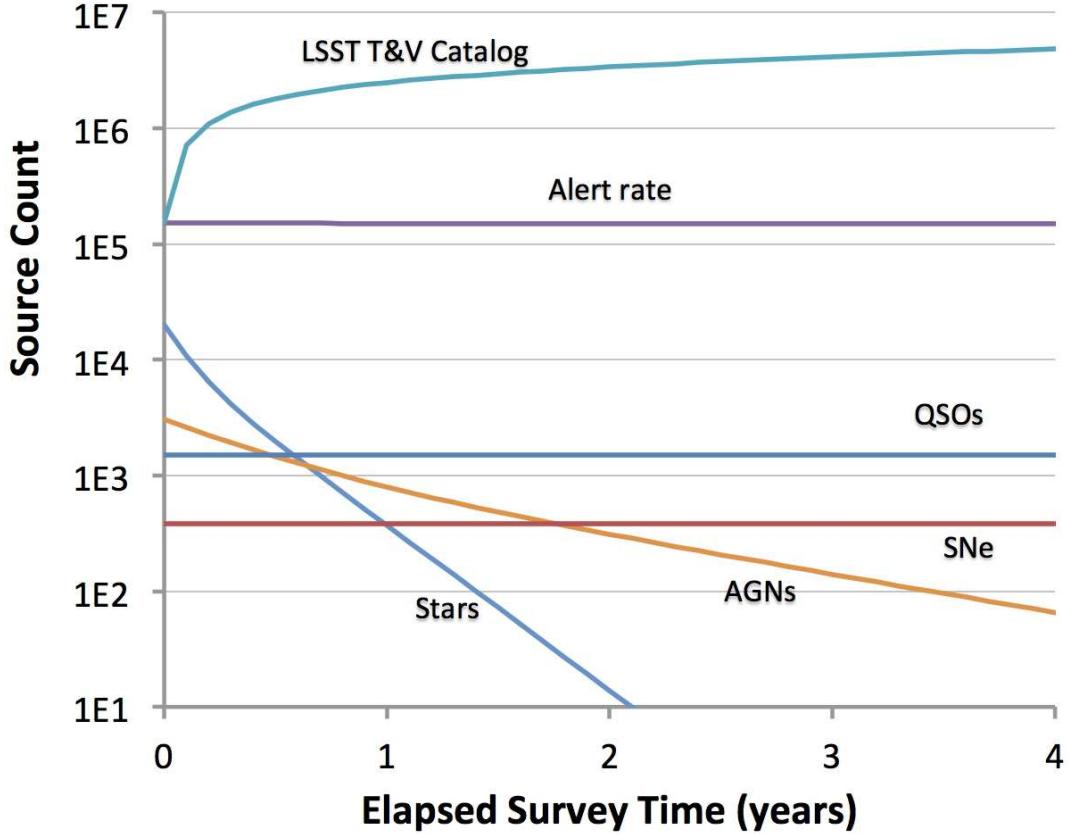


Figure 1: The LSST high latitude sky: discovery rates (per night) for stars, AGNs, SNe, and QSO's; alert rates (per night) for all of these source types combined; and transient and variable catalog sizes (total expected number of objects).

stars. The difference is that, while the GAIA survey will stop at $g=20$, it will reach lower levels of variability, thanks to the better photometry possible above the atmosphere. Approximately 90% of the variables detected by GAIA will have amplitudes of order 10 mmag or less.

We have presented this project at several conferences, and we are grateful for the feedback received. A longer version of this work, with details of assumptions and calculations, will be submitted for publication.

References

- [1] Bailey, S., Bernstein, J.P., Cinabro, D., et al. 2009, LSST Science Book, (Tucson, AZ: LSST), <http://www.lsst.org/lsst/scibook>, pp 383-384
- [2] Bourne, K.D., 2008, AstrN 329, 255
- [3] Gaskell, C.M., & Klimek, E.S. 2003, A&AT 22, 611
- [4] Gee, P. 2008, LSST Exposure Time Calculator (ver 4.3.2; Davis, CA: UC Davis), http://dls.physics.ucdavis.edu:8080/etc4_3work/servlets/LsstEtc.html
- [5] Klesman, A.J., & Sarajedini, V.L. 2012, MNRAS 425, 1215
- [6] <http://www.lsst.org/lsst/science/science-faq>
- [7] MacLeod, C.L., Ivezić, Z., Sesar, B. et al. 2012, ApJ 753, 106
- [8] Metcalfe, N., Shanks, T., Campos, A., McCracken, H.J., & Fong, R. 2001, MNRAS 323, 795
- [9] Palanque-Delabrouille, C., Yéche, C., Myers, A.D. et al. 2011, A&A 530, A122
- [10] Robin, A.C., Reyl, C., Derriè—re, S. and Picaud, S. A. 2003, A&A 409, 523

Hot-Wiring Flare Stars: Optical Flare Rates and Properties from Time-Domain Surveys

*Adam Kowalski
NASA Goddard Space Flight Center*

Abstract

Flares are thought to result from the reconnection of magnetic fields in the upper layers (coronae) of stellar atmospheres. The highly dynamic atmospheric response produces radiation across the electromagnetic spectrum, from the radio to X-rays, on a range of timescales, from seconds to days. Due to their high flare rates and energies combined with a large contrast against the background quiescent emission, the low-mass M dwarfs are the primary target for studying flare rates in the Galaxy. However, high-precision monitoring campaigns using Kepler and the Hubble Space Telescope have recently revealed important information on the flare rates of earlier-type, more massive stars. In this talk, I will focus on the properties of flares and flare stars in the optical and near-ultraviolet wavelength regimes as revealed from time-domain surveys, such as the repeat observations of the Sloan Digital Sky Surveys Stripe 82. I will discuss the importance of spectroscopic follow-up characterization of the quiescent and flare emission, and I will highlight new radiative-hydrodynamic modeling results that have enhanced our understanding of impulsive phase U-band flare emission.

Time-Domain Surveys: Transient Searches

Transient Alerts in LSST

*Jeffrey Kantor
Large Synoptic Survey Telescope*

1 Introduction

During LSST observing, transient events will be detected and alerts generated at the LSST Archive Center at NCSA in Champaign-Illinois. As a very high rate of alerts is expected, approaching 10 million per night, we plan for VOEvent-compliant Distributor/Brokers (<http://voevent.org>) to be the primary end-points of the full LSST alert streams. End users will then use these Distributor/Brokers to classify and filter events on the stream for those fitting their science goals. These Distributor/Brokers are envisioned to be operated as a community service by third parties who will have signed MOUs with LSST. The exact identification of Distributor/Brokers to receive alerts will be determined as LSST approaches full operations and may change over time, but it is in our interest to identify and coordinate with them as early as possible.

LSST will also operate a limited Distributor/Broker with a filtering capability at the Archive Center, to allow alerts to be sent directly to a limited number of entities that for some reason need to have a more direct connection to LSST. This might include, for example, observatories with significant follow-up capabilities whose observing may temporarily be more directly tied to LSST observing. It will let astronomers create simple filters that limit what alerts are ultimately forwarded to them. These user defined filters will be possible to specify using an SQL-like declarative language, or short snippets of (likely Python) code. We emphasize that this LSST-provided capability will be limited, and is not intended to satisfy the wide variety of use cases that a full-fledged public Event Distributor/Broker could. End users will not be able to subscribe to full, unfiltered, alert streams coming directly from LSST.

In this paper, we will discuss anticipated LSST data rates and capabilities for alert processing and distribution/brokering. We will clarify what the LSST Observatory will provide versus what we anticipate will be a community effort.

2 LSST Transient Science

The Large Synoptic Survey Telescope (LSST; <http://lsst.org>) is a planned, large-aperture, wide-field, ground-based telescope that will survey half the sky every few

nights in six optical bands from 320 to 1050 nm. It will explore a wide range of astrophysical questions, ranging from discovering killer asteroids, to examining the nature of dark energy.

The LSST will produce on average 15 terabytes of data per night, yielding an (uncompressed) data set of over 100 petabytes at the end of its 10-year mission. Dedicated HPC facilities will process the image data in near real time, with full-dataset re-processings on annual scale. A sophisticated data management system will enable database queries from individual users, as well as computationally intensive scientific investigations that utilize the entire data set.

LSST will support many areas of scientific research, as indicated in the LSST Science Book [1]. Of particular interest to the target audience of this paper are the sections on Transient Science and Solar System Science. LSST will detect and alert on an average of approximately 10 million transient events per night, where an event is defined as a significant, measured change in flux over a particular location.

LSST requirements are defined in the LSST Science Requirements Document (SRD) [2]. The following is an extract of the requirements related to transients and covers the contents, throughput, and filtering.

The fast release of data on likely optical transients will include measurements of position, flux, size and shape, using appropriate weighting functions, for all the objects detected above transSNR signal-to-noise ratio in difference images (design specification: 5). The data stream will also include prior variability information and data from the same night, if available. The prior variability information will at the very least include low-order light- curve moments and probability that the object is variable, and ideally the full light curves in all available bands. Specification: The system should be capable of reporting such data for at least transN candidate transients per field of view and visit (Table 1).

The users will have an option of a query-like pre-filtering of this data stream in order to select likely candidates for specific transient type. Users may also query the LSST science database at any time for additional information that may be useful, such as the properties of static objects that are positionally close to the candidate transients. Several pre-defined filters optimized for traditionally popular transients, such as supernovae and microlensed sources, will also be available, as well as the ability to add new pre-defined filters as the survey continues.

In normal survey mode, LSST will operate by capturing two back-to-back, 15-

Quantity	Design Spec	Minimum Spec	Stretch Goal
transN	10^4	10^5	10^6

Table 1: The minimum number of candidate transients per field of view that the system can report in real time.

second exposures for each pointing. The two exposures are referred to as snaps (aka exposures). They are combined to a visit, which is the basic input image product for transient alert processing, i.e. alerts are issued for each visit, not each snap. The primary purpose of the snaps is to enhance cosmic ray rejection. They are not to be confused with 30 to 90 minute revisits, scheduled to support Solar System science. The LSST Data Products Definition Document [3] is a readable description of LSST data products. Used to communicate with the science community, and to support the formal requirements flow-down. Describes the processing as well as the data products:

- Level 1 Data Products: Section 4
- Level 2 Data Products: Section 5
- Level 3 Data Products: Section 6
- Special Programs DPs: Section 7

Level 1 Data Products include the transient alerts.

LSST computing is sized for 10M alerts/night (average), 10k/visit (average), 40k/visit (peak). The DM System design includes, dedicated multi-gigabit/second networks for moving data from Chile to the US.

At the LSST Archive Center at the University of Illinois National Center for Supercomputing Applications (NCSA) dedicated computing infrastructure executes image differencing pipelines with improved algorithms for image calibration, detection, and alert generation. Solar System objects will be identified and linked together based on compatibility of their observed positions with motion around the Sun. An enhanced variant of the Pan-STARRS Moving Object Processing System (MOPS) algorithm has been used to develop an advanced prototype of the system. The fully developed algorithm will be used to identify and link observations of Solar System objects; measure their orbital elements; and measure their photometric properties. For each detected DIASource, LSST will emit an Event Alert within 60 seconds of the end of visit (defined as the end of image readout from the LSST Camera). LSST will measure and transmit with each alert:

- position
- flux, size, and shape
- light curves in all bands (up to a year; stretch: all)
- variability characterization (eg., low-order light-curve moments, probability the object is variable)
- cut-outs centered on the object (template, difference image)

Also, LSST will make available within 60 seconds fast moving objects (trailed) and known SSO's which suddenly develop activity (i.e. they show a non-point-source PSF). The goal is to transmit nearly everything LSST knows about any given event, enabling downstream classification and decision making without the need to call back into LSST databases (thus introducing extra latency).

We plan for VOEvent-compliant Distributor/Brokers (<http://voevent.org>) to be the primary end-points of the full LSST alert streams. End users will then use these Distributor/Brokers to classify and filter events on the stream for those fitting their science goals. These Distributor/Brokers are envisioned to be operated as a community service by third parties who will have signed MOUs with LSST.

The exact identification of Distributor/Brokers to receive alerts will be determined as LSST approaches full operations and may change over time, but it is in our interest to identify and coordinate with them as early as possible.

LSST will also operate a limited Distributor/Broker with a filtering capability at the Archive Center, to allow alerts to be sent directly to a limited number of entities that for some reason need to have a more direct connection to LSST. This might include, for example, observatories with significant follow-up capabilities whose observing may temporarily be more directly tied to LSST observing.

In conclusion, LSST will generate millions of transient alerts of interest to transient and solar system scientists every night, and will support public distribution of these alerts on 60 second time frames.

- Transient science
 - Nova, supernova, GRBs
 - Source characterization
 - Instantaneous discovery

- Nearby Solar System Objects

- NEOs, PHAs

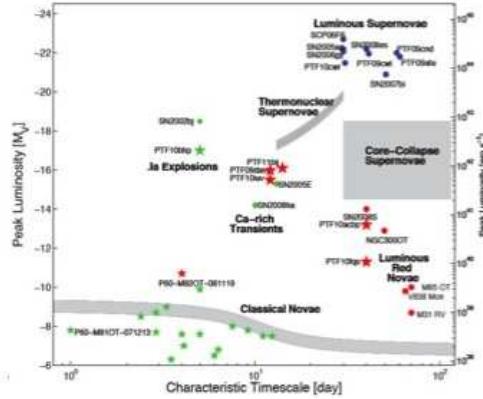
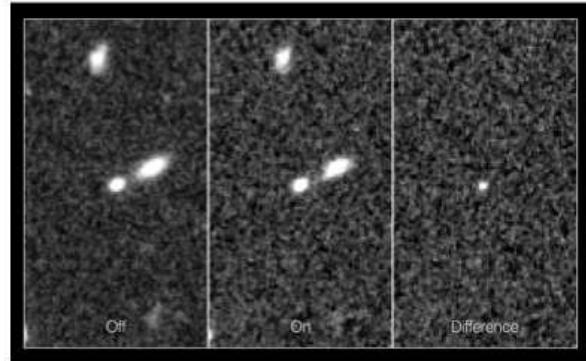


Figure 1: Transient Science with LSST

- Processing to enable rapid detection and follow-up of time-domain events
- Real-time image differencing as observing unfolds each night
- Measurement of position, brightness and shape for each detection
- *Alerts to detected changes transmitted within 60 seconds of observing, enabling rapid follow-up*



Transient Detection with Image Differencing (CANDELS; <http://www.spacetelescope.org/images/heic1306d/>)

Figure 2: Transient Detection with Difference Imaging (CANDELS:<http://www.spacetelescope.org/images/heic1306D/>)

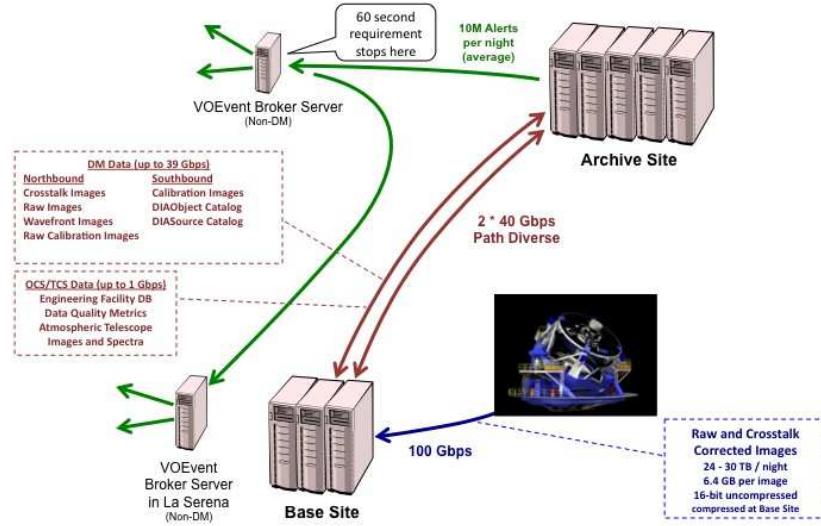


Figure 3: LSST Nightly International Data Flows

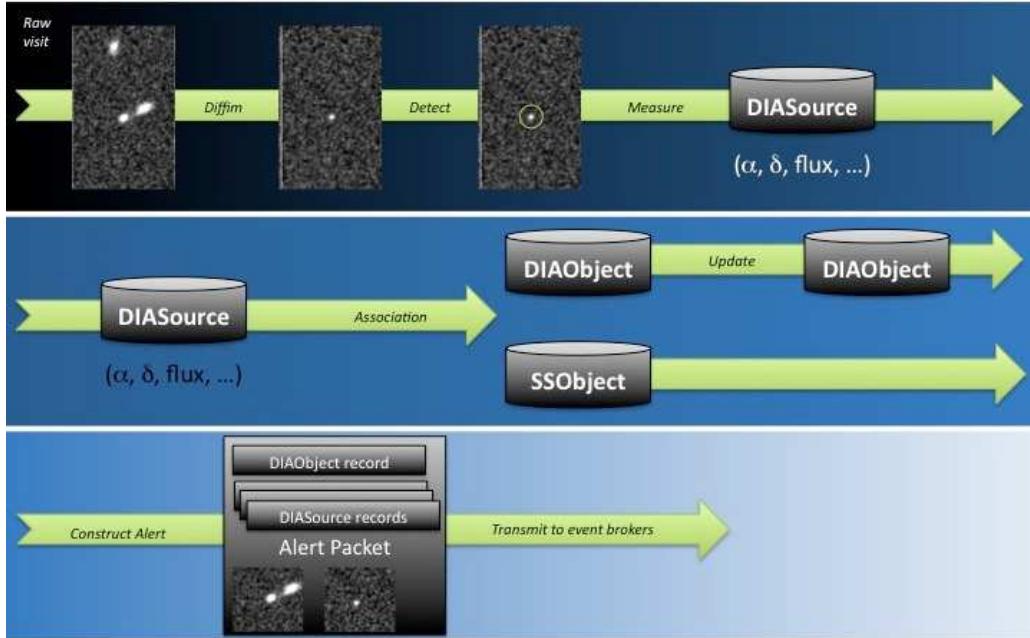


Figure 4: Level 1 Alert Production Outline

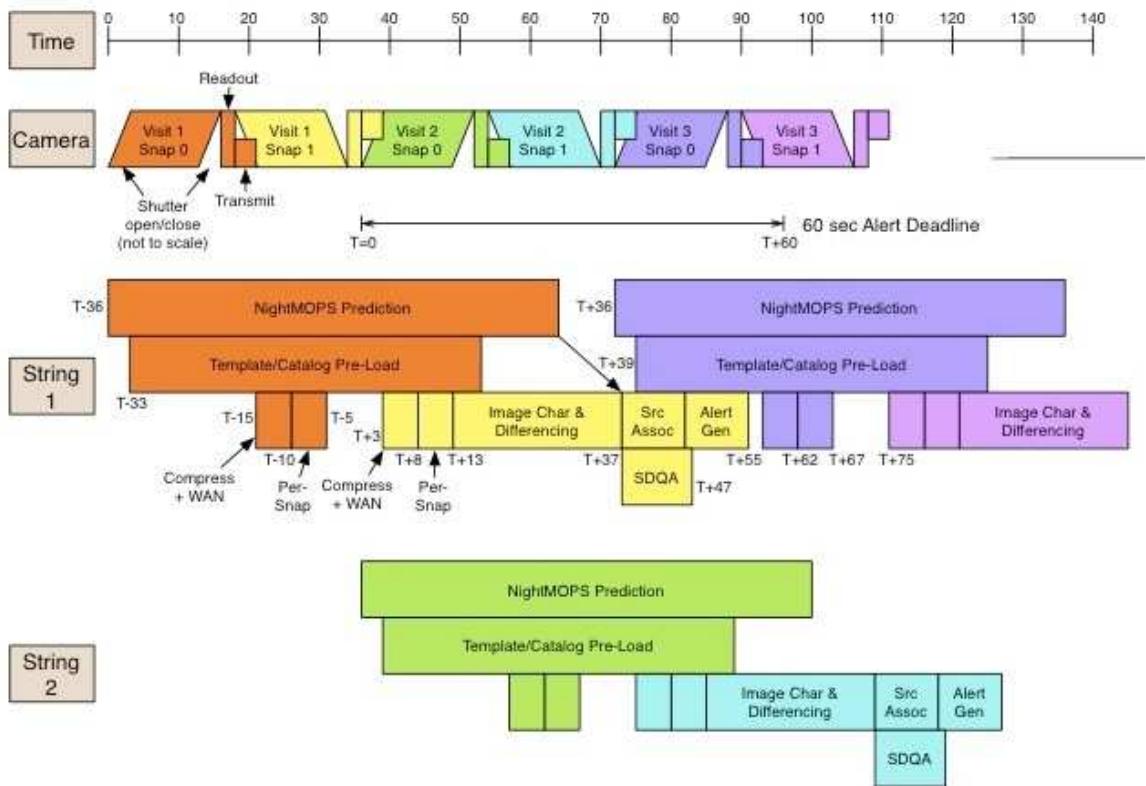


Figure 5: Level 1 Alert Production Timeline

References

- [1] LSST Science Collaborations. (2009). LSST Science Book (2nd ed.). Tucson, Arizona: LSST.
- [2] Ivezic, Z. and the LSST Science Collaboration. (2011). LSST Science Requirements Document (5th ed.). Tucson, Arizona: LSST.
- [3] Juric, M., Lupton, R.H., Axelrod, T., Bosch, J.F., Dubois-Felsmann, G., Ivezic, Z., Becker, A.C., Becla, J., Connolly, A.J., Freemon, M., Kantor, J., Lim, K-T, Shaw, D., Strauss, M., Tyson, J.A. (2013). LSST Data Products Definition Document. LSST Data Management. Tucson, Arizona: LSST.

The Zwicky Transient Facility

Eric Bellm
California Institute of Technology

1 Introduction

The Zwicky Transient Facility (ZTF; P.I. Shri Kulkarni) is a next-generation optical synoptic survey that builds on the experience and infrastructure of the Palomar Transient Factory (PTF) [12, 18]. Using a new 47 deg^2 survey camera, ZTF will survey more than an order of magnitude faster than PTF to discover rare transients and variables.

PTF (and its successor survey, the Intermediate Palomar Transient Factory, or iPTF) have conducted a transient-focused optical time-domain survey. PTF uses a 7.26 deg^2 camera on the Palomar 48-inch Oschin Schmidt telescope (P48) to survey the dynamic night sky in Mould-*R* and SDSS *g'* bands. Followup photometry and spectroscopy are provided by the 60- and 200-inch telescopes at Palomar and by other collaboration resources around the world.

PTF’s moderate-depth, followup-focused survey has yielded many notable successes. However, addressing leading-edge scientific questions (Section 4) requires a capability to survey at high cadence while maintaining wide areal coverage. Current facilities are inadequate for this purpose, but a straightforward upgrade of the PTF survey camera provides this capability while maintaining much of PTF’s demonstrably productive hardware and software infrastructure. ZTF will provide the best characterization of the bright to moderate-depth ($m \lesssim 21$) transient and variable sky and pave the way for LSST’s deeper survey.

2 Survey Design

The traditional measure of étendue (collecting area \times solid angle) is insufficient for characterizing the performance of time-domain surveys [22]. It relates most closely to the speed at which an instrument achieves a given coadded depth. Time domain surveys are often interested in the detection rate for specific classes of transients (e.g., Type Ia SNe or Tidal Disruption Events). These detection rates are a function of the intrinsic rate, brightness, and timescale of the transient; the cadence of the survey; and the spatial volume surveyed in each cadence period. For variability science,

the utility of time series data depends on the limiting magnitude, the photometric precision, the total number of observations, the cadence, and the bandpass(es) of the data.

This wide range of survey parameter space indicates the difficulty of optimizing a generic time-domain survey for a wide range of science goals. (It also suggests that specialized surveys will continue to be productive into the era of large time-domain facilities.) In consequence, single figures of merit are imperfect predictors of the performance of a time-domain survey, as much depends on the specifics of the chosen survey strategy in addition to the raw capabilities of the camera and telescope. However, optimization metrics are required to guide design studies and cost trades.

Building on the PTF heritage, we have chosen to optimize the ZTF camera design for the study of explosive transients. For any camera realization, we may trade survey cadence against the sky area covered per survey snapshot. We therefore seek to maximize the volumetric survey rate (\dot{V}), defined as the spatial volume within which a transient of specified absolute magnitude (e.g, $M = -19$) could be detected at 5σ , divided by the total time per exposure including readout and slew times. With appropriate choice of cadence, \dot{V} should be proportional to the transient detection rate. It implicitly incorporates the field of view of the camera, its limiting magnitude (which in turn includes the image quality, sky background, telescope and filter throughput, and read noise), and overheads [c.f. 22].

Notably, specifying the overhead between exposures implies an optimal exposure time to maximize \dot{V} . Exposures that are too short lead to an inefficient duty cycle, while exposures that are too long lead to smaller snapshot volumes, as the loss of areal covered is not offset by the increase in depth.

Guided by these considerations, our design for the ZTF survey camera (Section 3) maximizes the camera field of view, maintains PTF’s moderate image quality and depth, and minimizes the overhead between exposures and the number of filters.

3 The ZTF Camera

The 7.26 deg^2 field of view provided by the CFHT12k camera [17] currently used by PTF only covers a fraction of the available $\sim 47 \text{ deg}^2$ focal surface of the P48. By constructing a new camera that fills the focal surface with CCDs, we obtain a 6.5 times larger field of view. Modern readout electronics will reduce the overhead between exposures as well, providing a net improvement in survey speed of more than an order of magnitude relative to PTF. This speed boost will enable a transformative survey capable of simultaneously maintaining the high cadence and wide areal coverage needed to find rare, fast, and young transients.

The focal surface of the Schmidt telescope is curved, and during the Palomar Sky Surveys the photographic plates were physically bent on a mandrel to conform to

Specification	PTF	ZTF
Active Area	7.26 deg ²	47 deg ²
Exposure Time	60 sec	30 sec
Readout Time	36 sec	10 sec
Median Time Between Exposures	46 sec	15 sec
Median Image Quality (<i>R</i> band)	2.0" FWHM	2.0" FWHM
Median Single-visit Depth (<i>R</i> band)	20.7	20.4
Yearly Exposures per Field (3π)	19	290
Areal Survey Rate	247 deg ² /hr	3760 deg ² /hr
Volumetric Survey Rate ($M = -19$)	2.8×10^3 Mpc ³ /s	3.0×10^4 Mpc ³ /s

Table 1: Comparison of the PTF and ZTF cameras and survey performance metrics. Yearly exposures assume a hypothetical uniform 3π survey.

this focal surface. The PTF camera achieves acceptable image quality (median 2" FWHM in *R*) with a flat CCD focal plane, an optically powered dewar window, and flat filters. However, scaling a comparable design to the full ZTF focal plane produces unacceptable image quality.

We have developed an optical design that maintains PTF's image quality over the entire available field of view. An additional zero-power optic (to be fabricated from an existing blank) placed in front of the existing achromatic doublet Schmidt corrector provides a minor adjustment (10%) to its aspheric coefficient. A faceted CCD focal plane and individual field flattener lenses placed over each CCD correct the residual field curvature. An optically powered window maintains vacuum in the dewar. The optical design supports exchangeable flat filters, or the filter coatings may be deposited onto the field flatteners mounted over each CCD.

Improved yields for wafer-scale CCDs make large focal planes increasingly affordable. ZTF will use 16 e2v 6k×6k devices with 15 micron pixels. At 1"/pixel, the pixel scale is identical to that of PTF and adequately samples the median 2" image quality. The moderate pixel scale also mitigates the data volume. Six CCDs have been fabricated and delivered as of this writing. At 1 MHz readout, read time will be 10 sec; we require 15 sec net overhead between exposures to allow for slew and settling. With these shorter overheads, 30 sec exposures are optimal in maximizing \dot{V} . A compact dewar design minimizes mass and beam obstruction.

Table 1 compares the performance of the ZTF survey camera to that of PTF.

4 Selected Science Goals

4.1 Young SNe

Observations of SNe within the first 24 hours of explosion reveal key information about their progenitors and environments. Early photometric observations of SNe Ia constrain the radius of the progenitor and can distinguish single- and double-degenerate scenarios [7]. In core-collapse SNe, early observations probe the poorly-measured physics of shock breakout and shock heating [16]. Early-time “flash” spectroscopy of core-collapse SNe within hours of the explosion can directly measure the properties of the circumstellar medium and reveal the final stages of stellar evolution before the explosion [5].

Detecting, discovering, and following up young transients in a single night requires finely honed pipelines, procedures, and collaboration. The PTF and iPTF collaborations have demonstrated the ability to obtain these time-critical measurements on several occasions [6]. However, the total number of young SNe in the PTF data-stream is limited by the survey camera: obtaining the few-hour cadence observations needed to detect young SNe limits the survey to a much smaller area of sky. With ZTF’s wider, faster camera, the collaboration will be able to systematically study a true sample of SN progenitors rather than an isolated handful: we can detect twelve times more SNe with ZTF at any chosen cadence. In a high-cadence survey, ZTF will detect one SN within 24 hours of its explosion *every night*.

4.2 Fast-decaying transients

While PTF, CRTS, and Pan-STARRS1 have occasionally observed at relatively high cadences (images separated by less than a few hours), the correspondingly small areal coverage permitted by their survey cameras has limited the detection of fast transients to M-dwarf flares [1]. ZTF’s order-of-magnitude increase in survey speed will place much tighter constraints on the existence of fast-decaying explosive transients, exceeding published limits on areal exposure in less than one week of observations.

One intriguing event, PTF11agg [3], highlights the potential of ZTF in this area. Discovered by PTF during high-cadence monitoring of the Beehive Cluster for variable star studies, PTF11agg declined by almost two magnitudes over several hours. While its properties are consistent with an optical afterglow of a gamma-ray burst (GRB), there was no high-energy trigger from wide-field gamma-ray monitor. This raises the possibility that PTF11agg represents a new class of event, a baryon-loaded “dirty fireball” that would not show MeV emission. The inferred rate of such events would be about twenty times the GRB rate.

With ZTF’s faster survey speed, we expect to detect more than 20 PTF11agg-like events per year, as well as a handful of classical GRB orphan optical afterglows.

These measurements will place important constraints on the opening angles of GRB jets as well as the diversity of relativistic stellar explosions.

4.3 Gravitational Wave Counterparts

Beginning in 2015, advanced gravitational wave (GW) interferometers will begin operations. They are expected to detect the first GW signals from binary neutron star mergers. Detecting the electromagnetic counterparts to these events will provide vital physical information, including independent distance estimates and information about the merger progenitors and host galaxy. The mergers are predicted to produce optical counterparts, whether from afterglows of short-hard gamma-ray bursts or “kilonovae” powered by r-process nucleosynthesis [13, 11, 14, 8].

Unfortunately, the earliest GW detections will be very poorly localized, with error boxes of hundreds of square degrees with only two detectors and improving to tens of square degrees as more interferometers come online. Detecting a rapidly-decaying optical transient with unknown brightness in this large sky area is a monumental challenge. Success will require a well-tested technical stack, including all-sky reference images, fast and reliable image differencing, a complete local galaxy catalog to prioritize followup, and the ability to obtain rapid spectroscopy [9]. iPTF has proven this approach by successfully using its transient pipeline to localize the afterglows of Fermi-detected gamma-ray bursts within 70 square degrees [21]. ZTF’s wider field will be vital for achieving the same success with the larger search areas and fainter counterparts provided by GW detections.

4.4 Variability Science

The repeated observations provided by PTF and other surveys have built an increasingly valuable photometric variability catalog. Single-filter time variability information may be used to identify and classify variable stars [19], identify binary systems, and measure the mass of the supermassive black holes in AGN systems [10]. Variable stars may be used to trace Galactic structure and identify dwarf galaxies [4, 20], thereby mapping the gravitational potential of the Milky Way and testing predictions of Λ CDM cosmology [2]. Photometric variability may even predict stellar parameters, including effective temperature, surface gravity, and metallicity [15].

ZTF’s greater survey speed will provide an unprecedented variability catalog. If observations are spread evenly over the entire visible Northern sky, we will obtain nearly 300 observations per field each year. CRTS currently provides the most uniform photometric variability coverage. ZTF will provide a larger number of observations as well as improved cadence and depth, enabling a wide range of variability science on sources accessible to moderate-aperture telescopes and advancing community involvement in advance of LSST’s deeper survey.

E. B. is grateful for useful conversations with Shri Kulkarni, Tom Prince, Richard Dekany, Roger Smith, Jason Surace, Eran Ofek, Mansi Kasliwal, Branimir Sesar, and Paul Groot.

References

- [1] Berger, E., Leibler, C. N., Chornock, R., et al. 2013, ApJ, 779, 18
- [2] Bullock, J. S., & Johnston, K. V. 2005, ApJ, 635, 931
- [3] Cenko, S. B., Kulkarni, S. R., Horesh, A., et al. 2013, ApJ, 769, 130
- [4] Drake, A. J., Catelan, M., Djorgovski, S. G., et al. 2013, ApJ, 765, 154
- [5] Gal-Yam, A. 2014, in American Astronomical Society Meeting Abstracts, Vol. 223, American Astronomical Society Meeting Abstracts, #235.02
- [6] Gal-Yam, A., Kasliwal, M. M., Arcavi, I., et al. 2011, ApJ, 736, 159
- [7] Kasen, D. 2010, ApJ, 708, 1025
- [8] Kasen, D., Badnell, N. R., & Barnes, J. 2013, ApJ, 774, 25
- [9] Kasliwal, M. M., & Nissanke, S. 2013, ArXiv e-prints, [arXiv:1309.1554](https://arxiv.org/abs/1309.1554) [astro-ph.HE]
- [10] Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, ApJ, 698, 895
- [11] Kulkarni, S. R. 2005, ArXiv Astrophysics e-prints, [arXiv:astro-ph/0510256](https://arxiv.org/abs/astro-ph/0510256)
- [12] Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, PASP, 121, 1395
- [13] Li, L.-X., & Paczyński, B. 1998, ApJL, 507, L59
- [14] Metzger, B. D., Martínez-Pinedo, G., Darbha, S., et al. 2010, MNRAS, 406, 2650
- [15] Miller, A., Richards, J., Bloom, J. S., & on behalf of a larger team. 2014, in American Astronomical Society Meeting Abstracts, Vol. 223, American Astronomical Society Meeting Abstracts, #125.01
- [16] Nakar, E., & Sari, R. 2010, ApJ, 725, 904
- [17] Rahmer, G., Smith, R., Velur, V., et al. 2008, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7014, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
- [18] Rau, A., Kulkarni, S. R., Law, N. M., et al. 2009, PASP, 121, 1334
- [19] Richards, J. W., Starr, D. L., Miller, A. A., et al. 2012, ApJS, 203, 32
- [20] Sesar, B., Grillmair, C. J., Cohen, J. G., et al. 2013, ApJ, 776, 26
- [21] Singer, L. P., Cenko, S. B., Kasliwal, M. M., et al. 2013, ApJ, 776, L34
- [22] Tonry, J. L. 2011, PASP, 123, 58

SkyMapper and Supernovae

*Richard Scalzo
Australian National University*

Abstract

The SkyMapper Southern Sky Survey will be a wide-area digital survey of the southern sky, run from the robotic 1.3-m SkyMapper telescope at Siding Spring Observatory near Coonabarabran, NSW, Australia. The survey will include a supernova search run during poor seeing time, run as a rolling search to produce high-quality light curves for Hubble diagram cosmology. The search is currently taking data in science verification mode. I will briefly describe SkyMapper and then give an overview of supernova search activities, including pipeline design, operations, and interaction with the community.

The Catalina Real-Time Transient Survey (CRTS)

*Andrew Drake
California Institute of Technology*

Abstract

The Catalina Real-time Transient Survey (CRTS) is a completely open, VOEvent-enabled, optical transient survey that provides a model for the large synoptic surveys of the future. CRTS has so far discovered more than 7,000 highly variable and transient sources including 2,000 supernovae and 1,000 cataclysmic variables. I will highlight some of the rare and extreme types optical transients discovered by CRTS, as well as how increases in coverage and cadence of our second generation project, CRTS-II, will aid the discovery of new types of transient objects and phenomena. Lastly, I will discuss on-going efforts to characterize the variable sky using nine years of Catalina data for 500 million sources.

Pan-STARRS Transients, Recent Results, Future Plans

*Kenneth Chambers
Institute for Astronomy, University of Hawaii*

Abstract

The Pan-STARRS1 Surveys have discovered and provided precision photometry of large numbers of transients including SN Ia, new classes of ultra-luminous and under-luminous supernova, tidal disruption events and fast transients. Recent science results will be presented, together with plans for the public release of all PS1 data products, and for the upcoming PS1-PS2 Surveys starting March 2014, including the capability to respond to LIGO events.

Gaia – Revealing the Transient Sky

*Nicholas Walton
Institute of Astronomy, University of Cambridge, UK*

Abstract

The European Space Agency Gaia mission will launch 20 Nov 2013. It is set to perform a detailed census of a billion stars in our Milky Way. Through its on board astrometric, photometric, spectro-photometric and high resolution spectroscopic instrumentation it will be able to accurately determine the distances, positions, motions and astrophysical parameters to stars throughout the Milky Way. The impact of Gaia will be felt across all areas of astrophysics, primarily by revolutionising our knowledge of accurate stellar distances, through microarcsec level parallax measurements, across the Milky Way.

Gaia will also have a major impact in discovery and characterising of the 'Transient Sky'. Over its 5 year baseline mission operations - it will observe each point on the sky on average 70 times. It will discover many transient and variable objects, with a rich yield of objects ranging from rapidly moving near earth objects to distant supernovae and tidal disruption events.

This presentation, on the eve of the launch of Gaia, will describe the mission, and its potential for furthering our understanding of the transient sky. The alert data stream from Gaia will be described, noting the technical complexity involved in ensuring that science alerts from Gaia are rapidly distributed to the community. The nature of the processing chain of the alerts system will be noted, showing how the rich data from Gaia available for each alert can be utilised to enable the determination of a reliable source classification for each event. The formation of followup networks to effectively maximise the science from the alerts will be described - providing opportunities for all to participate in this.

Photometric Science Alerts from Gaia

*Heather Campbell, N. Blagorodnova, M. Fraser
G. Gilmore, S. Hodgkin, S. Koposov, and N. Walton
Institute of Astronomy, University of Cambridge
Madingley Road, Cambridge, CB3 0HA, UK*

*L. Wyrzykowski
Warsaw University Astronomical Observatory, 00-478 Warszawa, Poland*

1 Abstract

Gaia is a European Space Agency (ESA) astrometry space mission, and a successor to the ESA Hipparcos mission. The main goal of the Gaia mission is to collect high-precision astrometric data (i.e. positions, parallaxes, and proper motions) for the brightest one billion objects in the sky. This data, complemented with G band, multi-epoch photometric and low resolution (lowers) spectroscopic data collected from the same observing platform, will allow astronomers to reconstruct the formation history, structure, and evolution of the Galaxy.

In addition, the Gaia satellite is an excellent transient discovery instrument, covering the whole sky (including the Galactic plane) for the next 5 years, at high spatial resolution (50 to 100 mas, similar to the Hubble space telescope (HST)) with precise photometry (1% at G=19) and milliarcsecond astrometry (down to \sim 20mag). Thus, Gaia provides a unique opportunity for the discovery of large numbers of transient and anomalous events, e.g. supernovae, black hole binaries and tidal disruption events. We discuss the validation of the alerts stream for the first six months of the Gaia observations, in particular noting how a significant ground based campaign involving photometric and spectroscopic followup of early Gaia alerts is now in place. We discuss the validation approach, and highlight in more detail the specific case of Type Ia supernova (SNe Ia) to be discovered by Gaia. The intense initial ground based validation campaign will ensure that the Gaia alerts stream for the remainder of the Gaia mission, are well classified.

2 What is a Photometric Science Alert?

A photometric science alert is the appearance of a new source, or a change in flux, which suggests we could learn something from prompt ground-based follow-up. This

does not include: periodic variable stars (these sources may be better left to the end of the mission) and moving objects (however, astrometric microlensing would be an exception). The science alerts will be made public, within one to two days of Gaia detection, most of this time is due to downloading the data from the satellite.

3 Potential Triggers

Potential triggers for the the Gaia science alerts are objects of scientific interest which would benefit from fast ground based follow-up, as just discussed. Some examples of sources which maybe potential triggers include supernovae, super-luminous supernovae, tidal disruption events, cataclysmic variables, outbursts and eclipses from young stellar objects, X-ray binaries, microlensing events and other theoretical or unexpected phenomena. Figure 1 shows some of these potential triggers and the area of pars space they occupy for their brightness as a function of duration.

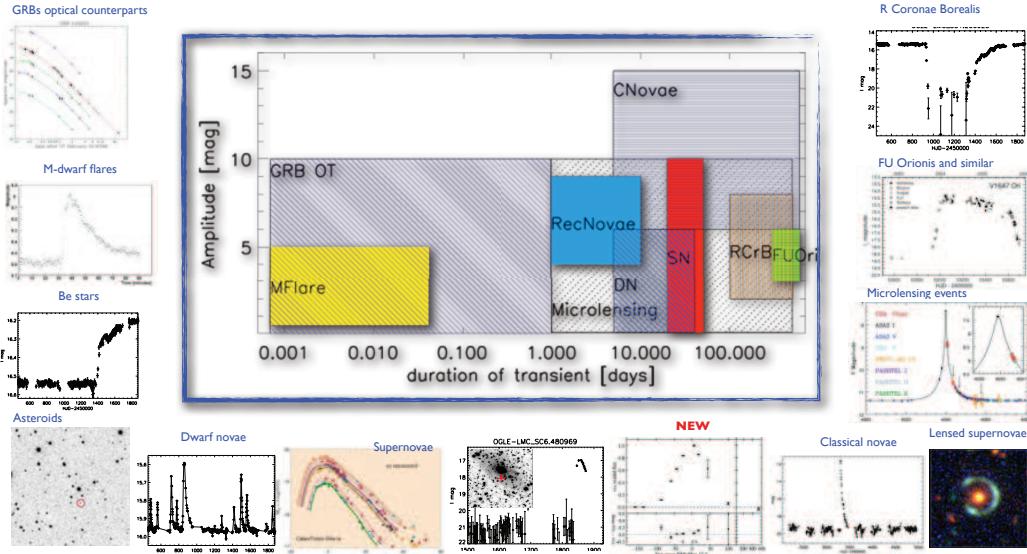


Figure 1: This shows the amplitude and duration of a range of potential triggers for the Gaia science alerts.

4 Gaia as a Transient Search Machine

Gaia is comparable to other transient search machines, such as the Catalina Sky Survey and the Palomar Transient Factory, as shown in Table 1, which covers similar

areas each day and similar limiting magnitude. The disadvantage of the Gaia survey is that the average cadence is only \sim 30days whereas transient surveys usually have a cadence of approximately 3 to 5 days. However, there is also a shorter cadence of 106.5 mins from the two mirrors in the satellite, also sometimes a 253.5 mins cadence, and sometimes 3 or more observations are thus obtained (when close to the 45 degrees ecliptic latitude zones for example). This 106.5 mins cadence is a huge advantage and means that changes in brightness should be detected quickly. Also, Gaia will cover the whole sky (including the Galactic plane), which is a significant survey area increase over other transient searches. The Gaia transient alerts will also have high spatial resolution with precise photometry (1% at G=19) and milliarcsecond astrometry (down to \sim 20mag), lowres spectra for all objects brighter than \sim 19mag and colours for fainter objects (see [5] for details of the photometry and lowres spectra).

Information	Gaia	Catalina Sky Survey	Palomar Transient Factory
deg ² day ⁻¹	~1230	1500	1000
Avg Cadence	\sim 30 days	14 days	5 days
Limiting mag	20	19.5	21
fsky	all sky	0.6	0.2

Table 1: Predicted numbers for the Gaia transient search compared to some ongoing surveys.

5 Time line

The Gaia satellite was launched on the 19th December 2013, and has now successfully been placed into orbit around the second Lagrange point. Over the next few months the telescope will undergo system shake-down and ESA commissioning (Figure 2). It is planned that in June the Gaia satellite will spend a month scanning the Ecliptic Poles internally verifying the data, and learning how to identify large amplitude variable stars (potential contaminants of the Gaia Science Alerts stream).

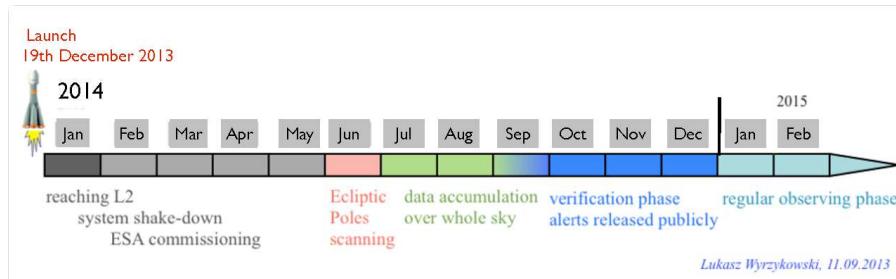


Figure 2: Current approximate timeline for Gaia operations and data accumulation.

Then in July Gaia will switch to nominal scanning and history of the whole sky will begin to be accumulated. In Figure 3 we show the expected coverage of Gaia by the end of July and then the end of September 2014. This will give some history of each patch of sky in the Gaia passbands and allow detection of transient objects. We propose to begin Gaia Alerts Spectroscopic Follow-up in the last weeks of August and the first week of September.

6 Scanning law

The Gaia satellite consist of two telescopes, which are projected onto one focal plane. The time between the two fields of view being observed is 106.5 mins and then the time between subsequent scans is 6 hours. After these initial observations the field will be revisited every \sim 10-30 days. Over the full mission each patch of sky will be measured, on average, approximately 70 times. The densest coverage is at 45 degrees to the ecliptic plane and this region is covered with approximately 200 epochs.

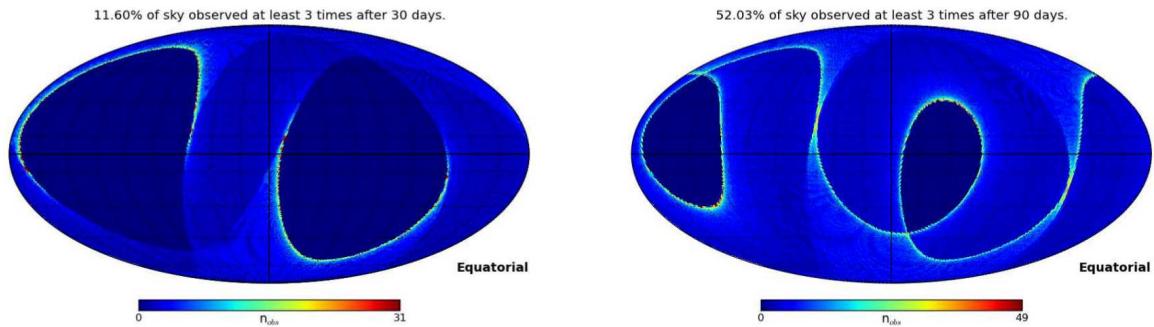


Figure 3: By 30 days 11.6% of the sky has been observed at least 3 times by Gaia. By 90 days, 52.03% of the sky has been observed at least 3 times by Gaia.

7 SN discovery rates

Simulations, [2] and updated by [1], predict Gaia will see \sim 6000 SNe down to G=19 (3/day), and twice this to G=20. One SN per day will be brighter than 18th magnitude (see Fig 4). For cataclysmic variables (CVs) the rate will be approximately similar, and Breedt, (priv. comm) predict Gaia should find 1000 new CVs. [3] predict that Gaia will find of order 20 Tidal disruption event's (TDE's) per year. Young stellar objects outbursts will be less common and Gaia will probably only find a few per year.

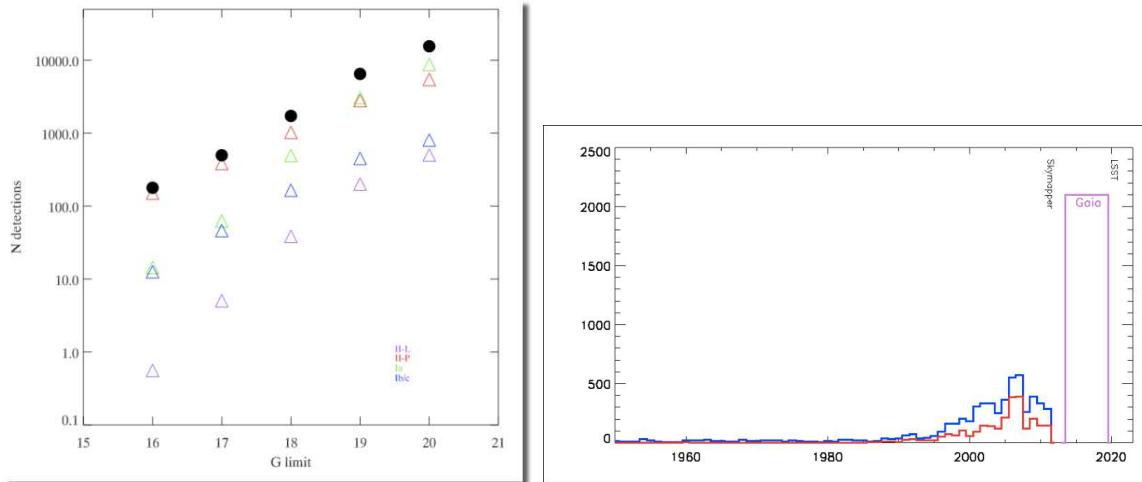


Figure 4: [Left]: Predicted SN detections with Gaia as a function of G-band magnitude. [Right]: Comparison between the Gaia SNe discovery rate (SN/yr) and current surveys. The open Gaia histogram is the number expected to 20th magnitude ($> 2000/\text{yr}$).

8 Alert Publication

Alerts are expected to be discovered and published to the world within \sim 24-48 hours of observation by the satellite. The Alert Stream will go live once Gaia has mapped at least 10% of the sky, a minimum of 3 times, which takes approximately one month (see Fig 3). Once the Gaia alert stream is fully operational all alerts will be made publicly available, and thus accessible for use by the community in their dedicated followup campaigns (see Section 10). During the commissioning, initialisation and early operations phase of Gaia (January - August 2014) - there will be systematic validation of the Gaia alerts, whereby the operational system will be assessed before going ‘Live’. The science alerts will be available to the community in web-based and email-based formats and will be produced in Virtual Observatory Event (VOEvent) - machine-readable format.

Each Alert package will consist of: coordinates, magnitudes, light curves, spectra, colours, proper motions, parallaxes (when available), astrophysical parameters (pars) (when available), features (random forest classifier see Section 9), classifier probabilities, cross match results.

9 Classification

Gaia is predicted to detect 44 million transits per day, which is \sim 150 - 800 GByte/day of data. Within this huge volume of data we expect 100s -1000s of potential interesting astrophysical triggers per day (real variables/moving objects). This precludes visual classification of a rich data stream and thus automated methods which are fast, repeatable and tuneable are essential. The Gaia alerts classification pipeline uses random forest classification. The random forest will use all the information available, and its features will include: light curve photometry (gradient, amplitude, historic rms, magnitude, signal-to-noise ratio, transit rms), lowres spectra (flux v lambda, colours, spectral shape coefficients, spectral type), auxiliary information (neighbour star, shape pars, motion pars, coords, crowding, calibration offset, correlations, QC pars) and crossmatch environment (near known star mags, near known variable class, near galaxy, near galaxy redshift and circumnuclear).

To build up a sufficient sample of classification labels in order to train the random forest classifier (e.g. [4]) we aim to observe \sim 500s homogenous high-quality spectra in the first year of the mission, spread across each broad class of transient phenomena (active galactic nuclei, core collapse SN, TDE, SN, Novae, CV and variable stars).

The light curve classification utilises the flux gradient of the transient object. The Gaia observations with 106.5 mins cadence are used to indicate the type of object. The lowers (BP/RP) spectra provide far more information to aid classification [3] and provide robust class for most objects, at >19 mag, when the classifier is fully trained on representative data. In addition, the transient object will be cross matched with archival catalogues, for example, Sloan Digital Sky Survey (SDSS), Two Micron All Sky Survey (2MASS), HST and Visible and Infrared Survey Telescope for Astronomy (VISTA). This will help remove known variable star contaminants and provide environmental information for the transient events, e.g. is there a host galaxy associated with the source and if so what is the type and magnitude.

10 Follow up

We are also co-ordinating a large program of photometric follow-up to improve the light curve sampling of Gaia transients. 47 x (7cm-2m) telescopes are listed as currently active (<http://bit.ly/1aHNXzy>) and 13 observatories are already doing tests (<http://bit.ly/17ViW7s>). All make use of our photometric calibration server (a tool developed to maximise the usefulness of the photometric followup data) to place the disparate data onto the same system (Wyrzykowski et al. 2013 *ATEL*#5245). Additionally, Las Cumbres Observatory Global Telescope Network (LCOGT) are expected to play a key role in the follow-up especially of μ lensing and young star transients. We point out the strong synergies with external facilities operating at different wave-

lengths. We will be able to confirm and characterise e.g. Low Frequency Radio Array (LOFAR) transients, and we may also trigger prompt *SWIFT* follow-up for particularly interesting events.

There is also a large educational (mostly utilising the Faulkes telescopes) and amateur involvement planned in the followup of these transient events, to assist in compiling light curves and increase the public evolvement and interest.

We need a large sample of well-exposed ($S/N \sim 20-50$), medium-dispersion ($R \sim 500 - 1000$) spectra, over a wide range of classes and magnitudes, to build classification training sets, in order for our (Random Forest) machine learning algorithms (discussed in Section 9) to perform well for the Gaia spectra for the remainder of the mission. Therefore we aim to obtain 1.5-4m telescope time to build this training set. It is important to invest time at the beginning of the Gaia mission to understand and characterise the transients that will be discovered with Gaia, so that we can optimise the process, and ensure that the rest of the mission is as productive as possible. We also intend to archive and release our spectroscopic classifications promptly after processing each night's observing.

11 Summary

The alert stream is non-proprietary and will be (some of) the first data from Gaia Summer 2014. We have planned an extensive follow-up program for classifying large numbers of transients: e.g. 10,000 SNe Ia over the whole sky. The alerts will be published one to two days after the event was initially detected (most of this time is due to the time taken for the data to be down linked from the satellite and processed). The alerts will be preliminarily classified using random forest classifiers based on the Gaia photometry and lowres spectra with additional cross match information from existing surveys. These classifications should improve after the first few months of ground based followup and retraining of the Bayesian classifiers. The alerts will be published in the VO format. For more information visit: <http://www.ast.cam.ac.uk/ioa/wikis/gsawgwiki>.

Acknowledgements Material used in this work has been provided by the Coordination Unit 5 (CU5) of the Gaia Data Processing and Analysis Consortium (DPAC). They are gratefully acknowledged for their contribution.

HCC, STH, GG, NAW, LW are members of the Gaia Data Processing and Analysis Consortium (DPAC) and this work has been supported by the UK Space Agency. NB has been supported by the The Gaia Research for European Astronomy Training (GREAT-ITN) network, funded through the European Union Seventh Framework Programme ([FP7/2007-2013] under grant agreement n° 264895.

References

- [1] G. Altavilla, M. T. Botticella and E. Cappellaro, **341**, 163-178 (2012)
- [2] V. A. Belokurov and N. W. Evans, MNRAS **341**, 569-576 (2003)
- [3] N. Blagorodnova, L. Wyrzykowski, and N. Walton et al., in prep
- [4] J. S. Bloom, J. W. Richards, P. E. Nugent, **124**, 921, 1175-1196 (2012)
- [5] C. Jordi, M. Gebran, J. M. Carrasco et al., **523**, A48, (2010)

Time-Domain Surveys: Moving Objects and Exo-Planets

Small Body Populations According to NEOWISE

*Amy Mainzer
Jet Propulsion Laboratory*

Abstract

The Wide-field Infrared Survey Explorer (WISE) surveyed the entire sky in four infrared wavelengths (3.4, 4.6, 12 and 22 microns) over the course of one year. From its sun-synchronous orbit, WISE imaged the entire sky multiple times with significant improvements in spatial resolution and sensitivity over its predecessor, the Infrared Astronomical Satellite. Enhancements to the WISE science data processing pipeline to support solar system science, collectively known as NEOWISE, enabled the individual exposures to be archived and new moving objects to be discovered. When the solid hydrogen used to cool the 12 and 22 micron detectors and telescope was depleted, NASA supported the continuation of the survey in the 3.4 and 4.6 micron bands for an additional four months to search for near-Earth objects and to complete a survey of the inner solar system. In total, NEOWISE detected more than 158,000 minor planets, including >34,000 new discoveries. This mid-infrared synoptic survey has resulted in range of scientific investigations throughout our solar system and beyond. Following one year of survey operations, the WISE spacecraft was put into hibernation in early 2011. NASA has recently opted to resurrect the mission as NEOWISE for the purpose of discovering and characterizing near-Earth objects.

The Catalina Sky Survey for Near-Earth Objects

*Eric Christensen
The University of Arizona*

Abstract

The Catalina Sky Survey (CSS) specializes in the detection of the closest transients in our transient universe: near-Earth objects (NEOs). CSS is the leading NEO survey program since 2005, with a discovery rate of 500-600 NEOs per year. This rate is set to substantially increase starting in 2014 with the deployment of wider FOV cameras at both survey telescopes, while a proposed 3-telescope system in Chile would provide a new and significant capability in the Southern Hemisphere beginning as early as 2015. Elements contributing to the success of CSS may be applied to other surveys, and include 1) Real-time processing, identification, and reporting of interesting transients; 2) Human-assisted validation to ensure a clean transient stream that is efficient to the limits of the system ($\sim 1\sigma$); 3) an integrated follow-up capability to ensure threshold or high-priority transients are properly confirmed and followed up. Additionally, the open-source nature of the CSS data enables considerable secondary science (i.e. CRTS), and CSS continues to pursue collaborations to maximize the utility of the data.

Time-Series Photometric Surveys: Some Musings

*Steve B. Howell
NASA Ames Research Center*

1 Introduction

Time-Series surveys are designed to detect variable, transient, rare, and new astronomical sources. They aim at discovery with a goal to provide large samples of "this and that". They are not designed to provide detailed study or analysis of individual objects. Detected sources are classified as variable if they show light and/or motion variations above some survey threshold. We ignore here changes due to uninteresting phenomena such as seeing or focus. What a survey delivers as a variable (or a constant) source critically depends on the photometric precision obtained and the process of data calibration and light curve processing. Observations in different Galactic locations and obtained with different filters (wavelengths) will find different populations of constant and variable sources.

Given the above complexities and nuances of time-series surveys, their interest lies in the sources they discover, especially the variable sources. Of these, the interesting sources are the prime driver of large efforts involving source classification, especially in near real-time. Note "interesting" can mean that a source is rare, highly variable, well understood, poorly understood, capable of follow-up, etc.

Source classification is a complex problem but can become manageable and even highly successful if one limits the total parameter space in which classification is attempted. For example, for a specific time sampling, certain classes of object will or will not be detectable. (This is not as clear cut as it sounds. For example, low amplitude periodic signals, not obvious in the data, can be teased-out of datasets that are long compared to the period.) Total time coverage is another example to consider. Attempts to classify sources from a survey of 30 days in total length with template models for many-hundred day semi-regular variables would be non-productive. Rise and fall times and light curve shape are additional temporal factors to keep in mind. Classification will always improve in accuracy as the number of samples for any given source increases.

Finally, the survey photometric precision will greatly limit the type of variable source that can be detected. Most modern large-area surveys will deliver a bright source single measurement photometric precision near 0.01-0.005 magnitude, with better results planned through co-addition. Controlling systematics using a standard

observing protocol and consistent data pipeline reduction procedures will be the keys to reaching these precision limits.

2 Two Illustrative Examples

To illustrate some of these points, I provide two examples. In the first you see the importance of time-sampling, both cadence-to-cadence as well as longevity. I doubt many could identify the source in the top panel of Figure 1 as a RR Lyrae star of ~ 0.5 day period. In fact, this source might be identified as a lower amplitude variable with a period near 1 day.

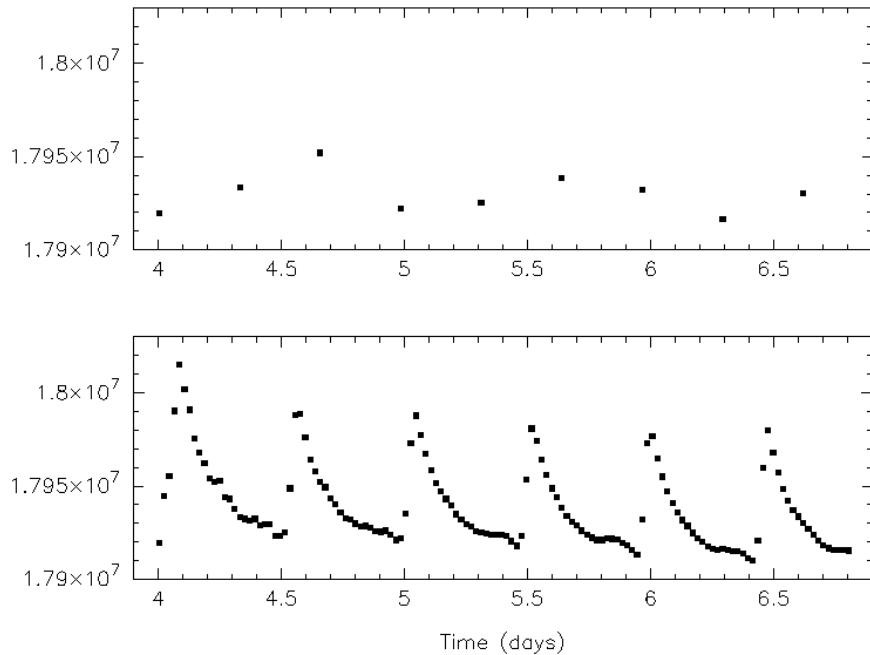


Figure 1: RR Lyrae star observed by the K2 mission during science verification tests. The top panel shows the light curve sampled every 7.5 hours - a proxy for a "once per night sampling". The bottom panel shows the full temporal resolution, sampled every 30 minutes, confirming a ~ 0.5 day RR Lyrae star.

In the second example, we are interested in transient sources. Transients are important as they often represent astrophysically interesting sources, such as supernovae, rare objects such as TOADs (Howell et al., 1995), or new astronomical discoveries. Figure 2 shows a recent example of a source that was believed to be interesting based on its very blue color selection and a past observation revealing outburst-like behavior. However, upon the onset of a detailed Kepler monitoring program, the source

(BOKS 45906, Feldmeier et al., 2011) was found to be very faint (near 22nd magnitude) and boring. That is, it showed essentially a complete lack of "interesting variations". After nearly a year of observation, BOKS 45906 redeemed itself, showing large amplitude, transient behavior and rapid flaring. This highly variable source fell back into obscurity about 1.5 years later, again becoming "boring". Today we believe the object is some sort of short period (56.5 min) interacting binary (Ramsay et al., 2013).

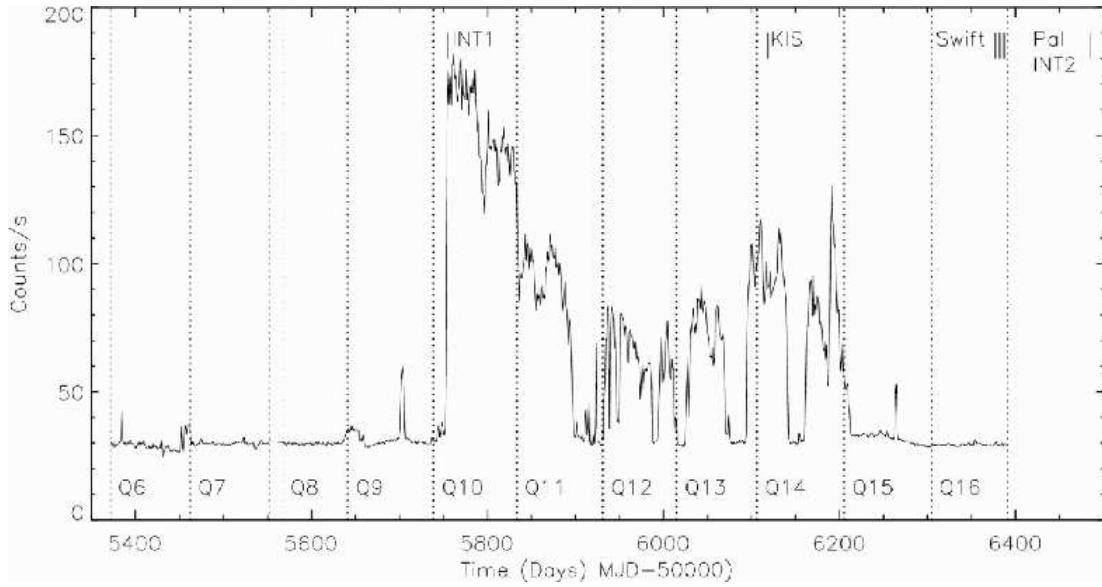


Figure 2: The Kepler light curve of the interacting binary star BOKS 45906, covering 1000 days, sampled every minute but plotted as 1 day bins. The time unit is in MJD - 50000.0. Note the long period (first year of data) showing effectively no variation - a boring source - followed by the rapid transient behavior (post day 5750). Suddenly, BOKS 45906 became very interesting!

3 Predicting Variable Sources in a Survey

Variability in a survey is dominated by low-amplitude, non-periodic sources. Periodic variables, such as pulsating stars or eclipsing binaries, make up only about 10% of all variables observed. This one fact alone has large ramifications for source classification, as non-periodic sources are tremendously difficult to categorize, especially the multitude with low modulation amplitudes. The number of variable sources, both periodic and non-periodic, that a survey will detect appears to be a universal function (see Howell 2008; Tonry et al., 2005) and, assuming relatively good sampling, is

related to the survey's photometric precision in magnitudes (σ) as follows (Fig. 3):

$$\% \text{Variable} = -23.95(\log \sigma) - 39.52.$$

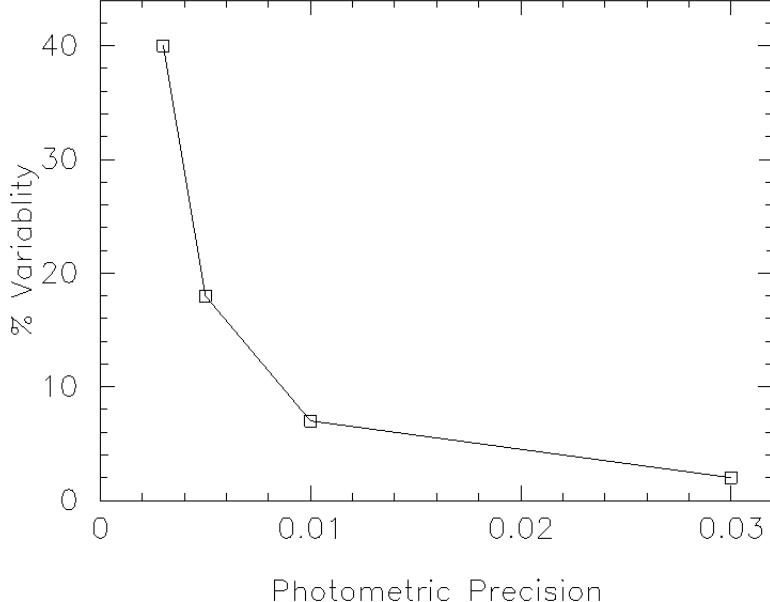


Figure 3: Survey variability fraction can be predicted. We show the universal relationship of the percentage of sources that will be found to be variable (both periodic and non-periodic) vs. the best photometric precision (mag) of the survey.

4 Conclusions

Our expectations for the findings and results from a survey, whatever they might be, are sometimes wrong. Surveys all have biases. Keep them in mind and try to avoid them in your thinking or at least realize they are present. Surveys are wonderful large-scale experiments. Some lessons learned, based on trial and error and their pitfalls as well as their successes, are as follows.

Variable objects are often highly useful probes of fundamental properties in astrophysics. Remember, just because a source is variable does not mean that it is periodic; don't confuse the two. Only $\sim 10\%$ of all variable objects will be periodic. The periodic sources, however, are much more easily classified and often much more astrophysically important and useful. The non-periodic variables, while in the vast majority, are the least understood. Perhaps they are full of potential, waiting to teach us much about the universe.

One can predict the percentage of variables that a survey will detect using an apparently robust relationship. Such predictions are highly useful in order to assess the intrinsic value of a survey and allow data collection and analysis requirements to be specified (e.g., Ridgway et al., 2014). Such a ubiquitous function is probably telling us something very important about the nature of variability.

Spectroscopy of discovered variables may not always provide an answer as to source identification. Some sources, especially those with small intrinsic photometric variations, may not be spectroscopically variable. Remember, traditional analysis techniques tend to yield traditional results. New and different analysis techniques, such as sonification of variability (Tutchton, R., et al., 2012) may help to reveal new insights.

Keep Watching the Skies!

I'd like to thank my many collaborators on the Kepler/K2 team as well as those many others with whom I have variable relations. I'd particularly like to acknowledge the *Hot-Wired* organizers for their initial idea and continued effort to bring together the diverse talents that are needed to study and understand the transient universe.

References

- [1] Feldmeier, J., et al., 2011, AJ, 142,2
- [2] Howell, S. B., Szkody, P., and Cannizzo, J., 1995, ApJ, 439, 337
- [3] Howell, S. B., 2008, AN, 329, 259
- [4] Ramsay, G., et al., 2013, MNRAS, 438, 789
- [5] Ridgway, S., et al., 2014, ApJ, in press
- [6] Tonry, J., Howell, S. B., Everett, M., et al., 2005, PASP, 117, 281
- [7] Tutchton, R., et al., 2012, JSARA, 6, 21

Passing NASA's Planet Quest Baton from Kepler to TESS

*Jon Jenkins
SETI Institute*

Abstract

Kepler vaulted into the heavens on March 7, 2009, initiating NASAs search for Earth-size planets orbiting Sun-like stars in the habitable zone, where liquid water could exist on a rocky planetary surface. In the 4 years since Kepler began science operations, a flood of photometric data on upwards of 190,000 stars of unprecedented precision and continuity has provoked a watershed of 134+ confirmed or validated planets, 3200+ planetary candidates (most sub-Neptune in size and many comparable to or smaller than Earth), and a resounding revolution in asteroseismology and astrophysics. The most recent discoveries include Kepler-62 with 5 planets total of which 2 are in the habitable zone with radii of 1.4 and 1.7 R_e . The focus of the mission is shifting towards how to rapidly vet the 18,000+ threshold crossing events produced with each transiting planet search, and towards those studies that will allow us to understand what the data are saying about the prevalence of planets in the solar neighborhood and throughout the galaxy. This talk will provide an overview of the science results from the Kepler Mission and the work ahead to derive the frequency of Earth-size planets in the habitable zone of solar-like stars from the treasure trove of Kepler data.

NASAs quest for exoplanets continues with the Transiting Exoplanet Survey Satellite (TESS) mission, slated for launch in May 2017 by NASAs Explorer Program. TESS will conduct an all-sky transit survey to identify the 1000 best small exoplanets in the solar neighborhood for follow up observations and characterization. TESSs targets will include all F, G, K dwarfs from +4 to +12 magnitude and all M dwarfs known within \sim 200 light-years. 500,000 target stars will be observed over two years with \sim 500 square degrees observed continuously for a year in each hemisphere in the James Webb Space Telescopes continuously viewable zones. Since the typical TESS target star is 5 magnitudes brighter than Kepler's and 10 times closer, TESS discoveries will afford significant opportunities to measure the masses of the exoplanets and to characterize their atmospheres with JWST, ELTs and other exoplanet explorers.

The ATLAS All-Sky Survey

*Larry Denneau
Institute for Astronomy, University of Hawaii*

Abstract

The Asteroid Terrestrial-impact Last Alert System (ATLAS) is a small project with an ambitious goal: early warning of asteroid impacts on Earth. We aim to provide one day warning for the smallest "town-killer" 30-kiloton asteroids up to three weeks for a 100-megaton impactor. ATLAS will execute a wide-field all-sky survey with four visits per footprint per night down to a sensitivity limit of V=20, suitable for detection dangerous asteroids and enabling other exciting time-domain astronomy. ATLAS is currently under construction and expects to be fully operational in late 2015. We provide an overview of the ATLAS system and discuss how ATLAS can participate in the emerging community of time-domain astronomy.

The Performance of MOPS in a Sensitive Search for Near-Earth Asteroids with the Dark Energy Camera

*David Trilling
Northern Arizona University*

*Lori Allen
National Optical Astronomy Observatory*

Abstract

We have been using LSST's Moving Object Pipeline System (MOPS) with data from NOAO's Dark Energy Camera and from the Spitzer Space Telescope, among other platforms. MOPS allows us to link moving object source detections and find moving objects in those data streams. I will report on the performance of MOPS for these projects and some lessons learned.

Time-Domain Surveys: Beyond Optical Photometry

The TeV Sky Observed by the High-Altitude Water Cherenkov Observatory (HAWC)

*Brenda Dingus
Los Alamos National Laboratory*

Abstract

The High Altitude Water Cherenkov (HAWC) gamma-ray observatory began science operations 1 Aug 2013 with $\sim 1/3$ of the full detector. The rest of the detector will be constructed within the next year. Even now, HAWC is the most sensitive, wide field of view (~ 2 sr), continuously operating, TeV gamma-ray observatory ever constructed. Every day HAWC observes the sky from a Declination of -30 deg to $+70$ deg. HAWC is searching for transient sources such as flares from active galactic nuclei and gamma ray bursts as well as measuring Galactic sources to the highest energies. I will present preliminary first results from HAWC.

Hearing & Seeing the Violent Universe

*Samaya Nissanke
California Institute of Technology*

Abstract

Joint gravitational-wave (GW) and multi wavelength electromagnetic (EM) observations of compact binary mergers should enable unprecedented studies of astrophysics in strong-field gravity, and of binary stellar evolution. Within the next decade, a worldwide network of advanced versions of ground-based GW interferometers should become operational from 10 Hz to a few kHz. At these frequencies, inspirals and mergers of neutron star binary mergers are expected to be amongst the most numerous and strongest GW-emitting sources. A subset of these events could be associated with a transient EM counterpart, and should be observable at different wavelengths, energies and timescales. In this talk, I will first discuss the EM counterparts that we expect to see from compact binary mergers and then describe how we can search and identify such EM counterparts using a slew of high-energy, optical and radio telescopes in the near future.

Follow-up of LIGO-Virgo Observations of Gravitational Waves

*Roy Williams
California Institute of Technology*

Abstract

In the next few years, the advanced LIGO and Virgo detectors will be operational, and hopefully detecting coalescences of compact objects such as neutron stars and black holes. The talk will review the science, the observational prospects, and how to get your telescope involved in this exciting science.

The Needle in the Hundred-Square-Degree Haystack: from Fermi GRBs to LIGO Discoveries

Leo Singer

California Institute of Technology

Abstract

Accurate localizations have driven and enriched our understanding of gamma-ray bursts. They could do the same for future gravitational-wave detections with LIGO and Virgo. We report the discovery of the optical afterglow of the gamma-ray burst (GRB) 130702A, identified upon searching 71 square degrees surrounding the Fermi Gamma-ray Burst Monitor (GBM) localization. Discovered and characterized by the intermediate Palomar Transient Factory (iPTF), iPTF13bxl is the first afterglow discovered solely based on a GBM localization. Real-time image subtraction, machine learning, human vetting, and rapid response multi-wavelength follow-up enabled us to quickly narrow a list of 27,004 optical transient candidates to a single afterglow-like source. The bright afterglow and emerging supernova offered an opportunity for extensive panchromatic follow-up. Furthermore, our discovery of iPTF13bxl represents an important step towards overcoming the challenges inherent in uncovering faint optical counterparts to comparably localized gravitational wave events in the Advanced LIGO and Virgo era.

ARTS – The Apertif Radio Transient System

Joeri van Leeuwen
ASTRON/U. Amsterdam, Netherlands

Abstract

Apertif is a highly innovative receiver system that is currently being constructed for the Westerbork Synthesis Radio Telescope. Its factor 30 increase in field-of-view allows astronomers to survey the entire sky at 1.4 GHz with an unprecedented combination of sensitivity and speed. ARTS, the Apertif Radio Transient System, will extend this wide-field Apertif system to high time resolution, enabling unique searches for millisecond transients and pulsars. Beam formers and transient detectors, powerful enough to cover the full 9 square degree field of view, will run in real-time. These will provide triggers to facilities ranging from radio to high-energy regimes, for follow up and localization of fast, enigmatic transients.

Radio Adventures in the Time Domain

Dale Frail

National Radio Astronomy Observatory

Abstract

The Transient Universe is one of the key science themes of the newly expanded Karl G. Jansky Very Large Array (VLA). In the study of transients, the VLA is both a powerful survey instrument and it is the preeminent follow-up telescope at meter to centimeter wavelengths. It offers a wide variety of capabilities including superb instantaneous sensitivity, wide frequency coverage, dynamic scheduling, quick response to external triggers, and fast temporal sampling. Pipeline processing has recently been implemented, with the goal at making the VLA data products more accessible to the multi-wavelength community. I will describe these time-domain capabilities, with examples drawn from fields as varied as solar physics, flare stars, supernovae, gamma-ray bursts and EM counterpart searches for gravitational waves.

The Karl G. Jansky VLA Sky Survey (VLASS): Defining a New View of the Dynamic Sky

*Steven Myers
National Radio Astronomy Observatory*

Abstract

NRAO recently announced the VLA Sky Survey (VLASS) initiative to design and prosecute a new generation centimeter-wave synoptic sky survey using the newly upgraded Karl G. Jansky Very Large Array (VLA). The parameters of the survey will be defined by a open process guided by a community-led Survey Science Group (SSG). A call for White Papers has been issued to the community, as input to the upcoming VLASS Science Planning Workshop to be held on Jan 5, 2014 adjacent to the AAS meeting in Washington, DC. All interested members of the physics and astronomy community are welcome to participate. In this presentation we summarize the capabilities of the VLA

VLA Searches for Fast Radio Transients at 1 TB hour⁻¹

C. J. Law, and G. C. Bower
University of California, Berkeley

S. Burke-Spolaor
Jet Propulsion Laboratory / California Institute of Technology

B. Butler
National Radio Astronomy Observatory

E. Lawrence
Los Alamos National Laboratory

T. J. W. Lazio, and C. Mattmann
Jet Propulsion Laboratory

M. Rupen
National Radio Astronomy Observatory

A. Siemion
University of California, Berkeley

S. VanderWiel
Los Alamos National Laboratory

1 Fast Radio Transients

Fast radio transients are pulses of dispersed radio emission lasting less than 1 second. Slower radio transients originate predominantly in synchrotron emission, while faster transients are caused by coherent processes. Furthermore, at timescales faster than 1 second, propagation through the Galactic plasma induces dispersion, the frequency-dependent arrival time quantified by dispersion measure (DM), that begins to be detectable at MHz through GHz radio frequencies.

1.1 Fast, Extragalactic Bursts

Fast transients surveys at the Parkes Observatory has revealed a new population of radio transients: the FRB [1, 2]. Their DMs range up to 1100 pc cm^{-3} , an order of magnitude larger than expected from the Galaxy and consistent with propagation through the IGM from distances up to $z \sim 1$.

Basic questions about FRBs remain open: *What are they?* and *How can we use them?* If they do in fact lie at cosmological distances, their dispersion can measure the baryonic mass of the IGM, much as Galactic pulsars of known distance have mapped the electron content of the Milky Way [3]. Beyond using FRBs as probes, understanding the origin of FRBs may have relevance to gamma-ray bursts and sources of gravitational waves [4, 5].

1.2 Pulsars, RRATs

Similar pulsar surveys have discovered a new class of Galactic radio transient: the rotating radio transient [6]. It is unclear whether extreme objects like magnetars or ordinary pulsars can generate pulses detected as RRATs [7].

Moving slightly beyond our Galaxy, the most distant radio transients associated with a host galaxy are in M31 [8]. The dispersion measure of any pulses known to be in M31 would make the first constraint on the Milky Way and M31 halo baryon content, which would help address the "missing baryon problem" [9].

1.3 Flare Stars, Ultracool Dwarfs, Exoplanets

Jupiter emits intense radio bursts that make it the brightest astronomical object in the solar system below 100 MHz. Coronal mass ejections (much as seen in the Sun), also drive radio fast, coherent radio flares. These processes could be used to measure magnetism and plasma properties of other stars [10] and should profoundly affect the habitability of orbiting exoplanets [11].

2 Fast Imaging with the VLA

While single-dish telescopes have pioneered the study of fast radio transients, interferometers are poised to transform the field. Interferometers form "synthetic" apertures many kilometers in diameter, which allows them to expand on every limitation of single-dish telescopes:

- **Precise localization:** Interferometers image with arcsecond precision, as shown in the image of a pulsar pulse shown in Figure 1.

- **High survey speed:** Interferometers have large fields of view and are powerful survey machines.
- **Robust calibration and interference rejection:** Interferometers can measure fluxes more accurately and reject interference that complicates single-dish fast transient searches.

Interferometers are technically more demanding than single-dish telescopes because their fundamental measurement is the correlation of pairs of antennas. Thus, where a single-dish telescope has a single data stream (or a few, if using a multi-beam receiver), a comparable interferometer like the VLA has 27 antennas and thus 351 data streams. An efficient algorithm for extracting transients from this massive data stream could revolutionize the study of fast transients by uniquely associating radio transients with multiwavelength counterparts (e.g., FRB host galaxies, RRAT NS hosts, stellar/planetary associations).

We have commissioned the Jansky Very Large Array (VLA) to observe with millisecond integrations and data rates of 1 TB hour⁻¹ [12]. We have also developed an extensive, parallelized software system to search visibility data for dispersed transients¹. The pipeline is written in Python/Cython and run within the NRAO software package CASA².

3 Survey for Fast Radio Bursts

We are now conducting a large VLA survey for the highly-dispersed radio transients known as FRBs. The goal of the survey is to detect at least one FRB, localize it to arcsecond precision, and uniquely associate it with other objects. Assuming that the FRB has a host galaxy, unique associations can be made with arcsecond localizations out to a redshift of 1 [13].

This survey uses the VLA correlator to write an integration each 5 ms. Faster integrations would be more sensitive to the ~ 1 ms FRB pulse widths, but those data rates are not sustainable. Assuming that FRBs uniformly populate a cosmological volume, we expect to detect one FRB in roughly 35 hours of observing. We have targeted five locations at high Galactic latitudes to avoid confusion with Galactic dispersion.

Our goal is to observe 150 hours to detect 1–10 FRBs or exclude the published event rate with 99% confidence. At the time of this writing (January 2014), we have observed 78 hours and processed roughly half of that. No events have been found brighter than our confidence threshold of 8σ , which is equivalent to a flux density of

¹Portions of the code base are available at <http://github.com/caseyjlaw/tpipe>.

²See <http://casa.nrao.edu>.

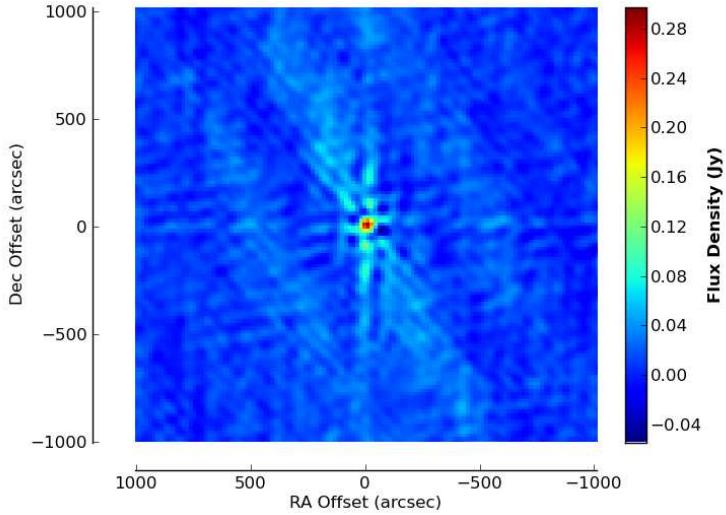


Figure 1: VLA image of a dispersed millisecond pulse from pulsar B0355+54. Our VLA FRB survey has mostly observed in "B" configuration, which has a resolution of 4 arcsec and localization precision roughly 10 times better.

130 mJy. At this threshold, we expect less than one false positive due to Gaussian noise in the entire survey.

The transient search pipeline is currently running on the NRAO Array Operations Center (AOC) cluster and on the "Darwin" cluster at Los Alamos National Lab (LANL). Data is transferred to LANL by mailing disks. We also have approved compute time and storage on the NERSC compute center. The search pipeline parallelizes DM trials over cores of a node and different time segments ("scans" in VLA parlance) are parallelized over nodes of the cluster.

The processing time and memory footprint are dominated by the FFT stage. The size of the image grid is determined by the VLA antenna configuration and ranges from 512–2048 on a side. In the more compact of these configurations (called "CnB"), the processing pipeline can search one hour of data in 70 hours on a single node, equivalent to roughly 340 images per second per node. The majority of our data were observed in a larger configuration (called "B") and processed several times slower.

4 Real-Time Detection as Solution to Big Data Challenge

The technical requirements for our VLA FRB survey are extreme in astronomy. However, planned observatories like the SKA and LSST are finding more science at high data rates. Lessons learned from our project will have increasing relevance to astronomers working to solve the "needle in a haystack" problem.

Extreme data rate science is limited by requirements of data production, distribution, processing, and curation. Our solution (snearkernet, three compute clusters) is manageable for a limited (and compelling!) science case, but it is not sustainable in the long term. We believe a sustainable solution will use real-time transient detection. Bringing computational support closer to telescopes ameliorates the distribution problem and lets us ignore data we know is uninteresting, a technique known as "data triage".

Figure 2 summarizes the concept of data triage in transient detection. In some applications, the process of measuring all information about a transient candidate may be substantially greater than simply detecting it³. The difference between the two can be critical for extreme data rate applications. Once a transient candidate is detected, the data associated with the candidate can be saved for more detailed analysis. Data triage is routinely employed in the particle physics community, where a well-defined theory predicts the interactions of a particle with the detector. A detailed theory is critical to define what the *absence* of a detection means.

³In the case of the VLA FRB project, most of our images contain zero-mean Gaussian distributed pixel values, so one can imagine a number of simple statistical tests of Gaussianity to determine whether a candidate transient is present.

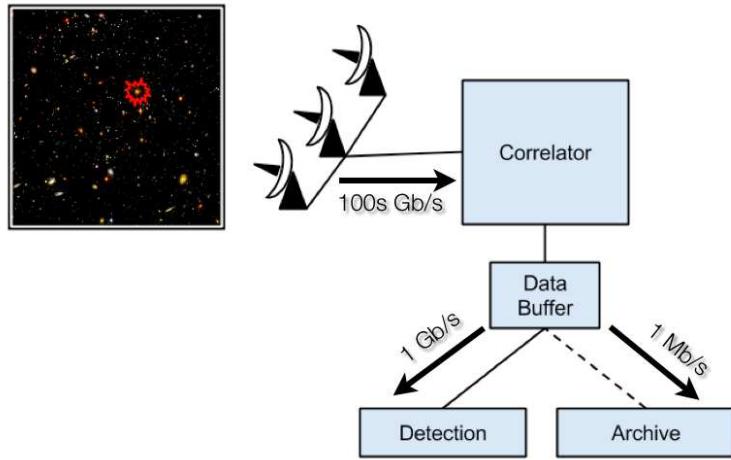


Figure 2: Schematic of the data flow of a real-time radio transient detection system utilizing data triage. For a radio telescope, the electromagnetic wave is Nyquist sampled (i.e., at GHz rates). A correlator then produces the fundamental data product at rates of order 1 Gb s^{-1} (equivalent to 450 MB hour^{-1}). An efficient transient detection algorithm [14] can reduce that data stream by orders of magnitude, making data management and detailed post-processing tractable.

References

- [1] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford. A Bright Millisecond Radio Burst of Extragalactic Origin. *Science*, 318:777–, nov 2007.
- [2] D. Thornton, B. Stappers, M. Bailes, B. Barsdell, S. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, D. J. Champion, P. Coster, N. D'Amico, A. Jameson, S. Johnston, M. Keith, M. Kramer, L. Levin, S. Milia, C. Ng, A. Possenti, and W. van Straten. A Population of Fast Radio Bursts at Cosmological Distances. *Science*, 341:53–56, jul 2013.
- [3] J. M. Cordes and T. J. W. Lazio. NE2001.I. A New Model for the Galactic Distribution of Free Electrons and its Fluctuations. *ArXiv Astrophysics e-prints*, jul 2002.
- [4] H. Falcke and L. Rezzolla. Fast radio bursts: the last sign of supramassive neutron stars. *ArXiv e-prints*, jul 2013.
- [5] K. Kashiyama, K. Ioka, and P. Mészáros. Cosmological Fast Radio Bursts from Binary White Dwarf Mergers. *ApJ*, 776:L39, oct 2013.
- [6] M. A. McLaughlin, A. G. Lyne, D. R. Lorimer, M. Kramer, A. J. Faulkner, R. N. Manchester, J. M. Cordes, F. Camilo, A. Possenti, I. H. Stairs, G. Hobbs, N. D'Amico, M. Burgay, and J. T. O'Brien. Transient radio bursts from rotating neutron stars. *Nature*, 439:817–820, feb 2006.
- [7] P. Weltevrede, B. W. Stappers, J. M. Rankin, and G. A. E. Wright. Is Pulsar B0656+14 a Very Nearby Rotating Radio Transient? *ApJ*, 645:L149–L152, jul 2006.
- [8] E. Rubio-Herrera, B. W. Stappers, J. W. T. Hessels, and R. Braun. A search for radio pulsars and fast transients in M31 using the Westerbork Synthesis Radio Telescope. *MNRAS*, 428:2857–2873, feb 2013.
- [9] J. N. Bregman. The Search for the Missing Baryons at Low Redshift. *ARA&A*, 45:221–259, sep 2007.
- [10] G. Hallinan, S. Bourke, C. Lane, A. Antonova, R. T. Zavala, W. F. Brisken, R. P. Boyle, F. J. Vrba, J. G. Doyle, and A. Golden. Periodic Bursts of Coherent Radio Emission from an Ultracool Dwarf. *ApJ*, 663:L25–L28, jul 2007.
- [11] J. C. Tarter, P. R. Backus, R. L. Mancinelli, J. M. Aurnou, D. E. Backman, and G. S. Basri. A Reappraisal of The Habitability of Planets around M Dwarf Stars.

Astrobiology, 7:30–65, mar 2007.

- [12] C. J. Law, G. C. Bower, M. Pokorny, M. P. Rupen, and K. Sowinski. The RRAT Trap: Interferometric Localization of Radio Pulses from J0628+0909. *ApJ*, 760:124, dec 2012.
- [13] J. S. Bloom, S. R. Kulkarni, and S. G. Djorgovski. The Observed Offset Distribution of Gamma-Ray Bursts from Their Host Galaxies: A Robust Clue to the Nature of the Progenitors. *AJ*, 123:1111–1148, mar 2002.
- [14] C. J. Law and G. C. Bower. All Transients, All the Time: Real-time Radio Transient Detection with Interferometric Closure Quantities. *ApJ*, 749:143, apr 2012.

Nuts and Bolts: Telescope Networking

Time to Revisit the Heterogeneous Telescope Network

Frederic V. Hessman

Institut für Astrophysik, Georg-August-Universität, Germany

1 Introduction

Time-domain astronomy demands a level of flexibility that is still not common in the organization and operation of normal astronomical observatories. With service observations, targets-of-opportunity (ToO), director's discretionary time, and other new-fangled features, many observatories have come a long way, and the increased flexibility has made it possible to perform some time-critical observations in a manner unthinkable within the classic observing-run paradigm. Nevertheless, these improvements have been made at the level of observatory operations and are largely carried out by humans and hence are not scalable to the scientific needs of the 21st century.

While the manner in which large cutting-edge telescopes are operated is inherently inflexible, the same should not be true of smaller telescopes (where "small" nowadays means, say, less than 2-3m). Since the number of potentially available "small" telescopes is large and the pressure factors are generally less than those of larger ones, there is a latent potential which could be used particularly well for classic follow-up purposes – e.g. ToO, surveys,... – if a means could be found to tap into that potential. There are, however, often local political / economical / sociological reasons for not wanting to yield up one's own telescope for a collaborative project. Ideally, then, observatories should be able to "donate" some fraction of the total amount of time available for the benefit of one or more external collaborations in which the observatory participates, knowing that the use of the data provided constitutes both a measurable scientific investment and reflects in some way positively on the donating institution. Given a means of accessing multiple resources within such a collaboration, an intelligent agent – principally a human but probably a piece of software – could allocate the observing time on a global scale and so optimize both the use of individual resources as well as the total scientific output of the collaboration.

This was the original idea behind the "Heterogeneous Telescope Network" (HTN) consortium: a loose interest-group of institutions, individual researchers, and even commercial companies interested in telescope networking. While several international HTN meetings were organized (e.g. Exeter 2005, Göttingen 2006) and various projects

have implemented sub-networks which have probed some of the possibilities, the HTN idea never got much further than to suggest a protocol for the exchange of observatory requests [1][10].

2 The Idea of a Telescope eMarket

The constraints on the operation of a truly flexible heterogeneous telescope network clearly define the sequence of communications between a server (the thing operating a telescope, whether a human or a piece of software) and a client (some intelligent agent trying to get some scientific project done). The analogy with an electronic market is very good, since some bargaining about the conditions of the “contract” is necessary and all of this ideally takes place semi-automatically.

- The client must have an idea of what resources are principally available (e.g. a telescope with a given aperture, camera or spectrograph, a particular filter inserted or removed, perhaps even a particular type of observational sequencing pattern). This information must be either broadcast or queried.
- The client then sends a request for a particular observation based upon the individual characteristics of each resource. Because the client is presumably querying several resources (not everyone has good weather!), the request is just a “would you in principle be willing and able to do this?” question.
- The server then responds and expresses a principle willingness to perform the observation or rejects the query.
- The client chooses the best resource based on the currently availability and resubmits the request as a formal request for a “contract” to perform the services.
- The server processes the request and either sends an acknowledgement or a refusal, along with a unique ID for the “contract” made between the server and client.
- Requests which take place over extended periods of time may require some sort of status info – if the observations are delayed enough, it could be that they are no longer needed and should be aborted.
- Finally, the successful observation is reported and the metadata and experimental data are transported.

The HTN consortium simply revised a previously extant protocol – “Remote Telescope Markup Language” [7][4] – and added on the additional transactional modes

needed to create a protocol perfectly sufficient to cover all of the steps above. Thus, the API is well-defined and built upon known technologies and tools.

None of this is rocket-science, but it is sufficiently complicated that the whole transaction must be recorded and processed at both ends. While a stripped-down system of

- I know what's out there.
- I know you are willing to be asked.
- When I ask, please do your best and send me an email if it works.

might work for very simple projects and very simple resources, anything more complicated than this requires a substantial subset of what an ideal eMarket for telescope time would require. Thus, it isn't surprising that an HTN didn't simply jump into existence. The reality was, that it simply takes a lot longer to get the hardware and software running and the resources working to a point where they are networkable. Also, the career paths of major players can easily disrupt the process by removing someone who has an important role in the development of a large informal consortium project. Frankly, even simple networks (e.g. our MONET "network" of just 2 telescopes) can be very hard to get and keep running sometimes. Finally, there is no point in doing all of this unless the science drivers are so strong that it becomes necessary to do the effort, thereby overcoming the various sources of sociological and institutional inertia and friction.

The HTN idea has thus been in a dormant state for almost a decade. Dormant doesn't mean comatose – the interest is there, and there are many relevant activities abuzz in the background. An initial HTN experiment was carried out by the RoboNet / eStar consortium but it depended upon a fairly homogenous software and control model. Thus, no coherent attempt at attaining a truly heterogeneous network has been made. There is a wonderful German word for such a state: "Dornrösenschlaf"¹. If there was some Prince(ss) Charming willing to work through all of the thorny problems², then a final kiss would open up a magical kingdom of totally new scientific possibilities.

The situation has changed. There are many massive surveys which are ongoing like *OGLE* [8], the *Catalina Sky Survey* [4], *Pan-STARRS* [6], and others that are in need of followups. The Las Cumbres Observatory's global network of telescopes (LCOGTN; [3]) is pushing the boundaries of fully homogenous telescope networks – a slightly different but also very similar challenge. The number of potential "minor" and "major" players has increased, not decreased. More ominous is the perspective that massive surveys like that of the *Large Synoptic Survey Telescope* will produce

¹The name "Dornröschen" or "little thorny rose" was translated into "Sleeping Beauty"

²In the original fairy tale, there are grisly descriptions of princely corpses hanging in the rose bush thorns – no sweat, no glory!

millions of events nightly from which hundreds or thousands need to be immediately looked at and for which there is, at present, no guaranteed follow-up system. Thus, there is a distinct need to create and exercise a homogenous telescope network *now* so that we will be able to handle the flood of data which will soon be spilled onto our astronomical doorstep.

3 The Problems

If one assumes that there is a widespread principle interest in providing/sharing resources without the intervention of too cumbersome institutional politics and that there exists a software package which makes it relatively easy to participate in a scientific collaboration using such resources is available, what would there be to prevent or even stop an HTN effort?

The worst show-stopper would be, of course, a lack of scientific benefit : if the scientific need for a HTN isn't large enough, then no one will want to create or maintain such a complicated system, even if it would be fun to try out. Fortunately, the science cases are clear and compelling: it should be obvious that some (but not all) kinds of interesting science would be made possible by an HTN.

Another problem is the heterogeneity of an HTN: unlike the LCOGTN, which has the goal of making the question which telescope actually makes an observation irrelevant, a truly heterogeneous telescope network will consist of resources with very different apertures, different types of instrumentation (cameras, spectrographs), filter systems, fields of view, local weather conditions, elevations, latitudes, and longitudes. The quality of the data coming from each resource will necessarily be highly variable, and the HTN client must be prepared to decide whether the data from a server at a particular time for a particular project is worth asking for.

On the other hand, too much relative homogeneity could also hinder the scientific output of an HTN: it may be great that you can access a range of telescopes capable of providing V-band images of similar fields of view and photometric quality, but if you need a spectrograph to do the next step of followups and your HTN doesn't have one, you may run out of the resources you actually need and not be able to continue profiting from access to a large number of telescopes.

The question of data and publication rights is subtle: some institutions might be simply willing to be acknowledged, but most would rightly like to have the option to be included in the entire process of data analysis, interpretation, and final publication. This is strictly a question of internal consortium politics and so isn't *directly* relevant to the operation and use of an HTN *per se*, but some political solutions may be easier to accept than others. Indeed, there are lots of examples of large astronomical consortia quite willing to pack dozens of consortium participants onto a paper, so a consortium should not have a problem with integrating a real – if sometimes perhaps

minor – contribution within the reasonable constraints of good scientific practice.

4 The Solution

Given the potential benefits and the acknowledged problems, a generic proposal – keeping the Prince(ss) Charming analogy – for kissing this magical idea into existence is clear:

- we need a suite of tools capable of making *any* telescope – manual, remote or robotic – principally HTN-compatible with a “minimum of effort” (whatever that means); and
- we need to encourage the creation of science-driven consortia interested in becoming the client of a HTN-connected network (either a general-purpose one or one created easily for the particular purposes of the consortium).

Neither of these goals is realistically attainable unless the generic goals can be expressed in quite concrete terms. I would like to suggest that it would realistically be possible to enable the creation and use of HTN networks if we somehow joined forces to

- create an open-source HTN-client software suite with
- an absolute minimum level of complexity (use and installation) so that observatories would be tempted to try it out;
- use the transport system already developed for VOEvent (e.g. see [9]);
- provide clear and simple examples of manual and automatic interfaces to the local resources “out of the box” – an important installation demo would, for example, be the capability of showing that the system as simply installed can be easily connected to a test network and that the test client is capable of doing something as simple and potentially immediately useful as putting up a pop-up window saying, “Your XXX consortium partners request you to perform the following observation - please press the ”YES” button if you are willing to contribute your time and effort now”;
- provide some help in interfacing with the local system (e.g. help in the simple mapping of schemata);
- help define minimum standards of client calibration/participation.

Even better would be to have an equally open-source example of an HTN-Server that potential HTN-client consortia could use as a starting point.

These are realistic goals given a minimum of participation and effort, since the payoff is potentially very great. A very good model for such a project is John Swinbank's VOEvent broker package, "Comet" (<https://github.com/jdswinbank/Comet>).

5 Conclusions

The original HTN idea was slightly ahead of its time: while the scientific benefit of being able to tap in a potentially large pool of astronomical resources was obvious, the effort required to implement a functioning system was too great relative to the on-going efforts of creating and maintaining increasingly automatic observatory systems operated outside of a heterogeneous network. Now, the potential participants have much more experience operating their hardware and software and many new players have appeared. The age of massive all-sky surveys has already started but our capacity to follow up astronomical events covering a wide range of scientific fields has not and will not keep up unless there is a paradigm shift in the way we utilize our telescopic resources. An open-source telescope eMarket package which is so observatory-friendly that there are few excuses left for *not* participating in an HTN would be the sociological game-changer needed.

References

- [1] A. Allan, F. Hessman, Bischoff, K., Burgdorf, M., Cavanagh, B., Christian, D., Clay, N., Dickens, R., Economou, F., Fadavi, M., Fraser, S., Granzer, T., Grosvenor, S., Jeness, T., Koratkar, A., Lehner, M., Mottram, C., Naylor, T., Saunders, E., Solomos, N., Steele, I., Tuparev, G., Vestrand, T., White, R., Yost, S., AN 327, 744 (2006).
- [2] T. M. Brown, et al., PASP 125, 1031
- [3] A. J. Drake, S. G. Djorgovski, A. Mahabal, E. Beshore, S. Larson, M. J. Graham, R. Williams, E. Christensen, M. Catalan, A. Boattini, A. Gibbs, R. Hill, R. Kowalski, ApJ 696, 870 (2009).
- [4] F. V. Hessman, AN 327, 751 (2006)
- [5] Z. Ivezic, J. Tyson, R. Allsman, J. Andrew, R. Angel, arXiv 0805.2366 (2008)
- [6] N. Kaiser, W. Burgett, K. Chambers, L. Denneau, J. Heasley, R. Jedicke, E. Magnier, J. Morgan, P. Onaka, J. Tonry, SPIE 7733, 12 (2010)
- [7] C. Pennypacker, M. Boer, R. Denny, F. V. Hessman, J. Aymon, N. Duric, S. Gordon, D. Barnaby, G. Spear, V. Hoette, A&A 395, 727 (2002)
- [8] A. Udalski, Acta Astron. 53, 291
- [9] P. B. Warner, R. L. Seaman, R. C. Smith, AN 329, 301 (2008)
- [10] R. R. White, A. Allan, AN 329, 232 (2008).

GCN/TAN: Past, Present & Future: Serving the Transient Community's Need

*Scott Barthelmy
NASA Goddard Space Flight Center*

Abstract

GCN/TAN (Gamma-ray Coordinates Network / Transient Astronomy Network) has been serving the transient astronomy community since 1993, and will continue do so. I will discuss some of the past, more about the current, and some expectations about future capabilities.

VOEvent: Where We Are; Where We're Going

*John Swinbank
University of Amsterdam, Netherlands*

1 Introduction

This meeting, *Hot-wiring the Transient Universe III*, is explicitly focused on the opportunities presented by “rapid, coordinated, multi-wavelength follow-up observations” of transient events. As amply demonstrated by the other manuscripts in this volume, the potential benefits are manifold. However, achieving the best possible scientific return requires addressing a range of technical challenges in terms of identifying transients, classifying them, disseminating news to the community and coordinating follow-up. Here we discuss VOEvent [4], which provides the basis of a solution to many of these issues.

There already exist mechanisms by which astronomers may rapidly distribute notifications of ongoing and recent events¹. However, the next generation of large-scale survey projects by telescopes such as Gaia, LSST and SKA promise a step change in transient detection rates: we are moving from an era of a few transient alerts per day to, perhaps, tens of millions. This presents a massive scalability challenge. New ways of working are required: it is obviously impractical for even large numbers of astronomers to write, read and understand millions of event descriptions, let alone to do so quickly enough to enable scientifically relevant follow-up observations.

VOEvent defines a standardized, machine-readable way of representing information about a transient event. VOEvent is flexible enough to usefully describe a wide variety of phenomena, while being appropriate for automatic generation, transmission and parsing. In this way, the human astronomer can ultimately be removed from the loop: transient detection software generates VOEvents describing the events observed, which are rapidly shipped to interested follow-up facilities worldwide, where intelligent systems can decide whether further observations are appropriate. VOEvent then provides a mechanism for those follow-up facilities to notify the community about their observations.

¹Of particular note are *The Astronomer's Telegram*, <http://www.astronomerstelegram.org/>, and the NASA *Gamma-ray Coordinates Network*, <http://gcn.gsfc.nasa.gov/>: both have long and distinguished track records of enabling transient astronomy.

VOEvent is developed by the Time Domain Interest Group² (TDIG) of the International Virtual Observatory Alliance³ (IVOA). This manuscript gives an overview of the VOEvent system, and summarizes the relevant ongoing and future work being undertaken. The TDIG actively solicits community participation in these activities: please do not hesitate to get involved and make your requirements known.

2 Structure and content of VOEvents

VOEvent defines an XML schema which describes how transient events may be described in a structured way. The VOEvent schema builds upon other IVOA standards. Each VOEvent document (or “packet”) describes a particular alert. Specifically, it may contain information on each of the following:

- The author; that is, the entity responsible for the contents of the packet.
- The event observed. A flexible notation is used to allow for a wide range of observations to be accurately represented.
- Where and when the observations were made.
- Instrument specific information about how the data was collected.
- The scientific assessment of the event. This provides scope for the author to describe why they believe follow-up observations are merited.
- Citations to other events. These may be used to both provide supplementary information—such as the results of follow-up observations—or to supersede or retract previous events.

All of the above information is presented in such a way that it is conveniently machine-parseable. In addition, it is possible to append human-readable textual descriptions to each section of the event, and to provide references to arbitrary URIs for further details or clarification.

We emphasize that the VOEvent packet describes a transient celestial event: it does not describe or request (other than by implication) a particular follow-up action. It is up to the recipient to determine whether any action is appropriate based on their capabilities and interests.

The specification deliberately leaves higher-level functionality which may be layered on top of VOEvent undefined (beyond some references to particular roles that entities interacting with VOEvents may perform).

²<http://www.voevent.org/>

³<http://www.ivoa.net/>

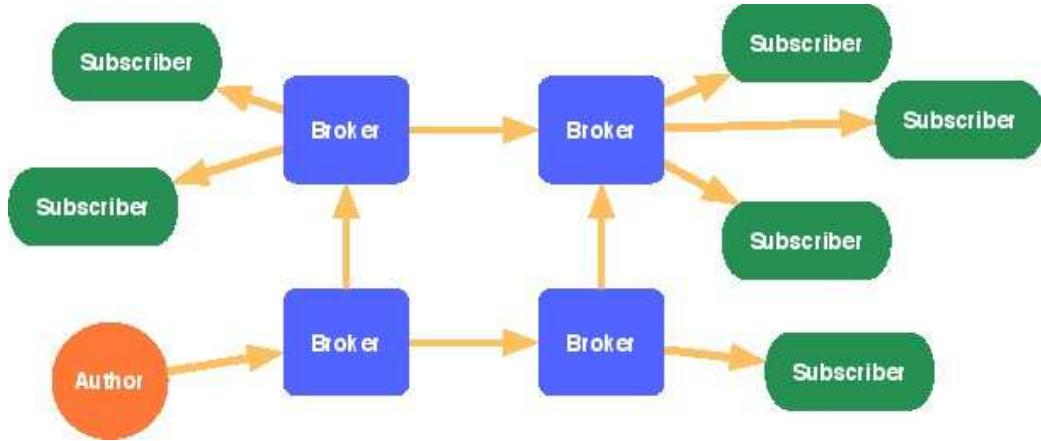


Figure 1: A schematic overview of the entities within a VTP network. The path of a VOEvent packet traversing the network is indicated by the arrows.

3 VOEvent transport

The VOEvent specification deliberately does not mandate a particular method of shipping events from their origin to interested recipients: rather, the end user is encouraged to choose a method well suited to their requirements. With time, though, it has come to be recognized that providing some baseline specification would be of significant benefit. This both provides a natural starting point for new users, and enables the construction of community wide infrastructure [6].

A basic design for a TCP-based protocol for distributing VOEvents was produced as an IVOA Note some years ago: this is the “VOEvent Transport Protocol”, or VTP [1]. As a Note, this protocol has not undergone the formal IVOA standardization process; rather, it simply represents the opinions of the authors. However, the protocol has the merits of simplicity, usefulness and ease of implementation. As such, there is now an ongoing effort to update, clarify and formalize the document, so that it can move towards full standardization over the coming months.

An overview of VTP is shown in Fig. 1. In brief, the system defines three independent roles on the network: that of “author”, “broker” and “subscriber”. The subscriber is the end user or facility which wants to receive a stream of VOEvents. They open a connection to a broker, which acts as a distribution hub. This connection is kept open continuously. The author is the individual or facility which writes VOEvents. When they have a packet to distribute, they connect to the broker and upload it; the broker then redistributes it to all connected subscribers.

One broker may subscribe to another. When the upstream broker receives an event, it is distributed to the downstream broker (along with the other subscribers), which, in turn, passes it along to its own subscribers. In this way, the protocol

enables the construction of extended networks of mutually-interconnected brokers. An author publishes to one broker, and the event is distributed to all subscribers across the network. Such a network is robust: a faulty broker can only disrupt traffic involving local authors and subscribers, rather than causing problems across the whole network.

VTP makes it convenient to build “added value” services atop the basic infrastructure as described. For example, a broker might perform server-side filtering on behalf of clients, only forwarding to them events which match criteria they have specified. Alternatively, a node on the network can record all the events it receives, making them available to clients via, for example, a Web-based interface. SkyAlert⁴ provides an excellent example of the possibilities.

Open-source software is available which can perform all of the roles within the VTP network⁵, and events from a number of sources are regularly being distributed. This infrastructure has already been used for published science [5], and it is hoped that this usage will grow as VTP heads towards standardization.

4 Future developments

The existing specification makes science with VOEvent possible today. However, the TDIG continues to evolve and enhance both the VOEvent definition itself and the surrounding infrastructure.

4.1 IVOA registry extension

The Registry acts as a directory of IVOA services available to the end user, be that user human or machine. A supplement to the existing registry specification is currently under development which will enable the registry to be used to describe facilities relevant to the publication, discovery and use of VOEvents: this is the “VOEventRegExt”. This will be intentionally generic: not limited to describing only the simple distribution model described in §3 above but rather able to represent a wide range of event handling services.

The VOEventRegExt standard is currently at an advanced draft stage; it is anticipated that it will move towards standardization over the coming months.

⁴<http://www.skyalert.org/>

⁵*Comet*, <http://comet.transientskp.org/>, provides a Python-based implementation, while the *Dakota VOEvent Tools*, <http://voevent.dc3.com/>, are written in C#.

4.2 Security

The term “security” when applied to VOEvent infrastructure is, perhaps, overloaded. It can refer either to *secrecy*—the idea that certain events should be available only to authorized recipients—or to *integrity*—a guarantee that a particular packet was genuinely produced by its claimed author. Within the context of the current document, we discuss only the latter: the former can generally be addressed by specialist transport mechanisms between authorized parties.

Mechanisms for ensuring the integrity of a particular event can take one of two forms: either the event is exchanged over an authenticated channel or a cryptographic signature is applied to the event itself. The former method is not a good match to the distributed nature of the network described in §3. Discussion has therefore focused on the latter.

Two proposals have been made to introduce cryptographic signatures to VOEvent [2, 3]. Both rely on standard public key cryptography algorithms, wherein the author signs the event packet in such a way that the recipient can verify their identity. To date, neither of the proposed systems has gained widespread acceptance. Broadly speaking, there are two reasons for this. One is that existing event networks are low traffic and, perhaps, low prestige: there is little motivation to compromise them. Secondly, the technical considerations are still under debate. In particular, the relationship between the set of bits which constitute a given VOEvent packet and the information contained within is not trivial. It is possible, and may in some circumstances be desirable, to mutate the structure of an event while leaving its signature intact, a process which is complex, and, under some proposed systems, impossible.

The first reason will become less relevant with time: as major facilities start publishing VOEvents, they will need to protect their reputation against forgeries, and as automatic response becomes increasingly commonplace, the potential consequences of false events become ever greater. For this reason, although there is no clear standard to adopt yet, it is anticipated that security concerns will be the focus of much of the effort around VOEvent in the future.

4.3 Bulk transportation and event containers

The transport protocol described in §3 is intentionally minimalist: it provides a basic level of functionality without attempting to address every possible use case. In particular, it might not be appropriate for transporting the many millions of events which are forecast by next-generation projects. Further, the basic VOEvent standard provides no capability for including supporting information such as cut-out images.

To address these use cases, a proposal has been made for a means of bundling a number of events, together with associated files, into a single package for convenient transmission. This concept, tentatively referred to as “VOEventContainer”, will be

the topic of future work within the TDIG.

4.4 Alternative serialization

The representation of events as XML documents builds on widely accepted standards and makes it possible to process them with a variety of off-the-shelf tools. However, XML is verbose and can be awkward to work with. Furthermore, per §4.2, the very flexibility offered by XML makes cryptographically signing events complex. Perhaps unsurprisingly, the request for an alternative representation of VOEvent is often repeated. This would involve defining a format which contains the same information as the current standard, but which avoids the disadvantages of XML. JSON⁶ is regularly cited in this context.

The TDIG is alert to the call for an alternative event representation, and welcomes proposals. However, no concrete development is currently underway in this area.

5 Conclusion

VOEvent provides a crucial piece of the infrastructure required to effectively respond to transients. It is already a mature and proven technology, and has played a key role in published science. It is anticipated that its importance will grow with the increasing data volume and consequent automation associated with current and future transient searches. The TDIG will continue to develop VOEvent to meet the challenges of the next generation: we invite the community to actively participate in this effort.

⁶<http://www.json.org/>

References

- [1] A. Allan and R.B. Denny, *VOEvent Transport Protocol, Version 2.1*, IVOA Note, 5 August 2009.
- [2] S.L. Allen, Astron. Nachr. 329, 298–300, 2008.
- [3] R.B. Denny, *A Proposal for Digital Signatures in VOEvent Messages, Version 1.10*, IVOA Note, 14 May 2008.
- [4] R. Seaman et al., *Sky Event Reporting Metadata, Version 2.0*, IVOA Recommendation, 11 July 2011.
- [5] T. Staley et al., MNRAS 428, 3114–3120, 2013.
- [6] R.D. Williams et al., Proc. SPIE 8448, id. 84480R, 2012.

Time Series Data Visualization in World Wide Telescope

Jonathan Fay
Microsoft Research

Abstract

WorldWide Telescope provides a rich set of timer series visualization for both archival and real time data. WWT consists of both interactive desktop tools for interactive immersive visualization and HTML5 web based controls that can be utilized in customized web pages. WWT supports a range of display options including full dome, power walls, stereo and virtual reality headsets.

RTS2 and BB – Network Observations

Petr Kubanek

Institute of Physics, Academy of Sciences of the Czech Republic

Abstract

On Hotwired I, I presented idea of observations running on RTS2. I would like to review the progress achieved to this goal, and provide an update on the current status. Along the path lies lot of changes in RTS2 structure, which I would like to present and explain.

Multi-Telescope Observing: the LCOGT Network Scheduler

*Eric S. Saunders
Las Cumbres Observatory Global Telescope*

*Sotiria Lampoudi
Liquid Robotics*

1 Introduction

Las Cumbres Observatory Global Telescope (LCOGT) is developing a worldwide network of fully robotic optical telescopes dedicated to time-domain astronomy. Observatory automation, longitudinal spacing of the sites, and a centralised network scheduler enable a range of observing modes impossible with traditional manual observing from a single location. These include continuous coverage of targets across sites, simultaneous observing with multiple resources, and cadenced time-series monitoring without diurnal gaps. The network also provides resource redundancy, with the potential for observations to be rescheduled in response to changing weather conditions. The scheduling model supports a wide variety of observing programs, which typically have very different constraints, goals, contingencies and timescales.

Heterogeneous requests in a networked environment present specific, unusual challenges for telescope scheduling that do not arise with single-resource schedulers. Here, we present a short overview of the LCOGT network scheduler. We describe our design goals, outline the scheduler's modular architecture, and highlight its current and planned capabilities.

2 Design goals

Las Cumbres Observatory Global Telescope (LCOGT) is a non-profit foundation dedicated to ground-based optical time-domain astronomy. Two 2m class and nine 1m class robotic telescopes, deployed across five sites and controlled and coordinated from a central headquarters in Santa Barbara, California, provide an automated observing network for a wide range of time-domain science (Fig. 1) [1].

Each telescope is fully robotic. All hardware is automated, from the weather sensors and dome to the telescope mount, instrumentation, and controlling computers.



Figure 1: The 2x2m and 9x1m telescopes making up the deployed LCOGT network, as of January 2014.

This level of automation is the enabler for the key feature of the network - its global scheduling. The control metaphor for the LCOGT network is that of a single global instrument, capable of observing an unusually large area of sky, and, for southern targets, almost always in darkness, so unfettered by the diurnal cycle. The network is homogeneous; within each telescope class, the telescopes are exact clones, with identical build, instrumentation and filters. Both spectra and CCD imaging may be performed from any site.

Most sites have more than one telescope present. This is desirable from both a cost and resource perspective, making maintenance easier and allowing the telescopes joint access to a shared site spectrograph. Multiple telescopes also provide important redundancy, increasing throughput (several observations can be concurrently made from the site), while mitigating against technical failures of individual resources.

The network is controlled and operated from a central headquarters, through a distributed “hub and spoke” software architecture. The software system is made up of a number of layers of increasing abstraction, from the low-level code that drives the hardware, through the sequencer for each telescope, and up to the agents that mediate access to each site [2]. Within the hub, the central databases that track the status of observing requests are interrogated and updated by the network scheduler in response to changing conditions and new observing requests, which are created by most users through the LCOGT website, and in some cases are submitted programmatically.

The modular architecture allows components to be tested or overridden in isolation, and also creates tolerance to network outages through the use of schedule caching at each site.

The purpose of the network is to support multiple distinct types of time domain science, across a large dynamic range of characteristic timescales. Many of these programs have competing requirements. Long-term continuous periodic monitoring, for example, is at odds with short timescale but highly urgent and potentially disruptive requests. Such requests may show up with very little warning. For example, transient survey programs generate previously unknown targets which require prompt follow-up. Thus an important design goal was that the network be highly responsive to new input, prioritise effectively between the many types of science, and allow requests to be made throughout an operational semester.

The other major goal was to fully utilise the potential arising from having a robotic network. Observations should be dynamically reschedulable in response to changing conditions at the sites. Schedule optimisation should take place across all telescopes together. Unique observing modes such as synchronous observing, hand-off between telescopes (for long unbroken time series), and cadences without diurnal aliasing should be possible. Finally, the abstraction of a single instrument should be preserved, so that observations can be described in a high-level way, independent of any particular observing site.

3 System overview

Requests are the principal unit of scheduling. They are abstract, generalised observing descriptions, which provide a way to ask for data from the network as a whole. Each request has one or more *observing windows*, which are temporal constraints indicating an outer bound within which an observation can occur. Note that an observing window may be much larger than the total duration of the observation. A window of one week, for example, means that the request may be performed at any time during that week, subject to target visibility. In general, larger windows provide more opportunities for a request to be successfully scheduled.

Requests also hold parameters such as target, autoguider, exposure and filter requirements, which are necessary to actually perform the observation. Optionally, constraints such as maximum acceptable airmass or lunar distance may also be provided.

If a request is selected for the final schedule, a *block* is created. A block is a description of an observation that has been bound to a specific telescope and time.

Figure 2 shows the flow of requests and blocks through the network. Requests enter the system through the web portal, or programmatically through the web service API. They are validated and stored in the request database. The network scheduler

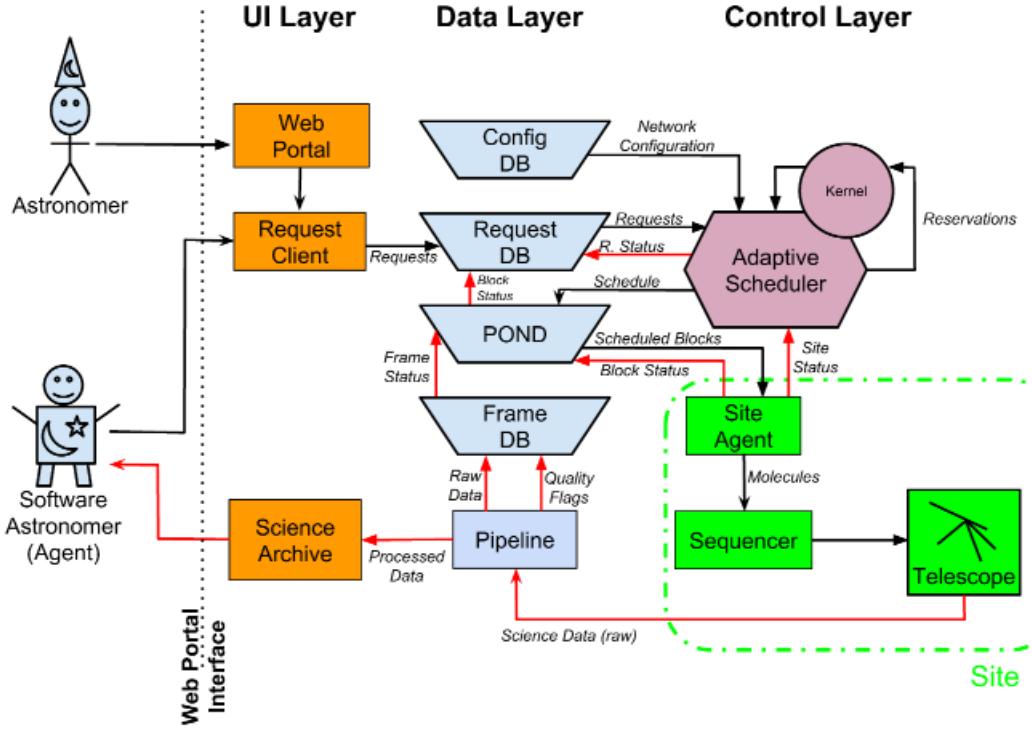


Figure 2: High-level architecture of the LCOGT observing software stack.

monitors the request database continuously. A scheduling run is triggered in response to changes in the pool of requests. A triggering change could be a new request arriving, an existing request being cancelled, or a scheduled block failing to be observed. Rescheduling is also triggered by any change in the set of available telescope resources. If a site becomes temporarily unusable as a result of cloud cover, for example, then the schedule is recalculated, and observations are reassigned among the remaining available sites.

Re-scheduling never pre-empts a block that is already executing at a telescope at scheduling time; interrupting such blocks is inefficient and causes significant “thrashing”. The only exception to this rule are requests flagged as time critical, called Target of Opportunity (ToO). These requests require immediate placement at the first available telescope. As long as ToO requests comprise only a small fraction of the total request pool, such interruptions are tolerable.

For each site, the windows of each request are reduced by the hard constraints of day and night, target visibility, and any other *a priori* constraints (e.g. airmass)

specified. Figure 3 depicts how this step reduces the range of possible times available to a request.

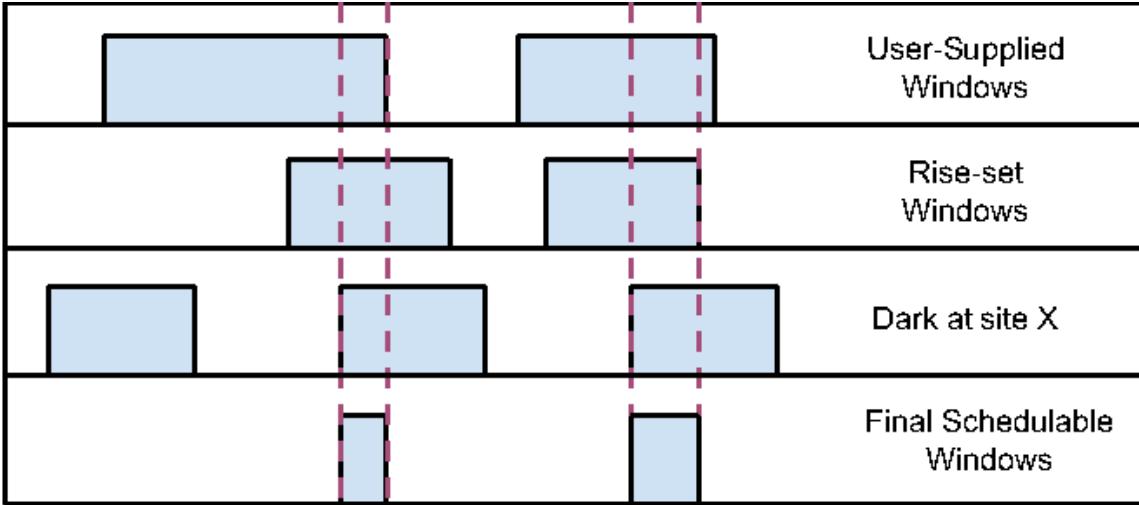


Figure 3: Intersecting request windows by hard constraints to obtain final schedulable windows.

With the windows now fully determined, the requests are mapped to an abstract formalism [3] that can be passed to the scheduling kernel. The scheduling kernel solves the well-defined mathematical problem laid out in [3] using a discretised integer linear program. The implementation of the kernel is beyond the scope of this paper; see [4] for a detailed discussion.

The computed schedule (optimised with respect to the implemented priority function) selects the requests which should be observed, and indicates where and when they should take place. Corresponding blocks are constructed, and stored in the Proposal and Observation Network Database (POND). This database is polled by the site agent at each site to determine what to observe.

Each site agent pulls several days of scheduling from the POND and stores a copy locally. This schedule caching allows the site to continue to observe for some time even in the event of a network outage between the site and the scheduler.

4 Version 1.0 capabilities

The LCOGT network begins official operations in April 2014. The central goal of the initial system is reliable operation and monitoring of the network with a small core feature set. From a scheduling perspective, it will include the following key capabilities:

- Globally optimised observation placement
- Simultaneous and cross-site observing
- Automatic weather rescheduling
- Automatic rescheduling of unsuccessful observations
- Target of opportunity observations, placed within 15 minutes of submission
- Both human and programmatic request submission interfaces
- Cadence-driven, multi-site sequences
- Support for solar-system objects
- Low-dispersion spectrographs at both 2m sites [1]

5 Future work

Many features are slated for implementation beyond version 1.0. Adding seeing and transparency constraints will allow just-in-time site selection based on current site conditions. Intra-project priorities will allow users to provide information on the relative importance of requests within their proposals. Evaluating the performance of different priority functions on simulated and historical data sets will be an ongoing process to tune the scheduling output. Automatic decomposition of large contiguous windows among sites will conveniently allow continuous observing for longer than a night. “Smart” cadences that autonomously adjust behaviour in response to missing observations are a largely unexplored area of research which promise to be powerful tools for specifying dynamic observing strategies [5]. Additionally, parts of the network may be used to monitor space junk and track satellites. Such objects present their own unique scheduling challenges.

References

- [1] Brown, T. M., Baliber, N., Bianco F. B., et al.: 2013, PASP 125, 931
- [2] Hawkins, E., Baliber, N., Bowman, M., et al.: 2010, SPIE 7737, 77370P
- [3] Lampoudi, S., Saunders, E. S.: 2013, ICORES 313
- [4] Lampoudi, S., Saunders, E. S.: 2014, in prep.
- [5] Saunders, E. S.: 2007, University of Exeter, "Optimal Observing of Astronomical Time Series Using Autonomous Agents", PhD thesis

Nuts and Bolts: Algorithms and Event Brokers

Novel Measures for Rare Transients

*Ashish Mahabal
California Institute of Technology*

Abstract

Data volumes in astronomy have been growing rapidly. Various projects and methodologies are starting to deal with this. As we cross-match and correlate datasets, the number of parameters per object—in other words dimensions we need to deal with—is also growing. This leads to more interesting issues as many values are missing, and many parameters are non-homogeneously redundant. One needs to tease apart clusters in this space which represent different physical properties, and hence phenomena. We describe measures that help to do that for transients from the Catalina Realtime Transient Survey, and project it to near future surveys. The measures are based partly on domain knowledge and are incorporated into statistical and machine learning techniques. We also describe the discriminating role of appropriate follow-up observations in near-real-time classification of transients. In particular such novel measures will help us find relatively rare transients.

The Modern Automated Astrophysics Stack

*Joshua Bloom
University of California at Berkeley*

PESSTO – The Public ESO Spectroscopic Survey for Transient Objects

Morgan Fraser

Institute of Astronomy, University of Cambridge, UK

Stefano Benetti

INAF - Osservatorio Astronomico di Padova, Italy

Cosimo Inserra

Astrophysics Research Centre, Queens University Belfast, UK

Andrea Pastorello

INAF - Osservatorio Astronomico di Padova, Italy

Stephen Smartt

Astrophysics Research Centre, Queens University Belfast, UK

Ken Smith

Astrophysics Research Centre, Queens University Belfast, UK

Mark Sullivan

School of Physics and Astronomy, University of Southampton, UK

Stefan Taubenberger

Max-Planck-Institut für Astrophysik, Garching, Germany

David Young

Astrophysics Research Centre, Queens University Belfast, UK

1 Introduction

Systematic searches for supernovae (SNe) and extra-galactic transients have historically been conducted via targeted surveys of nearby, and usually star-forming, galaxies (e.g. CHASE [1]; LOSS [2]). Such surveys have been successful in identifying core-collapse and thermonuclear SNe, luminous blue variable (LBV) outbursts, and flares from active galactic nuclei among other phenomena. However, targeted surveys of massive galaxies provide a biased sample of transients for statistical studies, and may

also miss phenomena such as transients at large offsets from their host galaxies, or which are found in low surface luminosity dwarf galaxies.

Modern surveys such as PanSTARRS [3], the Catalina Real-time Transient Survey [4], the La Silla-QUEST Survey [5] and the Palomar Transient Factory [6] have attempted to address this by surveying large areas of the sky in an unbiased and uniform fashion. However, the broad areas of sky these surveys cover each night (typically ~ 1000 s of square degrees per night, to a depth of $\sim 20 - 21$ mag) and consequent large numbers of transients found has led to a situation where the limiting factor is no longer discovery rates, but rather spectroscopic classification and followup.

To address this problem, the Public ESO Spectroscopic Survey of Transient Objects (PESSTO; PI: S.J. Smartt) is conducting a large scale spectroscopic survey for the classification and followup of SN candidates using the 3.6m ESO New Technology Telescope at La Silla [7]. The survey itself has two main goals - spectroscopic confirmation and classification of a large number of SN candidates which have been reported to the community on a rapid basis, and obtaining well-sampled and publicly available spectroscopic time series for a subset of targets which are of particular scientific interest.

2 Observations

PESSTO observes for approximately 10 nights per month, over 9 months of the year. Within each month, the 10 nights of observations are split into three separate observing runs of 2-4 nights to ensure a reasonable cadence for SN followup. 4-m class telescopes, such as the NTT, are well-suited to spectroscopic observations of targets which are between magnitude 16-20, and a low (17 \AA) resolution spectrum suitable for classification of a SN can be obtained in $\sim 20 - 60$ minutes.

PESSTO observations are conducted with both EFOSC2, the optical imager and spectrograph on the NTT, and SOFI, the NIR camera and spectrograph. All PESSTO classification spectra are taken with a fixed set-up - EFOSC2 with Gr#13. Followup spectra of scientifically interesting targets are taken with EFOSC2 using Gr#11 or Gr#16, and SOFI using either the blue or red grisms, while a small amount of time may also be spent on optical and NIR imaging of targets when they become too faint for spectroscopy. Gr#13 covers from $3685 - 9315 \text{ \AA}$ which is ideal for SN classification, while Gr#11 and Gr#16 cover $3380 - 7520 \text{ \AA}$ and $6015 - 10320 \text{ \AA}$ respectively.

The limited set of fixed observing modes used by PESSTO has several benefits: ensuring a homogenous data set, reducing the number of decisions which have to be taken by the observer, and facilitating and simplifying the rapid pipeline reduction of data. At the telescope, the observer only has to adjust the exposure time and choose the slit width to use for spectroscopy, depending on weather conditions.

3 The PESSTO Marshall and pipeline

PESSTO observations are coordinated and documented through the PESSTO Marshall (developed by D. Young). The PESSTO Marshall is a web-based application, with a MySQL backend database, containing details of targets, observations, and ancillary data such as the redshift of host galaxies, or the current magnitude of a SN as reported by the various surveys. Users can interact with the dynamic webpages to update data and comment on any of the objects contained in the Marshall.

Prospective targets which are publicly announced by surveys are automatically ingested into the Marshall. These candidates are then either selected as potentially interesting targets for a classification spectrum, or discarded. The Marshall is also used for the management of followup targets - with requests for observations being communicated to the observer at the telescope via the Marshall.

All PESSTO data are reduced using the PESSTO pipeline. This is a PYTHON-based pipeline developed by S. Valenti which uses PYRAF routines to reduce and calibrate the observations taken each night. All classifications made by PESSTO are announced within 24 hours of the end of the Chilean night, and the spectra made publicly available via WISEREP¹ [8]. To facilitate the rapid reduction of spectra and announcement of classifications, the PESSTO pipeline can be run in “rapid” mode on a single raw spectrum, with archival sensitivity curves and wavelength solutions used to calibrate the data. The rapid reduced spectra have been checked against full reductions of the same data (using calibrations from the same night), and found to have no systematic differences.

All data are fully reduced and released to the community at the end of each year via ESO². The final science-quality reduced spectra are also available via WISEREP, while in future all PESSTO data will also be available from the IA2 archive³.

4 Conclusions

The PESSTO survey has successfully implemented a model of observing which ensures homogenous, uniform quality data, with minimal effort from observers and data reducers. In large part, this is due to the adoption of a fixed set of observing modes, largely automated pipeline reduction of data, and good communication via the PESSTO Marshall. In its first year of operations, PESSTO data has been used to study over thirty SNe in detail, including the first days after explosion of a nearby Type II SN [9] for which a progenitor candidate was identified in archival Hubble Space Telescope images [10]; a H-poor core-collapse SN which exploded in a H-rich

¹<http://www.weizmann.ac.il/astrophysics/wiserep/home>

²http://www.eso.org/sci/observing/phase3/data_releases.html

³<http://ia2.oats.inaf.it/index.php/>

circumstellar environment [11], “faint and fast” H-poor SNe [12], a ~ 60 M_{\odot} star with three years of outbursts prior to its presumed demise [13], and one of the most intensively observed luminous Type Ia SNe [14].

The PESSTO survey has run for 1.5 years, and will continue for another 2.5 years (with an additional fifth year pending review by ESO). Over the full duration of the survey, PESSTO will classify some thousands of SNe, and provide detailed spectroscopic followup for ~ 100 of these. This rich, publicly available dataset will be of great value to the community in understanding the physics of SN explosions.

MF acknowledges funding through the European Union FP7 programme through ERC grant number 320360.

References

- [1] Hamuy, M., Pignata, G., Maza, J., et al., Mem. d. S. Ast. It., **83**, 388, (2012)
- [2] Leaman, J., Li, W., Chornock, R., & Filippenko, A. V., MNRAS, **412**, 1419, (2011)
- [3] Kaiser, N., Aussel, H., Burke, B. E., et al., Proc. SPIE, **4836**, 154, (2002)
- [4] Drake, A. J., Djorgovski, S. G., Mahabal, A. et al., ApJ, **696**, 870, (2009)
- [5] Hadjiyska, E., Rabinowitz, D., Baltay, C., et al., IAU Symposium, **285**, 324, (2012)
- [6] Law, N. M., Kulkarni, S. R., Dekany, R. G., et al., PASP, **121**, 1395, (2009)
- [7] Smartt, S. J., Valenti, S., Fraser, M., et al., The Messenger, **154**, 50, (2013)
- [8] Yaron, O., & Gal-Yam, A., PASP, **124**, 668, (2012)
- [9] Valenti, S., Sand, D., Pastorello, A., et al., MNRAS, **438**, L101, (2014)
- [10] Fraser, M., Maund, J. R., Smartt, S. J., et al., MNRAS, **L212**, (2013)
- [11] Inserra, C., Smartt, S. J., Scalzo, R., et al., MNRAS, **437**, L51, (2014)
- [12] Valenti, S., Yuan, F., Taubenberger, S., et al. 2014, MNRAS, **437**, 1519
- [13] Fraser, M., Inserra, C., Jerkstrand, A., et al., MNRAS, **433**, 1312, (2013)
- [14] Childress, M. J., Scalzo, R. A., Sim, S. A., et al., ApJ, **770**, 29, (2013)

Time Series Explorer

Jeff Scargle

NASA Ames Research Center / San Jose State University

1 Scope of the Time Series Explorer

A central theme of this conference is the need for advanced algorithms for automated analysis of massive amounts of time series data. To address this problem Tom Loredo and I are developing the *Time Series Explorer* – an analysis toolkit and automated pipeline. Here I discuss a few algorithms for this system, to be described in more detail elsewhere. The underlying goals are to detect and characterize periodicities, correlations, time delays, random activity, transient events, and other astronomically interesting features, and deliver these results in ways suitable for subsequent exploration, machine learning, data mining, and visualization. Application contexts range from exploratory data analysis, such as combing massive amounts of time series for unknown signals, through projects tightly targeted on specific measurements.

Desiderata for such tools include ability to handle the variety of data modes and irregular sampling characteristic of modern astronomy. Suitability for automated analysis, providing inputs to machine learning, data mining, and visualization systems, and other interface issues are less well defined but are subjects of much current research. Issues of computational efficiency will not be discussed here.

At the core of this setting is construction and processing of data structures to represent the information content of observations, measurements, or computations. One approach is this sequence, each step operating on the results of the previous one and feeding information to the subsequent one:

- 1 Represent the raw data in terms of some measure of intensity
- 2 Agglomerate and/or smooth item 1
- 3 Identify statistically significant time-domain features in item 2
- 4 Render item 3 into astrophysical meaningful features
- 5 Present item 4 in a form suited for the context

Examples of contexts in item 5 are machine learning, data mining, publication, archiving and visualization. The following sections briefly sketch a few key algorithmic ideas for the Time Series Explorer: non-parametric time-domain models, periodograms, and correlation functions. In all cases the representation in item 1 consists of assigning a data cell to each measurement [9] (including time-tags for individual photons)

2 Optimal Histograms as Time-Domain Models

We start with a problem that seems to have no connection with time series analysis, namely representation of the probability distribution of some quantity from repeated measurements. The goal is to quantify features of the distribution such as mean value, variance, or skewness. The standard approach is to count the number of times the measurements falls within each member of a set of pre-selected evenly spaced intervals. Figure 1 demonstrates the serious problem that the resulting histogram is quite dependent on the choice of these bins. Shown are results based on one principle or another for fixing the number of bins. In practice the choice is almost always arbitrary with some degree of adjustment to bring out desired sought-after features.

In fact these data are photon arrival times take from GRB 551 at
ftp://legacy.gsfc.nasa.gov/compton/data/batse/ascii_data/batse_tte/
Optimal generation of histograms and deriving light curves from photon time of arrival measurements are essentially identical problems [9]! Data-adaptive histograms (of the same data) shown in Figure 2 use the Bayesian Blocks algorithm [9]. The bins are not fixed in number or constrained to be evenly spaced. The figure shows two different cases to emphasize the slight but noticeable dependence on the value of the one parameter of the algorithm, from the prior for the number of bins. The solid line is based on fixing the false positive rate in random noise at 5%, but is very insensitive to this value. The dashed line adjusts for the presence of a signal as described in [9]; this captures the smaller pulse at 2.3 seconds at the cost of introducing a probably spurious one at around 0.7 seconds. Even if spurious the latter pulse has very little area and thus would not much affect many post-processing results. The ideas behind the histograms shown here and other attempts at codifying the choice of bins are discussed in [5] and at the companion website
http://www.astroml.org/book_figures/chapter5/fig_hist_binsize.html.

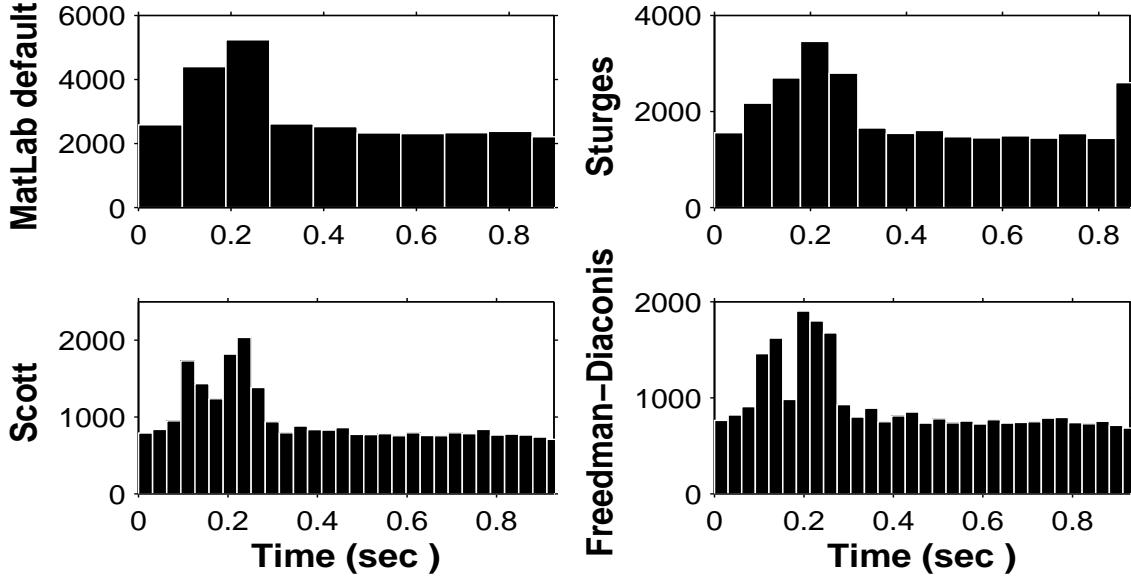


Figure 1: Histograms of 28,904 measurements with four different bin rules: (a) Matlab's default (10 bins). (b) Sturges: $1 + \log_2 N$ bins [12] (c) Scott: bins size = $\frac{3.49\sigma}{N^{1/3}}$ [11] (d) Freedman-Diaconis: bins size = $\frac{2IQR}{N^{1/3}}$; IQR is the interquartile range [4].

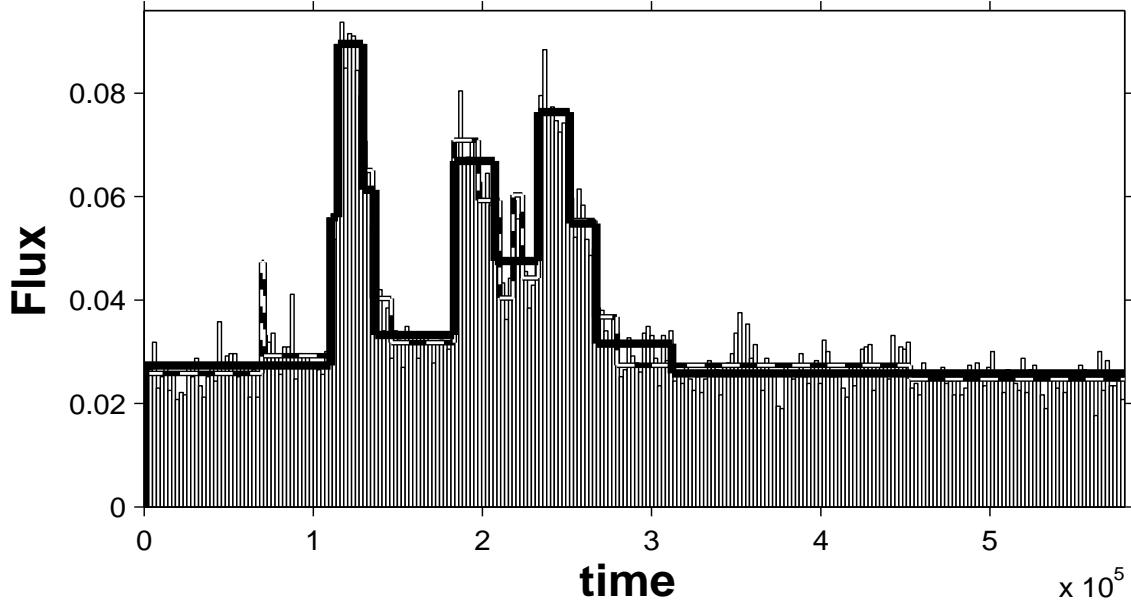


Figure 2: Bayesian Block histograms of the same data as in Fig. 1. Solid lines, and dashed sometimes hidden by the former, are with two different values of the bin-number prior parameter. Cf. the light histogram with evenly spaced bins.

3 Periodicities

In many of the contexts discussed in this conference detection of periodic signals in noise is of great importance. This topic has a huge literature. Here I only want to mention a very simple approach that does not seem to have been utilized very much in astronomy. The idea is simply to estimate the autocorrelation function using the Edelson and Krolik algorithm [2] for arbitrarily spaced data, and then apply the fast Fourier transform to yield an estimate of the power spectrum. These estimates are known as “slotted” correlation functions and power spectra in other fields (e.g. [1], apparently unaware this idea in astronomy, but making comparisons with a popular astronomical tool [6, 7, 8] for detecting periodicities). As demonstrated in [10] this approach is valuable for estimation of cross-correlation functions, cross-spectra, plus time-frequency and time-scale distributions.

As in many cases implementation of the grand idea is almost trivial but difficulties lie in details and small practical matters. While there is no pre-determined binning in time, one must establish bins in the *time lag* that need to be evenly spaced if the FFT is used to estimate the power spectrum. This fact necessitates care in bin selection and opens up the very dependence decried in Section 1. Depending on the sampling some bins can be empty of the cross-products fundamental to the Edelson and Krolik part of the algorithm; these can be handled by simple interpolation with little difficulty. Finally the power estimated with this technique is not guaranteed to be positive. This is somewhat rare, and can be ameliorated by simply taking the absolute value of the estimate. This ad hoc scheme seems to work well in simulations, but I do not know of a theoretical justification.

4 HOPing through the Time Domain

In this section we investigate the possible use for time series analysis of an algorithm developed for another domain entirely – topological analysis of density distributions. The group-finding algorithm **HOP** [3], developed to characterize the distribution of galaxies, is quite general and applies in any dimension. For any function \mathbf{f} and adjacencies defined for objects in a set \mathbf{S} , it yields a unique, parameter-free partition of \mathbf{S} into groups – one for each local maxima of \mathbf{f} , each being a connected set such that \mathbf{f} decreases monotonically away from the maximum. In short the algorithm finds all of the peaks of \mathbf{f} in \mathbf{S} and the connected structures flowing from them – mountain peaks and their watersheds. The basic idea of **HOP** is a simple hill climbing prescription, associating each object with its neighbors that have larger values of \mathbf{f} . In other words, given a set \mathbf{S} of spatially distributed objects assign a value of \mathbf{f} to each one and identify the objects (“neighbors”) adjacent to it. Then iteratively replace the index of each object with that of its neighbor with the largest value of \mathbf{f} . Rapid

convergence is obtained when no index value changes and each object is associated with a local maximum of \mathbf{f} . Now take the objects to be individual photons (instead of galaxies) described by a set of detection times t_i , and define the function for each such event as

$$f_i = \frac{2}{t_{i+1} - t_{i-1}} \quad (1)$$

This quantity is a convenient, if noisy, estimate of the intensity at the time t_i . **HOP** identifies peaks in the light curve and the photons naturally associated with them.

A plot of the intensity averaged over the groups obtained in this way is typically too noisy to be very useful. However applying **HOP** to a new time series where these output groups are treated as input objects yields a less noisy, more or less down-sampled representation. This iterative version, **IHOP**, is demonstrated in Fig. 3. This straightforward approach can obviously be applied to tasks such as pulse-hunting in gamma ray bursts and may well have further use in the same context as in Section 2.

I am grateful to Tom Loredo for his collaboration in this project.

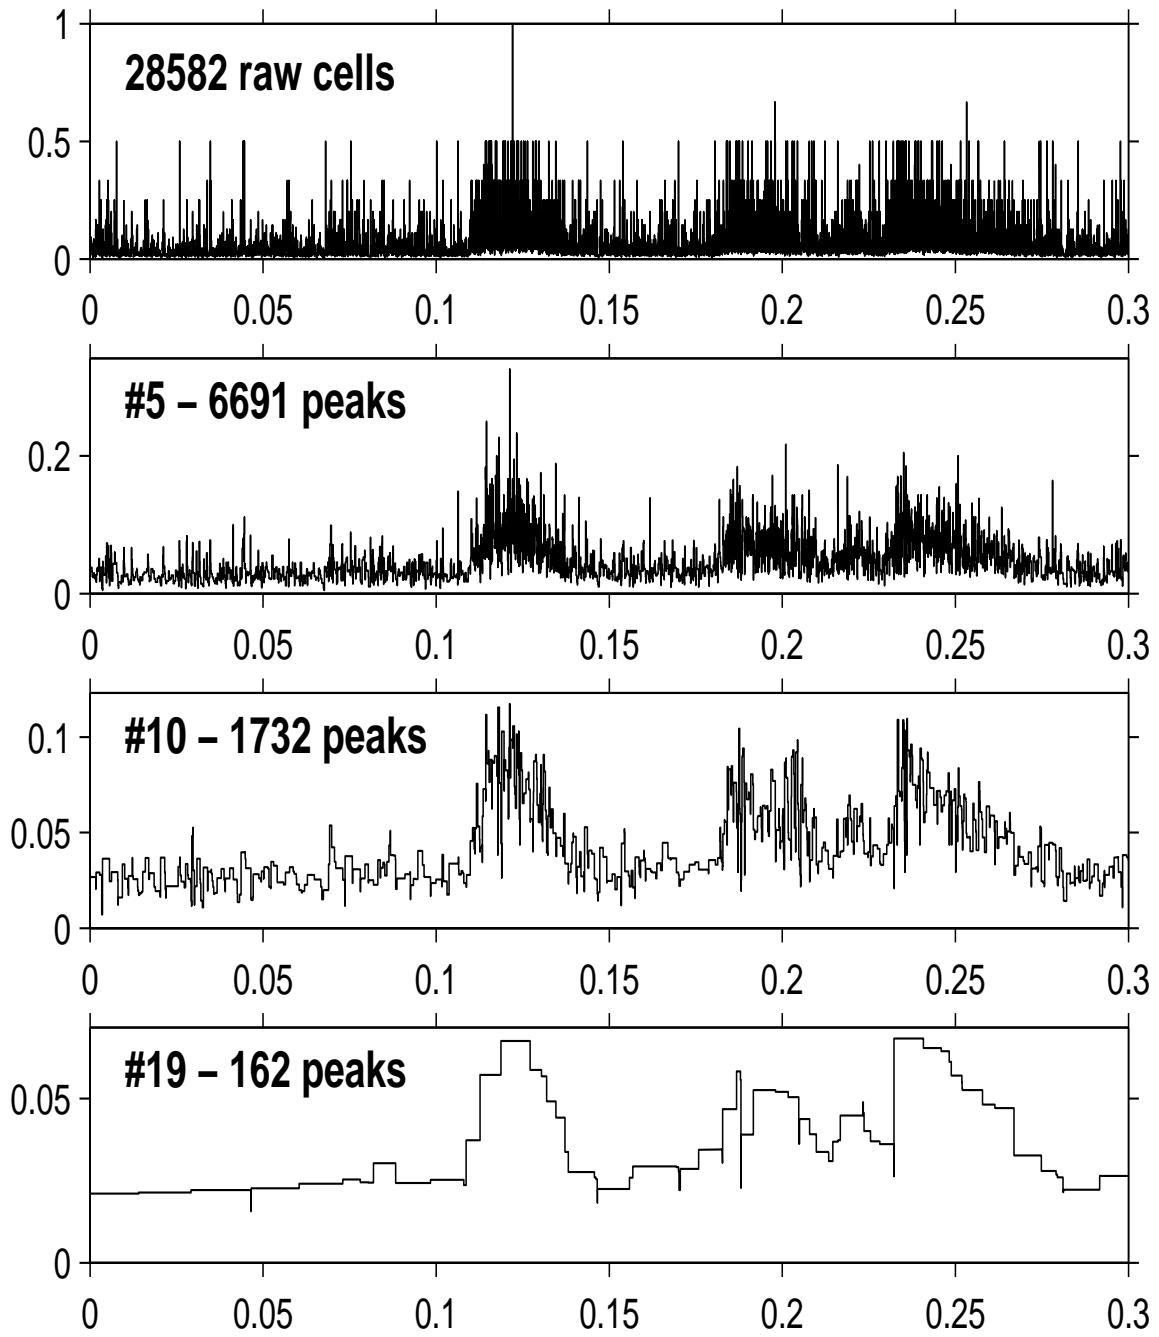


Figure 3: Raw data (top panel) represented as the value of f_i in Eq. (1) for the intervals between photons [9]. The next three panels are after 5, 10 and 19 iterations of IHOP.

References

- [1] P. M.T. Broerson, S. de Waele and R. Bos, The Accuracy of Time Series Analysis for Laser-Doppler Velocimetry, in 10th International Symposium on Application of Laser Techniques to Fluid Mechanics, Lisbon, July 10-13, 1 (2000)
- [2] R. Edelson and J. Krolik, J. H. 1988, ApJ 333, 646 (1988)
- [3] D. Eisenstein, and P. Hut, Ap.J. 498, 137 (1998)
- [4] D. Freedman, P. Diaconis, Probability Theory and Related Fields, **57**, 453 (1981)
- [5] Z. Ivezic, A. Connolly, J. VanderPlas and A. Gray, Statistics, Data Mining, and Machine Learning in Astronomy, (Princeton University Press, 2014)
- [6] N. Lomb, Ap&SS **39**, 447 (1976).
- [7] J. Scargle, J. ApJ, **263**, 835 (1982).
- [8] J. Scargle, ApJ, **343**, 874 (1989).
- [9] J. Scargle, J. Norris, B. Jackson, J. Chiang, ApJ, **764**, 167 (2013).
- [10] J. Scargle, S. Keil, and S. P. Worden, ApJ, **771**, 33 (2013).
- [11] D. W. Scott, WIREs Computational Statistics **2**, 497 (2010).
- [12] H. A. Sturges, Journal of the American Statistical Association, **65**, 66 (1926)

ANTARES: The Arizona-NOAO Temporal Analysis and Response to Events System

*Thomas Matheson, and Abhijit Saha
National Optical Astronomy Observatory*

*Richard Snodgrass, and John Kececioglu
Department of Computer Science, University of Arizona*

1 Introduction

The Arizona-NOAO Temporal Analysis and Response to Events System (ANTARES) is a joint project of the National Optical Astronomy Observatory and the Department of Computer Science at the University of Arizona. The goal is to build the software infrastructure necessary to process and filter alerts produced by time-domain surveys, with the ultimate source of such alerts being the Large Synoptic Survey Telescope (LSST). Such a tool is often called a broker [2], as it acts as the entity between producers and consumers. ANTARES will add value to alerts by annotating them with information from external sources such as previous surveys from across the electromagnetic spectrum. In addition, the temporal history of annotated alerts will provide further annotation for analysis. These alerts will go through a cascade of filters to select interesting candidates. For the prototype, ‘interesting’ is defined as the rarest or most unusual alert, but future systems will accommodate multiple filtering goals. The system is designed to be flexible, allowing users to access the stream at multiple points throughout the process, and to insert custom filters where necessary.

2 The Problem

The rapid growth of time-domain surveys produces discoveries at an ever-growing rate. Current optical surveys, such as the Lick Observatory Supernova Search¹, the Catalina Real-Time Transient Survey², the Panoramic Survey Telescope & Rapid Response System³, the Palomar Transient Factory (PTF and iPTF)⁴, and the La

¹http://astro.berkeley.edu/bait/public_html/kait.html

²<http://crts.caltech.edu/>

³<http://pan-starrs.ifa.hawaii.edu/public/>

⁴<http://ptf.caltech.edu/iptf/>

Silla-Quest Variability Survey⁵ generate transient alerts well beyond the available follow-up capacity. These projects have developed tools to filter their discoveries to focus on events of interest to each team. A good example of this is SkyAlert⁶, a system that has solved many of the astronomical issues associated with adding value to alerts. SkyAlert enables users to create filters on alerts, including ancillary information on these alerts, in order to find relevant events. The PTF system also employs tools to identify interesting alerts [1]. The scale of time-domain alert generation, though, is quickly increasing. The Zwicky Transient Facility [9] will have more than 6 times the field-of-view of PTF, while time domain surveys with DECam on the Blanco telescope benefit not only from the 3 deg^2 field-of-view, but the depth attainable with a 4m-class facility. Moreover, transients are generated across the electromagnetic spectrum, from radio facilities such as LOFAR⁷ to high-energy space-based observatories such as Fermi⁸, making the overall problem that much more complex.

On the horizon is LSST [8]. With its 10 deg^2 field-of-view and $\sim 6\text{m}$ collecting area, the transient detection rate leaps by orders of magnitude. LSST will detect (with 5σ significance) $10^3 - 10^4$ alerts per image, or $10^6 - 10^7$ per night. A good fraction of these will be known variable stars or moving objects [14, 5] (see also Ridgway’s contribution to these proceedings), but hidden among them will be rare and interesting objects that have relatively short lifetimes. Only with additional follow-up will these objects reveal their nature. These could range from short-lived phases of stellar evolution such as the final helium flash [6, 7] to superluminous supernovae [3] to electromagnetic counterparts of LIGO detections [15, 12]. Beyond these rare, but known or predicted, objects lies the great discovery space that awaits LSST. The superluminous supernovae were essentially unknown fifteen years ago and the discovery of dark energy was certainly surprising. Over its life, LSST will generate more than a billion alerts and some will be completely unknown and unanticipated objects. Without the ability to rapidly sort through millions of alerts each night and winnow them down to a reasonable number that can be studied in detail, we will lose these rare and potentially extraordinarily interesting objects. The astronomical community is becoming more aware of the necessity of such a tool [10].

3 ANTARES

The knowledge we have about an alert, such as brightness, change in flux, Galactic coordinates, ecliptic coordinates, distance to nearest galaxy, etc., constitute features that can probabilistically characterize alerts. We emphasize that this is a broad

⁵<http://hep.yale.edu/lasillaquest>

⁶<http://skyalert.org/>

⁷<http://www.transientskp.org/>

⁸<http://fermi.gsfc.nasa.gov/>

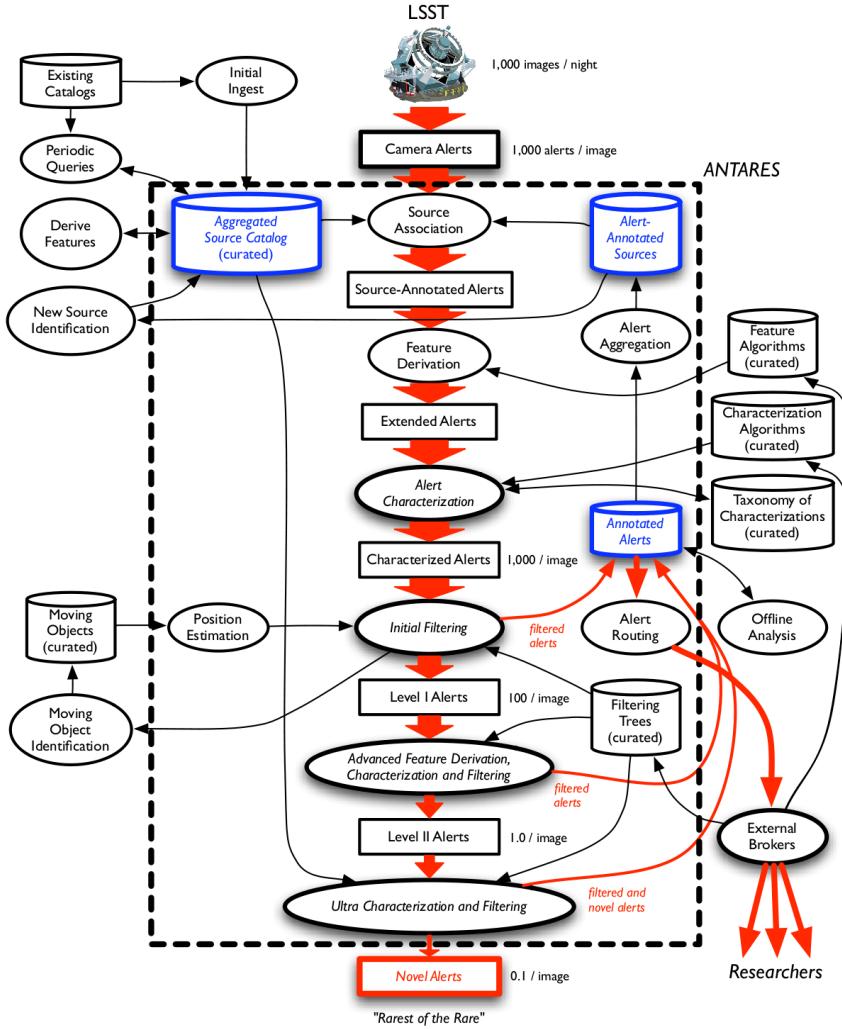


Figure 1: Basic architecture of the ANTARES system. The dashed box encompasses the processes that must keep up with the LSST frame-rate.

characterization, not a specific classification. Classification will have to come from software systems further downstream. Because of the time-scale of LSST exposures, with a new image every ~ 37 seconds, alerts must be processed rapidly to keep up with the data stream. Classification often requires more complex analysis and usually a more complete light curve [13, 4].

Figure 1 illustrates the main components of the ANTARES architecture. The overall design principles are open source and open access. The software will be available for anyone to implement and our implementation will be community driven. The

alert stream can be tapped at many points throughout the system.

The first stage is annotation that adds value to the alerts. Source association is a critical step to incorporate relevant astronomical knowledge for each alert. Catalogs of astronomical information, as well as the LSST source catalog will be the basis for this source association. Examples include the 2MASS All-Sky Data Release⁹, the Chandra Source Catalog¹⁰, the NRAO VLA Sky Survey¹¹, the Sloan Digital Sky Survey¹², the NASA Extragalactic Database¹³, and GAIA¹⁴, among many others. Even the proximity to known sources can provide useful constraints. In addition, the history of flux measurements at the position, such as a light curve, will be valuable annotation. An efficient database that can be updated regularly is an essential element of the system. This will be a valuable astronomical resource on its own. As mentioned before, the SkyAlert system provides a similar annotation. The problem for the future is the scale of alerts and the resulting necessity of this efficient database being integrated into the system brokering alerts.

For many alerts, there will only be a small number of features available for characterization, especially for an initial detection. If there are not enough features for discrimination by filtering, we can apply a probabilistic expectation of variability based on position on the sky and known distributions of variability [14]. For a position, we can construct a variability probability density function and predict the likelihood of the alert as observed. With more data, more features become available and more complex filtering algorithms can be used.

ANTARES will then use multiple layers of filters to sort the alerts and find the rarest or most interesting among them (this is the focus of the prototype project). The other alerts are not discarded. Rather, they are diverted from the main filtering stream but are still accessible to other filtering systems, including, potentially, copies of the ANTARES system itself that are tuned to specific goals. In this way, custom filters can be applied, allowing users to isolate exactly which of the alerts is of interest to them and thus address many different goals. These community-derived filtering algorithms will be applied in a multi-step process, allowing for better management of computational resources. By characterizing the alerts, the number of dimensions of feature space can be reduced. More complex filters can be applied to the smaller number of alerts after initial filtering stages.

The Arizona Machine-Experimentation Laboratory (AMELIE, Figure 2), provides a system for constructing and testing structural-causal models [11]. This essentially automates the scientific process and allows us to run experiments to test relationships

⁹<http://www.ipac.caltech.edu/2mass/releases/allsky/>

¹⁰<http://cxc.harvard.edu/csc/index.html>

¹¹<http://www.cv.nrao.edu/nvss/>

¹²<http://www.sdss.org/>

¹³<http://ned.ipac.caltech.edu/>

¹⁴<http://sci.esa.int/science-e/www/area/index.cfm?fareaid=26>

among features, including relationships that have not yet been apparent. It can observe the operation of ANTARES and make it more efficient.

The goal for the prototype is to distinguish rare and unusual objects. Once it is operational, the next stage is to expand the scope to allow users to find any type of alert of interest to them. In principle, there could be many stages of the ANTARES system itself, processing different data streams over different time scales. The overall alert ecosystem could accommodate multiple alert input streams and thus find a general way to serve the astronomical community's needs.

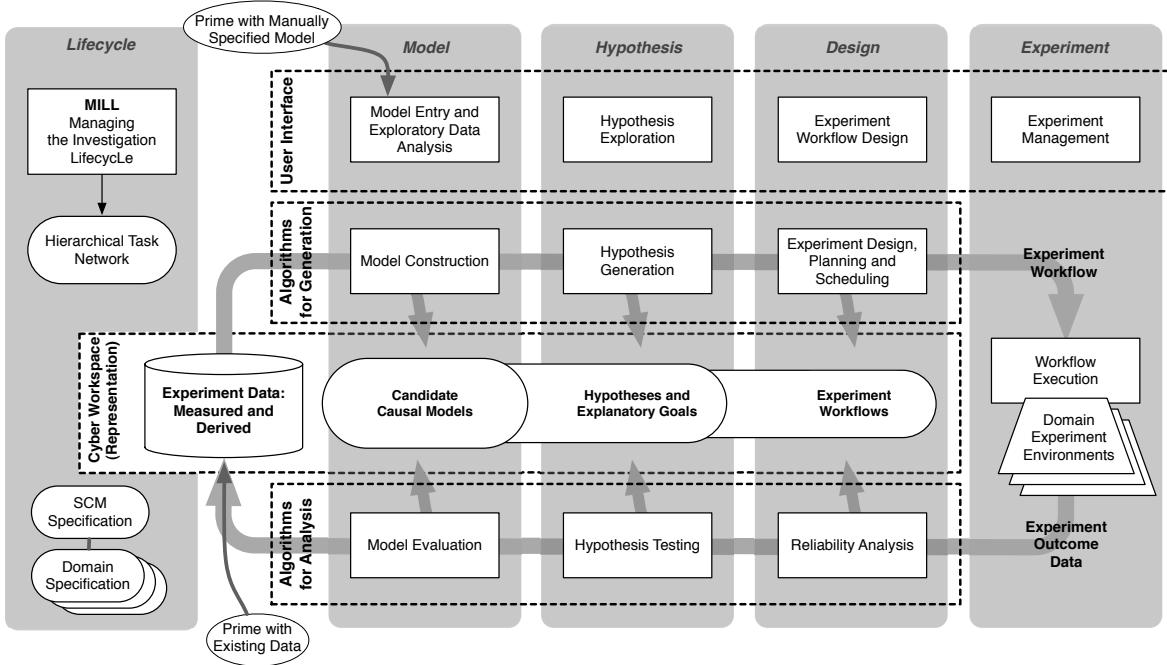


Figure 2: Basic architecture of AMELIE.

The authors are happy to acknowledge the NSF INSPIRE grant (CISE AST-1344024, PI: Snodgrass) that supports this work.

References

- [1] Bloom, J. S., et al. 2012, PASP, 124, 1175
- [2] Bourne, K. D., 2008, AstrN 329, 255
- [3] Gal-Yam, A. 2012, Science, 337, 927
- [4] Graham, M. J., et al. 2013, MNRAS, 434, 3423
- [5] Grav, T., et al. 2011, PASP, 123, 423
- [6] Herwig, F. 2005, ARAA, 43, 435
- [7] Iben, I., et al. 1983, ApJ, 264, 605
- [8] Krabbendam, V. L. & Sweeney, D. 2010, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 7733
- [9] Kulkarni, S. 2012, astro-ph/1202.2381
- [10] Matheson, T., et al. 2013, astro-ph/1311.2496
- [11] Morrison, C. & Snodgrass, R. T. 2011, Communications of the ACM, 36
- [12] Nissanke, S., Kasliwal, M., & Georgieva, A. 2013, ApJ, 767 124
- [13] Richards, J. W., et al. 2011, ApJ, 733, 10
- [14] Ridgway, S., et al. 2014, ApJ, submitted
- [15] Sigg, D., & LIGO Scientific Collaboration 2008, Classical and Quantum Gravity 25, 114041

Bayesian Time-Series Selection of AGN Using Multi-Band Difference-Imaging in the Pan-STARRS1 Medium-Deep Survey

Sidharth Kumar

Department of Astronomy, University of Maryland, College Park

1 Introduction

With the advent of the LSST era, the requirement of automated object identification in large volumes of data in existing catalogs, as well as in real-time data, is necessitated. With the knowledge of prior event types in telescopic surveys, it is possible to look for specific events in the data with a high degree of completeness and efficiency using variability [2, 3, 4], color based selection [3], and host-galaxy properties [5]. Although, a consummate and complete method would warrant the use of all these parameters in conjunction, time-series methods by themselves can contribute significantly to the selection process. Many time-series methods have been applied in the past to the identification of a broad spectrum of objects, as well as specific types; [4] discuss the identification of AGN via damped-random walk parameterization of difference-imaging light-curves, [6] on the applicability of single and multiple Ornstein-Uhlenbeck processes (hereafter OU process) to the identification of AGN, [3] on the separation of AGN from variable stars in photometric surveys through damped-random walk parameterization, and [7] on the photometric identification of specific supernovae (SNe) types. Particularly ubiquitous is the application of robust Bayesian methods [2] to the selection of sources using analytical deterministic and stochastic models for the light-curves. However, the applicability of these class of methods have been limited to single-band detections [4], or predominantly limited to using magnitude time-series data [3]. For the first time, we present multi-band difference-image selection of AGN and SNe, in the Pan-STARRS1 medium-deep fields in the g, r, i and z bands. Using Bayesian analysis, we estimate the likelihoods of a diverse range of SN models as compared to the that of the OU process. We then combine the model comparisons filter-wise using a K-means clustering algorithm [12] to provide a robust classification in each filter. The classifications are then combined across the g, r, i and z filters, to give a final classification including two measures to estimate the quality of the classification (§3).

The Pan-STARRS1 survey [9] has two operating modes, 1. The 3π survey which covers 3π square degrees at $\delta > -30$ degrees in 5 bands with a cadence of 2 observations per filter in a 6 month period, 2. Deeper multi-epoch images of 7 square degree fields in 5 bands, the so-called Medium Deep Field (MDF) Survey, for both extensive temporal coverage and depth. Depending on the weather, the accessible fields are observed with a staggered 3-day cadence in each of bands during dark and gray time (g_{P1}, r_{P1} on the first day, i_{P1} on the second day, z_{P1} on the third day, and then repeat with g_{P1}, r_{P1}), and in the y_{P1} band during bright time. On average, the cadence is 6 detections per filter in a 1 month period in g_{P1}, r_{P1}, i_{P1} , and z_{P1} , with a 1 week gap during bright time when the MDFs are exclusively observed in y_{P1} . While the 3π may detect millions of sources, in our studies we will exclusively use sources detected in the MDFs ($\approx 10^4$) since source classification entails dense time-series.

An exhaustive list of Pan-STARRS alerts are available in an online alerts database located in Harvard [9]. To derive the list of extragalactic transients, we cross-matched 18058 detected in the first 2.5 years of the Pan-STARRS1 medium-deep survey to within 3" of host galaxies detected in the deep-stack star-galaxy catalogs (Heinis 2014, In preparation), resulting in 8565 distinct extragalactic transients. These transients can be categorized broadly into stochastically varying, like AGN, or explosive-transient, like SNe. In the next section we discuss time-series models which can be used to assess their variability, and hence their categorization.

2 Time-Series Models

We assess the general shapes of the light-curves by comparing their similarities to SN-like bursting behavior, or with damped-random-walk type behavior like those of AGN. While an attempt to an exact SN or AGN fit is more tedious, the general shape of a SN light-curve could be approximated to certain analytical functional forms (Gaussian, Gamma, and generic analytic SN model (Analytic-SN)), and that of an AGN light-curve approximated by an Ornstein Uhlenbeck process [1] (OU process).

Model	Type	Equation
Gaussian	SN	$Flux(t) = \alpha + \beta e^{-(t-\mu)^2/\sigma^2}$
Gamma distribution	SN	$Flux(t) = \alpha + \beta \frac{(t-\mu)^{k-1} e^{-(t-\mu)/\theta}}{\theta^k \Gamma(k)}$
Analytic SN model	SN	$Flux(t) = \alpha + \beta \frac{e^{-(t-t_0)/t_{fall}}}{1+e^{-(t-t_0)/t_{rise}}}$
OU process	AGN	$dz(t) = -\frac{1}{\tau}z(t)dt + c^{1/2}N(t; 0, dt)$
No-Model	ALL	$Flux(t) = \frac{\sum_i(1/\delta_i^2)y_i}{\sum_i(1/\delta_i^2)}$

Table 1: Difference flux models used in the characterization of AGN and SN.

The models are described in Table 1. For SNe, although the asymmetry in the rising and falling limbs of the light-curve, cosmological redshift corrections, and extinction are not factored in, the attempt is to classify the source as a burst, or as stochastically varying, thereby not necessitating fitting specific inflections in the light-curves. Note, that since the models are compared with each other, only their relative fitnesses are important. Also, should the necessity arise of classifying the objects into particular sub-types of the major classifications, or that of extracting particular details about the parameters of the SNe, exact models [7] must be included in the comparisons, which although is beyond the scope of this paper, is a likely extension.

To assess the aptness of the models, we derive both the corrected Akaike information criterion (AICC) [13] and the leave one out cross-validation likelihood (LOOCV) [2].

$$AICC = 2k - 2\log\mathcal{L} + \frac{2k(k+1)}{n-k-1} \quad (1)$$

The AICC Eq.(1) measures the over-parameterization of a dataset by a model by correcting the maximum likelihood \mathcal{L} with the number of parameters k , as well as the finite size of the dataset n as compared to k . Although the *AICC* by itself is a good indicator of model fitness in the event that the data is representative of the general distribution of the source time-series, it may not take into account the variations in likelihood resulting from noisy data, thereby misrepresenting the actual goodness of fit. The LOOCV on the other hand is more robust to such variations which are especially common in difference-imaging. Since the AICC and the LOOCV are independent of each other, they can be used simultaneously to assess model fitness. This establishes a balance between the overall model fit via the AICC, and the robustness of the model to noise via the LOOCV. In this paper we use uniform priors for the SN and AGN model parameters, thereby not biasing the models toward particular regions of parameter space. We also use a Gaussian error model for all time-series models including the OU process, since our aim is to assess how well the mean time-series of each model fits the lightcurves.

3 Classification Method

To quantify the model fits to the data, we assess the leave one out cross-validation likelihood (LOOCV) and the corrected Akaike information criterion (AICC) for each model, filter-wise for each source. The LOOCV for each model, in each filter, are evaluated using a standard Metropolis-Hastings algorithm, and the AICC is computed from the sampled maximum likelihood. Sources which are best fit by the No-model are filtered out based on the No-Model having the highest LOOCV and the lowest AICC amongst the 5 models. We then construct the relative sign vector $RV_{i,f}$ Eq.(2)

for each object, in each filter, which is a measure of how well the data is described by the SN models as compared to the AGN model.

$$RV_{i,f} = \{sgn(LOOCV_{Gauss} - LOOCV_{OU}), sgn(LOOCV_{Gamma} - LOOCV_{OU}), \\ sgn(LOOCV_{Analytic-SN} - LOOCV_{OU}), sgn(AICC_{Gauss} - AICC_{OU}), \\ sgn(AICC_{Gamma} - AICC_{OU}), sgn(AICC_{Analytic-SN} - AICC_{OU})\} \quad (2)$$

where i is the object id, f is the filter, and sgn denotes the sign function, defined to be +1 for positive values and -1 for negative values. Ideally, for an SN the above will be $RV_{SN} = \{+1 + 1 + 1 - 1 - 1 - 1\}$ since the SN models should have a larger LOOCV, and a smaller AICC as compared to the OU process, while for AGN the signs should be reversed. However, it is possible that inherent biases in the data or the model cause one or more of the models to perform consistently worse as compared to the OU process in fitting the SN light-curves, due to noisy difference imaging resulting from astrometric errors. However, we demonstrate that our method is robust to such biases, due to redundancies in the use of multiple models in multiple filters, to describe the SN light-curves.

To test our classification we chose a diverse set of examples to reflect the spectrum of photometric properties of the dataset covering the entire gamut of SNe and AGN lightcurves. As examples for AGN, we considered 255 AGN selected from the GALEX Time Domain Survey [11]. For SNe, we identified 3300 extragalactic candidates based on their offsets ($0.4'' - 1''$) from their respective galactic hosts. We derived the offset limits by fitting a bi-modal distribution to the distribution of extragalactic alerts; a Gaussian distribution for the AGN parameterized by μ_{AGN}, σ_{AGN} , and for SN by μ_{SN}, σ_{SN} . All sources with offsets greater than $\mu_{AGN} + 2\sigma_{AGN}$ were designated as SNe. From this set we selected 100 SN by eye-balling their light-curves. We then constructed the vector of $RV_{i,f}$ for all verification set sources filter-wise, and use a K-means clustering supervised-machine-learning algorithm [12] that partitions the source vector into two distributions corresponding to SN and AGN. The algorithm obtains the centers of the two distributions or clusters by attempting to minimize the sum of squares of the distances of points x_j within each distribution or cluster S_i from the mean of the cluster μ_i .

$$\sum_{i=1}^k \sum_{x_j \in S_i} \|x_j - \mu_i\|^2 \quad (3)$$

Each source is then assigned a class $C_{i,f}$ as a SN, or an AGN, depending on the center μ_i it is clustered around. The squared-distance of the source point $d_{i,f}$ from the clustering center μ_i in filter f , is a measure of how reliably it is classified as the particular type, with a distance of $x_j - \mu_i = 0$ being the best, and larger distances

indicating less reliable classifications. This process is repeated for each source, in each of the g, r, i , and z bands independently. The advantage of classifying the sources filter-wise is that, a. the behavior of each source types in each filter could be very different, b. some filters may be more noisy than others thereby being less suited for classification, and c. the filters can reinforce the type of classification observed in other filters thereby leading to a more robust final classification. The final source classification C_i is decided by

$$C_i = \frac{\sum_f C_{i,f}}{4} \quad \& \quad d_i = \frac{\sum_f d_{i,f}}{N_{filters}} \quad (4)$$

where the sign of C_i indicates the type of the source (+1 for SN, -1 for AGN), and $|C_i|$ measures the quality of the classification. d_i is the average of squared-distances over all filters. Fig.1 and Fig.2 show the plots of C_i vs d_i for AGN and SNe, respectively. We classify AGN as having $C_i < -0.2$ and $d_i < 20$ or $C_i > 0.2$ and $d_i > 20$, while the SNe as having $C_i > -0.2$ and $d_i < 20$ or $C_i < -0.2$ and $d_i > 20$. A large value of d_i typically indicates that the source belongs to a class which is different from the one it is clustered around, *i.e.*, in this case it is the other class. The efficiency of our classification scheme on the training set is just over 90%, with 100% completeness.

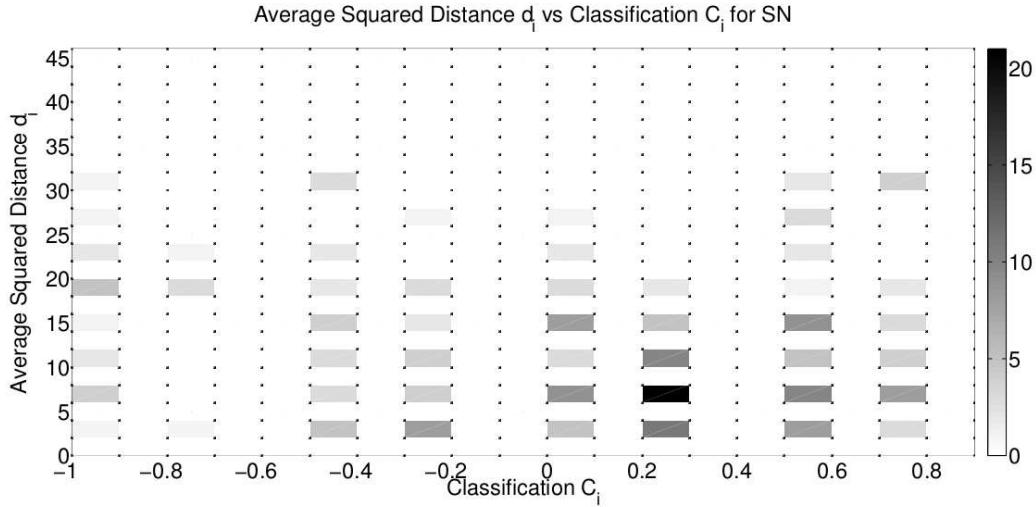


Figure 1: Distribution of 255 GALEX-TDS AGN in d_i vs C_i .

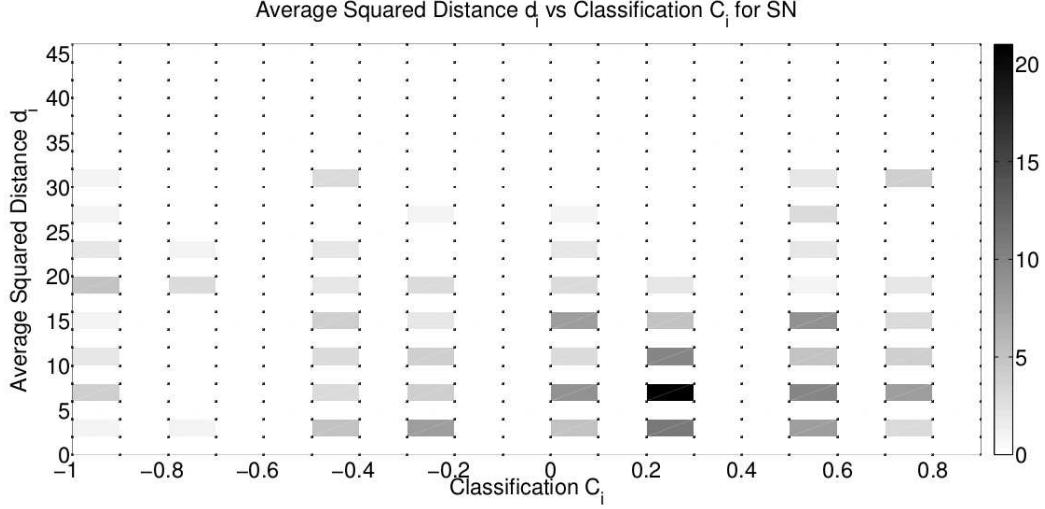


Figure 2: Distribution of 100 offset-selected SN in d_i vs C_i .

4 Conclusions and Future Work

We have discussed a Bayesian classification method to classify Pan-STARRS1 medium-deep transients identified with galactic hosts, using difference-image multi-band photometry, into SNe and AGN with 90% efficiency and 100% completeness. The methods herein can be applied to identifying AGN and SNe in existing catalogs, as well as providing real-time identification of sources in the era of Pan-STARRS2 and LSST. In addition, the method can be simply extended to the identification of particular sub-types of the broad source classes, provided their respective specific time-series models.

References

- [1] G. E. Uhlenbeck and L. S. Ornstein, Phys. Rev. **36**, 823 (1930).
- [2] C. A. L. Bailer-Jones, Astron. Astrophys. **546**, A89 (2012) [arXiv:1209.3730 [astro-ph.IM]].
- [3] N. R. Butler and J. S. Bloom, arXiv:1008.3143 [astro-ph.CO].
- [4] Y. Choi *et al.*, arXiv:1312.4957 [astro-ph.CO].
- [5] R. J. Foley and K. Mandel, Astrophys. J. **778**, 167 (2013) [arXiv:1309.2630 [astro-ph.CO]].
- [6] R. Andrae, D. -W. Kim and C. A. L. Bailer-Jones, arXiv:1304.2863 [astro-ph.CO].
- [7] R. Kessler *et al.*, Publ. Astron. Soc. Pac. **122**, 1415 (2010) [arXiv:1008.1024 [astro-ph.CO]].
- [8] M. Ganeshalingam *et al.*, arXiv:1107.2404 [astro-ph.CO].
- [9] R. Chornock *et al.*, Astrophys. J. **780**, 44 (2014) [arXiv:1309.3009 [astro-ph.CO]].
- [10] B. C. Kelly *et al.*, Astrophys. J. **730**, 52 (2011) [arXiv:1009.6011 [astro-ph.HE]].
- [11] S. Gezari *et al.*, Astrophys. J. **766**, 60 (2013) [arXiv:1302.1581 [astro-ph.CO]].
- [12] T. Kanungo *et al.*, IEEE Trans. Pat. Ana. Mach. Intel. **7**, 24 (2002)
- [13] K. P. Burnham and D. R. Anderson Statistical Theory and Methods (2002)

South African Astro-informatics Alliance (SA3)

Sudhanshu Barway
South African Astronomical Observatory

Abstract

South African astronomy is entering a new era of astronomical research with SALT and MeerKAT/SKA. These new facilities expected to produce huge amount of data and combined with multi wavelength databases that already exists, South African astronomers need to be equipped with latest technologies to deal with new challenges posed by SALT/MeerKAT/SKA. South African Astro-informatics Alliance (SA3) is a collaborative initiative lead by SAAO, SA-SKA and HartRAO to utilize the most recent advancements in IT technology to address many Terabytes of data volume generated by SALT/MeerKAT/SKA along with existing multi-wavelength data archives. In this talk, I present SA3 activities which include the development of a new generation of data archives and tools to address the many Terabytes of data that will be generated by the new South African observing facilities.

Predicting Fundamental Stellar Parameters from Photometric Light Curves

Adam Miller

Jet Propulsion Laboratory / California Institute of Technology

Abstract

We present a new machine learning based framework for the prediction of the fundamental stellar parameters, T_{eff} , $\log g$, and [Fe/H], based on the photometric light curves of variable stellar sources. The method was developed following a systematic spectroscopic survey of stellar variability. Variable sources were selected from repeated Sloan Digital Sky Survey (SDSS) observations of Stripe 82, and spectroscopic observations were obtained with Hectospec on the 6.5-m Multi-Mirror Telescope. In sum, spectra were obtained for $\sim 9,000$ stellar variables (including $\sim 3,000$ from the SDSS archive), for which we measured T_{eff} , $\log g$, and [Fe/H] using the Segue Stellar Parameters Pipeline (SSPP). Examining the full sample of $\sim 67,000$ variables in Stripe 82, we show that the vast majority of photometric variables are consistent with main-sequence stars, even after restricting the search to high galactic latitudes. From the spectroscopic sample we confirm that most of these stellar variables are G and K dwarfs, though there is a bias in the output of the SSPP that prevents the identification of M type variables. We are unable to identify the dominant source of variability for these stars, but eclipsing systems and/or star spots are the most likely explanation. We develop a machine learning model that can determine T_{eff} , $\log g$, and [Fe/H] without obtaining a spectrum. Instead, the random forest regression model uses SDSS color information and light curve features to infer stellar properties. We detail how the feature set is pruned and the model is optimized to produce final predictions of T_{eff} , $\log g$, and [Fe/H] with a typical scatter of 165 K, 0.42 dex, and 0.33 dex, respectively. We further show that for the subset of variables with at least 50 observations in the g band the typical scatter reduces to 75 K, 0.19 dex, and 0.16 dex, respectively. We consider these results an important step on the path to the efficient and optimal extraction of information from future time-domain experiments, such as the Large Survey Synoptic Telescope. We argue that this machine learning framework, for which we outline future possible improvements, will enable the construction of the most detailed maps of the Milky Way ever created.

State-Based Models for Light Curve Classification

*Andrew Becker
University of Washington*

Abstract

I discuss here the application of continuous time autoregressive models to the characterization of astrophysical variability. These types of models are general enough to represent many classes of variability, and descriptive enough to provide features for lightcurve classification. Importantly, the features of these models may be interpreted in terms of the power spectrum of the lightcurve, enabling constraints on characteristic timescales and periodicity. These models may be extended to include vector-valued inputs, raising the prospect of a fully general modeling and classification environment that uses multi-passband inputs to create a single phenomenological model. These types of spectral-temporal models are an important extension of extant techniques, and necessary in the upcoming eras of Gaia and LSST.

Follow-up Science, Opportunities and Strategies

Early Time Optical Emission from Gamma-Ray Bursts

Drejc Kopač, Andreja Gomboc, and Jure Japelj

Faculty of Mathematics and Physics, University of Ljubljana, Slovenia

1 Introduction

Early time optical emission (i.e., emission before $\sim 10^3$ sec after the GRB trigger) is an important tool to study the physics of Gamma-Ray Bursts (GRBs). Especially when it is detected during the ongoing gamma-ray emission, early time optical emission can help us understand true mechanisms behind prompt GRB emission and provide constraints on current and future GRB emission models.

In the last decade, the number of GRBs with early time optical detection has become large enough to allow various statistical studies (e.g., [1], [2]). This was made possible due to fast and accurate GRB localizations by the *Swift* satellite, immediate dissemination of GRB position via the Gamma-ray Coordinates Network, and due to growing number of rapid-response, fully autonomous optical telescopes. Furthermore, robotic telescopes or large field of view surveys are capable of detecting optical emission from GRBs during the still ongoing prompt gamma-ray emission, and are thus providing a growing sample of GRBs with contemporaneous gamma-ray and optical detection.

By selecting a heterogeneous subsample of 18 GRBs showing *optical peaks* during the ongoing gamma-ray emission, we performed an extensive temporal and spectral analysis of prompt optical emission. Here we discuss these results in a broader context of early time optical emission classification, and explore how early time optical polarimetry could serve as an important tool to study the mechanisms of prompt optical emission.

2 GRBs and early time optical emission

In the most commonly accepted theoretical interpretation, the fireball model is invoked to explain GRB emission. Central engine (presumably a black hole – accretion disk system) pumps out a pair of opposite jets, which are powered by the accretion of matter. Ultra-relativistic shells with different bulk Lorentz factors are ejected, leading to relativistic shocks which take place inside the outflow (internal shocks) and

give rise to the prompt gamma-ray emission. When relativistic shells hit the external medium (external shocks), they are decelerated and the result is emission at longer wavelengths, called the afterglow. Currently, mainly because of the rich variety of available experimental data, the external shock region is understood better than the internal shock region, where a consensus about the dominant dissipation mechanism has not yet been reached [3]. Various processes likely play an important role in providing the dissipation and emission, e.g., relativistic shocks, magnetic reconnections, inverse Compton scattering, etc. Central engine region is understood poorly, since we only have indirect evidence about the nature of progenitors.

Prompt optical emission offers a direct probe to better understand the internal shock region, and thus the acceleration and dissipation processes which power GRBs. By analyzing the temporal profiles of prompt optical light curves, we can compare the characteristic time scales and temporal structure with those of prompt gamma-ray light curves. Multi-wavelength spectral analysis can show if the spectral energy distribution is consistent with synchrotron emission, and whether optical and gamma-ray emission originate from the same emission processes.

3 Analysis of prompt optical peaks

The motivation for our study was GRB 090727, which was detected by the *Swift* satellite and observed promptly by the Liverpool Telescope (see [1] and references therein). This GRB showed a steep optical peak, which happened when the high-energy (gamma-ray and X-ray) emission was still active (see Figure 1, panel “090727”). Our modeling showed that the most likely scenario is that optical and high-energy peaks are not completely simultaneous, and more importantly, that the optical peak is not consistent with the interpretation in the context of the standard reverse external shock emission, because it is too steep. To put this result in a broader context, we selected a heterogeneous sample of 18 GRBs which show early time optical peaks occurring during the still ongoing gamma-ray emission (Figure 1).

By performing a detailed temporal analysis of optical peaks in our sample, we obtained rise and decay indices, peak times (t) and the peak durations (Δt). We found that in many, but not all cases, prompt optical peaks are very sharp, meaning that their durations are smaller than their peak times ($\Delta t/t < 1$), which implies an internal shock origin. Also, their rise (α_{rise}) and decay (α_{decay}) power-law indices are in many cases very steep and therefore inconsistent with the predictions of the standard afterglow emission theory. Figure 2 (left panel) shows power-law indices of these peaks, and the size of the circle around each point represents the relative value of $\Delta t/t$; smaller circle means lower $\Delta t/t$ and thus sharper peak. It is evident that steeper and sharper peaks likely do not correspond to afterglow emission (and do not populate the grey area in Figure 2, left panel).

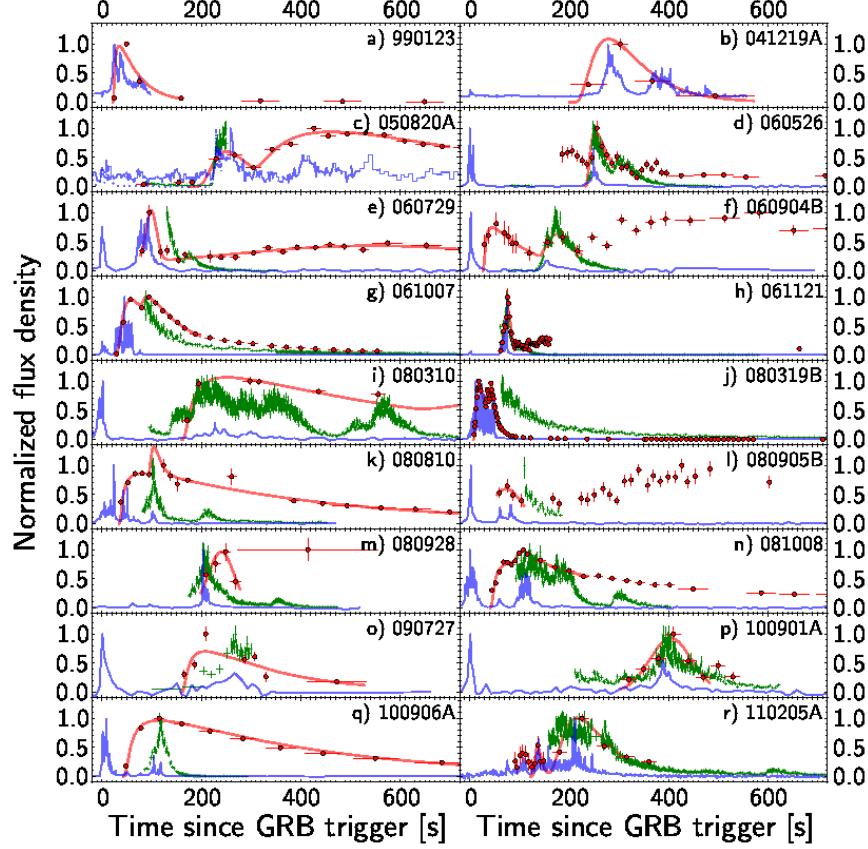


Figure 1: Optical (red), gamma-ray (blue) and X-ray (green) normalized light curves of GRBs showing prompt optical peaks [1].

Our spectral analysis showed that gamma-ray and optical-to-gamma-ray extrapolated spectral indices, obtained by fitting the spectrum with simple power-law behaviour ($F_\nu \propto \nu^{-\beta}$), are consistent with the synchrotron emission from relativistic electrons. Furthermore, the ratio of fluxes between gamma-ray and optical emission, $(\nu F_\nu)^\gamma / (\nu F_\nu)^{\text{OPT}}$, showed that the diversity is large, as the distribution of flux ratios spans over 5 orders of magnitude (Figure 2, right panel, green and blue histograms).

Such diversity in measured parameters, both temporal and spectral, implies some random processes and points towards prompt phase origin. We thus performed a simple internal shock dissipation simulation, where two ultra-relativistic shells collide and produce a pair of internal shocks, one forward and one reverse, similarly as in the external shock scenario. By simulating a number of such collisions within the parameter space of typical GRB properties, we predicted the observed flux in gamma-ray and optical regime using synchrotron emission theory (see [1] and references therein). The simulated distribution was consistent with $(\nu F_\nu)^\gamma / (\nu F_\nu)^{\text{OPT}}$

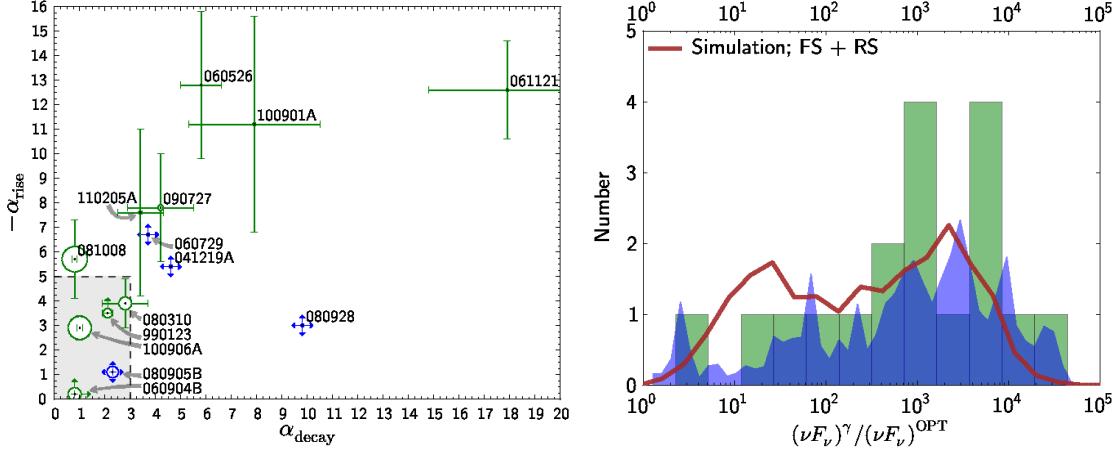


Figure 2: **Left:** $|\alpha_{\text{rise}}|$ versus α_{decay} diagram. Gray shaded area shows the region of power-law indices as predicted by the standard afterglow emission theory [4, 5]. The radius of the circle around each point represents the relative value of $\Delta t/t$. **Right:** Flux ratio distributions as obtained from observations (green and blue histograms, where blue histogram takes into account also the scatter of observed flux ratios) and from the internal shock dissipation model simulation (normalized brown solid line) [1].

distribution obtained from observations (Figure 2, right panel, brown solid line). We concluded that early time optical emission with sharp and steep peaks, occurring during the ongoing gamma-ray emission, could not be interpreted in the context of the standard external shock afterglow model, and could instead originate from dissipation within internal shock region, similarly as prompt gamma-ray emission [1].

4 Early time optical emission classification

Although sharp and steep prompt optical peaks, large diversity in temporal and spectral parameters, and consistency with internal shock dissipation simulation imply internal shock origin for early time optical emission, this is not the case for all optical peaks. Especially for GRBs where optical peaks are noticeably less sharp and steep, we can speculate that early time optical emission can represent standard afterglow emission, likely due to reverse external shock, which can manifest itself as a bright optical flash in the early time light curve¹ [6].

Furthermore, prompt optical and gamma-ray emission do not necessarily originate from the same internal shock component. We can have more collisions happening at various regions in the flow, one producing brighter gamma-ray emission and other

¹For example, this is evident for GRB 990123, where reverse external shock model can explain the early time optical light curve behavior (see [7]), and for which $\Delta t/t$ value is higher than 1.

producing brighter optical emission. In some cases optical and gamma-ray peaks could originate from just a slightly different region in the flow (different lab times and/or different radii of the expanding flow), while in others the difference between lab times is much higher, and so the temporal correlation is not a requirement. This means that in the majority of cases it is not straightforward to claim whether gamma-ray and optical peaks originate from the same region or not. Optical peak which seems “isolated” during the still ongoing gamma-ray emission could originate from internal shocks, but could as well be due to reverse external shock, other internal dissipation mechanism like magnetic reconnections, variable microphysics parameters, structured jets, etc.

To obtain more information and to better understand the early time optical emission, optical polarimetry serves as an important tool which can give more constraints on models and help determine the origin of early time optical emission.

5 Early time optical polarimetry

If GRB jets are baryonic and contain large scale and ordered magnetic fields, which could be advected from the central engine and survive long after the initial explosion, then any synchrotron emission from such medium should be highly (up to $\sim 70\%$) linearly polarized. Polarization at early times after the GRB explosion has been first measured for prompt gamma-ray emission, and especially recent results from GAP instrument proved that the degree of prompt gamma-ray polarization is indeed very high [8].

In optical regime, until 2006 polarization measurements have been performed typically around one day after the GRB trigger. In 2006, RINGO polarimeter, mounted on the Liverpool Telescope, provided the first early time optical afterglow polarization measurements for GRB 060418, and later for GRB 090102 [9, 10]. At early times, optical emission can be dominated by the reverse external shock afterglow contribution, and if the jet is magnetized, we expect to detect high degree of polarization from the reverse shock photons. In this case polarization is expected to decay with time as the reverse shock component fades and forward shock component starts to dominate the early time optical light curve. Decaying degree of polarization has been observed very recently with RINGO2 polarimeter in the case of GRB 120308A, where polarization degree of $P = 28 \pm 4\%$ has been measured at 4 minutes after the GRB trigger, and decreased over the next 10 minutes to $P = 16_{-4}^{+5}\%$ while polarization position angle remained stable [11].

The next step would be to measure polarization of prompt optical emission. RINGO3 polarimeter, currently mounted on the Liverpool Telescope, is capable of measuring polarization of the optical source brighter than 17th magnitude promptly after the GRB trigger. In case of sharp and steep optical peaks, which would occur

during prompt gamma-ray emission, polarization measurements would allow to test if the degree of polarization in optical regime is very high in comparison with the polarization degree of gamma-ray emission, and whether it decays when optical afterglow emission becomes dominant. Furthermore, since prompt emission originates from regions closest to the central engine, this would help us to study the spatial evolution of the GRB outflow.

6 Conclusion

GRBs with contemporaneous optical and gamma-ray detections are an important tool to study the GRB emission mechanisms. We show that prompt optical peaks with sharp and steep morphology could originate from internal shock region. However, early time optical light curves are in many cases complex and can show contributions from various emission components. This is often connected with experimental limitations, mostly due to faintness of early time optical emission, inadequate temporal resolution of optical observations, and too long response time of robotic optical telescopes.

Early time optical polarization measurements and its temporal evolution across various phases of optical light curve, coupled with polarization measurements in other wavelengths (gamma rays and radio), will provide vital information about the nature of GRB emission.

References

- [1] Kopač, D. et al., 2013, ApJ, 772, 73
- [2] Kann, D.A. et al., 2010, ApJ, 720, 1513
- [3] Gomboc, A., 2012, ConPh, 53, 339
- [4] Zhang, B., Kobayashi, S. & Mészáros, P., 2003, ApJ, 595, 950
- [5] Kobayashi, S. & Zhang, B., 2003, ApJ, 597, 455
- [6] Kobayashi, S., 2000, ApJ, 545, 807
- [7] Japelj, J. et al., 2014, ApJ, submitted
- [8] Yonetoku, D. et al., 2012, ApJL, 758, 1
- [9] Mundell, C.G. et al., 2007, Science, 315, 1822
- [10] Steele, I.A. et al., 2009, Nature, 462, 767
- [11] Mundell, C.G. et al., 2013, Nature, 504, 119

Burst of the Century? A Case Study of the Afterglow of Nearby Ultra-Bright GRB 130427A

*Daniel Perley
California Institute of Technology*

Abstract

GRB 130427A is the brightest gamma-ray burst observed by any satellite in almost 30 years and one of the most thoroughly observed. I will present a summary of the worldwide campaign to monitor the afterglow of this event from GHz to TeV energies and from seconds to years after the explosion. Remarkably, the entire data set can be described to good agreement using standard synchrotron afterglow theory, providing strong support for the validity the basic model in describing the evolution of this event and for GRB afterglows generally. Distinct forward and reverse shock components are resolved in both the SED and multifrequency light curves; the late-time high-energy emission seen by LAT is produced by the forward shock. We also infer a tenuous, wind-stratified medium surrounding this burst, suggesting a massive, low-metallicity progenitor. While GRB 130427A was an incredibly rare and fortuitous event its properties are probably not intrinsically unusual, and it provides lessons for what might be routinely achieved in the future with faster and deeper multiwavelength follow-up of gamma-ray bursts.

Optical Interferometry and Adaptive Optics of Bright Transients

*Florentin Millour¹, Olivier Chesneau, Anthony Meilland, and Nicolas Nardetto
Lagrange Laboratory, Université de Nice Sophia-Antipolis,
CNRS, Observatoire de la Côte d'Azur, Nice, France*

1 Introduction

Bright optical transients (i.e. transients typically visible with the naked eye) are populated mainly by novæ eruptions plus a few supernovæ (among which the SN1987a event). Indeed, usually one bright nova happens every two years, either in the Northern or the Southern hemisphere (see Fig. 1). It so happens that current interferometers have matching sensitivities, with typically visible or infrared limiting magnitudes in the range 5–7. The temporal development of the fireball, followed by a dust formation phase or the appearance of many coronal lines can be studied with the VLTI. The detailed geometry of the first phases of novæ in outburst remain virtually unexplored. This paper summarizes the work which has been done to date using mainly the Very Large Telescope Interferometer.

We invite the reader to have a look at the extensive review on the topic by Chesneau & Banerjee [1] for a complete description of the science on transients that can be achieved with optical/infrared long-baseline interferometers. We give a short summary of the content of this paper in the next section.

2 Why observing novæ with optical interferometers?

Optical interferometers represent a breakthrough in terms of spatial resolution, that can provide crucial information related to the nova phenomenon. All targets in the 3–5 kpc range can be potentially resolved by current interferometers (CHARA, VLTI, NPOI).

The VLT Interferometer can provide measurements of the angular diameters of the nova ejecta in continuum and lines, from the near-IR to the mid-IR in the very first moments of the outburst. The primary outcome of these observations is a direct

¹Send comments to fmillour@oca.eu

estimate of the expansion parallax, thus the distance to novae. Of importance is the possibility to spatially and spectrally resolve different near-IR emission lines to estimate the physical conditions throughout the wind and the ejecta. For those objectives, the medium spectral resolution of the near-IR instrument AMBER is an asset. If the nova appears to form dust (CO novae), an in-depth study of the dust forming regions can be carried out with the MIDI instrument. Using a set of flexible observing runs, we shall follow the outburst from the first days up to several months.

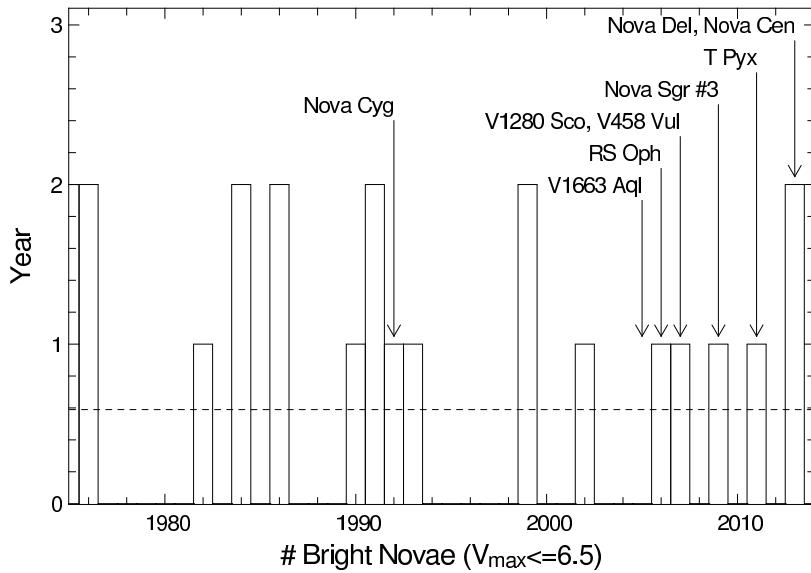


Figure 1: Typical frequency of bright novae from 1975 to 2013. About one nova every two years occurs, either in the Northern or the Southern hemisphere (data from http://www.cbat.eps.harvard.edu/nova_list.html, complemented by data from <http://asd.gsfc.nasa.gov/Koji.Mukai/novae/novae.html>. Some peak magnitudes were corrected using AAVSO data). V1663 Aql & V458 Vul do not appear in this plot as they are too faint in V (but bright in K). The average rate per year is plotted as a dashed line. The arrows show Novae observed with optical interferometers.

3 Nova as a spherical fireball

Up to now, the program of observations has focused on novae with magnitudes reachable by the VLTI (South declination & K magnitude ≥ 7). Past observations and theoretical work on the nova phenomenon have provided a substantial knowledge about the physical nature of these binary systems and the outburst. However, these investigations are naturally limited by the difficulty of estimating the distance, which

is usually inferred indirectly and with large errors. Spherical symmetry is a basic tenet adopted in the derivation of relationships that link the non-spatially resolved photometric and spectroscopic observations to the physical parameters of the system [2, Table 3]. Spherical symmetry is implicitly assumed when the uv coverage is not sufficient to perform a better analysis. This was the case for the Nova V1280 Sco [3] which was also observed exclusively using 2 telescope recombination. Two years after, observations with the AO system NACO mounted at the UT4 telescope revealed an impressive dusty bipolar nebula [4, see Fig. 2].

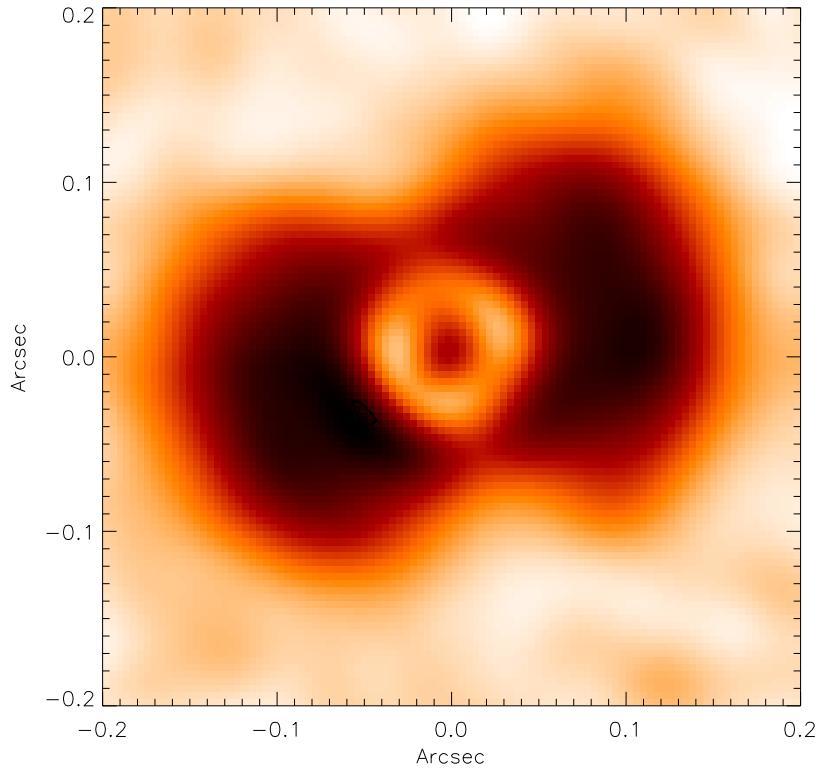


Figure 2: 2010 NACO K band image after a PSF subtraction revealing the impressive bipolar nebula. Mid-IR images also show that there is no dust emission in the equatorial plane.

4 A bipolar fireball from the first blink

An interferometer is mainly sensitive to the angular size of a nova in its early stages. Measuring the size of the fireball in different orientations on-sky allows us to infer the axis ratio and orientation of an individual nova shell. This is relatively easy to

obtain for a 3–6-telescopes interferometer. This was the case for the outburst of the recurrent nova RS Oph [5].

The highly collimated outflow from the RS Ophiuchi has been imaged by the HST [6] and in the radio [7]. The AMBER observations showed that the jet was already in existence 5.5 days after the discovery, and provided a unique view of radial-velocities which could afterwards complement the expansion rates derived by the HST and radio images [8].

The signature of a bipolar jet in interferometric data is now well identified provided that the emission lines are spectrally resolved ($R \sim 1500$). The famous nova T Pyx exhibited a spherical appearance in broadband PIONIER data, but the signature of bipolar kinematics was clearly detected in our spectrally resolved AMBER data [9]. The numerous peculiarities of the T Pyx eruptions can be explained in the frame of recurrent nearly face-on eruptions that launch fast material in the line-of-sight and slow material perpendicular to it, building up the slow expansion shell imaged by the HST.

5 Intermediate Luminosity Optical Transients

Intermediate-Luminosity Optical Transients (ILOTs), are eruptive stars with peak luminosity between those of novæ and supernovæ that have been also called Red Novæ or red Transient. The powering processes and whether they are due to binary interaction or are formed through single star evolution is debated. High angular resolution techniques can play a role by tracking bipolarity and the formation of disks. Furthermore, one can also study the remaining central star when a merger is highly suspected, for instance by detecting the deformation due to very high rotational rate.

5.1 Sakurai’s object

In 1996, Sakurai’s object (V4334 Sgr) suddenly brightened in the center of a faint Planetary Nebula (PN). This very rare event was interpreted as being the reignition of a hot white dwarf that caused a rapid evolution back to the cool giant phase. From 1998 on, copious amount of dust has formed continuously, screening out the star that remained embedded in this expanding high optical-depth envelope. Mid-IR interferometry performed in 2008 with the MIDI/VLTI instrument discovered an unexpectedly compact (30–40 milli-arc-second, 105–140 AU assuming a distance of 3.5 kpc), highly inclined, dust disk [10]. The major axis of the disk is aligned with an asymmetry seen in the old PN. This implies that the mechanism responsible for shaping the dust envelope surrounding Sakurai’s object was already at work when the old PN formed, a strong argument for binary interaction.

5.2 V 838 Mon

V838 Monocerotis erupted in 2002, brightening by 9 magnitudes in a series of outbursts, and eventually developed a spectacular light echo. A very red star emerged surrounded by copious amount of new dust that condensed from the expanding ejecta of the outbursts. V838 Mon is the close-by archetype of the ILOT sources which are triggering very active research currently. MIDI/VLTI observations obtained over the last few months showed that the dust resides in the form of a flattened structure ($\sim 15\text{-}50$ mas from 8 to 13 μm), i.e. a 90x300 AU flattened structure for a distance of 6.2 kpc (Sparks et al. 2008). The modelling of this extended structure is in progress but it is incomplete without a much better knowledge of the central source which is seen as a very cool M-L type super-giant. AMBER observations were also obtained in 2013 to measure the size of the central source, its shape (since it has potential to be a fast rotator) and study the cool photosphere/dusty disk transition. The AMBER data were essentially acquired with small baselines (≤ 30 m) and are quite noisy ($\sigma V^2 \approx 0.05$), leading to an angular size of 3 mas, but with a large uncertainty. We can basically only set an upper limit to the diameter of the HK-bands object, of 4.7 mas. This would make the HK-bands object smaller than 30 AU. This size (and shape) difference between the HK-bands and N band is striking and reminiscent of super-giant stars dusty disks [11, 12]. Further AMBER observations would enable us to pinpoint more precise properties of this intriguing object.

6 Conclusion and prospectives

We have presented here a few results from the VLTI campaigns on Novæ and ILOTs. These campaigns present a challenge in terms of scheduling and observatory response but provide unique insights on the early processes at stake when a nova explodes. With the upcoming infrared instruments like MATISSE [13] or GRAVITY [14], getting a finer idea of the geometry will be much faster than with current instruments. The development of visible interferometric instruments at the CHARA array or at the VLTI would also be an asset to get the sharpest multi-wavelength picture of these objects a few days after outburst before the advent of the ELT in the mid-2030s which will enable direct snapshot pictures of the fireball at tens of milli-arc-seconds resolution for a much larger number of novæ.

References

- [1] O. Chesneau and D. P. K. Banerjee. Interferometric studies of novae in the infrared. *Bulletin of the Astronomical Society of India*, 40:267, September 2012.
- [2] R. D. Gehrz, J. W. Truran, R. E. Williams, and S. Starrfield. Nucleosynthesis in Classical Novae and Its Contribution to the Interstellar Medium. *PASP*, 110:3–26, January 1998.
- [3] O. Chesneau, D. P. K. Banerjee, F. Millour, N. Nardetto, S. Sacuto, A. Spang, M. Wittkowski, N. M. Ashok, R. K. Das, C. Hummel, S. Kraus, E. Lagadec, S. Morel, M. Petr-Gotzens, F. Rantakyro, and M. Schöller. VLTI monitoring of the dust formation event of the Nova V1280 Scorpii. *A&A*, 487:223–235, aug 2008.
- [4] O. Chesneau, E. Lagadec, M. Otulakowska-Hypka, D. P. K. Banerjee, C. E. Woodward, E. Harvey, A. Spang, P. Kervella, F. Millour, N. Nardetto, N. M. Ashok, M. J. Barlow, M. Bode, A. Evans, D. K. Lynch, T. J. O’Brien, R. J. Rudy, and R. W. Russell. The expanding dusty bipolar nebula around the nova V1280 Scorpi. *A&A*, 545:A63, September 2012.
- [5] O. Chesneau, N. Nardetto, F. Millour, C. Hummel, A. Domiciano de Souza, D. Bonneau, M. Vannier, F. Rantakyro, A. Spang, F. Malbet, D. Mourard, M. F. Bode, T. J. O’Brien, G. Skinner, R. G. Petrov, P. Stee, E. Tatulli, and F. Vakili. AMBER/VLTI interferometric observations of the recurrent Nova RS Ophiuchi 5.5 days after outburst. *A&A*, 464:119–126, mar 2007.
- [6] M. F. Bode, D. J. Harman, T. J. O’Brien, H. E. Bond, S. Starrfield, M. J. Darnley, A. Evans, and S. P. S. Eyres. Hubble Space Telescope Imaging of the Expanding Nebular Remnant of the 2006 Outburst of the Recurrent Nova RS Ophiuchi. *ApJL*, 665:L63–L66, August 2007.
- [7] T. J. O’Brien, M. F. Bode, R. W. Porcas, T. W. B. Muxlow, S. P. S. Eyres, R. J. Beswick, S. T. Garrington, R. J. Davis, and A. Evans. An asymmetric shock wave in the 2006 outburst of the recurrent nova RS Ophiuchi. *Nature*, 442:279–281, July 2006.
- [8] V. A. R. M. Ribeiro, M. F. Bode, M. J. Darnley, D. J. Harman, A. M. Newsam, T. J. O’Brien, J. Bohigas, J. M. Echevarría, H. E. Bond, V. H. Chavushyan, R. Costero, R. Coziol, A. Evans, S. P. S. Eyres, J. León-Tavares, M. G. Richer, G. Tovmassian, S. Starrfield, and S. V. Zharikov. The Expanding Nebular Rem-

- nant of the Recurrent Nova RS Ophiuchi (2006). II. Modeling of Combined Hubble Space Telescope Imaging and Ground-based Spectroscopy. *ApJ*, 703:1955–1963, October 2009.
- [9] O. Chesneau, A. Meilland, D. P. K. Banerjee, J.-B. Le Bouquin, H. McAlister, F. Millour, S. T. Ridgway, A. Spang, T. Ten Brummelaar, M. Wittkowski, N. M. Ashok, M. Benisty, J.-P. Berger, T. Boyajian, C. Farrington, P. J. Goldfinger, A. Merand, N. Nardetto, R. Petrov, T. Rivinius, G. Schaefer, Y. Touhami, and G. Zins. The 2011 outburst of the recurrent nova `ASTROBJ`. Evidence for a face-on bipolar ejection. *A&A*, 534:L11, oct 2011.
 - [10] O. Chesneau, G. C. Clayton, F. Lykou, O. de Marco, C. A. Hummel, F. Kerber, E. Lagadec, J. Nordhaus, A. A. Zijlstra, and A. Evans. A dense disk of dust around the born-again Sakurai’s object. *A&A*, 493:L17–L20, jan 2009.
 - [11] F. Millour, A. Meilland, O. Chesneau, P. Stee, S. Kanaan, R. Petrov, D. Mourard, and S. Kraus. Imaging the spinning gas and dust in the disc around the supergiant A[e] star HD 62623. *A&A*, 526:A107, feb 2011.
 - [12] H. E. Wheelwright, W. J. de Wit, R. D. Oudmaijer, and J. S. Vink. VLTI/AMBER observations of the binary B[e] supergiant HD 327083. *A&A*, 538:A6, February 2012.
 - [13] B. Lopez, S. Wolf, S. Lagarde, P. Abraham, P. Antonelli, J. C. Augereau, U. Beckman, J. Behrend, N. Berruyer, Y. Bresson, O. Chesneau, J. M. Clausse, C. Connot, K. Demyk, W. C. Danchi, M. Dugué, S. Flament, A. Glazeborg, U. Graser, T. Henning, K. H. Hofmann, M. Heininger, Y. Hugues, W. Jaffe, S. Jankov, S. Kraus, W. Laun, C. Leinert, H. Linz, P. Mathias, K. Meisenheimer, A. Matter, J. L. Menut, F. Millour, U. Neumann, E. Nussbaum, A. Niedzielski, L. Mosonic, R. Petrov, T. Ratzka, S. Robbe-Dubois, A. Roussel, D. Schertl, F.-X. Schmider, B. Stecklum, E. Thiebaut, F. Vakili, K. Wagner, L. B. F. M. Waters, and G. Weigelt. MATISSE: perspective of imaging in the mid-infrared at the VLTI. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6268 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2006.
 - [14] F. Eisenhauer, G. Perrin, S. Rabien, A. Eckart, P. Lena, R. Genzel, R. Abuter, and T. Paumard. GRAVITY: The AO assisted, two object beam combiner instrument for the VLTI. *Astronomische Nachrichten*, 326:561–562, August 2005.

Multi-Color Robotic Observations with RATIR

*Nathaniel Butler
Arizona State University*

Abstract

We have finished the first year of science operations of a novel 6-color, simultaneous imaging camera – RATIR – mounted on an automated 1.5m telescope at San Pedro Martir Observatory, Baja, CA, MX. The camera/telescope system is designed for rapid (<5 min) observations of GRBs in the *riZYJH* bands, providing potential evidence for very high redshifts ($z > 6$) inferred through measurements of the IGM attenuation (dropouts). I will discuss the RATIR design and implementation as well as science results obtained to date.

The Robotic FLOYDS Spectrographs

*David Sand
Texas Tech University*

Abstract

I will discuss the twin FLOYDS robotic spectrographs, operating at the 2m Faulkes Telescopes North and South. The FLOYDS instruments were designed with supernova classification and monitoring in mind, with a very large wavelength coverage (\sim 320 to 1000 nm) and a resolution ($R \sim 300 - 500$, wavelength dependent) well-matched to the broad features of these and other transient and time domain events. Robotic acquisition of spectroscopic targets is the key ingredient for making robotic spectroscopy possible, and FLOYDS uses a slit-viewing camera with a $\sim 4' \times 6'$ field to either do direct world coordinate system fitting or standard blind offsets to automatically place science targets into the slit. Future work includes an 'all-electronic' target of opportunity mode, which will allow for fast transient spectroscopy with no human necessary, even for inputting information into a phase 2 GUI. Initial science highlights from FLOYDS will also be presented.

Dynamic Follow-up of Transient Events with the LCOGT Robotic Telescope Network

*R. A. Street, and Y. Tsapras
Las Cumbres Observatory Global Telescope*

*M.P.G. Hundertmark, K.D. Horne, and M. Dominik
School of Physics and Astronomy, University of St Andrews, UK*

*C. Snodgrass
Max Planck Institute for Solar System Research, Germany*

*D.M. Bramich
Qatar Environment and Energy Research Institute, Qatar Foundation, Qatar*

*R. Figuera Jaimes
European Southern Observatory, Germany*

1 Introduction

The demands of astronomical observations have changed dramatically over recent years. Traditionally, an observer would personally visit a single telescope with a static target list, to observe for a contiguous block of time awarded 6 months in advance. Modern multi-year surveys are breaking this paradigm. They observe continuously, from both hemispheres on the ground and from space with a range of different survey strategies, and produce a continuous stream of discoveries via online alerts. Some require an immediate response in order to extract new science. Most surveys are designed to provide new target detections only and cannot usually provide comprehensive characterization or even classification. Follow-up teams have therefore been organized to provide additional observations, but it is generally impossible for them to provide a target list in advance and fixed blocks of telescope time are often disadvantageous and inefficient. Teams of observers are often needed, creating a huge overhead in travel costs. The Las Cumbres Observatory Global Telescope Network (LCOGT) is a new and unique facility designed to address this issue.

2 The LCOGT Network

LCOGT is a network of 12 robotic telescopes, geographically distributed in both latitude and longitude: $2 \times 2\text{m} + 9 \times 1\text{m} + 1 \times 0.8\text{m}$ telescopes spread across 6 sites in both hemispheres. Telescopes are organized in clusters of between 1–3 at each site, and each aperture class of telescope supports an homogenous set of instruments and filters. The network is described in detail in [1]. During 2013, the 2m network was scheduled on legacy software, independently from the 1m network, but in the future all telescopes will be scheduled, crucially, *as a single facility*. This enables the maximum degree of flexibility in scheduling, allowing the network to compensate automatically for the loss of any site or individual telescope to weather or technical issues. It also brings new and unique science opportunities, being capable of targeting the same object from multiple sites, telescopes and instruments. In particular, with a multiple sites per hemisphere, greatly extended time series observations are possible at any cadence. For more details on the LCOGT Network Scheduler, see [2].

3 Exploiting Robotic Facilities for Astronomy

LCOGT has developed a web-based portal and Application Programming Interface (API) for the purpose of submitting observation requests. For large scale, highly responsive programs, efficiency and man-power limits make it desirable to have project-side software to compute the observations necessary for our science goals and interact with the LCOGT scheduler robotically. The system is designed to respond appropriately at any time to new survey alerts, without waiting for human approval.

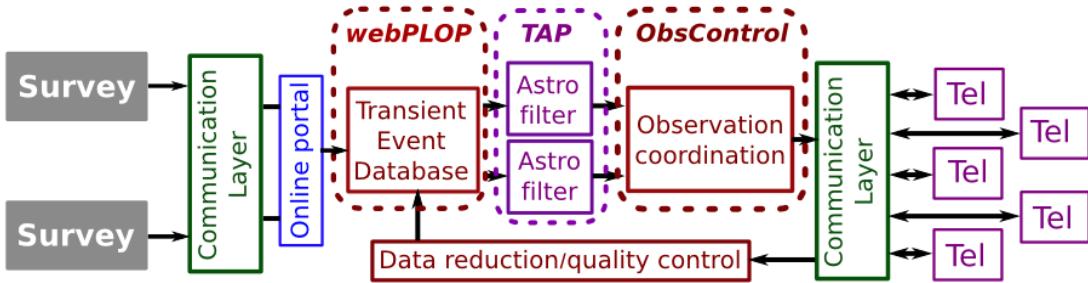


Figure 1: The outline of the robotic system designed to respond automatically to survey alerts and provide appropriate follow-up observations of targets identified to be high priority.

Figure 1 outlines the structure of this system. Its key features are i) survey alerts are gathered into a central database, ii) the alerts are classified and prioritized

according to one or more filters designed to identify specific classes of events of astronomical interest, iii) selected targets trigger observation requests according to a pre-determined ‘recipe’, which may vary according to the stage of an event and/or its parameters measured at a given time, iv) the data taken are automatically reduced and v) the target is automatically re-evaluated in the light of the new information and the observing program adjusted as necessary. We have built a system, following this structure, which runs a large-scale robotic observing program designed to characterize microlensing events.

4 Microlensing Planet Detection

A microlensing event occurs when a foreground star (the ‘lens’) crosses the observer’s line of sight to a background star (the ‘source’). This causes the source to brighten gradually as the gravity of the lens bends light around it, towards the observer, fading back to its normal brightness as the stars move out of alignment over the course of days to months. If the lensing star hosts a planet, then this can provide additional magnification of the source star, called an ‘anomaly’, if the planet happens to be close to the Einstein radius (θ_E) of its host star at the time of the event. This radius, $\theta_E = \sqrt{(4GM_L/c^2)(D_L^{-1} - D_S^{-1})}$, is determined by the lens star mass (M_L) and the distances of the lens and source from the observer (D_L and D_S). These events are transient, so it is necessary to obtain immediate follow-up observations. Furthermore, anomalies can occur at any time and last anything from just minutes to days. Microlensing events therefore have to be monitored around-the-clock, but different stages require different densities of sampling, as illustrated by Fig. 2. The priorities can change at a moment’s notice, particularly once an anomaly is detected. Two established surveys, OGLE¹ [3] and MOA² [4, 5], discover \sim 2000 microlensing events each year and provide prompt online alerts. Our RoboNet³ [6] microlensing follow-up program employs not just the LCOGT telescope network but also the Liverpool Telescope, Canary Islands, Spain.

5 Responsive Automated Observation Control System

The observation control system consists of three automated stages (see Fig. 1):

¹ogle.astrouw.edu.pl

²www.phys.canterbury.ac.nz/moa

³robonet.lcogt.net

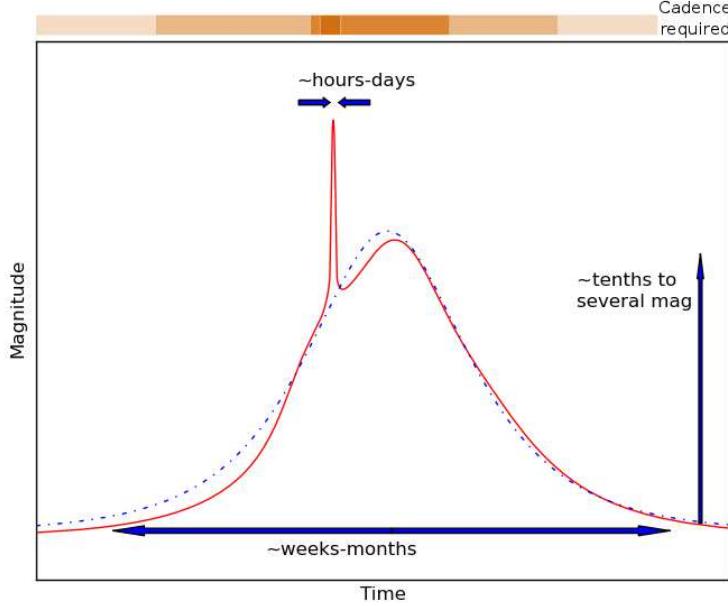


Figure 2: Schematic diagram of the light curve of a microlensing event. The density of shading of the bar at the top of the figure is directly proportional to the density of sampling necessary to properly characterize the critical stages of the light curve.

WebPLOP: the event database [7]:

This continuously harvests online alerts, downloading the location, finder chart and parameters of the initial model fit. It gathers all available light curve data by automatic query of the ARTEMiS system [8, 9], a public service that also provides alerts of anomalies. Updated data is gathered as the event progresses. This information is held in a database made publicly accessible through a web-based, interactive portal. WebPLOP performs its own model fitting procedure, but can also handle multiple ‘opinions’ on the same events from different sources, which feed into decisions on priorities for follow-up.

TAP: TArget Prioritization

In 2013 our team [10] developed a new algorithm to robotically prioritize events for follow-up. Typically \sim 50–100 events are ongoing but at different stages at any one time, and require a range of observing cadence (see Fig. 2). The parallel observing modes available on the LCOGT network require a further degree of optimization. TAP reviews the WebPLOP database at regular intervals, and returns a list of events to be observed, including recommendations for exposure time and, wherever possible,

ensures that a given event is observed with consistent instruments to simplify data reduction. This enables TAP to implement a pre-determined ‘observing recipe’ which dictates what observations should be made at each stage of an event, as a function of its peak magnification, current brightness, location on sky and existence of anomalous features. The software runs at frequent intervals under the cron to ensure a rapid response to new alerts, and also updates a human-readable HTML version of its output to assist coordination with other observers.

ObsControl.

This acts to implement TAP’s recommendations across the multiple telescope networks used by our program. In the 2013 season, while the LCOGT-network scheduler was at a primitive stage, it determined the optimum observing schedule for each of the 11 telescopes in use, factoring in hours of darkness, target visibility and telescope availability. These aspects continue to evolve to keep pace with the rapidly developing network-wide scheduler. ObsControl formulates and submits actual requests for observation using the different protocols currently required for the LCOGT-2m network, LCOGT-1m network and the Liverpool Telescope. The majority of the observations are conducted by the automated system, but there is also a web-based manual interface. ObsControl handles the download of the data obtained, which it filters through an in-built quality control process before preparing the data for automated pipeline reduction. All stages of our software maintain extensive, automated logging, crucial for both the identification and diagnosis of issues, and for evaluating the efficiency and effectiveness of its performance.

6 Conclusions

Figure 3 shows a sample of light curves obtained by our robotic follow-up program during the 2013 microlensing observing season. We observed 204 events in total, using 11 telescopes in multiple networks and achieved the dynamic, time variable cadence necessary to characterize these complex transients. This system demonstrated its ability to optimize a challenging observing program using the multiple modes of parallelization offered by LCOGT’s telescope network.

This has clear applications to many other fields in astronomy. The same system structure could be generalized to include other instrument types (eg. spectrographs) and to implement ‘observing recipes’ optimized for other classes of target.

We thank the Qatar National Research Fund, member of Qatar Foundation, for support from QNRF grant NPRP-09-476-1-078.

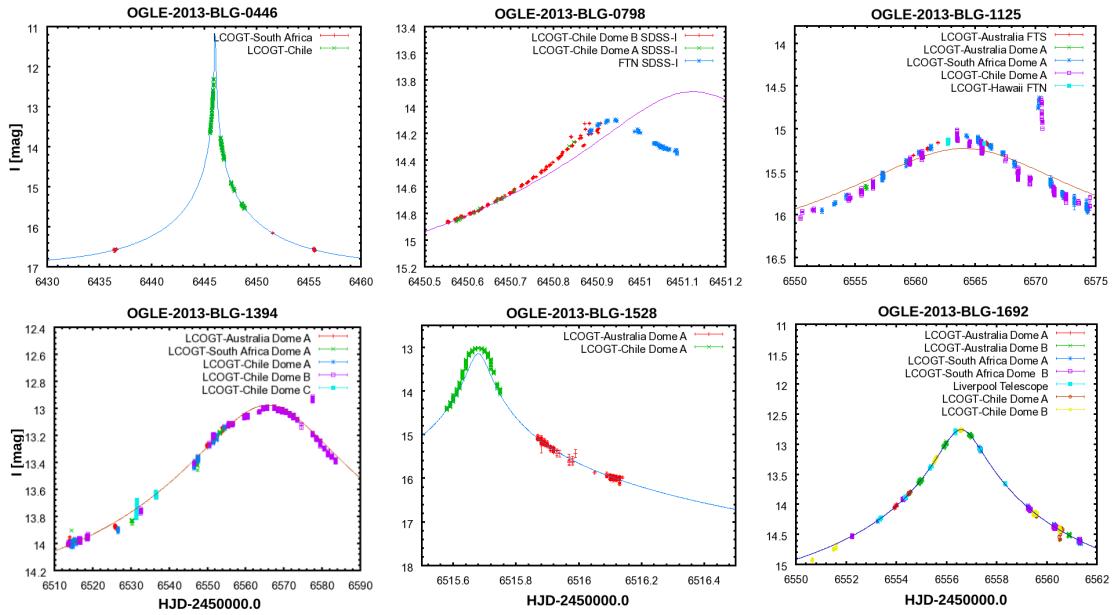


Figure 3: Sample of light curves obtained by our microlensing program in 2013, robotically coordinating 11 different telescopes.

References

- [1] T. Brown, et al., PASP, **125**, 1031 (2013).
- [2] E. Saunders, Proc. of Hotwiring the Transient Universe III, P. Wozniak, M. Graham, A. Mahabal, eds. (eConf, 2014).
- [3] A. Udalski et al., Acta Astron., **42**, 253, (1992).
- [4] I.A., Bond, F., Abe, R.J. Dodd et al., MNRAS, 327, 868 (2001).
- [5] T. Sumi, I.A. Bond, R.J. Dodd, ApJ, 591, 204 (2003)
- [6] Y. Tsapras et al., AN, **330**, 4, (2009).
- [7] C. Snodgrass, Y. Tsapras et al., Proc. of the Manchester Microlensing Conference, E. Kerins, S. Mao, N. Rattenbury, L. Wyrzykowski eds., 56, (2008).
- [8] M. Dominik et al. Proc. IAU Symp. 249, 35, (2008).
- [9] M. Dominik et al. AN, **329**, 248, (2008).
- [10] M. Hundertmark et al., ApJ, in prep. (2014).

Rapid Follow-up in iPTF and the Science it Enables

Iair Arcavi

Weizmann Institute of Science / Las Cumbres Observatory Global Telescope

Abstract

The intermediate Palomar Transient Factory (iPTF) is now routinely discovering supernovae within 24 hours of explosion. We have developed a pipeline for identifying interesting transients in real-time, sharing them among the collaboration and rapidly obtaining multi-wavelength followup observations. Such data is enabling us to "see" the types of stars which explode as supernovae and thus gain clues into one of the biggest unanswered questions in the field of explosive transients.

Transient Alert Follow-up Planned for CCAT

Tim Jenness

Department of Astronomy, Cornell University

Abstract

CCAT is a sub-millimeter telescope to be built on Cerro Chajnantor in Chile near the ALMA site. The remote location means that all observing will be done by remote observers with the future goal of fully autonomous observing using a dynamic scheduler. The fully autonomous observing mode provides a natural means for accepting transient alert notifications for immediate follow up.

1 Introduction

CCAT [28, 7, 13] is a 25 m diameter sub-millimeter telescope to be built on Cerro Chajnantor in Chile near the ALMA site at an altitude of 5600 m. The additional height above the ALMA array results in significant improvements in transparency across all observing bands (Fig. 1), and 1.64 times better than the ALMA site at 350 μ m [23]. CCAT will initially operate from 350 μ m to 2 mm but will be capable in the future of operating at 200 μ m in the very best weather.

The CCAT project has identified four major first generation instruments to achieve its science goals. SWCam [25] will be the first-light camera having of order 60,000 detectors operating mainly at 350 μ m with additional detectors out to 2 mm. CHAI [8] will be a large format heterodyne array operating in two bands with the backend able to process spectra with a bandwidth of 4 GHz and 64,000 channels. LWCam [9] is a dedicated long-wave camera operating in 5-6 bands between 750 μ m and 2.1 mm with a long-wavelength goal of 3.3 mm. X-Spec [2] is a multi-object spectrometer with \sim 100 beams on the sky, each covering a frequency range of 190-520 GHz in two bands simultaneously with a resolving power of 400 – 700.

2 The Case for Transient Follow up

In the sub-millimeter, variable sources are sometimes monitored regularly as part of general observatory operations of flux calibrators [e.g. 16] and of pointing sources such

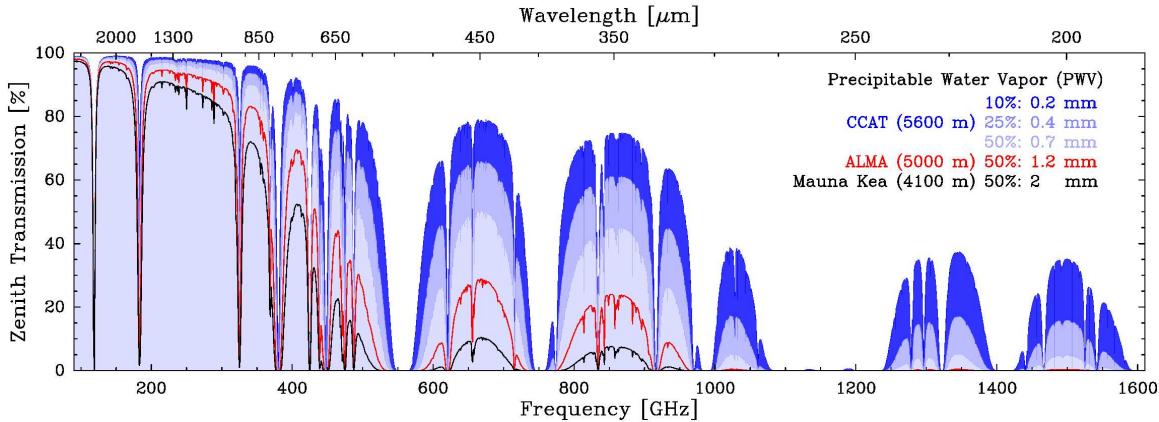


Figure 1: Atmospheric transmission for exceptional (10%), excellent (25%), and median conditions at CCAT and for median conditions at ALMA and at Mauna Kea [ATM model; 21]. Water vapor (PWV) distributions determined from $350\text{ }\mu\text{m}$ tipper measurements [23]. (figure credit: S. Radford).

as blazars [15]. Bright, time-varying sources do not generally require the ability to respond rapidly to time-sensitive alerts and can be observed as part of a monitoring program or as a general target of opportunity. Detecting the afterglows of gamma ray bursts (GRBs) in the sub-millimeter has proven to be difficult with the current generation of instrumentation [see e.g. 4] and the sooner that a telescope can get on target the more chance there is to see the peak of the light curve in the sub-millimeter. GRB 120422A [24] failed to detect any emission in the sub-millimeter despite being on source within 45 minutes and observing for nearly 2 hours with SCUBA-2 [11]. GRB 130427A, the brightest GRB in nearly 30 years [22], was not observed in the sub-millimeter but radio data and modeling suggests that the $850\text{ }\mu\text{m}$ flux would have been approximately 1 mJy after 2 days but more than 10 mJy if it had been observed within 4 hours of detection. First generation CCAT instruments such as SWCam will be able to observe an area of 0.15 sq deg to a depth of 1 mJy in only an hour in good weather. This is significantly better performance than current sub-millimeter instrumentation and indicates that the chances of detecting GRBs will increase considerably.

In addition to GRBs, LSST [12] will be coming online at around the same time as CCAT and will begin publishing millions of alerts per night. Some of these will be of interest to sub-millimeter astronomers and require reasonably fast follow up observations.

Once instrumentation has sufficient sensitivity to be useful, the main issue associated with time-sensitive alerts is how to respond to them in a timely manner. This is especially important for a common-user telescope designed for survey and P.I. observations.

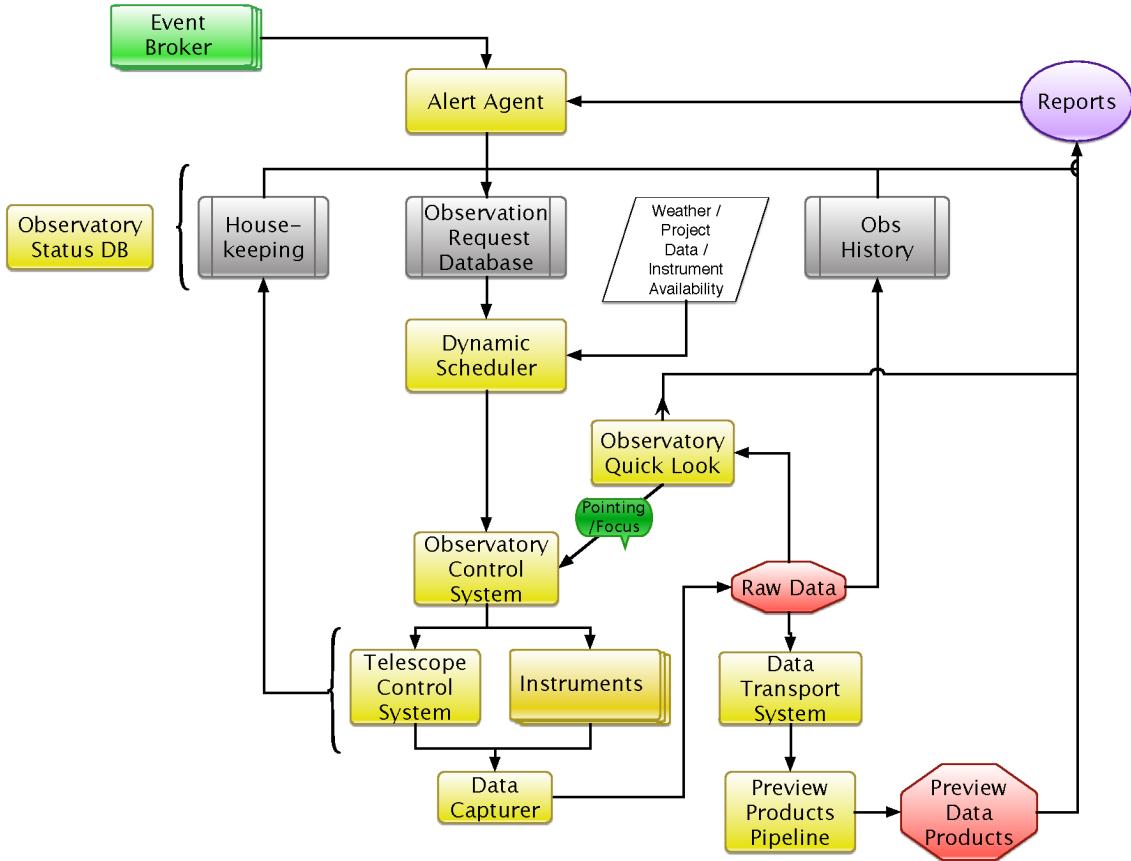


Figure 2: Flow chart of the CCAT software system from the perspective of processing transient alerts.

3 Reacting to Alerts

The CCAT observation scheduler will initially be a dynamic JIT ('just in time') scheduler determining the best observation block to observe at the current time. The system will be similar to that used by ALMA [18] and the James Clerk Maxwell Telescope [6]. A human operator, based either in San Pedro or at a remote observing location, will use the scheduler to guide the observing program and make the final choice of targets and associated calibrations.

The infrastructure being designed as part of the observation management system provides easy programmatic interface to the observation request database. The system architecture is shown in Fig. 2. The concept is that an alert broker, for example something like the ANTARES broker [20, 19], will send a VOEvent message [e.g. 27] to an alert agent. The alert agent will be run by an interested astronomer, possibly at their home institution. If the alert is of interest an observation will be submitted to

the observation request database at the telescope. Once this minimum scheduleable block has been submitted to the database the standard system will be used and the data will be processed in the normal way. Quality Assurance information and, possibly, flux measurements, will be fed back to the alert agent to allow the astronomer or agent to schedule follow up observations automatically.

This design is similar to that implemented at the UKIRT telescope [5, 14] which responded to a GRB alert within a few minutes [26] using the eSTAR system [1].

4 Post Commissioning

The goal, following telescope commissioning of the base system, is to upgrade the scheduler to fully autonomous operation [see e.g. 17, for background] where the observing queue will be monitored continuously and observations submitted as needed, calibrations will be scheduled when appropriate and observation blocks will be accepted or rejected automatically based on quality assurance data from the instrument pipelines. In the sub-millimeter it is sometimes the case that a flux calibrator will not be available until later in the night and so care must be taken to keep track of calibration data that are required for observations that have already been taken. This scheduling ability would allow time-sensitive followups to be inserted directly into the queue and observed without human intervention, similar to a fully robotic telescope such as LCOGT [3, 10].

Acknowledgments The CCAT Submillimeter Observatory (CCAT) is owned and operated by a consortium of universities and non-profit organizations located in the United States, Canada and Germany. Specifically the CCAT Consortium is comprised of: Cornell University, California Institute of Technology (Caltech), University of Colorado at Boulder, University of Cologne, University of Bonn, Dalhousie University, McGill University, McMaster University, University of British Columbia, University of Calgary, University of Toronto, University of Waterloo, University of Western Ontario and Associated Universities, Incorporated. The CCAT Engineering Design Phase was partially supported by funding from the National Science Foundation via AST-1118243.

References

- [1] Allan, A., Naylor, T., Steele, I. A., et al. 2004, in Proc. SPIE, Vol. 5496, Advanced Software, Control, and Communication Systems for Astronomy, ed. H. Lewis & G. Raffi, 313–322
- [2] Bradford, C., Hailey-Dunsheath, S., Shirokoff, E., et al. 2013, in American Astronomical Society Meeting Abstracts, Vol. 221, American Astronomical Society Meeting, #150.09
- [3] Brown, T. M., Baliber, N., Bianco, F. B., et al. 2013, PASP, 125, 1031
- [4] de Ugarte Postigo, A., Lundgren, A., Martín, S., et al. 2012, A&A, 538, A44
- [5] Economou, F., Jenness, T., Cavanagh, B., Adamson, A. J., & Allan, A. 2006, Astronomische Nachrichten, 327, 788
- [6] Economou, F., Jenness, T., Tilanus, R. P. J., et al. 2002, in ASP Conf. Ser., Vol. 281, Astronomical Data Analysis Software and Systems XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley, 488
- [7] Glenn, J., et al. 2013, in American Astronomical Society Meeting Abstracts, Vol. 221, American Astronomical Society Meeting, #150.06
- [8] Goldsmith, P., et al. 2012, Studying the Formation and Development of Molecular Clouds: with the CCAT Heterodyne Array Instrument (CHAI), http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/42645/1/12-2032_A1b.pdf
- [9] Golwala, S. R., et al. 2013, in American Astronomical Society Meeting Abstracts, Vol. 221, American Astronomical Society Meeting, #150.08
- [10] Hawkins, E., Baliber, N., Bowman, M., et al. 2010, in Proc. SPIE, Vol. 7737, Observatory Operations: Strategies, Processes, and Systems III., ed. D. R. Silva, A. B. Peck, & B. T. Soifer, 77370P
- [11] Holland, W. S., Bintley, D., Chapin, E. L., et al. 2013, MNRAS, 430, 2513
- [12] Ivezic, Z., Tyson, J. A., Acosta, E., et al. 2008, ArXiv e-prints, arXiv:0805.2366
- [13] Jenness, T., Brazier, A., Edwards, K., et al. 2014, in Astronomical Data Analysis Software and Systems XXIII, ed. N. Manset & P. Forshay, ASP Conf. Ser.(San Francisco: ASP), in press, arXiv:1401.8280
- [14] Jenness, T., & Economou, F. 2011, in Telescopes from Afar, ed. S. Gajadhar et al., 42, arXiv:1111.5855
- [15] Jenness, T., Robson, E. I., & Stevens, J. A. 2010, MNRAS, 401, 1240
- [16] Jenness, T., Stevens, J. A., Archibald, E. N., et al. 2002, MNRAS, 336, 14

- [17] Lampoudi, S., & Saunders, E. 2013, in 2nd International Conference on Operations Research and Enterprise Systems, 313–317, doi:10.5220/0004274203130317
- [18] Lucero, S., & Farris, A. 2007, in ASP Conf. Ser., Vol. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell, 673
- [19] Matheson, T. 2014, in Hot-Wiring the Transient Universe 3, ed. P. Wozniak, M. Graham, & A. Mahabal
- [20] Matheson, T., Saha, A., Snodgrass, R., & Kececioglu, J. 2014, in American Astronomical Society Meeting Abstracts, Vol. 223, American Astronomical Society Meeting, #343.02
- [21] Pardo, J. R., Cernicharo, J., & Serabyn, E. 2001, IEEE Transactions on Antennas and Propagation, 49, 1683
- [22] Perley, D. A., Cenko, S. B., Corsi, A., et al. 2014, ApJ, 781, 37
- [23] Radford, S. J. E. 2011, in Revista Mexicana de Astronomia y Astrofisica Conference Series, Vol. 41, Astronomical Site Testing Data in Chile, 87–90
- [24] Schulze, S., Malesani, D., Cucchiara, A., et al. 2014, ArXiv e-prints, arXiv:1401.3774
- [25] Stacey, G. J., Parshley, S., Nikola, T., et al. 2013, in American Astronomical Society Meeting Abstracts, Vol. 221, American Astronomical Society Meeting, #150.07
- [26] Tanvir, N., Levan, A., Kerr, T., & Wold, T. 2009, GRB Coordinates Network, 9202, 1
- [27] Williams, R. D., & Seaman, R. 2006, in ASP Conf. Ser., Vol. 351, Astronomical Data Analysis Software and Systems XV, ed. C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, 637
- [28] Woody, D., Padin, S., Chauvin, E., et al. 2012, in Proc. SPIE, Vol. 8444, Ground-based and Airborne Telescopes IV, 84442M

Data Triage of Astronomical Transients: A Machine Learning Approach

*Umaa Rebbapragada
Jet Propulsion Laboratory*

Abstract

This talk presents real-time machine learning systems for triage of big data streams generated by photometric and image-differencing pipelines. Our first system is a transient event detection system in development for the Palomar Transient Factory (PTF), a fully-automated synoptic sky survey that has demonstrated real-time discovery of optical transient events. The system is tasked with discriminating between real astronomical objects and bogus objects, which are usually artifacts of the image differencing pipeline. We performed a machine learning forensics investigation on PTF's initial system that led to training data improvements that decreased both false positive and negative rates. The second machine learning system is a real-time classification engine of transients and variables in development for the Australian Square Kilometre Array Pathfinder (ASKAP), an upcoming wide-field radio survey with unprecedented ability to investigate the radio transient sky. The goal of our system is to classify light curves into known classes with as few observations as possible in order to trigger follow-up on costlier assets. We discuss the violation of standard machine learning assumptions incurred by this task, and propose the use of ensemble and hierarchical machine learning classifiers that make predictions most robustly.

Lessons Learned and into the Future

Toward an Intelligent Event Broker: Automated Transient Classification

*Przemek Woźniak
Los Alamos National Laboratory*

Abstract

In order to succeed, the massive time-domain surveys of the future must automatically identify actionable information from the torrent of imaging data, classify emerging events, and optimize the follow-up strategy. To address this challenge, we are developing a fully autonomous, distributed event broker that will integrate cutting edge machine learning algorithms with high performance computing infrastructure. The talk will give an overview of this work and recent progress on image level variability detection and spectral classification using low resolution spectra.

How to Really Describe the Variable Sky

*Matthew Graham
California Institute of Technology*

Abstract

Whilst the classification of variable sources is one of the big challenges in the new era of time domain astronomy, it needs to be based on an effective characterization of the temporal behaviour of astronomical objects. Such attempts to date have been rather limited in scale and in scope. In this talk, we report on a systematic approach to describe the variability of 1.5 million CRTS sources in terms of the different kinds of phenomenology that they might exhibit. In particular, we are interested in ways in which characteristic timescales might reveal themselves and will address how this can aid object classification.

The Radio Transient Sky

Joseph Lazio

Jet Propulsion Laboratory / California Institute of Technology

Abstract

Radio transients are known on time scales from nanoseconds to years, from sources in the Galaxy and beyond, and with either coherent or incoherent emission mechanisms. Observations of this wide variety of sources are relevant to many of the highest profile questions in astronomy and astrophysics. As illustrations of the breadth of the radio transient sky, both coherent and incoherent radio emission has long been known from stars and stellar remnants and has informed topics ranging from stellar evolution to Galactic structure to relativistic jet dynamics to tests of fundamental physics. Coherent radio emission is now also known from brown dwarfs, and there are active programs to find similar emissions from extrasolar planets. Outside of the Galaxy, incoherent radio counterparts to supernovae, tidal disruption events, and gamma-ray bursts is well known and have contributed to topics such as understanding the cosmic star formation rate and the formation of relativistic jets. Excitingly, coherent radio bursts that appear to be at cosmological distances were recently discovered. I provide a survey of the radio transient sky, illustrating both how radio transients are part of the Hot-Wired Sky and are likely to help drive the Hot-Wiring. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Astrophysics in the Era of Massive Time-Domain Surveys

*George Djorgovski
California Institute of Technology*

Abstract

Synoptic sky surveys are now the largest data producers in astronomy, entering the Petascale regime, opening the time domain for a systematic exploration. A great variety of interesting phenomena, spanning essentially all subfields of astronomy, can only be studied in the time domain, and these new surveys are producing large statistical samples of the known types of objects and events for further studies (e.g., SNe, AGN, variable stars of many kinds), and have already uncovered previously unknown subtypes of these (e.g., rare or peculiar types of SNe). These surveys are generating a new science, and paving the way for even larger surveys to come, e.g., the LSST; our ability to fully exploit such forthcoming facilities depends critically on the science, methodology, and experience that are being accumulated now. Among the outstanding challenges, the foremost is our ability to conduct an effective follow-up of the interesting events discovered by the surveys in any wavelength regime. The follow-up resources, especially spectroscopy, are already and, for the predictable future, will be severely limited, thus requiring an intelligent down-selection of the most astrophysically interesting events to follow. The first step in that process is an automated, real-time, iterative classification of events, that incorporates heterogeneous data from the surveys themselves, archival and contextual information (spatial, temporal, and multiwavelength), and the incoming follow-up observations. The second step is an optimal automated event prioritization and allocation of the available follow-up resources that also change in time. Both of these challenges are highly non-trivial, and require a strong cyber-infrastructure based on the Virtual Observatory data grid, and the various astroinformatics efforts. Time domain astronomy is inherently an astronomy of telescope-computational systems, and will increasingly depend on novel machine learning and artificial intelligence tools. Another arena with a strong potential for discovery is a purely archival, non-time-critical exploration of the time domain, with the time dimension adding the complexity to an already challenging problem of data mining of highly-dimensional parameter spaces produced by sky surveys.

Poster Papers

Following up Fermi GBM Gamma-Ray Bursts

V. Connaughton, and M. S. Briggs
University of Alabama, Huntsville

A. Goldstein (NASA Postdoctoral Fellow)
NASA Marshall Space Flight Center

Abstract

The *Fermi* Gamma-Ray Burst Monitor (GBM) has been detecting 240 Gamma-Ray Bursts (GRBs) per year since 2008, 40-45 of them per year short GRBs. GBM is an all-sky transient monitor of the hard X-ray sky operating between 8 keV and 40 MeV. GBM localizes sources by triangulating the most likely source position based on observed count rates in detectors with different orientations to the sky. GRB locations are disseminated using GRB Coordinate Network (GCN) notices. We report here an analysis of over 300 GBM localizations for which more accurate positions are known. Systematic uncertainties of about $2 - 4^\circ$ affect about 90% of GBM localizations (68% confidence level), with larger systematic effects for the remaining 10%. These systematic components are added in quadrature to the statistical uncertainties of $1 - \sim 10^\circ$ and provided as probability maps to the follow-up community an hour or less after the GRB trigger. The intermediate Palomar Transient Factory (iPTF) has used these maps to detect three GRB afterglows using the GBM positional information.

1 The *Fermi* Gamma-Ray Burst Monitor

GBM is a collection of 14 uncollimated scintillators on-board the *Fermi* spacecraft [1]. Because of Earth occultation and passages through the South Atlantic Anomaly, each position on the sky is viewed by GBM with approximately a 50% duty cycle. Energy coverage from 8 - 1000 keV is provided by 12 Sodium Iodide (NaI) detectors, with two Bismuth Germanate detectors covering the range from 200 keV to 40 MeV. GBM triggers on a GRB when count rates in two or more of the NaI detectors significantly exceed background levels on one or more timescales from 16 to 4096 ms, usually in the 50 - 300 keV energy range. Other triggering energy ranges provide sensitivity to a variety of transient phenomena: solar flares and soft gamma-ray repeaters at the low end and terrestrial gamma-ray flashes at higher energies. The arrangement of GBM detectors on the *Fermi* spacecraft is shown in Figure 1. Source localization

is done by minimizing χ^2 on a grid of 41168 points on the sky, comparing the observed background-subtracted count rates in each of the 12 NaI detectors with model rates obtained by convolving the detector response with three representative GRB spectra. The spectrum that returns the lowest χ^2 in the minimization is assumed to best represent the GRB spectrum and the sky position minimizing χ^2 for that model spectrum is selected as the most likely position for the GRB. Locations are produced on-board, by the flight software, and on the ground, from data downlinked within seconds of the GRB trigger (ground-auto) or with a larger data set and human intervention (human-in-the-loop, or HitL). Flight software and ground-auto locations are distributed as GCN notices from 10– \sim 45 s following the trigger, with ground-auto locations being significantly more accurate owing to the availability of finer sky grid resolutions and more varied spectral shapes for model comparison on the ground. The HitL locations have latencies from 20 min to over an hour following the trigger and are also distributed as GCN notices. The suite of localizations, from flight software to human-in-loop, allows the follow-up community to optimize their observation strategy, with wide-field instruments, capable of rapid response but shallow coverage, more suited to chasing the automated locations, and more sensitive instruments that are less capable of covering large sky areas waiting for the HitL locations.

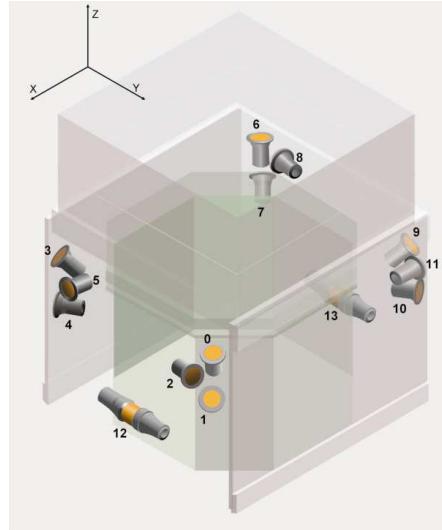


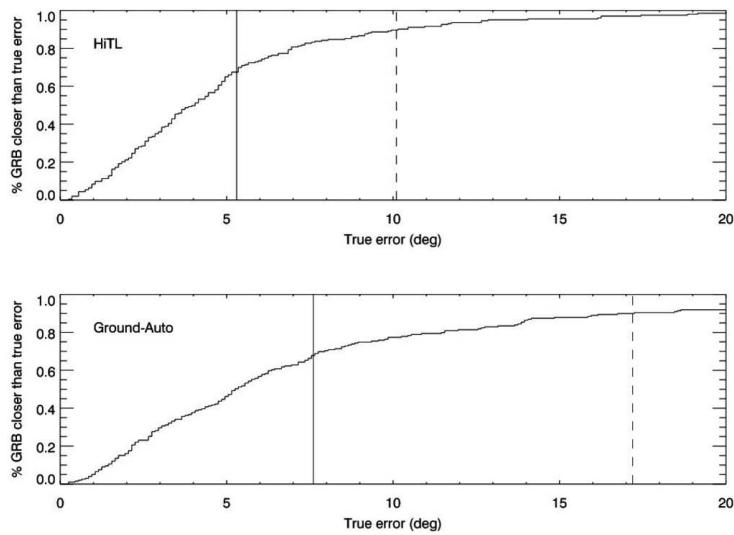
Figure 1: Arrangement of GBM detectors on *Fermi*. The 12 NaI (numbered 0 - 11) and 2 BGO (12 - 13) detectors are mounted on the +X and -X axes in the *Fermi* coordinate system. The Large Area Telescope (LAT) boresight is along the Z axis.

2 GBM location offsets from known positions

We use 203 GRBs localized by *Swift*, LAT, INTEGRAL, MAXI and SuperAGILE to analyze the accuracy of GBM GRB localizations. Figure 2 shows the fraction of GRBs as a function of the distance from the known position for the HitL and ground-

auto locations, regardless of the reported uncertainty on the localization. There can be several ground-auto locations, with updates for bursts showing brighter episodes following the initial localization. We use the last-reported ground-auto location, with the lowest statistical uncertainty. It can be seen that 68% of the ground-auto locations lie within 7.6° of the true position, with 90% within about 17° . The more refined HitL positions are significantly more accurate, with 68% within 5.2° of the true position and 90% within 10° , at the cost of a longer latency of at least 20 minutes. If we consider the reported statistical uncertainties, then for both the ground-auto and the HitL localizations, the 68% uncertainty regions contain the true position $\sim 40\%$ of the time and the 95% uncertainty regions 70% of the time. This implies there is a systematic component to the total uncertainty on the calculated burst position.

Figure 2: Accuracy of GBM localizations. The cumulative distributions show the fraction of GBM localizations as a function of offset in degrees from a known GRB location. The solid and dashed vertical lines show, respectively, the containment radius for 68% and 90% of the reference positions. The top panel is for the HitL locations, the bottom for the latest ground-auto location.

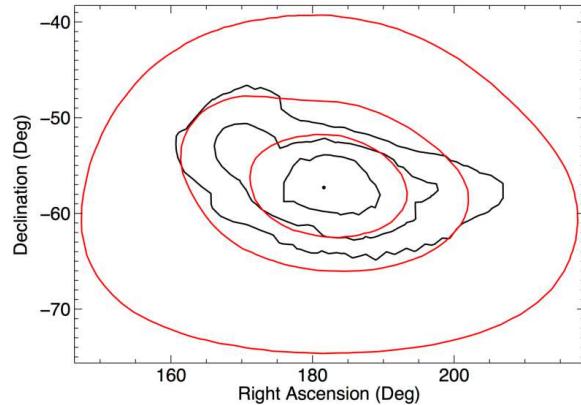


We augment the 203 reference positions with 9 intersecting annuli and 100 single annuli from InterPlanetary Network triangulations. A Bayesian approach used to characterize the systematic uncertainties in Burst And Transient Source Experiment (BATSE) GRB localizations [2] was then applied to the 312 GBM GRBs with reference locations. A preference for a model with two Gaussian components was found, with a core of about 90% of GRB localizations centered on $3.7^\circ \pm 0.2^\circ$ and a tail centered on $14.3^\circ \pm 2.5^\circ$. A core-plus-tail model was also favored for the BATSE GRB sample. We find a dependence for the systematic error in the core of the sample on the burst geometry in the spacecraft frame. Bursts from directions along the $\pm Y$ axes have smaller errors ($2.3^\circ \pm 0.4^\circ$) than those incident along the $\pm X$ axes ($4.1^\circ \pm 0.3^\circ$), with the fraction of GRBs in the core and the errors in the tail similar for both geometries [3].

3 Localization Contours

The core-plus-tail model has been convolved with the statistical uncertainty from the χ^2 -minimization process to produce probability maps for each GRB localized by GBM since January 2014. Figure 3 shows an example of the maps with the contours containing the 68%, 95%, and 99.8% probability for a GRB with a fairly large statistical uncertainty. These maps and the ASCII data for the probability contours that populate them are now uploaded to the *Fermi* Science Support Center upon production of the HitL location¹. Three GBM localization maps observed in tiling mode by iPTF resulted in afterglow detections and subsequently determined redshifts: GRB 130702A (GCN 14967), GRB 131011A (GCN 15324), and GRB 131231A (GCN 15653).

Figure 3: Probability map for the localization of GRB 080714745, with the statistical uncertainty contours (black) overplotted with the total uncertainty using the best model defined in [3]. The contours are 68%, 95%, and 99.8% confidence level regions around the χ^2 minimum from the GBM localization process.



¹http://heasarc.gsfc.nasa.gov/FTP/fermi/data/gbm/triggers/yyyy/bnnyymmddff/quicklook/glg_locplot_all_bnnyymmddff.png and [gll_loclist_all_bnnyymmddff.txt](http://heasarc.gsfc.nasa.gov/FTP/fermi/data/gbm/triggers/yyyy/bnnyymmddff/glg_loclist_all_bnnyymmddff.txt) where yyyy, mm, dd, fff are the year, month, day, and fraction of day for the GRB trigger, e.g., http://heasarc.gsfc.nasa.gov/FTP/fermi/data/gbm/triggers/2014/bn140122597/quicklook/glg_locplot_all.bn140122597.png.

References

- [1] C. A. Meegan et al., *Astrophys.J.* **702**, 791, (2009).
- [2] M. S. Briggs et al., *Astrophys.J.Supp.* **503**, 122 (1999).
- [3] V. Connaughton et al., submitted to *Astrophys.J.Supp.*

The TOROS Project

M.C. Diaz, and M. Benacquista

Center for Gravitational Wave Astronomy, The University of Texas at Brownsville

K. Belczyński

Warsaw University Observatory, Poland

M. Branchesi

Università di Urbino/INFN Sezione di Firenze, Italy

E. Brocato

Osservatorio Astronomico di Roma - INAF, Italy

D.L. DePoy

Texas A&M University

M.C. Díaz

Center for Gravitational Wave Astronomy, The University of Texas at Brownsville

M. Dominguez, and D. García Lambas

Universidad Nacional de Córdoba, Argentina

L.M. Macri, J.L. Marshall, and R.J. Oelkers

Texas A&M University

C.V. Torres

Center for Gravitational Wave Astronomy, The University of Texas at Brownsville

1 Searching for Optical counterparts of Gravitational Wave Sources

The first direct detection of gravitational waves (GW) may be possible with the Advanced Laser Interferometer Gravitational wave Observatory (aLIGO) and similar facilities within a few years. Binary neutron star (NS) mergers are expected to be among the most numerous and strongest GW sources [1].

The coalescence and merger of a neutron-star–neutron-star (NS-NS) or neutron-star–black-hole (NS-BH) binary is among the most energetic events in the Universe and has long been proposed as the process leading to short-hard gamma-ray

bursts (SGRBs) [2, 3]. NS-NS and NS-BH mergers are also some of the most promising candidates for producing gravitational-wave (GW) signals, detectable out to ~ 300 Mpc. Additionally it has been predicted [4, 5, 6] that the merger of NS-NS or NS-BH binaries should have an associated optical transient called a “macronova” or “kilonova”, powered by the radioactive decay of heavy nuclei synthesized in the merger ejecta through rapid neutron capture. It is also speculated that this mechanism may be the predominant source of stable r-process elements in the Universe (e.g., [7, 8]). These optical events are expected to have relatively low luminosities, with most of their emission at red and near-infrared wavelengths, and to last for a few days.

Recently reported optical and near-infrared observations of the transient associated with the short-duration gamma-ray burst GRB130603B [9, 10] exhibit the signatures predicted by “kilonova” models, indicating that this event could indeed confirm that compact-object mergers are the source of SGRBs. Kilonovae could then offer an alternative, unbeamed electromagnetic signature of the most promising sources for direct detection of gravitational waves. Simultaneous detection of electromagnetic signals associated with gravitational events seen by aLIGO and AdVIRGO could thus provide gravitational wave astronomers with crucial complementary information about these systems that is not directly accessible via gravitational waves.

But one key problem to consider is the following: GW detectors are all-sky detectors that allow localization via triangulation. They generally have poor pointing accuracy even for high signal-to-noise ratio events, leading to uncertainties of up to a few hundred square degrees in the target area for electromagnetic followup. Therefore, wide-field cameras and rapid follow-up observations will be crucial for the first EM counterpart detection. High quality, resolution and deep wide-field images are other crucial ingredients, which means that facilities with telescopes dedicated to the EM follow-up should be constructed at optimal locations.

In 2011, scientists from several European and North and South American institutions established a collaboration to develop such a facility, and called the project TOROS: Transient Optical Robotic Observatory of the South¹. In addition to the general scientific motivation outlined in the previous section, we were driven by the limited number of southern facilities with wide fields of view that would be capable of dedicated searches for aLIGO and adVIRGO released triggers during their first years of operation.

We decided that Cordón Macón, a mountain range in the province of Salta, Argentina ($W67^{\circ}19'41.6''$, $S24^{\circ}37'21.9''$, 4637m elev.) offered a unique opportunity for the development of a new astronomical facility. This site was first considered as a location to develop an astronomical facility by the European Southern Observatory, as part of their search for suitable locations for the European Extremely Large Telescope (E-ELT)².

¹<http://toros.phys.utb.edu>

²<http://www.eso.org/public/teles-instr/e-elt/>

2 TOROS characteristics

TOROS will have a primary mirror diameter of 0.6 m, a field-of-view of 9.85 sq.deg. and a very broad bandpass ($0.4 - 0.9 \mu\text{m}$, equivalent to a combination of the Sloan *griz* filters). It will be a fully robotic facility, driven by a priority-based intelligent agent/scheduler with four modes of operation in decreasing order of priority: (1) follow up of gravitational-wave triggers; (2) follow up of γ -ray burst triggers from *Fermi*, *Swift* and other missions; (3) baseline imaging of the entire surveyable area; (4) search for short-duration transient events, variable sources and moving objects within the DES and VISTA-VIKING fields.

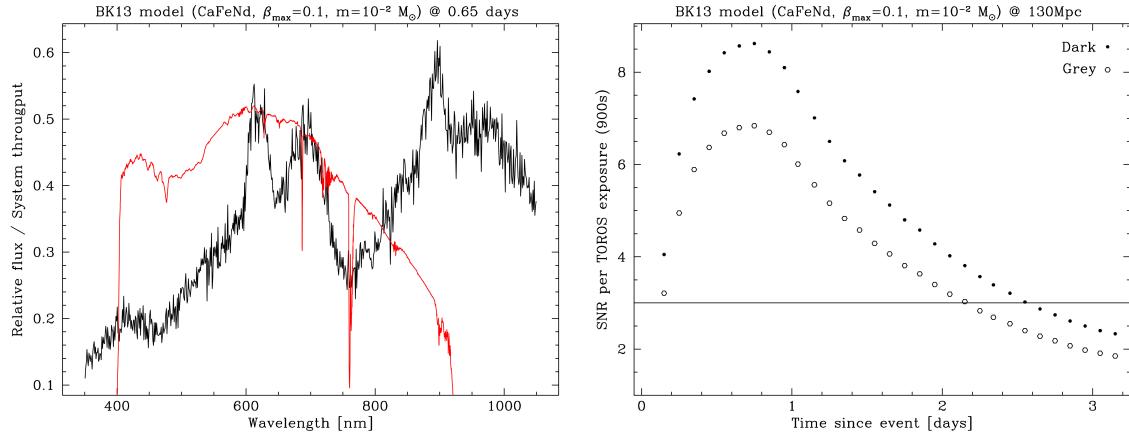


Figure 1: Left: Predicted spectrum at peak luminosity (in black) of the EM counterpart of a NS-NS merger, based on one of the latest models by Barnes & Kasen (priv. comm.); the expected total system throughput for TOROS is overplotted in red. The blue sensitivity will be helpful if the NS-NS merger produces a tidal tail of ^{56}Ni , as predicted by some models. Right: Expected SNR of an EM counterpart observed by TOROS as a function of time for the model on the left at the top-quartile distance predicted by the two-interferometer scenario of Kasliwal and Nissanke [12].

With TOROS, we expect to build a robust survey system that maximizes the probability of detecting the electromagnetic counterparts to GW events, taking into account: (a) the relatively poor localization of these events in the early years of aLIGO, prior to AdVIRGO operations, and (b) the large uncertainty in the expected optical/near-infrared luminosity and duration of these events [11]. The left panel of Figure 1, based on one of the latest models by Barnes & Kasen (priv. comm.), shows the predicted spectrum of the EM counterpart of a NS-NS merger at peak luminosity and the system throughput for TOROS.

Our survey strategy is based on the two- and three-interferometer coincident-trigger scenario (aLIGO only, 2015-2016; aLIGO+AdVIRGO, 2016/17 and beyond) considered by Kasliwal and Nissanke [12], based on simulations by Nissanke, Kasliwal

and Georgieva [13]. The most likely trigger rate for NS-NS mergers is $\sim 50 \text{ yr}^{-1}$, with a median localization area of 250 sq.deg. and $\langle D \rangle = 180 \text{ Mpc}$. The corresponding values for the top quartile of events are 170 sq.deg. and $\langle D \rangle = 130 \text{ Mpc}$. The addition of a third interferometer increases the event rate to $\sim 120 \text{ yr}^{-1}$ and significantly reduces the median localization area to 17 sq.deg. The area of extragalactic sky accessible to TOROS on any given night ($d < 35^\circ$, $|b| > 15^\circ$, elevation $> 30^\circ$ for > 3 hrs) ranges between $7 - 11 \times 10^3$ sq.deg. depending on the time of the year. This implies an observable trigger rate of $0.7 - 1.1 \text{ month}^{-1}$ for the two-interferometer scenario and $1.8 - 2.7 \text{ month}^{-1}$ for the three-interferometer case.

We plan to obtain 15-minute exposures so that we can fully cover the median localization area of the two-interferometer scenario even in the shortest nights of the year, while obtaining a sufficient SNR for the EM counterpart of a GW event at the median distance ($I = 21.7$ mag, based on the model shown in Fig. 1). We expect $4 - 5\sigma$ detections at peak magnitude under grey/dark sky conditions, improving to $7 - 9\sigma$ for the top quartile of events ($I = 21$ mag). Taking overheads into account, our system will be capable of covering the entire median localization area for the two-interferometer scenario (250 sq.deg. or 26 pointings) in 7 hours. Once AdVIRGO comes online (2016/17) we will be able to cover the median localization area in just two pointings, increasing our combined SNR at the end of night of imaging by $\sim 3.6\times$.

Except for the first four GW alerts from LVC, which will be governed by the aforementioned MoU, we will promptly release coordinates of potential transients and stacked images to the entire astronomical community. Additionally, we plan to execute our own photometric and spectroscopic followup.

References

- [1] J. Abadie *et al.* [LIGO Scientific and Virgo Collaborations], Class. Quant. Grav. **27**, 173001 (2010) [arXiv:1003.2480 [astro-ph.HE]].
- [2] D. Eichler, M. Livio, T. Piran and D. N. Schramm, Nature **340**, 126 (1989).
- [3] R. Narayan, B. Paczynski and T. Piran, Astrophys. J. **395**, L83 (1992) [astro-ph/9204001].
- [4] L. -X. Li and B. Paczynski, Astrophys. J. **507**, L59 (1998) [astro-ph/9807272].
- [5] A. Abramowski *et al.* [H.E.S.S. and VERITAS Collaborations], Astrophys. J. **746**, 151 (2012) [arXiv:1111.5341 [astro-ph.CO]].
- [6] E. Berger, arXiv:1311.2603 [astro-ph.HE].
- [7] C. Freiburghaus, S. Rosswog and F.-K. Thielemann, Astrophys. J. **525**, L121 (1999)
- [8] S. Goriely, A. Bauswein and H. -T. Janka, arXiv:1107.0899 [astro-ph.SR].
- [9] N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. Hjorth, K. Wiersema, R. Tunnicliffe and A. d. U. Postigo, arXiv:1306.4971 [astro-ph.HE].
- [10] E. Berger, W. Fong and R. Chornock, Astrophys. J. **774**, L23 (2013) [arXiv:1306.3960 [astro-ph.HE]].
- [11] J. Barnes and D. Kasen, Astrophys. J. **775**, 18 (2013) [arXiv:1303.5787 [astro-ph.HE]].
- [12] M. M. Kasliwal and S. Nissanke, arXiv:1309.1554 [astro-ph.HE].
- [13] S. Nissanke, M. Kasliwal and A. Georgieva, Astrophys. J. **767**, 124 (2013) [arXiv:1210.6362 [astro-ph.HE]].

Program

HTU-III Scientific Program

All talks allocated for 20 minutes including questions

DAY 1: WEDNESDAY, NOVEMBER 13

Coffee (7:30)

Welcome (8:00) P. Woźniak

SESSION 1

Massively parallel time-domain astrophysics: Challenges & opportunities
chair: M. Graham

R. Seaman	Autonomous infrastructure for observatory operations
T. Vestrand	The follow-up crisis: optimizing science in an opportunity-rich environment
J. Richards	Machine learning for time-domain discovery and classification
S. Ridgway	The variable sky
A. Kowalski	Hot-Wiring flare stars: Optical flare rates and properties from time- domain surveys

Break (20 minutes)

SESSION 2 (10:20)

Time-domain surveys I: Transient searches

chair: R. Street

J. Kantor	Transient alerts in LSST
E. Bellm	The Zwicky Transient Facility (ZTF)
R. Scalzo	SkyMapper and supernovae
A. Drake	The Catalina Real-time Transient Survey (CRTS)
K. Chambers	Pan-STARRS transients, recent results, future plans

LUNCH (12:00)

SESSION 2 (cont. at 1:20)

N. Walton	GAIA — Revealing the transient sky
H. Campbell	Transient astronomy with GAIA

SESSION 3 (2:00)

Nuts and bolts I: Telescope networking

chair: T. Matheson

- | | |
|--------------|---|
| F. Hessman | Time to revisit the "Heterogeneous Telescope Network" |
| S. Barthelmy | GCN/TAN: past, present & future: serving the transient community's need |
| J. Swinbank | VOEvent — Where we are, where we're going |

Break (20 minutes)

- | | |
|-------------|--|
| J. Fay | Time series data visualization in World Wide Telescope |
| P. Kubanek | RTS2 and BB — network observations |
| E. Saunders | Multi-telescope observing: the LCOGT network scheduler |

SESSION 4 (4:20)

Time-domain surveys II: Moving objects and exo-planets

chair: L. Walkowicz

- | | |
|----------------|--|
| A. Mainzer | Small body populations according to NEOWISE |
| E. Christensen | The Catalina Sky Survey for near-Earth objects |

GROUP DISCUSSION I (5:00)

**The science of transients and variable stars leading to the era of LSST,
moderated by L. Walkowicz and M. Kasliwal**

DINNER at Tomasitas (5:30 7:30)

Parallel evening breakout sessions (7:30 9:30)

- A) VOEvent/IVOA time domain interest group (organized by J. Swinbank and M. Fitzpatrick)

DAY 2: THURSDAY, NOVEMBER 14

Coffee (7:30)

SESSION 4 (cont. at 8:00)

Time-domain surveys II: Moving objects and exo-planets

chair: L. Walkowicz

S. Howell	High-precision time series photometry, what have we learned?
J. Jenkins	Passing NASA's planet quest baton from Kepler to TESS
L. Denneau	The ATLAS all-sky survey
L. Allen/D. Trilling	The performance of MOPS in a sensitive search for near-Earth asteroids with the Dark Energy Camera

SESSION 5 (9:20)

Time-domain surveys III: Beyond optical photometry

chair: T. Vestrand

B. Dingus	The TeV sky observed by the High-Altitude Water Cherenkov Observatory (HAWC)
S. Nissanke	Hearing & seeing the violent universe

Break (20 minutes)

R. Williams	Follow-up of LIGO-Virgo observations of gravitational waves
L. Singer	The needle in the hundred-square-degree haystack: from Fermi GRBs to LIGO discoveries
J. van Leeuwen	ARTS — the Apertif Radio Transient System
D. Frail	Radio adventures in the time domain
S. Myers	The Karl G. Jansky VLA Sky Survey (VLASS): Defining a new view of the dynamic sky
C. Law	VLA Searches for Fast Radio Transients at 1 TB/hour

LUNCH (12:20)

SESSION 6 (1:40)

Nuts and bolts II: Algorithms and event brokers

chair: J. Swinbank

- | | |
|------------|--|
| A. Mahabal | Novel measures for rare transients |
| J. Bloom | The modern automated astrophysics stack |
| M. Fraser | Near real-time discovery and classification of supernovae:
Lessons from PESSTO and prospects for GAIA |
| J. Scargle | Time Series Explorer |

Break (20 minutes)

SESSION 6 (cont. at 3:20)

- | | |
|-------------|---|
| T. Matheson | ANTARES — Arizona-NOAO Temporal Analysis and Response to Events System |
| S. Kumar | Bayesian classification of Pan-STARRS Medium Deep Transients |
| S. Barway | The South African Astro-informatics Alliance (SA3) |
| A. Miller | Predicting fundamental stellar parameters from photometric light curves |
| A. Becker | State-based models for light curve classification |

GROUP DISCUSSION II (5:00)

Technology and infrastructure for time domain astronomy in the third millennium, moderated by R. Williams & P. Woźniak

DINNER (5:30 7:30)

Parallel evening breakout sessions (7:30 9:30)

- C) Coordinated brokering of transient events: toward a unified follow-up system for LSST (organized by T. Matheson and J. Bloom)
- D) Supernova and GRB science (organized by M. Fraser)

DAY 3: FRIDAY, NOVEMBER 15

Coffee (7:30)

SESSION 7 (8:00)

Follow-up science, opportunities and strategies

chair: A. Mahabal

- | | |
|------------|--|
| D. Kopač | Gamma-Ray Bursts with early time optical emission |
| D. Perley | Burst of the century? A case study of the afterglow of nearby ultra-bright GRB 130427A |
| F. Millour | Optical interferometry and adaptive optics of bright transients |
| N. Butler | Multi-color robotic observations with RATIR |
| D. Sand | The robotic FLOYDS spectrographs |

Break (20 minutes)

- | | |
|-----------------|--|
| R. Street | Dynamic follow-up of transient events with the LCOGT robotic telescope network |
| I. Arcavi | Rapid followup in iPTF and the science it enables |
| T. Jenness | Remote operations and transient alert follow-up in CCAT |
| U. RebbaPragada | Data triage of astronomical transients: A machine learning approach |

Break (20 minutes)

SESSION 8 (11:40)

Lessons learned and into the future

chair: R. Seaman

- | | |
|---------------|--|
| P. Woźniak | Toward an intelligent event broker: automated transient classification |
| M. Graham | How to really describe the variable sky |
| J. Lazio | The radio transient sky |
| G. Djorgovski | Astrophysics in the era of massive time-domain surveys |

GROUP DISCUSSION III (1:00)

What comes next, moderated by J. Bloom

ADJOURN (1:30)

Participants

HTU-III Participants

Lori	Allen	NOAO
Iair	Arcavi	Weizmann Institute of Science / LCOGT
Scott	Barthelmy	NASA-GSFC
Sudhanshu	Barway	South African Astronomical Observatory
Dana	Baylis	New Mexico Tech
Andrew	Becker	University of Washington
Eric	Bellm	Caltech
Joshua	Bloom	UC Berkeley
Nathaniel	Butler	ASU
Heather	Campbell	Institute of Astronomy, Cambridge
Kenneth	Chambers	Institute for Astronomy, University of Hawaii
Eric	Christensen	The University of Arizona
Valerie	Connaughton	Univ. Alabama Huntsville
Mark	Cornell	MIT Lincoln Laboratory
Larry	Denneau	Institute for Astronomy, University of Hawaii
Mario	DIAZ	University of Texas at Brownsville
Brenda	Dingus	LANL
George	Djorgovski	Caltech
Andrew	Drake	Caltech
Jonathan	Fay	Microsoft Research
Mike	Fitzpatrick	NOAO
Dale	Frail	National Radio Astronomy Observatory (NRAO)
Morgan	Fraser	Institute of Astronomy, Cambridge
Matthew	Graham	Caltech
Xiaoyue	Guan	LANL
Frederic	Hessman	Univ. of Goettingen
Eric	Hooper	WIYN Observatory and University of Wisconsin-Madison
Steve	Howell	NASA Ames Research Center
Maohai	Huang	National Astronomical Observatories, China
Jon	Jenkins	SETI Institute
Tim	Jenness	Cornell University
Lynne	Jones	University of Washington
Jeffrey	Kantor	LSST
Mansi	Kasliwal	Carnegie Institution for Science
Khalid	Kayal	The University of Texas at Brownsville
Drejc	Kopač	University of Ljubljana, Faculty of Mathematics and Physics
Adam	Kowalski	NASA-GSFC
Petr	Kubanek	Institute of Physics
Sidharth	Kumar	University of Maryland
Casey	Law	UC Berkeley
Earl	Lawrence	LANL

Joseph	Lazio	JPL, Caltech
Ashish	Mahabal	Caltech
Amy	Mainzer	JPL
Thomas	Matheson	NOAO
Jose	McKinnon	The University of Texas at Brownsville
Adam	Miller	JPL/Caltech
Florentin	Millour	Lagrange
Daniela	Moody	LANL
Steven	Myers	NRAO
Samaya	Nissanke	Caltech
Knut	Olsen	NOAO
Etienne	Pallier	IRAP
Alin	Panaiteescu	LANL
Tania Marcela	Penuela Gomez	LMU-Munich
Daniel	Perley	Caltech
Yulei	QIU	National Astronomical Observatories, China
Umaa	Rebbapragada	Jet Propulsion Laboratory
Joseph	Richards	Lawrence Berkeley National Lab
Stephen	Ridgway	NOAO
David	Sand	Texas Tech University
Eric	Saunders	LCOGT
Richard	Scalzo	Australian National University
Jeff	Scargle	San Jose State University
Leo	Singer	California Institute of Technology
Rob	Seaman	NOAO
Rachel	Street	LCOGT
John	Swinbank	University of Amsterdam
David	Trilling	Northern Arizona University
Joeri	van Leeuwen	ASTRON/U. Amsterdam
Scott	Vander Wiel	LANL
Tom	Vestrand	LANL
Veljko	Vujcic	Astronomical Observatory Belgrade
Lucianne	Walkowicz	Princeton University
Nicholas	Walton	Institute of Astronomy, University of Cambridge, UK
Thomas	Wevers	RU Nijmegen
Roy	Williams	Caltech
Przemek	Woźniak	LANL
James	Wren	LANL
Chao	WU	National Astronomical Observatories, China

Author Index

Author Index

- Allen, L., 67
Arcavi, I., 197
Barthelmy, S., 103
Barway, S., 159
Becker, A., 163
Belczyński, K., 225
Bellm, E., 27
Benacquista, M., 225
Benetti, S., 131
Blagorodnova, N., 43
Bloom, J., 129
Bower, G.C., 85
Bramich, D.M., 189
Branchesi, M., 225
Briggs, M.S, 219
Brocato, E., 225
Burke-Spoloar, S., 85
Butler, B., 85
Butler, N., 185
Campbell, H., 43
Chambers, K., 39
Chesneau, O., 177
Christensen, E., 55
Connaughton, V., 219
Díaz, M.C., 225
Denneau, L., 65
DePoy, D.L., 225
Diaz, M.C., 225
Dingus, B., 71
Djorgovski, G., 215
Dominguez, M., 225
Dominik, M., 189
Drake, A., 37
Fay, J., 113
Frail, D., 81
Fraser, M., 43, 131
Garcia Lambas, D., 225
Gilmore, G., 43
Goldstein, A., 219
Gomboc, A., 167
Graham, M., 211
Hessman, F.V., 95
Hodgkin, S., 43
Horne, K.D., 189
Howell, S. B., 57
Howell, S.B., 9
Hundertmark, M.P.G., 189
Inserra, C., 131
Jaimes, R.F., 189
Japelj, J., 167
Jenkins, J., 63
Jenness, T., 199
Kantor, J., 19
Kececioglu, J., 145
Kopač, D., 167
Koposov, S., 43
Kowalski, A., 15
Kubanek, P., 115
Kumar, S., 151
Lampoudi, S., 117
Law, C.J., 85
Lawrence, E., 85
Lazio, J., 213
Lazio, T.J.W., 85
Macri, L.M., 225
Mahabal, A., 127
Mainzer, A., 53
Marshall, J.L., 225

- Matheson, T., 9, 145
Mattmann, C., 85
Meiland, A., 177
Mighell, K.J., 9
Miller, A., 161
Millour, F., 177
Myers, S., 83
Nardetto, N., 177
Nissanke, S., 73
Oelkers, R.J., 225
Olsen, K.A., 9
Pastorello, A., 131
Perley, D., 175
Rebbapragada, U., 205
Richards, J., 7
Ridgway, S.T., 9
Rupen, M., 85
Saha, A., 145
Sand, D., 187
Saunders, E.S., 117
Scalzo, R., 35
Scargle, J., 137
Seaman, R., 3
Siemion, A., 85
Singer, L., 77
Smartt, S., 131
Smith, K., 131
Snodgrass, C., 189
Snodgrass, R., 145
Street, R.A., 189
Sullivan, M., 131
Swinbank, J.D., 105
Taubenberger, S., 131
Torres, C.V., 225
Trilling, D., 67
Tsapras, Y., 189
van Leeuwen, J., 79
VanderWiel, S., 85
Vestrand, T., 5
Walton, N., 41, 43
Williams, R., 75
Woźniak, P., 209
Young, D., 131

