CrossMark

CAMBRIDGE
UNIVERSITY PRESS

Research Article

SkyHopper mission science case I: Identification of high redshift Gamma-Ray Bursts through space-based near-infrared afterglow observations

M. Thomas[1] , M. Trenti[1,2] , J. Greiner[3], M. Skrutskie[4], Duncan A. Forbes[5], S. Klose[6], K. J. Mack[7], R. Mearns[1], B. Metha[1,2], E. Skafidas[8], G. Tagliaferri[9] and N. Tanvir[10]

[1]School of Physics, The University of Melbourne, VIC 3010, Australia, [2]Australian Research Council Centre of Excellence for All-Sky Astrophysics in 3-Dimensions, Australia, [3]Max-Planck-Institut für extraterrestrische Physik, Giessenbachstr. 1, D-85740 Garching, Germany, [4]Department of Astronomy, University of Virginia, Charlottesville, P.O. Box 400325, VA 22904-4325, USA, [5]Centre for Astrophysics & Supercomputing, Swinburne University, Hawthorn, VIC 3122, Australia, [6]Thüringer Landessternwarte Tautenburg, Sternwarte 5, 07778 Tautenburg, Germany, [7]Physics Department, North Carolina State University, Raleigh, NC 27695, USA, [8]Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, VIC 3010, Australia, [9]INAF—Osservatorio Astronomico di Brera, Via Bianchi 46, I-23807 Merate, Italy and [10]School of Physics and Astronomy, University of Leicester, University Road, Leicester, LE1 7RH, UK

Abstract

Long-duration gamma-ray burst (GRB) afterglow observations offer cutting-edge opportunities to characterise the star formation history of the Universe back to the epoch of reionisation, and to measure the chemical composition of interstellar and intergalactic gas through absorption spectroscopy. The main barrier to progress is the low efficiency in rapidly and confidently identifying which bursts are high redshift ($z > 5$) candidates before they fade, as this requires low-latency follow-up observations at near-infrared wavelengths (or longer) to determine a reliable photometric redshift estimate. Since no current or planned gamma-ray observatories carry near-infrared telescopes on-board, complementary facilities are needed. So far this task has been performed by instruments on the ground, but sky visibility and weather constraints limit the number of GRB targets that can be observed and the speed at which follow-up is possible. In this work we develop a Monte Carlo simulation framework to investigate an alternative approach based on the use of a rapid-response near-infrared nano-satellite, capable of simultaneous imaging in four bands from 0.8 to 1.7 μm (a mission concept called SkyHopper). Using as reference a sample of 88 afterglows observed with the GROND instrument on the MPG/ESO telescope, we find that such a nano-satellite is capable of detecting in the H-band (1.6 μm) 72.5% ± 3.1% of GRBs concurrently observable with the Swift satellite via its UVOT instrument (and 44.1% ± 12.3% of high redshift ($z > 5$) GRBs) within 60 min of the GRB prompt emission. This corresponds to detecting ∼55 GRB afterglows per year, of which 1–3 have $z > 5$. These rates represent a substantial contribution to the field of high-z GRB science, as only 23 $z > 5$ GRBs have been collectively discovered by the entire astronomical community over the last ∼24 yr. Future discoveries are critically needed to take advantage of next generation follow-up spectroscopic facilities such as 30m-class ground telescopes and the James Webb Space Telescope. Furthermore, a systematic space-based follow-up of afterglows in the near-infrared will offer new insight on the population of dusty ('dark') GRBs which are primarily found at cosmic noon ($z \sim 1 - 3$). Additionally, we find that launching a mini-constellation of 3 near-infrared nano-satellites would increase the detection fraction of afterglows to ∼83% and substantially reduce the latency in the photometric redshift determination.

Keywords: gamma-ray bursts – near-infrared astronomy – infrared astronomical satellite

(Received 8 February 2022; revised 22 April 2022; accepted 29 April 2022)

1. Introduction

Long-duration gamma-ray bursts (GRBs) are some of the most luminous explosions in the observable universe. These events are relativistic, jetted explosions of very massive stars at the end of their lives[a] (Woosley & Bloom 2006) which in principle are detectable out to redshifts of $z = 10 - 20$ (Lamb & Reichart 2000) as they outshine their host galaxy by several orders of magnitude at their peak brightness. GRB afterglows decay in brightness rapidly and are generally visible only for a few days after the initial burst[b]

[a]Long GRBs are physically distinct from both short gamma-ray bursts, which are associated with the merger between two compact objects (Gehrels et al. 2005; Rueda et al. 2018) and the more recent class of 'ultra-long' GRBs (Schady 2017), neither of which are discussed in this paper.

[b]This statement does not apply to the radio wavelengths, where GRB afterglows can be observed for hundreds of days after the initial burst, for example, see Rhodes et al. (2020).

(Gehrels, Ramirez–Ruiz, & Fox 2009), but opportunistic afterglow observations offer many unique opportunities to both directly and indirectly study the high redshift universe, provided that rapid spectroscopic follow-up observations are performed.

Because long GRBs originate from the deaths of short-lived massive stars they serve to randomly sample sites of star formation throughout the universe, providing an alternative means to estimate the cosmic star formation rate (SFR), albeit with the possible presence of a metallicity bias in GRB production efficiency (Robertson & Ellis 2012; Trenti, Perna, & Tacchella 2013). Using GRBs to measure the SFR becomes particularly valuable at redshift $z > 5$, where other tracers of the SFR such as Lyman-α emitters and Lyman break galaxies become scarcer and prone to selection effects (Stanway, Bremer, & Lehnert 2008).

GRBs can also act as beacons for identifying faint high redshift galaxies with a precise redshift determination from afterglow spectroscopy. Their association with massive star death means that GRBs select star forming host galaxies independently from their host luminosities (Klose et al. 2004). By performing deep follow-up imaging on GRBs that have both high accuracy position and redshift measurements, it is possible for observers to investigate the luminosity function of high redshift galaxies (Trenti et al. 2012; Tanvir et al. 2012; McGuire 2016). Such measurements give insight into the fraction of star formation occurring in very faint galaxies, beyond the sensitivity limit of imaging surveys in blank fields, and have the potential to constrain the faint end of the galaxy luminosity function, which is critical to determine the contribution of such galaxies to cosmic reionisation.

The extreme luminosity of GRB afterglows make them powerful probes of interstellar gas in their host galaxies, which has been used on many occasions to characterise the interstellar medium in distant galaxies (e.g., Klose et al. 2004; Berger et al. 2005; Vreeswijk et al. 2007; Fox et al. 2008; Prochaska et al. 2008). Furthermore, GRB afterglows have a high intrinsic UV luminosity and exhibit a featureless power-law spectrum in their rest frame UV region (Sari, Piran, & Narayan 1998), making them natural probes of the neutral hydrogen fraction in the IGM through both observations of damping wings in the Lyman-α line and from analysis of the Lyman forest (Miralda–Escude 1998; Mesinger & Furlanetto 2008; McQuinn et al. 2008; Hartoog et al. 2015; Lidz et al. 2021). In this respect GRBs provide an opportunity to investigate the earliest stages of reionisation (a stage which is challenging to probe with quasar spectroscopy) as GRBs have been observed beyond the redshift of the most distant known quasars (Tanvir et al. 2009; Cucchiara et al. 2011; Wang et al. 2021a), and are possibly present at $z > 10$ (e.g., from massive Population III stars) at a time when quasars are not expected to be active yet.

The main challenge facing astronomers wishing to leverage the utility of GRBs during the epoch of reionisation is the observational difficulty of identifying them. Measuring the redshift of a GRB requires afterglow observations at optical/IR wavelengths, since there are no spectral features in the prompt emission phase (the initial high-energy photon emission phase of the GRB). Systematically identifying $z > 5$ GRBs through afterglow spectroscopy is challenging since those objects are rare, accounting for $\lesssim 6\%$ of the GRB population (Greiner et al. 2011; Wanderman & Piran 2010), and their typical brightness implies that 6–10 m class telescopes are needed, in particular for the most interesting objects at $z \gtrsim 7$ (e.g., Tanvir et al. 2009; Salvaterra et al. 2009, where the Lyman-α line is shifted beyond $1\,\mu$m). To further complicate the scenario, given the afterglow's power-law decay in brightness any

redshift determination must be done rapidly, within hours of the initial burst.

Since its launch in 2004 November the Neil Gehrels Swift satellite (Gehrels et al. 2004) (hereon referred to as Swift) has been the largest contributor to the rapid identification of GRB targets with its ability to detect and instantly localise GRBs in large numbers ($\sim$100 per year) using its Burst Alert Telescope (BAT) as well as to observe the GRB afterglow using its X-Ray Telescope (XRT) and UV/Optical Telescope (UVOT) within a few hundred seconds of the onset of the prompt emission. However, for $z \gtrsim 5$ sources the Lyman break is redshifted beyond Swift's UVOT wavelength coverage. Furthermore, one cannot infer the presence of a high redshift GRB from a Swift non-detection since such events are observationally degenerate with 'Dark' GRBs, which are undetected at optical–NIR wavelengths due to dust extinction along the line of sight and account for $\sim$25–40% of the GRB population (depending on the definition used) (Klose et al. 2000; Greiner et al. 2011; Perley et al. 2013). Therefore, observations of the GRB afterglow in the near-infrared wavelengths are critically required to establish that a GRB originates at $z \gtrsim 5$.

The Gamma-Ray Burst Optical/Near-Infrared Detector (GROND) is a 7-channel imager at the ground-based MPG/ESO 2.2 m telescope designed with the specific purpose of observing GRB afterglows in the visible and near-infrared, with demonstrated success (Greiner et al. 2008). However, ground-based observatories suffer several substantial disadvantages in the rapid follow-up of GRB afterglows: they only have access to the fraction of sky visible overhead, they can only make observations at night time, and they depend heavily on the weather. These factors mean that a single ground-based facility—even if optimally designed and located such as GROND—only has the capability to follow-up promptly a small fraction of the GRBs detected by a satellite such as Swift.

Nano-satellite missions are becoming more common across multiple fields of astronomy, from high precision photometric observations of exoplanet transits in the optical wavelengths (ASTERIA, see Knapp et al. 2020) to high-energy astrophysics. For the latter, several missions are underway to use nano-satellite missions to detect the prompt X-ray and gamma-ray emission from GRBs, for example, the GRID mission (Wen et al. 2019) which recently made its first detection of GRB210121A (Wang et al. 2021b), the CAMELOT mission (Werner et al. 2018) which made its first successful GRB detection of GRB210807A with its GRBAlpha instrument,[c] the HERMES Technologic and Scientific Pathfinder constellation (Fiore et al. 2020) scheduled to launch 6-satellites in 2023, with one additional nano-satellite (SpIRIT) carrying the HERMES instrument being developed at the University of Melbourne,[d] and the BurstCube mission (Racusin et al. 2017) which is anticipated to launch in 2022.

In this paper we propose the use of a near-infrared space telescope on a nano-satellite with rapid re-pointing and low-latency communication capabilities to promptly detect GRB afterglows beyond $\sim$1 μm and determine high-z GRB candidates from photometric redshift measurements. Given the high signal-to-noise near-infrared image quality afforded by observing from space (above the atmospheric foreground), a relatively small ($\sim$0.15 m aperture) telescope has the same point source sensitivity in the H-band as a $\sim$2 m class telescope at ground level.

Specifically, we investigate one of the science goals of the SkyHopper mission concept[e], which aims to systematically follow-up GRBs identified by *Swift* and future gamma/X-ray satellites (such as the SVOM mission; see Wei et al. 2016). By combining a GRB detection in the near-infrared with the UV/optical photometry from Swift, a fast photometric redshift estimate of the GRB redshift is possible by determining the location of the Lyman break in one of the observed bands (e.g., Krühler et al. 2011). In this context, a near-infrared nano-satellite such as SkyHopper would complement the existing and future observational infrastructure to allow for efficient detection of high redshift GRBs.

This paper focuses on modelling the expected performance of rapid-response near-infrared observations from low Earth orbit, thus quantifying the potential to use a near-infrared nano-satellite to observe GRB afterglows. In Section 2 we outline the models we use to construct GRB templates for simulated observations. In Section 3 we define the specific mission scenario and telescope parameters used in this work. In Section 4 we describe the simulation framework used for line-of-sight access calculations and performing mock afterglow observations. In Section 5 we explore a range of different approaches to identify an optimal exposure strategy to detect GRB afterglows for use on-board a rapid-response near-infrared nano-satellite. In Section 6 we discuss the key results and highlights the expected science opportunities that would be enabled by a SkyHopper-like space telescope. The main conclusions from this work are presented in Section 7.

2. Simulating Gamma-ray burst events

In order to compare the effectiveness of GRB afterglow observations with ground and space telescopes, we construct a simplified but effective procedure to generate realistic models of GRB events. The GRB observables relevant to this work are the co-moving redshift distribution of long GRBs, the near-infrared afterglow luminosity, and the early-time near-infrared afterglow light curve decay indices.

Due to the complexities in accurately simulating the detection of high-energy GRB photons by the Swift BAT we shall exclude models of GRB prompt emission from our simulation, despite the GRB prompt emission (both peak and isotropic energy release) being positively correlated with optical afterglow luminosity and negatively correlated with afterglow decay rate (Oates et al. 2015). By neglecting the correlation between prompt and afterglow luminosity we are likely under-estimating the performance of our nano-satellite, as those bursts detected by Swift must meet the minimum prompt emission flux threshold of Swift's BAT and therefore have a higher afterglow luminosity, making them easier to detect. Conversely, by neglecting the anti-correlation to afterglow decay rate we slightly overestimate the number of detections since brighter bursts decay more rapidly than fainter ones. We leave the inclusion of these correlations to future works, as simulating Swift's detection of high-energy photons is beyond the scope of this work.

We further choose to exclude the correlation between afterglow luminosity and light curve decay indices (Oates et al. 2015). Numerous factors make it difficult to precisely quantify this correlation at the time of writing (limited sample size of UV/Optical afterglow observations, assumptions regarding afterglow off axis emissivity modelling, k-correction type effects, etc.) and so for

[e]https://skyhopper.research.unimelb.edu.au/.

simplicity we treat the two observables as independent in our simulation.

By neglecting GRB prompt emission modelling and assuming that Swift is able to detect every accessible burst, our simulation will unrealistically increase the total number of bursts detected by Swift. We shall therefore present our results as the fraction of Swift triggers successfully detected by the near-infrared nano-satellite instead of the total number of GRB detections.

2.1. Redshift distribution

To generate a sample GRB population with a realistic redshift distribution we draw on the work of Wanderman & Piran (2010), who derive the differential co-moving space density of GRBs at redshift z:

$$R_{GRB}(z) = \begin{cases} (1+z)^{n_1} & z \leq z_1 \\ (1+z)^{n_1-n_2}(1+z)^{n_2} & z > z_1 \end{cases} \quad (1)$$

With $z_1 = 3.11$, $n_1 = 2.07$, $n_2 = -1.36$.

The observed distribution of GRB redshifts can then be calculated by multiplying this function by the co-moving volume element dV/dz and correcting for cosmological time dilation:

$$R(z) = \frac{R_{GRB}(z)}{(1+z)} \frac{dV(z)}{dz} \quad (2)$$

To generate the redshift of each simulated GRB event we randomly sample from the probability distribution described by Equation (2) over the domain $z \in \{0, 10\}$.

2.2. Afterglow light curve

Modelling the observed near-infrared light curve of a GRB afterglow for a high redshift ($z > 5$) GRB is an equivalent problem to modelling the rest frame UV/optical light curve due to the cosmological redshifting of light. As such, we draw on the work of Oates et al. (2009), who analysed the statistical properties of early-time GRB light curves observed by Swift UVOT. The Oates sample consists of 27 long GRB afterglows observed by Swift between 2005 and 2007 with redshifts ranging from 0.44 to 4.41, which corresponds to sampling rest frame wavelengths as low as ∼90 nm (any wavelengths shorter than this are absorbed by neutral hydrogen in the host galaxy and IGM). While rest frame wavelengths above ∼250 nm are redshifted beyond the *H*-band for bursts with $z > 5$, the sample exhibits no substantial evolution in the light curve decay index as a function of rest frame wavelength (though with the small sample size it is difficult to state this with confidence). We therefore adopt the results of Oates et al. (2009) under the assumption that there is no differential colour evolution within the intrinsic UV afterglow spectrum (Sari et al. 1998; Kumar & Zhang 2015), and that the afterglow light curve does not undergo significant evolution with redshift in the emission rest frame.

Motivated by the results from Oates et al. (2009) we model the afterglow in the observer frame using a broken power-law with a break at $t_{\rm obs} = 500$ s, separating the light curve into an 'early-time' and a 'late-time' decay phase, where the observer frame afterglow flux F varies as:

$$F(t) \propto t^{-\alpha} \quad , \quad \alpha = \begin{cases} 0, & t_{\rm obs} < 500 \text{ s} \\ 0.85, & t_{\rm obs} > 500 \text{ s} \end{cases} . \quad (3)$$

For the early-time decay curve we have simplified the results of Oates et al. (2009), who found that some events decay rapidly and some experience re-brightening during this phase. To take an average of this behaviour we model the decay slope as flat during this epoch because it also serves to impose a flux cutoff on the model, preventing the magnitude of the afterglow from diverging at early times (as it would in a power-law model with $\alpha > 0$) and thus potentially artificially inflating the number of GRBs detected by the nano-satellite at very early times. The value of the late-time decay index is set in accordance with the findings of Oates et al. (2009).

To test the validity of the modelling assumption that the light curve is flat for $t < 500\,$s, we will also test the case of a rising light curve at early times, that is, $\alpha_{t<500s} = -0.2$ in accordance with the maximum early-time decay index from Oates et al. (2009).

2.3. Afterglow brightness

Ideally, the observed H-band magnitude of the afterglow would be simulated by sampling from an early-time UV afterglow luminosity function, thereby calculating the observed brightness of the event by taking into account the jet orientation and off-axis emissivity, the extinction along the line of sight and cosmological redshifting.

However, this approach contains many points of uncertainty. Preliminary work has been carried out to derive an afterglow luminosity function (e.g., Oates et al. 2009; Wang et al. 2013), but the results are affected both by limited number of observed GRBs, and by assumptions on the GRB central engine theoretical modelling which affect off-axis emission (e.g., for a purely hydrodynamic jet, see van Eerten, Zhang, & MacFadyen 2010), an area which is not yet well understood.

Given these uncertainties we instead sample the H-band magnitude of the afterglow directly from a simplified probability distribution and apply a series of corrections to the observed magnitude in order to account for the effects of the distance to the event, k-correction, time dilation and host galaxy extinction.

To construct a simplified probability distribution for H-band magnitude we draw on a sample of 88 long GRB afterglows observed by GROND between 2007–2016 between redshifts $z \sim 0.3 - 9$. All GRBs in the sample were detected before rest frame $t = 4\,$h ($\sim 1.5 \times 10^4\,$s), and the data is recorded in the K-band and corrected into the H-band using $H = K + 0.20$ (which presumes the spectral energy distribution $F \propto \nu^{-\beta}$ with $\beta = 0.6$ as in the optical wavelengths (Kann, Klose, & Zeh 2006; Greiner et al. 2011) to match the bandpass of our simulated near-infrared telescope (see Section 3.1.1). An overview of the GRB detection times and corrected H-band magnitudes colour-coded by event redshift is shown in Figure 1, where the redshift distribution of the 59 events with measured redshift is consistent with the distribution described by Equation (2) (Figure 1, bottom panel).

In order to conveniently sample from this magnitude distribution, we rescale the magnitude of each event to the common observer frame time of $t_{obs} = 10^3\,$s post-burst using the light curve described in Section 2.2 (Equation (3)) and fit the resultant magnitude distribution with a Gaussian distribution with mean $H_{AB} = 18$ and $\sigma = 2$ (Figure 2).

To correct the magnitude of each GRB event to account for distance, k-correction, time dilation and host-galaxy extinction we undertake the following procedure:
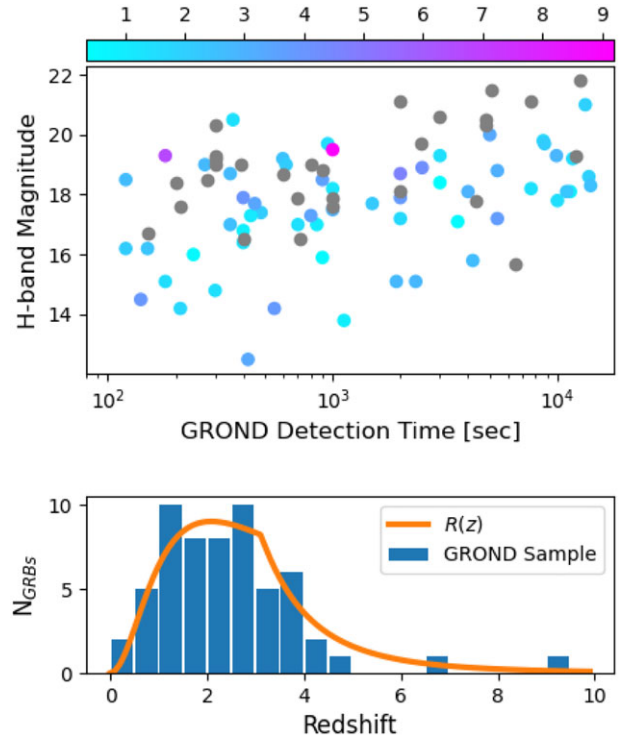
Figure 1. Overview of the sample of 88 published GRB afterglows observed by the GROND instrument with redshifts between $0.347 < z < 9.2$ used as baseline in this work. Top: The x-axis shows the time that GROND detected the GRB, and the y-axis represents the corrected H-band AB magnitude of the afterglow at the time of detection. The colour of each data point represents the redshift of the event, where grey data points indicate GRBs for which a redshift determination was not made. Bottom: Histogram of the redshift of the 59 GRBs in the GROND sample with a measured redshift. The yellow curve plots the GRB rate function described by Equation (2) (with arbitrary normalisation to match the scale of the data).
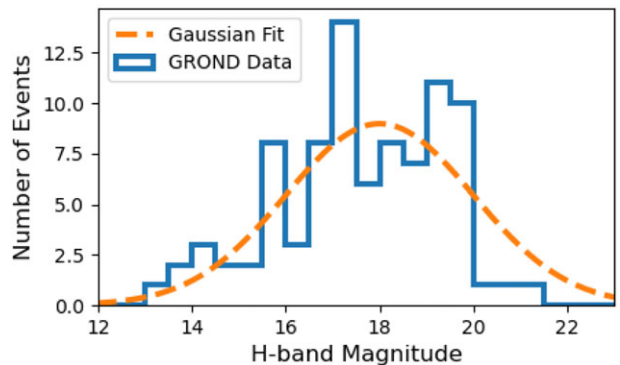
Figure 2. Rescaled H-band AB magnitude distribution from 88 GRB afterglows observed by GROND. Each GRB has been rescaled to a common time $t = 10^3$ s post-burst (in the observer frame) using the light curves described in Section 2.2.

We first generate a sample of $\sim 10^5$ afterglows each with a redshift (using Equation (2)) and H-band magnitude at $t = 1\,000\,$s (sampled as per Figure 2). For each event we calculate a correction to the observed magnitude resulting from distance (D), k-correction (k) time dilation (τ) and reddening effects, assuming that the magnitude sampled from the GROND distribution corresponds to a GRB at $z = 2$ (the mean GRB redshift[f]).

[f]https://www.mpe.mpg.de/jcg/grbgen.html.

$$\Delta m_D(z) = -5 \log \left(\frac{D_L(2)}{D_L(z)} \right) \qquad (4)$$

$$\Delta m_k(z) = (-\beta + 1) \cdot 2.5 \log \left(\frac{1+2}{1+z} \right) \qquad (5)$$

$$\Delta m_\tau(z) = \alpha \cdot 2.5 \log \left(\frac{1+2}{1+z} \right), \qquad (6)$$

where we adopt the notation for observed GRB flux $F \propto \nu^{-\beta} t^{-\alpha}$ (with $\alpha = 0.85$ as in Section 2.2 and $\beta = 0.6$ as before), and $D_L(z)$ is the luminosity distance to the source.

To simulate the effects of dust reddening as a function of redshift, we rescale the extinction curve of the Small Magellanic Cloud (Gordon et al. 2003) to arbitrary redshift z by normalising it to the value of the UV host galaxy extinction at that redshift ($A_{M_{UV}}(z)$ from Trenti, Perna, & Jimenez 2015) at $\lambda = 0.14 \,\mu$m. We then use the appropriately rescaled curves to subtract the extinction that would have occured if the burst had originated from redshift 2, and then apply the extinction for a burst originating from redshift z.

Using the sample of 10^5 afterglows we calculate the mean magnitude correction μ due to distance, k-correction, time dilation and reddening effects. Then, in order to preserve the mean of the GROND distribution we apply a magnitude correction to each event proportional to the difference from the mean for each of these quantities:

$$m_{\text{obs}} = m_{\text{sampled}} + (\Delta m_D(z) - \mu_D) + (\Delta m_k(z) - \mu_k)$$
$$+ (\Delta m_\tau(z) - \mu_\tau) + \left(\Delta m_{\text{reddening}}(z) - \mu_{\text{reddening}} \right), \qquad (7)$$

where we have calculated $\mu_D = 0.36$, $\mu_k = -0.06$ $\mu_\tau = -0.12$, $\mu_{\text{reddening}} = 0.03$ when sampling from the redshift distribution in Section 2.1.

3. Satellite mission scenario & telescope parameters

This section outlines the relevant details taken into account when simulating the operations of the proposed near-infrared nano-satellite, as well as our approach to simulating the Swift space telescope and GROND instrument.

3.1. Nano-satellite

Evaluating the performance of a nano-satellite in detecting GRB afterglows requires defining a specific mission scenario. This section will outline the SkyHopper[g] nano-satellite concept which we will use as the baseline mission scenario for this study, focusing on the key mission parameters used in our analysis and the motivations behind them.

3.1.1. SkyHopper satellite concept

The SkyHopper space telescope is a concept for a rapid-response near-infrared nano-satellite being developed by the University of Melbourne and several other Australian and international collaborators.

The spacecraft concept is based on a 12U CubeSat format ($\sim 36 \times 22 \times 24 \text{cm}^3$), which will house a $20 \times 10 \text{cm}^2$ rectangular telescope mirror (0.15 m equivalent aperture), feeding light to

[g]https://skyhopper.research.unimelb.edu.au/.

Table 1. SkyHopper reference signal-to-noise ratios calculated for an $m_{AB} = 19.5$ point source.

Exposure time (min)	Reference SNR
1	1.1
2	1.84
3	2.42
4	2.91
5	3.34
6	3.73
7	4.08
8	4.41
9	4.71
10	5

a $2\,048 \times 2\,048$ H2RG IR image sensor which is actively cooled to 145 K (a reference nominal operating temperature for near-infrared astronomical observations below 1.7 μm; e.g., see Dressel 2021 for Wide Field Camera 3 on the Hubble Space Telescope). SkyHopper is designed to have a $\sim 1.5 \,\text{deg}^2$ FOV, and will be capable of maintaining an RMS pointing stability of $<4''$ during observations.

SkyHopper is designed to image in four filters across the near-infrared (0.8–1.7 μm) simultaneously, with Swift UVOT-like celestial body avoidance angles (Table 3). The simultaneous four-band imaging capability allows the spacecraft to determine robustly a photometric redshift estimate for GRBs with $z \sim 6$–13 by measuring the Lyman break between two filters. This technique was first applied to the study of high redshift GRBs to estimate the redshift of GRB050904 at $z = 6.3$ (Tagliaferri et al. 2005) (in agreement with later spectroscopic measurements from the Subaru telescope; see Kawai et al. 2006) and has been used on several occasions over the following years (e.g., see Greiner et al. 2009; Krühler et al. 2011; Cucchiara et al. 2011). Given the rapid decay in afterglow brightness it is critical that these observations are performed simultaneously to ensure that reduced flux in one filter can be unambiguously associated with absorption by neutral hydrogen rather than the natural decay in brightness of the source.

The near-infrared detector on SkyHopper has a design goal to achieve point source sensitivity of $m_{AB} = 19.5$ [$t = 600$ s; 5σ; H-band], where the dominating noise contribution at 1.6 μm is assumed to be the thermal foreground from a relatively warm ($T \sim 250$ K) telescope baffle. For exposures longer than 10 min, the signal-to-noise ratio (SNR) of a single exposure can be approximated by the rescaling SNR $\propto \sqrt{t}$, however for shorter exposures this approximation breaks down as the readout noise becomes dominant.[h] Within this domain we have computed a set of reference signal-to-noise values for a $m_{AB} = 19.5$ point source with exposures ranging from 1 to 10 min, using an exposure time calculator for the SkyHopper telescope. Results are summarised in Table 1.

[h]Poisson noise from source photons may dominate for very bright targets, but in this case the afterglow would be clearly detected, hence we can neglect to take it into account to determine marginal S/N detections.

3.1.2. Orbit selection

The small dimensions of a nano-satellite limit the amount of on-board space available for batteries, making it desirable for it to fly in sun-synchronous orbit as it allows the spacecraft to receive a constant stream of electricity via its solar panels and thus circumvents the need to store large amounts of power.

For the analysis performed in this work, we model a nano-satellite in a Sun-synchronous dawn/dusk orbit:

- altitude $h = 550$ km;
- circular orbit (eccentricity $= 0$);
- dawn/dusk Sun-synchronous orbit (inclination $= 97.6°$).

3.1.3. Satellite TeleCommand and slew rate

Due to the power-law decay in afterglow brightness it is critical to understand how quickly a satellite can commence observations of the target, as the increased flux at early times gives the instrument a much higher probability of detecting the afterglow if it can arrive on target quickly. The two design aspects of the nano-satellite that are relevant to this calculation are the satellite's slew rate and the telecommunications latency in uplinking a re-pointing command to the spacecraft.

The SkyHopper spacecraft is capable of slewing at a nominal rate of $2 \deg s^{-1}$. In order to calculate the time delay between re-pointing command and target acquisition we assign the satellite a random pointing (one which obeys bright-source avoidance angles) at the time of each burst and calculate the slew time as the angular difference between the current pointing and the target coordinates divided by the slew rate.[i]

With regards to satellite TeleCommand there are several options to choose from when designing a space mission. Traditional Telemetry and TeleCommand (TTC) schemes use ground stations or government satellite relay networks such as NASA's Tracking and Data Relay Satellites (TDRS) to communicate in near real-time with satellites in orbit, but utilising this network is not only costly but requires the satellite to be equipped with a large antenna which often exceeds the volume/mass constraints on nano-satellite missions. Furthermore, a single ground station could only provide infrequent communications since it would only be seen ∼ twice per day by a satellite in low Earth Sun-synchronous orbit due to Earth's rotation.

A viable solution for budget-limited nano-satellite missions is to leverage existing machine-to-machine orbital telecommunication networks to send messages to satellites in low Earth orbit (LEO). One drawback to this solution is that such networks were not designed for space applications, and thus there can be delays in uplinking a command to the spacecraft as it passes between gaps of the telecommunications satellite coverage beams.[j] For our afterglow detectability study, this acts effectively as a stochastic-like source of TeleCommand latency for the space mission, where the distribution of latency times depends on the communications network being used and the satellite's orbit.

Table 2. Nominal Iridium TeleCommunications uplink latency for 550 km sun-synchronous orbit (as per Mearns & Trenti 2018). The 'Probability' column indicates the probability that contact is established with the satellite by the given time.

Time (mins)	Probability (%)
Immediately	31.8
1	41.4
2	50
6.6	68
10	74.2
25.4	95
40	99.7

In this work we presume that the nano-satellite mission is budget-limited and simulate TeleCommand using the Iridium machine-to-machine orbital telecommunications network, sampling directly from a statistical model of the nominal Iridium network performance derived by Mearns & Trenti (2018) for a 550 km Sun-Synchronous Orbit (Table 2).

3.2. Swift BAT

3.2.1. Prompt Gamma-ray detection

For simplicity we presume that Swift can detect every GRB that is not directly obstructed by the Earth, Sun or Moon. This deliberately ignores the fact that the prompt emission from the GRB must meet a certain flux threshold to be detected, and that BAT only has a 1.4sr FOV.[k] Ignoring these constraints artificially inflates the number of GRBs observable by Swift and hence the number of Swift triggers available for follow-up in a given year.

This simplification however will not influence our overall results because they are presented as the ratio of GRBs detected divided by the number of Swift GRB triggers, which is agnostic to the total number of GRBs detected by Swift.

3.2.2. Orbit details and downlink capabilities

We model the orbit of the Swift space telescope using its current orbital parameters at the time of writing:

- altitude $h = 561$ km;
- circular orbit (eccentricity $= 0$);
- inclination $= 20.6°$.

The Swift Observatory utilises NASA's TDRS System to automatically send GRB triggers to the GRB Coordinates Network (GCN). With regards to its source localisation capabilities, upon detecting the prompt gamma-ray emission from a GRB, Swift's BAT is able to autonomously localise the event to within 1–3 arcmin, downlinking this information as a GRB trigger distributed via the GCN approximately 20 s after the initial detection (Troja 2020). Following this, Swift automatically re-points the telescope (providing the burst meets visibility constraints) to observe the afterglow with its XRT, refining the localisation to within a radius of a few arcsec (Evans et al. 2009) and distributing this via the GCN at approximately 100 s post-burst.

---

[i]In actual operations the nano-satellite would need to maintain bright-source avoidance angles throughout its slew, but within the simplified framework of this simulation a direct slew yields a comparable result.

[j]The sat-com beams generally overlap with no gaps on the Earth surface, but this is not the case in low Earth orbit as constellations such as Iridium or Globalstar only orbit at ∼780–1400 km.

[k]https://swift.gsfc.nasa.gov/about_swift/bat_desc.html.

For the purposes of this simulation, we assume that every GRB trigger takes exactly 20 s to downlink from Swift to Earth. The prompt BAT localisation of 1–3 arcmin is sufficient for the nano-satellite to reliably observe the GRB due to its 1.5 deg^2 FOV (as per SkyHopper's design requirements) which makes a shift in source position of 1–3 arcmin irrelevant. This means the nano-satellite can slew and commence observations of the source without the need to wait for Swift's enhanced GRB localisation.

3.3. GROND

In order to compare space-based and ground-based near-infrared afterglow observations we employ a simplified model of the GROND instrument, taking into account the key points of difference between ground and space-based observations: sky visibility, weather and relative sensitivity at near-infrared wavelengths.

3.3.1. Modelling ground-based observations

GROND is one of the instruments on the MPG/ESO 2.2m telescope at La Silla observatory in Chile.[l] Simulating GROND's operations requires taking into account not only the line-of-sight (LoS) to the target but also the time of day and the weather in Chile.

With regards to the local time of day we presume that GROND is only able to observe GRB afterglows during the hours of 8pm and 6am local time (UTC-4). While in reality these times shift seasonally throughout the year, we consider them to be a sufficient representation of 'average' night time such that it will not substantially influence our results to neglect the seasonal variation.

To estimate the impact of weather on ground-based observations we draw on weather data collected by the astronomers on duty at La Silla between 1991 Jan–1999 May,[m] which records the fraction of nights each month where 'photometric' observing conditions (no visible clouds, transparency variations under 2%) were met, as well as the number of 'useless' nights where the weather conditions meant that the telescopes had to be closed due to wind and/or humidity. Both of these metrics are inaccurate representations of the impact of the weather on GROND's operations—often GROND would make observations (at a reduced sensitivity) even when nights were not 'clear' (less than 10% of the sky covered in clouds, transparency variations under 10%), let alone 'photometric'. Conversely it is an overestimate to say that GROND was able to observe on every night where the telescope dome was open.

Given the lack of more robust weather data and for the sake of comparing ground and space-based observing conditions we shall take these two scenarios—'photometric' observing conditions, and the telescope dome being open—as the lower and upper bounds on the impact of weather on GROND. Averaging the ~8 yr of data recorded at La Silla, 62% of nights are 'photometric', and only 15% are 'useless' (i.e., 85% are usable).[n] For both scenarios we presume that GROND maintains the same point source sensitivity described in Section 3.3.2, and that the classification of a given night will apply to the entirety of the night.

3.3.2. Telescope properties and detector specifications

GROND functions as part of the 2.2 m MPG/ESO telescope, which it shares with both the Wide Field Imager and the Fibre-fed Extended Range Optical Spectrograph (FEROS) spectrograph.

In the past, when a GRB trigger is picked up through the GCN, a mirror is folded into place (taking ~20 s) to re-direct the incoming light into the GROND instrument (Greiner et al. 2008). GROND software would then autonomously re-point the telescope to the target as soon as it becomes accessible in the sky over Chile. Upon a night trigger, the software interrupts the running exposure while reading out the data, which depending on the instrument can take 10-45 s. For simplicity, in this work we presume the entire instrument changeover and re-pointing of the observatory dome takes 120 s for every burst, meaning that the minimum possible time delay between Swift broadcasting a GRB trigger and GROND commencing its observations is 120 s (with extended time delays being related to the day/night cycle and LoS accessibility of the target).

In our modelling of afterglow observations we presume that GROND has a H-band 5-σ limiting magnitude of $H_{AB} = 20.1$ for an 8 min Observation Block[o] (an 8 min Observing Block typically consists of four 1.6 min exposures in the visual channels, and forty eight 10 s exposures in the J, H and K bands). This value reflects the sensitivity of GROND observing under a new Moon with typical 1.0″ of seeing and minimal airmass (1.0), and so it represents better than typical observing conditions from the ground. Given the spatially and temporally fluctuating atmospheric foreground noise GROND's afterglow observations consist of a number of short exposures of the target which are stacked together. To simulate this behaviour we presume that GROND takes a number of 8 min exposures of the target, and that the noise level is the same between exposures, and so SNR $\propto$ (number of exposures)$^{1/2}$ $\propto$ (exposure time)$^{1/2}$. Therefore, if GROND observes a $H_{AB} = 20.1$ source for a total duration t:

$$\text{SNR}_{\text{ref}}(t)|_{\text{GROND}} = 5 \cdot \sqrt{\frac{t}{8 \text{ min}}}. \tag{8}$$

In reality (i.e., taking into account the time-variable atmospheric foreground noise), Equation (8) is somewhat overestimating SNR_ref in the case of a long total exposure time, and under-estimating it for shorter exposures, but the overall impact on our analysis is not significant.

In terms of scheduling observations we presume that GROND observes the afterglow for as long as possible on the first night after the GRB trigger, but that it does not schedule further observations for subsequent nights. Note that until 2016 every GRB was (weather permitting) observed with GROND also in the following nights.

4. Afterglow observation pipeline

This section presents an overview of our simulation approach, and details the methods we use for line-of-sight access calculations and to determine the signal-to-noise ratio of the afterglow photometry.

4.1. Simulation strategy

To investigate the level to which a nano-satellite can perform follow-up observations on Swift GRB triggers we simulate the observational pipeline as follows:

- Using the models outlined in Section 2 we randomly generate GRB events assuming a uniform distribution on the celestial sphere, where each event occurs at a random time sampled

---

[l]https://www.eso.org/sci/facilities/lasilla/telescopes/national/2p2/overview.html.
[m]https://www.eso.org/sci/facilities/lasilla/astclim/weather/tablemwr.html.
[n]By definition, a night that is 'photometric' will also be 'usable'.

[o]https://www.mpe.mpg.de/jcg/GROND/operations.html

from a uniform temporal distribution throughout the year. We also randomly generate the number of GRB events which occur each year, sampling this from a Poisson distribution centred on 380 GRBs yr^{-1} (the average GRB detection rate of the Fermi Gamma-ray Burst Monitor if it could observe the whole sky von Kienlin et al. 2020). For each GRB we randomly generate the H-band magnitude and the redshift of the event, where the redshift-corrected H-band magnitude of the afterglow at observer-frame $t = 10^3$ s is generated as per Section 2.3, and the redshift of each event is generated according to Equation (2).

- For each GRB event, we determine whether Swift BAT is able to detect the GRB (Section 3.2). If it successfully detects the GRB, there is a 20 s delay in downlinking the GRB trigger to Earth using the NASA's TDRS network. Additionally, we determine whether the GRB is accessible for UVOT follow-up by checking whether Swift can achieve LoS access to the target within 90 min. If Swift BAT does not detect the GRB, no further action is taken.

- Upon receiving a GRB trigger in the nano-satellite control centre, a random uplink time and re-pointing time are sampled from their relevant probability distributions (Section 3.1.3) and we calculate the earliest time that the satellite is able to achieve LoS access to the target. We perform the same computation for observations with the GROND instrument, excluding the uplink communications latency from the calculation, and taking into account the additional variables of the day/night cycle and the probability of the weather prohibiting observations.

- As soon as an observatory achieves LoS access to the target it begins making observations of the afterglow for the entire duration of the observing window, meaning that observations are only cut short either by LoS obstruction, or (for ground-based observatories) dawn. The SNR of the entire observation is then calculated, taking into account the power-law decay of the afterglow brightness. Observations totalling SNR >5 are considered a successful afterglow detection.

4.2. Simulating afterglow observations

To calculate the flux captured in a single exposure we integrate the broken power-law light curve (Equation (3)) over the duration of the exposure (converting the magnitude of the GRB at the time of the observation into a reference flux in order to perform the integration), before converting the resultant fluence back into an observed AB magnitude 'm''. The SNR for an exposure of duration t is then computed from the reference SNR for the instrument considered, assuming that the signal scales as flux:

$$\text{SNR}(t) = \text{SNR}_{\text{ref}}(t) \cdot 10^{(m_{\text{ref}} - m(t))/2.5}, \tag{9}$$

where m_{ref} and $\text{SNR}_{\text{ref}}(t)$ are reference magnitudes and signal-to-noise values for an exposure of the same duration t. Note that (1) this equation assumes that the afterglow is faint, that is, the noise does not depend on the afterglow signal (e.g., background/foreground or detector noise limited); (2) these reference values are specific to the telescope making the afterglow observation; (3) the reference values and relevant noise sources for the near-infrared nano-satellite and the GROND instrument are discussed in Sections 3.1 and 3.3 respectively.

In the case where a telescope takes multiple exposures of the source each with their own individual durations t_i the SNR from each exposure can be combined as:

Table 3. Celestial body limb avoidance angles for each of the telescopes modelled in this work.

Satellite	Avoidance angle		
	Sun	Earth	Moon
Swift (BAT)	0°	0°	0°
Swift (UVOT)	46°	28°	23°
SkyHopper	46°	28°	23°
GROND	n/a	20°	9°

$$\text{SNR}\left(n \text{ exposures}\right) = \sqrt{\sum_{i=1}^{n} \text{SNR}_i(t_i)^2}, \tag{10}$$

where this equation assumes that the noise is uncorrelated between exposures. In this work we consider a 'detection' to be achieved when an afterglow observation reaches a cumulative SNR greater than 5.

4.3. Line of sight access modelling

The opportunistic observation of stochastic GRB events relies on having LoS access to the target. We simulate the influence of this directly, designing our own LoS access code and implementing it as follows.

We first generate a dataset of cartesian coordinates for each instrument relevant to our simulation—a nano-satellite, Swift and GROND—as well as the Sun, Earth and Moon (the most relevant astronomical bodies to Earth-based observations due to their apparent size and brightness). This dataset comprises of timestamped cartesian coordinates spanning the entirety of the year of 2023 (2023 Jan 1 00:00:00 UTCG—2024 Jan 1 00:00:00 UTCG) in 20 s timesteps. For simplicity, when simulating multiple years of observations we simply repeat the year of 2023 multiple times.

Using this set of cartesian coordinates we are able to determine the line-of-sight accessibility of an arbitrary point on the sky for each instrument, which enables us to determine the time that each instrument can access the GRB (after considerations of uplink and slew time are taken into account) as well as how long each instrument has to observe before the target coordinates become obstructed. To do so we take into account the radius and distance from the observer of each celestial body as well as the bright source avoidance angles for each telescope, which are quoted in Table 3. The operational constraints for Swift UVOT were taken from NASA's online UVOT digest.[p] A full discussion of each instrument listed in the table can be found in Sections 3.1 (Nano-Satellite), 3.2 (Swift BAT/UVOT) and 3.3 (GROND).

5. Optimising GRB observations in low earth orbit

When making early-time observations of GRB afterglows it is important to construct an exposure strategy that takes into account the power-law decay in the source brightness, since taking a single long exposure of the target risks washing out the increased signal from the bright early-time afterglow with background noise. It is possible to achieve a higher signal-to-noise observation by taking short exposures of the source at early times to capitalise on the increased source brightness before taking longer exposures

---

[p]https://swift.gsfc.nasa.gov/analysis/uvot_digest/numbers.html.

as the power-law decay begins to flatten off. We use Monte Carlo methods to compare the performance of several different exposure strategies in order to identify the optimal strategy for use on board a rapid-response space telescope.

5.1. Optimising an afterglow exposure strategy

We define an exposure strategy to be a sequence of consecutive exposures performed by the near-infrared nano-satellite, where for simplicity each exposure lasts for an integer number of minutes.

Presuming that the nano-satellite can achieve LoS access to the source, the amount of time available to observe the GRB afterglow depends on the satellite's position in orbit: it can be as short as a few seconds (if the GRB is about to be eclipsed by an astronomical body) and as long as $\sim$45 min (this depends on the bright source avoidance angles—for SkyHopper in LEO this corresponds to slightly less than half an orbit). As such, the exposure strategies we explore must be able to vary their total duration to match the duration of the observing window. Throughout this work we express exposure strategies as ordered lists where each number represents the duration of a given exposure in minutes. For example:

- A strategy that performs a 1 min exposure followed by 2 and 3 min exposures is notated as '[1, 2, 3]'.
- To construct variable length observing strategies we simply repeat the final exposure in the sequence for as long as the observing window allows. In our notation this is indicated with an ellipse, where the ellipse follows the exposure which will be repeated. For example, a strategy which consists of a 1 min exposure followed by consecutive 5 min exposures is notated as '[1, 5, 5, . . .]'.

The performance of each exposure strategy is judged with respect to two key criteria: the overall detection fraction and the speed at which it can make a detection.

Specifically, we focus on the detection of those GRBs which are also observable by Swift's UVOT,[q] so that it would be possible to distinguish between high redshift and 'Dark' GRBs by comparing the near-infrared photometry from SkyHopper with UV/optical data from Swift. While we do not directly model the Lyman-α dropout for high redshift bursts, we assume that all bursts with $z > 5$ will not be detected by Swift, and therefore a detection in the near-infrared would confirm the GRB as originating from high redshift.

For ease of reference we refer to a 'UVOT-observable' Swift trigger as one which can be observed by Swift's UVOT within 90 min (approximately one Swift orbit) of the GRB prompt emission (i.e., outside the Earth, Sun and Moon limb avoidance angles cited for Swift (UVOT) in Table 3). This does not mean that the burst was necessarily detected by UVOT—simply that UVOT could safely re-point to the burst without violating celestial body avoidance angles.

The detection fraction of a given strategy is defined as:

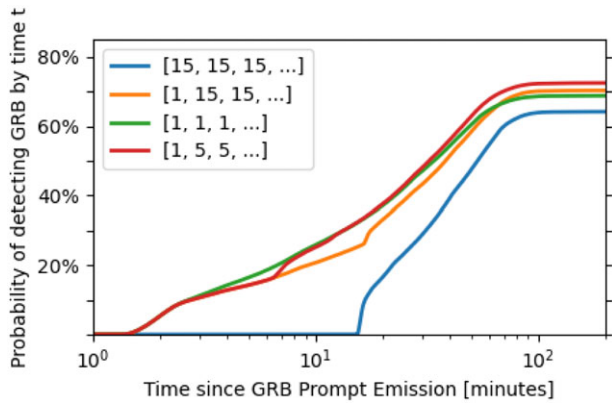$$\text{Det. Fraction} = \frac{\text{Follow-Up Detections}}{\text{UVOT-observable Triggers}}. \quad (11)$$

For ease of comparison GROND is also evaluated on the same definition of detection fraction.

5.1.1. Exposure strategy constraints

If one had access to limitless computing power and uplink TeleCommand capability then an afterglow exposure strategy could be tailored dynamically to observe each GRB, taking into account the time delay between the GRB occurring and observations commencing as well as the duration of the observing window. However, given the limitations in on-board computing power and the benefits in software development and testing from limiting TeleCommand complexity, we adopt a simplified approach in constructing an observing strategy: we presume that the exposure strategy is hard-coded into the satellite and the ground can only command the satellite to take an integer number of exposures from that strategy. One consequence of this approach is that the satellite will rarely be able to observe for the full duration of the exposure window (e.g., if an observing strategy consists of repeated 10 min exposures and the observing window is 19 min long, the satellite would only be able only perform a single 10 min exposure of the target).

A key consideration when designing an near-infrared observing strategy for use in low Earth orbit are cosmic rays (CRs) and 'snowball' events, which can compromise an exposure by depositing a large amount of charge onto the detector after a collision. Instead of directly simulating the impact of these events we simply place an upper limit of 15 min on any single exposure duration: CR events occur at a rate of 11 ± 5 CR/s for Hubble's WFC3/IR and impact 1–10 pixels on the detector (Dressel 2021) meaning that any given pixel has a <5% probability of being impacted by a cosmic ray for a 15 min exposure assuming the same sized near-infrared detector.[r] Considering that MgCdTe detectors such as Hubble's WFC3 and SkyHopper's can be read non-destructively multiple times during a single exposure, offering further mitigation against cosmic rays via fitting of the charge accumulation in a pixel versus time, we judge this impact as sufficiently low to not affect significantly our nano-satellite's ability to detect GRB afterglows, which only occupy a few pixels on the detector given they are point sources. The same argument applies to 'snowball' events, which occur at a much lower rate of $1\,\text{h}^{-1}$ (Green & Olszewski 2020). We also note that with these assumptions the impact of cosmic ray hits is smaller than or comparable to the probability that a GRB afterglow is along the same line of sight as a brighter foreground source (e.g., Galactic star), which is again estimated to be at the few percent level based on 2MASS number counts (estimating that a typical point source will occupy a few tens of pixels on the detector, and taking into account the density of point sources brighter than $H_{AB} = 20$ is of the order $\sim$$10^3\,\text{deg}^{-2}$ Skrutskie et al. 2006), and affects equally all observatories irrespective of their location on the ground or in space (note that brighter point sources will have diffraction spikes covering more pixels, but their number density in the sky drops more rapidly than the increase in the area covered, so for the estimate we can consider

---

[q]Here we wish to make a distinction between 'observable' (meaning the telescope is able to safely point to, and take exposures of, the target coordinates) and 'detectable' (meaning the observed flux from the afterglow is bright enough to be detected by the instrument). This distinction is critical, because while high redshift targets can be observable by Swift, they are not detectable due to the Lyman-α dropout, which is precisely what enables a photometric redshift estimation.

[r]SkyHopper's design is based on a similar detector.

Figure 3. Cumulative probability of making a follow-up detection on a UVOT-observable Swift GRB trigger for four different observing strategies, where $t = 0$ is the time the burst is detected by Swift BAT. Data is generated by simulating $\sim 10^5$ trial afterglow observations with each strategy.

Figure 4. PDF of the yearly follow-up detection fraction for four different observing strategies. The *x*-axis represents the fraction of UVOT-observable Swift triggers which were detected using the given exposure strategy. Data is generated by binning the results of $\sim 10^5$ trial afterglow observations into each year of simulated observations.

point sources of brightness comparable to the faintest afterglow we aim to detect).

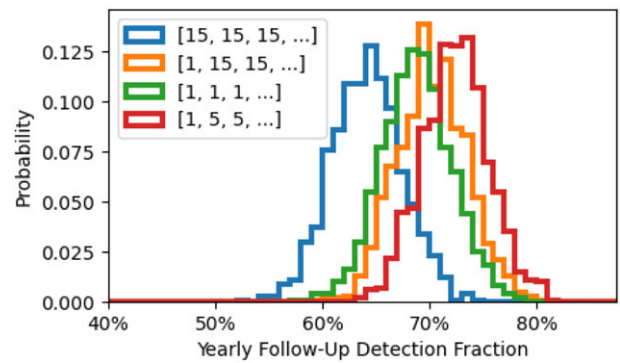### 5.1.2. Identifying an optimal observing strategy

We use a Monte Carlo approach to identify the optimal exposure strategy, testing each strategy on $\sim 10^5$ trial afterglow observations across 1000 yr of simulated observations. For each trial we record the overall GRB detection fraction as well as the time that the GRB was detected.

All of the results presented in this section (and those in Section 6) were generated for two cases of early-time $t < 500$ sec light curve decay index (as per Section 2.2): a flat light curve, and a rising light curve $F \propto t^{-\alpha}$, $\alpha = -0.2$. We found that the choice of early-time light curve had negligible impact on the detection statistics of a given strategy, and so here we present only the results from our 'flat light curve' simulations.

Figure 3 plots the cumulative probability that the nano-satellite will detect a Swift GRB trigger as a function of time after the trigger. This represents not only how quickly a given strategy can achieve a 5σ observation but also takes into account all stochastic sources of time delay in the observing pipeline: Swift's downlink time, the TeleCommand latency when uploading a re-pointing command to the satellite, the time associated with waiting for LoS access to the target and the satellite's slew time.

The worst performing strategy both in terms of timeliness and overall detection probability is one which takes consecutive 15 min exposures of the target. While one might expect longer exposures to yield more signal and thus a higher probability of detecting the afterglow, given LoS visibility constraints in low Earth orbit a strategy which performs shorter exposures can observe the target for longer on average (i.e., if the observing window is 14 min long, a [5, 5, 5, . . .] strategy can observe the target for a total of 10 min while a [15, 15, 15, . . .] strategy cannot observe the target at all).

The strategy which maximises the probability of an early-time ($t < 10$ min) GRB detection is one which makes consecutive 1 min exposures of the target since it is able to make the best use of the increased early-time source flux. This strategy is also able to achieve a very high ($\sim 70\%$) overall probability of detecting the afterglow. However in terms of overall detection capability it is marginally less effective than strategies that include longer exposures, which are better suited to late-time detections.

The best performing strategies at both early and late times are those which start with a 1 min exposure of the target then transition to increasingly longer exposures (a strategy which is similar to the approach taken by GROND). Comparing specifically the [1, 5, 5, . . .] and [1, 15, 15, . . .] strategies, we find that the increased signal-to-noise afforded by longer exposures mean that the [1, 15, 15, . . .] strategy is able to detect almost as many GRBs overall as the [1, 5, 5, . . .] strategy ($\sim 2\%$ fewer). However, by utilising shorter exposures the [1, 5, 5, . . .] strategy is able to make detections much more quickly, with its cumulative detection probability being consistently ~ 2–5% higher than the [1, 15, 15, . . .] strategy for $t > 8$ min.

It is worth noting that within our simulation, performing additional 1-min exposures at the start of an observing strategy would increase the probability of an intermediate-time detection ($t = 3$–7 min) by $\sim 5\%$ without substantially decreasing the overall detection probability. However, in this work we avoid over-tuning an observing strategy to maximise its effectiveness during this time period due to our simplifications in modelling the early-time GRB light curve (Section 2.2).

Figure 4 shows the variability in yearly follow-up detection fraction for several different observing strategies. We find that this variability is the same regardless of strategy choice (standard deviation $\sigma \sim 3\%$), indicating that the spread comes from the intrinsic variability of the GRB afterglows and/or stochastic visibility constraints in low Earth orbit. Therefore, we can identify the optimal observing strategy purely on the grounds of detection timeliness and mean overall detection fraction.
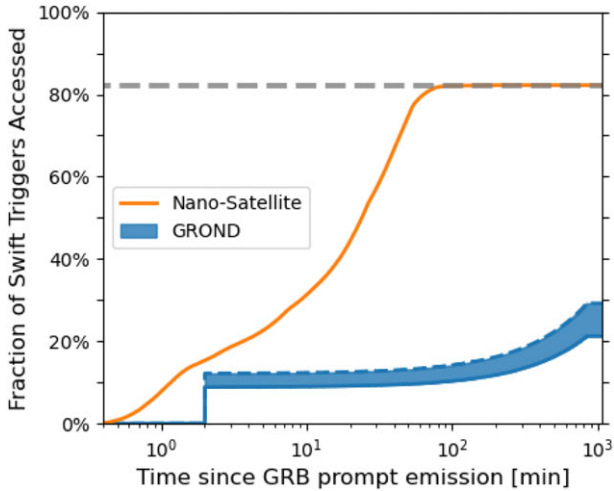
Table 4 shows the total fraction of UVOT-observable Swift GRB triggers successfully detected every year by a near-infrared nano-satellite employing a given strategy. In addition to those strategies listed in the table we made efforts to search for higher performing, more complex strategies such as a [1, 3, 5, 5, ...] strategy, but found that their detection statistics differed negligibly from the [1, 5, 5, ...] strategy. Therefore, we identify [1, 5, 5, . . .] as the preferred GRB observing strategy; not only does it have a high probability of an early-time afterglow detection, but it also achieves the highest probability of detecting the afterglow overall. We find that taking 5 min exposures of the GRB afterglow is more effective than taking 15 min exposures, likely because on average the telescope is able to spend more time taking exposures of the afterglow given the variability in the length of the observing window from low Earth orbit. Additionally, from a mission design

Table 4. Total percentage of successful follow-up detections for each strategy investigated in this work divided by the number of UVOT-observable Swift triggers (see Section 5.1 for a description of the strategy notation).

	Detection probability (%)	
Strategy	All GRBs	$z > 5$
$[1, 1, 1, \ldots]$	68.8 ± 3.3	39.3 ± 11.9
$[1, 5, 5, \ldots]$	72.5 ± 3.1	44.1 ± 12.3
$[1, 15, 15, \ldots]$	70.3 ± 3.2	41.7 ± 12.2
$[15, 15, 15, \ldots]$	64.2 ± 3.3	39.9 ± 12.1

Figure 5. Comparison of the cumulative probability that the instrument begins taking exposures of an arbitrary Swift GRB trigger within a certain time for the GROND instrument and a nano-satellite in a 550 km sun-synchronous orbit. The solid and dashed blue lines represent the worst case ('photometric') and best case ('usable') weather conditions on the ground, and the shaded blue region represents the possible range of GROND's performance. The jump in GROND's access fraction at $t = 120$ s is due to our simplified modelling assumption that GROND takes exactly 2 min to re-position its dome and begin observations on a target (see Section 3.3.2). The grey dashed line indicates the fraction of Swift GRB triggers that are UVOT-observable.

standpoint it is desirable to perform several shorter exposures in place of one long exposure to minimise the risk of a cosmic ray event disrupting an observation. In this work we elect not to further optimise the duration of the repeating exposure, leaving this to a future work which incorporates a more sophisticated model for GRB light curves.

6. Results and discussion

This section outlines the numerical results generated by performing $\sim 10^5$ trial observations of GRB afterglows within our Monte Carlo simulation framework. We present the results of a near-infrared nano-satellite using the $[1, 5, 5, \ldots]$ exposure strategy for observing GRB afterglows (as per Section 5), and for the GROND instrument using the observing strategy defined in Section 3.3.

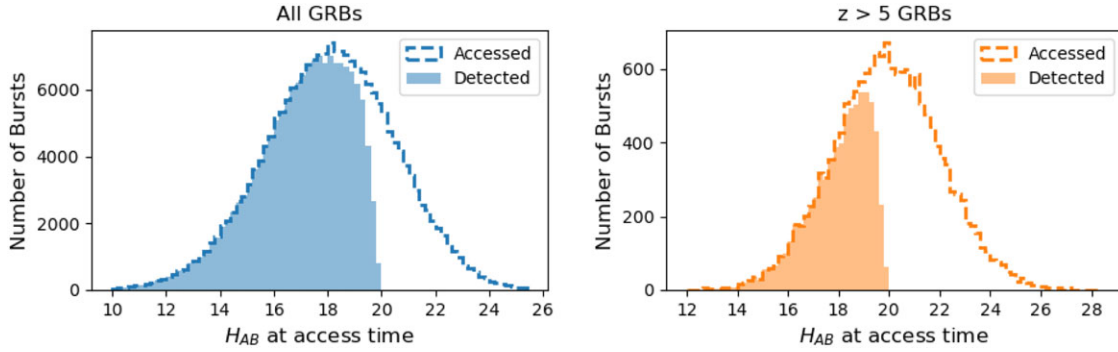### 6.1. Comparing ground and space-based observations

Figure 5 demonstrates the clear advantage that space-based afterglow observatories have over their ground-based counterparts: the nano-satellite can access every UVOT-observable Swift GRB within ~ 1 h of the burst whereas GROND is only able to access

Table 5. Number of events accessed/detected (in the H-band) divided by the number of UVOT-observable Swift GRB triggers for each year (%). The third row represents the percentage of $z > 5$ events detected divided by the total number of UVOT-observable $z > 5$ events which occurred each year. The labels 'phot.' (photometric) and 'usable' denote the weather modelling assumptions used for GROND, which are described in Section 3.3.1.

	Nano-Satellite	$GROND_H$ (phot.)	$GROND_H$ (usable)
Accessed	100.3 ± 0.1	26.1 ± 3.3	35.9 ± 3.8
Detected	72.5 ± 3.1	19.9 ± 2.9	27.3 ± 3.3
$(z > 5)$ Det.	44.1 ± 12.3	13.4 ± 8.9	18.1 ± 10.0

$\sim 30\%$ of such bursts after 14 h (under the assumption of optimal weather conditions).

Within our simulation, the nano-satellite is capable of accessing more than 100% of UVOT-observable Swift triggers (Table 5), meaning that there are GRBs which are inaccessible to UVOT but can be accessed by the nano-satellite. While the nano-satellite and UVOT share the same celestial body avoidance angles, UVOT's sky visibility is limited by its equatorial orbit (which passes between Earth and the Sun for a fraction of its orbit, obscuring two large sections of the sky simultaneously) whereas the sun-synchronous polar orbit of the nano-satellite affords it access to a much larger fraction of the sky throughout its orbit. Conversely a ground-based facility is limited by the weather, the day/night cycle and the limited fraction of sky accessible overhead, resulting in GROND only being able to access a relatively small fraction of Swift GRB triggers.

Despite having a slightly lower H-band point source sensitivity, a nano-satellite with the SkyHopper design requirements is comparable to GROND at detecting GRBs in the H-band given its ability to acquire the targets more rapidly (the nano-satellite detects $\sim 72\%$ of the events that it accesses while GROND detects $\sim 76\%$—Table 5). It is important to note however that this comparison only takes into account H-band detections—GROND's photometry from the bluer g', r', i', z' and J band are more sensitive and are able to increase the likelihood of detecting a GRB afterglow at redshift $z \lesssim 8$, but this modelling aspect is outside the scope of this work. For this reason, these numbers should not be taken to reflect the total number of GRBs that each instrument will observe per year. However, given H-band detections are critical for the identification of GRBs at the highest redshift, this comparison demonstrates the advantage that a near-infrared nano-satellite has in identifying GRBs during the epoch of reionisation compared to a ground-based observatory which is limited by the atmospheric thermal foreground. The advantage of space-based observatories would be even greater for observations in the K-band and at longer wavelengths, although in that case it would be challenging to meet thermal management requirements within a CubeSat form factor.

Figure 5 highlights the fact that UVOT-observable triggers only account for $\sim 80\%$ of the GRBs detected by Swift in our simulation. In reality, UVOT is able to follow up a higher fraction of BAT triggers since Swift generally tries to avoid pointing BAT close to the exclusion zone around the Sun in order to maximise the number of GRBs that it can follow up with its XRT and UVOT. Our simulation under-predicts this fraction since we model BAT as having a full-sky field of view and being capable of detecting bursts arbitrarily close to the Sun.

Figure 6. Histograms demonstrating the *H*-band afterglow magnitude of GRB afterglows at the time it was accessed by the near-infrared nano-satellite. The left plot depicts the full sample of ~10^5 afterglows, while the right plot shows only those events with *z* > 5. The dashed line indicates all afterglows generated in the simulation, and the shading indicates the GRBs which were detected by the near-infrared nano-satellite using the [1, 5, 5, ...] strategy.

Figure 6 compares the intrinsic distribution of *H*-band magnitudes at the time of access by the near-infrared nano-satellite with those events that were detected across the 10^5 trials. We find that the near-infrared nano-satellite is able to reliably detect bursts that have a *H*-band afterglow magnitude $H_{AB} \lesssim 19$ after re-pointing. This is true both for the full sample of afterglows as well as for the high redshift population, indicating that the primary reason for the reduced effectiveness in detecting high redshift bursts is because this population is on average ~2 mag fainter. We find that the near-infrared nano-satellite is able to detect every GRB with $H_{AB} < 17$ and that detection efficiency drop off gradually between $17 < H_{AB} < 20$, which is likely due to only having a short window of time to observe the burst before it is obstructed by the Earth, Sun or Moon.

6.2. High redshift GRB identification

From our simulation we find that a near-infrared nano-satellite using the [1, 5, 5, ...] observing strategy is able to detect 44.1% ± 12.3% of the UVOT-observable GRBs which originate at $z > 5$ (Table 5). To calculate the number of real-world detections this corresponds to, we calibrate the results from our simulation to the actual number of GRBs observed by Swift each year (our simulation over-predicts this number due to simplified modelling of Swift's observing pipeline—see Section 3.2.1). Presuming that Swift observes an average of 76.2 GRBs yr^{-1} with UVOT,[s] and that number of GRBs observed by Swift which have redshift $z > 5$ is approximately $f_{z>5} \sim 5\%$ (Wanderman & Piran 2010; Greiner et al. 2011; Perley et al. 2016) we can use the results from our simulation to compute the expected number of high-z GRBs observed by our nano-satellite as:

$$N_{GRBs}(z > 5) = 76.2 \times f_{z>5} \times (44.1\% \pm 12.3\%)$$
$$= 1.68 \pm 0.47$$

Therefore, we find that our near-infrared nano-satellite is able to detect 1.68 ± 0.47 high redshift GRBs per year to the 1σ confidence level. The uncertainty on this number in any given year is of course slightly higher, as there is Poisson uncertainty on the number of UVOT-observable Swift triggers that occur each year, which is not included in the above calculation.

For completeness we apply this same method to the nano-satellite's overall detection fraction, and calculate that a near-infrared nano-satellite should be able to detect on average ~55 ± 2 GRB afterglows yr^{-1} (excluding again Poisson noise on Swift triggers).

It is worth noting that these figures represent conservative estimates of the number of GRB detections per year by a near-infrared nano-satellite, as they presume that the telescope only performs follow-up on Swift GRB triggers. In reality, there are several other current and future all-sky observatories that could increase the yearly number of GRB triggers available for follow-up. Such instruments include the existing Fermi Gamma-Ray Space Telescope which detects ~240 GRBs per year (von Kienlin et al. 2020), as well as several future instruments including the Franco-Chinese SVOM Mission which is expected to detect ~60–70 bursts per year (Wei et al. 2016), or nano-satellite missions such as the CAMELOT (Werner et al. 2018) and HERMES (Fiore et al. 2020) projects. The only limitation to performing follow-up observations on GRB triggers from different instruments is that the source must be localised precisely enough to be reliably observed using a telescope with a relatively small field of view (SkyHopper's design requirements have a 1.5 deg^2 FoV).

### 6.3. Application of a nano-satellite constellation

As an extension to the results presented above, we test the afterglow follow-up capabilities of a small fleet of identical near-infrared nano-satellites. We presume that each satellite orbits in the same orbital plane as defined in Section 3.1.2, and are spaced uniformly around the orbit.

To simulate observations, a random initial pointing and uplink communications delay is generated independently for each satellite in the constellation, and each satellite begins observations as soon as it can safely slew to the afterglow coordinates. Each satellite uses the [1, 5, 5, . . .] strategy to observe GRBs, and the earliest detection time is recorded between all elements in the constellation.

With multiple satellites observing the same afterglow it becomes possible to stack the exposures from each satellite together to achieve an observation with higher signal-to-noise ratio overall, increasing the probability of detecting the source. We presume that the process of downlinking each of the images taken by each satellite in the constellation and stacking them would take several hours (each satellite must finish its full course of

---

[s]Data taken from the Swift catalog between 2005 and 2019 selecting only those bursts with UVOT observations.

Table 6. Total fraction of UVOT-accessible GRBs detected after stacking the signal between elements in a nano-satellite constellation. The number in brackets indicates the detection fraction for high redshift ($z > 5$) bursts.

	TeleCommand speed	
# Sats	Nominal	Instant
1	72.5% (44.1%)	74.8% (47.2%)
2	79.7% (52.5%)	82.9% (58.7%)
3	82.6% (57.4%)	85.4% (63.1%)
5	86.0% (63.3%)	88.1% (68.2%)
10	89.6% (70.4%)	91.0% (74.6%)
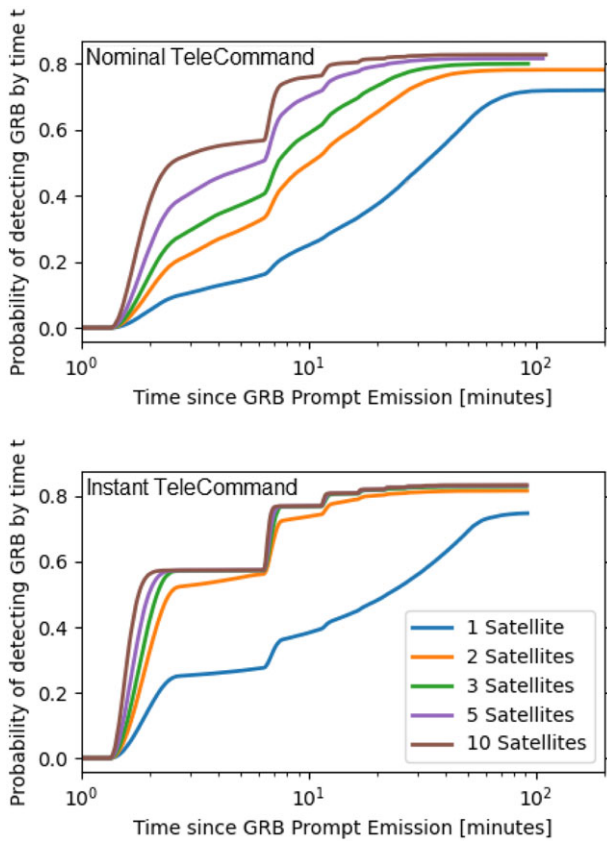
Figure 7. Cumulative probability of detecting a GRB afterglow when using a constellation of rapid-response near-infrared nano-satellites. Top: Satellite TeleCommand is modelled using the nominal telecommand latency of the Iridium network (Section 3.1.3). Bottom: Satellite TeleCommand is modelled as occuring instantaneously (no delay in uplinking a re-pointing command to any satellite in the constellation). The detection time represents the earliest time that a 5σ observation is achieved by any individual satellite in the constellation.

observations and downlink the final image over a limited telecommunication network) and so we do not include its effects in the prompt detection capabilities of the constellation (i.e., we exclude this effect from Figure 7), instead presenting the results of image stacking separately in Table 6. We approximate the signal-to-noise increase of exposure stacking using Equation (10)—in reality this process would be more complex and require inter-satellite calibration measures, but such modelling is beyond the scope of this work.

Figure 7 (upper panel) compares the cumulative probability of detecting a GRB over time for a range of constellation sizes. We find that in the case of nominal TeleCommand delay, launching a second satellite positioned on the other side of the globe doubles the probability of detecting a GRB at early times ($t \sim$ 3–11 min), which is unsurprising given the largest constraint on timely GRB observations in LEO is prompt LoS visibility to the source. However, launching a second satellite does not drastically increase the overall probability of detecting a GRB target, yielding a $\sim$6% increase in total detection fraction even after combining the signal from both two satellites. This behaviour is also true for larger constellations—making the constellation larger has a substantial impact on the early-time GRB detection capability without dramatically increasing the overall probability of detection. This indicates that the primary limitation on the total detection fraction of the constellation not its sky coverage, but rather the point source sensitivity of their instruments. Even an increase of the collecting area for the orbiting satellite would still miss the faint-end tail of the afterglow luminosity distribution.

Increasing the size of the nano-satellite constellation yields diminishing returns with regards to the early-time detection probability, since it only takes three satellites equally spaced around the globe to achieve 100% sky coverage. To illustrate this point, we simulated the behaviour of a satellite constellation that was able to communicate without any uplink latency (Figure 7, lower panel). We find that if one can achieve zero-latency uplink to each satellite, a 3-satellite constellation is able to effectively match the performance of a 10 satellite constellation with only a short $\sim$60 s delay in timeliness due to longer average re-pointing times for the 3 satellite constellation.

The benefit to a larger constellation in the case of zero uplink latency comes from combining the signal from each satellite in the constellation. Combining the signal from 3 satellites allows the constellation to detect a total of $\sim$83% of GRBs after signal stacking, whereas a 10-satellite constellation can detect $\sim$90% of UVOT-observable Swift GRB triggers after stacking (Table 6). The ability to combine signal between satellites becomes more valuable in the identification of high redshift events: a constellation of 3 satellites can detect $\sim$60% of high redshift GRBs, and a constellation of 10 can detect as many as 70% of high redshift bursts after combining the signal from the full constellation.

In terms of cost effectiveness launching additional satellites may not be the best option, as a smaller number of satellites capable of instantaneous uplink capability can match or exceed the performance of larger constellations operating with nominal TeleCommand delay. Before stacking, a constellation of 2 satellites communicating instantly can detect 81.7% of all bursts whereas a constellation of 10 satellites operating with nominal TeleCommand delay can only detect 82.7%. The larger constellations gain the benefits from combining signal between satellites, but even in this respect instantaneous communication is highly beneficial, as 5 satellites communicating instantly can achieve the same total detection fraction as 10 satellites with nominal communications delay (Table 6).

Given the lack of a substantial advantage to larger constellations, we conclude that a group of 2–3 satellites (including a third one helps with redundancy/down-time of instruments) is the largest number of rapid-response near-infrared telescopes one should launch to observe GRB afterglows. Given this conclusion, we generate Figure 8 in order to understand the performance increase gained by launching 3 satellites as compared to a single
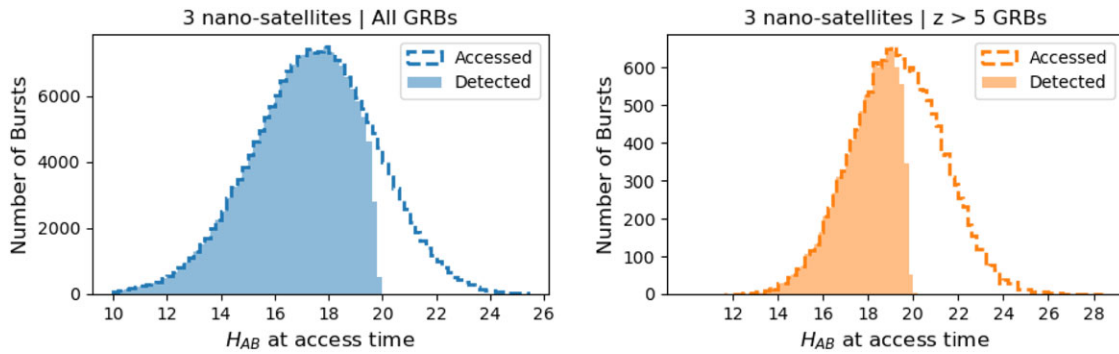
Figure 8. Histograms demonstrating the *H*-band afterglow magnitude of GRB afterglows at the time it was first accessed between a group of 3 near-infrared nano-satellites. The left plot depicts the full sample of $\sim 10^5$ afterglows, while the right plot shows only those events with $z > 5$. The dashed line indicates all afterglows generated in the simulation, and the shading indicates the GRBs which were detected (after stacking) by 3 nano-satellites using the [1, 5, 5, ...] strategy with nominal TeleCommand latency.

satellite. We find that a 3 satellite constellation can reliably detect every burst with a magnitude on access of $H_{AB} < 19$ (two magnitudes deeper than with a single satellite—Figure 6), when stacking the signal between satellites. Simultaneously, the mean magnitude of GRB afterglows at the time of access is $\sim$0.5 mag brighter in the case of 3 satellites compared to a single satellite, allowing the constellation to detect a greater fraction of the GRB population below the limiting magnitude of $H_{AB} = 20$ (the limiting magnitude to detection only increases as high as $H_{AB} \sim 20.6$ for a 10 satellite constellation).

7. Conclusion

We present the results from a simplified mission simulation testing the ability of a rapid-response near-infrared nano-satellite (based on the SkyHopper mission concept) to perform follow-up observations on Swift GRB triggers. The simulation developed for this work combines an orbital line-of-sight modelling framework with Monte Carlo methods to randomly generate GRB events and average an infrared nano-satellite's performance over many simulated observing runs. We find that a near-infrared nano-satellite is able to not only expand the existing catalogue of high redshift GRBs, but also to make a significant and unique contribution to the study of early-time near-infrared GRB afterglows. Using the simulation framework outlined in this work we demonstrate the following:

- A near-infrared nano-satellite (using the optimal observing strategy) in a 550 km polar Sun-synchronous orbit is capable of performing a successful follow-up afterglow detection on 72.5% ± 3.1% of UVOT-observable Swift GRB triggers within 2 h of the initial burst, which corresponds to a GRB detection rate of $\sim$55 ± 2 GRBs yr^{-1}. Furthermore $\sim$30% of Swift GRB triggers can be detected within $t = 10$ min of the GRB prompt emission, corresponding $\sim$20 GRBs yr^{-1} detected in this under-sampled region of GRB afterglow parameter space.

- A near-infrared nano-satellite can detect $\sim$1–3 high redshift ($z > 5$) GRBs yr^{-1}. This predicted performance represents a substantial advance for the field, as at the time of writing only 23 such GRBs have been discovered over the last $\sim$24 yr of observations.

- The optimal fixed observing strategy for early-time near-infrared afterglow observations is to first perform a 1 min exposure of the target before transitioning to performing 5 min

exposures of the target. This strategy maximises the overall probability of detecting a GRB afterglow while also achieving a very high chance of a timely ($<$11 min after GRB trigger) detection of the target. This strategy performs better than a strategy which attempts to take long exposures of the target given the variability in the observing window duration afforded from low Earth orbit.

- Launching a small constellation (2–3) of similar nano-satellites into orbit equally spaced around the same orbital plane is expected to double the early-time probability of detecting a GRB, and to increase the overall probability of detecting a GRB (after combining the signal from each satellite) from $\sim$72% for a single satellite to $\sim$83% for a constellation of 3. Similarly, a 3-satellite constellation is expected to detect of $\sim$58% of high redshift GRBs, compared to $\sim$44% for a single satellite. Launching more than 3 nano-satellites into orbit may not represent the best return on investment compared to reducing TeleCommand latency or increasing the point source sensitivity of the instrument (i.e., increasing telescope aperture). In the case of a zero-latency TeleCommand solution, a constellation of 3 equally spaced satellites can equal the performance of a 10 satellite constellation with respect to timely GRB detections, though 10 satellites yield $\sim$10% more GRB detections overall (i.e., $\sim$91% probability of afterglow detection) after stacking the observations from every satellite.

Overall, we find that using a near-infrared nano-satellite for rapid-response GRB afterglow observations is not only a viable application for GRB science, but would be able to contribute substantially to the field via the identification of high redshift GRBs in a way that is not easily achievable from the ground. While a cost/benefit analysis is outside the scope of this work, the estimated order-of-magnitude cost of a SkyHopper-like nano-satellite mission is in the range of AU \$10 million, making it competitive against the network of multiple 2 m-class facilities on the ground that would be required to yield a comparable efficiency in GRB afterglow imaging follow-up.

Collecting a large database of high redshift GRBs is critical for astrophysical progress. These GRBs are cutting-edge opportunities to characterise the star formation history of the Universe back to the epoch of reionisation (Trenti et al. 2012). They would enable the measurement of the chemical composition of intergalactic gas in the early Universe (Cucchiara, Totani, & Tanvir 2016) and the escape fraction of ionising radiation from galaxies

(Chen, Prochaska, & Gnedin 2007), which is one of the most challenging, yet fundamental, astronomical measurements. GRBs have distinct advantages as cosmological probes over quasars, as the latter carve out large ionised bubbles within their local environments (Chornock et al. 2013), and the future of the field is bright with the James Webb Space Telescope and 30-m class facilities coming online with spectroscopic follow-up capabilities.

References

Berger, E., et al. 2005, APh, 0512280, arXiv:astro–ph/0512280

Chen, H.-W., Prochaska, J. X., & Gnedin, N. Y. 2007, ApJ, 667, L125

Chornock, R., et al. 2013, ApJ, 774, 26

Cucchiara, A., Totani, T., & Tanvir, N. 2016, SSR, 202, 143

Cucchiara, A., et al. 2011, ApJ, 736, 7

Dressel, L. 2021, Wide Field Camera 3 Instrument Handbook, Version 13.0

Evans, P. A., et al. 2009, MNRAS, 397, 1177

Fiore, F., et al. 2020, Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, doi: 10.1117/12.2560680

Fox, A. J., Ledoux, C., Vreeswijk, P. M., Smette, A., & Jaunsen, A. O. 2008, A&A, 491, 189

Gehrels, N., Ramirez–Ruiz, E., & Fox, D. 2009, ARA&A, 47, 567–617

Gehrels, N., et al. 2004, ApJ, 611, 1005

Gehrels, N., et al. 2005, Natur, 437, 851

Gordon, K. D., Clayton, G. C., Misselt, K. A., Landolt, A. U., & Wolff, M. J. 2003, ApJ, 594, 279

Green, J. D., & Olszewski, H. 2020, Space Telescope WFC Instrument Science Report, 3

Greiner, J., et al. 2008, PASP, 120, 405

Greiner, J., et al. 2009, ApJ, 693, 1610

Greiner, J., et al. 2011, AAP, 526, A30

Hartoog, O. E., et al. 2015, A&A, 580, A139

Kann, D. A., Klose, S., & Zeh, A. 2006, ApJ, 641, 993

Kawai, N., et al. 2006, Natur, 440, 184

Klose, S., et al. 2004, A&A, 420, 899

Klose, S., et al. 2000, ApJ, 545, 271

Knapp, M., et al. 2020, AJ, 160, 23

Krühler, T., et al. 2011, A&A, 526, A153

Kumar, P., & Zhang, B. 2015, PhR, 561, 1

Lamb, D. Q., & Reichart, D. E. 2000, ApJ, 536, 1

Lidz, A., Chang, T.-C., Mas-Ribas, L., & Sun, G. 2021, ApJ, 917, 58

McGuire, J. T. W. 2016, in Eighth Huntsville Gamma-Ray Burst Symposium, Vol. 1962, 4021

McQuinn, M., Lidz, A., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2008, MNRAS, 388, 1101

Mearns, R., & Trenti, M. 2018, arXiv e-print, arXiv:1808.06746, arXiv:1808.06746

Mesinger, A., & Furlanetto, S. R. 2008, MNRAS, 385, 1348

Miralda–Escude, J. 1998, ApJ, 501, 15

Oates, S. R., et al. 2009, MNRAS, 395, 490

Oates, S. R., et al. 2015, MNRAS, 453, 4121

Perley, D. A., et al. 2013, ApJ, 778, 128

Perley, D. A., et al. 2016, ApJ, 817, 7

Prochaska, J. X., Dessauges–Zavadsky, M., Ramirez–Ruiz, E., & Chen, H. 2008, ApJ, 685, 344

Racusin, J., et al. 2017, arXiv e-prints, arXiv:1708.09292

Rhodes, L., et al. 2020, MNRAS, 496, 3326

Robertson, B. E., & Ellis, R. S. 2012, ApJ, 744, 95

Rueda, J., et al. 2018, JCAP, 2018, 006

Salvaterra, R., et al. 2009, Nature, 461, 1258

Sari, R., Piran, T., & Narayan, R. 1998, ApJ, 497, L17

Schady, P. 2017, arXiv e-prints, arXiv:1707.05214

Skrutskie, M. F., et al. 2006, AJ, 131, 1163

Stanway, E. R., Bremer, M. N., & Lehnert, M. D. 2008, MNRAS, 385, 493

Tagliaferri, G., et al. 2005, A&A, 443, L1

Tanvir, N. R., et al. 2009, Natur, 461, 1254

Tanvir, N. R., et al. 2012, ApJ, 754, 46

Trenti, M., Perna, R., & Jimenez, R. 2015, ApJ, 802, 103

Trenti, M., Perna, R., Levesque, E. M., Shull, J. M., & Stocke, J. T. 2012, ApJ, 749, L38

Trenti, M., Perna, R., & Tacchella, S. 2013, ApJ, 773, L22

Troja, E. 2020, The Neil Gehrels Swift Observatory Technical Handbook Version 17.0

van Eerten, H., Zhang, W., & MacFadyen, A. 2010, ApJ, 722, 235

von Kienlin, A., et al. 2020, ApJ, 893, 46

Vreeswijk, P. M., et al. 2007, A&A, 468, 83

Wanderman, D., & Piran, T. 2010, MNRAS, 406, 1944

Wang, F., et al. 2021a, ApJ, 907, L1

Wang, X.-G., et al. 2013, ApJ, 774, 132

Wang, X. I., et al. 2021b, ApJ, 922, 237

Wei, J., et al. 2016, arXiv e-prints, arXiv:1610.06892

Wen, J., et al. 2019, ExA, 48, 77

Werner, N., et al. 2018, in Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray, Vol. 10699, International Society for Optics and Photonics (SPIE), ed. J.-W. A. den Herder, S. Nikzad, & K. Nakazawa, 672

Woosley, S. E., & Bloom, J. S. 2006, ARA&A, 44, 507