Astronomy ESFRI & Research Infrastructure Cluster

ASTERICS - H2020 - 653477

Report on Algorithms and Standard Interfaces for Cross-Facility Scheduling

ASTERICS GA DELIVERABLE: D5.9

Document identifier:	ASTERICS-D5.9-final.pdf
Date:	**18 June 2018**
Work Package:	**WP5**
Lead Partner:	**STFC**
Document Status:	**Report**
Dissemination level:	**Public**
Document Link:	www.asterics2020.eu/documents/ASTERICS-D5.9-final.pdf

Abstract

Efficient observation planning and scheduling are critical elements for optimizing the exploitation of large facilities and space missions devoted to astronomical research. Many projects have developed and used scheduling tools with various algorithms, although this is still not common practice in classical observatories. This report describes a hierarchy of scheduling applications, ranging from a single telescope up to multi-observatory coordination, currently the most complex optimization problem.

Scheduling algorithms increase the return of scientific programmes by optimizing the operation of facilities as they follow the specified observation strategies. Scheduling tools are especially important for infrastructures where the operation is very complex, e.g. observatories with sub-arrays or multi-observatory coordination, or where programmes are heterogeneous, time critical, or require fast reaction to changing conditions, such as weather or transient events. Otherwise, the result is an inefficient operation or a very high cost in specialized human resources. New scientific cases are basically not affordable without using this kind of AI tool. Multi-messenger science is an example of the challenges that face efficient operation of large facilities, and we present an analysis of how to promote multi-messenger science through: building collaborations, sharing observatory activities and schedules, and providing AI tools for efficient local and global scheduling.

A scheduling framework called STARS is presented. It includes an abstraction of the scheduling problem and different algorithms (GA, MOEA, heuristics) that can be used at a range of facilities and for multi-observatory coordination.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/WP	Date
From	J. Colomé	IEEC/WP5	1/6/2018
Author(s)	A. Bridger J. Lightfoot A.Garcia-Piquer J. Colomé E.Oña D.F.Torres E.Díez P.Colomer	ATC/STFC ATC/STFC IEEC IEEC IEEC IEEC GTD GTD	1/6/2018
Reviewed by	A. Bridger J. Colomé	ATC/STFC IEEC	1/6/2018
Approved by	A. Bridger	STFC	1/6/2018
Approved by	Rob van der meer	ASTRON	18/6/2018

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	30 May 2018	First version	A. Bridger/STFC J. Lightfoot/STFC A.Garcia-Piquer/IEEC J. Colomé/IEEC E.Oña/IEEC D.F.Torres/IEEC E. Diez/GTD P. Colomer/GTD
2	1 June 2018	Final version	A. Bridger/STFC J. Lightfoot/STFC A.Garcia-Piquer/IEEC J. Colomé/IEEC E.Oña/IEEC D.F.Torres/IEEC E. Diez/GTD P. Colomer/GTD

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the ASTERICS "Document Management Procedure" will be followed: https://wiki.asterics2020.eu/wiki/Procedures

VI. TERMINOLOGY

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

A glossary of terms specific to this paper are given below:

- AAVSO - American Association of Variable Star Observers.
- AGN - Active Galactic Nuclei.
- ALMA - The Atacama Large Millimeter Array.

- AMON - The Astrophysical Multimessenger Observatory Network. www.amon.psu.edu
- ASAS-SN – A synoptic survey project using cameras to search for supernovae.
- ASTERICS - Astronomy ESFRI and Research Infrastructure Cluster.
- ATel - Astronomer's Telegram.
- BlackGEM - Three robotic 65cm telescopes dedicated to the follow-up of LIGO/VIRGO targets. astro.ru.nl/blackgem.
- CARMENES - A dual spectrograph on the 3.5m telescope at Calar Alto Observatory.
- Chandra – an X-ray satellite observatory.
- CLEOPATRA - Connecting Locations of ESFRI Observatories and Partners in Astronomy for Timing and Real-time Alerts. One of the five work packages in ASTERICS.
- CR - Cosmic Ray.
- CTA - The Cherenkov Telescope Array.
- DDT - Director's Discretionary Time
- ELT - ESO Extremely Large Telescope.
- EM - Electro-magnetic
- ESFRI - European Strategy Forum on Research Infrastructures.
- Fermi - γ-ray satellite.
- Gaia – astrometry satellite.
- GCN/TAN - GRB Coordinates Network / Transient Astronomy Network.
- GRB - Gamma-Ray Burst.
- GW - Gravity Wave.
- HST - Hubble Space Telescope.
- KM3NeT - neutrino detector.
- LCOGT - Las Cumbres Observatory Global Telescope.
- LIGO - The Laser Interferometer Gravitational-Wave Observatory.
- LT - The Liverpool Telescope. A 2m aperture robotic telescope operating on La Palma. telescope.livjm.ac.uk
- LSST - The Large Synoptic Survey Telescope (www.lsst.org). This will be an 8m telescope capable of imaging the entire visible sky every few nights. Up to 10 million transient alerts are expected to be generated each night.
- The `Marshall' - A bespoke web-app used by the PESSTO project to coordinate follow-up of transient targets to be classified.
- MeerKAT - A radio telescope array in South Africa. www.ska.ac.za.
- MeerLICHT - A 65cm optical telescope that will observe the same field as MeerKAT at all times. www.meerlicht.org.
- MoU - Memorandum of Understanding.
- NRAO - National Radio Astronomy Observatory.
- NTT - The ESO New Technology Telescope.
- OMS - Observation Management System.
- OPTICON – H2020 project promoting EU astronomy.
- Pan-STARRS – The Panoramic Survey Telescope and Rapid Response System, a synoptic survey.
- PESSTO / ePESSTO - The Public ESO Spectroscopic Survey of Transient Objects was an ESO long-term project allocated time on the NTT telescope for the spectroscopy of transient targets allowing them to be classified. ePESSTO is a successor long-term allocation.
- SB - Scheduling Block, a generic description of a proposed observation.
- Scheduler - Process that selects the next SB to be executed by each facility.

- SKA - The Square Kilometer Array.
- SNe - Supernovae.
- SNeX - The Supernova Exchange. supernova.exchange/public
- Swift - The SWIFT Gamma Ray Burst mission. SWIFT is a satellite that monitors parts of the sky for Gamma Ray Bursts.
- TAC - Time Allocation Committee.
- TAT - Transient and Alert Team. A group within PESSTO tasked with triaging incoming alerts.
- TDE - Tidal Disruption Event, a star is pulled apart by tidal forces as it falls into a black hole.
- TJO - Telescope Joan Oró, a 1m class robotic telescope operating at the Observatori Astronòmic del Montsec in Catalunya, Spain.
- ToO - Target of Opportunity.
- VIRGO – European laser interferometer gravity wave detector.
- VISTA – Visible and Infrared Survey Telescope for Astronomy. An ESO survey instrument.
- VLT - The ESO Very Large Telescope(s) at Paranal.
- VOEvent - A standardized language for reporting astronomical events (http://www.ivoa.net/documents/VOEvent).

VII. PROJECT SUMMARY

The EU funds a number of astronomical facilities that are members of the `European Strategy Forum for Research Infrastructures', ESFRI. The `ASTtronomy Esfri and Research Infrastructure CluSter', ASTERICS, is a €15 million project funded by the European Horizon 2020 framework, which aims to address the cross-cutting synergies and common challenges shared by the various astronomy ESFRI facilities (e.g SKA, CTA, KM3NeT and ELT).

CLEOPATRA, that is `Connecting Locations of Esfri Observatories and Partners in Astronomy for Timing and Real-time Alerts', is one of the five work packages in ASTERICS. Its aim is to develop scheduling schemes that maximise the scientific gain from the facilities. The problem domain ranges from scheduling multi-frequency, multi-messenger observations using several facilities to the scheduling of complex, many-element detector arrays at a single facility.

VIII. EXECUTIVE SUMMARY

The scope of this document includes all aspects of coordinated and multi-messenger observing, from the formation of the observing collaboration, through how the collaboration works, to the efficient scheduling of the planned observations at the facilities.

First, a description is given of the classical observing process on single facilities, noting features that are worth preserving. Next, the recent development of multi-messenger astronomy is outlined. Case studies of successful coordinated observing campaigns are presented and analysed. The process can be split into two parts: `observation planning' where the coordinated observations needed for the science are specified, and `observation scheduling' where these observations are executed efficiently at each facility.

Observation planning is itself split into two parts: the formation of the observing collaboration, and the working of the collaboration. How collaborations form is discussed; we suggest that Facebook could be used to make the process easier or a recently developed tool called SMARTNet. Efficient collaborations often communicate internally and order their process using web apps; we suggest that a framework and tools be developed to make the construction of such things easier.

Results are presented of investigations into candidate algorithms for use in facility observation scheduling. We describe a hierarchy of scheduling applications, ranging from a single telescope up to the multi-observatory coordination, currently the most complex problem for optimization. Algorithms have to be designed to suit the operational design of the facilities and the specific observation strategies of the scientific programmes. A scheduling framework called STARS is presented, which includes an abstraction of the scheduling problem, and different algorithms (GA, MOEA, heuristics) that can be used at different facilities.

LESSONS LEARNED

Task 5.4 is led by J. Colome (IEEC) and involves teams at ATC/STFC, IEEC and GTD. The task's activities were devoted to researching how the SKA and CTA could maximize their science return with AI scheduling solutions, and carry out programmes in a coordinated manner to do multi-messenger science. Scheduling solutions were explored and applied to different test cases in order to extract conclusions that would help extend their applicability to other ground and space-based facilities for space science research (e.g. LIGO, E-ELT, ALMA, ESA missions). Task 5.4 also incorporates multi-messenger astrophysics at the level of promoting consortium and individual collaborations, sharing programmes and schedules, and developing scheduling tools that optimize multi-observatory observations.

Several lessons have been learned so far and are listed hereafter:

- *Networking*: There have been meetings to promote the collaboration of the task partners, and also with partners in other tasks (e.g. task 5.2) to find synergies. In addition, ASTERICS has enabled participation in external meetings to disseminate the project and task activities. The latter has had a significant value, helping to collect experiences from various large facilities, and initiating steps towards the transference of the task outcomes to others. For instance, communication was initiated with teams involved in the scheduling tools for ALMA, LSST and ESO, in addition to the forum already available within ASTERICS for the SKA and CTA teams. Specifically, it is worth mentioning our participation in the following meetings: the SPIE meeting on Astronomical Telescopes and Instrumentation in Edinburgh in June 2016, the CLEOPATRA meeting on Transient Alerts in Amsterdam in September 2017, the ADASS XXVII conference (Chile, October 2017) where there was a special session on `Astronomical Scheduling in the Era of Big Observatories Scheduling', and the 6th International Conference on Space Mission Challenges for Information Technology (Alcalá de Henares, Spain, September 2017). We would also like to highlight two additional meetings that will take place in the coming months: the ESO conference on `Proposal Handling tools' and the SPIE meeting on `Astronomical Telescopes and

Instrumentation 2018' that will let us disseminate task results to audiences who will play a strategic role in future actions related to this task. The community of large facilities is concerned about the need to improve planning and scheduling procedures and we expect that the new contacts made could translate into future collaborations that will continue task 5.4 activities.

- *AI technologies, common features among facilities, and the importance of the partnership*: Research on AI technologies shows that Genetic Algorithms perform far better than Neural Networks, by a factor of 4 – 5. An additional increase in efficiency by a factor of 1.5 could be achieved by the use of Multi-Objective Evolutionary Algorithms (MOEA). Care is being taken to ensure that the solutions found are not limited to the CTA or SKA. Experience on several projects is shared among the different partners and the role played in SKA and CTA, in particular, has contributed to a productive collaboration, extracting common solutions for task execution that can be applied to other observatories. Constraints treated simultaneously are: the maximum time allocated for on-source tracking, minimum time spent on slewing, maximum number of completed programmes and their relative priority.
- *Scientific vs technological skills*: the 5.4 team is made up of technology experts on control software with huge experience in the operation of large astronomical facilities, and also researchers in astrophysics with expertise on operational models for observatories. This situation has avoided biasing the solution for efficient scheduling and multi-messenger strategies towards a pure technological solution and, at the same time, has helped focus on the conditions for managing astronomical observatories that ensure the expected scientific return. For instance, the observation strategies associated with different science cases have been modeled and incorporated in the proposed scheduling tools, and scientific use cases have been analyzed to select and work on those scenarios that will increase the return to the observatories. In particular, multi-observatory scheduling has been analyzed to validate the interest of the community, not only in transient follow-up, but in the coordinated execution of strategic surveys (results will be presented in deliverable 5.12). Therefore, the complementary skills that exist in the partnership have played a key role in the results obtained so far and presented in this document.
- *Risk from the maturity of the CTA and SKA projects*: The level of maturity of the CTA and SKA was identified as a risk for the proper execution of the task. Some of the tests were planned to be done using existing software for scheduling proposals. It was also required that the scientific objectives be defined to work with realistic test cases and demonstrate good performance of the technical solutions. It was clear that the risk was real and had to be mitigated. This was done by using non-consolidated definitions of the scientific programmes (in CTA and SKA), by using other facilities and science cases to check a subset of the proposed functionalities, by using operational conditions in precursor facilities (GASKAP for SKA) and by using an existing simulation platform at IEEC for testing (this platform was extended to test the multi-observatory coordinated scheduling, as is going to be described in deliverable 5.12). The available algorithms at IEEC, in particular the MOEA, were also adapted to cover the new optimization problem.
- *Overlapping interests with task 5.2*: Multi-messenger science including transient alerts is a key topic for multi-observatory scheduling. We realized there were important synergies with task 5.2 from the very beginning so the IEEC team participated in some discussions within that task. This resulted in the participation of the IEEC and ATC/STFC teams at the Transient Alert workshop organized by task 5.2, where a dedicated session was allocated to show the results obtained in task 5.4. The overlap

between the tasks is not sufficient to justify merging them, but there are clear benefits in following each other's progress and looking for common approaches. Interaction should have been promoted in an earlier stage and will happen in the future.

IX. TABLE OF CONTENTS

X. THE CONTEXT

Through most of history, the popular image of an astronomer has been of a person alone at night looking at the sky through a telescope. A cliché certainly, but one with a grain of truth; for a long time astronomy was done by groups working independently at facilities scattered across the world. In such circumstances, the scheduling of observations was done locally, based on the conditions and requirements at each site.

Today, we live in a period when astronomical technology has been advancing rapidly. Simple observations can now be made remotely using robotic telescopes. Some telescopes are antenna arrays that can be used as a single instrument or split into two or three, observing different targets simultaneously. New technology has enabled the detection of neutrinos, ɣ-rays and gravity waves. Another class of instrument aims to image large parts of the sky every few days, searching for small changes that signal something interesting.

Not surprisingly, this surge in capability has led to a great increase in our astronomical knowledge on all scales, from the structure of the Universe itself down to the detection of many planets around other stars. Greater knowledge and greater capability has led to ambitious plans for further observations, involving large surveys, or simultaneous observation of targets at different wavelengths, or quick reaction to transient phenomena.

The result is that today the use of our facilities is connected and interdependent as never before. Local observation scheduling by hand is no longer adequate, instead we need to find ways to arrange and coordinate observations efficiently across facilities, taking into account the full range of new use cases. This is a fresh and important field, offering great opportunities but also with room for mistakes and unfairness to occur. The goal of this document is to examine how observation scheduling is currently organized, expose any problems, and suggest improvements.

The document title is `Report on Scheduling Algorithms and Standard Interfaces for Cross-Facility Scheduling'. We have chosen a very broad interpretation of the word 'scheduling', one that includes all the elements involved in obtaining cross-facility observations; from the formation of collaborations, through how collaborations decide on the observations to be made, to the efficient scheduling of those observations at the facilities.

XI. PREVIOUS STUDIES

People have looked before at the problem of cross-facility observing. Two studies are of particular interest, though they both limit themselves to the observation of transients. First, `New windows on transients across the Universe' is an article associated with a Royal Society Discussion Meeting (O'Brian 2012), which gives an early look at the field. Second, 'Paving the Way to Simultaneous, Multi-wavelength Astronomy' (Middleton 2017) summarises the recommendations of a 2015 workshop on the subject. It describes in detail the range of astronomical targets that would benefit from coordinated observing, lists the difficulties in doing so with the present system, and suggests some solutions.

XII. CLASSICAL OBSERVING

We begin by looking at how astronomy is commonly done now. For many years, the typical astronomical facility was an observatory on a mountain top, with a telescope, instruments and staff. Though funds and observing time were allocated by a national agency, the observations were usually made in isolation. Electromagnetic waves (EM) were the only messenger studied.

The `Observation Management System' (OMS) is the overall high-level process by which science results are obtained, from the initial request for time, through the taking of the data, to the point where the results are published and archived. Over time, the process followed by observatories has converged to a broadly common model:

1. The facility publishes a 'Call for Proposals', advertising the capabilities offered and setting a deadline.
2. Scientists submit proposals, describing the data that they want and giving a scientific justification of why.
3. A technical review of each proposal is carried out by observatory staff, a scientific review by external scientists, a panel meets to discuss the applications. Proposals are ranked, and those above a certain point accepted.
4. A long-term scheduling plan is constructed, based on the accepted proposals and available resources.
5. A pool of `scheduling blocks' (SBs) is created, each containing a full description of the observation to be performed.
6. The observations are scheduled. There is much variation in how this is done. For example, the HST is very over-subscribed but operates under strict, predictable constraints, which encourages and allows time for the development of a carefully optimized schedule. Conversely, ALMA has a real-time constraint not known in advance, namely the weather, in which case the schedule must remain adaptable up to the last minute.
7. What happens after the data are taken also varies from observatory to observatory. The old approach was simply to give the raw data to the applicants and place a copy in a public archive after a proprietary period. Nowadays, many facilities have pipelines that reduce the data automatically.

Coordinated observations have often been carried out in the classical framework, though the process can be laborious. The observer (usually a collaboration) must make separate time applications to each required facility, making sure that the multi-facility linkage is clear to each, and hope they get the time they ask for. Problems with this are that the observing semesters of facilities may not match, and that possibly some applications will be successful but others not, resulting in a `patchy' observation. Once time has been obtained the applicants must then work with the facilities to put together the coordinated schedule.

In recent years, some facilities have tried to make the process easier by inviting applications for joint observations, e.g. ESO for VLT / XMM-Newton, NRAO for radio / Chandra, HST or Swift. However, these changes are piecemeal solutions to the broader problem.

Follow-up of `transient' events, which are short-lived and pop up unpredictably, is handled by observers asking for observations to be made in the event of a specific trigger - so called

`Target of Opportunity' (ToO) time. The observation is made as soon as is convenient after the trigger is pulled. In extreme cases, where an important but unforeseen transient occurs, observers can apply for `Director's Discretionary Time' (DDT) at a facility - with a response measured in hours. Obviously, this method cannot be used very often.

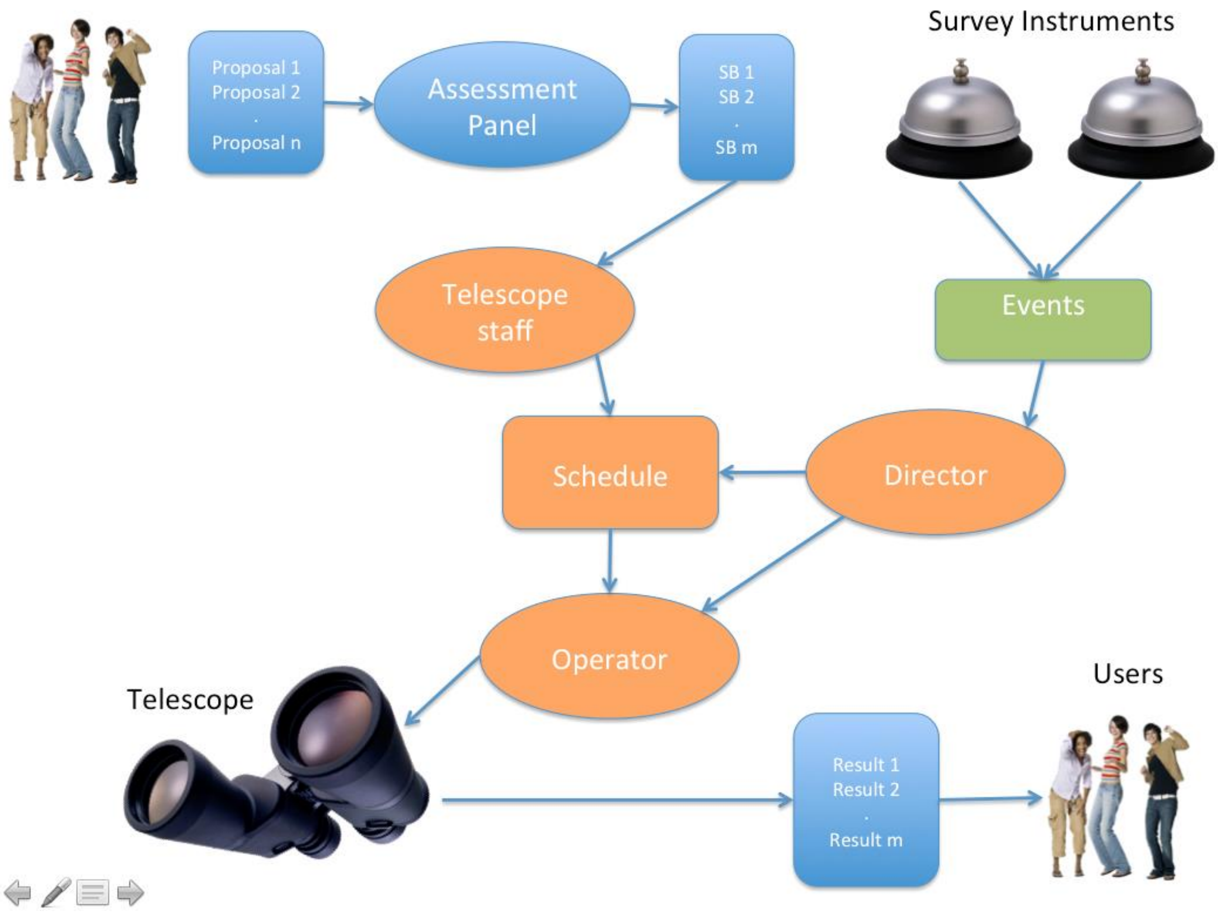

Figure 1 Diagram illustrating the current model for observation management.

XIII. COORDINATED OBSERVATIONS

Many targets can benefit from coordinated observation at multiple EM wavelengths. A good review of such sources forms part of 'Paving the Way …' (Middleton 2017)]. Fig. 2 is taken from that paper and shows the wavelengths and coordination timescales of the various classes of target.

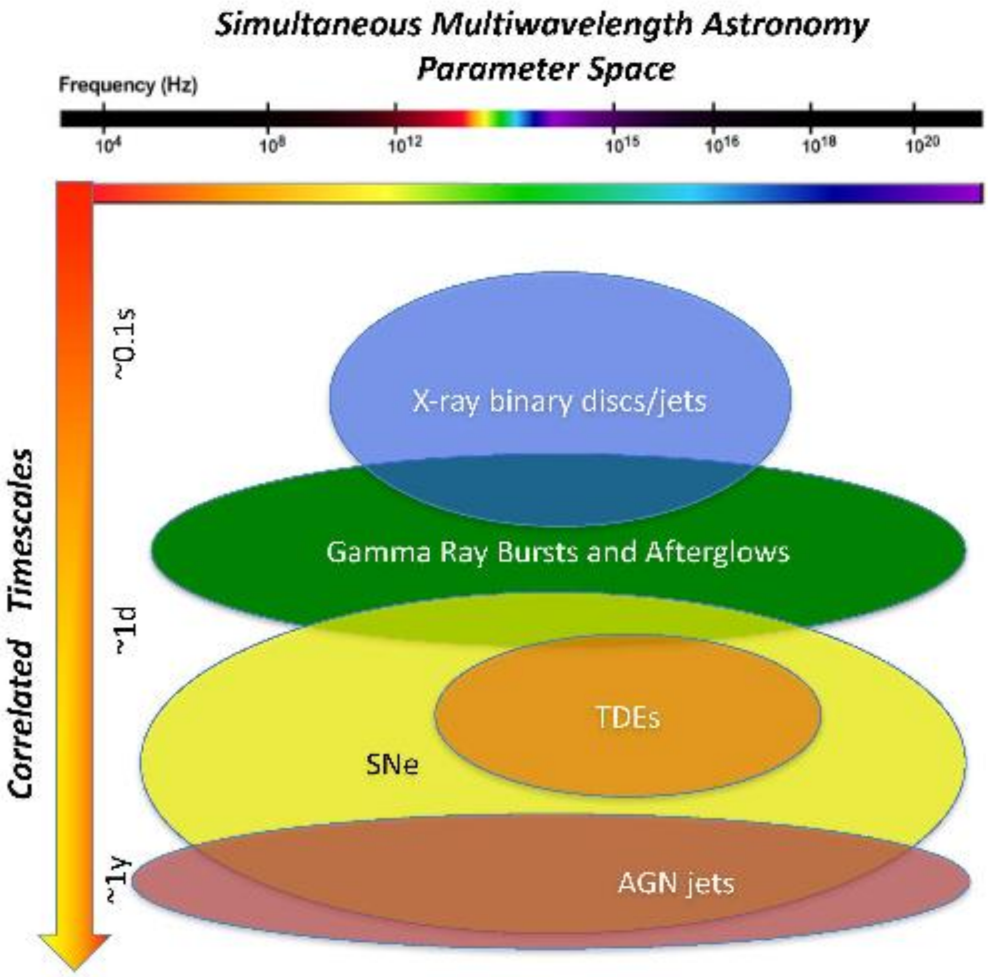

Figure 2 The above diagram, taken from 'Paving the Way…' (Middleton 2017), shows the range of observation wavelength and coordination timescale required for various types of object. TDEs are `Tidal Disruption Events', where a star is torn apart by a black hole, SNe supernovae, AGN `Active Galactic Nuclei'.

TRANSIENT ASTRONOMY

Many of the objects in Figure 2 that have been subjected to coordinated observation are transient targets. The coordination is driven by the fact that they fade after detection and must be observed while still bright enough to see. Supernovae and novae have been observed for hundreds of years, while γ-ray bursts (GRBs) joined the party in 1967, neutrinos in 1968, and gravity waves in 2017.

MULTI-MESSENGER ASTRONOMY – A NEW AGE

The 4 'messengers' of astronomy derive from the 4 fundamental forces of nature; electromagnetic waves (EM, electromagnetism), neutrinos (ν, the weak nuclear force), cosmic rays (CR, the strong nuclear force) and gravity waves (GW, gravity). Until the 20th century only the EM messenger was observed but today we can see them all, including ν from the Sun and supernovae, and GW from the merger of massive objects such as black holes or neutron stars.

Current ν and GW detectors are sensitive to a large area of sky but have poor directional resolution. This means that multi-messenger astronomy today usually consists of an exotic transient detection, perhaps correlated with alerts from different messengers, triggering follow-up EM observations to locate the source, then more detailed measurements of its EM characteristics before it fades.

XIV. CASE STUDIES

How is observation scheduling done presently? In this section, we look at some examples, concentrating on the more complicated use-cases because these are the main drivers of future work. We begin with a look at the coordinated observations of various types of object, many borrowed from 'Paving the Way…' (Middleton 2017), where they are described in greater detail.

ACTIVE GALACTIC NUCLEI

Studies have been made of the relation between accretion disk behaviour and radio jet production using monitoring X-ray and VLBA radio data, covering the same period and with observations simultaneous within a week or so (Lohfink 2013). This illustrates the fact that often data do not have to be obtained as part of a coordinated programme, but can be brought together after the event.

In principle, similar studies of the action closer to the central black hole could be made, but this would require simultaneity on the order of the light crossing time of the region involved - of order hours - requiring monitoring programmes with a higher cadence, or that specific coordinated observations be made.

BLAZARS

These are thought to be radio loud AGN with the jet pointing towards the observer, though the details of their behaviour are proving difficult to understand. Long term monitoring of a number of sources in radio, optical and ɣ-ray, with weekly cadence and day simultaneity would help resolve these difficulties.

SGRA*

The object at the Galactic Centre is our closest SMBH. As befits a unique and interesting source there have been many multi-wavelength studies of it. One notable, recent campaign centred around the first observing run of the Event Horizon Telescope (EHT) in April 2017. Coordinated observations were obtained using radio (EAVN), mm (GMVA, EHT, ALMA), near IR (VLT), X-ray (NuStar, Chandra, Swift) and ɣ-ray (HESS, MAGIC) facilities. Available details of the coordination process (Middleton 2017) suggest that some observations were simply scheduled as close together as could be managed, while others were ToO, triggered by the scheduled observations taking place or by the detection of transients associated with the target. This is an example of a case where a prominent observation, in this case the EHT, attracts `followers'.

STELLAR CORONAL ACTIVITY

Less energetic targets are also of interest. Coordinated, quasi-simultaneous optical and X-ray observations have been used to investigate links between coronal emissions from a star with the position of an orbiting planet (Scandariato 2013).

VARIABLE REFLECTION NEBULAE

Many young stars are associated with reflection nebulae, a proportion of which vary in shape, e.g. Hind's Nebula illuminated by T Tau, Gyulbudaghian's Nebula by PV Cephei. The prototypical object is `Hubble's Variable Nebula', NGC 2261, illuminated by the Herbig Be star R Mon (Lightfoot 1989).

The variability is from shadows cast by objects moving near the star, opening a unique window onto the behaviour of protoplanetary disks within 1au of their primary. Robotic telescopes, in particular the multi-instrument LCOGT, make possible the acquisition of long-term, high cadence image sequences of these objects. It would be interesting to obtain long-slit, medium resolution spectra of the reflection nebulae and, by correlating these with the shadows, seek to build up a 3d picture of events. Such observations need only be simultaneous on the order of days and could be arranged by ToO applications for the spectroscopy.

GW170817

The examples given so far have been of persistent sources, ones that have been known and observed for some time. From here on the example observations are triggered by transients; either pre-known objects that are suddenly doing something particularly interesting, or objects that were unknown until they brightened and were detected. Newly discovered transients usually fade after the initial flare-up so the push to coordinate is largely driven by the need to observe them while they are bright enough to see.

Gravity waves from the merger of 2 neutron stars were detected by LIGO and VIRGO on 17th Aug 2017. The `LIGO Scientific Collaboration', comprising over 100 institutions, moved to follow up the detection. The following is a brief description of the sequence:

1. The GW event is detected by LIGO/VIRGO. Its signature is consistent with the merger of 2 neutron stars, a type of event predicted to have an EM counterpart. Fermi-GBM (Gamma-ray Burst Monitor) detects a GRB emitted 1.7s later. Initially, the GW source location can only be pinned down to one hemisphere, twelve hours later re-analysis of the data narrows this to an area of 28 sq.deg. The GW luminosity and likely mass of the neutron stars implies a source distance of about 40 Mpc.
2. The first requirement is to locate the source accurately. A number of telescopes, including Pan-STARRS, VISTA and the LCOGT network, start searching for the optical counterpart. They work independently of each other but use similar strategies to speed up the search; imaging galaxies in the right area of sky and distance range, given that these will contain the great majority of stars in the search volume. The Carnegie Observatories team (Swope Supernova Survey) optimise their search further by prioritizing locations where as many galaxies as possible can be observed in a single field of view. In the event, the Swope 1m telescope is the first to report a new optical source in the galaxy NGC 4993, roughly 12.5 hours post event.

3. Is the NGC 4993 object really the GW source? More information is needed to be sure. Many telescopes start taking sequences of photometric measurements - the LCOGT is particularly well suited to this, having many identical instruments distributed around the world. In addition, several telescopes obtain spectra of the object, the crucial information needed to confirm the identification of the source as the GW target. The first spectrum is obtained by the Magellan telescope, subsequent ones by the NTT (working for the PESSTO project, described below), by the VLT and many others. This phase of the follow-up stretches over 10 days, continuing until the target fades.
4. Observation planning seems to have been carried out informally by a loose collaboration of scientists, many of whom knew each other and had worked together on the study of previous transients.

V404 CYGNI

This (BH)LMXB, acronym soup for an `X-Ray Binary involving a Black Hole and a Low Mass donor star', produced a series of strong X-ray outbursts in June 2015. The main phase lasted about 2 weeks and produced variable emission across the EM spectrum on timescales down to 10ms. The following description of events is taken from 'Paving the Way…' (Middleton 2017).

1. June 15, initial signs of brightening are detected by the Swift γ-ray satellite.
2. June 25, a public mailing list is established to facilitate communication between observers.
3. June 30, a public web interface is created to show details of planned/complete observations. The coordinated observing campaign is advertised via Astronomical Telegram AT 7735 and planned to centre around an HST observation which, however, is eventually cancelled because of the source fading.
4. Near continuous monitoring of the target is achieved in X-ray, γ-ray and radio after DDT and ToO applications to various facilities.
5. Members of the American Association of Variable Star Observers (AAVSO) make a mass of loosely coordinated photometric observations using small telescopes, which are collated later.

PESSTO and the `Marshall'

The `Public ESO Spectroscopic Survey of Transient Objects' (PESSTO), and its successor ePESSTO, are long-term programmes that use spectrographs on the New Technology Telescope (NTT) plus photometric monitoring by the LCOGT to observe and classify optical transients picked up by synoptic surveys (Smartt 2015).

The project uses a bespoke web app called the `Marshall' to assist the process of selecting suitable transients, the scheduling of follow up observations, and classification. A sample Marshall page, showing information for a candidate target, is shown in Figure 3. It presents in one place the current information for that source, telling the story of its progress through the follow-up observations, comments from team members, etc.

Transient alerts are sourced from several streams (e.g. Gaia, Pan-STARRS, Swift, ASAS-SN), filtered through selection criteria, and loaded into the Marshall `inbox', all automatically. Within hours, the project's `Transient and Alert Team' (TAT) triage them in order of likely interest - for PESSTO this means galaxy nuclear transients, high-luminosity variations in galaxy cores, TDEs, circum-nuclear SNe, blue hyper-variables, GRB, and GW. The TAT can choose to ignore them, `snooze' them i.e. wait for more information to come in, or schedule them for follow-up photometry or spectrography.

At the time of writing PESSTO had accepted roughly 40000 transient alerts, and classified nearly 8000.

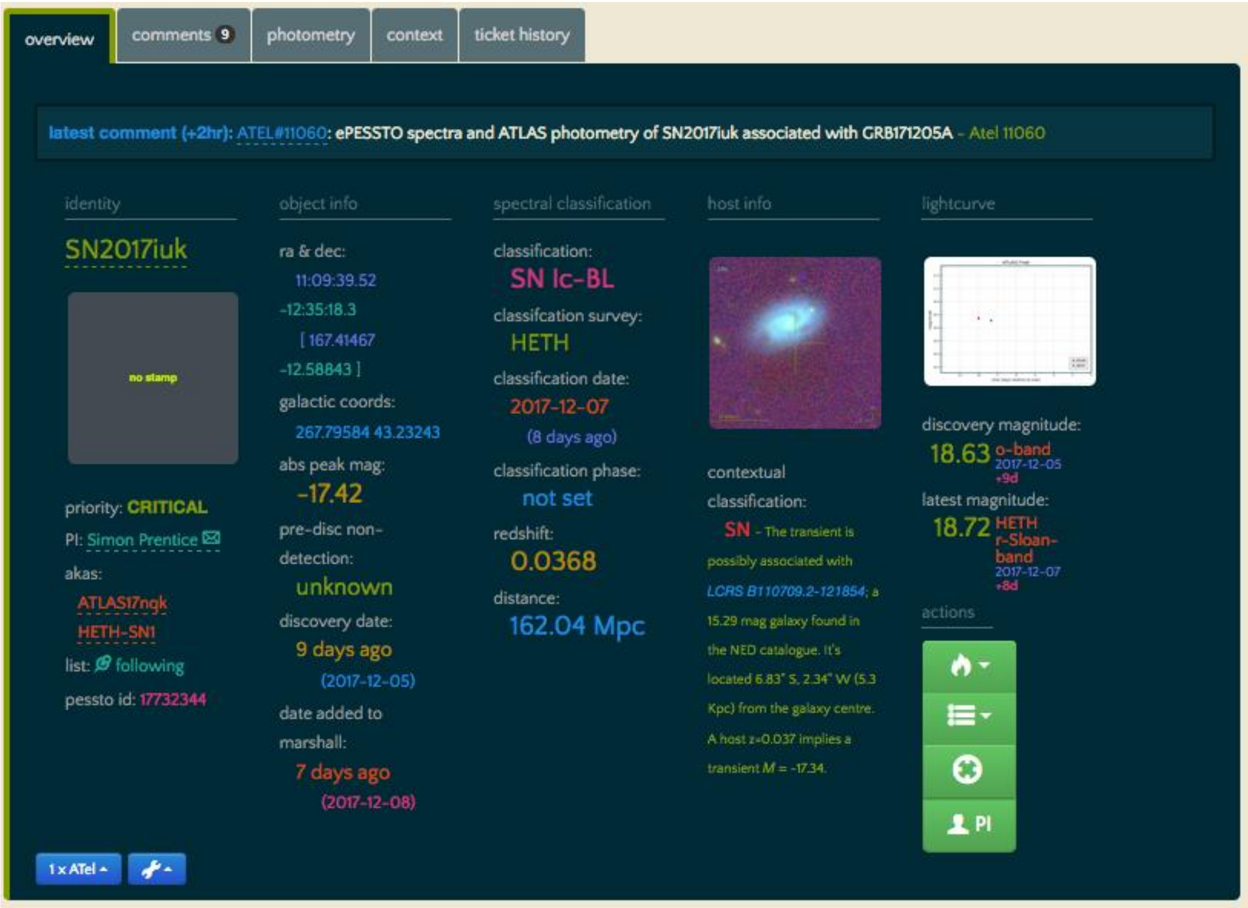

Figure 3 A screenshot from the PESSTO Marshall web app showing the information for a candidate transient.

XV. TRENDS IN INSTRUMENTATION

ROBOTIC TELESCOPES

These are telescopes that operate remotely without human attendance. They are well suited to general follow-up observations that are planned and subject to change in real time. Observations are generally requested via a web interface, scheduling is managed by a computer algorithm. General purpose instruments have been around for a while, e.g. the Liverpool Telescope (LT) since 2004, the Las Cumbres Observatory Global Telescope

(LCOGT) since 2005, extended in 2013. Both the LT and LCOGT have spectroscopes as well as imagers. Cutting edge or operator-intensive instrumentation is not offered. Some robotic telescopes are dedicated to particular pursuits, e.g. BlackGEM and MeerLICHT described below.

SYNOPTIC SURVEYS

Synoptic survey instruments obtain a broad picture of the sky and detect changes in it through time. The Large Synoptic Survey Telescope (LSST) is one such project, where an 8.4m telescope aims to survey the entire accessible sky every 2 nights, beginning in 2023. Analysis of the data is expected to generate tens of thousands of transient alerts each night. Care will be taken to ensure that the alerts are robust (LSST Science Book Version 2.0, ch.8 2009), and each will come with ancillary data. Most of the interesting and well-covered (by the LSST) transients will be bright, suitable for follow-up spectroscopy/photometry by telescopes in the 1-4m range.

The Square Kilometre Array (SKA) will survey a large part of the southern radio sky. In the course of these observations, dedicated data reduction pipelines will be searching the data stream for evidence of radio transients from pulsars, GRBs, etc. It is anticipated that a very large number of transients will be detected each night, though probably not as many as the LSST.

The Cherenkov Telescope Array (CTA) will conduct observations of the high energy (20GeV - 300TeV) ɣ-ray sky for the next decade and beyond. It is expecting to be used as a follow-up instrument triggered by high energy transients, and to generate its own alerts from a Very High Energy (VHE) transient survey (Acharya, Science with the Cherenkov Telescope Array 2017).

On a smaller scale, BlackGEM will use 3 robotic 65cm telescopes to carry out several surveys, one measuring 10000 square degrees (¼ of the sky), every two weeks, starting in September 2018. BlackGEM is funded by the LIGO and Virgo consortia specifically to locate GW targets within the broad `beam’ of the interferometers so, presumably, it will only do synoptic survey work when not chasing down a GW detection.

DEDICATED FOLLOW-UP

BlackGEM is an example of a trend where large facilities fund small optical telescopes to perform dedicated optical follow-up. Likewise, MeerLICHT is a single 65cm telescope that has been built to image at all times the same piece of sky as is being observed by the MeerKAT radio telescope. Both telescopes are robotic.

PRE-ALLOCATED FOLLOW-UP

A variation on `dedicated follow-up’ is where a large project has time pre-allocated on a range of standard facilities. ESA’s `PLanetary Transit and Oscillation of stars’ (PLATO) mission, scheduled for launch in 2026, will need many follow-up observations to confirm exoplanet candidates. Time for this will be allocated on a range of telescopes. Observations will be triggered by a scheduler that will know beforehand the availability and status of the participating facilities. The architecture is under discussion, but the need has been clearly identified.

THE EXPLOSION OF TRANSIENT NUMBERS

A most important point is that the advent of facilities such as the LSST, SKA and CTA will massively increase the rate at which transients are detected - estimates for the LSST alone range up to 10^7 alerts per night. Human intervention and actual follow-up observations will be possible for only a tiny proportion of this number. Most will only be followed up with data mined from other surveys.

Automatic alert curation will be essential, which makes high demands on the amount and quality of information delivered with each alert. Correlation of alerts and the accumulation of information on alert targets is a service offered by the `Astrophysical Multimessenger Observatory Network' (AMON) (Keivani 2017), and we think it likely that this service will be an important component in any alert curation system. AMON was constructed to make marginal alerts more robust, but further advantages from alert correlation will include updates to the target position, and information on follow-up observations, which would all be very useful for target filtering.

THE GCN/TAN NETWORK, VOEVENTS AND BEYOND?

Today, information on transients is communicated via the `GRB Coordinates Network' (GCN). GCN was created to disseminate the coordinates of $\gamma$-ray bursts; its name will be changing to `Transient Astronomy Network' (TAN) as it transitions to handle all types of transient and associated follow-up observations. GCN/TAN carries two types of message. `Notices' distribute transient locations via simple, token-value style e-mails, or internet socket packets for time-critical applications. `Circulars' are prose-style e-mails describing follow-up observations towards GRBs only; such information for other transient types must use the `Astronomer's Telegram' (ATel) system.

With the arrival of the LSST and other synoptic survey instruments, alert traffic is expected to increase by many orders of magnitude. GCN/TAN will be unable to carry this volume. The VOEvent schema and transport protocol were developed by the IVOA as the next generation carrier but the current view from the LSST is that even this will not suffice, and they are developing a new data format and transport protocol (D.Morris, co-chair IVOA working group, private communication).

XVI. RESEARCH CONFIDENTIALITY

Confidentiality is an important concern for all observers. Though research results are eventually published and the raw data made available to all, privacy is needed while the observations are taken and analysed; people often don't want their line of enquiry to become known before it bears fruit.

For classical observing this means that the details of applications for telescope time are kept secret, and data are usually kept private for a proprietary period, typically one year, after being taken.

In `transient' astronomy the alert itself is a valuable datum, useless if kept secret for any time, so the idea of a proprietary period will not work. A common policy is to share the alert only with institutions that have signed a `Memorandum of Understanding' (MoU), agreeing to act in concert with the generator of the alert until publication. Until that time, results and comments are shared privately between participating institutions.

XVII. IMPROVEMENTS IN COORDINATED / MULTI-MESSENGER OBSERVING

So far, we have described how coordinated observing works today. What lessons can be learned?

The easiest type of follow-up is to search the databases of other, contemporaneous, synoptic surveys for correlated alerts e.g. Swift, Fermi or CTA for -ray, LSST or BlackGEM for optical, SKA and LOFAR for radio. Indeed, the great majority of the forthcoming transient flood can *only* be followed up in this way.

The tide of transient alerts is rising fast and will soon require automated filtering to cut the number requiring human attention. Furthermore, initial alerts are often refined by later ones, which must be handled appropriately. Networking by GCN/TAN is adequate for overall coordination at the current rate of alerts but is not scalable to the higher rates expected. VOEvents or the new scheme being developed by the LSST will have to be used instead.

Where actual follow-up observations are required, collaborations are important. Coordinated and multi-messenger observations require a wide range of expertise, both to take the measurements and to analyse them. These demands are best met by a team, each member with experience of a specific domain, most likely at a different geographical location. Social and geographical barriers must be overcome to form a collaboration and, though several have formed with great success, it is worth asking if the process be made easier? Could social media help?

Another hindrance is lack of **continuity**. Established areas of research evolve comfort zones, with their own journals, conference series, etc., but there is currently no such zone for coordinated observing; no journal specialises in this work and conferences are irregular. When coupled with the geographical spread of the interested parties, this makes it hard to know what is happening in the field, to coordinate development, to keep track of the subject's history.

We can split the process of coordinated observing into two pieces; **observation planning**, the way the desired science translates to observations required, and **observation scheduling**, how the required observations are obtained efficiently.

For observation planning, a loose association of people can do well if they are motivated. This happened naturally for GW170817 because it was a seminal event. The V404 Cygni flare sequence commanded less attention but the collaboration leaders added buzz by getting HST time, then facilitated communication with their email list and web page. However, the observation planning process is often chaotic, for example the duplication of effort during the search for the optical counterpart of GW170817. Here, the PESSTO Marshall shows the way forward, where some supporting infrastructure and a defined process makes routine the triage of alerts and the scheduling of follow-up observations.

Research confidentiality is a general concern, handled differently in each of the case studies. For V404 Cygni the mailing list and coordination web page were public, as were some of the data obtained. For GW170817, members of the coordinated study group all signed the Ligo/Virgo MoU, which set out the rules to be followed. Access to the PESSTO Marshall is limited to team members, with results being published by ATel when ready.

CONTINUITY

We propose that a portal be established to act as a centre for all things to do with coordinated observing. Such a portal should have the following characteristics:

- It will persist beyond the ASTERICS project.
- Its content should be attractive and useful to visitors. Some possibilities:
  - News items on coordinated observing, articles on existing coordinated programs, telling the story of their development, lessons learned, etc. This would spread good practice.
  - A compendium of astronomical tools. Many countries and projects have developed a wide range of tools, most probably unknown outside their native user base. A page listing the tools, with links to their homepages, reader recommendations, and examples of their use would be interesting.
  - Links to public data archives and advice on how to access them, especially those belonging to synoptic survey instruments.
  - General useful links. For example, facility time application pages, etc.
  - A link to the `Coordinated Astronomy' Facebook group, if such is created. As described below, this would be a window on to existing collaborations and an aid to forming new ones.
  - A page linking to a framework and tools that can be assembled into a coordinating Marshall-like web app, for new collaborations.

OBSERVATION PLANNING

The planning is usually done by a collaboration. We look in turn at the **formation of a collaboration,** then the **working of the collaboration**.

THE FORMING OF A COLLABORATION

How do collaborations form? Until now, the process has taken place through social `magic', involving face-to-face contact and networking at conferences, workshops and the like. In today's connected age, contact is enabled by the web, the important functions of gossip, rumour and face-to-face conversation by social media platforms.

First, we note that there are several different types of collaboration, each with its particular aims and requirements, and each likely to interest a different kind of astronomer. These are:

- **Self-assembling collaborations,** which occur naturally as people chase major results, for example the LIGO/Virgo neutron star merger. The gravity waves attracted a wide range of astronomers, at all stages of their career, all keen to get a slice of the action and their name and expertise part of an important publication.

- **Architect designed collaborations**. These are usually put together by experienced, high status astronomers for a specific purpose, which can either be to await transients of a particular type and study them, like PESSTO, or to carry out a series of coordinated observations of a particular non-transient object. They resemble the `self-assembling' collaborations except that their target objects are less high profile. Their architects organize the effort to obtain follow-up observing time and money for infrastructure development. Others join the collaboration for a range of motives; because they are interested in the target obviously, but also to gain experience, to make contacts and raise their profile in the specialism, to have their name on some publications.
- **Single-object collaborations**. This simple form is used to observe a transient object that is of particular interest one person, who invites like-minded people to share the investigation. Such collaborations are small, involving just a few people, and there is little need for complicated infrastructure or process to make them work. They appeal to investigators at any career stage, but offer a rare chance for young astronomers to organise themselves and gain footholds in the profession - which should be encouraged!

All collaborations must bear in mind research confidentiality. Large collaborations, of the `self-assembling' and `architect designed' types, establish ground rules covering confidentiality and data privacy by requiring that members sign a `Memorandum of Understanding' (MOU) on joining. For the same reason, if web tools are used for data display by a collaboration, then these resources must be private.

How could social media help to form collaborations?

`Self-assembling' collaborations need no help as these accrete naturally around big scientific events. However, such collaborations should definitely have a presence on the web and would benefit from web-based organizational tools to allow messaging, data sharing, and coordinated planning to prevent duplication.

`Single object' collaborations observing transients are addressed by SMARTNet (Middleton 2017), which is described below. It seems a good, lightweight solution given the small number of people usually involved in each collaboration and, though it is intended for the observation of transients, it would work just as well for persistent objects.

For `architect designed' collaborations, we start by conducting a thought experiment. How would new collaborations develop if all potentially interested people could be brought together in one large room? The likely behaviour would be for people to congregate with others they know, then discuss and develop ideas from there. This would not be bad if it happened, but it would tend to reinforce pre-existing networks and freeze out newcomers.

A process improvement worth trying would be to follow the example of the 2018 `ASTERICS all-hands' workshop in Amsterdam, a meeting called with the express purpose of encouraging new linkages to form across the project. To this end, attendees were encouraged to write a brief summary of their interests on the registration site. Then, more importantly, there were several rounds of `pitch and discussion', each involving a short sequence of `pitches' where people gave a 5-minute description of an idea to the whole meeting, followed by hour long parallel `discussions' by groups attracted to each. This process did succeed in shaking people out of their comfort zones to look at new ideas and make new contacts.

Facebook FOR ASTRONOMERS

Today the `large room' online can be entered using any number of social media platforms. Which ones would suit our purpose? LinkedIn is an obvious option, being designed with the express purpose of linking professionals to their ideal job, though it may be overkill in this case. Facebook could be good enough if all we want to do is implement the simple `pitch/discuss' process described above. It is massively popular and has a simple interface coupled with sophisticated and proactive search capabilities.

A strawman design built on Facebook would involve:

- A closed Facebook group called `Coordinated Astronomy', which people would apply to join. This group, its members and a description of its purpose would be public but only members could see posts. Membership and behaviour would be the responsibility of the group admin and moderators.
- A set of standard hashtags for members to signal their interests and skills. By writing a public post on their own page or to the group with an appropriate collection of hashtags, a member's interests would become visible to Facebook's built-in search facility. This post would usefully contain human readable information as well; a link to the person's CV, homepage, publication list, etc., with a fuller explanation of their particular interests for collaborative work, and capabilities offered.
- A member seeking to establish a new collaboration would post a `pitch' to the general group page, briefly describing their idea to the rest of the group, and create a new closed group as the `table' at which interested parties would discuss how to proceed. The public description of the new group would describe the aims of the collaboration. An admin member of the `Coordinated Astronomy' group would have to link it to the new group to make the latter more visible. Once a collaboration forms it would establish its own private area on the web to coordinate work; the use of Facebook itself stops here.
- Public visibility of the group will give like-minded groups the chance to join forces rather than compete.

Facebook is designed to help people find others with similar backgrounds or interests, so it has powerful and flexible search facilities - we just have to give it data to work with. It is also a huge platform so will have no trouble handling the numbers involved in coordinated astronomy - perhaps a few thousand, given that in 2008 there were estimated to be 10,000 professional astronomers in the world (Forbes 2008). On the downside, use of Facebook brings along all the `noise' of the web, which will discourage some users.

.

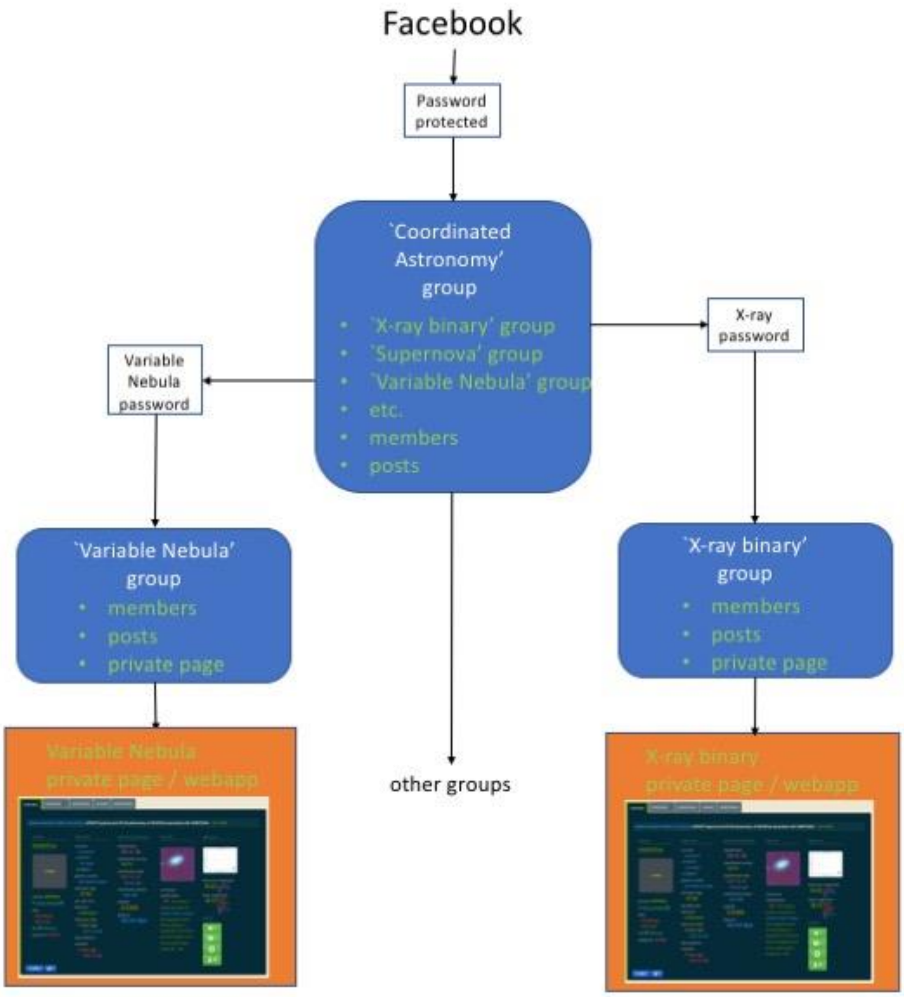

Figure 4 A strawman design for `Facebook for Astronomers'.

OR SMARTNet?

SMARTNet is a web app that developed out of the Leiden meeting described in `Paving the Way…' (Middleton 2017). Like `Facebook for Astronomers' its aim is to bring together people who have knowledge of and access to a wide range of instrumentation, and an interest in making coordinated observations, but no contact with each other. Unlike it, SMARTNet collaborations form around a particular target. Currently, its focus is on transients because that is where the need for rapid collaboration and coordination is most pressing. In principle, there is no barrier to using it for observations of persistent sources.

The design of SMARTNet is driven by a desire to keep it simple, both because the developers had limited resources, and because there is no clear idea yet of what will work best. An incremental design model has been adopted, and the current implementation is viewed as being in its testing phase, from which lessons will be learned for incorporation in future versions. The tool is straightforward and quick to use, reflecting the developers' belief that any interaction slower than writing an email will put people off. The tool is implemented using the Joopla Content Management System (CMS).

The SMARTNet process works as follows:

1. When a member becomes aware of an interesting target, they post an 'observation alert' on the page, with a description of the target, a description of planned observations and of observations that it would be good to have. This invitation is emailed to the other members, and a campaign timeline page is setup.
2. The timeline is a simple coordination aid. All observations contributing to a campaign are entered into the timeline spreadsheet by hand, whereupon they are automatically shown on a graphic representation, Figure 5.
3. Data exchange and decisions on the analysis and publication process are shuffled offline to be agreed there.

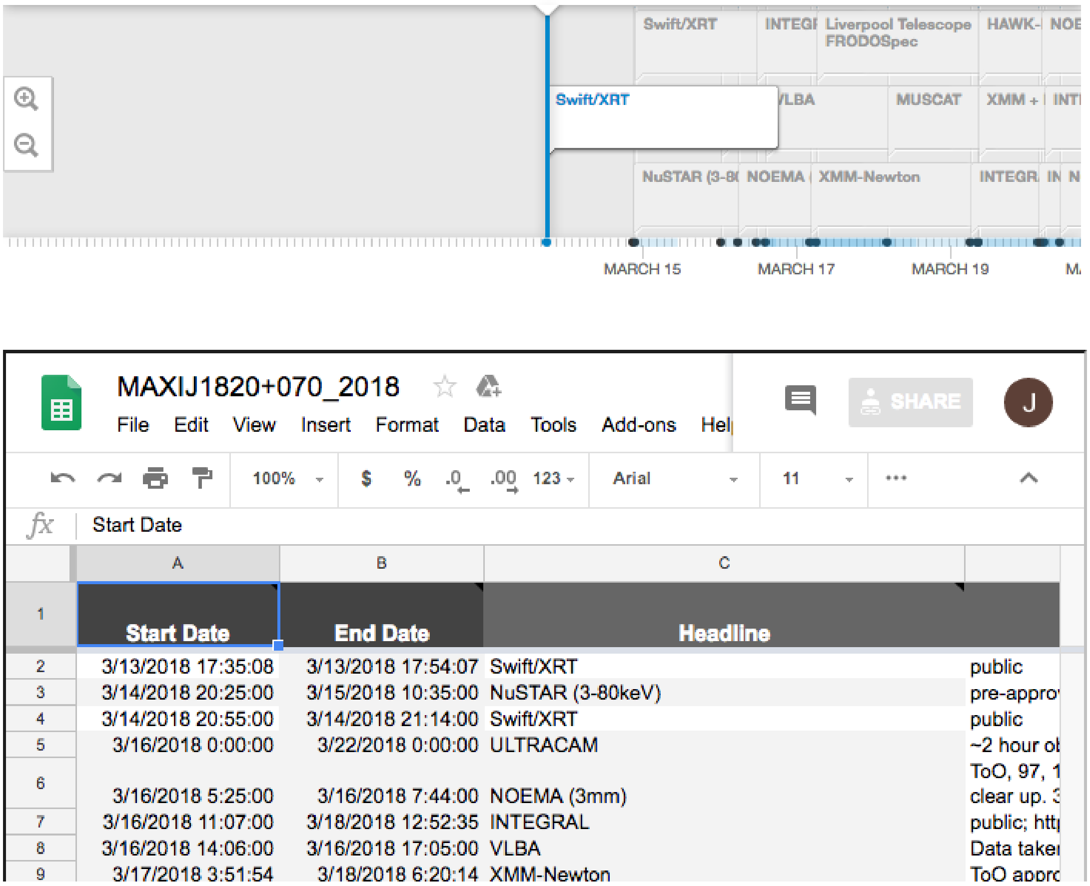

Figure 5 The coordination page for a SMARTNet campaign. Observations entered in the spreadsheet below are shown graphically in the timeline above.

The SMARTNet site also contains information that might be of use to members:

- A directory of members, each with a public profile including a brief description of interests. Currently there are only 50 or so members, so no sophisticated search facilities are needed to make the list useable.

- Links to useful places on the web, such as visibility tools and ToO forms for various facilities.

SMARTNet has some advantages over `Facebook for Astronomers’. The problems to be solved have been considered at greater length and by more people. The web app is in place and has been used successfully. As well as encouraging collaboration, the tool contains an element to help coordinate observations. Something both a blessing and a curse is that it is much more tightly focussed on its target users. At present this renders it invisible to many potential clients but also shuts out the noise and commotion of the web.

THE WORKING OF A COLLABORATION

How a collaboration works depends on its purpose.

`Single object’ collaborations may target either a persistent object at a specific time or a single transient - the latter case being the scenario handled by SMARTNet. If the trigger is private, such as a GW detection by LIGO, then the collaboration will be bound by the associated MoU. However, if the trigger is public, such as a GRB detected by Swift, or the observations are of a persistent object then the coordination need not be private. Indeed, it is simpler and more engaging for users to adopt an `open-science’ approach, where the descriptions of observations obtained or planned are displayed on a public timeline for all to see. Abuse of proprietary data can be prevented by conditions set downstream, at the point of data sharing.

`Self-assembling’ and `architect designed’ collaborations are set up to observe multiple instances of the same event type, nearly always transients. Here, an important aim is to create an efficient process that maximises the number of targets handled. Examples are PESSTO, already described, and the ‘Supernova Exchange‘ (SNeX), associated with the Las Cumbres Observatory Global Telescope (LCOGT), which is a platform for viewing, requesting and sharing observational data on supernovae. By their nature, such collaborations are long lived and can spend significant resources on developing processes and web tools to improve their effectiveness e.g. automatic filtering of transient streams, or tools to access synoptic survey databases. If the transient stream is steady then they can also apply for blocks of time on telescopes for follow-up observations. Paths for data sharing within the group are streamlined, requiring that this aspect be private and that group members be bound by a set of operating rules, perhaps codified in an MoU.

All collaborations that follow the same process multiple times would benefit from having their own `Marshall’. We should consider developing a framework to make this easier. Useful components might include:

1. A hosting service for web apps and web pages.
2. Example designs for web apps with tools that can be generally useful for collaborative work.
3. Software for filtering incoming transient streams according to a list of rules that can be configured by the users.
4. Most follow-up observations will be virtual, performed by harvesting information from synoptic survey databases, e.g. Gaia, Pan-STARRS, OGLE, ATLAS. A tool could be developed to do this automatically.
5. For transient alerts with poor localization, the first follow-up step is always to locate the target. Such searches are sped up by careful planning. In the hunt for the optical counterpart of the LIGO/Virgo neutron star merger, the search was concentrated

toward galaxies, with observed fields ordered according to number of galaxies covered. A software implementation of this scheme could easily be made available.
6. Tools to automatically reduce common types of observation e.g. LCOGT photometry.
7. Tools to ingest and display reduced follow-up data: photometry and spectroscopy.
8. Enable public access to synoptic images. BlackGem, for example, is funded to support follow-up LIGO/Virgo detections. The sky images of its synoptic surveys will be stored but there is no funding for public access.

ALLOCATION OF COORDINATED OBSERVING TIME – THE SuperTAC

If observations are required as part of a coordinated programme, the observers will have to follow some sort of application / time allocation process to obtain telescope time. Classically, applications have to be made to individual facilities, which is cumbersome for a number or reasons e.g. misaligned observing semesters, repetition of science case, patchy success, etc.

Small collaborations of the kind handled by SMARTNet try to overcome the problem by attracting people who either have time on facilities, or have the experience needed to credibly apply for it. Most applications will be for ToO or Director's time. Persisting collaborations can take a longer view, like PESSTO who have applied for blocks of spectroscopy follow-up time as parts of an ESO long-term project.

The underlying awkwardness remains, however. While several facilities are moving in the right direction by inviting combined observations for time on 2 telescopes, e.g. joint VLT/XMM-Newton to ESO, `Paving the Way…' (Middleton 2017) suggested that a proper solution would be to create a `SuperTAC', able to grant time on a range of facilities specifically for coordinated observing. Their version of the SuperTAC dealt only with transients and consequently was designed to react to requests within a few hours at any time of day, but the idea could be realigned to be more general. Indeed, this idea or similar is being actively considered by the ASTERICS policy forum.

The SuperTAC would seem a safe approach, derived from a tried and tested method for allocating time. However, there are difficulties. The panel would have to include a wide range of expertise to cover the facilities available. The applications ask for a large resource allocation, which would probably reinforce the advantage enjoyed by established groups and high-profile astronomers.

OBSERVATION SCHEDULING

The case studies show that transient follow-up planning is often a loose process, well suited to execution by a social collaboration unless a quick reaction time is needed. Such collaboration is less well suited, however, to handling the low-level detail of exactly when a particular observation should be run at a facility. The local constraints will not be known at the high level, and may interact with each other in all sorts of complex ways.

ONE SCHEDULER TO RULE THEM ALL…?

It is tempting to suggest a purely software solution, where lessons learnt are incorporated into the software of a mighty engine of scheduling wisdom, placed in charge of all ESFRI sites. But would this be a good idea?

The scheduling system has 3 basic parts:

1. A plan of the observations that need to be done, set as described in the previous section. This plan will change in real time as the observers react to events, change priorities or signal that data of sufficient quality has been obtained.
2. Information on the operational environment, including source positions, instrument availability, calibration requirements and perhaps the weather.
3. A scheduling algorithm, which must take into account constraints and requirements from all the many domains of the facilities it controls.

Point 3 is the problem. The facilities will range from premier instruments like the Hubble Space Telescope to robot workhorses such as the LCOGT. So various are the domain pressures that some of their demands will conflict. Others will be subtle and difficult to code, or will change on a short timescale. It is difficult to believe that any current scheduling algorithm could do a good job or be seen to do a good job.

Currently the great majority of facilities schedule their observations independently. Rather than replace the current system wholesale, it is perhaps safer to identify its weaknesses and remove them.

A MATRIX OF SCHEDULERS

In this design observation scheduling is carried out by local schedulers at each facility, whose job it is to obtain the requested observations efficiently, Figure 6

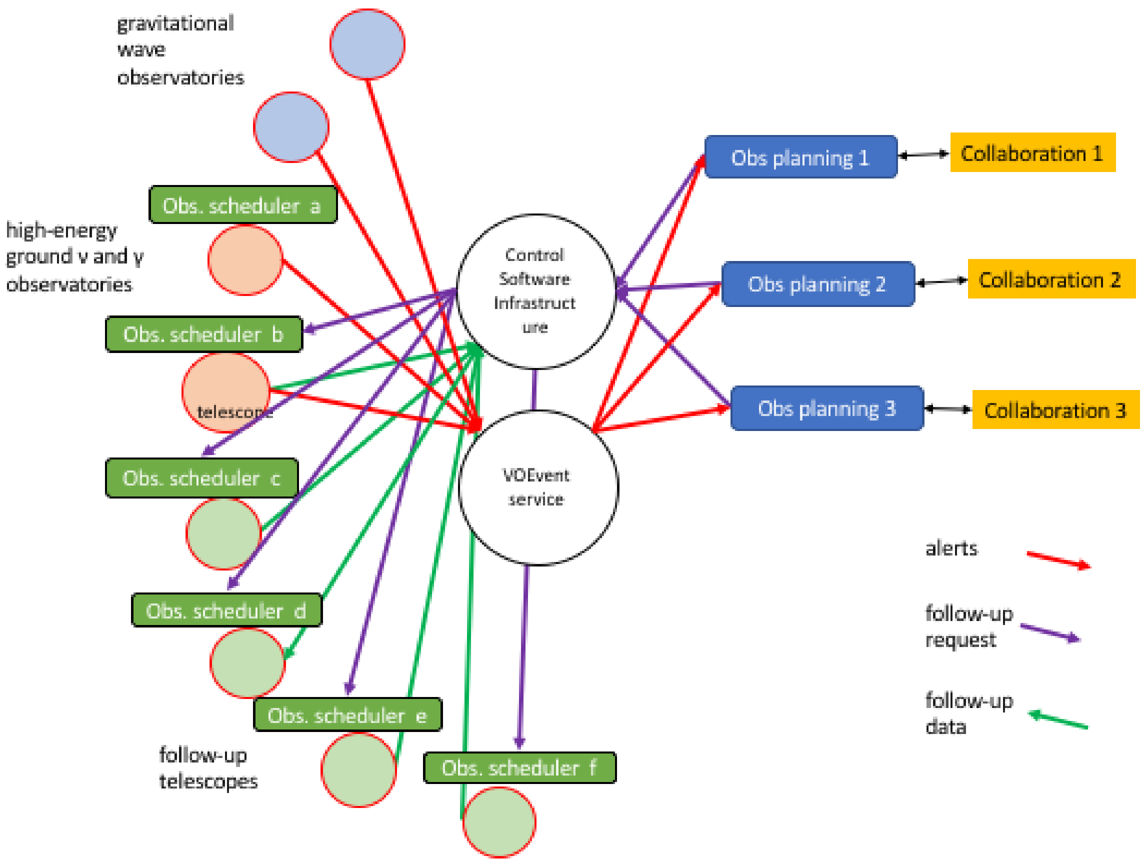

Figure 6 The overall design for a matrix of observation planners and schedulers.

XVIII. FACILITY SCHEDULERS

If observation planners set out the sequence of observations required, it is the job of the facility schedulers to perform the requested observations as efficiently as possible. The planned observations are requests to the facilities, not commands, each carrying within itself the priority level, observation description, and a list of constraints, such as timing, airmass, seeing, etc.

The best way to schedule a single facility will depend on local variables at the facility. Consequently, each facility scheduler will be maintained, developed and controlled by its home institution.

In general, facilities which operate an instrument simple enough (say a single telescope) and on a reaction timescale long enough (say a few hours) may well be happy to carry on being scheduled by people. Human driven systems are flexible and able to integrate many factors into the schedule that are difficult to program. On the other hand, humans are not so good at scheduling facilities that have complex demands for their time or which must react quickly to frequent changes.

The scheduler for each facility can take many forms but should present a common interface to higher elements of the observation management framework. In particular, they should all work with a ‘generic SB’ interface.

PERFORMANCE METRIC

The schedulers aim to maximise the total ‘value’ of the SBs executed, where the ‘value’ of each is calculated with some metric. The details of the metric may vary between facilities. However, if the common aim is to maximise scientific output then the simplest overall measure would be the number of publications that result from the observations. This information is not available to a scheduler at work, but probably maps quite closely to the number of proposed observations that are completed as requested by the applicants. This means that the metric should at least:

- Encourage good use of facility resources; for example, instruments which are available for only part of the scheduling period.
- Minimise overheads, e.g. slewing time.
- Make optimal use of good weather.
- Be biased towards maximising the number of completed projects.

SCHEDULING ALGORITHMS

A number of scheduling algorithms have been developed for operation in various environments. Two scheduling strategies are commonly needed: off-line and on-line (Rasconi, Proceedings of the 17th European Conference on Artificial Intelligence Frontiers 2006b). The off-line strategy includes two planning tools, long-term and mid-term schedulers, which are designed to plan the targets to be observed within a time interval according to the hard constraints that can be predicted. The on-line strategy considers a short-term scheduler that takes into account all constraints and adapts the previously computed mid-term plan to the immediate circumstances (Rasconi 2006a).

LONG-TERM AND MID-TERM

Long-term and mid-term schedulers are designed to plan observations within a time interval according to the hard constraints that can be predicted. These commonly take a long time to run, too long to respond to weather constraints, and operate in environments where humans have the final say in whether a schedule is 'good'.

These algorithms are appropriate when we have a list of SBs, a list of facilities, and the constraints are known beforehand and fixed in time. In this situation, it is possible to spend some effort trying to calculate an optimum schedule (i.e., assign each SB to a specific date and time) for a specific period of time that can range from years to a single night. Thus, the scheduling problem that we want to face-up is considered as NP-hard (Non-deterministic Polynomial acceptable problems) because it is impossible to solve in polynomial time (i.e., feasible time) due to the complexity in computing the large number of possible combinations in search of an optimum solution, and it is necessary to use strategies to guide the exploration in the search space of all feasible solutions with the aim of finding a solution near the optimal one in reasonable time. In this sense, different mathematical tools to solve automated planning and scheduling problems have been developed, ranging from simple heuristics to more complex Artificial Intelligence applications (Donati 2012) (Kitching 2013).

Some examples of ESA tools for computing long-term mission plans are AIMS (Pralet 2009) and Local Search Solution (Kitching 2011) for INTEGRAL, or XMAS for XMM-Newton (Castellini 2009). Also, NASA developed SPIKE, a well-known toolkit used to promote scientific research through the effective and efficient use of ground- and space-based astronomical observatories, which has been used in the long-term planning of the Hubble Space Telescope (Johnston 1994). All these tools use local search algorithms and specific heuristics to find a long-term schedule of the mission, so they may find a solution in a local optimum that is far from the global optimum (Tawarmalani 2002).

On the other hand, global search algorithms are powerful techniques that explore the search space deeply allowing solutions to be found near the global optima. There are several approaches related to global optimization, among which we highlight Evolutionary Computation. This family of algorithms is based on metaheuristics and stochastic optimization, and offers several techniques such as Evolutionary Algorithms, Ant Colony Optimization, Simulated Annealing, Particle Swarm Optimization and Differential Evolution. These kinds of methods have been successfully applied to various optimization problems in different situations, tackling disparate engineering problems. Furthermore, Genetic Algorithms - which belong to the Evolutionary Algorithms class - have been applied by ESA and NASA in long-term scheduling of space missions such as MrSPOCK for the Mars Express mission (Cesta 2009), the evolution of SPIKE that will be used in James Webb Space Telescope (Giuliano 2011), the scheduler of SOFIA mission (Civeit 2013), and the long-term mission planning of EchO (Garcia-Piquer 2014a).

SHORT-TERM

These algorithms run quickly and can respond to changing circumstances, e.g. alerts, weather.

1. *Dispatcher* (Granzer 2004). The algorithm ranks a list of candidate SBs according to merit values for each, and chooses the highest ranked SB to execute next. Merit values are the sum of a number of merit components, each looking at a different aspect - timeliness, weather, etc. The algorithm was developed and is used for

robotic telescopes where there is no human intervention. It is used in CARMENES (A. G. Garcia-Piquer 2014b), in the Automated Planet Finder Telescope (Burt 2015), and in the Las Cumbres Observatory Global Telescope (Brown 2013).

2. *Simplex method* (e.g. (Gharote 2009)). This is a well understood method for maximising some outcome while obeying a number of constraints. It has been criticised for use as a scheduler because it can only maximise one outcome.
3. *Squeaky wheel* (Joslin 1999). This is a simple but apparently effective algorithm that starts with an ordered list of candidate SBs then constructs an observing sequence by choosing SBs from the list in turn. The resulting sequence is examined, conflicts and gaps are noted. Problem SBs are demoted in the ordered list and the process repeated until the solution converges.

COMPOSITE

In some situations, the best scheduler may be a composite object combining the best qualities of long-term, mid-term and short-term schedulers (Granzer 2004). The long-term scheduler filters the available SBs based on hard constraints, for example removing SBs that are never at a high enough elevation, for which no instrument is available, etc. The mid-term scheduler looks at the remaining objects and, for each time period in the night, assembles a list of SBs that it is possible to observe, perhaps ordered to maximise observing efficiency. The long and mid-term schedulers are computationally expensive. The short-term dispatcher sorts the SBs in the target time period based on their merit: this scheduler component can respond to time variable conditions such as weather, alerts, human interaction, adapting the previously computed plan without compromising optimization.

XIX. `STARS' FRAMEWORK FOR TIME SCHEDULING

This section reports case studies done of the suitability of schedulers for particular telescopes and positions in the scheduling hierarchy. A general framework provides scheduling solutions for several facilities, using standard building blocks for optimization, operational constraints and algorithms.

IEEC has developed the framework STARS (Scheduling Technologies for Autonomous Robotic Systems) and has applied it to several space and ground observatories. This framework provides libraries for the definition of the survey (e.g. objects to be observed, features of the objects), the definition of the observatory (e.g. location, number of telescopes, type of telescopes), astronomical calculations (e.g. object coordinates, object elevation, Sun and Moon position, Moon phase), long and mid-term schedulers based on Evolutionary Algorithms, and a dispatcher short-term scheduler that uses astronomy-based heuristics.

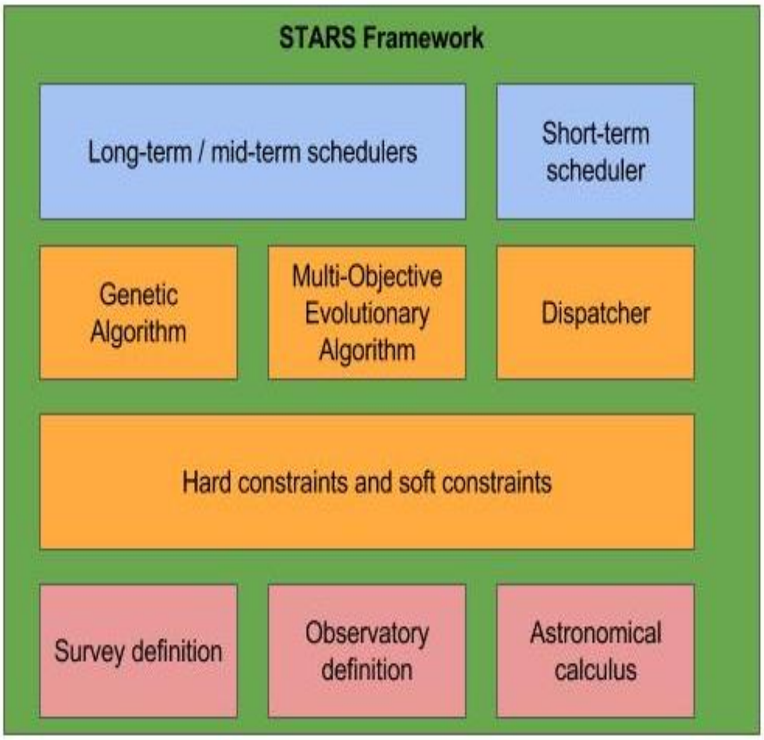

Figure 7 STARS framework modules

SCHEDULING CONSTRAINTS

One of the major challenges of an astronomical survey is the efficient scheduling of the observations of the numerous objects in the sample. In general, any kind of astronomical survey requires the execution of a large number of observations fulfilling several constraints. Some of these constraints can be predicted (e.g. visibility and elevation of the object) and have to be necessarily satisfied, and others are unknown until the time of execution of the observations (e.g. weather, availability of resources). In addition, there are some scientific constraints that should be optimized, such as the proposal completeness, the number of targets that have to be observed and the number of observations of each target. The optimization of these constraints is a key factor for obtaining a suitable schedule with an adequate exploitation of the resources and with a high scientific return. STARS provides the definition of general constraints to be considered in the optimization process and allows the definition of new specific constraints for each survey. STARS divides the constraints into two different categories: hard constraints and soft constraints. The first have to be strictly satisfied, while the second express a preference of some observation combinations over others. Thus, the final scheduling solution must fulfill all hard constraints and should optimize soft constraints.

HARD CONSTRAINTS

Hard constraints are mainly related to the visibility of the targets from the observatory site, the operation overhead times, and the environmental conditions. These restrictions can be adapted for any kind of survey or facility, including space-based telescopes (see e.g. Garcia-Piquer 2014c). The most common hard constraints are described below:

1. Night. The object shall only be observed from afternoon twilight to morning twilight. The coordinates of the targets on the sky and twilight times are computed according to the date of observation and location of the observatory. Additionally, if needed, the observability window of each target can be computed according to given ephemerides for the case of objects with periodic variability. Night-time is usually computed taking into account the astronomical twilight (when the Sun is at -18 deg elevation).
2. Elevation. The elevation of each object is calculated according to its equatorial coordinates and the geographic coordinates of the observatory. The objects shall only be observed if they exceed a specific elevation for at least a certain amount of time. The elevation and time are two parameters that can be introduced in the global configuration of STARS and are typically based on the minimum pointing altitude of the telescope and the estimated integration time.
3. Moon (or bright objects) influence. Targets shall be observed when 1) the Moon is below the horizon or 2) the Moon is sufficiently far away that the observation is not significantly contaminated by background light. The latter can be managed, for instance, with a distance threshold (e.g. a minimum acceptable distance to the Moon of 20 deg) and a condition to select objects that fulfil the observation strategy for the particular science case (e.g. beyond the minimum distance, a hard constraint function on the Moon is evaluated to select only targets that are at least M mag brighter than the background).
4. Visibility duration. The total time during which the night, elevation, and moon influence constraints are fulfilled shall be equal or higher than the minimum visibility time required for a target observation. This minimum time corresponds to the exposure time, which is computed based on the object, the instrument and the observing conditions.
5. Pointing. In case of pointing restrictions, targets shall only be observed if they are between minimum and maximum elevations as defined by the survey requirements. There can be additional restrictions in case the dome/roof configuration reduces the visible area of the sky.
6. Overlapping. In operational terms, there are three kinds of tasks to be considered: 1) a target observation, 2) a readout of the previous observation, and 3) slewing to acquire a new target. Only tasks 2) and 3) can be executed in parallel.
7. Overhead time. Pointing to a particular object requires a specific telescope and instrument configuration. The time between consecutive observations considers both the telescope slew and acquisition time and the instrument readout time. The former includes the time needed to move the dome (if any), and telescope, and an overhead slew time for acquisition, while the latter is defined by the detector properties. The total overhead time between observations is the duration of the process that takes the longer time.
8. Environmental conditions. An observation can be programmed when the environmental conditions permit. In operation mode, the OMS informs STARS whether the environmental conditions are suitable for observation. For schedule calculation, a weather model must be used based on the logged observatory environmental conditions, thus taking the influence of seasonal weather into account. These models contribute to a better estimation of the available night-time for scientific operations.

SOFT CONSTRAINTS

Science requirements are identified as soft constraints. Some of the potential soft constraints are given here below, although this may vary according to the observatory mission.

1. Observing time. The integrated global observing time (i.e., the time that the telescope is observing), especially that of high-priority objects, should be maximized. This guarantees that the most interesting targets are sufficiently observed.
2. Observation deviation. The variance of the number of times that objects of the same priority have been observed in the complete survey should be minimized. This constraint should ensure that all targets will have a proper share of assigned observing time.
3. Observing cadence. It is possible to select the number of observations per night required for each target. As an optional constraint, the planning tool includes functionality to observe the targets at appropriate times. Target ephemerides are also considered.

Some of the constraints can be computed in advance but others, such as weather conditions (environmental conditions), can only be determined in real time during observations and the scheduler must be reactive to their variations. For this reason, although it is an independent system, STARS is connected with the OMS control system from which it receives environment parameters and returns an observation request optimized according to current conditions. In order to reduce waiting times, one of the STARS requirements is that it should invest less than X seconds in selecting the next target to be observed (where X is defined at the facility operations plan; and can also depend on the observation strategy of a specific type of object, like the transient case). The simultaneous fulfilment of the hard constraints and the optimization of the soft constraints should provide a scheduling solution that maximizes the scientific return of the survey.

SCHEDULING OPTIMIZATION

Evolutionary Algorithms (EAs) are an Artificial Intelligence paradigm that includes the learning algorithms based on the way nature solves the problem of living entities (Freitas 2002) by means of natural selection (Darwin 1859) and evolution (Mendel 1866). Genetic Algorithms (GAs) are among the techniques included in the EAs, and are based on finding candidate solutions to an optimization problem considered NP-hard (Holland 1975). They are theoretically and empirically proven to provide a robust search in complex spaces, thereby offering a valid approach to problems requiring efficient and effective searches (Goldberg 1989). The GA process is roughly based on selection, reproduction (crossover) and mutation operators. These genetic operators must be adapted to the particularities of the problem to be solved in order to obtain suitable results (Garcia-Piquer 2012). STARS allows the use of GAs in a transparent way and provides the definition of the genetic operators for the scheduling problem, making it possible to redefine them if necessary.

Generally, a scheduler for astronomical observations has more than one parameter that needs to be optimized, resulting in a Multi-objective Optimization Problem (MOP) that can be defined as the problem of finding a vector of decision variables satisfying constraints and optimizing a vector function whose elements represent the objective functions (Osyczka 1985). These functions form a mathematical description of performance criteria that are usually not disjoint (i.e., they are in conflict with each other). Hence, the term "optimize"' refers to finding a solution that yields acceptable values for all objective functions (Coello Coello 1999). Usually there is not a single point that simultaneously optimizes all the objective functions of a MOP. Therefore, in these problems it is necessary to look for trade-offs, rather than single solutions. The concept of Pareto Optimality (Pareto 1897) defines that we can consider a Pareto optimal when no feasible vector of decision variables exists that would decrease some criterion without causing a simultaneous increase in at least one other criterion. Thus, this concept almost always does not yield a single solution but a set of solutions called the Pareto optimal set. All solutions included in the Pareto optimal set are non-dominated (i.e., there is no solution better than the rest) and they have a different trade-off between objectives (Garcia-Piquer 2012).The plot of the objective functions whose non-dominated vectors are in the Pareto optimal set is called the Pareto front (see for further details Coello Coello 1999 and Coello Coello 2001). Multi-Objective Evolutionary Algorithms (MOEAs) (Coello Coello 2007) are recognized as one of the most valuable and promising approaches to addressing complex and diverse problems of multi-objective optimization. STARS allows the use of MOEAs in a transparent way providing the same functionalities as in GAs.

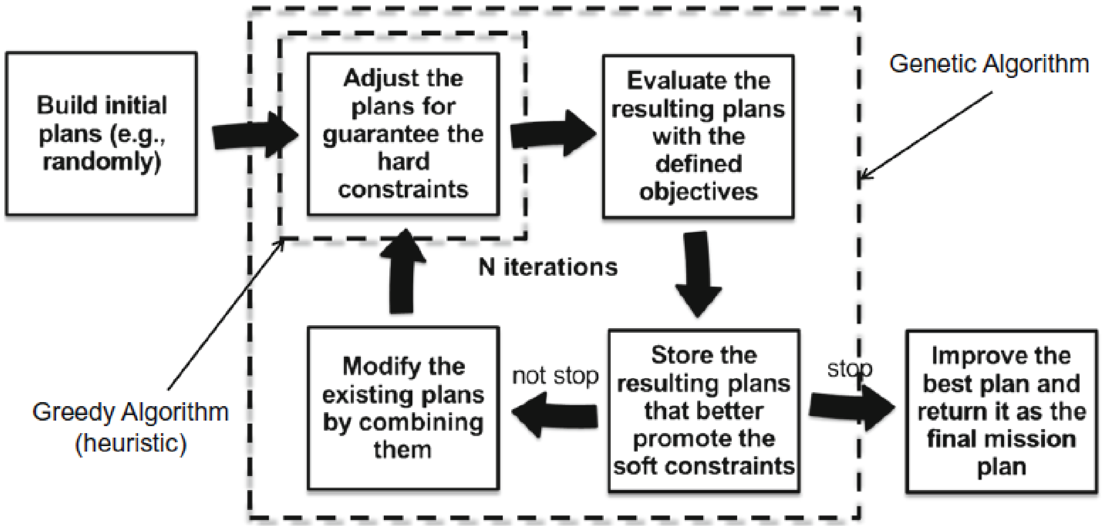

Figure 8 GA optimization cycle

Two scheduling strategies are included in STARS: off-line and on-line (Rasconi 2006b). The off -line strategy includes two planning tools, long-term and mid-term schedulers, which are designed to plan the targets to be observed within a time interval according to the hard constraints that can be predicted. The on-line strategy considers a short-term scheduler that takes into account all constraints and adapts the previously computed mid-term plan to the immediate circumstances (Rasconi 2006a). The following table summarizes the constraints considered by each of the three schedulers.

Constraint	Long-term	Mid-term	Short-term
Hard Constraints			
Night	X	X	X
Elevation	X	X	X
Moon influence	X	X	X
Visibility duration	X	X	X
Pointing	X	X	X
Overlapping		X	X
Overhead time		X	X
Environmental conditions			X
Soft Constraints			
Observing time		X	X
Observation deviation		X	X
Observing cadence	X		

Their combination is illustrated in Figure 9 and we provide additional details on the three scheduling tools in the next section.

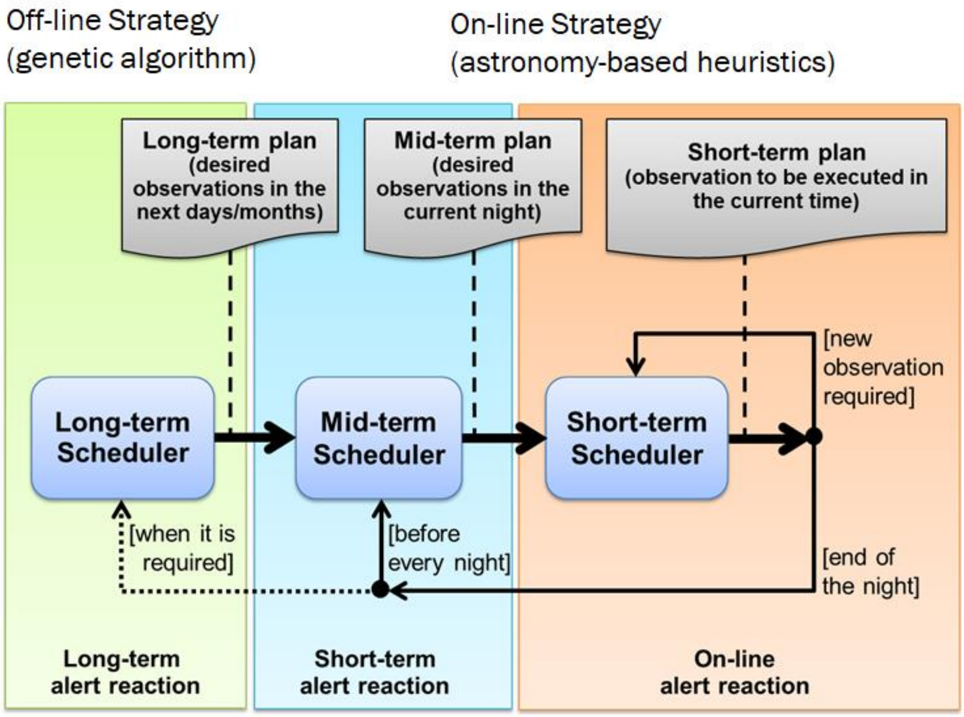

Figure 9 Combination of the three schedulers of STARS

LONG-TERM SCHEDULER

The long-term scheduler plans object observations with a time scope of several months. It takes into account the constraints that can be predicted beforehand, except those related to conflicts between objects (e.g. overlapping). A procedure based on GAs or MOEAs is applied to identify the best nights when each object should be observed by optimizing a specific soft constraint. The resulting plan provides a list of potential observation dates for each object. The long-term scheduler is re-run periodically to take into account observations previously done, thus counteracting any effect that unexpected situations may have on the optimization. The execution of this scheduler is not time-critical because it is run daily before the start of telescope operations and it can be used as a standalone planning tool for any observatory.

The design of the MOEA used in the long-term scheduler is defined as follows:

- The individual genotype uses a binary encoding that represents whether the target is planned on a given night. Each individual consists of N genes $\{o_1,...,o_N\}$, where N is the cardinality of the set of nights (N) that the long-term scheduler considers (e.g. the number of nights in three months) and o_i corresponds to night i. Moreover, the o_i value has to be 0 or 1, where 0 indicates that the target is not planned in the corresponding night and 1 indicates that the target is planned. The order of the targets in the genotype indicates a temporal sequence from the first night to the last night considered. The initial population is built by creating N_i new individuals assigning to each allele a value of 0 or 1 with a probability of 0.5. This representation does not allow unfeasible individuals.

- The selection, crossover, and mutation operators are defined as follows: in the case of the long-term scheduler, a mutated gene g' is obtained by negating g (i.e. alleles that are 0 become 1 and those that are 1 become 0).
- The long-term scheduler has the optimization goal of identifying the nights when an object should be observed according to the observing cadence constraint. It is commonly desirable to observe the targets when they are near their culmination. The optimization objectives promote the observation of an object near meridian crossing at the same time that maximizes the number of observations carried out. Thus, two objective functions are defined, respectively, to promote the observations of an object near meridian crossing and to promote the number of observations of the object.
- The most suitable solution is the individual that has the lower average of objectives from all the individuals in the first Pareto front, i.e., non-dominated individuals; the defined objectives have to be minimized to be optimized.

MID-TERM SCHEDULER

The mid-term scheduler plans the observations that should be executed during a specific night by optimizing some suitable soft constraints, and according to the results of the long-term plan. Moreover, the resulting mid-term plan fulfills all the hard constraints that are predictable, including those related to conflicts between objects. The execution of this scheduler is not time critical because it can be run before the start of the nightly telescope operation, so it can be used as a standalone scheduling tool for any observatory.

The MOEA used in the mid-term scheduler is designed as follows:

- The proposed individual genotype is made up of double numbers that represent the starting time of the observation of the targets. Each individual consists of T genes $\{o_1;... o_T\}$, where T is the cardinality of the set of targets to be planned (T), and o_i corresponds to target i. Moreover, the o_i value has to be between the range $[w_{st}, w_{et}-d_t]$, where w is a random uniform window in W_t, which are all the visibility windows in the night for target t; w_{st} is the Julian day of the starting time of window w for target t; w_{et} is the Julian day of the ending time of window w for target t; and d_t is the estimated integration time in Julian days for target t. Moreover, o_i can have a value of -1 indicating that target i does not have a starting time assigned (i.e., it is not planned). The order of the targets in the genotype does not indicate a temporal sequence, but it is only the order of the targets in the input data. The temporal sequence of targets is defined by the alleles because they indicate the starting time assigned to each target. For instance, a target in position i of the genotype can be planned in a time window previous to the time window of target i-1. The initial population is built by creating N_I new individuals assigning to each allele o_i a -1 value or a value in the range between $[w_{st}, w_{et}-d_t]$ following a uniform distribution. The process to build each individual is based on placing the observations of the targets, selected in random order, and avoiding overlaps. In case of overlapping, the target is unplanned (i.e., a value -1 is assigned).
- The GA process is roughly based on applying selection, reproduction (crossover), mutation, and replacement operators for several iterations (Goldberg 1989) (Freitas 2002). In the mid-term scheduler, the mutation operator alters a gene g' by changing its allele with a value inside the potential time windows of the corresponding target. Thus, a mutated gene g' changes its allele with a random uniform value in the range $[w_{st}, w_{et}-d_t]$ and -1 (i.e., g'=μ). In this case, the crossover of two feasible individuals can generate unfeasible offspring due to overlapping and the mutation of a feasible

individual can also generate an unfeasible solution. This is solved by a repairing procedure devoted to obtain feasible new individuals, as the next point explains.

- An individual represents the time windows assigned to target observations, but it does not consider the slew time between two observations. Thus, this aspect has to be considered to obtain the final planning codified by each individual. This modification can produce an unfeasible individual because it can have conflicting observations (i.e., presence of overlaps in the observations). There are two ways to obtain an unfeasible individual that requires repair during the GA process: 1) the individual is overlapping between two or more observations, and 2) there is overlapping between two or more observations when slew time is added to each observation. We may find that it is necessary to repair the individuals after the mutation process in order to obtain feasible individuals. Thus, the main idea of the repair operator is to solve all overlaps in the individual by `unplanning' conflicting targets. The unplanning of one target can solve overlaps between several targets.
- The optimization goal of the mid-term scheduler is to plan the selected objects according to two objectives related to the observing time and observation deviation soft constraints with the aim of minimizing the instrument idle time (time at night during which the instrument is not acquiring scientific data) weighted with the priority completeness of the targets, and mitigating the problem of scheduling the objects that require longer observations. Consequently, the mid-term scheduler optimizes two functions: one function to promote the time scheduled for observations of objects near meridian crossing according also to their priority, and a second function to promote a proper distribution of the observations of the objects with the same priority.
- The same strategy as in the long-term scheduler is applied to select the most suitable solution: the most suitable solution is the individual that has the lower average of objectives from all the individuals in the first Pareto front, i.e. non-dominated individuals; the defined objectives have to be minimized to be optimized.

SHORT-TERM SCHEDULER

The short-term scheduler computes the next observation to be executed during the night by optimizing some soft constraints and by considering all previous observations. Moreover, the selected observation fulfills all the hard constraints, Thus, this scheduler reacts to immediate conditions (weather, errors, delays, events). Unlike the long-term and mid-term schedulers, the short-term scheduler is time critical because it has to select an observation in a short time (in the order of a few seconds). For this reason, in order to avoid intensive calculations, it repairs the night schedule obtained by the mid-term scheduler using astronomy-based heuristics instead of using a GA or MOEA. For this purpose, it uses a dispatcher algorithm based on target ranking. This ranking is based on astronomical heuristics, and the targets are sorted according to the first rule, then the second rule, and so on.

The following is an example on how these astronomical heuristics are defined in the case of the CARMENES instrument (García-Piquer 2017). In the CARMENES Scheduling Tool (CAST), the short-term algorithm is called after the end of an observation and it repairs the night schedule obtained by the mid-term scheduler (Akturk 1999) using astronomy-based heuristics (Giuliano 2007) instead of using a GA, in order to avoid intensive calculations.

First, the algorithm removes all objects whose assigned observation period ended before the current time from the midterm plan and selects the next target. For this target, the code computes the slew time of the telescope. The target observation obtained for the mid-term plan is adapted (i.e., advanced or delayed) according to the current time, computed slew time, and integration time. This observation is only selected if it fulfills the hard constraint requirements until the end of the observation. Otherwise, it is discarded and the gap between the current time and the start of the next observation in the mid-term plan is filled. The filling process sorts all the observations that (1) are not already in the mid-term plan, (2) fulfill the hard constraints during the entire window, and (3) can be completed in the available time. This ranking is performed according to several criteria described below. Finally, the first observation in the sorted list is selected as the next observation. The filling process is repeated until the gap is filled or there are no target observations left. Each target selected by the short-term scheduler is sent to the instrument control system, and the information on the success of the observation is stored in the database for use in subsequent scheduler runs. Besides, the sorted list of objects can be provided to the operator for override in case of need. The ranking of the targets is key in the process of repairing the mid-term plan by filling a gap between the last executed observation and the next observations recommended by the mid-term plan. This ranking is based on astronomical heuristics, and the targets are sorted according to the first rule, then the second rule, and so on. The defined rules are:

1. The number of times that the target is observed during the current night (smallest to largest).
2. The target is not in the remaining mid-term plan.
3. The priority of the target (largest to smallest).
4. The number of times that the target has been observed in the survey (smallest to largest).
5. The proximity to meridian crossing (largest to smallest).

The main idea of this process is to fill the gaps with interesting objects at the current time, according to the times that they have been observed, their priority, and the proximity to meridian crossing. Rule 2 is key to fill gaps without excessively affecting the mid-term plan, which has been globally optimized.

SCHEDULING APPLICATIONS

There are various different flavours of the scheduling problem, reflecting the diverse nature of the facilities to be coordinated and the many ways in which they can be expected to work together. Each flavour may require a different scheduling algorithm. This section describes some of the different types of application, and how scheduling has been implemented using STARS, or could be in the future.

SINGLE TELESCOPE

This is the classical situation, with an isolated facility comprising a simple telescope, possibly carrying a range of instrumentation. The achieved performance using STARS is described hereafter for three different projects. Two specific examples cover the efficient scheduling of a predefined survey; the CARMENES ground-based instrument (carmenes.caha.es) and the ARIEL ESA space mission (ariel-spacemission.eu). An additional example of a robotic telescope, the TJO (www.oadm.cat) illustrates the scheduler response to a multi-purpose open observatory that is constantly updated with new user requests.

CARMENES

The CARMENES instrument is an optical and infrared high-resolution spectrograph that started a survey of about 300 M-dwarf stars in search of terrestrial exoplanets in September 2016. Targeted spectroscopic exoplanet surveys face the challenge of maximizing their planet detection rates by means of careful planning. For a large planet survey, the number of possible observation combinations, i.e., the sequence of observations night after night, both in total time and number of targets, is enormous.

Sophisticated scheduling tools and the improved understanding of the exoplanet population must be employed to investigate an efficient and optimal way to plan the execution of observations. Evolutionary computation techniques in STARS are used to create an automatic scheduler that minimizes the idle periods of the telescope and distributes the observations among all the targets using configurable criteria. The case of the CARMENES survey was simulated with a realistic sample of targets, and the efficiency of the planning tool was estimated both in terms of telescope operations and planet detection.

The scheduling simulations produce plans that use about 99% of the available telescope time (including overheads) and optimally distribute the observations among the different targets. Under such conditions, and using current planet statistics, the optimized plan using this tool should allow the CARMENES survey to discover about 65% of the planets with radial-velocity semi-amplitudes greater than 1 m/s when considering only photon noise.

In general, the scheduler for CARMENES is focused on optimizing the three aforementioned soft constraints: observing time, observation deviation, and observation sequence. The first two constraints maximize the use of the telescope and the instrument while the last constraint is included as an optional condition to increase the scientific return.

A set of simulations of the CARMENES survey are used with the aim of analyzing the efficiency of STARS in the use of resources. For a quantitative analysis, different metrics are defined related to the use of the telescope and the instrument:

- The fraction of targets that are planned by the scheduler.
- The total number of observations.
- The fraction of available time that the telescope is operating.
- The fraction of time during which the instrument is performing science observations.
- The fraction of overhead time.

To follow the real survey as closely as possible, the following procedure is adopted. The long-term scheduler has a scope of six months and is executed every three months during the survey. The mid-term scheduler is executed every day taking the observations acquired during the previous night into account. The short-term scheduler is executed "on the fly" each time a new observation is required and takes into account the mid-term plan, the observations already carried out during the night, and any variation on the weather or instrument conditions.

In terms of parametrization, the long-term and mid-term schedulers have several parameters related to the GA as explained before. In particular, the configuration parameters used in the experiments are related to the number of generations of the evolutionary algorithm, the number of elements in the initial set of solutions and in subsequent generations and to the probabilities of selection, crossing, and mutation.

Because GAs are stochastic methods, STARS is executed 50 times with different random seeds with the aim of avoiding any bias in the results due to convergence to local minima. Hereafter, each of these executions is referred to as a trial. Table 1 summarizes the parameters of the simulations and the results of the metrics used to evaluate the efficiency.

STARS parameters[a]	
Days planned	1096
Total targets	309
Total observable time	10 703.05 h
Unfavorable weather time	4300.23 88 h
Available time for observations[b,f]	59.82 0.82% (6402.81 h)
Execution time	23.85 0.11 h
Metrics[a]	
Planned targets[c]	100 0%
Observations done	20827 293
Working time[d,f]	99.05 0.06% (6342.03 h)
Tracking time[e,f]	84.18 0.03% (5338.77 h)
Overhead time[e,f]	15.82 0.03% (1003.22 h)

Notes. (a) The uncertainties are computed as the standard deviation of 50 random trials. (b) Ratio of time available for observations, excluding bad weather time. (c) Percentage of the targets that are planned. (d) Ratio of time with scheduled telescope operations with respect to the available time for observations. (e) Ratio of time with respect to the working time. (f) The number of hours is indicated in parentheses.

Table 1 STARS parameters and mean values of the metrics.

From the simulated weather statistics, the average usable time for observations at the Calar Alto observatory (CARMENES is hosted at the 3.5 telescope of CAHA) is around 60% of the total night time. Observatory statistics actually indicate that 70% of the nights are useful according to meteorological variables (Sánchez 2007 & Sánchez 2008). However, an additional up to 20% of lost time from cloudiness or technical issues was considered. Thus, in three years, there are about 6400 h during which observations can be scheduled. The results from the simulations presented in Table 1 show that all required targets are always planned. Around 21000 observations are scheduled, occupying 99% of the good weather time. The breakdown of this time indicates that 84% corresponds to the telescope collecting photons and 16% is spent during slews to new positions. This means that the instrument is collecting photons during approximately 1700h per year. The simulations show that the scheduler can optimize the observing time of the telescope by selecting the best targets to observe according to environmental conditions at each time.

Additionally, it is important to distribute the observations of all the targets equitably. This is the second soft constraint that the scheduler must optimize. Figure 10 shows the number of times that each target is observed. On average, each M dwarf in the sample is observed 67 times during the three-year survey. The standard deviation of the number of observations between targets is 3, which indicates that the resources are reasonably distributed among the different targets. Only a few of the targets have a number of observations significantly below the average, but this is because of their very limited visibility during the year (e.g. low declination and faintness). This means that all targets have a high number of observations and, if necessary, this number can still be increased if some of the targets are discarded during the initial sample clean-up (very fast rotators, active stars, spectroscopic binaries that passed the filters, etc.). On the other hand, the simulations with STARS could also help in the optimization of the sample since any target for which a minimum number of observations is not reached can be rejected, or upper limits to the number of observations per target can be set.

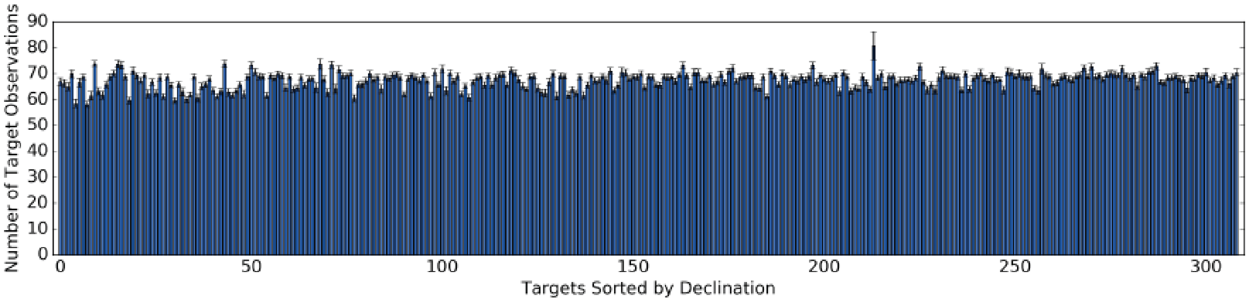

Figure 10 Number of observations scheduled for each target. The horizontal axis represents the identifier of each one of the 309 targets used in the simulations sorted by increasing declination. Declination ranges from -21 to +83 deg in the sample. The bars show the average and standard deviation values of the 50 executions.

The simulations using the STARS scheduling tool for CARMENES show that the instrument can carry out a spectroscopic survey of a large sample of M-dwarf stars in three years, and it is possible to optimize the survey planning by minimizing idle instrument periods and fulfilling the science objectives in an efficient manner to maximize the scientific return. The simulations with a list of 309 M dwarfs show that STARS optimizes the use of the instrument and can yield over 60 observations per target, fulfilling all constraints. Besides, an important advantage of using GAs in the automated scheduling process is that they guarantee a feasible, consistent, and near-optimal solution according to the constraints defined in the problem. STARS can be adapted to the needs of the astronomers in different situations during the survey and the night operation. Moreover, because it is based on a multi-objective optimization problem paradigm via multi-objective evolutionary algorithms, it is able to find optimal solutions with a trade-off between several criteria that may conflict with each other. This aspect provides an important advantage to STARS in order to guarantee the efficiency of the solutions in terms of use of resources (e.g. telescope operations) and in terms of science (e.g. planet detection).

ARIEL

The Atmospheric Remote-Sensing Infrared Exoplanet Large-Survey (ARIEL) (Tinetti 2016) is the fourth medium-class mission in ESA's Cosmic Vision programme, to be launched in mid 2028. ARIEL will be the first dedicated mission to investigate the physics and chemistry of exoplanetary atmospheres, by observing their transmission and emission spectra during the transits and occultations of exoplanets, respectively. The mission goal is to survey a few hundred exoplanets orbiting nearby stars, with sizes from Jupiter to a few Earths, with a precise low-resolution spectrograph.

The strategy of the mission, i.e. the observations of exoplanet transits and occultations, is a major challenge. Planning such observations, which are time constrained following precise ephemerides, for hundreds of exoplanets that should be observed several times to reach high-precision, demands a thorough planning process (Garcia-Piquer 2017).The efficiency of the mission must be maximised under multiple constraints; the target visibility, the duration of the transit or occultation event, the number of repetitions needed. Although past surveys have generally used a manual approach to planning, the process of analyzing all the survey information and computing all the possible combinations becomes unaffordable for human planners.

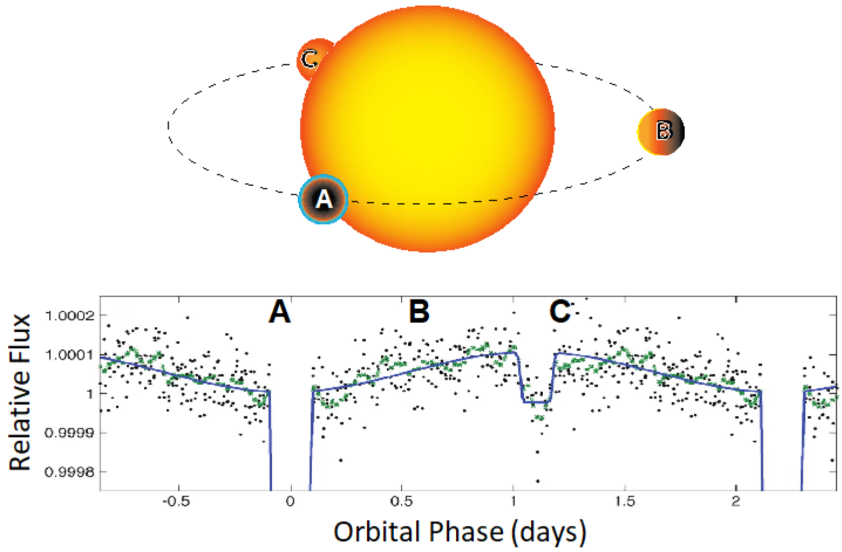

Figure 11 Transit and occultation events considered by the ARIEL mission.

The definition of the scheduling problem for ARIEL includes both the operations that must be performed and the constraints. Regarding operations, the scheduling of an astronomical survey does not consist only of planning observations of the stellar targets, but also of other necessary satellite operations. An efficient mission planner should take them all into account, especially those that cannot be carried out simultaneously. In the case of the ARIEL mission, the additional operations include science and calibration observations, and station keeping tasks. Downlink communications with ground-based stations are expected to be simultaneous with observations.

- Science observations consist of the spectroscopic follow-up of exoplanet events: transits and occultations. The duration of an event is the time between the first and fourth contacts (T_{14}) between the star and the planet. Observations are centered on the transit but they are extended before and after the contacts in order to precisely measure the transit depth; a total of $2.5T_{14}$ is assumed for each event. The total duration and its ephemerides depend on the physical and orbital properties of the planetary system, but they can be measured beforehand and are provided as an input parameter to the scheduler.

- Calibration tasks consist of the spectroscopic follow-up of some specific very stable G-type stars. The goal is the monitoring of the instrumental response on the typical timescales of transits. Following the ARIEL design, the mission planning tool has to plan 1 hour observations of these targets (hereafter, short calibration) every 1.5±0.5 days and 6 hour observations (hereafter, long calibration) every 15±5 days, giving priority to long calibrations over short ones. Period ranges are allowed in order to increase the flexibility of the planning of these tasks. A list of 536 G-type stable stars is provided so that there is one of these stars within 5 deg of any planet target.
- Station keeping operations are defined to keep the spacecraft in the assigned orbit. The sequence of these operations cannot be established precisely in advance, but as a conservative approach, time slots of 8 hours every 28±3 days are reserved in the planning.

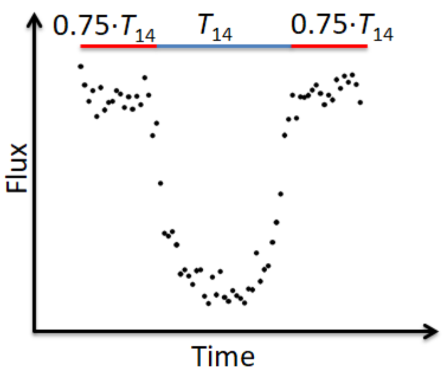

Figure 12 Transit event duration

According to the ARIEL design, five hard constraints are identified for the scheduling problem.

- *Orbital* constraints are the satellite orbit and attitude which set the visibility of targets. In the case of the ARIEL mission, an L2 orbit with a maximum wobble angle of 25 deg is assumed. Therefore, at any time the telescope can point to any star farther than 65 deg from the Sun-Earth direction.
- *Transit* constraints are defined by the ephemerides of the planetary system events. Targets must be planned only when the planet transits or is occulted by the star.
- The *Target Completeness* constraint is related to science observations. In order to satisfy the precision and signal-to-noise ratio requirements of the planetary spectra, piling-up of several transits or occultations will be needed for each target. Consequently, only the observation of all required events for a given target is useful, and any target that cannot be completed during the mission lifetime is not considered in the scheduling.
- The *Slewing* constraint accounts for the time to point to a particular target and acquire data. This mainly depends on the slewing rate of the satellite, that according to

mission requirements is assumed to be 4.5 deg/min. Besides, 5 minutes are added for stabilization of the system.
- The *Overlapping* constraint takes into account that none of the operations defined above can be done simultaneously. This means that the scheduler must plan the operation tasks avoiding overlap between different target observations, calibrations and station keeping operations, including the slew time to point to a new target.

On the other hand, the goal of the soft constraints is to optimize the survey planning. These constraints are not required to be fulfilled in order to obtain a valid plan, but they are used to prefer some solutions over others. In the case of ARIEL, these constraints are related to survey and scientific efficiency.

- Observing time optimization is one of the goals of the mission planning. The time in the plan during which the telescope is observing objects should be maximized.
- Number of completed targets should be maximized in order to increase the scientific efficiency of the mission, taking into account the priority of each kind of planet.

The optimization of these soft constraints will produce a long-term mission plan (LTMP) with a large number of observations as well as a large number of completed targets. Nevertheless, the soft constraints can be modified through the characteristics of the optimization method used in order to provide more flexibility to the planning tool.

In scheduling the ARIEL survey, we identify two main aspects based on the problem conditions: 1) the optimization of the positioning of satellite calibrations or station keeping operations so that only lower priority targets are restricted, and 2) the optimization of the observation scheduling of each target event, avoiding overlapping and optimizing specific objectives.

Taking into account these considerations, different steps for the mission planning optimization are considered, together with the necessary techniques implemented in the STARS framework. The steps are:

1. Compute the time windows where the targets can be observed.
2. Remove targets that cannot be observed a certain fraction of the requested times.
3. Optimize the satellite operations by minimizing the potential overlap of calibrations and station keeping maneuvers with science targets. This optimization is based on GAs.
4. Allocate exoplanet observations, avoiding as much as possible overlap between different targets, and maximizing the number of completed targets. In this step, the slew time of the telescope is taken into account. It is also solved with a GA. At the end of this process, the successful `individual' is a feasible plan. However, in this process the fulfilment of the Target Completeness constraint is not forced in order to allow the GA to satisfy it during the optimization process. If at the end of the genetic process this constraint is not satisfied, the resulting plan will be modified to fulfil it. The fitness function that evaluates each individual considers two objectives; minimizing the number of priority target events that are not planned, and minimizing the time that the telescope is not working.
5. The long-term plan so obtained may include targets with a number of re-visits smaller than the requested observations given in the target list. In order to increase the efficiency of the survey, these targets can be removed from the mission plan and the

time used for further observations. Therefore, this phase of the scheduling is devoted to filling the gaps in the long-term plan after uncompleted targets are dropped. It tries to plan as many new target events as possible distributing them between all the remaining targets in the mission plan. It must be emphasized that if this process is not applied, calibrations and station keeping could be planned in this step to fill these gaps. After this step, a final mission plan is suggested. It includes only completed exoplanet targets, and calibrations and station keeping operations with no overlap. The survey optimization functions are then recomputed to be used as indicators of the survey efficiency.

The list of targets is one of the key inputs of the scheduling problem because the sky visibility or the time of planetary events limits the windows when they can be scheduled. To simulate the ARIEL survey a sample of 1102 exoplanet systems provided by the ARIEL Science Team (Zingales et al., priv. comm.) are used. This list includes 122 known planets and 980 simulated following a conservative assumption on the rate of exoplanets per star, and future facilities. It provides the coordinates, the ephemerides of transits and occultations and the number of events necessary for each target. Actually, three different subsamples are defined according to their level of precision and resolution. From the scheduler point of view, only the number of events required is important. The list of targets also ranks all exoplanet systems with a priority normalized to 1 according to their properties. A total of 459 objects have the highest priority.

Using the sample of ARIEL targets, we have computed a mission plan fulfilling all hard constraints and optimizing the soft ones. Following the mission design, we assume a satellite lifetime of 3.5 years after the commissioning phase, during which target observations, calibrations and housekeeping manoeuvres must be scheduled. As a measure of the survey efficiency, we compute the total number of surveyed exoplanet systems (including only completed targets), the fraction of time allocated to each operation and the time remaining between observations. Since the GAs compute an optimized plan starting from initial random solutions, a total of 25 simulated mission plans are obtained and average values calculated. We have run different simulations checking the impact of selected parameters. In terms of computational time, each plan is obtained in about 7 minutes. Table 2 shows the results for the different cases studied.

Survey case	Planned targets	Working time			Waiting time
		On targets	Slewing	Other	
Case 1 (459 targets)	$99.3 \pm 0.2\%$ 458 ± 1 systems	$52.08 \pm 0.64\%$ 15967 ± 196 h	$2.44 \pm 0.03\%$ 747 ± 9 h	5.46% 1673 h	$40.02 \pm 0.67\%$ 12273 ± 204 h
Case 2 (1102 targets)	$86.6 \pm 0.6\%$ 953 ± 6 systems	$64.66 \pm 0.45\%$ 19825 ± 137 h	$3.23 \pm 0.04\%$ 990 ± 11 h	5.44% 1669 h	$26.67 \pm 0.47\%$ 8176 ± 144 h
Case 3 (1102 targets)	$86.7 \pm 0.5\%$ 955 ± 5?systems	$73.85 \pm 0.21\%$ 22640 ± 65 h	$3.91 \pm 0.04\%$ 1199 ± 11 h	5.44% 1669 h	$16.80 \pm 0.23\%$ 5152 ± 69 h

Table 2 Results obtained for the ARIEL mission and considering different test cases.

In Case 1 we took into account only the 459 highest priority planetary systems. This simulation reveals that almost all the highest priority targets can be completed in 3.5 years up to the highest precision level. This means a total of 211718 events are scheduled, taking about 60% of the mission plan, and leaving 40% of free time.

In Case 2 we used the whole list of 1102 targets with their respective priorities. The completion of all the observations in the list would require 31250 hours, which actually is about 2% more time than that available in 3.5 years. Obviously, the final mission plan will not complete the whole sample because of the difficulty in avoiding overlap between targets requiring several observations, but it produces an efficient mission plan including 278122 events observed. A total of 9536 exoplanets would be surveyed. Operations would take in total about 22500 hours of the mission lifetime, leaving only 27% of waiting time between targets. This waiting time is due to the removal of targets that cannot be completed at the required precision level.

In Case 3, we show that these waiting time periods can be filled with observations of other targets either uncompleted or with more observations than required. In that case, the waiting time is only about 17% of the mission lifetime, and it is distributed in intervals of less than 6 hours (planetary events last longer than that for 80% of the sample). This fraction of waiting time is due to the difficulty of finding consecutive but not overlapping planetary events. However, it can be used to increase the follow-up of out-of-transit phases, which could be very useful to better constrain transit depth, stellar variability or study exoplanet phase curves.

The results obtained support the conclusion that the proposed scheduler technology is robust and can function in a variety of scenarios, offering a competitive performance which does not depend on the collection of exoplanets to be observed. We demonstrate that STARS makes it possible to fulfill the ARIEL mission requirements by observing more than 500 exoplanets, 440 of them with high resolution. The proposed process is based on GAs but other global search algorithms can be applied following the same steps. Automatic scheduling tools have other advantages as well. From the operational point of view, they guarantee a feasible, consistent and near-optimal solution that fulfills all constraints, and which can be obtained with a few hours of computation. They can also be used to estimate the efficiency of the survey and, therefore, to study the impact of different parameters or find which targets are most restrictive.

Joan Oró robotic Telescopes

The Joan Oró Telescope (TJO) at the Montsec Astronomical Observatory (www.oadm.cat) is a 1m-class open and multi-purpose telescope working under completely unattended control since 2007. The TJO has operated using the STARS framework since early 2017 and the autonomous scheduling tool is fully integrated into the proposal handling procedures.

The scheduler for TJO is focused on optimizing the telescope operations (science and calibration tasks) and the soft constraints (observing time, observation sequence and completed targets/proposals), in order to maximize the use of the telescope and the instrument while increasing the scientific return. STARS takes into account some of the hard constraints described above (visibility, target completeness, slewing and overlap of scientific and calibration tasks), but it also covers slightly more complex observational strategies. In particular, the STARS optimization algorithm considers observation sequences, where observations of different targets or of the same target must follow a predefined pattern. The latter means that all required events for a given sequence must be scheduled, and any sequence that cannot be completed during the night, week, month, etc. is not considered.

The open-time proposals, DDT programmes and strategic surveys are combined by considering suitable prioritization and optimization constraints, in order to provide the most suitable scheduling block whenever the telescope control system requests it from the scheduler. For a quantitative analysis of the STARS performance, different metrics are defined related to the use of the telescope and the instrument:

- The fraction of targets that are planned by the scheduler.

- The total number of observations.
- The fraction of available time that the telescope is operating.
- The fraction of time during which the telescope is performing science observations.
- The fraction of overhead time.

The long and mid-term scheduler has a scope of six months and is executed every day, taking the observations acquired during the previous night into account. The short-term scheduler is executed "on the fly" each time a new observation is requested and takes into account the mid-term plan, the observations already carried out during the night, and any variation on the weather or instrument conditions.

The scheduling procedure starts by computing the visibility windows for the uncompleted proposals, and this is based on the target altitude, Moon conditions and predefined event ephemeris, among other constraints. Figure 13 illustrates the time windows for different targets, whereas Figure 14 shows the distribution of Scheduling Blocks from different proposals as distributed by STARS.

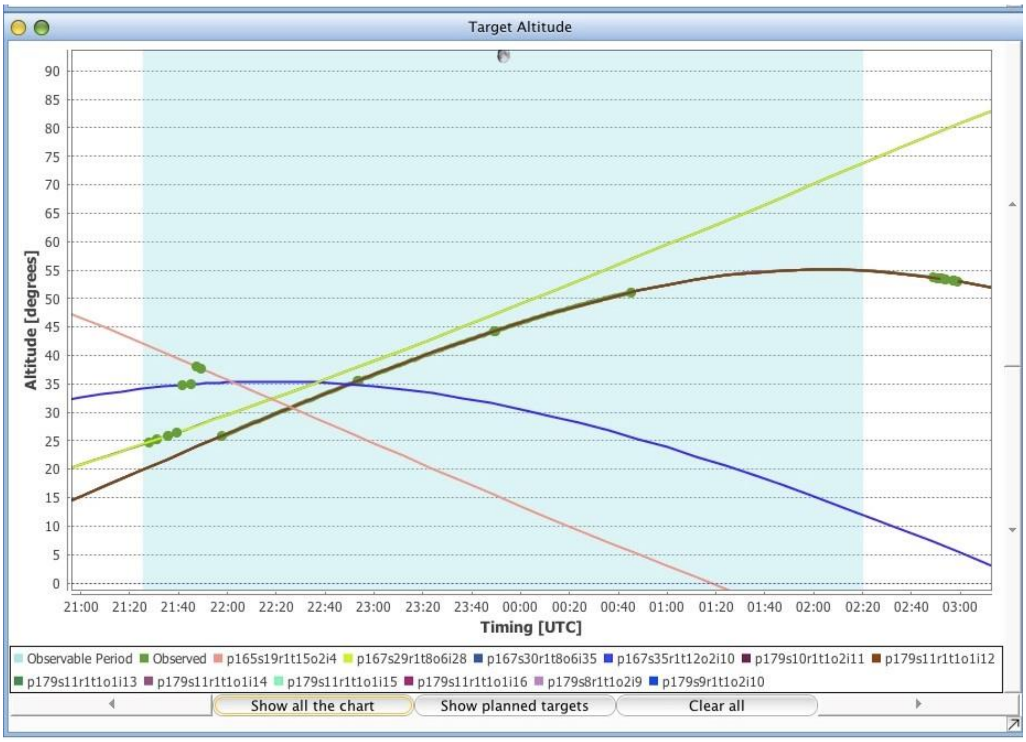

Figure 13 Observable time windows for different targets

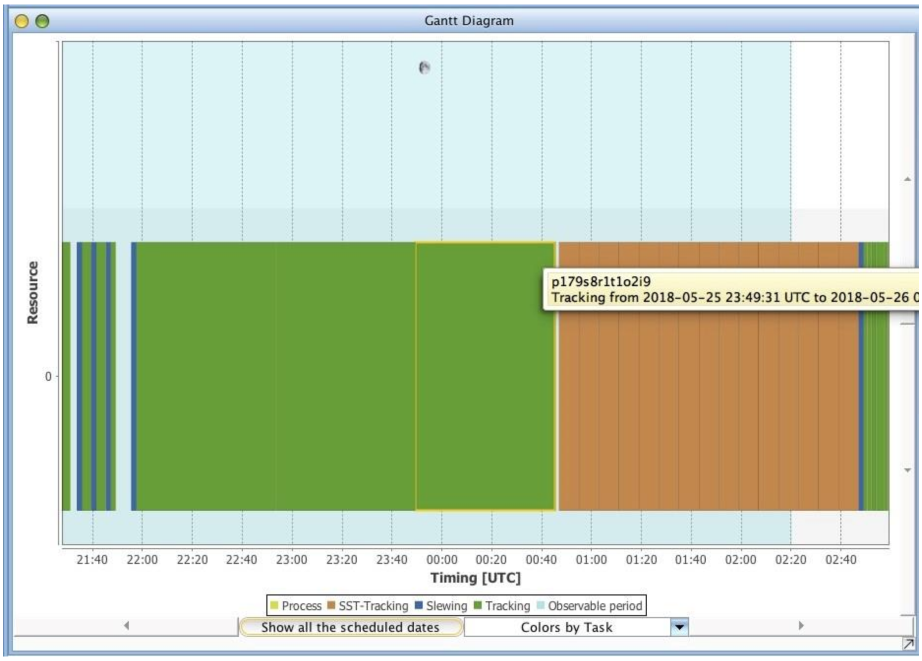

Figure 14 Night schedule

SINGLE OBSERVATORY WITH SUB-ARRAYS AND MULTIPLE SITES

CTA

The specific example of the Cherenkov Telescope Array (CTA) illustrates how STARS is used to extend the scheduling optimization to a facility that is composed of sub-arrays and has more than one site.

CTA (Acharya 2013) is the next-generation ground-based very-high-energy gamma-ray observatory and will deploy about 100 (20) telescopes on an area of about 4 km^2 (0.4 km^2) on a southern (northern) site. The CTA consortium is currently in negotiations to establish the location of the southern and northern installations in Armazones, Chile, and La Palma, Spain, respectively. CTA is a big step forward in the field of ground-based gamma-ray astronomy, not only because of the expected scientific return, but also due to the very large scale of the instrument to be controlled. In the southern array, at least three types (two types in the northern array) of imaging atmospheric Cherenkov telescopes will be deployed. These telescopes will have a range of sizes (with typical reflector diameters of 23 m, 12 m, and 4 m) and their cameras will comprise 1000-10000 pixels, acquiring data in the kHz domain. The CTA performance requirements and the inherent complexity of operating, controlling and monitoring such a large, distributed, multi-telescope array leads to new challenges in designing and developing the CTA control software.

The CTA scheduler must provide an automatic schedule of observations to each subarray that satisfies hard constraints and is highly optimized with regard to soft constraints. A near-optimal solution is obtained with STARS, coordinating the multiple facilities CTA North and CTA South. It is also highly configurable to be used at different sites, under different weather conditions, with any particular sub-array definition, following particular cadence and observation strategies, proprietary time and observational constraints. Long-term, mid-term and short-term cycles are used to simulate the performance required for the future operation of this large infrastructure.

The following is a description of the scheduling conditions for CTA. Different operation tasks are considered (science, calibration and maintenance), observation types (including transient alert follow-up), and modes with scenarios specific to this project (e.g. compact/full array and sub-array modes, or convergent/divergent modes), two sites (North/South) and tens of telescopes to be coordinated. At the current definition phase, CTA considers that each subarray contains a number of telescopes of the same or different types, and telescopes can be shared between subarrays. Those subarrays will then exclude each other from simultaneous observation, i.e. Subarray 1: all telescopes vs Subarray 2: large and medium-size telescopes - LST and MST, respectively.

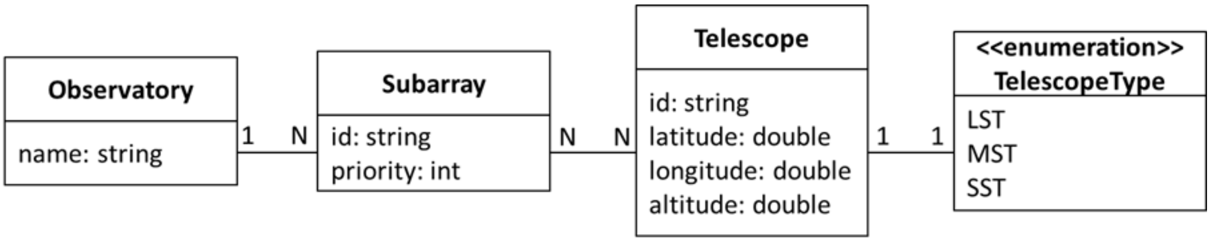

Figure 15 CTA model describing the possible sub-array configurations.

The CTA has some hard constraints in common with other projects and available in the STARS general design: visibility, dark hours (global) as a function of the Sun altitude and Moon conditions (solar horizon (-18º), Moon phase and Moon distance to the target position), maximum Zenith Angle for the target/SB. Sub-array configurations add complexity that translates to the following resource hard constraints: subarrays that share telescopes cannot observe at the same time and each target is observed by the assigned subarray. Soft constraints are very similar to those defined in other projects: maximize observation time each night, maximize completion of proposals, minimize slew time each night (time blocks will increase due to consecutive observations). Two optimization objectives are used: maximize observation time or proposal completion, and minimize slew time.

The following scenarios are used to simulate the operation of CTA and check the STARS performance in the optimization process:

- Three key programmes and surveys (their details cannot be disclosed for CTA Consortium confidentiality reasons) in 1 and 3 years full proprietary time are used as the simulation sample.
- Real weather conditions are used based on archival data using 2 years monitoring. The (configurable) conditions used to allow observations are the wind speed (< 36 km/h), the humidity (from 4% to 95% operational range), cloud conditions (only cloudless hours) and temperature (between -10 and 25 deg C).
- Observations with partial Moon light are not considered in the time computation.
- Time distribution (hours/yr): T_S = 2062.25 h, $T_{Moonless}$ ~ 1676 h, $T_{UseableMoonless}$ ~ 1284 h (after weather selection), $T_{Available}$ ~ 1220 h (CTA availability, according to reqs: 95% for science, 5% others).

The results obtained for a specific simulated scenario are given below. It has the following elements: site location is CTA South, 464 targets (requiring 3692.11 hours to complete all of them), 1 year scheduling and 3 years scheduling testing, 4 subarrays and 3 types of telescopes (LST, MST, SST), and targets are assigned to subarrays according their type. Table 3 shows the distribution of target sub-samples to different sub-arrays.

Target Type	#Targets	Subarrays Assigned
SURVEY 1	170	0 (LST + MST + SST)
SURVEY 2	231	1 (LST + MST)
SURVEY 3	20	2 (MST + SST)
SURVEY 4	5	2 (MST + SST)
SURVEY 5	38	3 (LST)

Table 3 Distribution of the simulated sample in different surveys

Simulations are used to check also the different performance of the GA and MOEA algorithms. MOEA proves to be more efficient based on those metrics that are more relevant for CTA, as is shown in Table 4 and Table 5 (where WT is the working time: slew time + observation time):

Subarrays	Required Time (h)	#Observations	Working Time (h)	Slew (% of WT)
0 (LST + MST + SST)	1557.75	865	313.88	8.14
1 (LST + MST)	200.01	602	220.28	8.90
2 (MST + SST)	550.00	719	261.81	8.46
3 (LST)	1384.35	1423	514.33	7.78

Table 4 Results obtained using the GA algorithm

Subarrays	Required Time (h)	#Observations	Working Time (h)	Slew (% of WT)
0 (LST + MST + SST)	1557.75	1108	376.86	2.00
1 (LST + MST)	200.01	660	230.13	4.40
2 (MST + SST)	550.00	800	271.45	1.76
3 (LST)	1384.35	1703	581.14	2.32

Table 5 Results obtained using the MOEA algorithm

A distribution of the SBs in the three different sub-arrays (1, 2 and 3) and the full array (compact configuration) considered in the test case is given in Figure 16 for illustration.

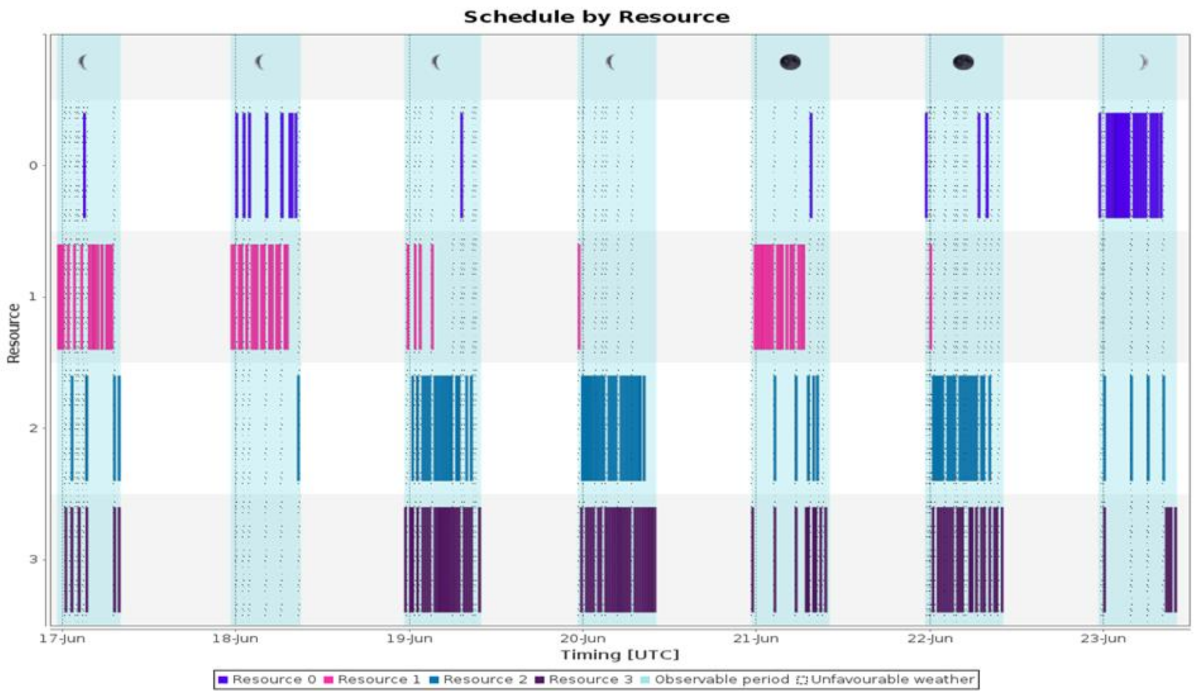

Figure 16 Distribution of SBs to the different sub-arrays (labelled as resource 0 for the compact array, and resources 1, 2 and 3 for the sub-arrays considered).

A different simulation for the same test case, 3 years of operation (2993.71 hours available) and based on MOEA algorithm, shows a performance that achieves the required science return with an optimum operation. Table 6 and Table 7 summarize the surveys simulated and the performance obtained respectively.

Target Type	#Targets	#Planned (#Completed)
SURVEY 1	170	170 (38)
SURVEY 2	231	231 (231)
SURVEY 3	20	20 (20)
SURVEY 4	5	5 (5)
SURVEY 5	38	38 (31)

Table 6 Survey configurations

Subarrays	Required Time (h)	#Observations	Working Time (h)	Slew (% of WT)
0 (LST + MST + SST)	1557.75	2012	685.76	2.20
1 (LST + MST)	200.01	697	243.80	4.70
2 (MST + SST)	550.00	1653	559.65	1.55
3 (LST)	1384.35	2913	992.24	2.14

Table 7 Results of the 3 year simulation (where WT is the working time: slew time + observation time)

Scenarios that consider scheduling proposals and surveys using CTA North and South sites in a coordinated manner have also been tested with STARS. These simulations are carried out by extending the applicability of the MOEA algorithms and the fitness function.

We simulated a complete decade of operation (from 1/1/2021 to 1/1/2031) with the full array/observatory (including North and South sites). Figure 17 illustrates the results obtained, where the fields of view from the North and South sites are drawn, together with all the requested SBs (pending, observed and in progress). The overlapping area of the two fields of view shows the region of the sky where observations can be taken with either array. This flexibility increases the chances to schedule high-priority programmes in case of bad weather conditions or technical issues affecting one site. On the other hand, the ability to schedule the same task at two different sites makes it necessary to have a global schedule view when running long and mid-term planning. Short-term reaction is a bit more complex, as it requires coordinating the local scheduling with the other site. Simulations are obtained assuming instant knowledge of the situation on both sites. However, real operations will be significantly different due to communication delays and a suitable model should be used in STARS to properly simulate this scenario where short-term coordination is required.

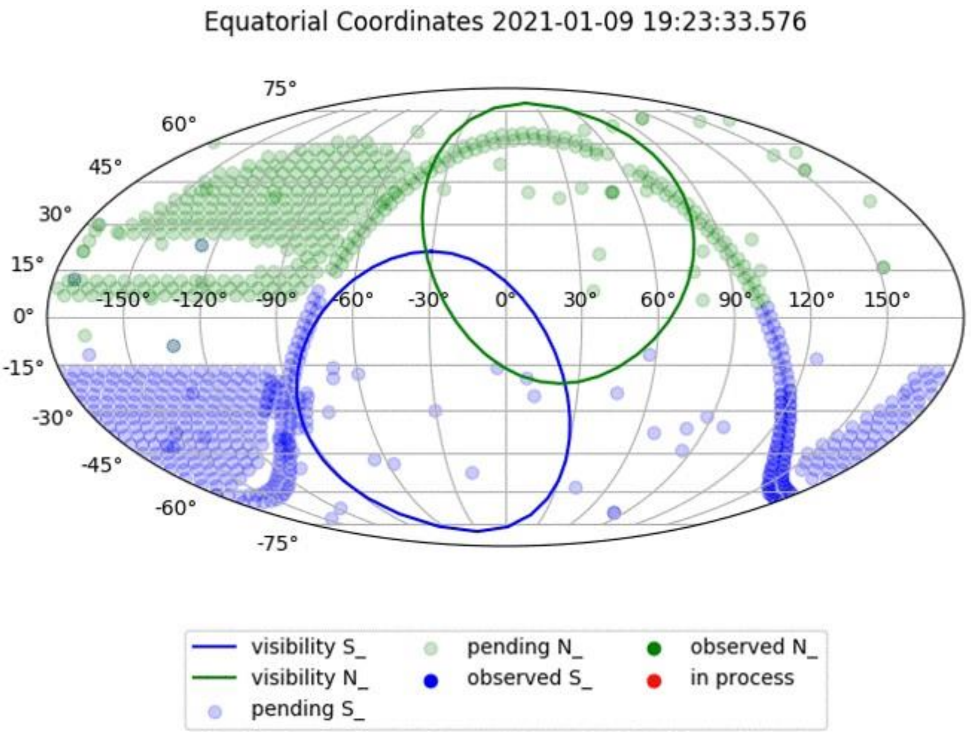

Figure 17 Simulation results.

The results obtained for each site and, in particular, the metric that measures the available time used for science observations, is illustrated in Figure 18. The ratio of time used per year decreases as the surveys are being completed, and remaining observations are extended several years due to the challenging scheduling conditions for some of the SBs.

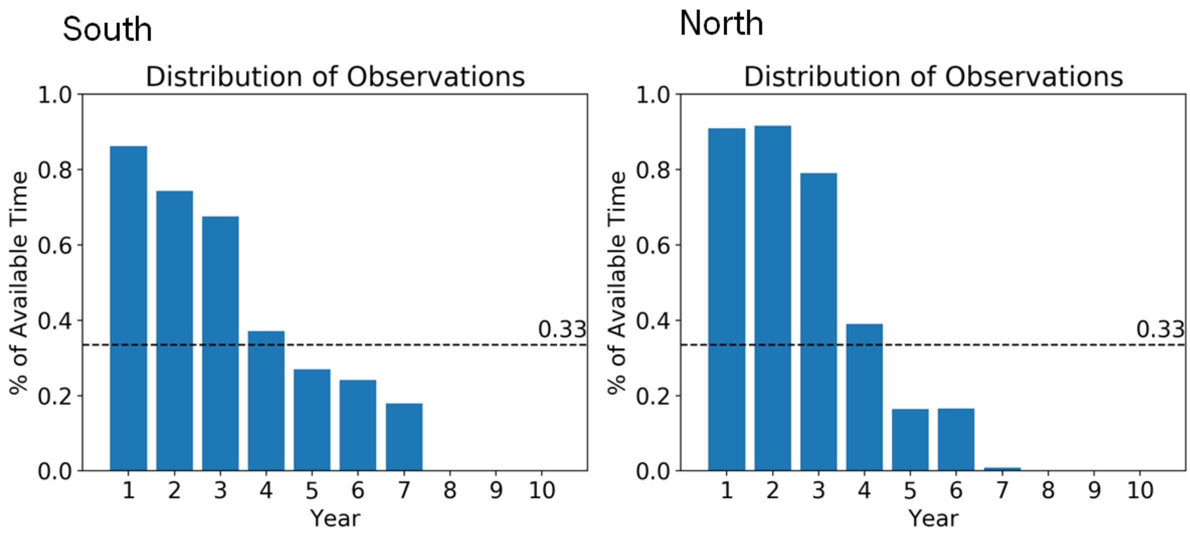

Figure 18 Available time used for science observations in CTA simulations using North and South sites.

MULTI-OBSERVATORY SCHEDULING

BILATERAL CROSS-FACILITY COORDINATION

The specific example is a scheduler that has been implemented to coordinate the execution of CTA and SKA surveys. The work is focused on proving the feasibility of multi-observatory scheduling coordination, and measured in terms of the capability to promote multi-messenger science based on ɣ-ray and radio astronomy data. The results of this work are not discussed here and will be presented in the ASTERICS deliverable 5.12.

BILATERAL CROSS-FACILITY COORDINATION WITH `ALERTS'

The CTA is a ɣ-ray observatory that will conduct surveys of the sky, generating alerts when it spots new sources. However, it will also be sensitive to alerts from other instruments and conduct follow-up observations in a routine way. The problem of addressing transient events is covered by the CTA description given in the previous section with the response to 'alerts' an added responsibility of the short-term scheduler.

The most important requirement is to filter the flood of incoming alerts (~ 1 million per night from the LSST alone) down to the small number that can be managed without disrupting the telescope's survey work. It is possible that much of the filtering will be done by AMON or similar, a service that correlates the alert stream from a single target. As correlated alerts accumulate, AMON will itself generate alerts that give a growing and more specific description of the nature of the target. This refinement process would be an important part of the necessary filtering as the CTA would actually be able to look out for alerts from AMON that describe a very specific type of target. Initial testing based on STARS and an estimation of the incoming transients has given a good performance, as shown in Figure 19. However, if alert filtering is aided by AMON or similar, then this will need to be considered as a mandatory complementary service for STARS to solve the scheduling of transients.

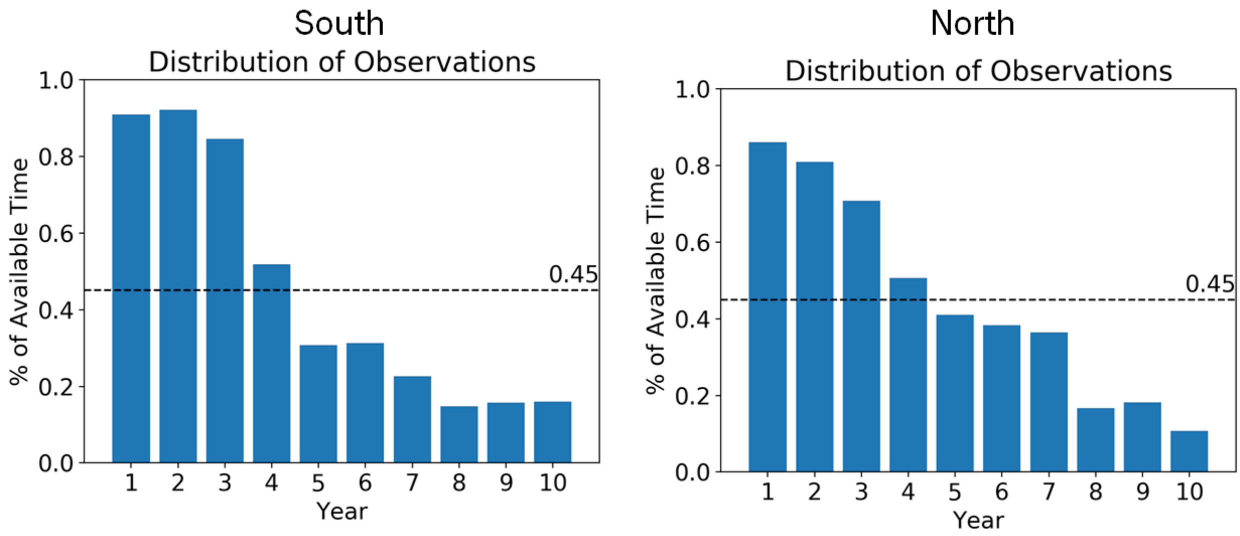

Figure 19 Available time used for science observations in CTA simulations using North and South sites, including transient alert estimation.

COORDINATION OF PRE-ALLOCATED OBSERVING TIME ON MULTIPLE FACILITIES

The ESA PLATO mission is an example of an observatory that will require a large number of ground-based facilities to confirm exoplanet candidates. These facilities will range from ESO large telescopes to smaller instruments with a heterogeneous night-time availability dedicated to follow-up activities. These will be triggered from a single service that, a priori, will know the availability and status of the full set of facilities committed to participate in the network. The architecture is under discussion, but the need has been clearly identified.

MULTILATERAL CROSS-FACILITY COORDINATION

An extensive analysis of this scenario will be evaluated in deliverable 5.12. We will examine various fundamental questions. What is the best way to optimize communication? Is loose coupling sufficient, where each facility would simply publish where it will be looking (and with which partners)? Or would it better to have a high-level infrastructure above the individual facilities, partly to consider submitted proposals and partly to implement the scheduling coordination required? How and where would such a high-level infrastructure be implemented? The Cloud?

XX. CONTROL SOFTWARE INFRASTRUCTURE

The effect of control software infrastructure on scheduling performance has been examined.

The ALMA Common Software (ACS), developed by ESO in collaboration with its partners, provides the infrastructure of the distributed software system of ALMA and other projects. ACS, built on top of CORBA and DDS middleware, is based on a Component-Container paradigm and hides the complexity of the middleware allowing the developer to focus on domain specific issues.

The scheduling function could fit in the infrastructure, interfaced by means of the publisher/subscriber pattern provided by ACS, as it allows specific events (e.g. Scheduling blocks, alerts) to be published and received through the whole system. The ACS framework supports three languages (C++, Java, Python) which can be used to implement components, publishers/subscribers, distributed error logging, alarm management, etc. Therefore, in the multi-facility observation coordination scenario, a component could be implemented using ACS to interface the scheduler to all facilities, where each facility had its own component.

In that scenario, the ACS provides the scheduler with the following services:

- The ACS can feed the scheduler, making use of different notification channels to receive and send different type of events as required.
- ACS distributed logging allows monitoring of the whole system from anywhere on the system, as it is developed on top of the CORBA Notify Service. This subsystem comes with a powerful GUI tool that allows log filtering in real time or offline.
- The ACS alarm system could be used to submit facility changes that should be taken into account by the scheduler. This system uses the CORBA Notify Service to distribute the alarms, which could be received and processed by a purpose-built scheduler client.

The ACS has some drawbacks and weaknesses when considering its potential use for the multi-facility operations:

- The publisher/subscriber pattern presents problems regarding the Notify Service asynchronous behavior, since the publisher is never aware of the subscribers connected. As a consequence, during operation, problems of missed events have been reported.
- There are some reported performance problems with the Notify Service in ACS, such as slow deliveries or Notify Service crashes, that could be worse in a multi-facility scenario due to scalability.
- The community around ACS is rather small since it has not been used for too many other projects, and this should also be taken into account if one expects to have new features in the future. A major drawback could be then the potential interest in the development and upgrade of such framework. On the other hand, ACS will not be discontinued as it is being used by ALMA, so its maintenance and operational life is guaranteed.
- Network security is not implemented in the ACS framework since its original design was conceived for isolated networks. Considering the level of criticality of the communications between infrastructures and the scheduler, that subject shall be addressed more in detail.

Regarding the publisher/subscriber pattern, some improvements could be proposed. The publisher/subscriber pattern presents an unbalanced information flow: while the ACS facilities send all kind of events to the scheduler with a high data rate expected, the scheduler calls facility components directly, or uses the ACS Notify Service, to deliver the selected SBs to the observatories at a much lower data rate:

- From scheduler to facilities: a priori, it would seem better not to use the Notify Service to send data from the scheduler to a facility. Because the amount of data transferred is expected to be much lower than the amount received, a point-to-point communication from scheduler to facilities could ensure reliable delivery of the information, instead of having to deal with information loss due to the asynchronous nature of the publisher/subscriber pattern.
- From facilities to scheduler: the notification channels could saturate due to the high data rate. An approach that might solve that problem would be to create a dedicated channel between both ACS components, using a well-known high-performance messaging library such as ZeroMQ or RabbitMQ in place of the CORBA Notify Service.

XXI. CONCLUSIONS

Efficient proposal planning and time scheduling are critical elements in optimizing the exploitation of large facilities and space missions devoted to astronomical research. Multiple projects have addressed the observation planning and scheduling problem with specific algorithms, although in classical observatories humans usually lead the decision making process still.

We have described a hierarchy of scheduling applications, starting from a single telescope and ranging up to multi-observatory coordination, which is currently the most complex problem for optimization. Algorithms have to be designed to suit the operational design of the facilities and the observation strategies of the scientific programmes. A specific framework called STARS has been presented, which includes an abstraction of the scheduling problem, and different algorithms (GA, MOEA, heuristics) that can be used at different facilities.

In addition, we have evaluated multi-facility coordination from the perspective of multi-messenger science. Multi-messenger science is opening a new era in astrophysical research and is identified as a major challenge for the operation of large facilities under construction. New trends in instrumentation are going to detect many thousands of new transient phenomena each night, and large facilities are preparing strategies to carry out follow-up activities that contribute to multi-messenger data exploitation. Finally, observatory surveys can also be coordinated between facilities with the same purpose: promoting simultaneous or synchronized observations.

Based on current experience of observation planning and scheduling, three pieces make up the puzzle that must be solved to scale up the single facility scheduling problem to that of multi-observatory coordinated scheduling for multi-messenger science:

- *Forming collaborations*: We need a platform to promote collaboration among individuals or among large consortiums. A `Facebook' or SMARTNet type of solution would help researcher-researcher collaboration, whereas larger agreements on Access Policies are necessary to extend the collaboration between consortia for shared use of their infrastructures and the exploitation of key scientific programmes.
- *Dashboard of observatory schedules*: Once a collaboration is in place, a platform to share the detailed observation patterns and strategies or the confirmed/tentative

schedule would make it easier to synchronize observations. This could be done by individual researchers willing to share proposal details or their observation schedule (given by the observatory) to encourage a multi-messenger approach with other people. Publication of proposals and specific schedules could be seen as a way to promote `open multi-messenger science' with a kind of GPL license for joint exploitation programmes. The scheduling data in the platform could be updated by all the subscribed observatories if suitable interfaces and protocols were defined (perhaps the VOEvent format), making it a live dashboard showing the schedules of them all. This will help exploit the combined capabilities of ground and space-based observatories, and covers not only transient follow-up (leader vs followers, which is already quite extended), but also more general multi-observatory planning.

- *Framework for time planning and scheduling*: Suitable scheduling algorithms are key to computing the best schedule for a specific proposal (one proposal and one observatory), the best schedule for a synchronized observation (one proposal and two/multiple observatories) or the best schedule for synchronized multi-observatory observations (multiple proposals and multiple observatories). Multi-observatory scheduling can be optimized based on a global scheduler, or on a hierarchy of layers to ensure local (single telescope) and global (array of observatories) optimization are achieved. Different AI algorithms can be used to take into account not only target visibility constraints, but many others (Moon, priority, elevation, simultaneity, etc.), and obtain an optimized solution for different objectives.

These three ingredients must be scaled and evolved to end-up with a global solution. Therefore, we propose that steps be taken to promote collaboration formation and the efficient scheduling of observations using software tools. The latter must cover the range from single facility to array observatories, as described in the document.

Multi-observatory scheduling, using the CTA and SKA as a test case, has also been evaluated and the results obtained will be the focus of the 5.12 deliverable of ASTERICS. We will describe how synchronized data can be collected to obtain `multi-messenger science' at no additional cost to each facility. This is because we can schedule both facilities to promote the quasi-simultaneous observation of the same targets while keeping the same performance in terms of number of proposals completed. These tests are going to be extended to include the follow-up to gravitational wave events.

The control software infrastructure is an issue for multi-facility scheduling. The ACS could form a good basis but will require development to overcome performance, reliability and security issues in its messaging network.

FUTURE WORK

Particular actions are proposed for the near future in order to achieve the aforementioned goals:

- Construct a `Coordinated Observing News' web portal through which information on multi-messenger and coordinated observing can be distributed. This should persist beyond the end of ASTERICS.
- Coordinated observations are generally controlled by collaborations. We should investigate ways in which it could be made easier to form new collaborations, or to find and join existing collaborations.
- Help collaborations plan efficient observation sequences by providing components for the construction of coordinating web apps.

- Gauge interest in holding a meeting in the last year of ASTERICS to bring together people with experience, e.g. PESSTO and SMARTNet, to chart the way forward in developing useful tools for coordinated observing.
- Advertise SMARTNet. At the time of writing (Summer 2018), the site is being used by only 50 people or so. This may be due to a lack of visibility. That said, its developers are currently content to let it grow at its own pace, so we should consult with them before proceeding.
- Consider extending SMARTNet to cover non-transient targets.
- Continue developing the STARS scheduling framework and its library of algorithms so that they are useful in optimizing schedules across the complexity scales of ground and space-based observatories.
- Develop metrics to evaluate the optimal operation of cases ranging from of an individual observatory up to the coordination of multiple-observatories.
- Develop a dashboard to display the schedules of subscribed observatories across the range of messengers. This will require the definition of standard interfaces to let observatories publish schedules for consenting individual researchers or programmes.

XXII. REFERENCES

A. Garcia-Piquer, J. C. Morales, J. Colomé and I. Ribas. 2017. "Evolutionary Computation for the ARIEL Mission Planning Tool." *6th International Conference on Space Mission Challenges for Information Technology (SMC-IT-17).* 101-106.

Acharya, B.S. et al. 2013. "Introducing the CTA concept." *Astroparticle Physics* 43: 3.

—. 2017. "Science with the Cherenkov Telescope Array." arxiv.org/abs/1709.07997.

Akturk, M. S., & Kılıç, K. 1999. *J. Intelligent Manufacturing* 10: 387.

Brown, T. M., Baliber, N., Bianco, F. B., et al. 2013. *PASP* 125: 1031.

Burt, J., Holden, B., Hanson, R., et al. 2015. *J. Astron. Telesc., Instrum. Syst.* 1: 044003.

Castellini, F., Lavagna, M. 2009. "Advanced planning and scheduling initiative XMAS tool: AI for automatic scheduling of XMM-Newton long term plans." *The 6th international workshop on planning and scheduling for space, IWPSS-09.*

Cesta, A., Cortellessa, G., Fratini, S., & Oddi, A. 2009. *The 6th Int.Workshop on Planning and Scheduling for Space, IWPSS-09.*

Civeit, T. 2013. *Proc. 2013 IEEE Aerospace Conf.* 298.

Coello Coello, C. A. 1999. *Knowledge and Information Systems* 1: 269.

Coello Coello, C. A. 2001. "Evolutionary Multi-Criterion Optimization." Edited by L. Thiele, K. Deb, C. A. Coello, & D. Corne E. Zitzler. *Lect. Notes Comput. Sci.* 1993: 21.

Coello Coello, C. A., Lamont, G. B., & Veldhuizen, D. A. V. 2007. *Evolutionary algorithms for solving multi-objective problems.* Springer.

Darwin, C. R. 1859. *On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life.* John Murray.

Donati, A., Reinhold, B., Martinez-Heras, J. A., & Policella, N. 2012. " The 12th International Conference on Space Operations."
Forbes, D. 2008. *Mercury* Spring: 25.
Freitas, A. A. 2002. *Data Mining and Knowledge Discovery with Evolutionary Algorithms* . Springer.
—. 2002. *Data Mining and Knowledge Discovery with Evolutionary Algorithms.* Springer-Verlag New York.
G. Tinetti, P. Drossart, P. Eccleston et al. 2016. "The science of ARIEL (Atmospheric Remote-sensing Infrared Exoplanet Large-survey)." *Space Telescopes and Instrumentation: Optical, Infrared, and Millimeter Wave, Proc. of SPIE.* 99041X.
García-Piquer, A. et al. 2017. *Astronomy and Astrophysics* 604: A87.
Garcia-Piquer, A. 2012. *Ph.D. Thesis, Enginyeria i Arquitectura La Salle, Universitat Ramon Llull, Barcelona.*
Garcia-Piquer, A., Fornells, A., Bacardit, J., et al. 2014a. *IEEE Transactions on Evolutionary Computation* 18: 36.
Garcia-Piquer, A., Guàrdia, J., Colomé, J., et al. 2014b. *Proc. SPIE* 9152: E21.
Garcia-Piquer, A., Ribas, I., & Colomé, J. 2014c. *J. Exp. Astron.* 40: 671.
Gharote, M.S. et al. 2009. "Automated Telescope Scheduling." *Proceedings of Low Frequency Radio Universe Conference ASP Conference Series.*
Giuliano, M. E., Hawkins, R., Rager, R. 2011. *The International Workshop on Planning and Scheduling for Space, IWPSS-11.*
Giuliano, M. E., Rager, R., & Ferdous, N. 2007. *Proceedings of The International Conference on Automated Planning and Scheduling (AAAI).* 160.
Goldberg, D. E. 1989. *Genetic Algorithm in Search, Optimization, and Machine Learning* . Addison-Wesley.
—. 1989. *Genetic Algorithm in Search, Optimization, and Machine Learning.* Addison-Wesley.
Granzer, T. 2004. "What makes an automated Telescope Robotic?" *Astronomical Notes* 325: 513.
Holland, J. H. 1975. *Adaptation in natural and artificial systems* . The University of Michigan Press.
Johnston, M., Miller, G. 1994. "SPIKE: intelligent scheduling of hubble space telescope observations." In *Intelligent Scheduling* . Morgan Kaufmann.
Joslin, D.E. et al. 1999. "Squeaky Wheel Optimization." *Journal of Artificial Intelligence Research* 10: 353.
Keivani, A. et al. 2017. arxiv.org/abs/1708.04724.
Kitching, M., Policella, N. 2013. *The 12th Symposium on Advanced Space Technologies in Robotics and Automation.*
—. 2011. "A local search solution for the INTEGRAL long term planning." *Proceedings of the 7th international workshop on planning and scheduling for space, IWPSS11.*
Lightfoot, J.F. 1989. *Monthly Notices of the Royal Astronomical Society* 239: 665.
Lohfink, A.M. et al. 2013. *Astrophysical Journal* 772: 83.
2009. "LSST Science Book Version 2.0, ch.8." www.lsst.org/sites/default/files/docs/sciencebook/SB_8.pdf.
Mendel, G. 1866. "Experiments in Plant-Hybridization." *Verhandlungen des naturforschenden Vereines in Brünn* d. IV für das Jahr 1865, Abhand-lungen B: 3-47.

Middleton, M.J. et al. 2017. "Paving the Way to Simultaneous, Multi-wavelength Astronomy." arxiv.org/pdf/1709.03520.pdf.

O'Brian, P.T. and Smartt, S.J. 2012. "Interpreting signals from astrophysical transient experiments." (Phil. Trans. R. Soc. A) 371: 20120498.

Osyczka, A. 1985. *Design Optimization* 1: 193.

Pareto, V. 1897. *Course d'Économie Politique .* Vol. 2. F. Pichou, Lausanne and Paris.

Pralet, C., Verfaillie, G. 2009. "AIMS: a tool for long-term planning of the ESA INTEGRAL Mission." *The 6th international workshop on planning and scheduling for space, IWPSS-09 .*

Rasconi, R., Policella, N., & Cesta, A. 2006a. *Proceedings of the ICAPS Workshop on Constraint Satisfaction Techniques for Planning and Scheduling Problems.* 46.

—. 2006a. *Proceedings of the ICAPS Workshop on Constraint Satisfaction Techniques for Planning and Scheduling Problems.* 46.

Rasconi, R., Policella, N., & Cesta, A. 2006b. "Proceedings of the 17th European Conference on Artificial Intelligence ." *Frontiers in Artificial Intelligence and Applications* (IOS Press) 141: 845.

Rasconi, R., Policella, N., & Cesta, A. 2006b. "Proceedings of the 17th European Conference on Artificial Intelligence Frontiers." *Artificial Intelligence and Applications* (IOS Press) 141: 845.

Sánchez, S. F., Aceituno, J., Thiele, U., Pérez-Ramírez, D., & Alves, J. 2007. *PASP* 119: 1186.

Sánchez, S. F., Thiele, U., Aceituno, J., et al. 2008. *PASP* 120: 1244.

Scandariato, G. et al. 2013. *Astronomy and Astrophysics* 552: A7.

Smartt, S.J. 2015. *Astronomy and Astrophysics* 579: A40.

Tawarmalani, M., Sahinidis, N.V. 2002. "Exact algorithms for global optimization of mixed-integer nonlinear programs." In *Handbook of Global Optimization*, 65-86. Kluwer Academic.