



This is a repository copy of *Developing the GOTO telescope control system*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/171510/>

Version: Accepted Version

Proceedings Paper:

Dyer, M.J. orcid.org/0000-0003-3665-5482, Dhillon, V.S. orcid.org/0000-0003-4236-9642, Littlefair, S. orcid.org/0000-0001-7221-855X et al. (7 more authors) (2020) Developing the GOTO telescope control system. In: Guzman, J.C. and Ibsen, J., (eds.) Proceedings of SPIE : Software and Cyberinfrastructure for Astronomy VI. SPIE Astronomical Telescopes + Instrumentation 2020, 14-18 Dec 2020, Online conference. SPIE - Society of Photo-optical Instrumentation Engineers . ISBN 9781510636910

<https://doi.org/10.1117/12.2561506>

Copyright 2020 Society of Photo Optical Instrumentation Engineers (SPIE). One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this publication for a fee or for commercial purposes, or modification of the contents of the publication are prohibited.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Developing the GOTO Telescope Control System

Martin J. Dyer^a, Vik S. Dhillon^{a,b}, Stuart Littlefair^a, Danny Steeghs^c, Krzysztof Ulaczyk^c, Paul Chote^c, Joseph Lyman^c, Duncan K. Galloway^d, Kendall Ackley^d, Yik Lun Mong^d, and the GOTO Collaboration

^aDepartment of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, UK

^bInstituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain

^cDepartment of Physics, University of Warwick, Coventry CV4 7AL, UK

^dSchool of Physics & Astronomy, Monash University, Clayton VIC 3800, Australia

ABSTRACT

The Gravitational-wave Optical Transient Observer (GOTO) is a wide-field telescope project focused on detecting optical counterparts to gravitational-wave sources. The GOTO Telescope Control System (G-TeCS) is a custom robotic control system which autonomously manages the GOTO telescope hardware and nightly operations. Since the commissioning the GOTO prototype on La Palma in 2017, development of the control system has focused on the alert handling and scheduling systems. These allow GOTO to receive and process transient alerts and then schedule and carry out observations, all without the need for human involvement. GOTO is ultimately anticipated to include multiple telescope arrays on independent mounts, both on La Palma and at a southern site in Australia. When complete these mounts will be linked to form a single multi-site observatory, requiring more advanced scheduling systems to best optimise survey and follow-up observations.

Keywords: telescopes – gravitational waves – autonomous observation – multi-site observatories – automatic control – software

1. THE GOTO PROJECT

The Gravitational-wave Optical Transient Observer (GOTO) collaboration* was founded in 2014, with the aim of detecting optical counterparts to gravitational-wave sources through the construction of a network of robotic, wide-field telescopes.¹ Each GOTO telescope consists of an array of unit telescopes (UTs) attached to a single, fast-slewing robotic mount, and with a complete array of eight UTs each GOTO telescope will have an instantaneous field of view of ~ 40 square degrees. Each GOTO site will ultimately host two independent mounts in separate domes. Under normal operations each telescope will carry out an all-sky survey, building up an archive of deep reference images. When a gravitational-wave alert is received the telescope will switch to covering as much of the visible probability skymap as possible, using the same grid as the all-sky survey. Through the process of difference imaging, new transient sources will be detected by comparing the new images to the reference image taken at the same location.

The first GOTO prototype telescope was constructed at the Observatorio del Roque de los Muchachos on La Palma in the Canary Islands in 2017, and featured four 40 cm unit telescopes (Figure 1, left). In 2020 the prototype unit telescopes were replaced with new, second-generation models with enclosed tubes (Figure 1, right). An additional mount with eight unit telescopes is under construction, and will be deployed to a second neighbouring dome on La Palma in early 2021. Funding has also been secured for a southern node at Siding Spring Observatory in New South Wales, Australia, which will host an additional two domes each with independent mounts and eight UTs. When complete the GOTO network will be able observe and react to alerts with near 24-hour coverage, with an instantaneous field of view of ~ 80 square degrees at each site when utilising both mounts. It is planned that the entire network will be operational by the time that the LIGO-Virgo-KAGRA gravitational-wave detectors resume operations in 2022.²

Send correspondence to MJD - email: martin.dyer@sheffield.ac.uk

*<https://goto-observatory.org/>



Figure 1. The GOTO prototype on La Palma: with the initial four unit telescopes (left) and after the 2020 upgrade (right). The four outer tubes shown on the right are temporary stand-ins while the original OTAs are refurbished.

2. ROBOTIC TELESCOPE CONTROL

A custom software package called the GOTO Telescope Control System (G-TeCS)[†] was developed to control all operations of the robotic telescope.^{3,4} G-TeCS controls all aspects of the telescope, including the physical hardware, observation scheduling and automated operations such as ensuring the dome closes in poor weather conditions. Initial development of G-TeCS focused on the hardware control software and the pilot, and was described previously in [3]. Sections 2.1 and 2.2 summarise the previous work, before detailing the more recent developments in the alert follow-up system in Sections 2.3 and 2.4.

2.1 Hardware control

Each category of hardware has an associated daemon, which monitors the hardware and allows commands to be sent to it. There are eight primary hardware control daemons as shown in Figure 2: the camera, exposure queue, filter wheel, focuser, OTA, dome, mount and power daemons. The dome and mount daemons are solely for controlling and monitoring the clamshell dome and robotic mount respectively, while the power daemon monitors and provides a unified interface for multiple networked power supply and distribution units. Due to the design of GOTO, with multiple unit telescopes on each mount, the camera, filter wheel, focuser and OTA daemons interact with their hardware units through simple interface daemons running on small Intel PCs attached to the main boom-arm. Using these interfaces the daemons each control the four appropriate hardware units (eight in the final design) as one. For example, the filter wheel daemon can change the filter of each unit telescope individually or together in parallel. Likewise exposures on each camera are started, downloaded and saved simultaneously. The exposure queue daemon provides a common interface for taking sets of multiple exposures by issuing commands to the camera daemon in sequence, and also allows for filters to be set between exposures. The OTA daemon is a recent addition to control the mirror covers on the new enclosed tubes.

The hardware daemons offer a simple interface to the physical hardware, and they regularly fetch and store any relevant status information. These daemons however are not designed to operate autonomously, and will not issue commands to the hardware unless told to by a human user or the robotic pilot (see Section 2.2.1). The exception to this is the dome daemon, which has a limited amount of autonomy enabling it to automatically close the dome if the weather conditions turn bad (hence the direct connection between it and the conditions monitor shown in Figure 2). This provides a critical backup to ensure the telescope is protected in case the order to close from the pilot is delayed or the system is in manual mode (see Section 2.2.2).

[†]<https://github.com/GOTO-OBS/g-tecs>

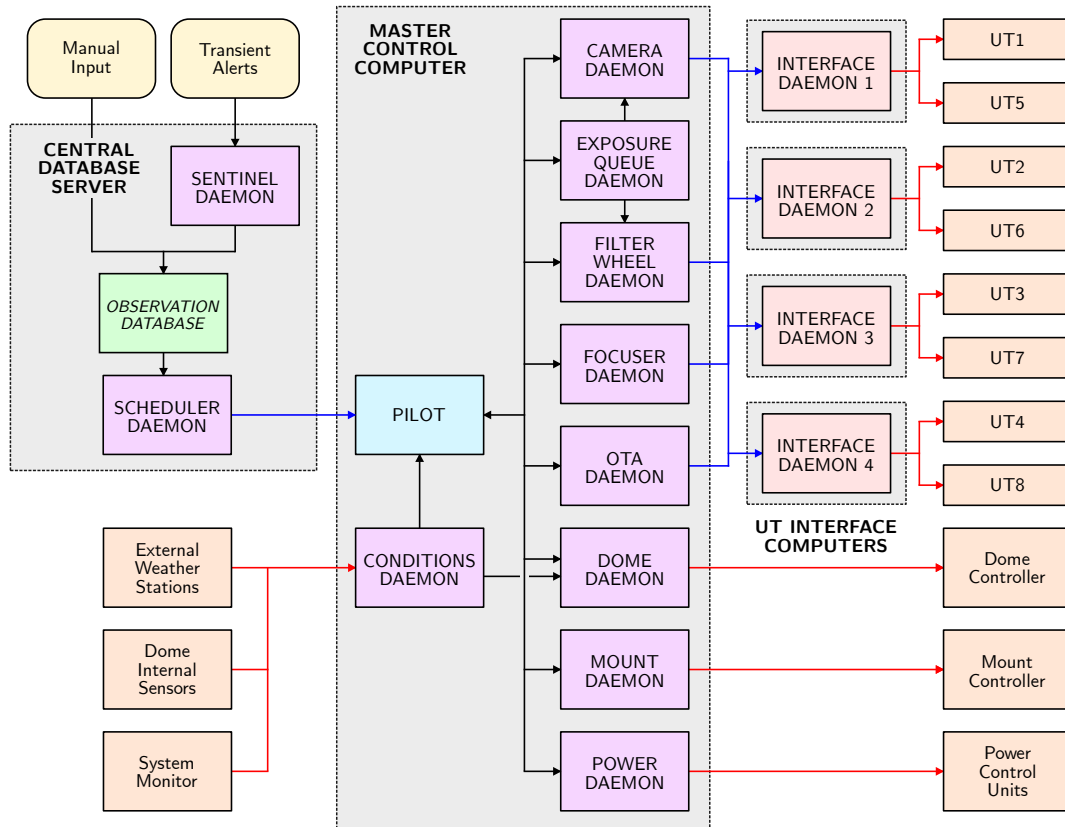


Figure 2. The G-TeCS software architecture.

2.2 Autonomous systems

The hardware daemons are used to interface with and monitor the telescope hardware, however aside from the dome closing in bad weather they have no autonomy of their own and need to be sent commands by either a human observer or by the *pilot* master control script. The pilot is an asynchronous Python script containing multiple coroutines, which together monitor and operate the telescope hardware.

2.2.1 Pilot routines

The primary routine within the pilot is called the night marshal, and its purpose is to run predefined observing scripts at set times during the night based on the altitude of the Sun. At the start of the night the pilot powers on the hardware and takes bias and dark frames. Once the Sun has set below the horizon, and the conditions are safe to do so, the night marshal issues the command to open the dome and starts the script to take sky flats in each of the four *LRGB* filters. When taking flats the telescope is pointed at the anti-sun position,⁵ and exposure times are calculated based on the sky brightness.⁶ After the flats script has finished the night marshal starts the autofocus routine, which focuses each unit telescope in parallel using the half-flux diameter V-curve method.⁷ Once the telescopes are focused the pilot switches to observing mode. For the rest of the night the pilot fetches the highest priority pointings from the scheduler (see Section 2.3), slewing the mount to each target and taking the requested exposures, until the Sun rises the next morning. Then the night marshal runs the flats routine again to take additional twilight flats, before closing the dome and shutting down the system. For security there is an independent night countdown timer within the pilot, which will trigger the shutdown when the Sun rises above the horizon if the night marshal hasn't already done so. There is also an entirely separate script, called the day marshal, which runs every morning and ensures that the dome is closed in case of an error in the pilot.

Aside from the night marshal, the pilot also contains multiple check routines that monitor different aspects of the control system. One routine fetches the latest pointings from the scheduler (see Section 2.3) every 10 seconds, interrupting the current observation if the new target has a higher priority. Another routine monitors the hardware daemons for any errors, which are defined by comparing the expected status to the actual one. If the two differ (e.g. the cameras are not at the requested temperature, or dome is open when it should be closed) the pilot is able to issue a series of predefined recovery commands to the daemons to attempt to fix the problem. If this is unsuccessful then the pilot shuts down and sends out alerts for human assistance. A third routine is in charge of pausing and resuming the night marshal if required: if a hardware error is detected, the system is in manual mode (see Section 2.2.2), or the conditions flags are bad (see Section 2.2.3).

2.2.2 Observing modes

Although GOTO is capable of operating fully autonomously, it is sometimes necessary for a human operator to take control. An example of the need for a human observer was when GOTO was used to observe the asteroid Phaethon,⁸ as the existing scheduling system was not designed for moving objects. It is also important for safety that the automated systems can be reliably disabled if there are workers or observers on site, as it could be dangerous if the system still tried to move the mount or dome unexpectedly. G-TeCS therefore includes three operating modes: robotic, manual or engineering. Robotic mode is the default, and while the system is in this mode the pilot will be in complete control of the telescope and the hardware will moved autonomously (for example the dome will always automatically close in bad weather). Manual mode is intended for cases when a human is in charge of the telescope, either on-site or remotely. In this mode the pilot will be paused so it will not interrupt commands sent by the observer, and while the dome will still close automatically if the conditions are bad this, and all the automated systems, can be disabled if the observer chooses. Finally engineering mode is used when there are workers on site. In this mode all autonomous systems are disabled and cannot be enabled, and the pilot will refuse to start.

2.2.3 Conditions monitoring

Before GOTO could be trusted to run without human supervision it was necessary for it to have a robust and reliable system for closing the dome in bad weather. G-TeCS therefore contains a conditions daemon, shown in Figure 2, which is tasked with monitoring the site conditions.

The conditions daemon monitors standard weather values reported by the on-site weather stations located just a few metres from the GOTO domes. This includes the rain, temperature, humidity, dew point, wind speed and wind gust values. Other sources of conditions information includes the internal temperature and humidity sensors, the network link, the amount of free disk space, the power level of the dome UPS and the status of the dome access hatch. The daemon processes these inputs into a series of output flags, which have a value of 0 (good), 1 (bad) or 2 (error). Each conditions flag has a limit below or above which the flag will turn from good to bad, as well as delay times that have to be exceeded before the flag changes (therefore smoothing out any sudden spikes). If any of the flags are marked as not good (i.e. the sum of all flags is > 0) then the overall conditions are bad, which will trigger the dome to close and the pilot to pause its current task.

A recent addition was a manual override flag, which can be set while the system is still in robotic mode for cases that the conditions monitor is not able to detect. This is most commonly used to pause observing during cloudy skies when the weather is otherwise good, as there is no automatic cloud sensor connected to the conditions monitor. The flag can also be used for rarer weather conditions on La Palma, such as smoke from forest fires or high levels of dust from the Sahara.

		Highest priority pointing is . . .			
		same as current pointing	a new, valid pointing	a new, invalid pointing	(None)
Current pointing is . . .	valid	Continue current pointing	Interrupt and start new pointing if it is a ToO and the current pointing is not, otherwise continue current pointing	N/A	N/A
	invalid	Park	Interrupt and start new pointing	Park	Park
	(None)	Park	Start new pointing	Park	Park

Table 1. Actions the pilot will take based on the scheduler results. The scheduler will return the current pointing and the highest priority pointing, which may be the same or different (and either could be None, indicating the pilot is not currently observing or the queue is empty, respectively).

2.3 Observation scheduling

In order for the pilot to function during the night it needs to know what targets to observe. GOTO operates under a “just-in-time” scheduling model: rather than creating a plan at the beginning of the night of what to observe, each time the pilot queries the scheduler the priority of every valid pointing in the queue is calculated and the highest priority pointing at that time is returned. This system is very reactive to any incoming alerts, as the new pointings will immediately be included in the queue at the next check. This method also naturally works around any delay in observations due to poor conditions, unlike a fixed night plan. During the third LIGO-Virgo observing run GOTO was able to process gravitational-wave alerts and begin observations within 30 seconds of them being received.⁹

The scheduler carries out several steps to determine the highest priority pointing. First the queue of pending pointings is fetched from the observation database, and any that are invalid (e.g. below the horizon, or too close to the Moon) are filtered out. The pointings are then sorted by several parameters. The first is the effective rank (R), which is the combination of the starting rank assigned when the pointing was entered into the database (R_s) and the number of times that pointing has been observed since (n_{obs}) multiplied by 10 ($R = R_s + 10 \times n_{\text{obs}}$). This parameter was chosen to account for observing large numbers of pointings for different events simultaneously, and is discussed in Section 2.4.2. For pointings of equal effective rank, any that are designated as targets of opportunity will be prioritised (and are able to interrupt observations that are in progress). In cases where there are still multiple pointings with equal priority then a tiebreak parameter (V) is calculated by combining any weighting (W) assigned to the pointings (typically from the enclosed skymap probability, see Section 2.4.3) and the airmass of the target ($X(t)$) at the time of the calculation. These are combined in the ratio 10:1 using

$$V(t) = \frac{10}{11}W + \frac{1}{11}(2 - X(t)). \quad (1)$$

The combination of tile weighting and airmass is adjusted to maximise the chance of detecting a counterpart while not neglecting data quality. The 10:1 ratio was determined using preliminary strategy simulations detailed in [4], and work is ongoing to further improve and optimise this algorithm.

The check routine within the pilot queries the scheduler every 10 seconds, and the scheduler returns both the highest priority pointing at the query time and the current pointing with its priority recalculated. The pilot compares the two and decides between three options: carry on with the current observation, switch to a new observation, or park the telescope (in the case that there are no valid targets). The different possible cases are summarised in Table 1. Most of the time if the pilot is currently observing a pointing it will remain the highest priority in subsequent checks, so the pilot will continue observing it. Even if the scheduler finds that a different pointing now has a higher priority the pilot will not change targets unless the new pointing is marked as a target of opportunity and the current pointing is not. Otherwise the pilot will wait until it has finished the current pointing, mark it as complete in the database and ask the scheduler for the next target to observe.

GCN packet type	Source	Notice type	GOTO-alert Event subclass
61	NASA/ <i>Swift</i>	SWIFT_BAT_GRB_POS	GRBEvent
115	NASA/ <i>Fermi</i>	FERMI_GBM_FIN_POS	GRBEvent
150	LIGO-Virgo	LVC_PRELIMINARY	GWEvent
151	LIGO-Virgo	LVC_INITIAL	GWEvent
152	LIGO-Virgo	LVC_UPDATE	GWEvent
164	LIGO-Virgo	LVC_RETRACTION	GWRetractionEvent

Table 2. GCN notices and corresponding classes recognised by the GOTO-alert event handler.

2.4 Transient alert processing

In order for targets to be observed by the pilot they must have entries defined in the observation database. These can be added manually, but for automated follow-up observations they have to be inserted whenever an alert is received. As shown in Figure 2 this is the job of the G-TeCS sentinel daemon. The sentinel alert listener continuously monitors the transient alert stream output by the 4 Pi Sky event broker,¹⁰ primarily the NASA GCN/TAN network,¹¹ using functions from the PyGCN Python package[‡]. Alert information is received in the form of VOEvents,¹² and are parsed using the voevent-parse Python package.¹³ Each event is then processed using a Python package developed for the GOTO project, GOTO-alert[§].

2.4.1 Event classes

GOTO-alert is designed around a single `Event` object class, and “interesting” events are processed using individual subclasses. There are currently three subclasses defined within GOTO-alert, as given in Table 2: gamma-ray burst (GRB) detections from *Fermi* GBM¹⁴ or *Swift* BAT¹⁵ use the `GRBEvent` class, while the different stages of gravitational-wave (GW) detections from the LIGO-Virgo Collaboration¹⁶ use the `GWEvent` class (aside from retraction notices, which use a unique `GWRetractionEvent` class). Notices with a packet type not listed in Table 2 are ignored, although the flexibility of the GOTO-alert code makes it easy to define new event subclasses to process them if desired. Likewise events that do have the correct packet type but are not marked as observations (e.g. test notices) are also rejected.

2.4.2 Observing strategy

Once an event has been assigned the correct subclass the sentinel passes it to the GOTO-alert event handler, which determines the correct observing strategy for the event and inserts pointings into the observation database. Observations can be tailored to the properties of the triggering event by modifying the pointing rank or observing constraints, as well as the specific exposures requested (e.g. exposure time, filters) and the desired cadence for any repeat observations. For LIGO-Virgo detections subsequent notices can contain updated localisation skymaps, overwriting those from previous alerts. In the future the observing strategy will also need to consider the different GOTO observing sites and how best to utilise the telescopes at each site (see Section 3).

To determine the priority for a given event the event handler uses the decision tree shown in Figure 3. The codes at the end of the branches (`GW_CLOSE_NS`, `GRB_FERMI`, etc) correspond to individual strategies, and event sources and parameters contained within the alert are used to decide which strategy is appropriate for a given event. First, as they are GOTO’s primary science goal, gravitational-wave alerts are always prioritised over GRB alerts. Gravitational-wave signals that are detected by the CBC (Compact Binary Coalescence) pipeline can be split between those that are predicted to contain neutron stars (either neutron star-neutron star binaries or neutron star-black hole binaries) and those that are not (black hole-black hole binaries). As the former are considered more likely to produce electromagnetic counterparts they are prioritised for follow-up. Events are then further split between “near” and “far” for each branch, which is roughly determined based on the distance provided in the notice skymap. The highest priority event is therefore a gravitational-wave signal from a nearby neutron star-neutron star or neutron star-black hole binary, which is most likely to produce a counterpart that GOTO can observe. Gravitational-wave burst detections are rarer and have less certain origins, so are prioritised

[‡]<https://pypi.org/project/pygcn>

[§]<https://github.com/GOTO-OBS/goto-alert>

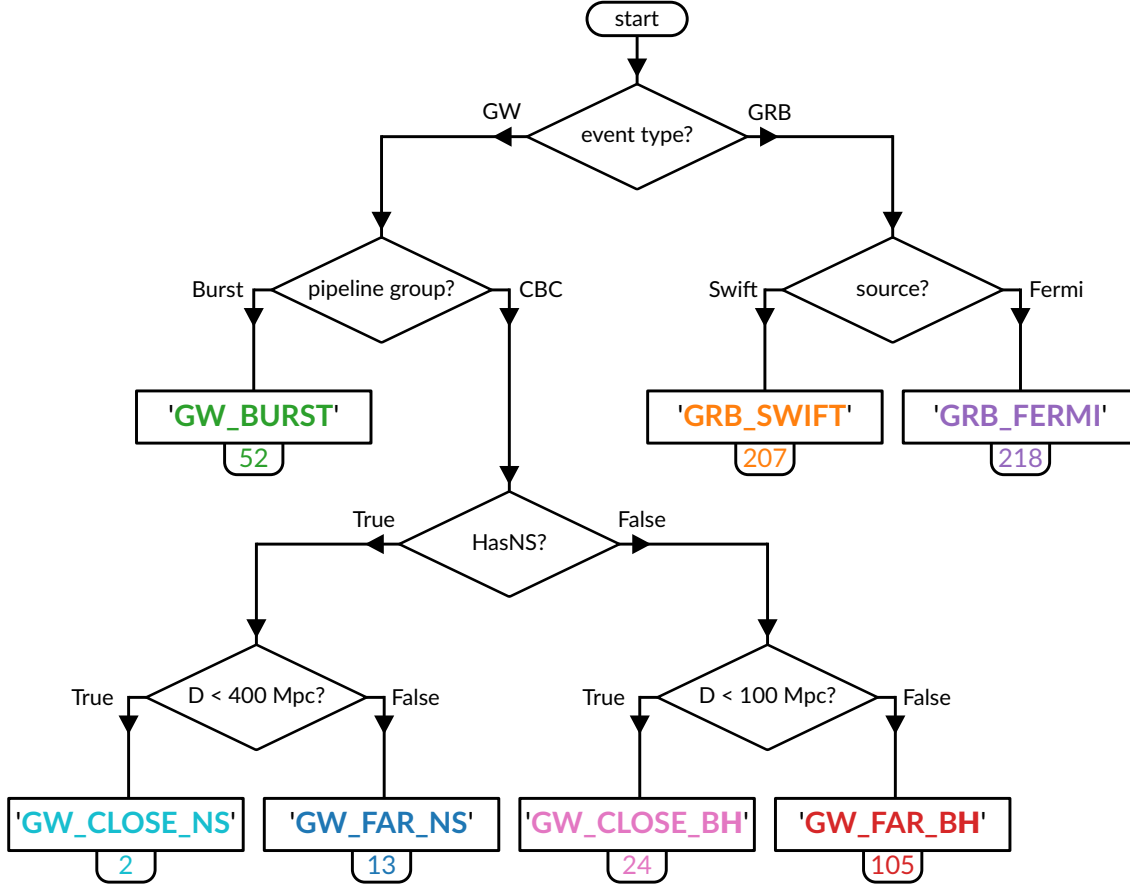


Figure 3. Decision tree for determining event strategy. The number below each strategy code is the starting rank used when adding pointings for that event to the observation database (lower ranks = higher priority).

at a medium rank, while GRB alerts from *Swift* are prioritised over those from *Fermi* as they usually have smaller localisation regions.

The primary value used when scheduling pointings, as described in Section 2.3, is the starting rank. This is assigned by the event handler and is the same for all pointing from a given event. Due to the method of sorting pointings by their effective rank, pointings from different events can be differentiated and prioritised for follow-up should GOTO need to follow-up multiple events at the same time. For instance, if two alerts are received simultaneously, one from a close neutron star binary and one from a distant one, then the pointings for the first event will be inserted into the database at rank 2 while the others will be inserted with starting rank 13 (the ranks are given in Figure 3). When the queue is sorted by the scheduler the visible pointings from the first event will always be sorted higher than the second. However, as pointings from the event are observed and are marked completed by the pilot their effective rank starts to increase: from 2 when $n_{\text{obs}} = 0$ to 12, 22, 32 etc. Once a pointing has been observed twice its effective rank is 22, which is now below the starting rank of pointings for the second event (13). Therefore once all visible pointings from the first event have been observed twice, pointings from the second event will be sorted higher, and so will start to be observed. Subsequent pointings from this event will then also increase in effective rank from 13 to 23, 33 etc, and from then on priority will alternate between pointings from each event. It is worth emphasising that this only matters when multiple events have pending pointings in the scheduler queue at the same time, and as pointings only remain in the queue for a maximum of three days this is currently fairly rare. However, as new gravitational-wave detectors come online, and existing ones are upgraded, the requirement to decide which event to prioritise may be more common.

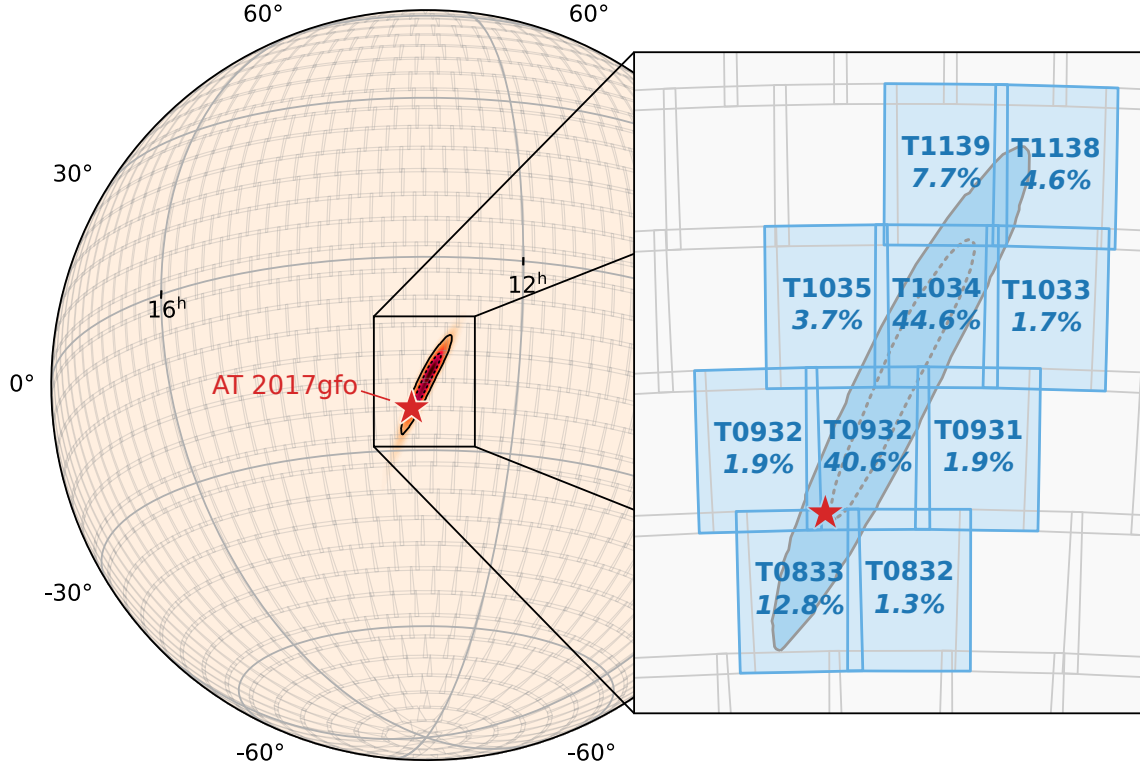


Figure 4. GOTO-4 tiling applied to the skymap for GW170817. The blue tiles are those containing more than 1% of the skymap probability. Note that overlaps between tiles means the total probability can add up to more than 100%.

2.4.3 Mapping onto the all-sky grid

GOTO carries out observations aligned to an all-sky grid, made up of tiles equal to the combined field of view of all the unit telescopes on the mount. For the GOTO-4 prototype each tile covered 3.7 degrees in right ascension and 4.9 degrees in declination, and the all-sky grid was defined by dividing the sky into a series of 2913 equally-spaced 18.1 square degree tiles. These tiles are used as the basis of the all-sky survey that GOTO carries out when not following-up transient alerts. Survey images taken of each tile are stacked to make master reference frames, and when follow-up observations are taken aligned to the same grid new images are compared to the reference frames to detect transient sources through difference imaging (subtracting the reference from the image).

It is therefore necessary that any follow-up observations are aligned to the same grid used by the all-sky survey. LIGO-Virgo gravitational-wave notices contain probability skymaps encoded using the HEALPix framework,¹⁷ and the GOTO-alert event handler maps the skymap onto the all-sky grid by summing the values of all HEALPix pixels that are contained within each tile. An example of this mapping is shown in Figure 4, using the initial three-detector skymap for GW170817.^{18,19} Once the probability value is known for each tile, those with the highest probability values are added as pointings to the observation database, and the enclosed probability is used as the weighting factor when choosing between tiles of equal rank (see Section 2.3). The number of tiles to add as pointings needs to be limited to ensure that tiles are revisited regularly: with the “effective-rank” method outlined in Section 2.4.2, the second observation of a tile, regardless of the probability contained within, will only take place once all other valid pointings added at the same rank have also been observed at least once. The GOTO-alert event handler selects tiles with a mean contour value of 90% (i.e. the mean contour value of all the pixels contained within is 90% or more). However, more quantitative simulations of different skymaps could be used to determine if 90% is the optimal choice for all cases. For example, the contour value limit could be modified depending on the overall size of the skymap, and as more GOTO telescopes are built the limit could be lowered to add more tiles to the database, as more telescopes will be able to cover the skymap faster.

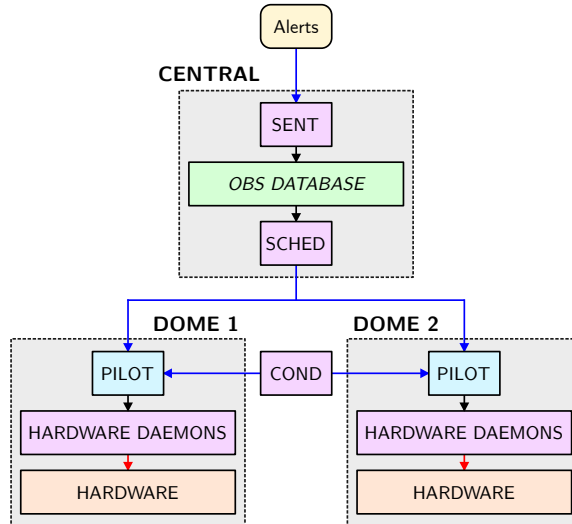


Figure 5. The proposed future architecture for G-TeCS, expanding Figure 2 to cover multiple telescopes.

3. FUTURE EXPANSION

The first GOTO prototype telescope has been successfully observing on La Palma for several years, using the software systems described in Section 2 to automate every aspect of its operations. During 2021, several new GOTO telescopes will be constructed: first a second mount on La Palma next to the existing telescope, then two more mounts at Siding Spring Observatory in Australia. As part of the commissioning work, G-TeCS will be expanded to control the new hardware, ultimately operating all four mounts as a single, global, autonomous observatory.

3.1 Controlling multiple telescopes

The addition of the second telescope on La Palma will require the expansion of the existing G-TeCS software. The general architecture, as shown in Figure 2, with the pilot receiving targets from the scheduler and sending commands to the hardware daemons, will remain the same for both telescopes. Aside from a new direct-drive mount, the hardware will be the same as the current system on La Palma, so the hardware control daemons will be copied across with minimal changes. As shown in Figure 5, it is planned that each telescope will have an independent pilot, and will therefore carry out normal operations completely independently. This will ensure that, for example, a hardware fault with one telescope will not prevent the other from observing. It is also anticipated that the two pilots will share a single conditions monitor as, aside from the internal sensors in each dome, both telescopes will be monitoring the same weather stations and other system properties described in Section 2.2.3.

The biggest opportunity presented by having multiple connected telescopes, and therefore the most important change to the overall control system, is having them share a common scheduling system. This was always planned to take the form of each pilot querying a central scheduler, which is one of the reasons the existing system already has the sentinel, database and scheduler running on a separate server to the rest of the control system (as shown in Figure 2). Figure 5 shows the natural extension of this, with both pilots receiving targets from the same scheduler daemon. This architecture allows the pilots to continue to operate independently, without any knowledge of the other, while any decisions relating to operating the telescopes in parallel are taken at the scheduler level.

There are numerous opportunities and challenges in developing this multi-telescope scheduling system. As described in Section 2.3, every few seconds the existing scheduler takes the current queue and recalculates the highest priority pointing, which is then sent to the pilot. Initially, this system could be expanded relatively easy by having the scheduler simply calculate the top two highest priority pointings, and passing them on to

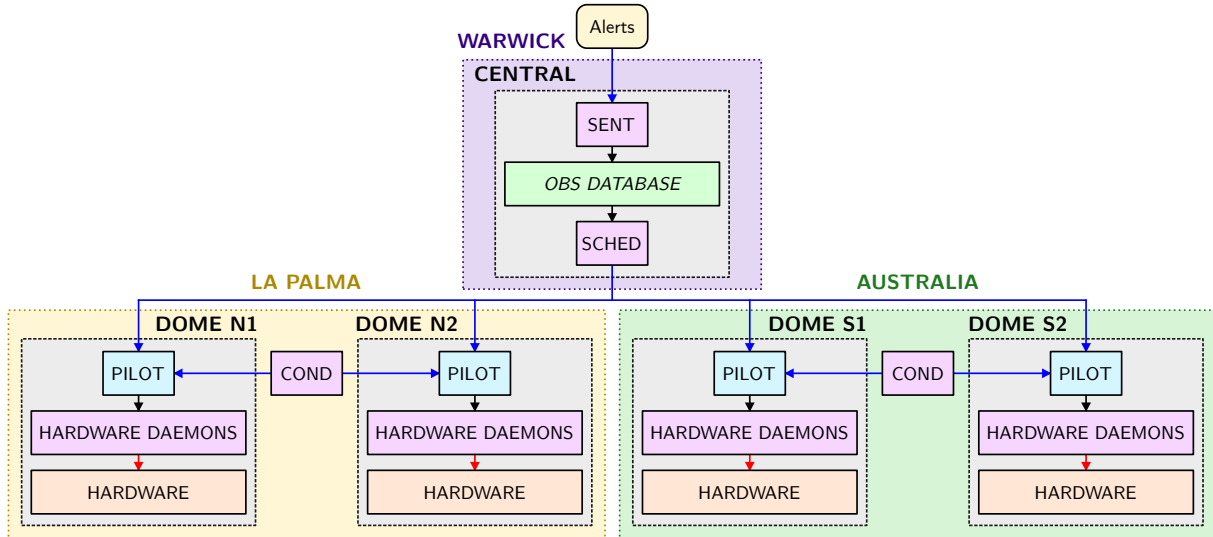


Figure 6. The final G-TeCS software architecture, with multiple telescopes at multiple sites each receiving targets from the central scheduler.

the two telescopes so that one always observes the highest-priority pointing and the other always observes the second-highest. However in reality truly independent telescopes will never remain perfectly in sync, as the slew time taken to move to each target will not be the same, and so the scheduler will have to be able to handle asynchronous requests from either telescope at any time. The upgraded scheduler could also take the time to slew from the current position, as well as any time remaining in the current observation, into account when assigning targets. In addition, while most scheduling constraints will be the same for both mounts (e.g. the Moon phase or Sun altitude) it is possible that the two telescopes could have different artificial horizons and therefore altitude limits, thereby forcing the choice if the highest priority target is only visible by one telescope.

Having multiple mounts also opens up opportunities for more advanced observing strategies, which would have to be determined by the scheduler. The two telescopes could both observe different parts of the sky to cover the survey or gravitational-wave skymaps faster, or they could observe the same tiles to achieve a greater depth when the images are stacked. The presence of the coloured filters adds even more possibilities: it should be possible to have the two telescopes observe the same tile simultaneously but using different filters, to get immediate colour information on all sources across the field. It might also be desirable to accept the impact on survey cadence and have each telescope carry out an independent survey in different filters, or perhaps have one taking rapid 60 s exposures while the other surveys the sky more slowly but to a greater depth. These decisions will ultimately be made depending on the science requirements of the GOTO collaboration, but it is important that the future G-TeCS scheduling system should be able to handle whatever strategy is chosen.

3.2 Controlling multiple sites

As described in Section 1, the GOTO project will ultimately include sites in both the northern and southern hemispheres. The same principle described previously for controlling multiple telescopes will apply wherever they are situated in the world: each will have an independent pilot which will receive targets from a central scheduling system. Figure 6 expands Figure 5 to cover four telescopes spread across two different sites: La Palma in the Canary Islands and Siding Spring in Australia. The server hosting the observation database and scheduler is currently situated on La Palma, but it is anticipated that this will be relocated to the GOTO data centre at Warwick University, UK, where it will be accessible from both sites.

When upgrading the scheduler to account for multiple sites it will have to take into account not just different horizons but also the position of the Sun, to know which telescopes are available to observe. For the planned GOTO sites there is a fortunate simplification: as the Canary Islands and eastern Australia are near-antipodal,

there will never be a case when both sites are observing at the same time (when the Sun is below the horizon for one it will be above the horizon for the other). This drastically simplifies the calculations for the central scheduler: as it will only ever have to determine targets for telescopes based at the same site, it should not need to include any extra logic to that already required for the two La Palma telescopes described in Section 3.1, aside from tailoring its decision based on which site is active at the current time.

Although the two sites will never be observing simultaneously, there is still a benefit to them sharing a scheduler, and in particular having a shared observation database. La Palma is located at 28° north and Siding Spring is located at 31° south, meaning each telescope can see approximately 75% of the sky and 50% of the sky is visible from both sites. With a shared database, the sites will be able to effectively share the load of observing the all-sky survey tiles within this region. Likewise, if a gravitational-wave probability skymap is visible from both hemispheres whichever site starts observing second (i.e. the site where it was daytime at the time the event was received) should take into account which tiles have already been observed by the other site, and then prioritise those which have not yet been observed. Therefore the best way to utilise both sites efficiently is with a fully-connected system, making the entire GOTO network a single global observatory.

4. CONCLUSION

The GOTO Telescope Control System (G-TeCS) has been successfully commissioned with the GOTO-4 prototype telescope on La Palma. The software controls all aspects of the robotic telescope, including hardware control, conditions monitoring, alert listening and observation scheduling. In particular, the system allowed GOTO to successfully participate in the third LIGO-Virgo gravitational-wave observing run in 2019-20.^{9,20}

As more telescopes are added to the GOTO network, G-TeCS will evolve from controlling a single telescope to overseeing the entire global network. Each GOTO node will run semi-independently, with a central scheduler prioritising observations and assigning targets to each telescope. By the start of the next gravitational-wave observing run GOTO should be a fully operational network, providing near 24-hour coverage of the visible sky and able to begin follow-up observations of any new detections in less than a minute.

Acknowledgements

The Gravitational-wave Optical Transient Observer (GOTO) project acknowledges the support of the Monash-Warwick Alliance; the University of Warwick; Monash University; the University of Sheffield; the University of Leicester; Armagh Observatory & Planetarium; the National Astronomical Research Institute of Thailand (NARIT); the Instituto de Astrofísica de Canarias; the University of Turku; the University of Manchester; the University of Portsmouth and the UK Science and Technology Facilities Council (grant numbers ST/T007184/1 and ST/T003103/1).

This project makes use of Astropy, a community-developed core Python package for Astronomy.^{21,22}

REFERENCES

- [1] Steeghs, D., Galloway, D. K., Ackley, K., et al., “The Gravitational-wave Optical Transient Observer(GOTO): prototype performance and prospects for transient science,” (in prep.).
- [2] Abbott, B. P., Abbott, R., Abbott, T. D., et al., “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA,” *Living Reviews in Relativity* **21**, 3 (Apr. 2018).
- [3] Dyer, M. J., Dhillon, V. S., Littlefair, S., et al., “A telescope control and scheduling system for the Gravitational-wave Optical Transient Observer (GOTO),” in [*Observatory Operations: Strategies, Processes, and Systems VII*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **10704**, 107040C (July 2018).
- [4] Dyer, M. J., *A telescope control and scheduling system for the Gravitational-wave Optical Transient Observer*, PhD thesis, University of Sheffield (Jan. 2020).
- [5] Wei, P., Shang, Z., Ma, B., et al., “Problems with twilight/supersky flat-field for wide-field robotic telescopes and the solution,” in [*Observatory Operations: Strategies, Processes, and Systems V*], *SPIE Conference Series* **9149**, 91492H (Aug. 2014).

- [6] Tyson, N. D. and Gal, R. R., “An Exposure Guide for Taking Twilight Flatfields With large Format CCDs,” *AJ* **105**, 1206 (Mar. 1993).
- [7] Weber, L. and Brady, S., “Fast auto-focus method and software for ccd-based telescopes,” *Minor Planet Amateur/Professional Workshop*, 104–113 (2001).
- [8] Borisov, G., Devogèle, M., Cellino, A., et al., “Rotational variation of the linear polarization of the asteroid (3200) Phaethon as evidence for inhomogeneity in its surface properties,” *MNRAS* **480**, L131–L135 (Oct. 2018).
- [9] Gompertz, B. P., Cutter, R., Steeghs, D., et al., “Searching for electromagnetic counterparts to gravitational-wave merger events with the prototype Gravitational-Wave Optical Transient Observer (GOTO-4),” *MNRAS* **497**, 726–738 (July 2020).
- [10] Staley, T. D. and Fender, R., “The 4 Pi Sky Transient Alerts Hub,” *arXiv e-prints*, 1606.03735 (June 2016).
- [11] Barthelmy, S. D., Butterworth, P., Cline, T. L., et al., “The GRB coordinates network (GCN): A status report,” in [*Gamma-Ray Bursts, 4th Hunstville Symposium*], *American Institute of Physics Conference Series* **428**, 99–103 (May 1998).
- [12] Seaman, R., Williams, R., Allan, A., et al., “Sky event reporting metadata version 2.0,” *Recommendation, International Virtual Observatory Alliance* (June 2011).
- [13] Staley, T. D., “voevent-parse: Parse, manipulate, and generate VOEvent XML packets,” *Astrophysics Source Code Library*, 1411.003 (Nov. 2014).
- [14] Meegan, C., Lichti, G., Bhat, P. N., et al., “The Fermi Gamma-ray Burst Monitor,” *ApJ* **702**, 791–804 (Sept. 2009).
- [15] Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al., “The Burst Alert Telescope (BAT) on the SWIFT Midex Mission,” *Space Sci. Rev.* **120**, 143–164 (Oct. 2005).
- [16] Abbott, B. P., Abbott, R., Abbott, T. D., et al., “Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo,” *Living Reviews in Relativity* **19**, 1 (Feb. 2016).
- [17] Górski, K. M., Hivon, E., Banday, A. J., et al., “HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere,” *ApJ* **622**, 759–771 (Apr. 2005).
- [18] Abbott, B. P., Abbott, R., Abbott, T. D., et al., “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Phys. Rev. Lett.* **119**, 161101 (Oct. 2017).
- [19] Abbott, B. P., Abbott, R., Abbott, T. D., et al., “Multi-messenger Observations of a Binary Neutron Star Merger,” *ApJ* **848**, L12 (Oct. 2017).
- [20] Ackley, K., Amati, L., Barbieri, C., et al., “Observational constraints on the optical and near-infrared emission from the neutron star-black hole binary merger candidate S190814bv,” *A&A* **643**, A113 (Nov. 2020).
- [21] Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al., “Astropy: A community Python package for astronomy,” *A&A* **558**, A33 (Oct. 2013).
- [22] Astropy Collaboration, Price-Whelan, A. M., Sipócz, B. M., et al., “The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package,” *AJ* **156**, 123 (Sept. 2018).