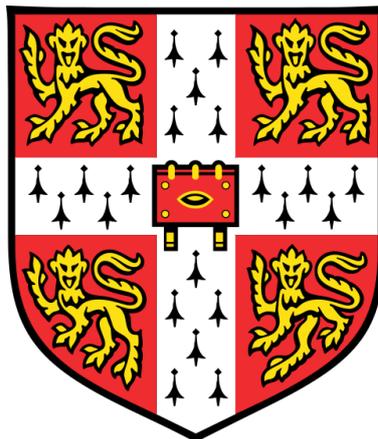


**Optimised ground-based near-infrared instrumentation
for robotic exoplanet transit surveys**

Peter Pihlmann Pedersen

Corpus Christi College



Supervisor: Prof. Didier Queloz

Assessors: Prof. Daniel Bayliss, Dr. John Young

2023 January

This thesis is submitted for the degree of Doctor of Philosophy

DECLARATION

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration, except as declared below and specified in the text and acknowledgements. It is not substantially the same as any that I have submitted, or is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution, except as declared below and specified in the text and acknowledgements. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution, except as declared below and specified in the text and acknowledgements. This dissertation does not exceed the prescribed word limit for the Degree Committee of Physics and Chemistry. It contains fewer than 60,000 words including abstract, tables, footnotes and appendices, and has fewer than 150 figures.

The introduction chapter to SPECULOOS, Chapter 2, was published as ‘Development of the SPECULOOS exoplanet search project’ in *SPIE, Ground-based and Airborne Telescopes VIII* (Sebastian, Pedersen, Murray, Ducrot, Garcia, Burdanov, Pozuelos, Delrez, Wells, Dransfield et al. 2020b), where Daniel Sebastian, I, and the SPECULOOS consortium contributed to the ‘Introduction’ and ‘Network of ground-based observatories’, Daniel Sebastian contributed to the ‘Target selection’, Elsa Ducrot to ‘Scheduling’, Catriona Murray and I to ‘Data reduction and quality’, Lionel Garcia and I to ‘SPECULOOS Portal’.

Chapter 3 has been published as ‘Precise near-infrared photometry, accounting for precipitable water vapour at SPECULOOS Southern Observatory’ in *Monthly Notices of the Royal Astronomical Society* (Pedersen et al. 2022).

Throughout the thesis, parts were submitted for my First Year report titled ‘Feasibility of near-infrared instrumentation for ground-based exoplanet transit survey SPECULOOS’, particularly from Chapters 1 and 4.

SUMMARY

Optimised ground-based near-infrared instrumentation for robotic exoplanet transit surveys
by Peter Pihlmann Pedersen

In this PhD, I have advanced the study of ground-based photometric observations in the near-infrared. Specifically, I have worked in the context of robotic exoplanet transit surveys in collaboration with the SPECULOOS Southern Observatory (SSO). Here, I targeted optimising the photometric precision of observing late M, L type stars by developing correction methods and a new instrument called SPIRIT.

My first original contribution was the development of a correction method for the induced effects from varying precipitable water vapour (PWV) in our atmosphere, specifically on differentially resolved light curves from SSO. This work succeeded in reducing false variability of time-series data from late M, L type stars on both long and short timescales, to the extent of removing false transit features.

In parallel, I performed a feasibility study of introducing new near-infrared instrumentation to SSO. One which would permit better photometric precision than the existing CCD Si based instrumentation, and likewise minimise the induced effects of PWV variability. An InGaAs based instrument, sensitive up to 1.62 μm was identified, and a custom wide-pass filter called zYJ was designed and manufactured to form SPeculoos' Infra-Red photometric Imager for Transits (SPIRIT). It proved to be a significantly lower-cost alternative to the traditionally used HgCdTe based instrumentation, as well as being better suited to robotic observatories.

On sky results of SPIRIT at SSO successfully demonstrated better photometric precision for stars below 2550 K than the existing instrumentation. It similarly demonstrated the benefit of seeing further into the infrared for minimising the observed variability of M, L type stars. Finally, the custom designed wide-pass filter, zYJ, successfully demonstrated a significantly lower sensitivity to PWV variability.

These results pave a new avenue for ground-based near-infrared robotic exoplanet transit surveys, as well as similar time-series focused astronomy. I conclude my work by suggesting viable routes to further improve the photometric precision of such new instrumentation.

ACRONYMS

| | |
|---------|---|
| 2MASS | Two Micron All-Sky Survey |
| ACP | Astronomy Control Program |
| ADC | Analogue to Digital Converter |
| ADU | Analogue to Digital Units |
| ALC | Artificial Light Curve |
| API | Application Programming Interface |
| ASCOM | AStronomy Common Object Model |
| CCD | Charged-Coupled Device |
| CDF | Cumulative Distribution Function |
| CDS | Correlated Double Sampling |
| CMOS | Complementary Metal-Oxide-Semiconductor |
| COTS | Commercially-Off-The-Shelf |
| ELT | Extremely Large Telescope |
| EMI | Electro-Magnetic Interference |
| ESO | European Southern Observatory |
| FPA | Focal Plane Array |
| FWHM | Full-Width-Half-Maximum |
| GPS | Global Positioning System |
| HPP | Half-Power Point |
| IAU | International Astronomical Union |
| JD | Julian Date |
| JWST | James Webb Space Telescope |
| LAMP | Linux, Apache, MySQL, PHP/Python |
| LC | Light Curve |
| LHATPRO | Low Humidity And Temperature PROfiling microwave radiometer |
| NGTS | Next Generation Transit Survey |
| NIR | Near Infra-Red |
| OTA | Optical Tube Assembly |
| PIRT | Princeton Infrared Technologies |
| PORTAL | Pipeline Output inteRacTion Analysis Layer |
| ppt | Parts Per Thousand |
| PSF | Point Spread Function |
| PTC | Photon Transfer Curve |
| PWV | Precipitable Water Vapour |

| | |
|-----------|---|
| QE | Quantum Efficiency |
| REST | REpresentational State Transfer |
| RMS | Root Mean Squared |
| ROIC | Read Out Integrated Circuit |
| SAINT-EX | Search And characterIsationN of Transiting EXoplanets |
| SDK | Software Development Kit |
| SNAP | Super Nova Acceleration Probe |
| SNO | SPECULOOS-Northern Observatory |
| SNR | Signal-to-Noise-Ratio |
| SOC | Specific Operating Conditions |
| SPECULOOS | Search for habitable Planets EClipsing ULtra-cOOl Stars |
| SPIRIT | SPeculoos' Infra-Red photometric Imager for Transits |
| SPOCK | SPeculoos Observatory sChedule maKer |
| SQL | Standard Query Language |
| SSO | SPECULOOS-Southern Observatory |
| TESS | Transiting Exoplanet Survey Satellite |
| TOI | TESS Object of Interest |
| TRAPPIST | Transiting Planets and Planetesimals Small Telescope |
| UCD | Ultra Cool Dwarf |
| UKIRT | United Kingdom Infra-Red Telescope |
| VISTA | Visible and Infrared Survey Telescope for Astronomy |
| VLT | Very Large Telescope |

CONTENTS

| | |
|---|-------------|
| Declaration | iii |
| Summary | v |
| Acronyms | vii |
| Contents | ix |
| Acknowledgements | xiii |
| 1 Introduction | 1 |
| 1.1 Introduction to exoplanets and the potential selection bias | 1 |
| 1.1.1 The transit method | 3 |
| 1.2 Infrared detectors in Astronomy | 5 |
| 1.2.1 The visible domain | 6 |
| 1.2.2 The infrared domain | 7 |
| 1.3 Our atmosphere’s influence on observations | 10 |
| 1.3.1 Differential photometry | 12 |
| 2 SPECULOOS | 15 |
| 2.1 Introduction | 16 |
| 2.2 Network of ground-based observatories | 18 |
| 2.3 Targets and Observations | 21 |
| 2.3.1 Target selection | 21 |
| 2.3.2 Scheduling | 22 |
| 2.4 Data reduction and quality | 24 |
| 2.4.1 Data Reduction Pipeline | 24 |
| 2.4.2 Photometric quality | 25 |
| 2.5 SPECULOOS Portal | 26 |
| 2.5.1 Backend | 26 |
| 2.5.2 Frontend | 27 |
| 3 Precise NIR photometry, accounting for PWV at SSO | 31 |
| 3.1 Introduction | 32 |
| 3.2 Method of correction | 34 |
| 3.2.1 The PWV grid | 36 |
| 3.2.2 PWV measurements – estimating line-of-sight PWV | 36 |
| 3.2.3 Accounting for altitude difference | 37 |

| | | |
|----------|--|------------|
| 3.3 | Results and discussion | 38 |
| 3.3.1 | PWV variability at Paranal | 38 |
| 3.3.2 | Evaluation of the PWV correction | 42 |
| 3.3.3 | PWV correction in action | 44 |
| 3.4 | Conclusions | 48 |
| 4 | Feasibility study | 51 |
| 4.1 | Fundamentals | 51 |
| 4.1.1 | Photometric precision | 53 |
| 4.1.2 | Available NIR detectors and camera systems | 55 |
| 4.1.3 | Photometric precision model verification | 56 |
| 4.2 | Filter optimisation | 57 |
| 4.3 | Precision grid models | 61 |
| 4.3.1 | Precision grid results | 63 |
| 4.3.2 | Calibrating to the SPECULOOS target list | 64 |
| 4.4 | SPECULOOS target list precision | 65 |
| 4.4.1 | Sensitivity analysis | 68 |
| 4.5 | Field of view suitability | 77 |
| 4.5.1 | Upper magnitude limit | 77 |
| 4.5.2 | Lower detectable magnitude | 78 |
| 4.5.3 | Evaluating SPECULOOS' field of views | 78 |
| 4.6 | Conclusions | 81 |
| 5 | Camera characterisation | 83 |
| 5.1 | Camera operation | 84 |
| 5.2 | Manufacturer's report | 86 |
| 5.3 | Parameter verification in Cambridge | 88 |
| 5.3.1 | Read noise | 89 |
| 5.3.2 | Dark current | 92 |
| 5.3.3 | Point source tests | 93 |
| 5.4 | Gain | 95 |
| 5.5 | Linearity | 99 |
| 5.6 | Usable full well | 100 |
| 5.7 | Bad pixels | 101 |
| 5.8 | Final characterised specification | 103 |
| 6 | First light | 105 |
| 6.1 | Preparation and integration of SPIRIT at SSO | 106 |
| 6.1.1 | Instrument mounting mechanism | 106 |
| 6.1.2 | Cooling system | 106 |
| 6.1.3 | Software | 109 |
| 6.1.4 | Filters | 113 |
| 6.1.5 | Setbacks | 115 |
| 6.1.6 | Final installed system | 116 |
| 6.2 | Early results and observations | 117 |
| 6.2.1 | Quantifying the true dark current | 118 |
| 6.2.2 | Cool star simultaneous observation | 123 |

| | | |
|----------|---|------------|
| 6.2.3 | Noise models | 127 |
| 6.2.4 | PWV insensitivity | 132 |
| 6.3 | Example exoplanet photometry | 132 |
| 6.3.1 | TOI-3884.01 | 134 |
| 6.3.2 | Trappist-1b | 137 |
| 6.4 | Conclusions and future improvements | 138 |
| A | 1280SciCam procurement document | 141 |
| B | Filter procurement document | 143 |
| C | Example image and light curves from SPIRIT | 147 |
| | References | 151 |

ACKNOWLEDGEMENTS

An immense gratitude and thank you goes towards my supervisor, Didier Queloz, for first, accepting me as his student, second, enabling this incredible project and the avenues it has taken me, and third, for providing great mentorship and setting a great example of how to approach work and life.

I would also like to thank members of the SPECULOOS consortium for providing the necessary support to enabling the successes in this PhD. Likewise, an extended thank you to past and present members of the Exoplanets group, mechanical and astrophysics workshops, and Sensor CDT in Cambridge, for the various discussions and aid that have been necessary in order to complete this work. Similarly, a huge appreciation to the ESO engineers at Paranal who have helped SPECULOOS and I in times of need.

Lastly, an immeasurable thank you goes towards my friends and family for supporting me along this journey and being patient with my sometimes crazy ideas.

INTRODUCTION

Understanding the origin of life is an endeavour which humanity may never answer. Be it by design or due to the unfortunate inaccessibility of information. Nonetheless, we pursue the realms of science and philosophy in the hope for answers. One promising domain is the field of exoplanets. Here, the discovery and characterisation of planets beyond our own solar system contributes to a population analysis which may reveal answers to our origins. This study however is often bounded by instrumentation sensitivities, atmospheric and space-science limitations.

This chapter is dedicated to a description of exoplanet science and topics related to their discoveries and this PhD.

1.1 Introduction to exoplanets and the potential selection bias

The working definition of an exoplanet, as set by the IAU's Commission F2: exoplanets and the Solar System, in August 2018^a, is as follows:

“Objects with true masses below the limiting mass for thermonuclear fusion of deuterium (currently calculated to be 13 Jupiter masses for objects of solar metallicity) that orbit stars, brown dwarfs or stellar remnants and that have a mass ratio with the central object below the L4/L5 instability ($M/M_{central} < 2/(25 + \sqrt{621}) \sim 1/25$) are ‘planets’ (no matter how they formed).”

^aDate accessed 2022-02-20: https://www.iau.org/science/scientific_bodies/commissions/F2/info/documents/

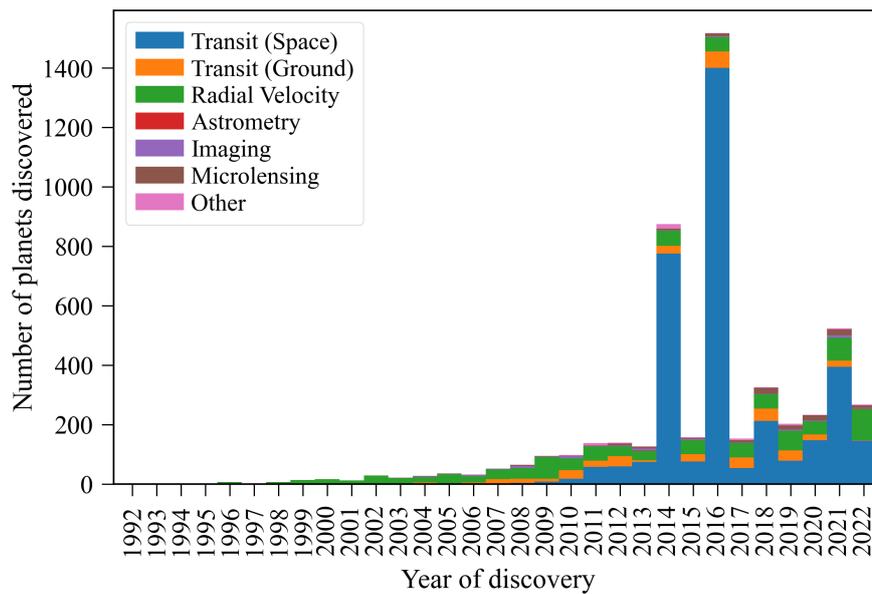


Figure 1.1: The number of exoplanets confirmed as a function of time, distinguished between their discovery method, with the transit method distinguished by locale. Data from the NASA Exoplanet Archive, as of 2022-09-30.

Exoplanets have been postulated by many over the past centuries, but it took the discovery of Michel Mayor and Didier Queloz (Mayor & Queloz 1995) to convince the world of their existence. Their discovery was of a hot Jupiter sized planet orbiting closely to a Sun-like star, 51-Pegasi, with an orbital period of 4 days. Since this discovery 27 years ago, 5,171 exoplanets have been confirmed in 3,870 separate star systems, with 861 systems having been found with more than one planet.^b

Over the years, a multitude of methods to detect exoplanets have been developed, refined, and applied to a range of ground and space based observatories. Namely, methods based on measuring the radial velocity and photometry of host stars have been the most successful to date, in both detecting and characterising exoplanets (Charbonneau et al. 2000). Other methods, such as direct imaging, astrometry, and micro-lensing, have been less prominent due to the inherent increased difficulty/limitations in their methods (Traub et al. 2010; Perryman et al. 2014; Gaudi 2012). In Figure 1.1, I show the proportion of confirmed detections over time, divided by their initial detection method. The dominant method to date, by far, has been the photometric transit method, observed from space.

The Kepler Space Telescope (Koch et al. 2010) made the majority of these discoveries (Burke et al. 2014), who has now been retired. A successor to Kepler, the Transiting Exoplanet Survey Satellite (TESS) (Ricker et al. 2014) has had similar success in discovering numerous

^bAs of 2022-09-30, <https://exoplanetarchive.ipac.caltech.edu/>

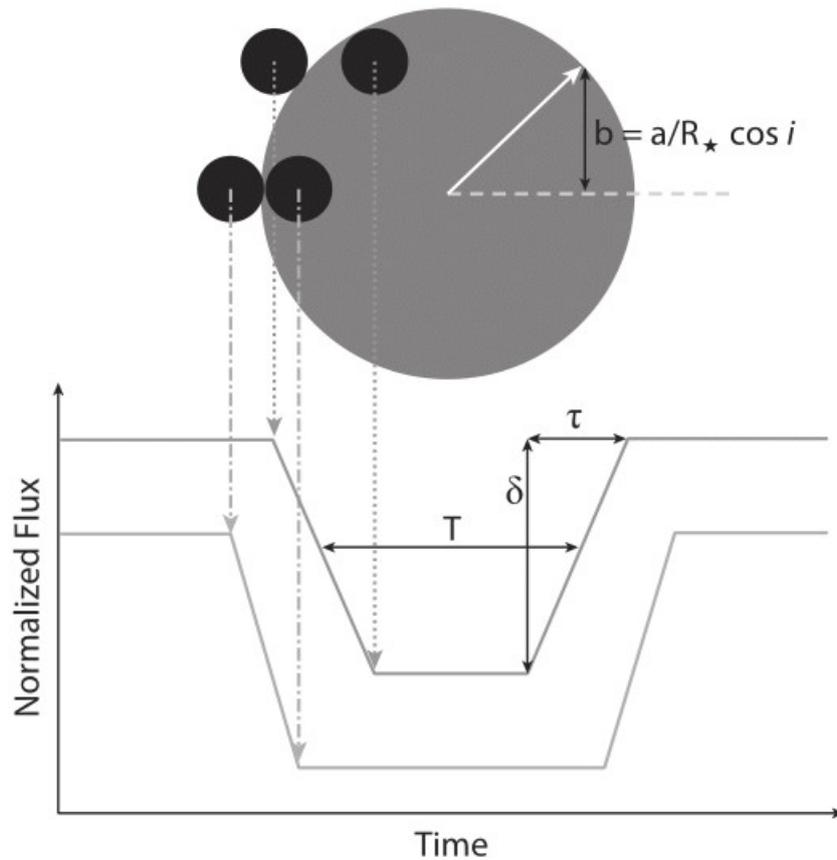


Figure 1.2: On the top, a simplified diagram illustrating a transiting planet across the disc of its host star, with two example transit scenarios for different impact parameters, b . On the bottom, the effect on the observed flux as a function of time, with a transit depth, δ , transit duration, T , and ingress/egress duration, τ . Credit: John Asher Johnson.

exoplanets (Guerrero et al. 2021). Whilst ground-based observatories often have similar or better instrumentation than space-borne observatories, we are limited by our atmosphere inducing false variability and noise into the data. It is therefore important to probe these limitations faced by ground-based observatories to increase the detection probability of current and future transit surveys.

1.1.1 The transit method

If a planet's orbital path, radius, and its relative position to the observer permit light from the planet's host light to be obscured, then a transit can be observed. During a transit, the relative change of light and the duration it is observed for is dictated by a handful of geometric properties of the planetary system. Figure 1.2 attempts to illustrate them. Assuming the planet's disc is

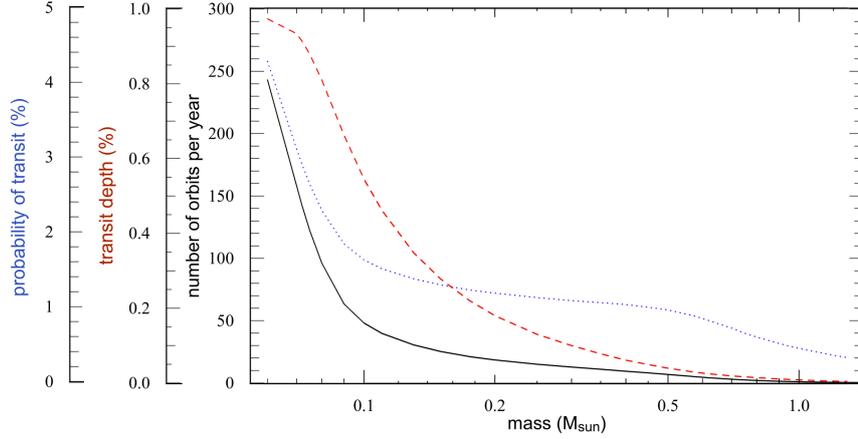


Figure 1.3: For an Earth-like planet orbiting within the habitable zone, the number of orbits per year (black), the transit depth (dashed red), and the probability of transit occurring (dotted blue) as a function of the host’s mass. Figure adapted from [He et al. \(2017\)](#).

completely within the host star’s disc from the observer’s perspective, then the fractional change of flux, the transit depth, δ , between the beginning of the transit to the middle of the transit can be approximated by the ratio of the planet’s to star’s disc area:

$$\delta \approx \left(\frac{R_p}{R_\star} \right)^2, \quad (1.1)$$

where R_p and R_\star are the radii of the planet and star respectively. The approximation in Equation 1.1 is due to the fact that the transit depth will also be a function of other variables. Such as variations in the star’s brightness across the disc (for example due to faculae, spots, and flares) and other objects in the orbital plane will affect the transit depth over the transit’s duration. Naturally, this means that in systems with active stars for example, it may become difficult to distinguish when a transit has occurred in reality. However, if components of its variation are stably periodic and at periods not synchronised with the planet’s orbital period, one can model them out to better reveal the transit events.

For an Earth-sized planet transiting a star, the amplitude of the transit depth is expected to increase for lower mass stars, given the smaller stellar radii ([He et al. 2017](#); [Mann et al. 2015](#)). For these systems, there’s also an expected increased geometric probability of observing a planet orbit in the supposed habitable zone. Lower mass stars in the main sequence have cooler effective temperatures, and thus the equilibrium temperature for the habitable zone is met at much closer orbits to its host star. Figure 1.3 quantifies the expected transit depths and the associated probabilities for an Earth-sized planet.

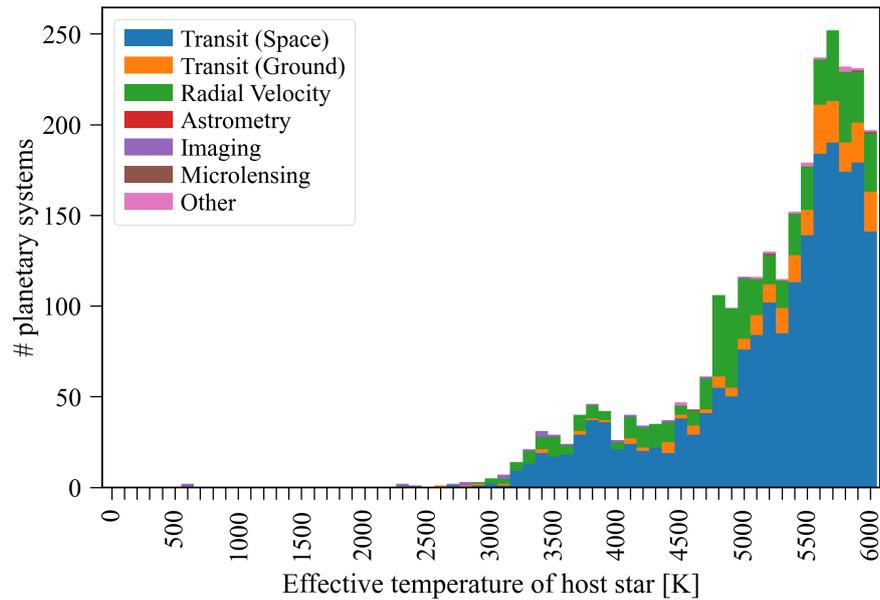


Figure 1.4: Histogram of the number of planetary systems discovered to date as a function of the effective temperature of their host star (cut off at the temperature of our Sun). Distinguished between their discovery method, with the transit method distinguished by locale. Data from the NASA Exoplanet Archive, as of 2022-09-30.

Whilst these higher probabilities and larger transit depths at lower masses seem promising. Stars at the lower masses are also considerably dimmer and have notably less flux in the regions of the spectrum where common detectors are sensitive to (discussed in the following sub-section). Stellar spectra shift to redder wavelength the cooler the effective temperature. Consequently, this leads to more photon derived noise when observing lower mass stars, which in turn leads to lower detection rates. This is evident when we look at the population of detected planetary systems discovered to date, as a function of their host star's effective temperature, see Figure 1.4. The distribution diminishes at lower temperatures, and as suggested, this is possibly a selection bias due to the lower sensitivities of past surveys in this low mass/low effective temperature regime.

1.2 Infrared detectors in Astronomy

The infrared is broadly defined between the wavelengths of 750 nm to 0.3 mm, and can be subdivided into multiple regions. The division of these regions are dependent on which field you ask, in astronomy and for the context of this thesis, we will define the Near-Infrared (NIR) as between 0.9 to 2.5 μm , as has been done in Heidt (2022). In astronomy, there are numerous motives to observe in the infrared over the visible (380 nm to 750 nm). For example, one

can identify more chemical species with the increased number of absorption/emission lines present in the infrared. Infrared light is also less attenuated by interstellar dust, permitting one to probe further into certain regions. Similarly, for an equally wide bandpass in the infrared as the visible domain, one can record more flux emitted by some cooler objects due to their peak radiation shifting into the infrared. Finally, specifically in the exoplanet domain, one can see sharper defined transits (Claret & Bloemen 2011). However, whilst these benefits are highly sought after, to observe in the infrared is traditionally very costly and comes with a handful of disadvantages that are significantly more prominent than observed in the visible domain.

1.2.1 The visible domain

Today, most astronomical detectors used at professional observatories are silicon based Charged-Coupled Devices (CCDs). A detector type invented by members at Bell labs over 50 years ago (Boyle & Smith 1970), which has revolutionised the field and indeed the world that we know today. When cooled appropriately (~ -100 °C), they present exceptionally low detector noise characteristics. In addition to the low noise, other detector properties such as quantum efficiency, linearity, persistence, amongst others, have been equally well optimised.

In CCDs, however, to achieve low read noises readout mechanism can take in the order of a few seconds to complete. The mechanism requires all the accumulated charge from each pixel to be converted to a voltage through a few output nodes on the detector, buffered, and then converted into a digital signal via an analogue-to-digital converter (ADC). An alternative detector type, the Complementary Metal-Oxide-Semiconductor (CMOS), attempts to alleviate the charge-to-voltage bottleneck. Here, each pixel instead has its own set of electronics to permit charge-to-voltage conversion, which enables significantly faster readouts. One advantage of this faster readout mechanism is the removed need for a mechanical shutter, which CCDs often require. However, with embedded electronics within each pixel, CMOS based detectors have a significantly increased design complexity. Thus, CMOS technology, whilst conceived at a similar time as the CCD, the required lithography technology to produce them had not been advanced enough to compete with CCDs in astronomy, until the past few years. The rise of smartphones, and other low cost imaging devices, the research and development into the CMOS detector has accelerated in the last decades to now meet and exceed the noise performances of the CCD (Karpov et al. 2020), while displaying advantages in readout speeds, lack of mechanical shutters, and up to ~ 100 x less power consumption than the equivalent CCD (~ 50 mW:5 W).

The majority of the developments into detectors has been made with silicon, which is capable of absorbing photons from the ultraviolet into the early NIR. Silicon can be tuned to best absorb more or less of certain regions within that range depending on the application, by increasing the substrate thickness or doping of the material to better absorb lower energy photons for example.

However, they are fundamentally limited by silicon's 1.1 eV band gap, limiting the absorption of photons beyond 1100 nm. CCD's are comparatively easier to modify for increased NIR absorption, as suggested by manufacturer's of CCDs.^c Modifying a detector's properties are much easier to handle in a CCD's analogue and digital circuits than in the equivalent CMOS circuitry. To push further into the NIR, however, one must consider materials that go below silicon's band gap, such as InSb, HgCdTe, InGaAs, and colloidal quantum dots.

1.2.2 The infrared domain

Military and remote sensing applications have primarily driven the research and development of NIR and beyond infrared detectors. Unsurprisingly, this line of development has seen the CMOS readout architecture adopted, primarily due to its faster and lower power readout modes when compared to CCDs. Due to the specialised applications of these detectors however, there is yet to be one packaged commercially-off-the-shelf (COTS) suitable for the exoplanet science domain, whereas there are many for the visible-NIR wavelengths.

In astronomy, NIR detectors have dominantly been made with HgCdTe (Beletic et al. 2008), a material photosensitive from the visible to wavelengths beyond the NIR, where its sensitivity can be tuned based on the relative compositions of each element, with the common tuned cut-offs of 1.75 μm , 2.5 μm , and 5.3 μm . To form a detector, this material is then hybridised onto a silicon CMOS readout integrated circuit (ROIC). A similar hybridisation process can be done with InGaAs (Pearsall & Hopson 1978), with commonly seen cut-offs of 1.7 μm and 2.2 μm (Ettenberg et al. 2018). These tunable cut-offs are dependent on the stability of respective elements in the composition, often reflected by the respective dark currents produced. The cut-offs and the overall quantum efficiency of the detectors are also seen to shift to lower wavelengths as a function of lowering temperatures (Seshadri et al. 2006; Krishnamurthy et al. 2019).

To produce a camera, the detector is often packaged in a vacuum environment and cooled to lower the noise induced by the dark current. Unfortunately, present infrared detector mediums are plagued with dark current, and thus require extensive cooling to temperatures lower than normally seen with silicon based detectors. This usually results in the use of liquid based coolant, such as liquid nitrogen, which makes for heavier maintenance than air based cooling. As such, there is a low frequency of infrared detectors in astronomy due to the inherent difficulty and maintenance costs in integrating them into a moving telescope. In turn, there is also a higher cost associated with the lower volume production of infrared detectors (when compared to the visible domain). When purchased, they are then integrated into bespoke enclosures and software

^cDate accessed 2022-02-23: <https://www.teledynedalsa.com/en/learn/knowledge-center/ccd-vs-cmos/>

environments and sometimes with multiple detector configurations, which further drives the cost and time in producing the infrared instrumentation.

Nonetheless, multiple well known observatories have HgCdTe based instrumentation, such as Hubble, JWST, UKIRT, and VISTA. However, HgCdTe's cost and maintenance requirements are not suitable for smaller and robotic ground based facilities. InGaAs on the other hand, a cheaper alternative, has been used in many non-astronomy detectors over the past couple of decades, where InGaAs has presented noise characteristics comparable to the traditionally utilised HgCdTe material. The feasibility of using InGaAs for astrophysical instrumentation was first investigated with US military funded work in [Hoffman et al. \(2005\)](#), where they tested two detectors. Both larger than $1k \times 1k$ pixels with $20 \mu\text{m}$ pixel pitch, and high quantum efficiencies between 0.9 to $1.7 \mu\text{m}$, with one of the detectors tailored for astrophysical purposes. This detector was operated at 140 K , and achieved dark currents of $5 \text{ e}^-/\text{s}/\text{pixel}$, and a readout noise of 15 e^-_{rms} with Correlated Double Sampling (CDS, [Jun et al. \(2004\)](#)) or 5 e^-_{rms} using Fowler-8 sampling ([Fowler & Gatley 1990](#)). Other important detector parameters, such as linearity, persistence, inter-pixel cross-talk, and quantum efficiencies were not recorded accurately or not published. Nonetheless, from the recorded parameters, the performances were comparable to a detector made with HgCdTe from around the same time (e.g. [Chuh et al. \(2006\)](#)).

The following year, [Seshadri et al. \(2006\)](#) studied the performances of InGaAs detectors for the Super Nova Acceleration Probe (SNAP) dark energy space mission concept. Similar low temperature results at 140 K on a $1k \times 1k$ InGaAs detector led to the same conclusion that InGaAs based detectors were comparable to HgCdTe arrays of the time. However, they state that the slightly better performance provided by their own $1.7 \mu\text{m}$ cut-off HgCdTe based detectors, led to their decision to use them as the baseline technology for SNAP. Nonetheless, they stressed that their findings still suggested that InGaAs based detectors are a suitable alternative for future astronomy applications.

When [Beletic et al. \(2008\)](#) was published by the sensor manufacturer Teledyne, work on InGaAs for astronomy mostly stagnated. The paper demonstrated the vast capabilities of Teledyne with respect to HgCdTe based detector production and ROICs. It detailed the performance gains with their CdZnTe substrate removal technology ([Piquette et al. 2008](#)), where CdZnTe is necessary for the production process of HgCdTe based detectors. Substrate removal extended the QE to the visible region, as well as increasing the overall QE, without affecting the noise characteristics of the detector. It also removed the effects of cosmic ray events that produce clouds of detected signal due to particle-induced flashes of light in the CdZnTe substrate. InP substrate used during InGaAs detector production was also removed by [Martin et al. \(2005\)](#), and demonstrated similar performance gains, however not as much as with

the CdZnTe substrate removed HgCdTe based detectors.

Since [Beletic et al. \(2008\)](#), work on InGaAs detectors has primarily shifted towards the production of warm (room/ambient temperature) cameras, mainly for military use. InGaAs exhibits lower dark currents than HgCdTe when operated at room temperatures ([Sullivan et al. 2014](#)). As a result, many companies have sprung up in the past decade to manufacture warm cameras. This has driven the cost of InGaAs cameras down to a fraction of the cost when compared to HgCdTe cameras. In recent years, this has motivated those working on astronomy instrumentation, such as in [Sullivan et al. \(2013\)](#), to build a low-cost InGaAs camera for astronomy. The work in [Sullivan et al. \(2013, 2014\)](#) led to on-sky tests performed in 2016, detailed in [Simcoe et al. \(2019\)](#), where they operated a camera with a 640×512 pixel InGaAs based detector with the InP substrate still attached, on a 2.5 m telescope. They achieved a dark current of $113 \text{ e}^-/\text{s}/\text{pixel}$ at a temperature of $-40 \text{ }^\circ\text{C}$, with a readout noise of 43 e_{rms}^- using up-the-ramp sampling ([Glendinning et al. 1990](#)) using a custom read-out board, maintained by a simple thermoelectric cooler with water cooling. A comparable HgCdTe based detector would have dark currents of $\sim 100,000 \text{ e}^-/\text{s}/\text{pixel}$ at $-40 \text{ }^\circ\text{C}$. The camera's optics, for re-imaging the focal plane to a plate scale of $0.4''/\text{pixel}$, were kept at ambient temperature with no cold stop, which is usually required by NIR instrumentation as thermal emissions can become a concern. However, as the InGaAs cut-off wavelength is $\sim 1.6 \text{ }\mu\text{m}$ at cooler temperatures, a cold optical set-up was not deemed necessary for their application. Measurements of the night sky showed that the sky flux in the Y band is comparable to the dark current. In J, the sky exceeds their dark current by a factor of four. The read noise falls below sky + dark noise for exposures of $t > 7 \text{ s}$ in Y and 3.5 s in J. They were able to perform time-series photometry in the J band by observing a known exoplanet with a transit depth of 13.4 mmag (1.2%), with a J magnitude of 12.5 (5380 K) host star. [Strausbaugh et al. \(2018\)](#) similarly observed a known transit ($\sim 2\%$ depth), except with a cooled off-the-shelf InGaAs camera. Their precision was sufficient to resolve each transit, suggesting once again that InGaAs could be suitable for astronomy.

To date, no one has actively developed a commercially available, astronomy grade, camera using an InGaAs detector that has characteristics comparable to a substrate removed HgCdTe based camera, nor at the resolutions that HgCdTe based detectors such as Teledyne are capable of delivering ($4\text{k} \times 4\text{k}$). The closest that comes to this is a 2017 1024×1280 pixel substrate removed InGaAs detector manufactured by Princeton Infrared Technologies^d (PIRT), offered in a camera build called the 1280SciCam. It operates at $-60 \text{ }^\circ\text{C}$, with water cooling, and claims to have dark currents of $300 \text{ e}^-/\text{s}/\text{pixel}$, with a readout noise of 90 e_{rms}^- (35 e_{rms}^- was stated on their website in 2018^e) using CDS. The dark current is notably lower than seen in

^dDate accessed 2022-02-27: <https://www.princetonirtech.com/products/1280scicam>

^e<https://web.archive.org/web/20180227174252/http://www.princetonirtech.com/products/>

HgCdTe at the equivalent temperature, however the read noise is 18 times higher than what was achieved in [Hoffman et al. \(2005\)](#) and 6 times higher than what is offered by Teledyne. There are other InGaAs detectors on the market, such as those offered by Raptor Photonics (who use a sensor from Semiconductor Devices), and Sony, offering the same resolution as PIRT. These can be purchased off-the-shelf and embedded into their own system, however their sensor characteristics are not as competitive as PIRT's overall offering (see Table 4.1). Sony's offering presents one interesting characteristic however, a higher QE in the visible domain than the other substrate removed InGaAs detectors. This is due to its unique approach of using copper instead of indium bumps when hybridising onto the silicon CMOS ROIC, however it was not considered formally as part of this work.

While we can optimise detector properties, ground based observatories also have to consider the effects of the atmosphere, and whether they pose a limit on the signal-to-noise ratio on their observations.

1.3 Our atmosphere's influence on observations

Photometry in the NIR presents more challenges than in the visible. Atmospheric radiance and transmission, in addition to its variability ([Hecht et al. 1995](#); [Otarola et al. 2019](#)), are more noticeable in the NIR and can therefore lead to many difficulties when processing time-series photometry. Figure 1.5 shows the median atmospheric transmission and radiance spectra from the Paranal Observatory in Chile.

Atmospheric and environmental radiance can dominate the noise budget of a single exposure, particularly when the target is of low flux. Where radiance is a sum of multiple components, namely; scattered moonlight, scattered starlight, zodiacal light, thermal emission by the telescope and any instrumentation attached, molecular emission in the lower atmosphere, airglow emission lines of the upper atmosphere, and the residual airglow continuum. For observations in the NIR, the largest concern derives from the variable emission lines from free radicals in the upper atmosphere. This variability places background estimation as a critical aspect of time-series photometry. However, with differential photometry and techniques linked to it, this background variability is slightly less of a concern, as is discussed in the following subsection.

A background effect that is similarly difficult to overcome is fringing. Depending on the amount and type of emission lines, internal reflections can occur within a detector that lead to interference patterns in an image. The pattern itself depends on the thickness of the detector medium and is effectively unchanged over the lifetime of the detector. However, the intensity of the fringing pattern varies with the amount and type of emission lines. Radiance is observed

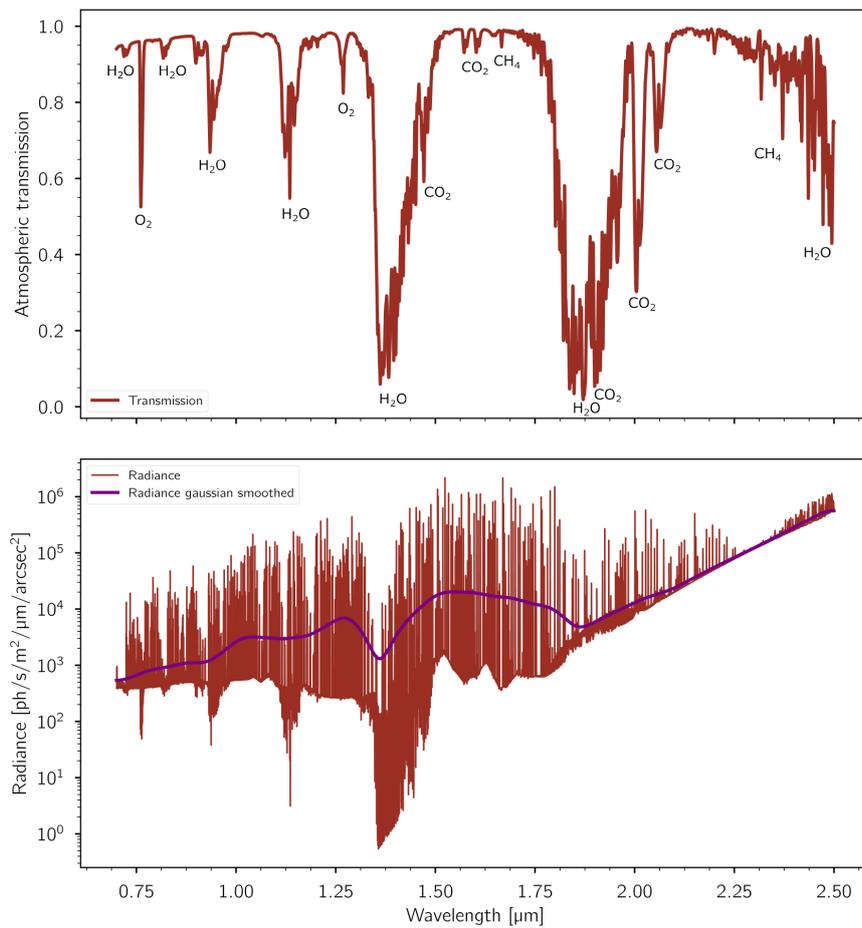


Figure 1.5: Top: The atmospheric transmission spectrum at airmass 1 with molecular absorption regions labelled (Smette et al. 2015). Bottom: Sum of all radiance components, with the spectrum Gaussian smoothed superimposed. Spectra modelled for Paranal, Chile, generated using the The Cerro Paranal Advanced Sky Model (Noll et al. 2012; Jones et al. 2013) for a median night with 2.5 mm of precipitable water vapour.

to be worst at the start and end of the night due to the sun providing energy to form free radicals in the upper atmosphere. Thermal emission from the environment, telescope and the attached instrumentation, as well as molecular emission in the lower atmosphere dominate from $\sim 2 \mu\text{m}$.

Atmospheric transmission is usually above 95% in the NIR, at zenith. There are a few distinct regions that are of lower transmission due to molecular absorption lines, most notably by H_2O . In the NIR, absorption due to O_2 , CO_2 , and CH_4 are also present, but are less of a concern when compared to H_2O (Smette et al. 2015). There are similar changes to atmospheric transmission due to ozone absorption, Rayleigh scattering by air molecules, and Mie scattering by aerosols, however their effects are negligible in the NIR. Transmission decreases as a function of increasing airmass due to the increasing thickness of the atmosphere at higher zenith angles that light has to transverse. The various sources of absorption can usually be handled with differential photometry techniques. Such as, observations made through uniform thin clouds, as it is assumed clouds behave as grey absorbers. However, small differences can arise in differential photometry due to the variability of molecular absorption lines throughout the night. Long and short term variability in time-series photometry can appear due to changes in atmospheric composition, as demonstrated and corrected for precipitable water vapour variability in Chapter 3.

Another ground based limitation derives from turbulence in the atmosphere, where spatially and temporally varying refractive indexes of the different layers in the atmosphere adds additional noise to an observation. In addition to this, the angular resolution of the telescope, if not limited by the diffraction limit of the optics, is worsened by the blurring induced by turbulence. This blurring is quantified by the ‘seeing’, the full-width-half-maximum (FWHM) of an imaged point-like source in the field, where the two-dimensional profile often resembles a Gaussian distribution. Seeing generally improves into the infrared (Glass & Glass 1999), where the seeing scales with square inverse of the wavelength.

1.3.1 Differential photometry

Observing a transit in the NIR from a ground based observatory utilises the same post-processing technique as in the visible, image reduction (with the addition of bad pixel treatment), estimated background removal, followed by a technique called differential photometry. It is designed to minimise any of the atmospheric variabilities that may affect a target’s observed change of flux; the atmospheric effects described in the previous subsection. The process involves comparing the relative changes in brightness between the target and other stars in the field of view, called comparison stars, by taking images periodically throughout the entire course of a night. A differential light curve is then formed by dividing time-series values of normalised fluxes from the target star with one or more comparison stars. The comparison stars are vetted depending

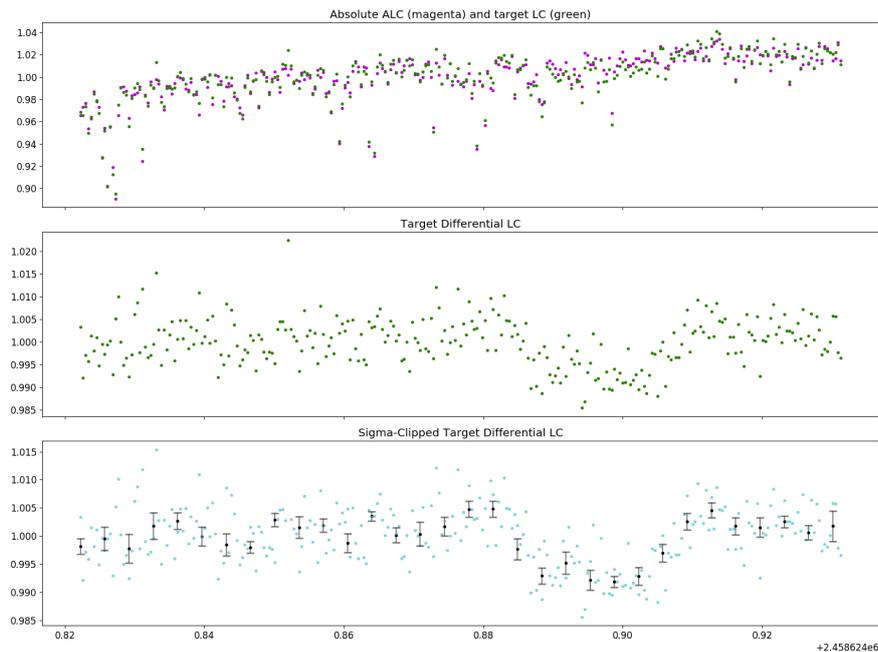


Figure 1.6: Trappist-1b transit in relative flux, as observed by SPECULOOS South (see Chapter 2) on the night beginning 2019-05-20. Time units in Julian date. Top: Comparison between the generated artificial light curve (ALC) from all the vetted comparison stars (magenta), and the light curve (LC) from the target, Trappist-1 (green). Middle: Resulting differential LC of Trappist-1. Bottom: Differential LC clipped based on a set threshold, with 5 minute bins overlaid. Credit: Catriona Murray.

on their brightness, colour, and angular distance to the target star. Their normalised fluxes are weighted and combined to form an artificial light curve (ALC), detailed in [Murray et al. \(2020\)](#). The process is demonstrated for an observation of a Trappist-1b transit in Figure 1.6.

The flux from a star is usually measured using aperture photometry, where the flux is taken as the summed counts of a circular region of pixels around a star. Aperture photometry offers the fastest method of measuring stellar fluxes, however it struggles when two or more stars are within the same aperture, or when the aperture is too small and the star's point spread function (PSF) extends beyond the aperture's radius and PSF's FWHM is variable in the night. Therefore, it's important for the angular resolution of the camera, making the observations, to be sufficient enough to distinguish the target from an optical neighbour. However, to perform accurate differential photometry, one must have a sufficiently wide enough field of view to obtain enough comparison stars. There is therefore a balance one has to make between detector size and pixel size, which are usually dictated by cost and/or availability of detectors that meet the requirements. Guiding is similarly of importance, such that the same pixels are used throughout the differential photometry process. Without sufficient guiding, one risks

introducing systematic noise from, for example, pixel-to-pixel sensitivity differences or flat fielding errors.

The remainder of this thesis includes further introduction topics throughout the text. The following chapter presents a detailed introduction to a robotic transit survey, SPECULOOS, and its operations. The remaining chapters cover a range of topics related to SPECULOOS; including precipitable water vapour, photometric modelling, and finally, characterisation of a new NIR instrument.

CHAPTER

SPECULOOS

The majority of the contents in this chapter are from [Sebastian et al. \(2020b\)](#), where I was second contributing author. This chapter serves as an introduction to SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars) and its operations. The text has been partially updated since publication and only sub-sections relevant to my thesis included. My specific contributions are detailed in the declaration of this thesis.

SPECULOOS aims to perform a transit search on the nearest (< 40 pc) ultracool (< 3000 K) dwarf stars. The project's main motivation is to discover potentially habitable planets well-suited for detailed atmospheric characterisation with present and upcoming giant telescopes, like the James Webb Space Telescope (JWST) and European Large Telescope (ELT). The project is based on a network of 1 m robotic telescopes, namely four located at SPECULOOS-Southern Observatory (SSO) in Paranal, Chile, one telescope at SPECULOOS-Northern Observatory (SNO) in Tenerife, and the SAINT-Ex telescope in San Pedro Mártir, Mexico. The prototype survey of the SPECULOOS project on the 60 cm TRAPPIST telescope (Chile) discovered the TRAPPIST-1 system, composed of seven temperate Earth-sized planets orbiting a nearby (12 pc) Jupiter-sized star. In this chapter, I review the current status of SPECULOOS, its first results, the plans for its development, and its connection to the Transiting Exoplanet Survey Satellite (TESS) and JWST.

2.1 Introduction

The past decade of atmospheric studies on close-in giant exoplanets has enabled the first constraints on planetary exospheres, molecular and atomic species, elemental ratios, temperature profiles, clouds, and even atmospheric circulations (Madhusudhan 2019). The main laboratories driving this revolution are thousands of transiting exoplanets that have been detected by ground and space-based transit search surveys. The field has developed to perform similar atmospheric studies for smaller and more temperate planets. Examples are GJ 1214b (Charbonneau et al. 2009) (M4.5 V) or more recently K2-18b (Tsiaras et al. 2019; Benneke et al. 2019) (M2.5 V). For the latter, a temperate ‘mini-Neptune’, the detection of atmospheric water vapour marks an important step towards constraining the atmospheric composition of potentially habitable worlds. Nevertheless, for temperate, rocky exoplanets that transit solar-type stars, such atmospheric characterisations are far beyond reach with current and upcoming instrumentation.

This is not true for the latest-type dwarf stars. First, the signal-to-noise ratio (SNR) for eclipse spectroscopy measurements depends on the size-ratio between a planet and its host star and therefore increases for late-type dwarf stars. Second, planets with shorter orbital periods can receive the same stellar irradiation as Earth due to the lower luminosities of their host star, resulting in more frequent planetary transits for temperate, rocky planets.

Studies (Kaltenegger & Traub 2009; de Wit & Seager 2013; Morley et al. 2017) have shown that the combination of both effects brings the atmospheric characterisation of temperate Earth-sized exoplanets within reach of the James Webb Space Telescope (JWST). This applies under the condition that they transit stars that are both very nearby (< 15 pc) and very small ($< 0.15 R_{\odot}$), with spectral type M6 or later. The first detection of terrestrial planets orbiting a M-dwarf was Kepler-42 (Muirhead et al. 2012) (spectral type M4), with its compact system of sub-Earth-sized planets. Other transiting rocky planets were detected around mid-to-late-type M-dwarfs by MEarth (Nutzman & Charbonneau 2008) (GJ1132, a M4 V in Berta-Thompson et al. (2015); LHS 1140, a M4.5 V in Dittmann et al. (2017) and Ment et al. (2019)) and by TESS (Ricker et al. 2015) (LP 791-18, a M6 V in Crossfield et al. (2019); LHS 3844, a M6 V in Vanderspek et al. (2019); TOI 540, a M4 V in Ment et al. (2020); LP 890-9, a M6 V in Delrez et al. (2022)).

Although, short-period, rocky planets are more common around low-mass stars (Howard et al. 2012; Hardegree-Ullman et al. 2019) and show deeper and more frequent transits than for earlier-type host-stars, not many of such transiting planets have actually been detected around late-type M-dwarfs. This is especially true for ultracool dwarfs. The classical definition of ultracool dwarfs (hereafter, UCDs) includes dwarfs with spectral type M7 and later, as well as brown dwarfs (BDs) (Kirkpatrick et al. 1997; Kirkpatrick 2005).

The main reason for this apparent lack of planets is the faintness and activity of these stars. Transit search surveys using telescopes with relatively small apertures (< 50 cm), like M_{Earth} and TESS, have a high detection potential for mid-type M-dwarfs (M3 to M5), but this potential drops sharply for later-type objects (Sullivan et al. 2015; Barclay et al. 2018). The unique system of seven transiting planets that has been detected orbiting TRAPPIST-1 (Gillon et al. 2016; Gillon et al. 2017; Luger et al. 2017), a nearby (12 pc) M8V star, is up-to-date the most optimal, laboratory for atmospheric characterisation with JWST in the temperate Earth-sized regime (Lustig-Yaeger et al. 2019; Macdonald & Cowan 2019).

Due to the scarcity of such detections, not much is known about the structure of planetary systems of late-type M-dwarfs. Until now, TRAPPIST-1 has remained the only transiting system discovered around an UCD, while some statistical constraints could be inferred from the null results of several recent projects (Demory et al. 2016; He et al. 2017; Sagar et al. 2019; Lienhard et al. 2020).

The kinds of planetary system could be formed around UCDs have been analysed by several theoretical studies (Payne & Lodato 2007; Raymond 2007; Lissauer 2007; Montgomery & Laughlin 2009; Alibert & Benz 2017; Coleman et al. 2019; Schoonenberg et al. 2019; Miguel et al. 2020; Liu et al. 2020), which spanned a variety of outcomes from water-rich to water-poor planets, along with several orbital architectures. Notably, the inferred period distribution of planets for typical UCDs. While UCDs are commonly expected to have one or more planets with a period distribution that peaks for short-period planets (Coleman et al. 2019), a bimodal period distribution also seems to be likely (Miguel et al. 2020) with one group with very short periods and another group with larger periods, located beyond the snow line. A transit search survey, only focusing on short-period planets of UCDs can already distinguish between these various outcomes.

A photometric survey specifically designed to search for transits of Earth-sized planets, or smaller, among the nearest UCDs will allow first, the detection of Earth-like planets amenable for atmospheric characterisation, and second, to draw constraints that could be used to differentiate between different models of planet formation among the lowest-mass stars. Given the relative faintness of those objects and the lessons learned from the TRAPPIST-UCDTS prototype survey (Gillon et al. 2013; Burdanov et al. 2018; Lienhard et al. 2020), we developed SPECULOOS (Search for habitable Planets Eclipsing ULtra-coOL Stars). The SPECULOOS project has been initiated and led by the University of Liège and is done in collaboration with the Universities of Cambridge, Birmingham, Bern, and Massachusetts Institute of Technology. It is based on a network of 1 m class, robotic telescopes that have a significant detection efficiency for planets as small as the Earth – and even smaller – for a large fraction of UCDs within 40 pc. The goals of SPECULOOS are: (i) the search for rocky planets that are well-suited for atmospheric



Figure 2.1: SPECULOOS Southern Observatory (SSO) with four 1 m class telescopes, located at ESO Paranal Observatory (Chile). Telescope’s names from left to right: Callisto, Ganymede, Europa, and Io.

characterisation with facilities like JWST and (ii) more globally, to perform a volume-limited (< 40 pc) transit search of UCDs to derive a statistical census of their short-period planet population (Delrez et al. 2018; Gillon 2018).

In this chapter, I detail the current status of this survey. In Section 2.2, I describe the observatories that form the ground-based SPECULOOS network, and their remote operations. Section 2.3 details the target selection, observing programs and scheduling. In Section 2.4 I review the data reduction, photometric quality control and impact of atmospheric effects on near-infrared (NIR) observations of UCDs. And finally, in Section 2.5 I detail the data archive and user interface for the SPECULOOS survey.

2.2 Network of ground-based observatories

The SPECULOOS network (see Table 2.1) is composed of six telescopes at three different sites. The SPECULOOS Southern Observatory (Delrez et al. 2018; Jehin et al. 2018; Murray et al. 2020) (SSO) with four telescopes named after the four Galilean moons Io, Europa, Ganymede, and Callisto at ESO Paranal Observatory (Chile) became fully operational in January 2019, see Figure 2.1. The SPECULOOS Northern Observatory (Niraula et al. 2020) (SNO) is currently composed of one telescope (Artemis), which is located at the Teide Observatory (Canary Islands, Spain) and operational since June 2019. SAINT-EX (Demory et al. 2020) (Search And characterisation of Transiting EXoplanets) with one telescope in San Pedro Mártir observatory (Mexico) became operational in March 2019. Each of these observatories is devoting 70%

Table 2.1: Observatories of the SPECULOOS network of telescopes

| Observatory | Telescopes | Host observatory |
|-------------|------------|---|
| SSO | 4 | ESO Paranal Observatory, Chile |
| SNO | 1 | Teide Observatory, Tenerife, Spain |
| SAINT-EX | 1 | Observatorio Astronómico Nacional, Mexico |

of its usable observational time to the SPECULOOS survey. The remaining time is used for different programmes. For example, a fraction of this time of the Artemis telescope is dedicated to educational and outreach programs. These programs are carried out with collaborators in the Canary Islands and from the Boston area. In addition, the two 60 cm TRAPPIST (Gillon et al. 2011; Jehin et al. 2011) robotic telescopes (one in Chile, the other in Morocco), while not officially part of the SPECULOOS network, have devoted a fraction of their time to supporting the project and focus on its brightest targets.

Each of the SPECULOOS observatories is composed of identical robotic Ritchey-Chrétien telescopes with 1 m aperture, designed and built by the German company ASTELCO^a. The optical design contains a 1 m primary mirror (f/2.3) coupled with a 28 cm diameter secondary. The combined system’s focal ratio is f/8. Each telescope is equipped with a robotic equatorial ASTELCO New Technology Mount NTM-1000. This mount uses direct-drive torque motors, which allow fast slewing (up to 20 degrees per s), accurate pointing (better than 3 arcseconds) and tracking accuracy better than 2 arcseconds over 15 minutes without any guiding. It is designed for continuous tracking without meridian flip. The telescope is installed in a 6.25 m dome designed and built by the Gambato^b company, which is fully integrated to the robotic telescope control system.

Each telescope is equipped with an Andor iKon-L^c thermoelectrically cooled camera with a near-infrared optimised, deep depletion $2k \times 2k$ e2v CCD detector (13.5 μm pixel size). The field of view on the sky is 12×12 arcminutes, yielding a pixel scale of 0.35 arcseconds/pixel. Exposure control is realised by a mechanical shutter, using overlapping iris blades. Despite the fact that these shutters are generally very durable, they have a limited lifetime. To increase this shutter lifetime, we restrict our remote observations to exposure times larger than 10 seconds. The camera is usually operated at -60 °C (via five-stage Peltier cooling) with a dark current around $0.2 \text{ e}^-/\text{s}/\text{pixel}$. The detector provides high sensitivity in a wide range (350-950 nm), with a maximum quantum efficiency of 94% at both 420 and 740 nm. Each camera has its own filter wheel providing the Sloan g' , r' , i' , z' , and two special exoplanet filters; the near-infrared

^a<http://www.astelco.com>

^b<http://www.gambato.com>

^c<https://andor.oxinst.com/products/ikon-xl-and-ikon-large-ccd-series/ikon-l-936>

luminance I+z' filter (transmittance > 90% from 750 to beyond 1000 nm, which is mostly used for the SPECULOOS core program); and a blue-blocking filter called Exo (transmittance > 90% from 500 to beyond 1000 nm).

Every aspect of the remote observations is controlled via ACP Expert Observatory Control Software^d. It communicates to the camera, the telescope control system via an ASCOM^e compatible interface, and provides an inbuilt weather server that allows to continuously monitor data from a weather station. The human intervention is reduced to a quick status check and manual starting of the night's observing plans which is carried out remotely via a secure Virtual Private Network (VPN) connection. We developed a planning tool, named the SPeculoos Observatory sSchedule maKer (SPOCK, which is optimised to the 'long stare' observation strategy of SPECULOOS (see Section 2.3)). Its primary aim is to automatically generate and submit daily observing scripts (ACP observing plans) for the SSO, SNO, and SAINT-EX observatories. These are simple text files, which are linked to each other and allows to pre-define the nights observing, flat and calibration schedule. Its source code is available on Github^f. These pre-defined observation plans allow us to automatically start the observations at any time in the night when weather conditions are fine. In case of unsafe weather, the ACP weather server allows to stop the observations and close the dome in order to avoid damage to the equipment. Additionally, the dome is directly connected to the weather station which offers a redundant way to automatically close the dome in such conditions.

To allow automatic continuation of observations, I have developed a script that monitors the ACP weather server after the observations have been terminated due to high winds, clouds or high humidity. It implements specific conditions that need to be met during a specific waiting time to apply lower wind, cloud, or humidity limits than we use for the termination of observations. It allows us automatically to determine whether the conditions have improved and if it's safe to continue observations, using the pre-defined plans. In case rain or snow has been detected, an automatic continuation will not be possible for safety reasons. This automated script has been tested and adopted for the SSO and SNO observatories and is in regular use for these sites. In principle, it can be used for any observatory which is controlled by ACP. All telescopes require a minimum of on-site maintenance. The host observatories provide emergency help in case of technical difficulties as well as regular check-ups to ensure continuous robotic operations.

^d<https://acpx.dc3.com>

^e<https://ascom-standards.org>

^f<https://github.com/educrot/SPOCK>

2.3 Targets and Observations

2.3.1 Target selection

The SPECULOOS target list contains a homogeneous selected sample of close-by low-mass stars and UCDs (Sebastian et al. 2020a). The targets have been selected as low-mass dwarfs starting from the Gaia DR2 point source catalogue (Gaia Collaboration et al. 2016, 2018) which has been cross-matched with the 2MASS point-source catalogue (Skrutskie et al. 2006). During this cross-match we enforced the agreement between the two catalogues not only in terms of position but also in terms of effective temperatures inferred from different photometric indicators. The final selection of the SPECULOOS target catalogue is based on each target’s individual parameters and its match to the SPECULOOS science goals. These are (i) finding rocky Earth-like planets around close and bright targets, that are well-suited for atmospheric characterisation with JWST and (ii) to provide a census of the short-period planet population of UCDs. For the brighter targets in this catalogue, we leverage the efficiency to detect small planets within TESS. That means, whether these targets are monitored with the SPECULOOS network depends on whether the photometric quality of the TESS data allows the detection of small transit signals. Such signals will be followed up by the SPECULOOS telescope network. The 40 pc list of late-type targets and SPECULOOS target list are publicly available (Sebastian et al. 2020a) together with a detailed description of the target selection and survey strategy. The target list is divided into three non-overlapping observing programmes:

- Programme 1: includes 365 targets for which a transiting Earth-like planet (same size, mass, irradiation, and atmospheric composition as the Earth) will allow atmospheric characterisation by means of transit transmission spectroscopy with JWST. For the brighter and earlier targets from this list SPECULOOS telescopes follow up any transit signal within TESS. Each of the remaining targets will be observed for at least 200 hr in order to effectively survey any possible period up to the mid-habitable zone (Kopparapu et al. 2013) of these stars with an average integrated phase coverage of 80%.
- Programme 2: includes 171 targets with spectral type M5 and later, which are not in Programme 1, but bright enough to allow a detection of Earth-sized temperate (here: irradiation of $4 S_{\oplus}$) planets – like TRAPPIST-1b – with TESS. The SPECULOOS telescopes follow up any transit signal detected within TESS.
- Programme 3: includes 1,121 targets with spectral type M6 and later and aims to explore the occurrence rate for temperate planets of UCDs within 40 pc. In order to detect short-period up to temperate planets – like TRAPPIST-1b – each of these targets will be observed for at least 100 hr. This allows us to effectively survey any possible period smaller than 6 days with an average integrated phase coverage of 80%.

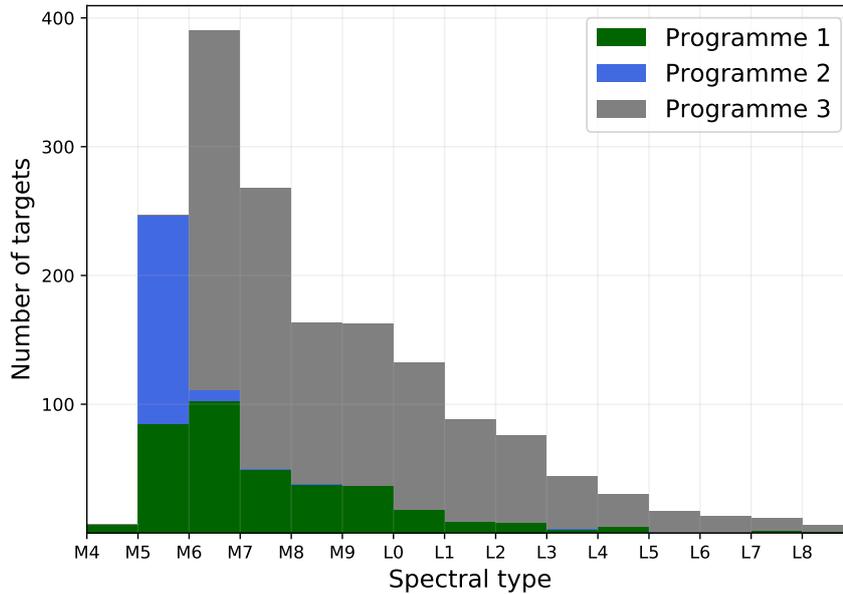


Figure 2.2: Spectral type distribution of all targets within the three SPECULOOS programmes.

The full target catalogue comprises of 1,657 targets. As shown in Figure 2.2, the spectral type distribution of the SPECULOOS catalogue peaks at M6V with about 400 targets and decreases with later spectral types. The catalogue is incomplete for targets earlier than M6V because of the varying spectral type cut of the different programmes. Nevertheless, no cuts were introduced for later-type stars to maximise the catalogue completeness for UCDs.

2.3.2 Scheduling

Each night, we continuously observe one or two targets per telescope. Despite strategies with non-continuous observations being in use (Nutzman & Charbonneau 2008; Tamburo & Muirhead 2019), this observation approach has been shown to be more reliable for the detection of short transit events (down to 15 min for very-short-period ≤ 1 d planets) among UCDs (Delrez et al. 2018; Gibbs et al. 2020).

The SPECULOOS planning tool SPOCK schedules targets automatically depending on their priority and observability. The priority is defined as the value of the SNR that has been used to select the targets for the different observing programmes. The SNR for JWST transmission spectroscopy for an ‘Earth-like’ planet in Programme 1, the TESS detection SNR for a temperate planet in Programme 2, and the SPECULOOS detection SNR for a temperate planet in Programme 3. The higher the corresponding SNR the higher is the targets priority. The priority of each target also depends on the completion factor. The completion ratio is

defined as

$$r_c = \frac{hours_{observed}}{hours_{threshold}},$$

which embodies the fraction of hours of observation completed versus the number of hours required for each target. The required number of hours for each target is defined by the observing programme it belongs to. This completion ratio is used to rank targets to favour the quick completion of ongoing targets, as opposed to starting new ones continually.

The observability is defined as the best visibility window for each target, based on its coordinates. For each night, the selection process ranks all targets that are at their optimum visibility (respecting constraints imposed by the observatory such as Moon angular distance and altitude) and selects the one with highest priority. To implement those constraints, SPOCK makes use of the `ASTROPLAN` package (Morris et al. 2018), which is a flexible Python toolbox for astronomical observation planning and scheduling. It also optimises the annual period of the year for which the target is the most visible at a relatively low airmass, based on the completion ratio of the target and the optimisation of available observatories.

Targets that are observable all night long will be scheduled for the complete night. However, some targets have latitudes that do not allow for observation during all the available night time for the given site – even at their peak of visibility – such that observational gaps at the beginning or end of the night can appear for those targets. In that regard, an additional target is added to complement the schedule and avoid losing observing time. The observation durations of the two targets are set to be comparable, to prevent having overly short observation blocks (1 hr or less). We say comparable because we do not exactly split the night in half; instead, we adopt a nightly observing duration adapted to each target’s visibility (e.g. depending on Moon angular distance), which will shift from night to night during the observation period. This situation of observing two targets per night is rather frequent and thus the most common observation mode for this survey.

One of the main roles of SPOCK is to handle the coordination of multi-site observations. For instance, between two targets with similar priorities but one observable only from one site and the other from several sites, SPOCK will choose the target that yields the most coverage. In addition, when possible, a one-hour overlap between observations from two different sites is scheduled to aid in the combination of the light curves.

2.4 Data reduction and quality

2.4.1 Data Reduction Pipeline

The raw images from the SSO observatory are automatically transferred at the end of each night to the online ESO archive[§]. These images are then downloaded to a server at the University of Cambridge (UK) the next day. In parallel, the raw images from the SNO observatory are directly transferred each day to the same Cambridge server.

Images from the SSO and SNO observatories are then processed daily by the automatic SSO Pipeline (Murray et al. 2020). The SSO Pipeline is custom-built for the calibration and photometry requirements of the SPECULOOS survey, following a modular architecture similar to the Next-Generation Transit Survey (NGTS) pipeline (Wheatley et al. 2018a), and utilising the CASUTOOLS (Irwin et al. 2004) package of image processing tools (namely IMCORE, WCSFIT, IMSTACK and IMCORE_LIST). The science images are reduced using standard methods of bias and dark subtraction and flat-field correction. We then use a local version of ASTROMETRY.NET code (Lang et al. 2010) to cross-match our science images with reference catalogues built from the 2MASS catalogue to generate initial World Coordinate System (WCS) solutions. To refine these WCS solutions we perform source detection on our images and then, using the initial WCS solutions, cross-match these sources with the Gaia DR1 catalogue (Gaia Collaboration et al. 2016), to further correct each image for translations, skews, scales, and rotations. For each field that is observed the pipeline requires an input catalogue, containing the positions of every source from which to extract aperture photometry data. We choose to have one, unique catalogue for every observed field to reference for all nights that field is observed. This allows us to monitor the photometry of all sources in this field consistently over long periods of time. To create this catalogue we perform source detection on a stacked image, generated from 50 images taken in the middle of the first night of observation of a field. We also cross-match this catalogue with Gaia DR2 (Gaia Collaboration et al. 2018) to apply proper motion corrections on a night-by-night basis. There is also the facility to cross-match with other catalogues, such as 2MASS (Skrutskie et al. 2006). This catalogue then defines the central positions of apertures from which we measure raw photometry for each object in the field, for 13 different aperture radii.

Once the raw, aperture photometry has been extracted, the SSO Pipeline corrects for ground-based systematics shared by other objects in the field by using an automated differential photometry algorithm. This iterative algorithm calculates an ‘artificial’ comparison light curve by weighting the sufficiently bright comparison stars according to their variability and distance to the target. Due to the survey’s design, the redder comparison stars in the field are significantly

[§]http://archive.eso.org/eso/eso_archive_main.html

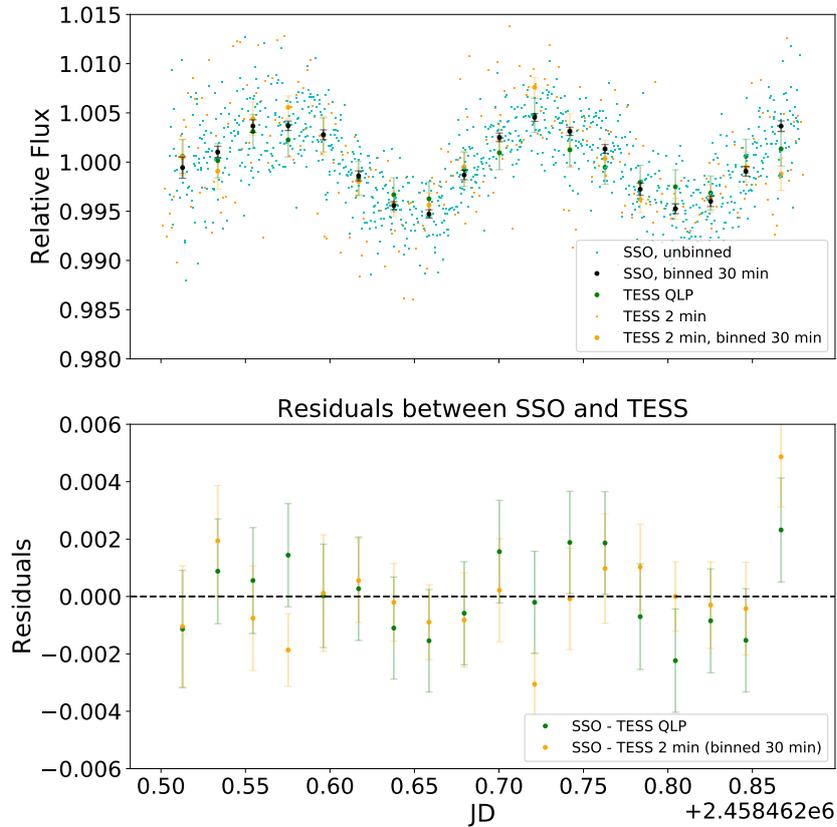


Figure 2.3: Top: Comparison of light curves from the SSO observatory and TESS for an M6V object ($J = 10.3\text{mag}$) on 2018 December 10. The differential light curve from a single SSO telescope is in cyan, with 30 minute binned data in black, the light curve from TESS 2 minute cadence data is in orange and MIT Quick Look Pipeline 30 minute cadence data is in green. Bottom: The residuals between the SSO light curve and the two TESS light curves.

fainter than the target, therefore we do not input any colour information into the weightings, to avoid introducing noise into the target’s light curve. This mismatch in spectral type between the target and comparison stars leads to second-order differential extinction residuals imprinted on the target’s light curve. To mitigate these residuals we implement a correction for precipitable water vapour from first principles, detailed in Chapter 3.

2.4.2 Photometric quality

To assess the photometric quality of the SSO observatory, we measured the fractional RMS of our target light curves for 30 minute bins, from January 1st (start of official scientific operations) to September 18th, 2019 (Murray et al. 2020). This data sample contained 98 targets observed over 179 nights of observation with the SSO telescopes. We performed no detrending, correction

for intrinsic variability or removal of sub-optimal observing conditions. In this analysis, we demonstrated that we are reaching sub-millimag precisions for $\sim 30\%$ of light curves (with a median precision of 1.5 mmag), and up to 0.26 mmag for the brightest objects. There have been recent developments to the SSO Pipeline to clean the final light curves, including flagging frames where the observing conditions significantly affect the photometry. This process also includes ongoing work to identify and remove flares (Murray et al. 2022), and modelling of the stellar variability of our targets from long term photometric monitoring. Removing the activity of our targets will allow us to further refine the photometric quality of our facilities and improve our detection efficiency for small planets.

In addition, we compared simultaneous observations of a variable M6V object with a single SSO telescope and TESS 2 minute and 30 minute cadence observations (Figure 2.3) (Murray et al. 2020). There was excellent agreement between the three datasets, with the SSO light curve exhibiting the least white noise. As TESS is not optimised for UCDs, we expect that the quality of the SSO light curves will exceed TESS when we observe cooler and redder objects, however for the bright M5V and M6V objects the quality of the photometry will be comparable. This demonstrates how we can utilise the synergy between SPECULOOS and TESS to optimise the detection of Earth-sized planets, as in Programme 2.

2.5 SPECULOOS Portal

To enable easy access to the SSO Pipeline’s output, a web-based service and interface was designed, called PORTAL (Pipeline Output inteRacTion Analysis Layer). It was built using a common backend stack – LAMP (Linux, Apache, MySQL, PHP/Python), with a Vue.js, Plotly, and D3.js based frontend. Its main operation permits quick and interactive visualisation of nightly light curves from the SSO and SNO observatories. It also serves as a RESTful API, which allows members of the SPECULOOS consortium to download and analyse any of the pipeline’s outputs and connected metadata.

2.5.1 Backend

The backend bridges the link between SSO Pipeline data products and the user-friendly frontend interface, by behaving as a RESTful API. The pipeline’s output directories are navigated by the backend, using Gaia DR2 source IDs as the main search parameter, with the option to further refine by specifying date, telescope, and filter – with arrays of values and wildcards permitted in a search query. To retrieve nightly observation data, the backend processes multiple text files (one per aperture) containing nightly differential light curves (PWV and non-PWV corrected)

and metadata (relative RA and DEC movement, FWHM, PSF, sky level, and airmass), and formats them into a single JSON structure.

For larger queries, a SQL database was implemented. It contains all the nightly observation data, as well as environment/telescope specific metadata. The tables are updated whenever a directory is updated by the pipeline. The database also stores SPECULOOS' observation history (based on the pipeline's output directories) and the observation schedule (provided by SPOCK's output files). The SQL database also stores user submitted flags and comments on nightly data provided by the frontend.

The total number of hours a target has been observed for, as processed by the pipeline, can be also queried. The backend also produces low resolution videos from the raw $2K \times 2K$ images acquired during an observation, which forms part of a target's detailed observation view produced by the frontend.

2.5.2 Frontend

The frontend is a user-centred designed interface for displaying differential photometry data. It allows navigation and interaction with observations made by all of SPECULOOS' facilities, on a target and nightly basis. The user-friendly interface, as shown in Figure 2.4, has a vertical menu panel and an interchanging main pane.

The navigation panel contains a search function with autocomplete capabilities. A user can search with either a target's Gaia DR2 source ID or the shorter SPECULOOS target ID. By default, the search function will return all the light curves for a particular target, ordered by newest date first. Nightly light curves are presented within interactive scatter plots, with raw and binned data. In its default view, the scatter plots are displayed with PWV corrected light curves (if available) using an aperture value pre-determined by the pipeline, within the differential flux range of $[0.98, 1.02]$, and a binned period of 0.005 JD. A user can quickly interact with the data by zooming in/out and toggling between apertures, binned periods, and the PWV applied correction. If a user spots an interesting feature, such as a flare, a transit feature, a type of variability, or an issue with the data, they can tag the data, as well as submit a comment to be connected with the observation. On hover, one can view previous submissions, or one can view all submissions for a target on a separate tab.

To vet the data further, one can access a target's detailed observation view for a night, as exemplified in the bottom of Figure 2.4. Here, one can explore metadata, such as relative RA and DEC movement, FWHM, sky level, and airmass. Some of the extended metadata are quality checked for anomalies, which are used to inform the user. A video of the observation is also present, which allows one to navigate images of the night, triggered by hovering over

Exploration tools
latest observations or random ones

An extensive **search bar** with autocomplete features

Quick links

Centralised **documentation** for download and analysis

An interactive calendar to explore our **schedules**

Target header with access to related **Slack discussions**

Interactive light-curves featuring collaborative tools (submissions)

Detailed observation view

Easy vetting of common observation measurements

Live view of observation images when hovering plot(s)

Figure 2.4: A descriptive view of PORTAL’s web interface. Top: A target specific view is shown for Trappist-1, with an example set of interactive light curves and a user submitted flag. Bottom: A detailed observation view for a particular night for Trappist-1b. The collection of tools PORTAL offers allows for an intuitive and collaborative exploration of differential photometry data, which can be downloaded for extensive offline analysis via the backend’s RESTful API.

different parts of the light curve. Lastly, one can automatically open Slack channels to allow for further discussions of the target.

To promote further vetting of data, a user can view the latest light curves, a random target's light curves, or a detailed observation view for a random target and night. The schedule can also be navigated by date, and if the target has been observed, it links to its respective detailed observation view. The final links on the navigation panel are to the telescopes' status pages, and a timelapse view of SSO's widefield night-sensitive webcam.

In future, I will continue to improve the usability and ease of the platform. One planned improvement is to integrate analysis tools that are currently being developed by consortium members to further vet nightly observations.

PRECISE NEAR-INFRARED PHOTOMETRY, ACCOUNTING FOR PRECIPITABLE WATER VAPOUR AT SPECULOOS SOUTHERN OBSERVATORY

The work in this chapter was published in [Pedersen et al. \(2022\)](#). It describes the photometric effects induced by the variability of precipitable water vapour (PWV). Specifically, its variability can heavily affect the accuracy of time-series photometric measurements gathered from the ground, especially in the near-infrared. I present here a novel method of modelling and mitigating this variability, as well as open-sourcing the developed tool – Umbrella. In this study, I evaluate the extent to which the photometry in three common bandpasses (r' , i' , z'), and SPECULOOS' primary bandpass ($I+z'$), are photometrically affected by PWV variability. In this selection of bandpasses, the $I+z'$ bandpass was found to be most sensitive to PWV variability, followed by z' , i' , and r' . The correction was evaluated on global light curves of nearby late M and L-type stars observed by SPECULOOS' Southern Observatory (SSO) with the $I+z'$ bandpass, using PWV measurements from the LHATPRO. A median reduction in RMS of 1.3% was observed for variability shorter than the expected transit duration for SSO's targets. On timescales longer than the expected transit duration, where long-term variability may be induced, a median reduction in RMS of 53.8% was observed.

3.1 Introduction

Ground-based photometric observations are affected by atmospheric variability. A major source of this contamination comes from the multitude of molecular absorption lines which are known to affect atmospheric transmission. Predominantly in the near-infrared, time-varying amounts of H₂O in different layers of the atmosphere affect ground-based observations across a wide range of wavelengths – with the amount of H₂O in a column of our atmosphere quantified as the amount of ‘precipitable water vapour’ (PWV), normally quoted in millimetres. O₃, O₂, CO₂ and CH₄, amongst other molecular absorption lines, likewise play a role, however, often to a lesser extent (Smette et al. 2015).

To mitigate the majority of atmospheric effects, techniques such as differential photometry are commonly adopted (Howell 2006). They involve simultaneous observations of multiple objects in a field of view to estimate first-order changes of atmospheric transmission and instrumental effects over the course of an observational period. In differential photometry, objects of similar brightness and spectral energy distribution are used to calculate an ‘artificial’ comparison star. For example, in Murray et al. (2020), the artificial comparison star was created by applying a weight to all the objects in the field of view, accounting for their effective temperature, noise, variability and distance to the target of interest. However, depending on the observational bandpass and an object’s spectral energy distribution, the net flux observed on the ground can be seen to vary from object to object if the atmosphere is time-variable in its composition – leading to a second-order differential effect which cannot be corrected by differential photometry.

This effect is a particularly significant problem for high-precision ground-based photometric studies of cool stars, such as M dwarfs and later types (e.g., Blake et al. 2008; Irwin et al. 2008; Tamburo et al. 2022). This is due to the large disparity in spectral energy distribution frequently observed in a field of view, as cool stars are typically much redder than the comparison stars suitable for differential photometry in any given field of view. To reduce this effect, one could observe with a narrow-band filter (e.g. Garcia-Mejia et al. 2020), but at the cost of instrumental precision due to the reduction in photons collected. Instead, most of these studies have relied on post-correction methods to reduce the nature of this effect, with varying degrees of success.

Irwin et al. (2011) developed a “common mode” approach, which they applied to MEarth’s northern survey (715 nm long-pass filter). This approach involved using median values over 30-minute periods from multiple simultaneous observations of similar type objects, from 8 independent telescopes. The common mode, scaled per target via a least-squares optimisation, was then used to correct for atmospheric effects over the course of a target’s observation run. They noted that the scaling values correlated with stellar type.

Some observational bandpasses can be extremely sensitive to PWV changes, such as MEarth’s bandpass and the primary bandpass used in this study ($I+z'$, 700 – 1100 nm). Such bandpasses induce false variability, including structures able to mimic a transit feature or to hide a real one in differentially resolved light curves when subject to PWV changes over the course of an observation. To correct for this variability, atmospheric transmission profiles in the near-infrared can be modelled with tools such as Molecfit (Smette et al. 2015). These models can be used to correct photometric observations when the spectra of the observed objects are also known. Methods of quantifying PWV changes are therefore necessary for photometric surveys which use sensitive bandpasses, with temporal resolutions at least half the minimum expected transit duration to resolve such features, to follow Nyquist sampling.

In astronomy, on-site atmospheric transmission profiles are often inferred by spectrographs or by multi-band photometers with strategically placed narrow-band filters. The aTmcam multi-band instrument (Li et al. 2012, 2014) for instance, located on Cerro Tololo at ~2200 m, was able to quantify atmospheric PWV with a stated precision ~0.6 mm. A similar instrument called CAMAL (Baker et al. 2017), located on Mount Hopkins at ~2600 m (same site as MEarth’s northern facility), had a stated precision of better than 0.5 mm in dry conditions (PWV < 4 mm). The use of a spectrograph on the other hand, such as in Li et al. (2017), gained a precision of 0.11 mm when evaluating high-resolution near-infrared H-band spectra of hot stars from the APOGEE spectrograph (located on the Apache Point Observatory at ~2800 m), calibrated with GPS-derived PWV values.

PWV values derived from timing delays in GPS signals have historically been used for meteorological studies, with large networks of GPS-PWV derived data in the public domain, such as the SuomiNet project (Ware et al. 2000) – typically with a precision of ~1 mm. Remote-sensing satellites have likewise enabled wide spatial and temporal coverage of a multitude of atmospheric parameters globally. Marín et al. (2015) were able to estimate PWV in very dry conditions at the Chajnantor plateau (at ~5100 m) with historical observations made by the now-decommissioned GOES-12 satellite and validated by an on-site radiometer. They attained relative errors of 51% and 33% over the ranges 0 – 0.4 mm and 0.4 – 1.2 mm, respectively. Similar work was achieved in Valdés et al. (2021), yielding better uncertainties for Cerro Paranal at around 27%, validated similarly with an on-site radiometer.

Radiometers derive PWV values from water vapour emission lines in the GHz region. At high altitudes, or low PWVs, the 183 GHz emission line is often observed. At the Paranal Observatory (with an elevation of ~2600 m), a 183 GHz based radiometer was commissioned in October 2011, the Low Humidity and Temperature Profiling microwave radiometer (LHATPRO; Kerber et al. 2012), located on the VLT platform. It has a quoted accuracy of better than 0.1 mm (when PWV between 0.5 – 9 mm) and a precision of 0.03 mm, with a saturation limit of 20 mm.

The motivation for this study derives from the SPECULOOS (Search for habitable Planets EClipsing ULtra-cOOl Stars; Gillon 2018; Delrez et al. 2018; Murray et al. 2020; Sebastian et al. 2020a,b) project, a ground-based photometric survey targeting nearby (< 40 pc) late M, L-type stars, with its primary aim to discover transiting terrestrial planets. Exoplanets found with the SPECULOOS survey will enable the unique opportunity to observe their atmospheres for potential biological signals with future large observatories, like TRAPPIST-1 (Gillon et al. 2017) found by SPECULOOS' prototype study. However, to maximise the probability of finding such planets, one must minimise the red noise the survey is subjected to, including the noise induced by atmospheric PWV variability.

In this work, I have developed methods to photometrically correct for PWV-induced variability on differential light curves. Our study has been evaluated on photometric data from SPECULOOS' Southern Observatory (SSO), which consists of four 1 m class telescopes located at the Paranal Observatory (Jehin et al. 2018). I utilised the standard observing modes of the LHATPRO instrument (located ~ 200 m above and a 1.8 km lateral distance away from SSO), in addition to ground relative humidity and temperature measurements. These measurements have been used to estimate the PWV experienced by SSO observations by including an altitude difference correction and line-of-sight estimate. I have also assessed the impact of photometric contamination by the temporally-varying atmospheric PWV in several commonly used filter bandpasses, extending the work on the PWV correction described in Murray et al. (2020). In the following sections, I describe the methodology, quantify the extent of PWV variability at Paranal, its effect on common red-visible – near-infrared bandpasses, and evaluate our correction method on photometric observations from SSO performed with the $I+z'$ bandpass.

3.2 Method of correction

The observed flux of an object through our atmosphere can be described as:

$$f(X, \text{PWV}, T_{\text{eff}}, t) = \int W(\lambda, X, \text{PWV}, t) R(\lambda) S(\lambda, T_{\text{eff}}, t) d\lambda, \quad (3.1)$$

where $W(\lambda, X, \text{PWV}, t)$ is the atmospheric transmission as function of wavelength (λ), airmass (X), and PWV value with time (t). $R(\lambda)$ is the overall bandpass response as a function of wavelength. $S(\lambda, T_{\text{eff}}, t)$ is the flux density distribution of an observed object as a function of wavelength, effective temperature (T_{eff}), and time.

In Figure 3.1, the amount with which different PWV values affect visible and near-infrared atmospheric transmission is shown. When Equation 3.1 is applied, the change of flux as a function of PWV can be observed to differ significantly as a function of effective temperature; this is illustrated in Figure 3.2 for the main bandpasses used by SPECULOOS (r' , i' , z' , and

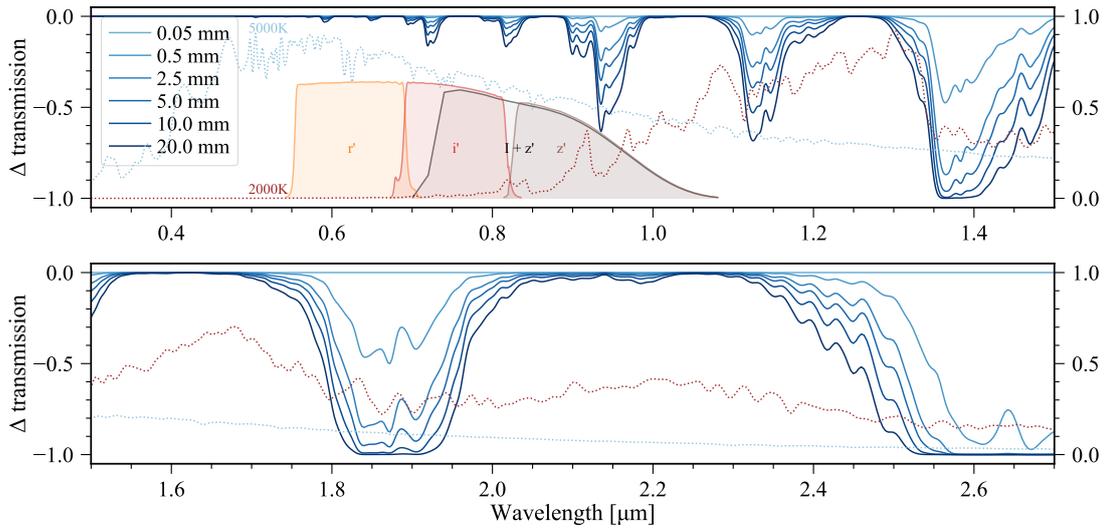


Figure 3.1: The fractional change of atmospheric transmission (left-hand axis) from the visible to near-infrared, at airmass 1, from a PWV of 0.05 mm to 20 mm is shown via a series of low resolution atmospheric spectra (from the SkyCalc Sky Model Calculator, for 2400 m; Jones et al. 2013; Noll et al. 2012). 5000 K and 2000 K stellar spectra are superimposed as dotted lines, from PHOENIX BT-Settl models (Allard et al. 2012). Four observational bandpasses (right-hand axis), r' , i' , z' , $I+z'$, with instrumental efficiencies of a telescope equipped with a deeply-depleted CCD accounted for are also shown.

$I+z'$). For differential photometry, a second-order differential effect is thus induced when PWV is time-variable.

One can divide out this effect by modelling the expected differential light curve produced by PWV variability, as per Equation 3.1. The difficulty in this approach stems from acquiring line-of-sight PWV values, in addition to having representative flux density distributions of the objects observed.

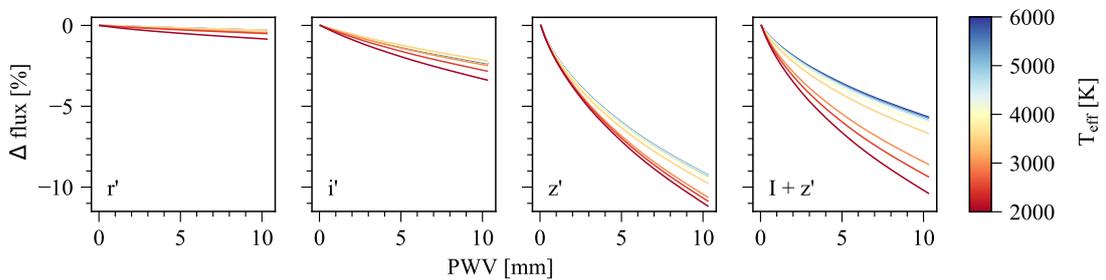


Figure 3.2: The change in flux as a function of PWV for different-temperature stars (from 6000 K to 2000 K in steps of 500 K) as modelled through four observational bandpasses, r' , i' , z' , and $I+z'$ (profiles shown in Figure 3.1), with respect to a 0.05 mm PWV atmosphere at airmass 1.

3.2.1 The PWV grid

For each filter bandpass I have developed a grid which outputs a time-independent value from Equation 3.1, when fed in observational parameters. These parameters are airmass (between 1 – 3), effective temperature (between 2000 K – 36500 K), and PWV (between 0.05 mm – 30 mm).

Each grid was built using atmospheric transmission profiles (0.5 μm – 2.0 μm) from the online SkyCalc Sky Model Calculator by ESO (Jones et al. 2013; Noll et al. 2012), for 2400 m (the closest available altitude to SPECULOOS’ site, using the ‘Entire year’ ‘Entire night’ profiles). In addition to these profiles, PHOENIX BT-Settl stellar models (Allard et al. 2012) provided by the Python Astrolib PySynphot Python package (Laidler et al. 2008) were used, with 2000 K as the lowest available temperature in the package. To build a stellar spectrum, one requires three parameters: effective temperature, metallicity, [Fe/H], and surface gravity, $\log(g)$. A representative set of stellar models were built using the parameters from Pecaut & Mamajek (2013)^a, assigning a metallicity index of 0 to each spectra.

I generated 273 atmospheric profiles, a permutation of airmasses between 1 – 3 at 0.1 intervals, and PWV values between 0.05 mm – 30 mm ([0.05, 0.1, 0.25, 0.5, 1.0, 1.5, 2.5, 3.5, 5.0, 7.5, 10.0, 20.0, 30.0] mm). I then included a set of 91 stellar spectra to produce a grid of 24,843 $f(X, \text{PWV}, T_{\text{eff}}, t)$ values from Equation 3.1 to interpolate in between. I have made the PWV grid code, Umbrella, publicly accessible on GitHub^b.

3.2.2 PWV measurements – estimating line-of-sight PWV

I estimated line-of-sight PWV values by linearly interpolating over airmass between two LHATPRO measurements, one at zenith and another at an airmass of 2 (altitude of 30°) from its cone scan measurements. Cone scan measurements were an average of 4 measurements at different bearings at a fixed altitude of 30°, with its value scaled as if observed at zenith. Zenith and cone scan values were measured at ~2 minute and ~15 minute cadences respectively. The LHATPRO also produces an all-sky scan value every ~6 hours, but these were not considered here.

LHATPRO PWV measurements, for the period of this study, were acquired via ESO’s Ambient Query Form^c in two stages. The first stage was acquired before the online archive was updated on 2020 Aug 28, the second stage was acquired after the update. Before the update, the data downloaded included zenith, cone scan, and all-sky measurements. Unfortunately

^aSpecifically, the updated values from https://www.pas.rochester.edu/~emamajek/EEM_dwarf_UBVIJHK_colors_Teff.txt version 2019.3.22, where the 1960 K effective temperature was rounded to 2000 K.

^b<https://github.com/ppp-one/umbrella>

^chttps://archive.eso.org/wdb/wdb/asm/lhatpro_paranal/form

for these measurements, a running average over 5 measurements and a ~ 1 minute smoothing time-average was done by the service prior to downloading from the online archive. As a result, the different measurement types were not labelled or regularly spaced to be easily differentiable. To differentiate between the observing modes I used a peak detection method. Large peaks spaced at ~ 6 hours intervals (assumed to be all-sky measurements) were first removed, then peaks spaced at ~ 15 minute intervals were registered as the average of 4 cone scan values and 1 zenith value. Thus, only cone scan measurements which were larger or equal to zenith measurements of this dataset were subsequently identified.

The data acquired after the update only included zenith measurements, without the issue of the 5 point running average and the ~ 1 minute time-average. The other measurement types were not made public at the time of acquisition. Thus, with the previously acquired data, and the new zenith values, I could estimate non-moving averaged cone scan measurements by using the new zenith measurements from around the same measurement period.

3.2.3 Accounting for altitude difference

The majority of water vapour resides close to the ground, with a scale height between 1 and 3 km (Kerber et al. 2017). The altitude difference of ~ 200 m between the LHATPRO at the VLT platform and SSO (located at the lower altitude) will therefore introduce an additional amount of PWV affecting our observations.

I estimated the missing vertical column of water vapour by integrating over the altitude difference, Δh , the estimated change in the density of water vapour, ρ (in kg/m^3), between the respective sites,

$$\text{PWV} = \int \rho \, dh \approx \frac{1}{2} \Delta h (\rho_{vlt} + \rho_{sso}). \quad (3.2)$$

To estimate the PWV, a linear change in water vapour density was assumed between SSO, ρ_{sso} , and the VLT platform, ρ_{vlt} , where the LHATPRO is located. This yields a value in kg/m^2 , equivalent to PWV in millimetres when liquid density of water is 1000 kg/m^3 .

To estimate the density of water vapour (Sensirion 2009), one can use ambient temperature, T in $^\circ\text{C}$, and relative humidity, RH , measurements from the respective sites:

$$\rho = 0.2167 RH \frac{6.112 \exp\left(\frac{17.62 T}{243.12 + T}\right)}{273.15 + T} f_w(P), \quad (3.3)$$

where $f_w(P)$ is “water vapour enhancement factor” as a function of pressure, P , in hPa:

$$f_w(P) = 1.0016 + 3.15 \times 10^{-6} P - 0.074 P^{-1}, \quad (3.4)$$

where pressure was assumed fixed over time for the respective altitudes.

I used existing temperature and humidity sensors from the respective sites to produce density estimates. At SSO, I used the temperature and relative humidity sensor (Sensirion SHT15) on board a Boltwood Cloud Sensor II, with an assumed accuracy of ± 1 °C and $\pm 4\%$ on relative humidity.^d At the VLT platform, values of temperature and relative humidity were measured by a VAISALA METEORological station 2m above the platform, with a quoted accuracy of ± 0.2 °C and $\pm 1\%$ respectively (Sandrock et al. 1999).

The altitude difference was determined using Google Maps Elevation service (Google 2022), with LHATPRO's position on the VLT platform returning an altitude of 2633 m and SSO an altitude of 2446 m. For comparison, the altitude of VISTA's platform from Google Maps Elevation service was found to agree with ESO's stated value within 2 m. GPS-derived altitude values are also available at SSO, yielding a median value of 2482 m. The neighbouring facility to SSO, NGTS, has a quoted altitude of 2440 m (Wheatley et al. 2018b), which led us to disregard the GPS-derived value. It was decided to use the altitude difference given by Google Maps Elevation service of 187 m, with an assumed error of ± 10 m. The lateral distance between the LHATPRO and SSO of 1.8 km was ignored.

3.3 Results and discussion

The following sections detail the results of applying the PWV correction to differential light curves observed by SSO over the course of approximately a year (2019 Feb 17 – 2020 Jan 31).

3.3.1 PWV variability at Paranal

We experienced a median zenith value of 2.3 mm, a median cone scan value of 2.9 mm (pre-scaled by the LHATPRO service, as if observed at zenith), and a median PWV value of 0.26 mm calculated from the altitude difference, throughout our nightly observations.

During our study, I assessed the absolute difference between 30° (airmass 2) cone scan values and zenith values, as shown in Figure 3.3. I found a median percentage difference of 26%, as shown by the 0.5 CDF mark on Figure 3.3. However, due to the nature of retrieving LHATPRO data (see Section 3.2.2), the majority of the detected cone scan values were always to be larger than zenith values, and as such the variation during this period may be different in reality.

A similar difference was observed in further detail in Querel & Kerber (2014), where they analysed 21 months of periodic all-sky scans, performed by the LHATPRO every 6 hours. From

^dAccuracy taken from the sensor's datasheet: <https://sensirion.com/us/products/catalog/SHT15/> for the low humidity conditions seen in Paranal, Chile.

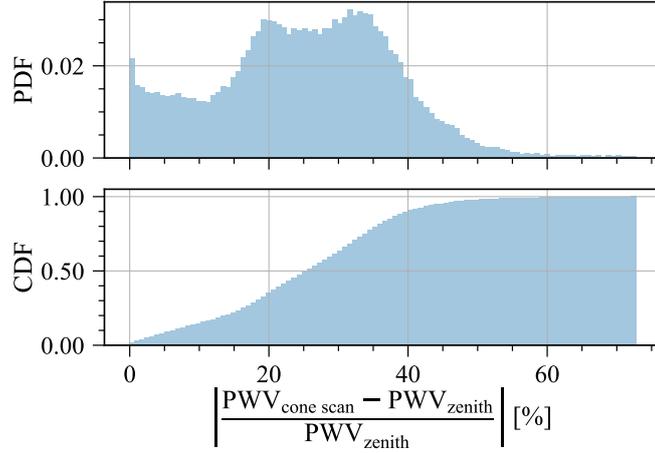


Figure 3.3: Probability density function (PDF) and cumulative distribution function (CDF) of the percentage difference between cone scan and zenith PWV measurements. Data from all publicly available nights, taking only dusk to dawn measurements.

Table 3.1: Following Figure 3.4, total proportion (in percent) of consecutive blocks of one hour of PWV zenith data, only between the hours of dusk to dawn, which display a maximum change to a light curve’s flux (ΔLC) in mmag, per bandpass.

| ΔLC [mmag] | r' | i' | z' | $I + z'$ |
|--------------------|-------|-------|-------|----------|
| 0 – 1 | 100.0 | 100.0 | 84.23 | 62.99 |
| 1 – 2 | 0.0 | 0.0 | 11.72 | 26.10 |
| 2 – 5 | 0.0 | 0.0 | 3.84 | 9.72 |
| 5 – 10 | 0.0 | 0.0 | 0.20 | 1.06 |
| > 10 | 0.0 | 0.0 | 0.00 | 0.12 |

Table 3.2: Proportion (in percent) of transit-like structures over the evaluated dataset, within a range of depths in mmag, induced by PWV variability at zenith, per bandpass.

| Depth [mmag] | r' | i' | z' | $I + z'$ |
|--------------|-------|-------|------|----------|
| 0 – 1 | 10.17 | 10.17 | 8.71 | 6.35 |
| 1 – 2 | 0.00 | 0.00 | 0.98 | 2.77 |
| 2 – 5 | 0.00 | 0.00 | 0.49 | 0.90 |
| 5 – 10 | 0.00 | 0.00 | 0.00 | 0.16 |
| > 10 | 0.00 | 0.00 | 0.00 | 0.00 |

their study, they found a median all-sky variation of 10–26% peak-to-valley, down to 27.5° , with respect to the all-sky’s zenith value. From this, they argued that zenith observations of PWV are sufficient for the general analysis and correction of astronomical data at Paranal. However, with the increasing interest in cooler stars, photometric surveys may rely on line-of-sight PWV observations, depending on their desired photometric accuracy.

Figure 3.4 and Table 3.1 detail the extent to which common near-infrared bandpasses (for silicon based detectors) will induce second-order differential variability with PWV changes

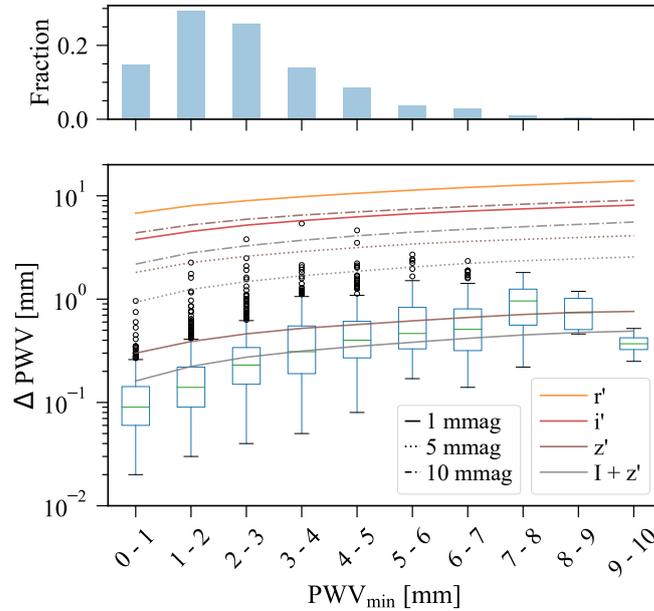


Figure 3.4: Sensitivity of common bandpasses and the $I+z'$ bandpass to PWV changes. Bottom: Measured PWV changes from evaluated consecutive blocks of one hour intervals of PWV zenith data (only between the hours of dusk to dawn). Here, I calculated the max-min change (Δ PWV) of PWV within each evaluated hour. Grouped the measured changes into 1 mm intervals of PWV_{\min} , and plotted the respective box plots of the PWV changes from PWV_{\min} in each interval, where each box plot shows the standard median (green line), inter quartile range (IQR) of the lower (Q1) and upper (Q3) quartiles, and lower whisker ($Q1 - 1.5 \cdot IQR$) and upper whisker lines ($Q3 + 1.5 \cdot IQR$), with black circles denoting outliers of the whisker range. The coloured lines show the required Δ PWV from the middle of each interval to induce a 1 mmag, 5 mmag, and 10 mmag change in a light curve (when observing a 2700 K target star and a 5200 K comparison star), represented by the solid, dotted, and dash-dotted lines respectively for each r' , i' , z' , and $I+z'$ bandpasses, in orange, red, brown and grey respectively. Changes at 5 mmag and 10 mmag for r' and i' were outside the model's PWV range of 30 mm. Top: The fraction of hours contributing to the evaluated PWV interval ranges.

experienced in Paranal during this study. Here, the respective change in a light curve's flux, per bandpass, of observing a late M-dwarf (with a temperature of 2700 K and a 5200 K comparison star – median temperatures of the sample in this paper) were simulated, using nightly PWV zenith data from the LHATPRO service. The $I+z'$ bandpass was found to be the most sensitive to PWV variability, followed by z' , i' , and r' . On the hourly scale, the approximate expected transit duration for a temperate, rocky planet around a late-M or L-type star (Traub & Cutri 2008), the majority of PWV variability was found to induce a 0 – 1 mmag change in a light curve, within the typical photometric precision achieved by SSO. However, with z' , and more significantly with the $I+z'$ bandpass, changes greater than 1 mmag were observed at 16% and 37% occurrence rates respectively, mimicking transit structures from time to time.

The accuracy on Δ PWV was calculated to be 0.042 mm, propagated from LHATPRO's stated single-measurement precision of 0.03 mm, assuming the accuracy of better than 0.1 mm of a single measurement to be systematic. For the $I+z'$ bandpass, this accuracy level limits the correction to 0.4 mmag when PWV_{\min} is at 0.05 mm (the lower limit of the atmospheric models used), 0.2 mmag at Paranal's median value of 2.3 mm, and 0.1 mmag at a rarely seen 10 mm. This therefore suggests PWV measurements from the LHATPRO PWV measurement are sufficient for correcting sub-mmag-level changes in our $I+z'$ light curves, and even more so in the other band passes considered here.

The proportion of transit-like structures induced by zenith PWV variability over the evaluated dataset was estimated in Table 3.2, using the same data as Table 3.1. This was calculated by finding the proportion of consecutive hours that displayed a dip in one hour and followed an equivalent rise, within $\pm 25\%$, in the following hour. This method may miss some structures that occurred within an evaluated hour. Nonetheless, it suggests the occurrence rate of a significant transit-like structure, greater than ~ 5 mmag, will not occur for the r' and i' bandpasses, and is a rare occurrence for both the z' and $I+z'$ bandpasses. The $I+z'$ bandpass will display a greater depth, given it has approximately twice the sensitivity to PWV changes than the z' bandpass when observing late M and L-type stars. If such an event occurs, it is likely to be visible in co-current observations of similar temperature targets by other telescopes on site in the same bandpass, if they exist.

SSO is subject to additional PWV-induced effects due to the altitude difference and observing through a variety of airmasses. The additional previously unaccounted for PWV from the altitude difference between SSO and the LHATPRO, calculated with the method described in Section 3.2.3, was seen to generally scale with zenith PWV, however was not seen to have a direct relationship with zenith values. The derived altitude difference PWV values were dependent on the accuracies of the respective temperature and humidity sensors, and the accuracy of the altitude difference in meters. Propagating the quantified errors on the derived altitude

difference PWV dataset, provided a median fractional error of 18%. The fractional error is likely higher in reality due to the approximation made in Equation 3.2. The effect of the altitude derived PWV is evaluated in the following sub-section.

Similarly, the validity of linearly interpolating cone scan values over time, as described in Section 3.2.2, from a 15 minute to a 2 minutes time base (to match the cadence of zenith measurements), was assumed to be acceptable. To test this assumption, I used the ~ 2 minute cadence zenith values, and every 7th value from the dataset linearly interpolated back onto the same time base as the original zenith observations, to simulate the cadence of the cone scan measurements. The difference between the values were found to be within $\pm 2\%$, under 0.1 mm in most cases. The validity of interpolating between zenith and cone scan values over airmass is addressed in the following sub-section.

3.3.2 Evaluation of the PWV correction

I used observations made by SSO with its primary bandpass, $I+z'$. SSO, with its four telescopes, made 1193 unique observations (divided by telescope/target/night) of 103 targets observed between 2019 Feb 17 to 2020 Jan 31 with the $I+z'$ bandpass. This amounted to a sum of 5420 hours of on-sky data. Differential light curves of these observations were produced with the SSO pipeline described in [Murray et al. \(2020\)](#).

For the correction, the target star was assigned an effective temperature derived in [Sebastian et al. \(2020a\)](#), with an assumed systematic error of ± 100 K. The range of target temperatures evaluated was 2000 – 3000 K, with a median temperature of 2700 K. The comparison light curve, behaving as an artificial star, had an effective temperature assigned by the SSO pipeline – a weighted sum of effective temperatures from the Gaia DR2 ([Gaia Collaboration et al. 2018](#)) catalogue, of the stars used in the field. The range of temperatures for the artificial star evaluated was 4000 – 6000 K, with a median temperature of 5200 K. For the $I+z'$ bandpass, the second-order effect induced by the artificial star behaves very similarly for any temperature above 4000 K, as illustrated in Figure 3.2. The target star was often the coolest star in the field, and as such, forming an ideal artificial star with a temperature equivalent to the target star was not possible. Consequently, for the fields evaluated, this disparity in temperature between the target star and the effective comparison star would always yield a PWV induced effect under changing atmospheric conditions. One could potentially further improve the correction by fine-tuning the assigned effective temperature of the target, or by using the real spectra of the target – this was not considered here.

When searching for transits, we seek to minimise any atmospherically induced variability on differentially resolved light curves. This is to maximise the likelihood of detecting real transit events. In this context, the observed dataset was evaluated in two regimes – with a

Table 3.3: Evaluation of the RMS percentage change, $((\sigma_{\text{LC corrected}} - \sigma_{\text{LC}})/\sigma_{\text{LC}}) \times 100$, of 103 global light curves (LC) with low and high pass filtering at 120 minutes. Showing the [10,25,50,75,90]th percentiles using the respective PWV derived values for correction, from zenith, zenith + altitude difference (ΔAlt), estimated line-of-sight (Est. LoS), and Est. LoS + ΔAlt .

| Percentiles [%] | Low pass [%] | | | | High pass [%] | | | |
|-----------------|--------------|-----------------------------|----------|-------------------------------|---------------|-----------------------------|----------|-------------------------------|
| | Zenith | Zenith + ΔAlt | Est. LoS | Est. LoS + ΔAlt | Zenith | Zenith + ΔAlt | Est. LoS | Est. LoS + ΔAlt |
| 10 | -0.2 | -0.2 | -0.5 | -0.1 | 1.7 | 1.7 | 4.3 | 4.0 |
| 25 | -14.4 | -14.4 | -14.3 | -16.0 | 0.2 | 0.2 | 0.4 | 0.5 |
| 50 | -51.8 | -52.9 | -52.7 | -53.8 | -1.2 | -1.3 | -1.2 | -1.1 |
| 75 | -63.8 | -64.5 | -65.6 | -66.2 | -4.3 | -4.3 | -4.0 | -3.9 |
| 90 | -71.9 | -72.1 | -73.0 | -73.5 | -9.1 | -8.6 | -6.6 | -6.1 |

low and high pass temporal filter, with a dividing period of 120 minutes. Where the low pass temporal filter maintained variability that was greater than 120 minutes, and high pass filter maintained the variability shorter than 120 minutes. This was to demonstrate the correction's effect around and below the transit timescale, and on timescales where long-term variability may be induced. Table 3.3 details the effectiveness of the correction. Here, I evaluated the percentage change in the root-mean-square (RMS) on global light curves (30 minute binned) observed by SSO before and after the correction for the low and high pass temporally filtered domains. The correction was calculated with combinations of each of the PWV derived values (zenith, estimated line-of-sight, and altitude difference).

On timescales longer than 120 minutes, a zenith PWV based correction demonstrated a large median RMS percentage change of -51.8% , when compared to an uncorrected differential light curve. This large percentage change in RMS is attributed to the long-term multi-millimetre variability of PWV, and as such demonstrates the correction's importance for variability studies on late-M and L-type stars, when using a bandpass like $I+z'$. The addition of the altitude difference derived PWV to both zenith and estimated line-of-sight PWV made a marginal improvement, at -52.9% and -53.8% respectively. Using the estimated line-of-sight PWV alone yielded -52.7% .

Whilst the estimated line-of-sight PWV presented a marginally better median percentage change, the assumption of circular symmetry around zenith in some instances may introduce false variability post-correction. To yield a more accurate correction, a new measurement mode for the LHATPRO is subsequently suggested: a continuous monitoring mode which constructs a low resolution all-sky map, maintaining a PWV accuracy of 0.1 mm, at a cadence better than 30 minutes (half the expected transit duration around a late M, L-type star, to follow Nyquist sampling). Since we only seek to minimise long-term variability and false transit features, the current 2 minute cadence for zenith is faster than we currently need.

The estimated line-of-sight PWV also minimised a previously uncorrected airmass effect: an observed decrease of differential flux when transitioning to higher airmasses. This effect is

illustrated in Figure 3.5. Without the correction, a -3.0 mmag differential flux change from the median differential flux at airmasses between 1.0 and 1.2 to the median differential flux at airmasses between 1.8 and 2.0 was observed. This change minimises slightly with the correction derived with zenith PWV, and noticeably improves with the estimated line-of-site PWV with a median differential flux increase of 0.1 mmag observed from airmasses between 1.0 and 1.2 to airmasses between 1.8 and 2.0, versus the -3.0 mmag observed without line-of-sight correction. I performed a two-sample Kolmogorov–Smirnov test to assess the statistical significance of this result, where I assessed the cumulative distributions of differential flux at high airmasses (1.8 – 2.0) corrected with zenith PWV values (Figure 3.5, middle plot, orange CDF), $F(x)$, to the correction derived from estimated line-of-site PWV values (Figure 3.5, bottom plot, orange CDF), $G(x)$. Our null hypothesis being that $F(x) \geq G(x)$ for all airmasses. The resulting Kolmogorov–Smirnov test yielded a p-value of 0.992, leading us to accept the null hypothesis, confirming our observation that line-of-sight PWV data helps to minimise the decreasing of differential flux at higher airmasses. From the distributions seen in Figure 3.5, we can also see the reduced spread of differential flux when the correction is applied, as quantified in Table 3.3.

On timescales shorter than 120 minutes, a median RMS percentage change of around -1% for all PWV derived values was observed. However, without the ground-truth of stellar variability at this scale, it’s difficult to argue if one PWV derived correction is better than another when solely based on the RMS percentage change. Whilst it has been seen to correct transit-like features in Murray et al. (2020) and one example in the next sub-section, the correction at shorter time scales was seen to increase the RMS in about a quarter of the instances evaluated. Once again, without the ground truth at this scale, or extended trend modelling, it’s currently difficult to validate the true extent of improvement at this scale. It is therefore recommended to manually vet where possible. Nonetheless, with the dataset evaluated here, we have not seen instances where the correction has induced significant false variability. The cases where this may occur would be likely due to PWV measurement failure, or significant line-of-sight differences from observation and PWV measurement.

3.3.3 PWV correction in action

In this sub-section, examples of the correction with the $I+z'$ bandpass are shown, on the global and nightly scale. The first example, in Figure 3.6, shows five consecutive nights of observations of three targets of similar temperature (2500 - 2700 K), with the global light curves normalised over the assessed period. Between the first and second night, a large PWV change was observed. The resulting second-order effect is evident in the uncorrected light curves, showing around a 30 mmag change for all the targets. The behaviour on shorter timescales, however, was not

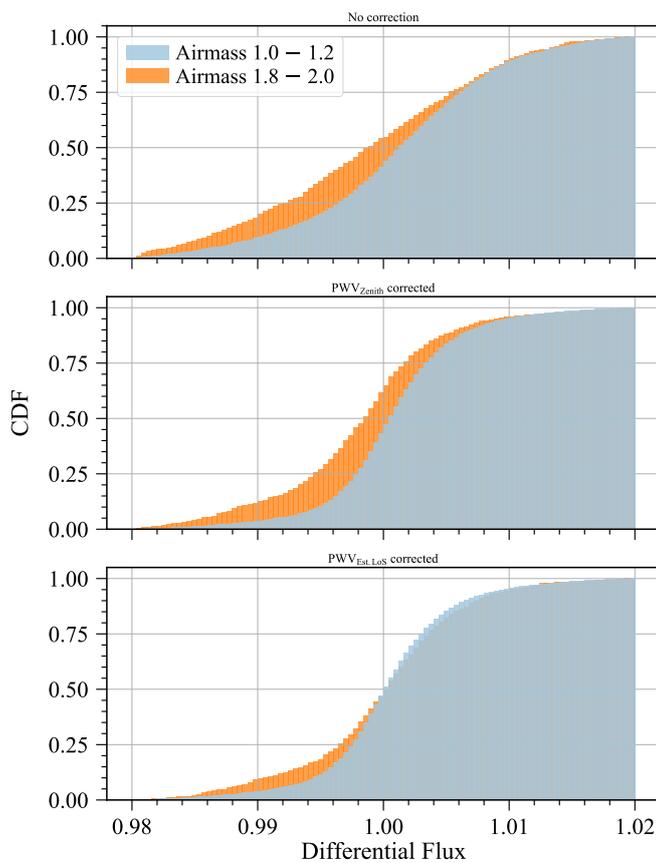


Figure 3.5: Assessing the correction’s effect on differential flux at different airmass ranges. Cumulative distribution function (CDF) histograms of differential flux between 0.98 and 1.02 (of 0.005 JD binned on-sky data) observed with the $I+z'$ bandpass. Global dataset assessed for when observations were in two different ranges of airmasses 1.0 – 1.2 and 1.8 – 2.0, with different levels of PWV derived corrections applied. First without any correction (top plot), zenith PWV correction (middle plot), and then estimated line-of-sight (Est. LoS) PWV correction (bottom plot).

always comparable between targets due to the respective stellar variability that is often seen with M dwarfs (Günther et al. 2022). Beyond the second night on Figure 3.6, the PWV changes were less significant on the light curves by eye. However, a level of difference is still visible.

The respective PWV values from zenith and cone scan followed a similar trend for the five nights. A similar behaviour was seen for the remaining dataset. The altitude difference derived PWV stayed relatively constant with the exception of the third night, where a small increase was observed at the end of the night, with an opposing change seen at zenith around the same time. A layer of water vapour likely transitioned from above the LHATPRO to the layers below it. A sharp dip in flux in all the light curves was observed at the same time. The correction removed the majority of the structure, however some residual in the shape was left in all the

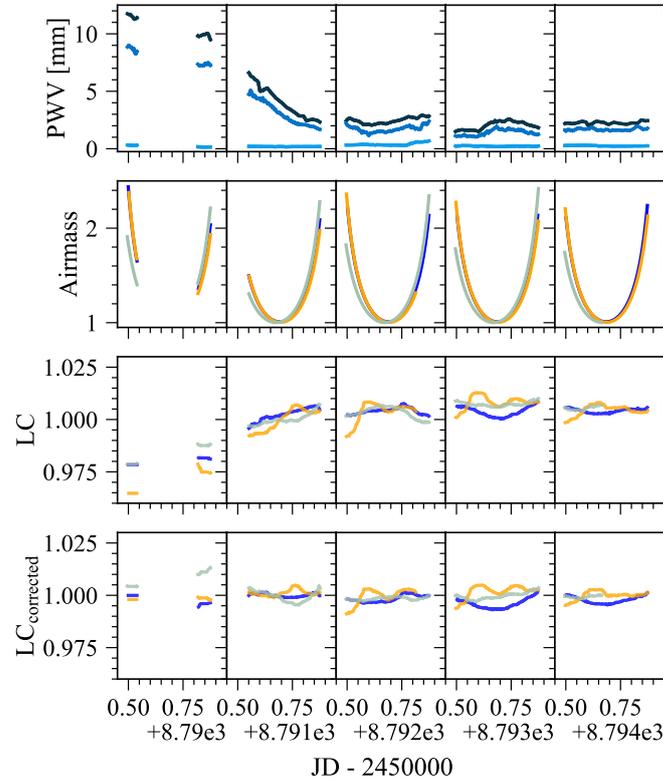


Figure 3.6: Five consecutive nights observed by three SSO telescopes of three different targets, one per telescope - distinguished by the different colours in the lower three plots. The first night is partially missing data due to bad weather. First row: PWV values from cone scan (dark blue), zenith (blue), altitude measurements (light blue) are shown. Second row: Airmass of respective observations, to show the transition between zenith and the airmass cone scan PWV measurements are made (airmass 2). Third row: Differential light curves (LC) without PWV correction. Bottom row: Differential light curves with PWV correction, using estimated line-of-sight + altitude difference PWV. The global light curves were normalised over the assessed period. The effective temperatures of the target/comparison stars were 2500/5000, in orange; 2600/4700, in light green; 2700/4300, in blue, respectively. The light curves shown are the trends from 120 minute windows using a median method, where the average 120 minute binned error was sub-mmag.

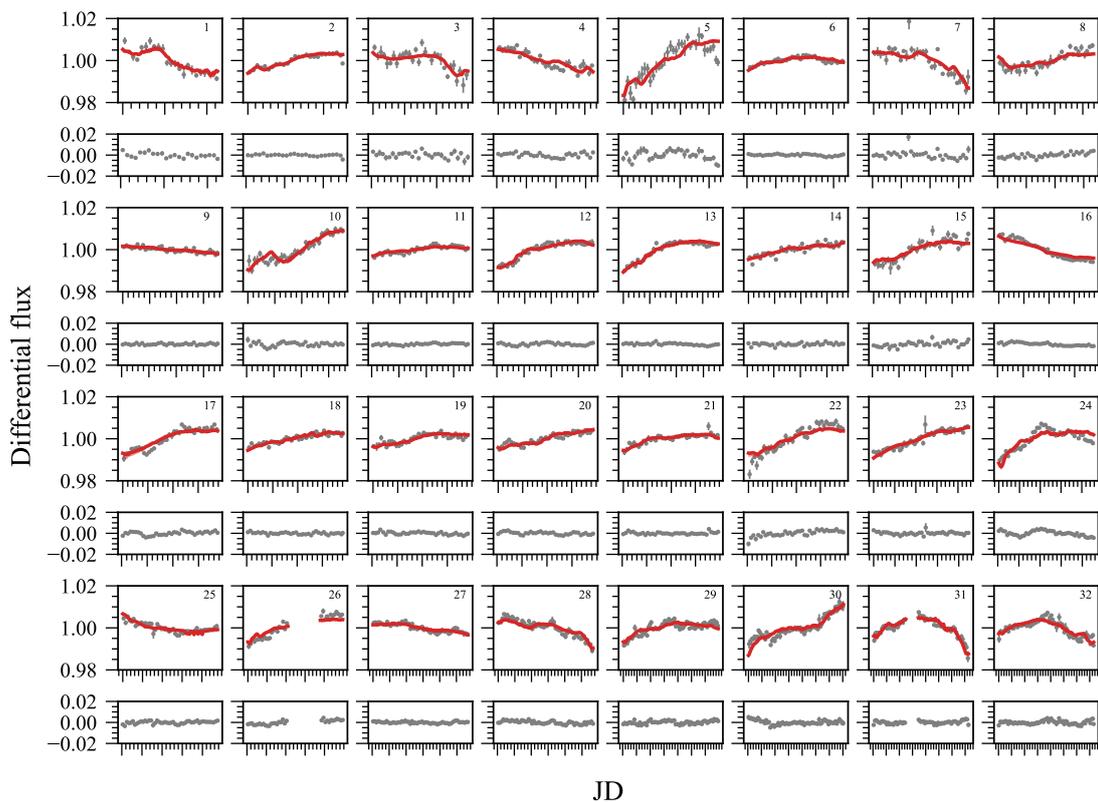


Figure 3.7: A selection of 32 observations of 19 different targets observed with the $I+z'$ bandpass which matched closely with the expected trend from the correction. Top row: 0.005 JD binned (7.2 minutes) uncorrected differential flux light curves in grey. Expected trend from the correction in red, using estimated line-of-sight + altitude difference PWV and knowledge of the target and comparison stars effective temperatures. In shaded red, although the effect is not visible for the majority of examples, the expected trend from the correction also plotted using ± 100 K from the target's effective temperature. Second row: Residual of the correction (observed data - expected trend) of the above subplot. Row order then repeats. Ordered from shortest to longest timescales, where the major ticks on the x-axis are 0.05 JD (72 minutes).

light curves, suggesting some amount of water vapour was unaccounted for in the line-of-sight and altitude estimates. The residual shape could likewise be attributed to inaccuracies in the target and comparison stars temperature estimates, where a higher target temperature estimate would have under corrected the variability. This could be similarly argued for the light green target between the first and second night. Co-current observations of this target could not be found to rule out stellar variability.

If one were to adopt the methodology in [Irwin et al. \(2011\)](#), briefly described in Section 3.1, for timescales greater than 120 minutes, one would need at least 9 co-current observations of similar temperature objects to yield a common mode with variability at the sub-mm scale at

30 minute bins. At timescales below 120 minutes, however, one would need an unrealistic 90 co-current observations. This was calculated by evaluating the overall RMS SSO experienced (with correction) at the respective scales, and assuming a RMS scaling of root sum the total number of co-current observations. With SSO having only 4 telescopes, this method was thus deemed not suitable.

In Figure 3.7, 32 examples of single observations are shown. In these examples, the uncorrected differential light curves displayed a close resemblance to the expected trend modelled by the PWV grid, with PWV values from the estimated line-of-sight and altitude difference. These were found by finding the nights where the standard deviation of the night's data was significantly reduced by the correction. Such examples often occur on quiet targets, where the second-order effects are very evident on both short and long timescales. For example, on the shorter timescales, false transit features have been mostly corrected for in subplots labelled 1, 3 and 12. On the longer timescales, a range of other variabilities are closely matched, such as an inverse airmass like shape in the subplot labelled 31. There is an instance in Figure 3.7 where a transit like feature is induced by the correction (subplot labelled 10). The exact origin of this feature is unknown, most likely a line-of-sight induced feature. Manual vetting is therefore recommended when such events occur. The remaining light curves matched closely the expected trend modelled by the PWV grid, with many examples beyond the 32 presented light curves which similarly match the modelled trend.

3.4 Conclusions

I have developed a method of modelling and mitigating the second-order effect induced by PWV on time-series photometric data. This has been enabled by leveraging the accurate measurement modes provided by an onsite radiometer, the LHATPRO, and local environmental sensor data. The developed tool, the PWV grid, has proven to be an essential for SSO, and we believe it can help other studies who are likewise sensitive to PWV and have access to accurate PWV data. The PWV grid code, Umbrella, has been open-sourced on GitHub^e.

I found, for removing transit-like structures and long term variability on late M- and early L-type stars, the LHATPRO's single measurement PWV accuracy of better than 0.1 mm, and precision of 0.03 mm, is sufficient to eliminate sub-mmag level PWV induced photometric effects for the $I+z'$ and z' bandpasses, and more than sufficient for the i' and r' bandpass. The $I+z'$ bandpass was shown to be exceptionally sensitive to second-order effects induced by PWV, and without aid of the correction, the bandpass significantly limits ones ability to do variability studies on late M, L-type stars. On the transit timescale, the bandpass is sensitive to

^e<https://github.com/ppp-one/umbrella>

variability which may mimic transit-like structures on the rare occasion with Paranal's level of PWV variability.

PWV data from zenith was found ample to support the majority of the correction needed for the four telescopes at SSO. However, through our use of zenith and cone scan measurement modes, there are residual second-order effects induced at higher airmasses which would require line-of-sight measurements to accurately correct for. I have therefore recommended a continuous all-sky observing mode for the LHATPRO, such to support more accurate line-of-sight estimates for our multiple telescopes at Paranal.

The additional PWV derived from the altitude difference between the LHATPRO and SSO was accounted for through the use of local environment sensors, and was shown to improve the correction on timescales longer than 120 minutes. On shorter timescales, a more accurate method of accounting for the altitude difference may be needed. If one does not have access to PWV data, then optimising the bandpass for the survey is necessary.

FEASIBILITY STUDY

In this chapter I evaluate the feasibility of introducing new NIR instrumentation to SPECULOOS. One which would permit better photometric precision than the existing CCD Si based instrumentation, and likewise minimise the induced effects of PWV variability. Here, I assess and compare the expected photometric precision of different commercially available detectors with standard and custom NIR filters. Each system was optimised in the context of observing the coolest objects in SPECULOOS' target list ([Sebastian et al. 2020a](#)), whilst taking into account PWV variability discussed in Chapter 3. Finally, I evaluate the optimal field of view to perform differential photometry for our targets using catalogues from the 2MASS ([Skrutskie et al. 2006](#)) and Gaia surveys ([Gaia Collaboration et al. 2018](#)).

4.1 Fundamentals

First, we must understand the elements which contribute to modelling photometric precision of an observation. Here, we begin with the astronomical objects of interest, which can radiate light from all regions of the electromagnetic spectrum. For cool stars (< 2700 K), their radiation predominantly resides in the NIR regions. Once their light reaches the Earth's atmosphere, its flux has reduced as expected by the inverse square law, excluding the potential absorption from interstellar medium. It then traverses and is attenuated by the atmosphere's composition till it finally reaches our telescope.

Figure 4.1 shows the overall system efficiencies for the detectors in consideration here, after accounting for all the instrumental efficiencies that are present in one of SPECULOOS'

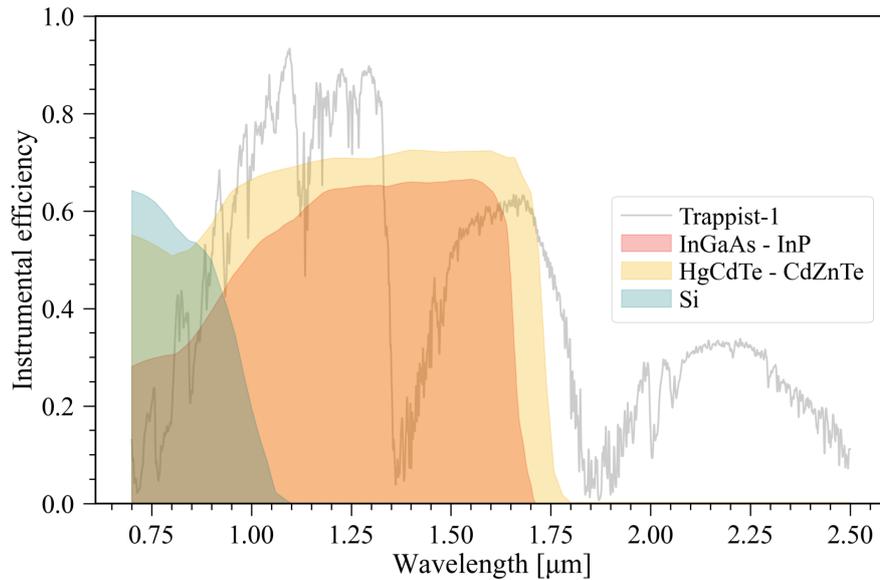


Figure 4.1: Overall system efficiency as viewed with an InGaAs-InP CMOS detector (shaded red-orange), with a HgCdTe-CdZnTe CMOS detector (shaded yellow), and with a deeply-depleted Si based CCD (shaded blue). The overall efficiency consists of the detector’s quantum efficiencies (sourced from manufacturers of presented detectors), the aluminium coated mirror reflectance of primary and secondary mirror, and the fused silica transmission of the two correcting lenses present on a SPECULOOS telescope. A Trappist-1 spectrum (grey line) obtained with the SpeX instrument (Bardalez Gagliuffi et al. 2019) with the median atmospheric transmission at Paranal applied, superimposed for reference.

telescopes. A proportion of the incoming flux is lost at each stage of the telescope’s optical design due to the respective efficiencies. Light is first reflected off the primary and then secondary mirror, before travelling through a pair of correcting lenses made from fused-silica. The light then transmits through a filter (not shown in Figure 4.1) and then finally another fused-silica element, the camera window, before being absorbed by the detector. At the detector, the remaining photons are converted into electrons, as dictated by the detector’s quantum efficiency, for the ROIC to digitise. Figure 4.1 also shows the spectra of Trappist-1, where one can see the significant increase in flux as we progress into the NIR. Trappist-1 has an effective blackbody temperature of 2560 K (Gillon et al. 2017). SPECULOOS’ target list (Sebastian et al. 2020a) includes a large proportion of targets below this temperature, the spectra of these targets shifts towards redder wavelengths the cooler they are, as illustrated in Figure 4.2.

SPECULOOS utilises a deeply-depleted CCD (as described in Chapters 1 & 2 and shown in blue in Figure 4.1) in their cameras. As is evident from Figure 4.2, there is a larger portion of flux from cool stars that is currently unobserved by the existing system due to the lack of the current detector’s quantum efficiency beyond $\sim 1 \mu\text{m}$. Thus, the route considered here

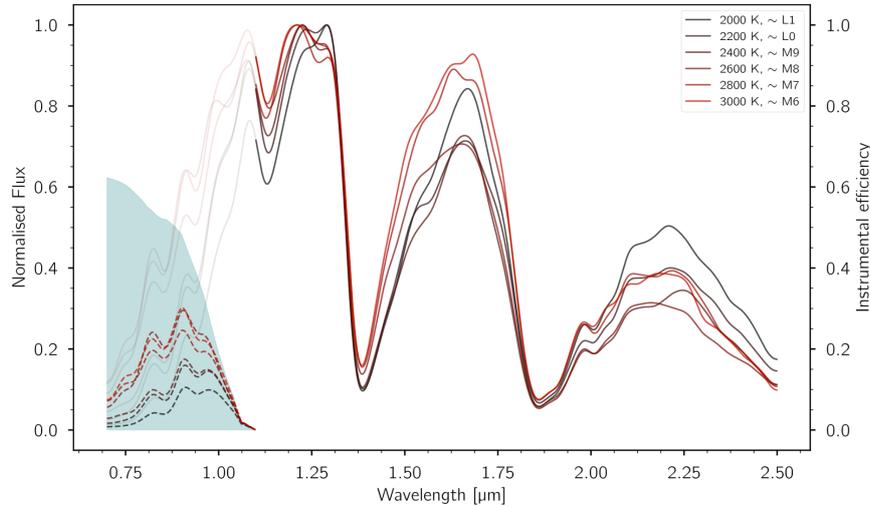


Figure 4.2: Example stellar spectra (solid lines), observed through the median atmosphere present in Paranal, of objects that are targeted by SPECULOOS - generated using the Phoenix BT-Settl atmospheric models through Astrolib PySynphot (Lim et al. 2015; Allard et al. 2012), and the Cerro Paranal Advanced Sky Model (Noll et al. 2012; Jones et al. 2013). They have been Gaussian smoothed to remove high resolution features to show the differences in flux more clearly. The spectra are normalised by their peak emission. Between 0.7 to 1.1 μm , the total instrumental efficiency present in SPECULOOS with the $I+z'$ filter (shaded blue) has been multiplied by the normalised spectra (dashed lines).

to increase the overall system efficiency is to evaluate the potential of another detector type sensitive in the NIR. However, as discussed in Chapter 1, there are consequences in moving into the NIR we must evaluate.

One could also consider changes to other aspects of the telescope to squeeze further efficiencies, such as using a silver or gold coating on the primary and secondary mirrors instead of the currently used aluminium. An operation to re-coat the mirrors however was not within the plan of this PhD. There are fixed aspects of the instrumentation that are not possible to change, such as the size of the mirrors. This thus places an upper limit on the amount of photons we can collect. The photometric precision will thus fundamentally be a function of the sky conditions, environment, the filter of choice, and the detector's properties, which are discussed in the following subsections.

4.1.1 Photometric precision

After an exposure, the total sum of electrons in each pixel of a detector will be composed of a number of components. Namely, the target's photoelectrons (N_{\star}), the sky/environmental photoelectrons ($N_{\text{sky}} + N_{\text{thermal}}$), and dark current electrons inherent to a detector (N_{dark}) –

all of which follow Poisson statistics and increase linearly with time. During the readout, a time-independent noise is added to an exposure, N_{read} , which is often defined as the standard deviation of a series of bias frames. Additionally, scintillation noise due to the Earth's turbulent atmosphere, σ_{scint} , has to be accounted for, where spatially and temporally varying refractive indices of the atmosphere adds additional noise to an observation. σ_{scint} is usually the dominant noise source for bright targets. The precision of observing the flux of a target thus follows the simple form given by [Nutzman & Charbonneau \(2008\)](#),

$$\sigma_{precision} = \frac{\sqrt{N_{\star} + \sigma_{scint}^2 + n_{pixel} \times (N_{sky} + N_{thermal} + N_{dark} + N_{read}^2)}}{N_{\star}}, \quad (4.1)$$

where n_{pixel} is the number of pixels that encircles the point spread function of the target. For aperture photometry, with the currently implemented instrumentation, SPECULOOS tends to use an aperture radius of 11.3 pixels ($0.35'' \times 11.3 = 4.0''$), equating to 401 pixels for n_{pixel} . Taking the form from [Osborn et al. \(2015\)](#), σ_{scint} can be assumed to follow

$$\sigma_{scint}^2 = 10^{-5} C_Y^2 D^{-4/3} t^{-1} (\cos \gamma)^{-3} \exp(-2h_{obs}/H) N_{\star}^2, \quad (4.2)$$

where, C_Y is an empirical coefficient (location dependent), D is the diameter of the telescope, t is the exposure time of the observation, γ is the zenith angle, h_{obs} is the altitude of the observatory and H the scale height of the atmospheric turbulence, which is generally accepted to be approximately 8000 m. All parameters are in standard SI units. The observatory altitude was set to 2440 m, C_Y was set to the median value of Paranal of 1.56 using the parameters determined by [Osborn et al. \(2015\)](#). The diameter, D , was set to 1 m. σ_{scint} is suggested to decrease slightly as one progresses further into the infrared, but a relationship with wavelength has not been established. There is a seasonal dependence on C_Y , as shown by [O'Brien et al. \(2021\)](#), where the scintillation noise is expected to improve in the winter months.

Finally, the target's photoelectrons N_{\star} can be calculated to be

$$N_{\star} = t \times \pi(D/2)^2 \times \int_0^{\infty} E(\lambda) T(\lambda) F(\lambda) \frac{\lambda}{hc} d\lambda, \quad (4.3)$$

where $F(\lambda)$ is the stellar flux density of the target, $E(\lambda)$ is the overall system efficiency shown in Figure 4.1, and $T(\lambda)$ is the transmission spectra of the atmosphere, as shown in the top of Figure 1.5. N_{sky} was similarly calculated using the radiance spectrum components shown in the bottom of Figure 1.5, multiplied by the field of view on the sky of a single pixel.

Table 4.1: Core detector parameters and Rough Order of Magnitude (ROM) cost as quoted by the respective manufacturers for ‘typical’ specifications in 2018, when this portion of the work was completed. I include the offerings from Teledyne, Princeton Infrared Technologies (PIRT), Raptor Photonics, with the last two offering complete camera packages (housing, cooling, data readout). Andor’s iKon-L-936 (part of the existing instrumentation - working primarily in the visible domain) is included for comparison. There are other manufacturers of HgCdTe and InGaAs based detectors (Bhan & Dhar 2019), but we restricted our assessment to those who could produce substrate removed detectors (to permit higher quantum efficiencies), cut off before 1.75 μm (to avoid major thermal noise), and were above $1\text{K} \times 1\text{K}$ in array size (to maintain a good balance of angular resolution and field of view - assessed in Section 4.5.3).

| | Teledyne | PIRT | Raptor | Andor |
|-------------------------------|------------------------------------|--------------------|--------------------|--------------------|
| Material | HgCdTe-CdZnTe | InGaAs-InP | InGaAs-InP | Si |
| Model | HxRG | 1280SciCam | Ninox 1280 | iKon-L-936 |
| DC [e/s] | ≤ 0.05 @ -150 °C | 300 @ -60 °C | 2000 @ -15 °C | 0 @ -100 °C |
| DC [e/s] @ -60 °C | ~ 6000 (15 μm) | 300 | n/a | 0.2 |
| Read noise [e] | ≤ 30 | 35 | 160/28 | 6.2 |
| Well depth [e] | $\geq 80\text{k}$ | 45k | 450k/10k | 64k |
| Array size | 1K, 2K, 4K | 1280×1024 | 1280×1024 | 2048×2048 |
| Pixel pitch [μm] | 10, 15, 18 | 12 | 10 | 13.5 |
| ADC bit depth | 16 | 14 | 12 | 16 |
| ROM [\$] | 0.3M (1K), 0.5M (2K) | 65k | 50k | 60k |

4.1.2 Available NIR detectors and camera systems

There are currently no NIR dedicated detector or camera systems available on the market which would permit a plug-and-play experience with our system at SPECULOOS. Of the available detectors/cameras, some would require extensive engineering and/or modification to permit operation and operation at appropriate operating temperatures for NIR observations. Each would at least demand considerable software development to work with our telescope control system, the connected data flow, and light curve production pipeline. Table 4.1 details some of the currently available detector and camera systems on the market.

The HgCdTe based solutions were priced significantly higher than the InGaAs offerings, as previously explained in Chapter 1. In addition to this, one would need to couple a HgCdTe detector to a cryogenic cooler solution, which would either require liquid nitrogen or a Stirling cycle cryocooler to maintain it at the temperatures needed for operation. Given the noticeable difference in price, the additional engineering, and engineering time that HgCdTe would require to bring to operation, we chose to continue with InGaAs for this feasibility study. To simplify the analysis, we will present the feasibility study on comparing the existing SPECULOOS silicon based instrumentation with the InGaAs based offering from PIRT. PIRT’s offering presents the best overall package of the InGaAs selection.

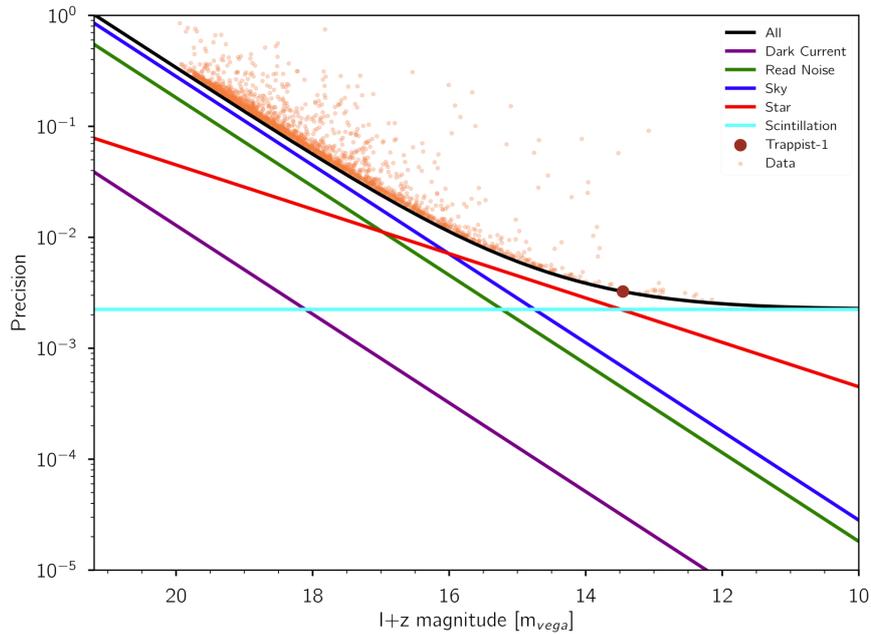


Figure 4.3: Precision of an observation with $I+z'$ for 10 s exposures as a function of stellar magnitude. The orange dots show the measured precision, for single exposures, of the differential light curves of all the stars in a field observed during an example night. The black curve is the precision model presented in Equation 4.1. The relevant contributions of the star (red), read noise from the detector (green), noise from the sky (blue), noise from dark current (purple), and the atmospheric scintillation (cyan) are also shown. The expected precision for Trappist-1, if observed under similar conditions, is shown as well.

4.1.3 Photometric precision model verification

To verify Equation 4.1, the photometric precision of all the stars in a single field of view, during a night previously observed by SPECULOOS with the existing instrumentation, was evaluated. This was done by taking the standard deviation of a star's differential light curve and dividing it by its median flux. This particular night was performed with 10 s exposures observing through SPECULOOS' primary filter $I+z'$. The $I+z'$ filter has a mean efficiency of above 90% between 0.7 to 1.1 μm , and results in a total system efficiency shown in the blue shaded region in Figure 4.2. The Vega magnitude of each star, in the $I+z'$ filter, was also calculated using a synthetic spectrum of Vega from Astrolib PySynphot's spectrum library (Lim et al. 2015). The corresponding camera parameters for the existing instrumentation, provided by the manufacturer, are listed in Table 4.1. The results, shown in Figure 4.3, show a close agreement with the theory, where the sky flux per pixel was taken from the average of the measured background counts from the reduced images. The scintillation matched closely to Equation 4.2 taking the average zenith angle of the observations, equalling 50° .

In Figure 4.3, the points that reside above the black curve representing the total precision model are likely due to the inherent stellar variability over the night, or red noise that has not been identified. The other coloured lines in Figure 4.3 break down the relative components which contribute to the noise model in Equation 4.1. For a Trappist-1 magnitude, we can see we are star photon limited, followed on by scintillation, sky, read noise, and dark current. The photon limitation could be reduced by maximising the sensitivity of the camera and filter combination. For the moment, the existing system has been maximised for its application. One could argue increasing the detector's operating temperature could push its quantum efficiency further into the red, and thus receive more photons from cooler targets. However, with the compromise of having higher dark and sky noise, this was not explored here. One related comment is the issue of observing with the $I+z'$ filter, a filter that utilises the tail-end of the detector's quantum efficiency (with the tail-end position having a temperature dependence).

If the detector's cooling solution is unable to sustain sufficient temperature stability, this can contribute to the red noise of an observation. The extent of this red noise has not been quantified specifically in the context of SPECULOOS' instrumentation. The Transiting Exoplanet Survey Satellite (TESS) consortium however, who are using a CCD with a similar quantum efficiency curve, looked at the quantum efficiency changes from $-70\text{ }^{\circ}\text{C}$ to $-50\text{ }^{\circ}\text{C}$ at 1000 nm, where they saw a 4% increase in efficiency (Krishnamurthy et al. 2019). This increase in temperature induced a 10 nm shift to the tail-end of the quantum efficiency curve. One can then approximate that a $1\text{ }^{\circ}\text{C}$ change at SPECULOOS' operating temperature of $-60\text{ }^{\circ}\text{C}$ will yield a $\sim 0.2\%$ change in efficiency at 1000 nm, and a 0.5 nm shift. Under differential photometry, the second order effect induced by quantum efficiency increase should be minimal, the shift however may induce some issues if spectral features from the target come into view that are greater than that from the comparison stars. In Paranal, the ambient temperature can rapidly change by a few degrees celsius in a night, and so it's essential that the detector's housing can accommodate these levels of changes.

With this verification in place, the aim, therefore, with this feasibility study is to see if a new InGaAs based camera and filter combination can minimise Equation 4.1 to an extent better than the existing instrumentation, whilst also minimising the effects induced by varying atmospheric PWV. The following sections describe the approach taken in optimising the precision of the available NIR system, in the context of the SPECULOOS target list (described in Chapter 2).

4.2 Filter optimisation

Traditionally in the infrared, astronomers have utilised filter bandpasses to work in between water absorption bands in our atmosphere, as shown in Figure 4.4. This is done to minimise

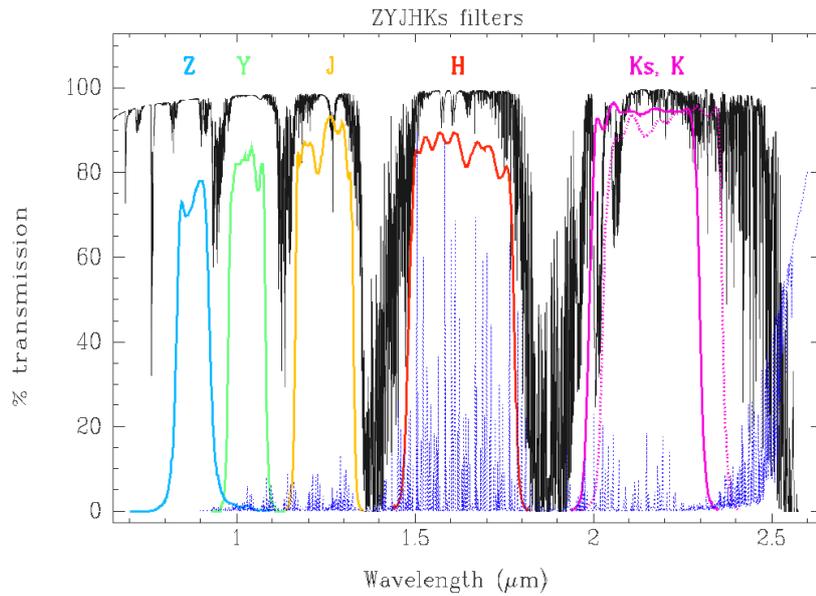


Figure 4.4: Traditional NIR filter bandpasses (labelled z, Y, J, H, Ks, K) used in astronomy. A typical atmospheric transmission and radiance are also superimposed in black and dark blue respectively. Figure adapted from Nordic Optical Telescope web archive.

the potential photometric variability induced by changing atmospheric conditions, as discussed in the previous chapter. However, to maximise the signal-to-noise ratio of an observation we will consider wider bandpass ranges in this evaluation.

One previous approach to minimise PWV effects was done by [Simons & Tokunaga \(2002\)](#), where they produced optimised variants of the J, H, K, L, M filter set (originally defined in [Johnson \(1966\)](#)), known as the Mauna Kea Observatories (MKO) infrared filter set. In this approach, they narrowed down the widths of the respective bandpasses until they avoided the majority of the surrounding water vapour absorption features. In Figure 4.5, using the methods developed in the previous chapter, I show the percentage change in flux observed through the J_{MKO} for different temperature stars as a function of PWV. The J_{MKO} filter passes light between 1.17 and 1.33 μm . Since the change of flux with increasing PWV is similar for the range of effective temperatures evaluated (where the warmer temperatures are typically seen for comparison stars in SPECULOOS field of view, and our target stars being below 2700 K), differential photometry with this filter will be less susceptible to second order effects induced by PWV variability in the context of a SPECULOOS observation.

The aim therefore, is to minimise the spread in the change of flux, Δflux , induced by a change in PWV (as shown in Figures 3.2, 4.5), between different temperature stars. One way to provide a metric to optimise, is to measure the spread in Δflux across the range of effective

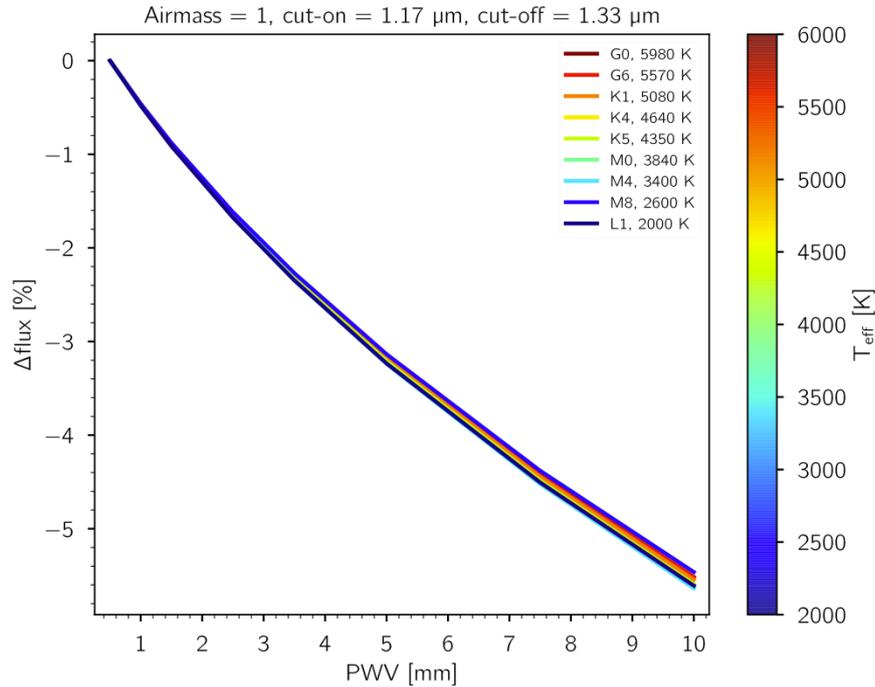


Figure 4.5: Change in flux in the J_{MKO} filter, with respect to 0.5 mm PWV atmosphere, as a function of PWV for different type stars. Modelled using Phoenix BT-Settl stellar atmospheric models (Allard et al. 2012) with known stellar parameters, and the Cerro Paranal Advanced Sky Model (Noll et al. 2012; Jones et al. 2013) to produce the different atmospheric transmission functions at different PWVs.

temperatures, say, across a PWV change from 0.5 mm to 10 mm. I define this metric to be called $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$, which is measured as the standard deviation of the $\Delta \text{ fluxes}$ experienced by a range of effective temperatures that are likely to be used as part of the differential photometry technique for SPECULOOS (2000 K to 6000 K). This metric can be used to measure the extent PWV will affect the differential photometry technique for a particular filter. One can thus model the various cut-ons and cut-offs for a filter that is used in combination with a particular detector, in this case an InGaAs-InP based detector. Figure 4.6 shows how $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$ changes as one uses different cut-ons and cut-offs in between $0.7 \mu\text{m}$ where most of SPECULOOS target's begin the majority of their flux, and cutting off at $1.7 \mu\text{m}$, where the quantum efficiency of an InGaAs-InP detector tends to 0.

Figure 4.6 was produced using the same methods developed in the previous chapter for assessing the PWV-flux change of standard filters. We can see the three major water absorption regions which appear in the evaluated range, at around 0.95 , 1.15 , and $1.40 \mu\text{m}$. We can also see that their influence is apparent whenever they are included in a filter's bandpass. However, at certain cut-ons and cut-offs, even with their introduction, the behaviour of $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$

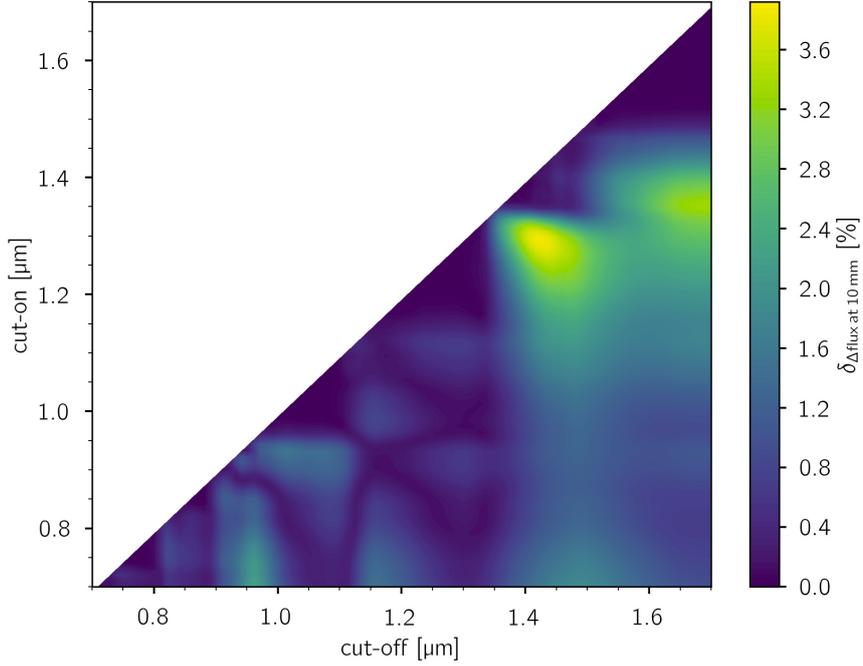


Figure 4.6: The range of cut-ons and cut-offs for a filter observed with a substrate removed InGaAs camera, with the colour map representing the regions where the water vapour absorption poses the greatest problems to the differential photometry technique, quantified by $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$.

Table 4.2: $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$ (defined in text) for the local minimas from Figure 4.6 in addition to a similar calculation for the $I+z'$ filter on the existing instrumentation.

| Name | Cut-on [μm] | Cut-off [μm] | $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$ [%] |
|----------------|--------------------------|---------------------------|--|
| J_{MKO} | 1.17 | 1.33 | 0.06 |
| YJ | 0.96 | 1.33 | 0.09 |
| zYJ | 0.81 | 1.33 | 0.11 |
| $I+z'$ (Andor) | 0.70 | 1.10 | 1.29 |

is minimised. These minimisations are due to a unique combination of stellar spectral flux distributions and the relative positions of absorption features included in the filter range, which permits the distributions to experience similar overall percentage changes in flux with changing PWV. These regions cut-ons and cut-offs are near symmetric around large water absorption features.

To identify the best bandpasses, we considered bandpasses which were wider than J_{MKO} and were within the local minimas shown in Figure 4.6. I began by filtering the regions that had a $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$ less than the lower 10th percentile of the resulting range of $\delta_{\Delta \text{ flux at } 10 \text{ mm}}$, this allowed us to isolate a few bandpasses of interest, as detailed in Table 4.2. These cut-on

and cut-off ranges will now serve as the basis from which we can model the precision of cool stars through a SPECULOOS telescope.

4.3 Precision grid models

To model the precision of the SPECULOOS target list with different detector types and bespoke filters, a direct approach was taken. Here, we obtained model spectra of stars and our atmosphere to integrate over the system response of the instrumentation. A tool called the Precision Grid was developed for this purpose. Feeding in the system response of the telescope + filter + camera combination, a pair of interpolatable grids of stellar fluxes and sky radiances, as a function of different atmospheric PWV and airmass were generated. The values from these grids were used to calculate the model precision detailed in Equation 4.1. The method to generate these grids followed a methodology similar to the PWV grid developed in the previous chapter. Here, Equation 4.3 was used to calculate the stellar fluxes and sky radiances per arcsecond using atmospheric transmission and radiance profiles ($0.5\ \mu\text{m} - 2.0\ \mu\text{m}$) from the online SkyCalc Sky Model Calculator by ESO (Noll et al. 2012; Jones et al. 2013), for 2400 m (the closest available altitude to SPECULOOS' site). In addition to these profiles, PHOENIX BT-Settl stellar models (Allard et al. 2012) provided by the Python Astrolib PySynphot Python package (Laidler et al. 2008) were used, with 2000 K as the lowest available temperature in the package. To extend below the 2000 K limit, ATMO 2020 (Phillips et al. 2020) chemical equilibrium models were used to provide spectra till 450 K, where the authors stated a range of validity of $< 2000\ \text{K}$. This approach provided the flexibility to model the precision of different systems and targets with greater ease.

To build a stellar spectrum from the PHOENIX BT-Settl stellar models, it required four parameters: effective temperature, metallicity ($[\text{Fe}/\text{H}]$), surface gravity ($\log(g)$), and stellar radius. A representative set of stellar models were built using the respective parameters from Pecaut & Mamajek (2013)^a, assigning a metallicity index of 0 to each spectrum. To form spectra below 2000 K with the ATMO 2020 chemical equilibrium models, the model archive allowed a selection of temperatures with steps of 100 K for $T_{\text{eff}} > 600\ \text{K}$, 50 K for $T_{\text{eff}} < 600\ \text{K}$, and with 0.5 steps in $\log(g)$ between 2.5 and 5.5. A temperature-radius relationship was made using Pecaut & Mamajek (2013)'s list, and a fixed $\log(g)$ of 5.5 was chosen to make the final stellar models.

A grid was generated for each filter bandpass and system efficiencies I had identified in the previous section, which outputs a value from Equation 4.3 when fed in observational

^aSpecifically, the updated values from https://www.pas.rochester.edu/~emamajek/EEM_dwarf_UBVIJHK_colors_Teff.txt version 2021.03.02

parameters. These parameters were airmass (between 1 – 3), effective temperature (between 450 K – 36500 K), and PWV (between 0.05 mm – 30 mm). I generated 273 atmospheric profiles, a permutation of airmasses between 1 – 3 at 0.1 intervals, and PWV values between 0.05 mm – 30 mm ([0.05, 0.1, 0.25, 0.5, 1.0, 1.5, 2.5, 3.5, 5.0, 7.5, 10.0, 20.0, 30.0] mm). I then included a set of 107 model stellar spectra to produce a grid of 29,211 values from Equation 4.3 to interpolate in between.

To obtain a value from Equation 4.1 we must also know the pixel scale on the sky, exposure time and seeing of an observation, where the optimal exposure time is a function of the seeing. The seeing/pixel scale will also dictate the size of n_{pixel} needed for the aperture photometry used by SPECULOOS. To estimate the exposure time of an observation, we set the target star's brightest pixel well to fill 70% of the well depth. Here, we assume the brightest target star pixel to be its central pixel when in focus and not subject to any optical imperfections. Traditionally, 70% is the recommended nominal value to aim for. Later, I explore the effects on precision for different well fill depths in addition to varying other parameters. For the target's star central pixel to reach 70%, we must understand the spatial distribution of the star's flux on the detector. The rapidly varying atmosphere will often spread the flux across a few arcseconds. For a very short exposure one will often see a random structure, but exposures longer than a few seconds one will start to see a circularly symmetric distribution. In the case of SPECULOOS' 1 m class telescopes and many other seeing limited telescopes, the distribution closely resembles a 2D gaussian distribution once in focus. A 2D gaussian distribution can be used to model the flux's distribution, where the seeing (quoted as the full-width-half-maximum) is equal to the gaussian's distribution standard deviation / 2.355. To fill 70% of the well depth, we also take into account the dark current and sky/thermal background that accumulates over an exposure. For simplicity, since the detectors we are evaluating are not sensitive to wavelengths where thermal noise becomes dominant, we have assumed the thermal component of the noise model to be part of the sky background or dark current values. Thus, with knowledge of the well depth, we have all the information needed to calculate the exposure time. Lastly, to estimate the size of n_{pixel} from the seeing/pixel scale, we need to collect enough representative signal from the target star's distribution on the detector whilst minimising the size of noise terms which n_{pixel} enlarges. In [Howell \(2006\)](#), they suggest a maximum of 3 times the seeing to be equal to the aperture photometry radius. One could use smaller apertures, but I have assumed this upper limit for the models.

With all these components together and grids generated, we can build the expected precision performance feeding in the following parameters to the Precision Grid tool: PWV, airmass, seeing, plate scale, dark current, read noise, well depth, well fill aim, read time, primary and secondary mirror size, target effective temperature, and target distance. This information will

Table 4.3: 10 minute binned modelled precision (in ppt) from observing a range of cool stars at 10 pc using different NIR bandpasses with PIRT’s camera offering cooled to -60 °C, in addition to the existing silicon CCD based instrumentation with its $I+z'$ bandpass.

| | Target effective temperature [K] | | | | | |
|--------------|----------------------------------|------|------|------|-------|---------|
| | 3000 | 2500 | 2000 | 1500 | 1000 | 500 |
| zYJ | 0.33 | 0.43 | 0.94 | 2.20 | 7.26 | 316.87 |
| YJ | 0.33 | 0.44 | 0.97 | 2.28 | 7.33 | 319.92 |
| J | 0.37 | 0.58 | 1.45 | 3.40 | 10.07 | 439.58 |
| I+z' (Andor) | 0.46 | 0.63 | 1.86 | 4.65 | 27.85 | 1048.53 |

provide the exposure time and precision for a single exposure, and for a binned period taking the precision of a single exposure divided by square root of N exposures in the desired period.

4.3.1 Precision grid results

The expected precision of new NIR instrumentation mounted on a SPECULOOS telescope compared with the existing silicon CCD based instrumentation is evaluated here. I begin by narrowing the filter options made in Table 4.2. Each filter option presented great theoretical invariability to PWV changes during differential photometry. Here, I isolate one filter to continue the comparison of precision using the precision grid described in the previous subsection. Table 4.3 shows the expected 10 minute binned precision for each filter (assuming an 80% transmission efficiency) and camera combination for a range of temperatures from 3000 K to 500 K at 10 pc. We take the properties from Table 4.1 and use the precision grid with stellar fluxes of different effective temperatures, considering an observation done at zenith with a seeing of 1.34" and 2.5 mm of PWV. At SPECULOOS-South, we have a median seeing of 1.34" on the telescope (named Callisto) that the new NIR instrument will be installed on. One may expect to see a marginal improvement in seeing by a few percent observing with the zYJ band, a seeing-limited telescope is expected to improve proportional with $\lambda^{1/5}$ (Boyd 1978) of the observation, however for the modelling I have assumed them to be the same for all instruments.

From Table 4.3, it’s evident that the widest NIR offering with PIRT’s camera presents the best modelled precision. The additional flux observed outweighs the additional noise from the detector and the sky. Likewise, the faster readout time permits more images to be captured within a binned period. Thus, with this evaluation, I will continue the modelling comparison with SPECULOOS’ target list.

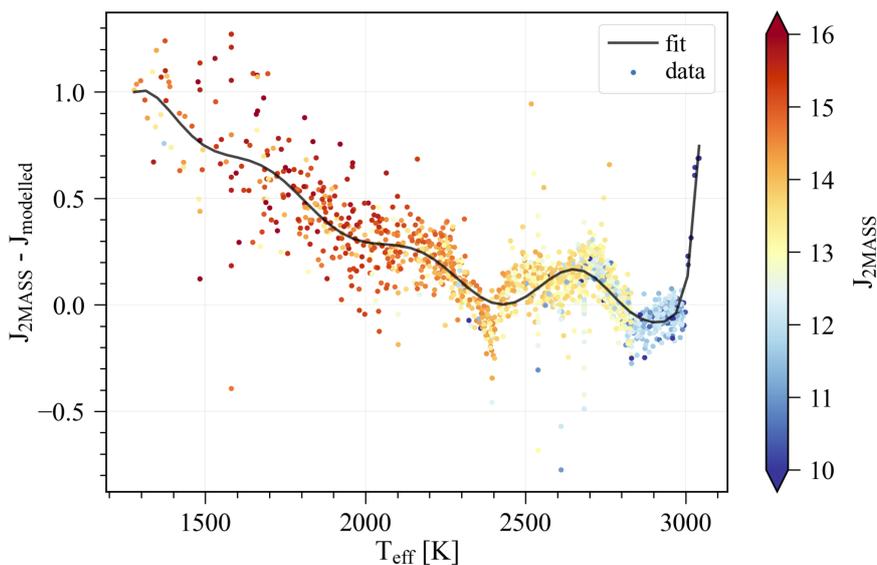


Figure 4.7: Calculated synthetic 2MASS J magnitude (using the spectral response curve from (Cohen et al. (2003)), and the Vega spectra from the Python Astrolib PySynphot Python package (Lim et al. 2015)) subtracted from the J magnitude measured in the 2MASS survey for all the targets in the SPECULOOS target list as a function of effective temperature. A 16th order polynomial fit was used to model the trend. Each datapoint is colour coded with J 2MASS’ magnitude.

4.3.2 Calibrating to the SPECULOOS target list

As described in Chapter 2, the SPECULOOS target list comprises the nearest (< 40 pc) known ultra cool dwarfs. We will use its aggregated information of effective temperatures, distances, and measured magnitudes from the 2MASS survey (Skrutskie et al. 2006) to model the expected precision of SPECULOOS’ current instrumentation compared with the NIR offering from PIRT’s 1280SciCam and the proposed zYJ bandpass. However, prior to beginning the comparison, we will calibrate the precision grid to the SPECULOOS target list using the 2MASS’ J bandpass. In Figure 4.7, I show the calculated synthetic J magnitude (using the spectral response curve from Cohen et al. (2003)), and the Vega spectra from the Python Astrolib PySynphot Python package (Lim et al. 2015)) subtracted from the J magnitude measured during the 2MASS survey for all the targets in the SPECULOOS target list as a function of effective temperature. Since J 2MASS bandpass is the closest bandpass to the proposed zYJ bandpass, it will provide a relevant scaling to more accurately estimate the flux from SPECULOOS’ target stars.

In Figure 4.7, we can see the majority of the synthetic magnitudes are brighter than what was measured by 2MASS, likewise we can see a varying trend as one decreases in effective temperature. The variations from 0 are likely attributed to deviations from the chemical com-

Table 4.4: Same as Table 4.3 after target flux calibration, taking the temperatures within the calibration fit made in Figure 4.7.

| | Target effective temperature [K] | | | |
|--------------|----------------------------------|------|------|------|
| | 3000 | 2500 | 2000 | 1500 |
| zYJ | 0.33 | 0.43 | 1.18 | 4.25 |
| YJ | 0.34 | 0.45 | 1.22 | 4.42 |
| J | 0.37 | 0.59 | 1.86 | 6.63 |
| I+z' (Andor) | 0.46 | 0.63 | 2.35 | 8.94 |

position assumptions made in the stellar models. Similarly, errors in temperature classification (± 100 K), radius estimates ($\pm 0.01 M_{sun}$), and errors in distance estimates from Gaia (Gaia Collaboration et al. 2018) will play a role. To attempt to correct the average deviation, a 16th order polynomial was used to fit the data. One can see deviations from the fit increases as we progress to dimmer stars, where magnitude errors in the 2MASS survey (± 0.03 mag for stars < 13 mag, but progresses to ± 0.10 mag at 16 mag)^b increases, and the errors previously highlighted will contribute to the spread around the fit. There is an evident increase in the deviation at 2000 K mark, when transitioning to the ATMO 2020 models. This fit can be then used to scale the flux output from the flux grid developed in the previous section, tailoring the precision grid to the SPECULOOS target list within a temperature range of 3042 K and 1278 K. The fit is an overfitting, and therefore can not be used outside this range.

With this calibration in place, Table 4.3 can be re-calculated for the valid temperature range of the calibration, shown in Table 4.4. As expected, the modelled precision has worsened for the cooler targets. The decrease in flux will require longer exposures, which has a greater impact on the NIR instrumentation given the higher sky and dark components. Nonetheless, it still maintains a better expected precision than the existing instrumentation.

4.4 SPECULOOS target list precision

The precision understandably decreases for lower temperatures, due to their lower intrinsic brightness, their smaller stellar radii (where there is a 50% decrease in radii from 3000 K to 2000 K, as reported by SPECULOOS target list Sebastian et al. (2020a)) allows for a higher signal-to-noise-ratio (SNR) of seeing changes in flux due to a transiting planet. This is assuming the same size of planets are to orbit independent of their host star's stellar temperature, within the assessed SPECULOOS target temperature range. Thus, to compare the precision of the target list in the context of searching for transits, we will assess the SNR of measuring a dip in

^bInformation from (Last accessed 2022-04-15): <https://old.ipac.caltech.edu/2mass/releases/allsky/doc/figures/secii2f7.gif>

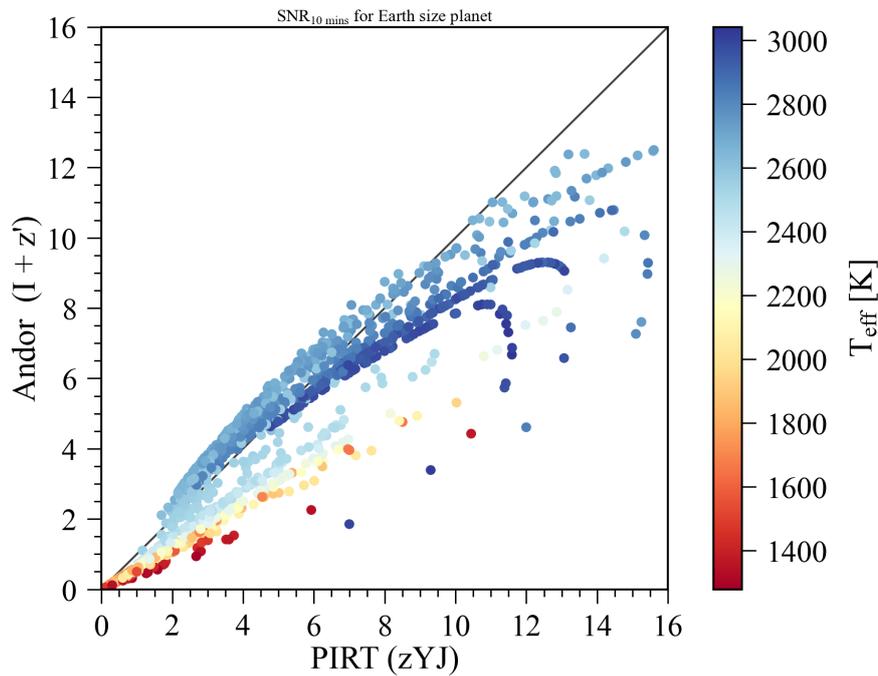


Figure 4.8: The 10 minute binned modelled SNR for a change of flux induced by a transiting Earth-sized planet as observed by the proposed NIR instrumentation compared with the existing instrumentation. Each point is colour coded as a function of stellar temperature.

flux proportional to an Earth-sized planet transiting across the disc of their host star, taking the stellar radius value from the aggregated target list.

Figure 4.8 shows the 10 minute binned SNR for a change of flux induced by a transiting Earth-sized planet between the existing and the proposed NIR instrumentation. When modelling the precision, I limited the exposure period to a maximum value of 120 s, to maintain a minimum cadence to the observations, such to capture a potential short variability (Murray et al. 2022). Likewise, high frequency red noise can instead affect a subset of shorter exposures rather than a single longer exposure (Pont et al. 2006). In the following subsection, I explore the effect on the photometric precision with respect to well fill, quantifying the effect for the instances where the exposure times required to fill 70% of the well are longer than 120 s. In Figure 4.9, I show the modelled exposure times for both evaluated instruments. For the proposed NIR instrument, the expected exposure times are notably shorter than the existing instrument, with median values being 36.5 s vs 120 s respectively. This is for a handful reasons - increased stellar flux, smaller well depth, higher dark currents and sky background filling the central pixel faster. The exposures at the tail end of the histogram are mostly dominated by dark current and sky flux rather than target flux.

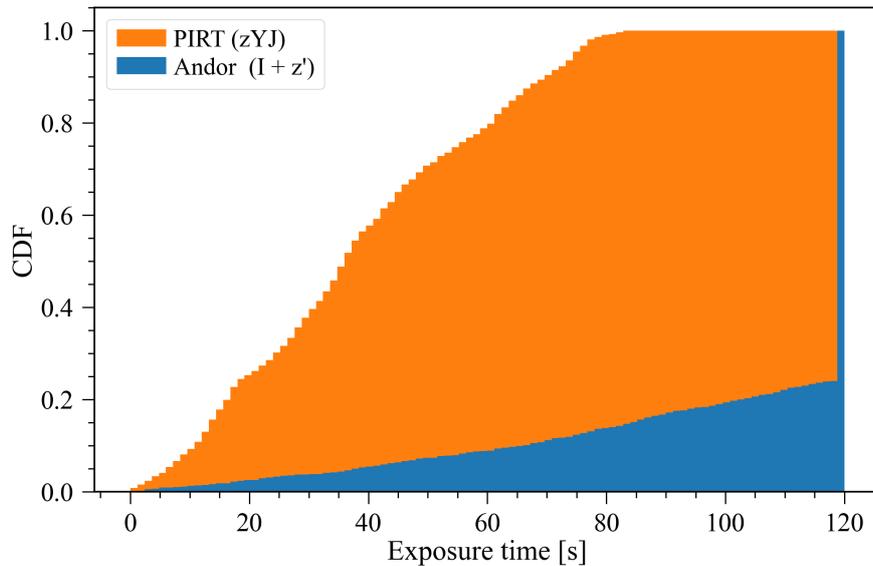


Figure 4.9: Modelled exposure times observed by the proposed NIR instrumentation (in orange) compared with the existing instrumentation (in blue), displayed as a cumulative histogram with 100 bins. The exposure time was calculated as the time required to fill 70% of the target’s central pixel’s well depth at a seeing of $1.34''$.

As can be seen in Figure 4.8, the majority (68%) of the targets lie below the 1:1 black line, which presents an SNR advantage to the proposed NIR instrumentation. Some targets still present an advantage to the existing instrumentation, primarily targets above the Trappist-1 effective temperature of ~ 2550 K. Since SPECULOOS is a transit-hunting survey, the SNR of observing a transit has to be ideally sufficiently large to see in a single light curve. In the instance of a 10 minute bin, we can set an arbitrary cut-off of an SNR of 3 to remove the targets which will be too noisy to detect a terrestrial Earth-like transit, to minimise wasting valuable telescope time. 60% of the $\text{SNR} > 3$ targets for the proposed NIR instrument present a better SNR than the existing instrument. 88 targets are in the realm where $\text{SNR} < 3$ for the existing instrument and $\text{SNR} > 3$ for the proposed NIR instrument, with a median temperature of 2263 K within this set. One may notice a path of curvature to some of the warmer temperature ranges, and thus brighter targets. This is due to the longer read time (10.5 s vs 0.1 s) associated with the CCD based instrumentation, where both instruments will have very short exposure times (less than a couple of seconds) for these bright targets. As a result, the number of images within a binned period will be lower in the existing instrumentation in comparison to the proposed NIR instrument, which is based on the much faster CMOS readout architecture. In these cases, we would likely observe out of focus for the CCD based camera, or with a different bandpass to increase the exposure time such to increase the photometric precision and to minimise the impact on the longevity of the mechanical shutter.

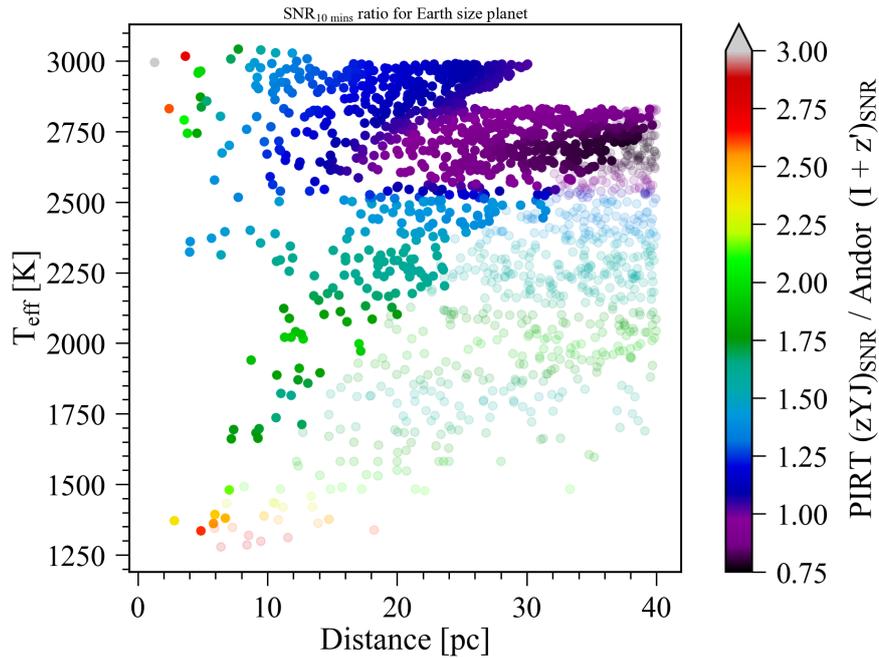


Figure 4.10: SNR ratio between the proposed NIR instrument and existing instrumentation for the entire SPECULOOS target list, plotted as a function of effective temperature against distance. Solid coloured points represent the cases where the $\text{SNR} > 3$ for the proposed NIR instrumentation.

Presenting the information from the previously computed models as a ratio of SNRs in Figure 4.10, plotted as a function of effective temperature and distance, we can see how the majority of the interesting targets in the target list for the proposed NIR instrument lie below 2550 K. In this domain, we can see some targets presenting up to $\sim 3x$ higher SNR than seen by the existing instrumentation. Looking closely, we can see some targets of the same temperature having an SNR lower than more distant targets. These cases can be attributed to errors in stellar radius in the target list that were used to calculate the 10 minute binned SNR for a change of flux induced by a transiting Earth-sized planet. One can also see a discontinuity around 2000 K, where the ATMO 2020 models were introduced. Nonetheless, from the modelled values, we are limited to the closest targets if we wish to have an $\text{SNR} > 3$, going from < 30 pc at 2550 K to $< \sim 6$ pc at 1250 K. Progressing from these models, we can then assess the effect of varying the various components contributing to instrumental precision as part of a sensitivity analysis.

4.4.1 Sensitivity analysis

Here, taking into consideration the scenarios of observing a range of targets of different brightnesses, and thus different exposure times, I assess the contributions of each noise component

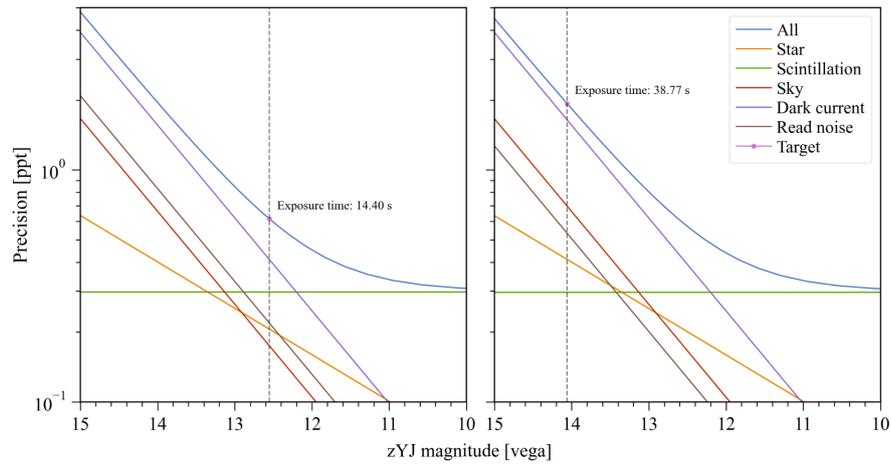


Figure 4.11: 10 minute binned precision vs the calculated zYJ magnitude for a range of brightnesses for a fixed exposure time. Each plot also shows the contribution from each noise component in Equation 4.1, excluding thermal contribution which was assumed to be negligible for our wavelength range. On the left we have the noise contributions for a bright target, and on the right for a dimmer target. Each exposure time is calculated to meet 70% of the well depth.

from Equation 4.1 to the final precision of an observation. Focusing on the cooler targets, I take the median temperature value of 2263 K from the previously identified 88 targets who had a $\text{SNR} < 3$ for the existing instrument, and $\text{SNR} > 3$ for the proposed NIR instrumentation. I modelled two distances, one at 20 pc, and another at 10 pc to vary the brightness. This yielded exposure times of 14.4 s and 38.8 s (a value close to the median value of the target list) respectively for the proposed NIR instrument, and 120 s (hitting the 2 minute cadence set limit) for both distances for the existing instrumentation.

In Figure 4.11, we can see which noise component dominates at the specific magnitude of the target star. For longer exposures, we are dark limited, followed on by sky, read noise, star, and scintillation. For the shorter exposure, we are still dark-limited, but the read noise and scintillation take a comparatively greater role in reducing the precision of a binned observation. Focusing on the longer exposure scenario, which matches closely with the median modelled exposure time of the target list, in Figure 4.12, I vary each independent variable that contributes to the proposed NIR instrument's precision such to assess the specification that could be realistically met if we were to procure such an instrument.

In Figure 4.12, I varied the pixel scale (the angular resolution each pixel has on the sky, 0.1 to 2") as if we were to re-image rather than to change the physical size of the pixels. As can be seen, the behaviour of the existing instrument rapidly improves in precision as we progress from small plate scales to the value it currently operates at, 0.35". A low pixel scales, a higher number of pixels are required to fill an aperture, and as a result there's a higher summed contribution

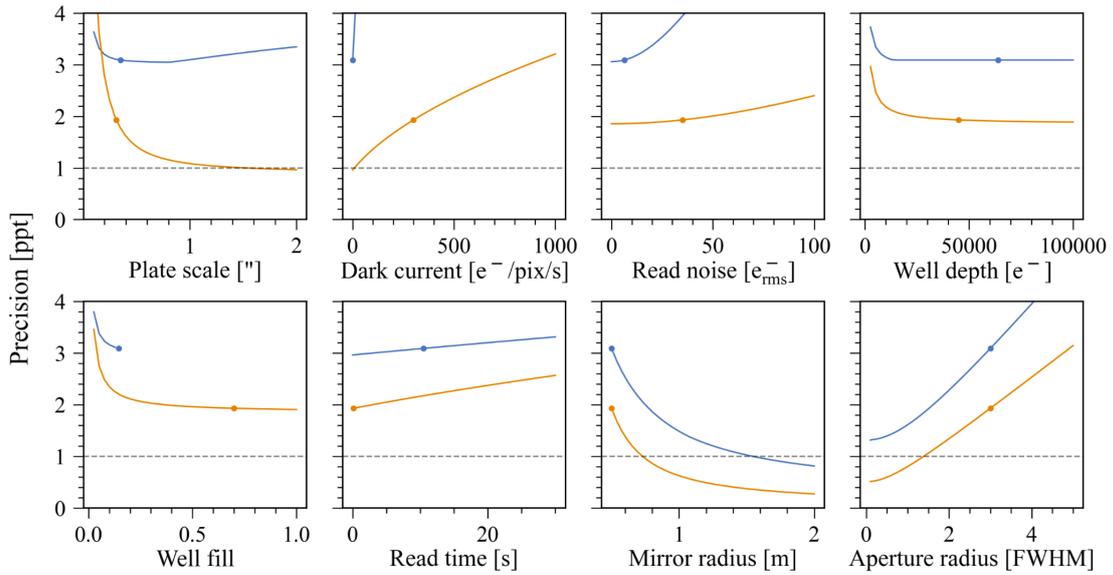


Figure 4.12: Modelled 10 minute binned precision for a 2263 K target at 20 pc for the proposed NIR instrument (in orange) and existing instrumentation (in blue), varying controllable (by engineering or data processing) independent variables that contribute to the precision equation, Equation 4.1. The round markers on each plot indicate the fixed value the instrument is known/modelled to operate at. The dashed grey line is the 10 minute binned precision to achieve 1 ppt.

of read noise and dark current noise. From there it plateaus, and gradually increases as we reach 2". The plateauing is partially due to the exposure capping at 120 s but primarily the system reaching optimal precision for this target, the gradual increase can be attributed to the higher proportion of sky background in each pixel at those scales which don't outweigh the additional number of images per bin resulting from the shorter exposure times. The proposed NIR instrument behaves in a similar manner, but the improvement in precision occurs over a larger range given the higher relative amounts of the stellar flux and detector noise components. It has a marginally smaller plate scale than the existing instrumentation (0.31" vs 0.35"), and it would benefit greatly from having a larger plate scale to reach its optimum precision. However, in order to resolve stars from optically neighbouring ones and avoid the blending of flux, one should have an angular resolution at least twice of lowest resolvable objects in the field to follow Nyquist sampling (Willey 1992), which is often limited by the seeing of the site. Moreover, a higher number of pixels covering a target will minimise any potential flat-fielding errors (Baumer et al. 2017) - however, this was not looked into as part of this study.

The seeing at SPECULOOS-South is $\sim 1.34''$, the value used for the modelling. A pixel scale smaller than $0.67''$ would thus match the Nyquist condition for the median seeing case, at this scale the 10 minute binned precision would improve from 1.9 ppt to 1.2 ppt, owing to the

increased number of images from the shorter exposure time of 12.6 s and fewer n_{pixels} used. However, re-imaging the focal plane of a SPECULOOS telescope was not within the scope of this project. Nonetheless, the pixel pitch of the proposed NIR offering that was assessed here from PIRT (12 μm vs 13.5 μm of the existing camera) was of a size acceptable for use in Paranal. An alternative route would be to consider a detector with a larger pixel pitch. However, the dark current per pixel would also increase since its value is primarily dependent on the surface area of the photosensitive medium. Increasing the pixel scale via this route would not improve the binned precision. Ideally, one should have a small pixel pitch to minimise dark current whilst re-imaging to maintain a plate scale suitable for ground-based photometry. At the time of writing, the smallest pixel pitch commercially available in NIR detectors comes at 1.82 μm (Lee et al. 2020), based on Lead sulphide (PbS) quantum dots rather than InGaAs. Whilst cheaper to produce, their QE is $\sim 5\times$ lower than the assessed InGaAs based detector for the zYJ region. Sony has succeeded in producing an InGaAs detector with a pixel pitch of 5 μm . Unfortunately, their smaller pixel pitch offering comes with significantly higher read noise ($150 e^-_{rms}$)^c, and without re-imaging the focal plane, the precision would become similar to the existing instrumentation for this target. The pixel scale also affects the total field of view on the sky, as a function of the total number of pixels forming the detector. The total field of view will dictate the number of comparison stars available for differential photometry, a topic covered for the next subsection.

Moving onto dark current sensitivity, considering a range per pixel from 0 to 1000 $e^-/\text{s}/\text{pixel}$ - a range which encompasses no dark current to a value $\sim 3\times$ higher than the typical value quoted by PIRT for -60°C operation. As can be seen in Figure 4.12, the existing instrument's precision does not cope well with dark currents in the order seen in proposed NIR instrumentation. If we were to consider a comparable range of 0.0 to 0.6 $e^-/\text{s}/\text{pixel}$ instead, we would see a negligible binned precision change. I did not consider the change in quantum efficiency that the existing instrument would experience at higher operating temperatures/higher dark currents. The proposed NIR instrument on the other hand, which is dark-limited for this observation, shows a notable sensitivity that would yield 30% change in precision if we were to double the dark current (for the InGaAs medium, the dark current is estimated to double every 6.6 $^\circ\text{C}$ as found with data from Sullivan et al. (2014)). Given the dark current is a controllable parameter, a function of the stability of the detector medium and operating temperature, we seek to minimise this value as much as possible. For comparison, the sky background, for a median night in Paranal, is estimated to be $\sim 62 e^-/\text{s}/\text{pixel}$ for the proposed NIR instrument with the zYJ filter, and $\sim 14 e^-/\text{s}/\text{pixel}$ for the existing instrument with the $I+z'$ filter, 880 and 152 $e^-/\text{m}^2/\text{arcsec}^2/\text{s}$ respectively. This means that the proposed NIR instrument has $\sim 4.4\times$ higher sky background

^cDate accessed 2022-09-30: https://www.qhyccd.com/scientific-camera-qhy990_qhy991/

than the existing instrument, and $\sim 5x$ higher dark current than sky background per pixel. It would potentially require cooling the proposed NIR instrument offering by $20\text{ }^{\circ}\text{C}$ to $-80\text{ }^{\circ}\text{C}$ to be sky limited, which would halve the binned precision in this example. However, this is assuming that the quoted dark current value is entirely sourced from the detector and not also a combination of thermal background (which was not modelled here separately) - this is discussed in detail in Section 6.2.1.

For the brightest targets, read noise becomes a more dominant factor. For the median case shown in Figure 4.12, increasing the read noise, whilst having an effect, is less dominant than the effect seen with the dark current. Doubling the read noise would see a 10% change in precision for the proposed NIR instrument, and a 3% change in the existing instrument. Reducing the read noise has a minimal effect, with other factors dominating the noise budget at this exposure time. In theory, to benefit observing brighter targets, the read noise should be reduced through the use of Fowler Sampling or similar methods, which could see the read noise decrease significantly as shown in [Hoffman et al. \(2005\)](#) to values comparable to those seen in low-noise CCDs. Unfortunately, PIRT's detector could only be used in a Correlated Double Sampling (CDS) mode, and they stated in private communications that it was not possible with their current camera offerings.

The well depths provided by the manufacturers were more than sufficient for the target considered here, with the existing instrumentation not being able to use the majority of the well depth given it reached the imposed 2 minute exposure limit, reaching 15% of the well. The target fill for the brightest pixel of the target star was set to reach 70% of the well depth. This is to deal with the variability in target brightness, but also variability in seeing which would change the shape of the star on the detector, and thus its peak brightness. If we were to observe at the exposures estimated in Figure 4.11, and the seeing were to change from SSO's median value of 1.34" to 1.11" (25th percentile), the well would reach 87% of well fill for the proposed NIR instrument and 20% for existing instrument. At the 5th percentile, 0.85" seeing, the well fill for the proposed NIR instrument would exceed saturation, reaching 125%. Thus, one should be careful when setting the exposure time, and should perhaps accommodate for smaller seeing such as filling to 50% rather than 70%. One should avoid reaching saturation since the detector's behaviour near the saturation limit is often non-linear. On the other end, if the seeing was to worsen and reach 95th percentile, 2.14", the well fill would reach 46% with the exposure time set for this target.

Image read time will also have an effect on the binned precision, by reducing the number of images used in a binned period. The proposed NIR offering has significantly faster readouts given its CMOS architecture. The benefit of decreasing the read time for the existing instrument for this target would be minimal, given the already long exposure period. A parameter that

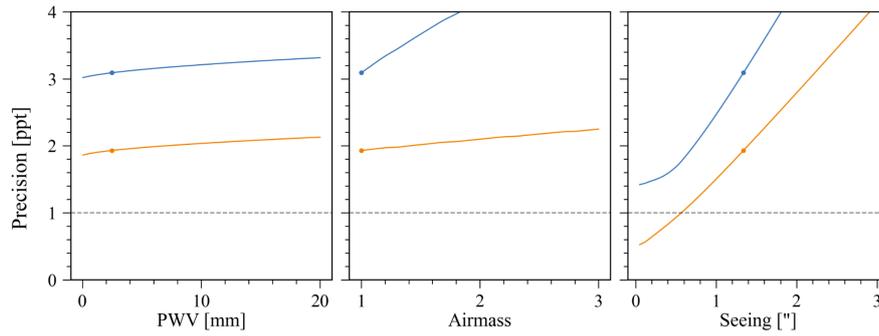


Figure 4.13: Modelled 10 minute precision for a 2263 K target at 20 pc for the proposed NIR instrument (in orange) and existing instrumentation (in blue), varying the sky parameters of an observation. Specifically, precipitable water vapour (PWV), airmass, and seeing. An increase in airmass would normally increase the seeing - however, this was not accounted for in the models, but rather varied as independent parameters. The round markers on each plot indicate the fixed value the atmosphere was modelled at. The dashed grey line is the 10 minute binned precision to achieve 1 ppt.

would benefit greatly both instruments would naturally be increasing the size of the telescope(!). For the target assessed here, it would require a ~ 1.5 m class telescope to reach sub-1 ppt binned precision for the proposed NIR instrument, and a ~ 3 m class telescope for the existing instrument. The final independent parameter under our control, affecting the precision of an observation, is the aperture radius used for encapsulating the flux from the target star. Here, it was assumed that the radius to be 3 FWHM, equivalent to 7 sigma within a gaussian, certainly enough to capture all the flux from a 2D gaussian distribution. It is possible however to use smaller apertures, where with the existing instrumentation we often find the best aperture to be between 2-3 FWHM. However, there is a limit to how small an aperture can be used, where the variability in seeing throughout the night places a lower limit on the aperture size. Currently, the pipeline (Murray et al. 2020) used to produce light curves from an observation uses a fixed aperture for the entire night. Hence, if the seeing were to vary, more or less of the flux will appear in the aperture and give rise to changes in the light curve due to atmospheric variability rather than variability from the target source.

Ideally, one should find a site which minimises the seeing of an observatory, such to achieve the best precision from the ground. Figure 4.13 shows the effect that changes in atmospheric parameters/conditions can induce on the binned precision. Changes in seeing have the greatest effect on precision due to changes in the total number of pixels used in the differential photometry process, where fewer pixels that contribute to the overall noise are used to measure stellar flux. The models assume fractional pixels, in reality the discrete nature of pixels would mean that subplots on varying the aperture (Figure 4.12) and seeing (Figure 4.13)

would have a step-like nature rather than a smooth continuous curve. Nonetheless, it indicates that a priority in minimising the seeing should be made. At the moment, the median seeing at the telescope available for proposed NIR instrumentation is $\sim 1.34''$, with similar values seen at the three other identical telescopes at SSO and also at the one other identical telescope located at the Teide Observatory (Burdanov et al. 2022). Whereas, at the nearby VLT platform, a median seeing of $0.66''$ is achieved.^d The expected diffraction limit of one of SPECULOOS' 1 m class telescopes, operating at 1 micron, is $\sim 0.25''$. Thus, we can assume that the differences in seeing can be attributed to differences in the site's wind/thermal patterns, dome seeing, or issues with the optical design. However, it was not within the scope of this PhD to assess the source of the issue and scope potential solutions, but it is recommended for future work. An area that may be of interest to pursue would be the domain of adaptive optics (AO) (Wilson & Jenkins 1996), where recent developments have permitted the application of AO on smaller telescopes. In Quintavalla et al. (2020), they applied multiactuator adaptive lenses (MALs) and managed to reduce the FWHM of an example star by $\sim 25\%$ in one field (from $1.83''$ to $1.42''$ for a 30 s exposure on a 1 m telescope). However, with the more advanced setups, such as that used on the VLT, milli-arcsecond FWHMs can be reached (Riccardi et al. 2022). It would ideally however require to be applicable to the entire field of view ($12'$ on the existing instrumentation), since current aperture photometry methods use the same aperture size for all stars in the field. Although, with a sufficient number of comparison stars, multiple apertures can be potentially used to reach the desired SNR to form the artificial comparison star as part of the differential photometry process. Another potential advantage of using AO, would be to have greater FWHM stability throughout a night, which would minimise the effect of false features being induced by sudden changes of FWHM.

Other atmospheric parameters such as airmass and PWV will affect the transmission of a bandpass. As can be seen in Figure 4.13, an increase in PWV will decrease the target flux by $\sim 16\%$ for both filters from 0 to 20 mm, which has a small effect on the binned precision. An increase in airmass on the other hand has a much greater effect on the existing instrument, given its greater sensitivity to sky background in its noise budget, which scales with airmass. The proposed NIR instrument has a dark current approximately 5 times higher than the median sky background at zenith for Paranal, and as such it has a lower sensitivity to higher airmasses. Thus, with the previously assessed variables in Figures 4.12-4.13 and Figure 4.11, we can deduce which parameters will generally present the greatest sensitivity to improving/reducing the photometric precision of a binned observation. However, without considering changes to the optics of the telescope, such as re-imaging the focal plane to change the pixel scale, or

^dParanal Site Information, date accessed: 2021-09-15 <https://www.eso.org/sci/facilities/paranal/astroclimate/site.html>

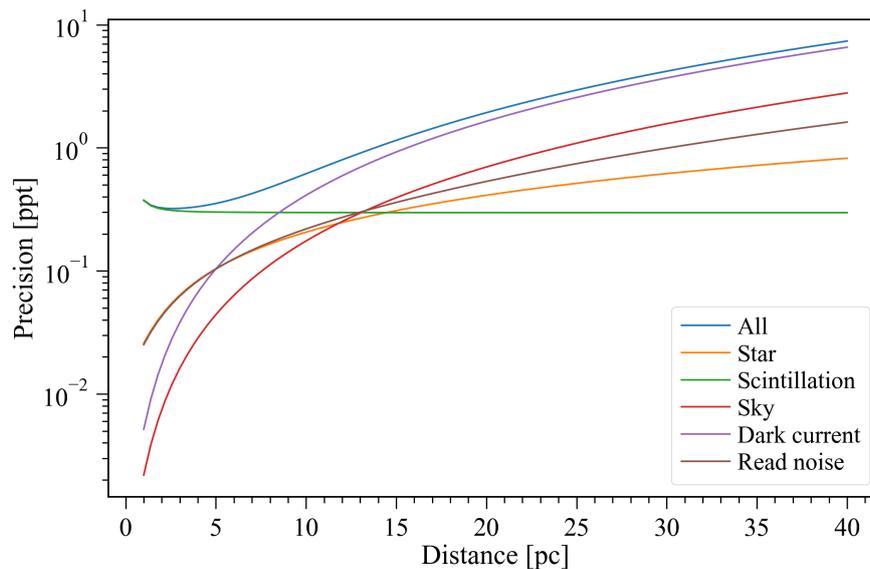


Figure 4.14: The relationship between the modelled 10 minute binned photometric precision (separated into its components from Equation 4.1) of a fixed effective temperature target (2263 K) at varying distances is shown. At each distance, the expected exposure time to fill 70% of the target’s central pixel’s well depth was modelled.

investigating means of minimising the seeing, the parameter which presented the greatest scope for affecting the binned precision for the exposure time evaluated here was the dark current.

At other exposure times however, the dark current’s sensitivity to limiting an exposure will be different, where other parameters will instead take the limiting role. To evaluate this, I considered the range of exposure times the proposed NIR instrumentation will be subjected to if it were to observe the entire SPECULOOS target list (see Figure 4.9 for the cumulative histogram of exposure times). We can recreate Figure 4.11 in the context of the target’s precision at a range of distances which would yield the range of exposure times expected for SPECULOOS target list. The resulting modelling, in Figure 4.14, shows dark current as the limiting factor for the majority of the distances until 8.5 pc (11 s of exposure), where scintillation is then limiting. At 4.9 pc (4 s), read noise and star photon noise exceeds the noise contribution given by the dark current, while scintillation maintains limiting. At these close distances and thus short exposures, the observations will be mostly sensitive to read time, particularly in situations where read time is in the order of the exposure time, resulting in a larger portion of the binned period being occupied by the read time rather than integration time. Changing the other variables considered in Figures 4.12-4.13 did not affect the binned precision significantly at these shorter exposure times. If we wish to have more targets at better precision, we need to reduce the dark current, after which we will be limited by the sky background. At this point we would need to have a significantly larger primary mirror, which would also help reduce the

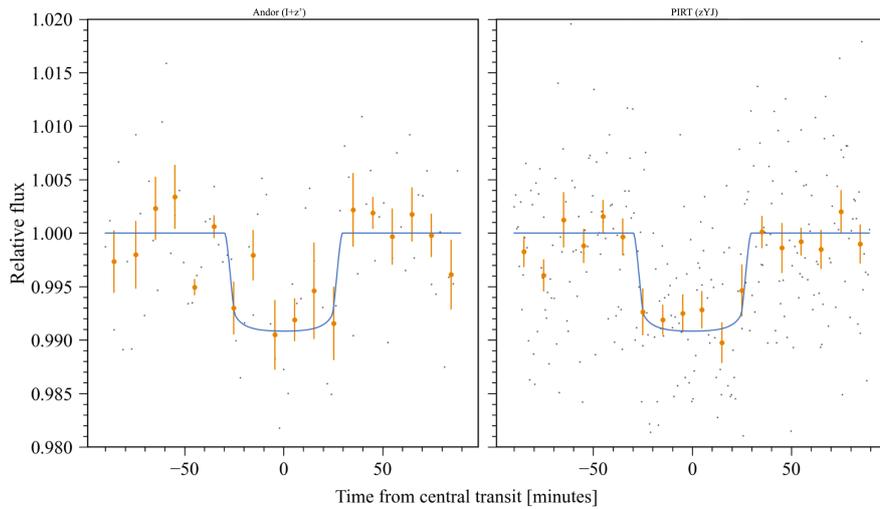


Figure 4.15: Modelled transit (in blue) for an Earth-sized planet around a 2263 K star (at 20 pc) with a transit duration of 1 hour and a period of 10 days. On the left, modelled scenario for the existing instrumentation, and on the right, the modelled scenario for the proposed NIR instrumentation. For each subplot, the grey points represent the relative flux from single images using the modelled single exposure precision as the error on a normal continuous random variable to model its variability, and the orange points represent the 10 minute binned values with the error bar as the standard deviation of the binned period.

scintillation noise and observe targets that are currently too faint to yield useful photometry. To yield an $\text{SNR} > 3$ (with the metric defined for Figure 4.8), we ideally need to be below 2 ppt at 10 minute bins (taking the median binned precision for SPECULOOS targets between SNR of 2.9 and 3.1), which is approximately achieved by exposure times below the median exposure time (36.5 s) for the evaluated SPECULOOS target list. If we were to have a dark current of $0 \text{ e}^-/\text{s}/\text{pixel}$, the achieved binned precision would be expected to reach 0.64 ppt for the same exposure time.

Thus, with all the photometric precision information and their relative sensitivities, we can model a single transit event with the expected cadence and photometric precision of the respective instruments. Continuing with the example 2263 K target at 20 pc evaluated in this subsection, we can use transit models (from the `batman` Python package (Kreidberg 2015)) for an Earth-size planet inducing a transit around such a system, as shown in Figure 4.15. The binned precision is notably better for the proposed NIR instrument, where the transit of 9.2 ppt is better distinguished. However, whilst the proposed NIR system is theoretically showing promise, we must also consider the field of view the detector size will provide and the amount of comparison stars it will yield necessary for the differential photometry process.

4.5 Field of view suitability

The best field of view contains a sufficient number of bright, but not too bright, comparison stars for the differential photometry process. They should ideally be of the same spectral energy distribution, but with the proposed zYJ bandpass, spectral energy distribution is expected to be less of a concern with it optimised for PWV invariability in the differential photometry process. To understand if the proposed NIR instrument's field of view of $6.6' \times 5.3'$ will be sufficient for SPECULOOS targets, I evaluate the 2MASS (Skrutskie et al. 2006) catalogue at different field of views to identify the minimum field of view which will satisfy the SPECULOOS target list. From this analysis, I set the upper brightness threshold relative to the target star for each field and conclude the number of fields deemed suitable for differential photometry that yield an artificial star with a 10 minute binned precision better than 1 ppt.

4.5.1 Upper magnitude limit

The upper magnitude limit for a comparison star with respect to the target star is dependent on the elements filling the target's central pixel well depth (wd , in electrons) after an exposure. Namely, the target's integrated flux (N_{target}), sky background (N_{sky}), and dark current (N_{dark}) in the target's central pixel. The target flux we receive will be spread over many pixels, due to the seeing. At a seeing of $1.34''$, the central pixel will receive $\sim 4.6\%$ (A) of the target's total flux. The percentage fill the central pixel will thus be equivalent to

$$central\ pixel_{star} = \frac{fAN_{target} + N_{dark} + N_{sky}}{wd} = \begin{cases} 0.7, & f = 1 \\ \text{otherwise,} & \end{cases} \quad (4.4)$$

where the aim is to reach 70% of the well depth. For the comparison stars in the field, their central pixel will fill to other percentages depending on their relative brightness to the target star (f). Their relative brightness to the target can be related by the difference in magnitude (m)

$$m_{target} - m_{star} = 2.5 \log_{10} \frac{F_{star}}{F_{target}} = 2.5 \log_{10} f. \quad (4.5)$$

At the target star's well fill, we can produce a useful relationship to remove our dependence on N_{sky} and N_{dark} as

$$N_{dark} + N_{sky} = 0.7wd - AN_{target}. \quad (4.6)$$

We can then set an absolute upper limit of brightness for a comparison star as when the central pixel reaches saturation, giving

$$central\ pixel_{star} = \frac{fAN_{target} + 0.7wd - AN_{target}}{wd} = 1. \quad (4.7)$$

Rearranging for f:

$$f = \frac{0.3wd}{AN_{target}} + 1, \quad (4.8)$$

Which enables us to calculate the magnitude difference using Equation 4.5:

$$\Delta m_{upper} = m_{star} - m_{target} = -2.5 \log_{10} \left(\frac{0.3wd}{AN_{target}} + 1 \right). \quad (4.9)$$

The size of N_{target} will vary depending on the integration time, given the dark current and sky background will also contribute to reaching the 70% target well fill. For the example target star considered in the previous subsection (2263 K star at 20 pc, with a J magnitude of 13.3), the upper magnitude limit for a comparison star in the field will be -0.6 from the target star's magnitude. For a brighter target star example, a star with a J magnitude of 7.5, will have an upper magnitude limit of -0.4 in the field, given less of the well will be filled with dark current and sky background.

4.5.2 Lower detectable magnitude

The lower magnitude is dependent on our ability to detect the star, this does not mean it will be used for the differential photometry process however. A crude detection threshold criteria is set as

$$3\sqrt{N_{sky} + N_{dark} + N_{read}^2} < AN_{star}, \quad (4.10)$$

where the peak flux of a star has to be at least three times the background noise. Once again using Equation 4.5, we can define the magnitude difference to the target star to be

$$\Delta m_{lower} = -2.5 \log_{10} \left(\frac{3\sqrt{N_{sky} + N_{dark} + N_{read}^2}}{AN_{target}} \right). \quad (4.11)$$

For a 13.3 J magnitude target star, the lower detectable magnitude will be $+4.2$, and for a brighter 7.5 J magnitude star, it will have a lower limit of $+6.1$.

4.5.3 Evaluating SPECULOOS' field of views

This is an indicative study, as opposed to a quantitative calculation of the viability of a field of view, as there may be variable comparison stars in the field that make it unsuitable for

differential photometry. Nonetheless, I began by querying the 2MASS archive for each of the SPECULOOS targets that presented a $\text{SNR} > 3$ (for the metric defined for Figure 4.8), starting from 2' up to a 24' square field of view (such to evaluate a field of view greater than the existing instrumentation and the proposed NIR instrumentation), yielding 963 fields. For this analysis, I assumed that differences in J magnitude will be equivalent to the differences observed in the zYJ bandpass. For each field of view I applied an upper and lower magnitude cut, using the definitions from the previous subsections with the calculated N_{target} integrated flux and background noise parameters from the precision models, calculated in Section 4.4. The target stars were also removed from the queried 2MASS fields.

From here, I incrementally evaluated the field of view from each target star in 1' increments, taking the brightest stars in the field and calculating the 10 minute binned precision in zYJ (using the difference in magnitude to estimate N_{star} for Equation 4.1) for the produced artificial star used as part of the differential photometry process, only adding stars if they improved the precision. I set a minimum of 3 comparison stars in the field before considering the field potentially usable. Whilst 1 comparison star is sometimes considered enough, there is a risk that its variability is above the noise level for the duration of the observation and will influence the final light curve and present false variability in the target star. A minimum value of 3 would minimise this risk, although not entirely, and would also permit one vet to the comparison stars from one another. One further limitation of this analysis is that I did not consider the effect of blending of stars in the field. Figure 4.16 shows the resulting calculation, displaying the fraction of usable fields that could potentially yield artificial light curves with 10 minute binned precisions which still maintains a $\text{SNR} > 3$ for the target star's differentiated light curve.

The results from Figure 4.16 indicate that approximately half the fields will be observable with the proposed NIR instrument's field of view of $6.6' \times 5.3'$ if we wish to maintain a $\text{SNR} > 3$. This result will be different if we were to have a different metric of SNR, for example, the SNR for a flux change for a larger planet than Earth size, or if the camera had better noise characteristics. The overall result is similar to having an upper limit on the artificial light curve's precision of around 1.3 ppt at 10 minute bins. Nonetheless, this indicates that if we wish to have a higher proportion of usable fields we should seek to observe at a larger field of view. 92% of the fields would become observable if we were to match the field of view of the existing instrumentation of 12'. However, without re-imaging the focal plane, or having a greater number of pixels or larger pixels, this route was not possible within the scope of the project.

Figure 4.17 shows a mollweide projection of the night sky of all the observable fields. The majority of the usable fields, at the field of view of the proposed NIR instrument, lie within the galactic plane. There is an increased risk of blending when observing within the galactic plane

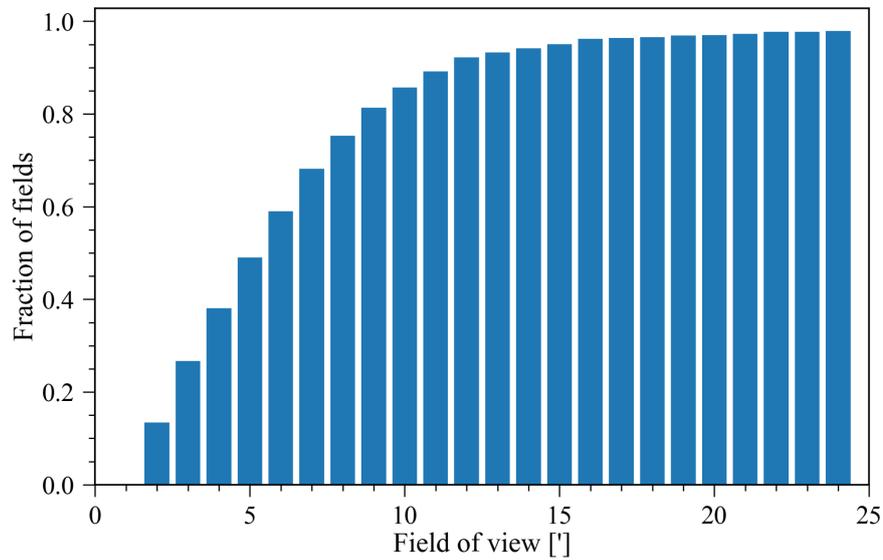


Figure 4.16: Fraction of usable fields from the SPECULOOS target list, at different square fields of view, capable of producing an artificial light curve with a 10 minute binned precision which still maintains a $\text{SNR} > 3$ (for the metric defined for Figure 4.8) for the target star's differentiated light curve.

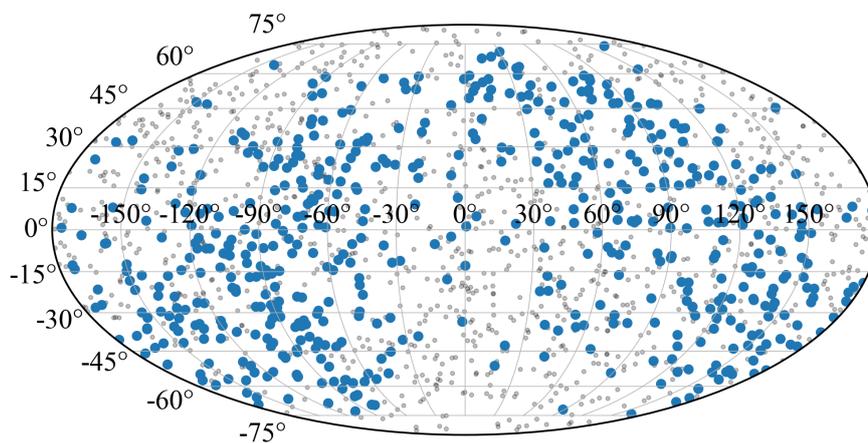


Figure 4.17: A mollweide projection of the night sky for all the SPECULOOS targets' coordinates shown as small grey points, with the larger blue points showing the usable field of views at a field of view of 6' (from Figure 4.16)

given the higher angular density of stars. Additionally, SPECULOOS targets have high proper motion, and whilst they were previously identifiable via surveys like 2MASS, their visibility may have changed to be obscured/blended today. This work was not extended to probe issues with blending. In future, if one were to be faced with blending issues, vetting with more detailed catalogues created by Gaia will be necessary, and potentially exploring the use of PSF photometry (Bradley et al. 2016) which is not currently used by the SPECULOOS-South pipeline.

4.6 Conclusions

Through the modelling presented in this chapter, the proposed NIR instrumentation is expected to provide better binned photometric data than the existing instrumentation for targets below ~ 2550 K, in some cases by $\times 2-3$. A camera from Princeton Infrared Technologies (1280SciCam model, discussed in Section 4.1.2) was identified, combined with a wide pass, PWV effect suppressing filter (zYJ, described in Section 4.2). The proposed NIR instrumentation is expected to provide dark limited photometric precision. For the median case, discussed in Section 4.4.1, its modelled photometric precision was found to be most sensitive to variations in dark current, site seeing, and pixel plate scale on sky.

With the modelled performances, we expect to be able to detect Earth-sized induced transits for coolest subset of the SPECULOOS target list, as discussed in Section 4.4. However, with a smaller field of view, than the existing instrumentation by 50%, the availability of suitable comparison stars for the differential photometry process is then reduced, as discussed in Section 4.5.

Nonetheless, the proposed instrumentation presents a strong argument as a successor, or at the very least a complement, to the existing instrumentation, not only through showing better photometric performances, but likewise as instrument which can probe further into the infrared - as demonstrated at the end of this thesis. It similarly presents itself as a lower-cost, more appropriate for robotic operations, than the traditionally used HgCdTe based detectors. Similarly, this detector and the proposed filter presents an interesting avenue for the future of infrared ground-based astronomy.

As a result of this work, the 1280SciCam from PIRT and a zYJ filter set were procured. The procurement for the filters required quotes from multiple manufacturers, with Brinell Vision^e finally selected as they offered the best average transmission over the other companies who responded to the tender. We also requested a J_{MKO} and an H-long filter as part of the procurement, detailed in Appendix B. For the 1280SciCam, it took around 18 months to

^e<https://brinellvision.com/>

negotiate a contract (Appendix A) which we, the University of Cambridge's finance department, and PIRT were all content with. It then took a further 6 months for delivery to Cambridge. The agreed upon specifications shifted slightly from the modelled values in this chapter, but not to an extent where the conclusions drawn here would have differed.

The following chapters detail a lab based characterisation of the 1280SciCam, preparation for on-sky, and finally some first light results.

CHAPTER 5

CAMERA CHARACTERISATION

The procurement process of the camera and filters began in October 2019, but was severely delayed due to administrative and pandemic related issues. It took just under 2 years before we received equipment from the camera manufacturer Princeton Infrared Technologies (PIRT), and the first filter set from Brinell Vision. The delays were sufficiently long enough that PIRT's science camera offering had changed to their 2nd generation of the 1280SciCam (where we were expecting the 1st generation), which they were kind enough to upgrade our order to without additional charge. The fundamental sensing hardware was the same, but with more human friendly upgrades and increased reliability to the software control.

The camera arrived in late August 2021, and as per the procurement contract (Appendix A), I had one month to verify the core parameters of the camera to ensure the equipment arrived without damage from shipment from the United States of America. Prior to the delivery, a report of the core parameters and a few additional non-standard tests were performed by the manufacturer which verified our requested specifications, specifically in the "LOWNOISE" mode. However, as was found out during our verification period, the "LOWNOISE" mode had issues which were not caught by the manufacturer. The "default" operating mode was subsequently chosen as primary operation mode for use after accepting a higher read noise. The results from this investigation are detailed in this chapter.

The work in this chapter was performed over the course of 3 consecutive months, in parallel with preparations for on sky tests at the Paranal Observatory, Chile, detailed in the next chapter.

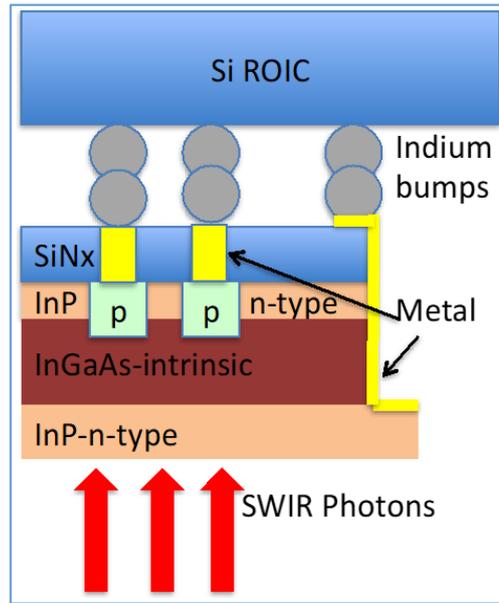


Figure 5.1: InGaAs detector array structure attached to Si ROIC circuit. Diagram provided by PIRT.

5.1 Camera operation

Before delving into characterisation, I will describe the camera and its modes of operation. The camera contains a 1280x1024 12 μm lattice-matched indium gallium arsenide ($\text{In}_{0.53}\text{Ga}_{0.47}\text{As}$) focal plane array (FPA), hybridised onto a silicon CMOS readout integrated circuit (ROIC). The detector array is backside illuminated with a >99% fill factor, meaning the light comes in the side opposite of where the photodiodes are formed. The hybrid imager structure is shown in Figure 5.1. The InGaAs detector array has had an InP substrate removed, allowing a larger visible response than most commercially available InGaAs detectors.

Given the CMOS architecture, the camera does not require a mechanical shutter as it permits fast global readout. Interestingly, in its default operation, images from the camera are produced continuously as opposed to on an external trigger, at a frequency of $1/(\text{frame time})$, where the frame time is defined as the time between the beginning of an exposure, to the time the frame is read out. The exposure time on the other hand is the integration time of photons on the detector. It was assumed that this operation mode was default given PIRT's largest customers are in machine vision and military applications. The camera could be programmed to start an exposure on a trigger, but this was found to have a higher read noise (discussed in Section 6.1.3).

The time difference between the frame time and exposure time was user settable, as long as the frame time was greater than the exposure time plus the read time. The time resolution

of setting the exposure and frame time was $1/(\text{detector's clock frequency})$, where the clock frequency was set at 15 MHz for our unit. The manufacturer provided three “Specific Operating Conditions” (SOCs) for our unit, which had differences in the bias voltage and undisclosed differences in the use of the read out electronics. The default mode operated at a bias voltage of 3.1 V, the second provided mode operated at 2.8 V, which was suggested to have a lower dark current than the default mode (disputed in Section 5.3.2), but at the cost of slightly higher read noise than the default mode. The last mode, titled “LOWNOISE”, was suggested to have the same dark current as the default mode but with read noise that met our requirements. As mentioned previously, the “LOWNOISE” mode came with problems that were not caught by the manufacturer.

Communication with the camera was performed via the Camera Link protocol, a protocol which has presented a standard for the communication between cameras and frame grabbers (specialised computer hardware, often as a PCIe card). It defines the methods for data transfer, camera timing, serial communications, and real-time signalling to a camera. A single Camera Link cable (called base configuration) permits a bandwidth of 255 MB/s and up to 850 MB/s for two cables (medium-full configuration). For our use case, the Camera Link base configuration was more than sufficient given the expected exposure times for the SPECULOOS targets. Each image, from the 1280SciCam was expected to be 2.2 MB (14 bit/pixel), and with the read time set to ~ 0.1 s (see Section 6.1.3) the fastest rate needed would be 22 MB/s. With 14 bits per pixel, the digital output was within the range 0 to 16383 ADUs. As discussed in Section 5.6, with ~ 2800 ADU taken by the bias level and ~ 2000 ADU taken before the pixels reached saturation, giving an effective range of ~ 11500 ADU, 13.5 bit depth. The serial communication was used to set the operating SOC, exposure and frame time, and used to set and read the FPA temperature, amongst other settings and sensors not discussed here. Setting the FPA temperature was limited to a lower bound of -60 °C by the manufacturer due to the chosen temperature sensor used at the FPA having a sensitivity limit to -65 °C.

Cooling the FPA (which was enclosed in a vacuum) was performed by a 4 stage stacked peltier cooler attached to a copper cold finger. The copper cold finger could have its heat dissipated via two methods, first via an attached heatsink and high speed fan, and second via liquid cooling, or a combination of both. Fan only cooling would permit the FPA temperature to comfortably hold at -40 °C at room temperatures, with the camera requiring <60 W of power to reach this temperature differential, where <5 W is used for the non-cooling portion of the camera. To reach -60 °C however, it required either an ambient temperature below 15 °C or the introduction of liquid cooling.

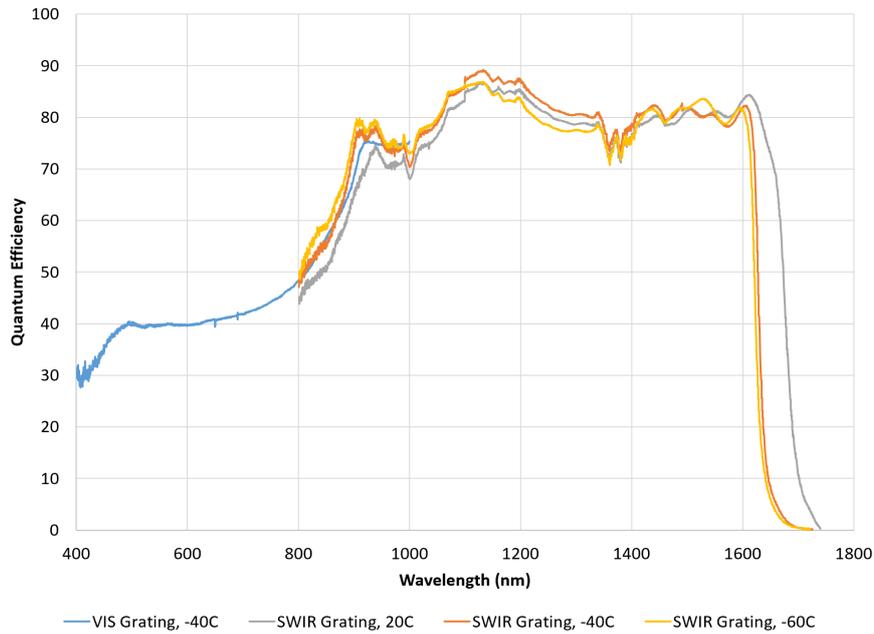


Figure 5.2: Measured QE curve of the 1280SciCam by the camera manufacturer at three different FPA temperatures, +20 °C, -40 °C, and -60 °C. Figure credit: PIRT.

5.2 Manufacturer’s report

Prior to verification in our lab, we received a verification report from the manufacturer. It met all (read noise value was disputed later) the parameter requirements detailed in Appendix A. They performed 3 additional measurements which we did not have the facility to perform within our lab due to a lack of resources or time. These measurements were quantum efficiency (QE) curves, pixel persistence, and environmental temperature dependence of cooling performance.

The high resolution QE curve shown in Figure 5.2 was obtained by the manufacturer, which I did not have the resources to do in our lab. The QE was measured at three different focal plane array (FPA) temperatures (+20 °C, -40 °C, and -60 °C), for the complete camera build, i.e. not corrected for the fused silica window in front of the detector. The accuracy of the QE measurements was assumed to be within $\pm 10\%$ based on discussions with the manufacturer. One known feature of decreasing the FPA temperature was the QE cutoff shifting to lower wavelengths, which has been observed with other InGaAs based detectors (Seshadri et al. 2006) and similarly with Si based detectors (Krishnamurthy et al. 2019). At -60 °C, the cutoff (measured at a QE of 40% to the closest nanometer on Figure 5.2) was 1624 nm, whereas at +20 °C the cutoff was 54 nm higher at 1678 nm, and at -40 °C the cutoff was 1630 nm. A decreasing cutoff provides the benefit of lowering any environmental thermal influence on the detector, but at the disadvantage of lowering the sensitivity within the H band. A similar

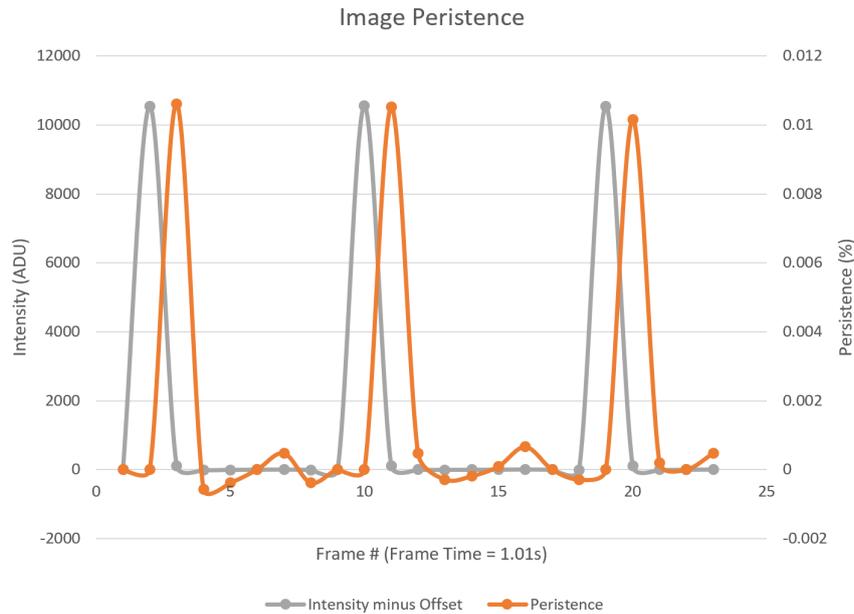


Figure 5.3: Persistence measurements performed by the camera manufacturer, where a light source was triggered to be on for the one frame and then shut off for the next 7 subsequent frames, and then repeated. In grey, median ADU value of the frame (bias average subtracted from each pixel). In orange, the persistence as a percentage of flux with respect to the previous frame. Figure credit: PIRT.

shifting effect was seen at ~ 900 nm, the point where InP substrate removal permits QE into the visible domain. Without the InP substrate removal, the QE of an InGaAs detector would instead cut on at ~ 900 nm (Seshadri et al. 2006). The sharp decline observed with our unit at this wavelength may suggest not all the InP substrate was removed. There is no conclusive evidence on the overall QE changing as a function of temperature given the error in the QE measurements. The mean QE across the zYJ filter range was 77%, around the value used as part of the modelling in the previous chapter.

Figure 5.3 shows the persistence of the whole frame as measured by the manufacturer, where a light source was triggered to be on for the one frame to reach $\sim 80\%$ of the frame's wells and then shut off for the next 7 subsequent frames, and then repeated 2 times more. In each case, the persistence of the frame after the frame exposed was around 0.01%, and returned near baseline in the following frame. The manufacturer did not detail the nature of light exposure. Nonetheless, it was indicative enough that the persistence was not going to cause an issue for the use case of transit photometry, or for the case of long term variability studies. For events such as flares, one may need to account for the effect of persistence. An unexplained repeating shape was seen in between light exposures, but at a level which did not pose a concern for our use case.

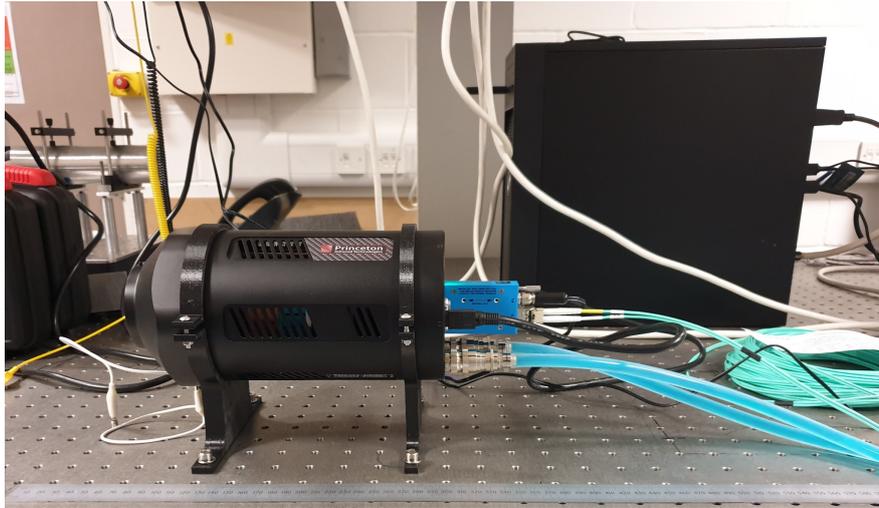


Figure 5.4: Initial lab set up with the 1280SciCam, held in place with 3D printed mounts on an optical bench. Cooled with liquid (set at 15 °C) from a Solid State Cooling Systems UC160 chiller, connected with the provided 12 V power supply and Camera Link fibre converters.

The final set of tests were to verify that the equipment could operate within the ambient temperature range expected in Paranál. Here, the manufacturer had access to a temperature controlled enclosure which could simulate a temperature range of -8 °C to $+30\text{ °C}$, which the camera was shown to operate within and able to maintain FPA temperature stability.

5.3 Parameter verification in Cambridge

The equipment arrived undamaged at our laboratory in late August 2021. The delivered equipment consisted of the PIRT 1280SciCam 2nd generation camera, compatible liquid tubing for external cooling with a liquid chiller (Solid State Cooling Systems UC160, acquired separately), a computer with a Camera Link compatible frame grabber (Bitflow Axion 1xE) and demo software, Camera Link fibre converters (EDT VisionLink RCX pair), in addition to an array of cables for power and communication. Once all set up as per PIRT's instructions, their demo software was used to acquire images of appropriate exposure lengths to conduct tests.

Figure 5.4 shows the initial lab set up for verifying the read noise and dark current of the camera. The camera was fixed onto an optical bench with 3D printed mounts, with the body of the camera's body grounded to the table (which was grounded to earth). Communication with the camera was performed via the Camera Link fibre converters, to have the same hardware as if we were to operate on sky. Images were acquired with PIRT's demo software, a Windows desktop based interface. The software was unfortunately limited in its ability to perform image

acquisition, where the process to set the exposure or save images could not be automated, with the exception of one mode which could increase the exposure time over a set number of images (but not frame time). More annoyingly, it could not save any metadata with the images (for example exposure time, FPA temperature, or time of acquisition). This made the process of verifying basic camera parameters unnecessarily cumbersome. On top of this, the software would occasionally crash for long exposures and had a handful of bugs which took a few weeks of my time to correct for. Nonetheless, eventually all the necessary data was acquired to measure the read noise and dark current of the camera under its various operating modes. Later, an ASCOM based driver was developed for use on sky, described in the following chapter.

The following subsections use the gain values provided by the manufacturer ($5.0\text{--}5.8\text{ e}^-/\text{ADU}$, different depending on the mode of operation), as measured with the Photon Transfer Curve method (Janesick 2007).

5.3.1 Read noise

When a pixel is read out, there is an associated noise with reading its voltage value. The noise primarily derives from spurious electrons introduced by the analog to digital conversion and connected circuitry. The 1280SciCam utilises Correlated Double Sampling (CDS) as its method to minimise the effect of read noise, where a pixel's voltage is read after reset and once again after an exposure. The difference between these values gives the pixel's value, and minimises thermal derived noise. Other methods, such as Fowler sampling (Fowler & Gatley 1990), as mentioned in Section 1.2.2, can significantly reduce the read noise by reading multiple times during an exposure, but was unfortunately not utilised by the manufacturer.

To measure the read noise, I took the standard deviation of each pixel from 100 bias frames taken with an exposure time of $111\text{ }\mu\text{s}$ and frame time of 1.11 s , with the lens cap on, taking the median value of all the standard deviations as the read noise. It was originally intended to expose for $100\text{ }\mu\text{s}$ and at a frame time of 1 s , however a bug (discovered after the fact) in the acquisition software increased the exposure and frame time by a factor of $(10/9)$ due to the software assuming a higher sensor clock speed than our unit. There was not an explicit bias exposure time suggested by the manufacturer, nor could we expose for 0 s . We instead decided to repeat as close to the conditions as performed in the manufacturer's report. The read noise measurements were repeated for the different SOCs at both $-40\text{ }^\circ\text{C}$ (shown in Table 5.1) and $-60\text{ }^\circ\text{C}$ (shown in Table 5.2).

As the 1280SciCam is CMOS based, each pixel will have its own read noise given there's a digitization circuit tied to each pixel. Here, I report the median value for the detector converted to electrons, with the gain values (different for each SOC and temperature) provided by the manufacturer. At $-40\text{ }^\circ\text{C}$, the read noise measured in our lab for the default mode and the 2.8 V

Table 5.1: Read noise and background flux values measured at a FPA temperature of $-40\text{ }^{\circ}\text{C}$ for different SOCs.

| | Read noise (our results) [e^-] | Read noise (from report) [e^-] | Background (our results) [$e^-/\text{s/pixel}$] | Background (from report) [$e^-/\text{s/pixel}$] | Gain (from report) [e^-/ADU] |
|---------------------|--|--|---|---|---|
| Default (3.1V bias) | 94.96 | 98.0 | 662.84 | 668 | 5.8 |
| LOWNOISE | 41.52 | 39.2 | | | 5.0 |
| 2.8V bias | 95.43 | 99.2 | 706.32 | 579 | 5.6 |

Table 5.2: Read noise and background flux values measured at a FPA temperature of $-60\text{ }^{\circ}\text{C}$ for different SOCs.

| | Read noise (our results) [e^-] | Read noise (from report) [e^-] | Background (our results) [$e^-/\text{s/pixel}$] | Background (from report) [$e^-/\text{s/pixel}$] | Gain (from report) [e^-/ADU] |
|---------------------|--|--|---|---|---|
| Default (3.1V bias) | 86.45 | 91.4 | 251.09 | 260 | 5.5 |
| LOWNOISE | 39.66 | 40.7 | | 260 | 5.1 |
| 2.8V bias | 92.06 | 93.3 | 260.24 | 185 | 5.5 |

bias mode was lower than that was measured by the manufacturer. These values however were over twice what was asked in the procurement requirements. The LOWNOISE mode yielded a read noise which was within our requirements, but marginally higher than what was stated in the manufacturers' report. This mode however had issues at both $-40\text{ }^{\circ}\text{C}$ and $-60\text{ }^{\circ}\text{C}$, which were discovered in the dark current measurements, described in the following sub-section.

At $-60\text{ }^{\circ}\text{C}$, the read noise values for each SOC were lower than that seen at $-40\text{ }^{\circ}\text{C}$, as were the gain values for the non-LOWNOISE modes. Later, described in Section 5.4, I measured the gain value for the default mode at $-60\text{ }^{\circ}\text{C}$ in our lab, where I found a lower gain of $5.092\text{ }e^-/\text{ADU}$ versus the $5.5\text{ }e^-/\text{ADU}$ that was reported by the manufacturer. Nonetheless, it was enough to indicate that only the LOWNOISE mode met our requirements, even with the differences in the gain. Characterisation of another 1280SciCam unit, performed by [Birch et al. \(2022\)](#), noticed a similar trend of decreasing read noise with temperature at a fixed gain of 4.3, and attained a lower read noise of $79.9 \pm 0.3\text{ }e^-_{rms}$ at their lowest operating temperature of $-40\text{ }^{\circ}\text{C}$.

The lower read noise values achieved in the lab when compared to the manufacturer's values could not be explained. A more stable mains voltage was perhaps attributed. However, [Birch et al. \(2022\)](#) recorded the read noise of their unit with a conditioned and unconditioned power source, but did not notice a difference at their coolest operating temperature ($-40\text{ }^{\circ}\text{C}$). The second possible source was due to the exposure time and frame time used - whilst very similar to the manufacturer's value, it was noticed that the bias structure would change slightly if I varied the relationship between the frame time and exposure time. The manufacturer did not have recommendations on the best exposure to frame time relationship to use to minimise the read out noise.

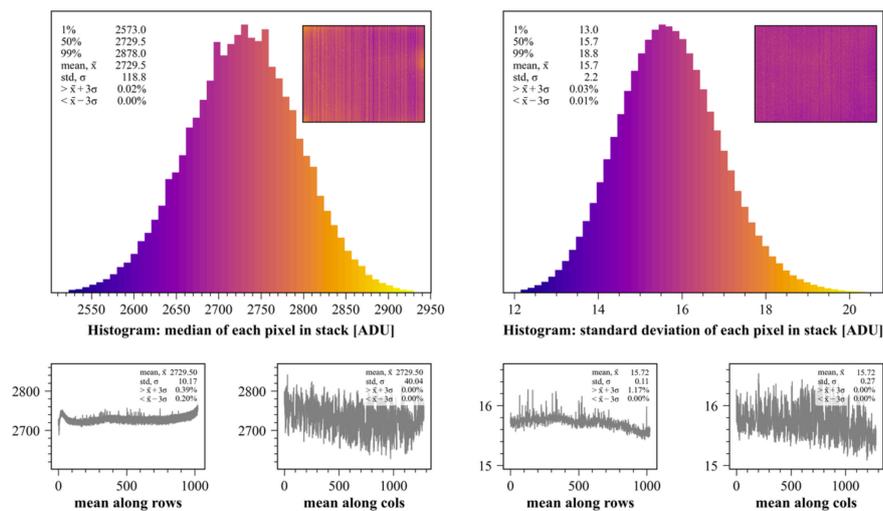


Figure 5.5: Master bias and read noise histograms (showing the 0.1 to 99.9 percentiles) measured with the default SOC at -60°C . The mean along rows and columns (cols) of the respective master bias and read noise image, inset into the histogram plots with matching colour scale.

Figure 5.5 shows histograms of the master bias frame and the read noise of each pixel for the default SOC at -60°C , from the measurements described above. We can see the median bias value was at 2729.5 ADU, with the 1% and 99% percentile at 2573 and 2878 ADU respectively. There was a notable banding structure in the master bias along the columns, with a slight general decrease in value from left to right (direction from inset image in Figure 5.5). A concentrated glow, similarly seen in [Birch et al. \(2022\)](#), on the right side of the detector is seen, assumed to be from the ROIC. Likewise, a glow from the bottom and top left was visible. The variation along the rows was smoother when compared to the column variation, but with a notable fast rise and dip across the first ~ 50 rows from the top. There were spikes in some of the row means, attributable to hot pixels, discussed in Section 5.7. Similar spikes were seen along the columns but masked by the banding structure.

The variation in read noise, shown on the right of Figure 5.5, showed a similar banding structure along the columns. The read noise had a median value of 15.7 ADU (86.45 e^- with an assumed gain of 5.5, or 79.91 e^- with our measured gain of 5.092), with the 1% and 99% percentile at 13.0 and 18.8 ADU respectively. Unexpectedly, around the assumed ROIC glow regions, the read noise was observed to be slightly lower than the surrounding areas. The other SOC's displayed similar bias structures and histogram distributions, but shifted due to slightly different median bias and read noise levels. However, the LOWNOISE SOC's bias distribution was wider by 44%, whilst only having a 3% higher bias.

5.3.2 Dark current

The relative proportions of the elements that make up the 1280SciCam's detector, indium gallium arsenide, dictate the cut off and stability of the material, which in turn affects the dark current produced per pixel. PIRT have chosen a stable configuration of InGaAs ($\text{In}_{0.53}\text{Ga}_{0.47}\text{As}$) which minimises the strain in the material, thus minimising the dark current that these elements together would likely produce. In this configuration, the cut off wavelength was found to be around 1624 nm at $-60\text{ }^\circ\text{C}$, as shown in Figure 5.2. At a fixed detector temperature, the dark current was expected to be a constant value. If not measured to be constant over a series of measurements at different environmental temperatures and a fixed detector temperature, then a varying thermal flux is likely to be present in the readings. The 1280SciCam was found to be sensitive to thermal radiation (but only in later work, see Section 6.2.1). Consequently, during dark exposures it was measuring a combination of dark current and thermal flux, in combination I have named this as the detector's "background flux".

To measure the detector's "background flux", I recorded seven images at increasing exposure times (0.00 to 77.78 s, the time at which saturation would begin at $-40\text{ }^\circ\text{C}$) at a fixed frame time 1 s longer than the longest exposure time. I then fit the trend of the average frame value with exposure time before saturation. Tables 5.1 and 5.2 detail the results. These measurements were performed with the lens cap on at room temperature.

The default SOC was found to have the lowest background flux followed by the 2.8 V bias mode under similar environmental temperatures. For the default SOC, I recorded a background flux marginally lower than the manufacturer's value, whereas the 2.8 V bias mode was notably higher than the manufacturer's value. The discrepancy here was unexplained. The scaling between the $-40\text{ }^\circ\text{C}$ to $-60\text{ }^\circ\text{C}$ did not follow the halving of dark current expected every $6.6\text{ }^\circ\text{C}$ found with data from [Sullivan et al. \(2014\)](#). Where we should have expected $\sim 80\text{ e}^-/\text{s}/\text{pixel}$ as opposed to $251\text{ e}^-/\text{s}/\text{pixel}$ with the default SOC at $-60\text{ }^\circ\text{C}$, following the value of $663\text{ e}^-/\text{s}/\text{pixel}$ at $-40\text{ }^\circ\text{C}$. It was later confirmed that this discrepancy was indeed from a thermal background, more in Section 6.2.1, and the real dark current was closer to $\sim 60\text{ e}^-/\text{s}/\text{pixel}$. Closer matching the dark current of $68\text{ e}^-/\text{s}/\text{pixel}$ found in [Birch et al. \(2022\)](#), who instead measured the dark current within a cooled chamber, maintaining a lens cap temperature of $-10.6\text{ }^\circ\text{C}$. For comparison, a HgCdTe based detector (with a tuned cut off of $1.75\text{ }\mu\text{m}$) at the same FPA temperature and pixel pitch would have yielded $\sim 4000\text{ e}^-/\text{s}/\text{pixel}$ ([Beletic et al. 2008](#)) without thermal contribution. The thermal contribution for a HgCdTe would be $\sim 5\text{-}10$ higher than an InGaAs detector given the longer cut off wavelength based on calculations made in Section 6.2.1.

Figure 5.6 shows a histogram of the background flux value of each pixel on the detector. There was no distinct region on the detector which had a higher count rate, with no banding or

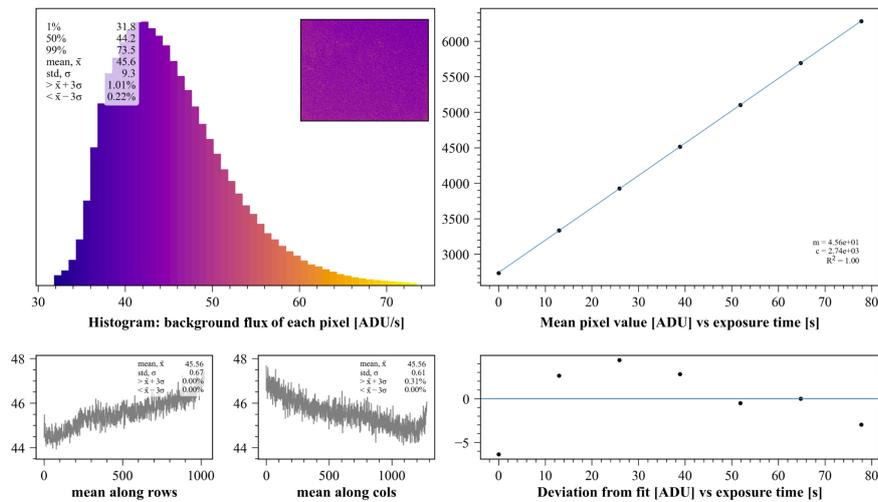


Figure 5.6: Background flux measurement of the default SOC at $-60\text{ }^{\circ}\text{C}$. Top left: Histogram (showing the 1 to 99 percentiles) of background flux per pixel in ADU/s, with inset image with matching colour scale. Bottom left: The mean along rows and columns (cols) of the background flux per pixel inset image. Top right: Mean frame values in ADU against the respective exposure times with linear fit. Bottom left: Deviation from the linear fit made in the top right plot.

ROIC glow visible as seen in the bias image (inset in Figure 5.5). It instead had a number of randomly scattered dead pixels and pixels with higher dark current. The histogram distribution of background fluxes was not gaussian, but instead skewed to have a tail towards higher values. In the inset image of Figure 5.6, showing the background flux distribution, there was a slight gradient from the bottom left towards the top right, possibly due to a thermal gradient from the environment or cooling mechanism. The mean increase of the background flux was seen to be linear with time, with the deviation from the fit being minor with the exception of the 0 s bias exposure. The fit was seen to be significantly less linear for the LOWNOISE mode, where the background flux was seen not to have a relationship with exposure period, as if the mode was not subject to the effects of background flux. This behaviour was not seen by the manufacturer as they did not properly test all the SOC's for the length of exposures I did here, discovered after the discussions about our results. Given the non-linear behaviour, it was necessary to then assess the detector's SOC's behaviour with light.

5.3.3 Point source tests

To help us discern which SOC would be most suitable for observing on-sky, I had to simulate an on-sky like scene to image in the lab. The resulting set up (shown in Figure 5.7) produced a point-like source of light in the centre of the image. This was achieved by using a non-NIR

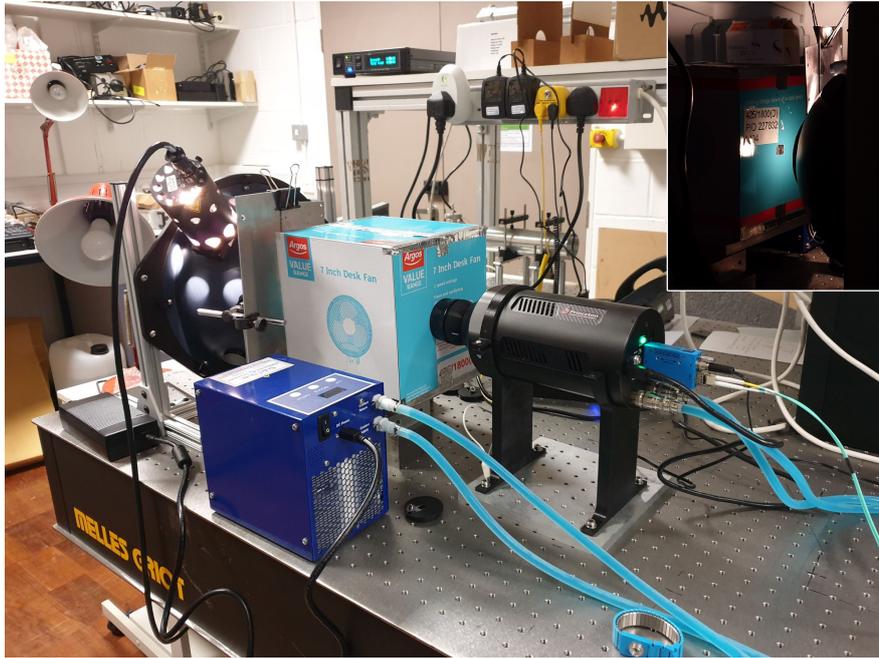


Figure 5.7: Main image: Lab setup to image a point source of light. 1280SciCam camera connected as in Figure 5.4, mounted on a taller set of 3D printed mounts to be inline with the centre of the integrating sphere. A 28 mm F/8 lens (not optimised for NIR) was mounted and connected to a boxed enclosure. The box was sealed with aluminium tape to minimise stray light. The opposite side of the box (top right inset image) had a small cut out covered with aluminium tape which had a < 1 mm hole made with a pinhead. The inside of the box, facing the camera, was covered in matt black tape, to minimise any reflection of light from the point-like source. The lens was focused onto this point-like source where the point-like source covered width of 8 pixels.

lens to focus onto a pinhole which was illuminated with a stable light source. The scene was enclosed in a box, so as to minimise the influence of stray light. The imaged scene was covered in matt black tape, to minimise any reflection of light from the point-like source.

The light source was a 12 V 100 W halogen lamp illuminating a 12" Labsphere integrating sphere, output through a 1" hole to illuminate the pinhole. The output of the halogen lamp into the integrating sphere was attenuated using a variable mechanical aperture to allow exposures of the point-like source up to ~ 25 s without saturation. The halogen lamp was left to thermally stabilise for 50 minutes before taking exposures in a dark room. Beginning with the default SOC, with the FPA at -60 °C, I exposed the scene for 11 exposure times between 1.11 s and 33.33 s, changing to the next SOC once the last exposure was completed. Measuring the point's integrated flux between 1.11s and 10.78 s, with the surrounding background subtracted, showed similar flux levels for the default and 2.8 V bias voltage SOC. The LOWNOISE SOC recorded a flux rate a factor of 5 lower than recorded with the default and 2.8 V bias SOC's,

but more concerningly it began to deviate from the linear trend about 10 s earlier despite none of the pixels saturating, where it had been seen to saturate with brighter scenes. Behaving as if flux was being lost at longer integration times. This then firmly ruled out the LOWNOISE mode. This led us to conclude that the default mode provided the best and most reliable results for science operations after demonstrating a lower dark current and read noise than the 2.8 V bias level SOC. The remainder of the measurements in this and the following chapter were performed with the camera operating with its default SOC.

5.4 Gain

The gain dictates the overall conversion factor between the analogue to digital units to the electrons held within a pixel. The manufacturer provided gain values for each of the SOCs, but we wanted to have a more accurate understanding in our lab, specifically for the default SOC we will be operating on sky with.

I measured the gain using a photon transfer curve (PTC) with the mean-variance method, as done in [Birch et al. \(2022\)](#), which involves the measurement of signal variance, σ_{tot}^2 , as a function of changing flux. Here, leveraging the fact that the flux and dark current is Poisson distributed, the variance from the difference between two flat field frames of the same exposure will be a function of twice the Poisson noise, σ_P , and read noise, σ_{RN} , leading to:

$$\sigma_{tot(electrons)}^2 = \sigma_{P(electrons)}^2 + \sigma_{RN(electrons)}^2, \quad (5.1)$$

$$\sigma_{tot(electrons)}^2 = gS + \sigma_{RN(electrons)}^2, \quad (5.2)$$

$$\sigma_{tot(ADU)}^2 = \frac{1}{g}S + \sigma_{RN(ADU)}^2, \quad (5.3)$$

where the S is the measured signal (bias offset removed) in ADU and g is the gain in e^-/ADU . The difference was taken because it removes any fixed pattern noise present in the frame. Thus, with enough flat fields, one can measure the gain by taking the inverse of the gradient of various signal means against their variances. However, if the detector is non-linear, this method can lead to erroneous estimations of these parameters ([Pain & Hancock 2003](#)). Likewise, the read noise may have a signal dependence ([Levski et al. 2021](#)), similarly leading to incorrect estimations of parameters using this method.

To produce flat fields, I projected a near-uniform light distribution onto the detector's surface using the integrating sphere described in the previous subsection. Here, to permit integrations to the lengths expected to be used on sky, the halogen bulb was changed to the lowest power

rating locally available (Osram branded 12 V 20 W). It was also set to operate at the lowest current level where the lamp was still visibly outputting light by eye. This was achieved using Labsphere's stabilised DC power supply, where a current limit of 0.5 A at 12 V was chosen (the lamp eventually stabilised at 1.363 V and 0.541 A at this current limit, as measured by Labsphere's power supply control software). Its output into the sphere was further lowered using the manual mechanical attenuator in front of the lamp. The detector was also placed at a fixed distance away from the integrating sphere's 1" output port, which eventually permitted ~25 s integrations without saturation.

The uniformity of an integrating sphere's output varies as a function of the output port diameter, the distance from it, and size of the detector. Following Labsphere's documentation^a, the uniformity at the output port is uniform, but rapidly decreases at small distances away from the sphere's output but returns to uniformity at larger distances. In our instance, greater than 10 times the output port diameter for our size detector (13 × 16 mm) presented a uniformity of 100% by their metrics. However, to place the detector at this distance, an enclosure was thus required to shield the detector from any stray light from the environment.

To build a suitably designed enclosure, which also minimised reflections that could affect the uniformity, was a struggle given the time constraints at this point of the characterisation. Nonetheless, the resulting enclosure (described in Figure 5.8) was good enough to produce a flat field suitable for gain measurements, which placed the camera at 400 mm from the output port, 15.7 times the output port diameter. The uniformity could have likely been improved with a larger enclosure and the inclusion of a baffle, as done in [Astier et al. \(2019\)](#). To minimise any influence of light leaks, room lights were switched off, and any visible lights in the room (e.g. LEDs from equipment) were taped over. The setup was aligned using the mounting holes on the optical table.

Once everything was set up, I took flat fields of 100 different exposures, from 0 to 30 s, taking 25 images at each exposure period, resulting in a total of 2500 images over the course of 12 hours. By this point, I was no longer using the manufacturer's demo software, but instead using my own developed camera driver (described in next chapter) which overcame all the shortcomings of the demo software. Here, the frame time was set to be the exposure time + 100 ms. To produce the mean-variance PTC, I took the variances and means of consecutive frame pairs (after removing bad pixels - described in Section 5.7) of the same exposure period, yielding 24 pairs per exposure period. The resulting plot is shown in Figure 5.9.

The relationship shown in the mean-variance PTC was not completely linear, with large deviations from linearity seen at low flux as well as high flux levels close to saturation. Where

^aPage 16 of <https://labsphere.com/wp-content/uploads/2021/09/Integrating-Sphere-Theory-and-Applications.pdf>



Figure 5.8: Main image: Side view of the setup used for gain and linearity measurements. 12" integrating sphere with 1" output port connected to a cardboard box (internal dimensions: $W333 \times D390 \times H285$ mm), which was lined with Thorlabs black aluminium foil (4.5-6.5% reflectivity across the sensitivity range of the 1280Sci-Cam) to minimise internal reflections. Camera connected to the opposite side of the box at approximately 400 mm away from the output port. Top right image: Top down view, without the box's lid, showing the inside of the box.

the linear region was defined to be along the central flux range in between two features in the curve, where the line of best fit was made. Five distinct features were seen deviating from the line of best fit, one near the centre of the usable ADU range, then a pair nearly equidistant from the centre but with opposing deviations, followed on by another similar pair again with opposing deviations. The sources of these distinct deviations are unknown, but thought to have a signal dependence rather than an exposure period dependence.

Repeating the mean-variance PTC with a smaller set of exposures, but this time keeping a fixed exposure period of 1 s, and instead increasing the amount of flux output from the integrating sphere by varying the manual mechanical attentator in front of the halogen lamp. Since this was a smaller set of measurements, I by chance matched the mean flux level at the centre of one of the distinct features, where we could see a similar deviation. However, we can also see that the overall variance is marginally lower when compared to the previously determined line of best fit with the fixed flux-varying exposure dataset. The 1 s exposures still maintained a similar gradient, suggesting the gain was similar along the central flux region. Another notable difference was the region of saturation, which was ~ 500 ADU larger than

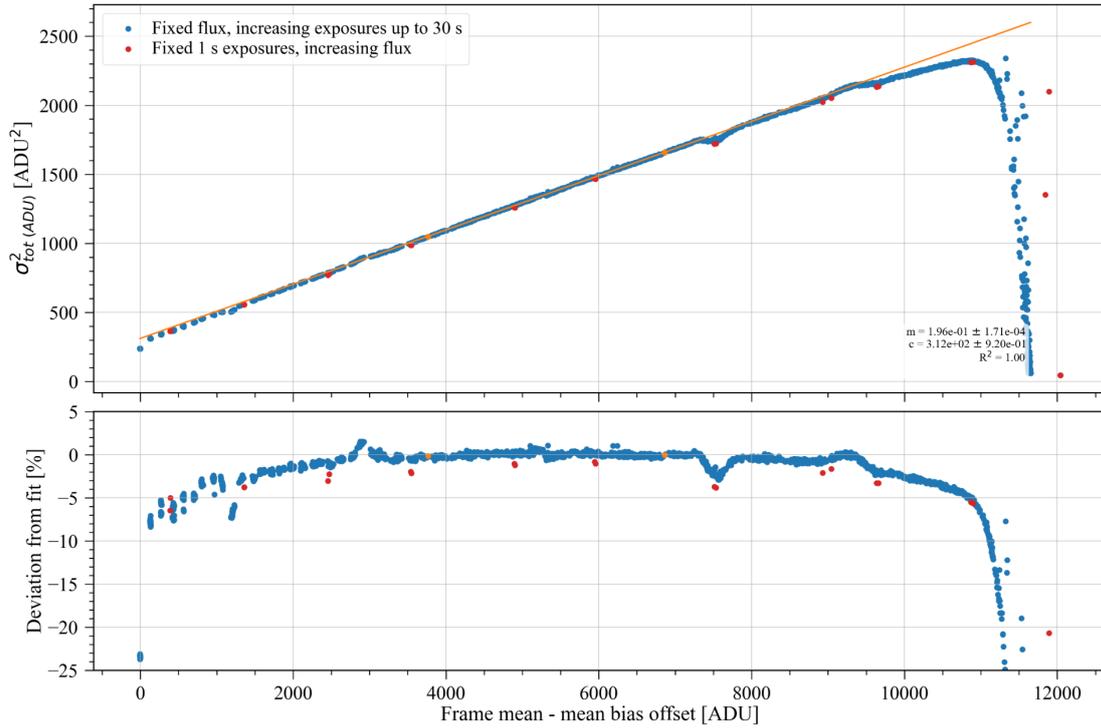


Figure 5.9: Mean-variance photon transfer curve, in blue for 2400 pairs of flat fields of increasing exposure up to 30 s after removing bad pixels of a constant flux source, and in red for 24 pairs of flat fields of 1 s fixed exposure period and increasing flux. Top: Variance of flat field pair difference (divided by 2) against flat field pair mean minus mean bias offset, with the line of best fit along a linear region of blue data points (denoted between the orange points). Bottom: Deviation from line of best fit for both datasets

the previous dataset - suggesting the well depth may be smaller for larger exposure periods. To note, the distinct features were also seen in other 0-30 s exposure gain tests, not shown in Figure 5.9.

Sources of deviations from linearity are unclear, whether it is a change in the gain, or change in the overall read noise. Changes in the gain with signal level are expected with CMOS detectors. According to [Pain & Hancock \(2003\)](#), the gain is expected to increase with increasing electron accumulation due to the reduction in the reverse-bias voltage of the detector. This would explain the deviations seen at low flux as well as high flux levels. There are methods to quantify the gain with non-linear characteristics ([Bohndiek et al. 2008](#)), but these methods require a calibrated photodiode which we did not have. A second method, in [Rest et al. \(2016\)](#), suggested a differential PTC algorithm, but assumes a fixed read out noise measured with bias frames, which we assume not to be true in our case.

Despite these deviations from the fit, we can calculate the gain along the seemingly linear

region to be $5.092 \pm 0.004 \text{ e}^-/\text{ADU}$, with a read noise of $89.950 \pm 0.154 \text{ e}^-$ ($17.664 \pm 0.026 \text{ ADU}$). The gain value I found was lower than reported by the manufacturer by 7.4%, and the read noise was higher in electrons and in ADU than the previously measured value but still lower than reported by the manufacturer in electrons. To see if these non-linearities pose an issue to linearity in recording flux, we can plot the received flux as a function of exposure time, described in the next subsection.

5.5 Linearity

Whilst a PTC can be non-linear, signal linearity can be unaffected (Downing et al. 2006). The manufacturer utilised a ROIC based on a capacitive trans-impedance amplifier (CTIA), which are more complex to produce but come with higher linearity than the source follower architecture traditionally used in other IR detectors (Beletic et al. 2008; Dangi et al. 2022). We attempted to quantify this, but unfortunately the error spread in the measurements made it difficult to draw a conclusive value for linearity.

I reused the setup and data used for the gain measurements. Plotting the recorded flux as a function of the exposure duration, shown in Figure 5.10, we can see the non-linearity did not follow the same extent as seen in Figure 5.9. However, the dataset in blue (of 100 different exposure periods, with 25 frames per exposure period) may have been subjected to varying flux levels from lamp variations and/or thermal background changes over the 11.3 hours data was recorded for. Shortly before this dataset, a smaller set of exposures (10 different exposure periods, with 6 frames per exposure period), shown in red, was acquired with the same set up over a shorter 0.3 hours, but the propagated error suggests these results did not yield much further information. If the negative deviations from the fit are to be true, the low flux deviations would not line up with the changes potentially seen with the gain changes (if a higher gradient at the lower fluxes can be assumed to be due to a lower gain) observed in the PTC curve - where for lower fluxes/lower gains, we would expect to see a positive deviation from the fit, and for higher fluxes/higher gain, a negative deviation as is somewhat seen.

The manufacturer produced a linearity curve with significantly shorter exposure times, with the highest exposure period of 30 ms, with 100 frames per exposure time, claiming an average non-linearity of 0.074% before saturation over 100 different exposure times. A slight positive deviation from linearity was observed at low exposures, but the deviation was smaller than expected for the gain estimated with the gradient at lower fluxes from the PTC curve. To form more conclusive answers, for when a camera like this is next in a lab, it is recommended to use a thermally insulated set up, and do a smaller set of different exposures (but with enough frames per set to minimise Poisson noise) so as to be less susceptible to lamp and/or thermally

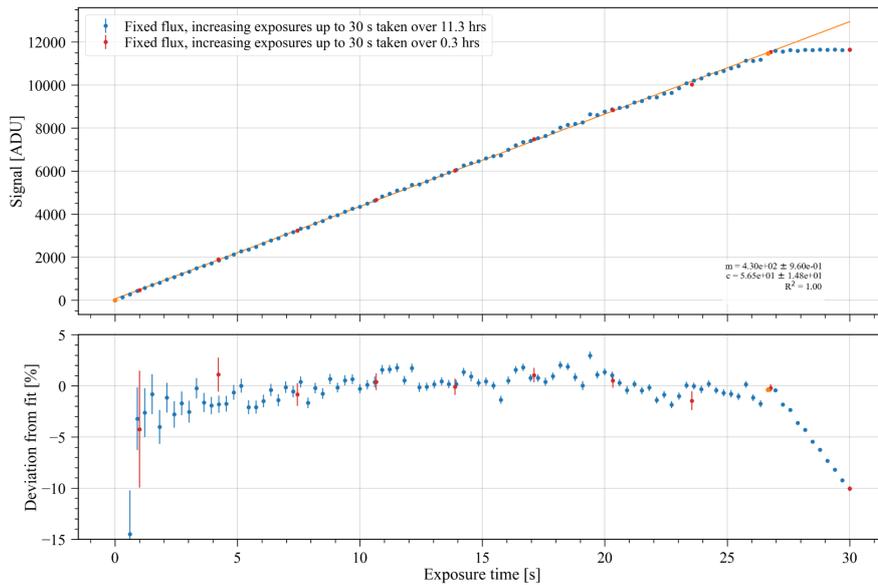


Figure 5.10: Linearity measurement attempts. Top: Signal (mean of frames minus mean bias, bad pixels removed) against exposure time. Line of best fit of blue data points, covering the range 0 s exposure to the exposure before saturation (denoted by the orange points). Bottom: Percentage deviation from line of best fit.

induced variations to produce more accurate measurements of linearity. In addition, one should use a calibrated photodiode to utilise non-linear gain calculation methods (e.g. [Bohndiek et al. \(2008\)](#)), testing a range of exposure times expected to be used on sky, but also shorter exposure times to saturation.

5.6 Usable full well

From the measurements made for gain and linearity, we can determine the usable full well depth. In this detector and gain configuration, saturation of the well occurred before ADC saturation, allowing us to quantify the average electron well capacity of a pixel. However, from the PTC measurements, we saw the level of zero variance differ between a 30 s and 1 s exposure. Unfortunately, the relationship between saturation level and exposure time was not extensively measured in the lab. Thus, we will stick to quantifying the values for these two exposures.

To quantify the full well, we take the difference between bias level and the point where the camera reaches saturation. Definitions for saturation are vague in literature, some suggesting it to be the point where the frame's mean deviates by 5% from the line of best fit from linearity measurements ([Hopkinson et al. 2004](#)). This method yielded a value of $\sim 11,610$ ADU for the data set which exposed up to 30 s, or $59,118 e^-$ with the gain previously calculated. However, at

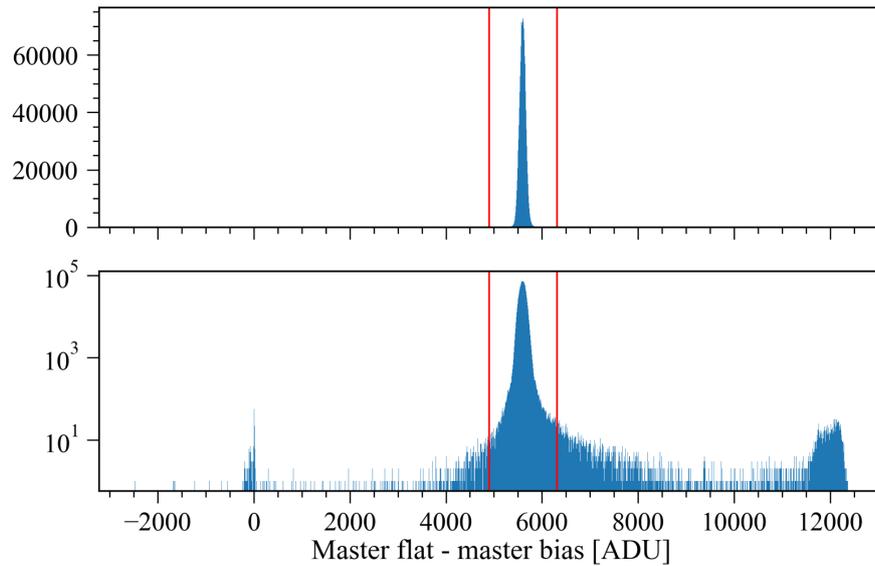


Figure 5.11: Histogram of master flat minus master bias, with 1638 bins (to give a width of ~ 10 ADU per bin). Red vertical lines show the region of the frame's mean $\pm 3\sigma$.

this point, the mean-variance PTC data set tells us the variance of the frame's signal at this level is close to 0 - suggesting the usable well is in fact at a lower ADU value. Instead, considering the point on the PTC curve where the variance starts to decline (i.e. when the local gradient becomes negative), this will give us an indication of when the average pixel starts to decline in usability. This method yields a saturation value of ~ 11090 ADU, or $56470 e^-$. For the 1 s mean-variance PTC, the 0 value variance point was ~ 500 ADU higher, suggesting an additional $2500 e^-$ for the average well depth. Although, with the gain changing as a function of signal level, this electron value is likely to be different in reality. Independently, with these ADU ranges, it suggests the true usable bit-depth of the detector is 13.44 bits for the long exposures measured here, or 13.50 bits for 1 s second exposures.

5.7 Bad pixels

A small proportion of the 1280SciCam's 1.3M pixels will unfortunately underperform. The cause for underperformance can be due to various reasons. In [Smith et al. \(2014\)](#), they list the following factors which may contribute individually, or in combination, to a pixel's underperformance: poor pixel bump bond connectivity, vignetting, addressing faults in the mux, severe sensitivity deficiency, poor signal linearity, low full well, excessive noise and high dark current. Here, we didn't strive to find the underlying cause, but instead to categorise if a pixel was not falling within 3σ of a frame's mean. I identified the bad pixels under two

Table 5.3: Percentage value of the frame’s underperforming pixels ($\pm 3\sigma$ outside of the frame’s mean) for a master flat and master dark (master bias subtracted). Results divided between the assessed frames, and the proportions below and above the usability thresholds, and combined total of underperforming pixels from both frames.

| | Lower [%] | Upper [%] | Total [%] |
|----------|-----------|-----------|-----------|
| Flat | 0.046 | 0.225 | 0.271 |
| Dark | 0.052 | 0.254 | 0.306 |
| Combined | | | 0.325 |

conditions, under the condition of exposure to light and darkness. A master dark frame was formed from 100 frames at 30 s exposures, taking the median of each pixel in the stack. A master flat frame, where the median pixel value of the master frame reached $\sim 50\%$ of the usable well defined in the previous subsection. The master flat was similarly formed from 100 frames at 14 s exposures, once again taking the median of each pixel. Each master frame was subtracted from a master bias frame (formed from 100 $100 \mu\text{s}$ exposures).

Figure 5.11 shows the histogram distribution of the final master flat, with the distribution looking similar for the master dark frame. Here, we can see a large proportion of hot pixels beyond the usable ADU well, and a smaller proportion of non-responsive pixels below 0 ADU, with tails of the main distribution larger towards larger ADU. This larger tail on the right of the distribution, beyond the 3σ line, was similarly seen with the dark histogram distribution, suggesting perhaps it was due to pixels with higher dark currents, and/or pixels with increased sensitivity. This could be diagnosed with the dark current measurements made in Section 5.3.2. Similarly, with the read noise measurements made in Section 5.3.1, we could similarly diagnose the pixels which have larger than expected read noises and mark them as underperforming. However, this was not looked into as part of this analysis.

Table 5.3 shows the proportion of underperforming pixels for the respective assessed frames, and the combined total number of identified bad pixels. More hot pixels were found with this detector than dead pixels. The exact number of hot pixels and dead pixels is difficult to define given the overlapping $>3\sigma$ distributions for the respective frames, where the some of the $\mu - 3\sigma$ for the flat frame overlapped with the some of the $\mu + 3\sigma$ of the dark frame for example. Nonetheless, with identified bad pixels we can plot their distribution on the detector, shown in Figure 5.12.

The distribution of the bad pixels were not clustered in unique regions, but rather near-uniformly distributed across the detector. This distribution is undoubtedly going to pose issues for the differential photometry process, discussed in the next chapter. Whilst the total percentage of 0.325% met our a requirement of $<0.5\%$, for future procurements, it is recommended to have a metric that quantifies the distribution of bad pixels as part of the requested specification.

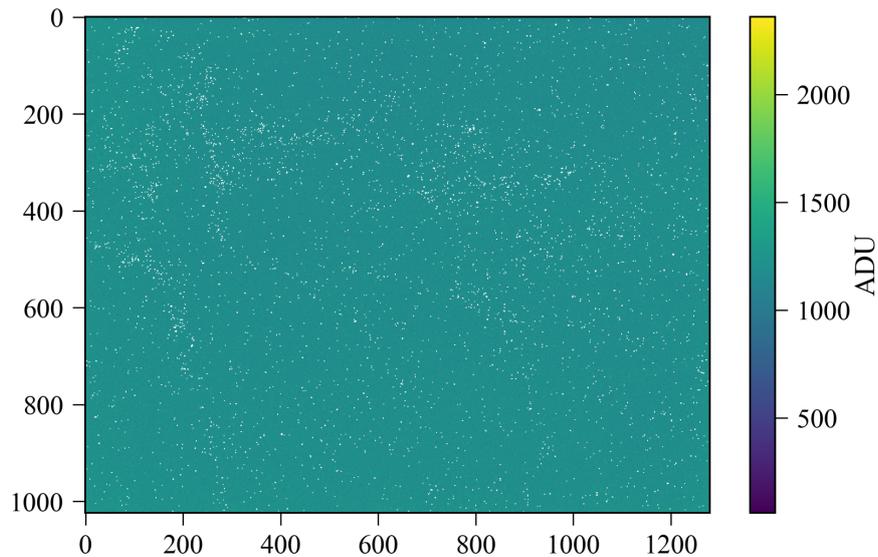


Figure 5.12: Master dark frame with master bias subtracted with the colour scale between $\pm 3\sigma$ of the frame's mean. White pixels denote the identified bad pixels.

5.8 Final characterised specification

In this chapter, I have characterised the core parameters of the 1280SciCam provided by Princeton Infrared Technologies, shown in Table 5.4. They met or exceeded all our requested specifications with an exception on the read noise, which was double than what was requested in our procurement document (Appendix A). The read noise was met in the provided 'LOWNOISE' mode, but this mode was found to have severe issues in Section 5.3.3 that would yield it unsuitable for photometric applications. This unfortunately meant that the camera will be read noise limited for the majority of observations made with a SPECULOOS telescope, as detailed in the following chapter.

The gain was found to be lower than reported by the manufacture by 7.4% in Section 5.4, which may have repercussions on the overall quantum efficiency reported for the camera unit we received. The non-linear shape of the mean-variance PTC curve in Figure 5.9 suggests that there are variations in gain as a function of signal level. However, despite this non-linearity, it was not reflected in linearity measurements made by the manufacturer, and similarly by the work in Section 5.5. However, the accuracy of our lab measurements were not sufficient to draw conclusive values. A more accurate linearity measurement is thus desired the next time a similar unit is in a lab environment.

Despite the lower gain, the read noise was suggested to be higher from the measurements made in Section 5.4 than the measurements inferred from a series of bias frames in Section 5.3.1.

Table 5.4: PIRT 1280SciCam characterised specifications for the default SOC.

| Parameter | Unit | Values | Notes |
|------------------------------|-------------------------|---------------|---------------------------------------|
| Array format | pixel | 1024 × 1280 | By design |
| Detector material | | InGaAs | By design |
| Detector substrate | | InP - removed | By design |
| Pixel pitch | μm | 12 | By design |
| Pixel fill factor | % | > 99 | By design |
| Data output | bits | > 13.44 | See Section 5.6 |
| Gain | e ⁻ /ADU | 5.092 ± 0.004 | Across central range, see Section 5.4 |
| Readout duration | s | ~0.1 | Set by driver, see Section 6.1.3 |
| Median dark current @ -60 °C | e ⁻ /s/pixel | ~60 | See Sections 5.3.2 and 6.2.1 |
| Median thermal flux | e ⁻ /s/pixel | < 200 | See Sections 5.3.2 and 6.2.1 |
| Median read noise | e ⁻ /pixel | 89.95±0.15 | See Section 5.4 |
| Usable well capacity | e ⁻ /pixel | ~56470 | See Section 5.5 |
| Inoperable pixels | % | 0.325 | See Section 5.7 |
| Non-linearity | % | Undefined | See Section 5.5 |
| Persistence | % | < 0.1 | See Section 5.2 |

Indicating perhaps there to be an exposure time or signal dependence to the read noise. There was at least an exposure time dependence on the well depth found in Section 5.4, where 1 s exposures were found to have a ~500 ADU higher saturation value than 30 s exposures.

The uniformity was assessed with the setup shown in Figure 5.8, with a variation of ~ ±1% found across the frame, with images taken with camera orientated at two different angles along its axis. However, with the enclosure not ideally made, there was doubt in the accuracy of the measurements.

In conclusion, despite the higher read noise, the camera specification derived in this chapter still indicates a good photometric precision for the targets within the SPECULOOS target list, as discussed in the next chapter.

FIRST LIGHT

In late September 2021, the Chilean government announced it was re-opening its borders to non-resident foreign nationals who had met vaccination, testing, and isolation requirements. This news subsequently prompted the planning of a long overdue service mission to SPECULOOS' Southern Observatory (SSO) for early January 2022 (where the previous service mission occurred in December 2019, with an expected service period of 6-12 months). This sparked the opportunity to test the InGaAs camera and developed filters on sky, with the combined instrumentation named SPIRIT (SPeculoos' Infra-Red photometric Imager for Transits, with the instrument's banner shown in Figure 6.1). However, for this to occur, the camera mounting mechanism, cooling system, software, and export licences had to be finalised prior to shipping to Chile.

The integration of SPIRIT on one of SSO's telescopes (on the telescope named Callisto)



SPECULOOS' Infra-Red photometric Imager for Transits

Figure 6.1: Banner and logo for the new SPECULOOS NIR instrument, with Blinky (also known as Akabei, the leader of the Ghosts and arch-enemy of Pac-Man) chosen as the mascot for the instrument.

occurred over 2 missions. The first mission, from the 5th to 19th of January 2022, was to ensure the instrument and cooling system could be successfully integrated with the telescope and control software. The second mission, from 2nd to 16th of March 2022, after ironing out any software or mechanical issues found during the previous mission, was to formally integrate SPIRIT at SPECULOOS to proceed with longer term on sky observations which are detailed in this chapter.

6.1 Preparation and integration of SPIRIT at SSO

6.1.1 Instrument mounting mechanism

While we were waiting for the 1280SciCam to arrive in Cambridge, we were provided with a 3D model and technical drawings of the exterior of the camera by the manufacturer. Similarly, from the telescope manufacturer, we were provided with a 3D model of the existing instrumentation connected to the back of the telescope's optical tube assembly (OTA). These drawings plus information about the focus position allowed us to design an appropriate mount for the instrument. It was decided to use the existing filter wheel, this therefore permitted us to reuse the majority of the mounting mechanism already designed for the existing instrumentation. The only portion that required designing and manufacturing was a coupling adaptor from the camera to a mounting plate which previously had the existing instrument's shutter and camera mated to it. The resulting designed coupling adaptor is shown in Figure 6.2.

Figure 6.2 and the corresponding 3D file were passed onto the University of Cambridge's Physics engineering workshop to manufacture. The design was submitted in early September 2021 after seeing the camera in person and we were satisfied it was the best solution to mount the camera to the telescope. The backlog in the workshop, due to pandemic delays, delayed the completion of the piece to mid-November 2021.

6.1.2 Cooling system

As confirmed in the lab, the camera demanded liquid cooling to operate the FPA at -60 °C, where the manufacturer recommended running a liquid chiller at 15 °C to achieve this. Since we were not in a position to modify the camera's cooling system to permit air only cooling for -60 °C operation (as a modification was not agreed with the manufacturer and it would have voided the warranty for us to implement), we had to engineer a solution to introduce liquid cooling on a moving telescope. Such introductions of liquid cooling on a German equatorially mounted telescope we had not seen in literature before, especially one suitable for robotic operation.

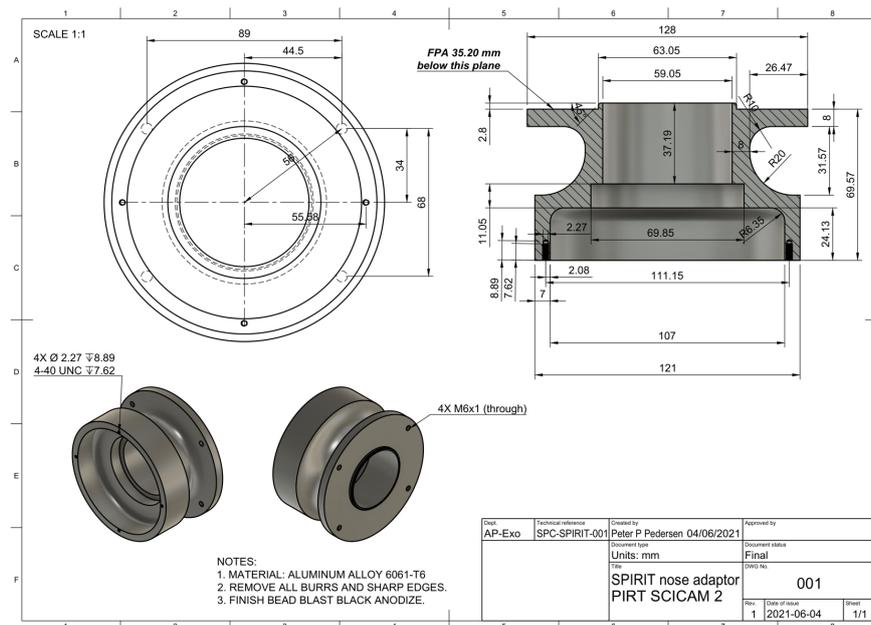


Figure 6.2: Technical drawing of the coupling adaptor for the 1280SciCam to the mounting mechanism at the back of the telescope's OTA. The final manufactured piece was made of a different, but similar in strength, aluminium alloy (6082 HE30 instead of 6061-T6) given the availability of local materials.

After contacting numerous commercial-off-the-shelf liquid chiller providers, none had the functionality to fully operate in a moving environment. This was due to the type of liquid reservoirs used, where it would let air into the system if placed incorrectly and subsequently reduce its cooling capacity or damage the internal pump. If the chiller was to be placed on the OTA, the chiller would be required to function in nearly all orientations. The avenue of placing a chiller on a non-moving portion of the telescope was explored, with the aid of a 3D printed model of the telescope (made with the aid of two 2D profile drawings and onsite photos as the manufacturer was not willing to share an accurate 3D model), but none of the possible positions were deemed safe for robotic operation due to the risk of damaging the liquid tubing during the telescope's movement.

We consulted specialised liquid cooling engineers, where they proposed an untested solution of using a bladder pressure tank to permit placing the chiller on the OTA. But in the end, it was the inspiration from fuel tanks used in remote controlled model planes which presented the solution. These fuel tanks have a normal inlet and an outlet which is connected to an internal flexible tube with a weighted end. This simple reservoir design was ideal for use on a moving equatorially mounted telescope. The weighted outlet would always be submerged in liquid for the orientations the telescope would point under normal operations (above the horizon, with

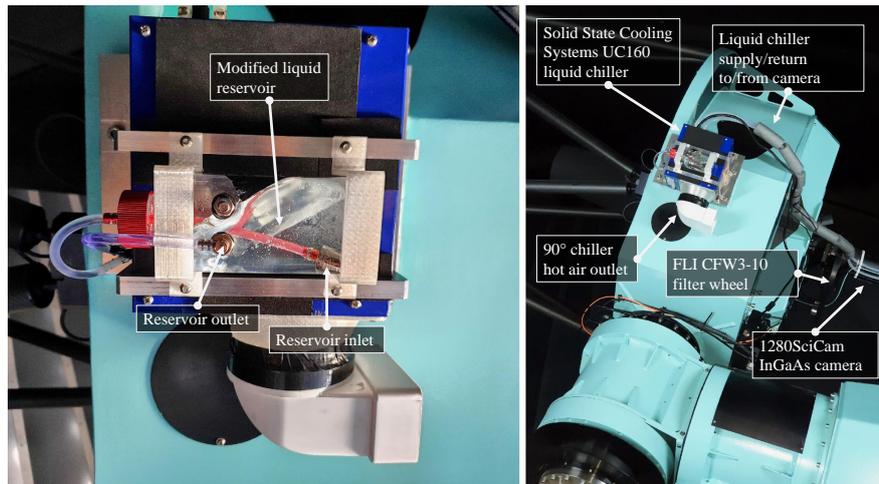


Figure 6.3: Mounted modified liquid chiller on the side of the OTA. Left: Closeup of mounted liquid chiller. Right: Relative positioning of the chiller on the OTA showing the tubing routed to the back of the camera, partially insulated with foam.

the telescope above or below the mount depending on the target's apparent sky position).

Here, after understanding the basic operations of the liquid chiller used in the lab (Solid State Cooling Systems UC160), I replaced its internal 75 ml reservoir with an external 250 ml reservoir made of a polyethylene plastic bottle designed for use in remote controlled model planes. This modification required disabling the chiller's liquid level sensor (by shorting the sensor's terminals), previously used as a safety measure to switch off the chiller for when the old reservoir ran out of liquid, which unfortunately could not be reused with the new reservoir. To fit the new reservoir, the chiller's internal tubing had to be extended and adapted to match the internal diameters of the in/outlets of the new reservoir.

The next step was to engineer a mounting plate to fix the chiller to the telescope. Unfortunately, as we did not have detailed drawings of the telescope, other than information about the available mounting holes, we had to make provisions to manufacture a mount which could allow securing the chiller in two possible locations on the telescope. The first and preferred location was on the side of the OTA, away from the main instrumentation. In this location, if there was to be a leak, it would most likely occur away from the most expensive portions of the instrumentation/telescope. However, to successfully mount in this position, it was dependent on the distance between the side of the OTA and the main mount, which was not confidently known until we were on site. The second and least preferred option was to mount the chiller at the back of the OTA with the instrumentation. Once on site, we were able to confirm the first option was possible.

Figure 6.3 shows the mounted chiller with the modified reservoir on the side of the OTA.

It was mounted with an aluminium plate (with insulating 3 mm foam in between the telescope and the plate to minimise any thermal conduction to the OTA) with holes to secure the plate to either the side or back of the OTA. On the plate were four smaller holes which lined up with the chiller's outer screws, and four threaded holes for brackets to secure the chiller. On these brackets were threaded holes to attach 3D printed brackets for the new reservoir for either back or side OTA mounting options. All the aluminium pieces were designed and manufactured in the Astrophysics workshop with great help from Peter Doherty. The chiller operated with a peltier cooler, with its heat dissipated by a large heatsink and fan. To potentially minimise the effects of the hot air dissipated by the chiller on dome seeing, it was recommended to have the air redirected away from the OTA. This was done by 3D printing an adapter at the chiller's air outlet, which would allow us to attach standard 4" tubing parts to redirect the air. However we are unsure if the implemented 90 degree turn had any effect on the dome seeing. The coolant used was a 30:70 mix of propylene glycol and water by volume respectively.^a At this ratio, the coolant would not freeze until approximately $-16\text{ }^{\circ}\text{C}$, $8\text{ }^{\circ}\text{C}$ lower than the lowest temperature value stated by ESO's historical climate information for Paranal, Chile.^b

Once the chiller, camera + filter wheel were installed, and the liquid tubing connected (with quick release connectors, such that if the tubing were to be disconnected, minimal liquid would be leaked), the telescope was moved to different extremes to assess the performance of the chiller maintaining $15\text{ }^{\circ}\text{C}$ and the FPA maintaining $-60\text{ }^{\circ}\text{C}$. At the extremes assessed, (pointing at zenith, north and south horizon, and one position with the OTA below the mount), the chiller and camera were left stationary for 15 minutes at each position and the temperature of the chiller and camera recorded. In these positions, the chiller was seen to maintain $15\text{ }^{\circ}\text{C}$ and the FPA maintained at $-60\text{ }^{\circ}\text{C}$. Similarly, under tracking conditions, the chiller and FPA maintained at their set temperatures.

6.1.3 Software

For the camera to have compatibility with the telescope control system used by SPECULOOS, described in Chapter 2, the best route forward was to have an ASCOM^c compatible driver developed for the camera. The ASCOM driver platform is used by SPECULOOS, where ASCOM has established standard interface communications between common astronomy applications and hardware.

^aDate accessed 2021-12-01: https://insist2.com/admin/upl/freezing_point_of_propylene_glycol.pdf

^bDate accessed 2021-12-01: <https://www.eso.org/sci/facilities/paranal/astroclimate/site.html>

^c<https://ascom-standards.org/>

6.1.3.1 Camera driver

We received a quote from the manufacturer to develop an ASCOM driver, for \$60,000 USD. Under our contract we could enact the option if I was unable to develop one myself. An ASCOM driver would permit plug-and-play with the image acquisition software used by SPECULOOS, MaxIm DL, and the connected telescope/camera control layer managed by ACP and guiding software DONUTS (McCormac et al. 2013), all of which operated on Windows. By late October 2021, I successfully developed a basic ASCOM driver, where I had interfaced the C# software development kit (SDK) of the BitFlow camera link frame grabber with the methods and properties expected by an ASCOM camera driver.

The driver consisted of three operating threads. The main thread managed all the ASCOM camera methods, a second thread handled image acquisition whenever the camera produced an image, and a final thread handled the sending and receiving of serial commands with the camera (used to set exposure and frame times, and set/read FPA temperature). Figure 6.4 shows the ASCOM camera states and the routes taken to start and stop an exposure from each state. Traditionally, for image acquisition, a request is made to start an exposure, and an image is returned after the exposure period and read time (in total, approximately equivalent to the frame time on 1280SciCam) have been completed. I began developing the driver with this route in mind, however the images produced with this mode of operation led to noisier images than the default operation mode of producing images continuously in a cycle, with an internally timed trigger with a period equivalent to the frame time. Here, after a frame time was finished the next exposure would immediately begin, without an external trigger. There was a level of additional stability provided by this mode which led me to design the driver around this mode of operation.

To set a new exposure period and receive an image with this internally timed trigger mode would require the following sequence: first, sending the exposure time and frame time setting commands via the driver's serial interface, second, waiting for and discarding the current exposure to finish with the old exposure period, then discarding the following incoming image with the new exposure time as this image would contain more noise/artefacts like the single trigger mode previously described. Then finally, the third incoming image would be saved as the image with the newly set exposure time. This consequently meant that setting a new exposure time and then finally receiving an image would take at least two frame times, as opposed to one. The frame time in the driver was set to be a fixed 100 ms longer than the exposure time. As mentioned in the previous chapter, the structure of the bias would vary slightly with the relative length of frame time with exposure time. The exact behaviour of varying the frame time with exposure time was not explored here. We noticed that for one short exposure time

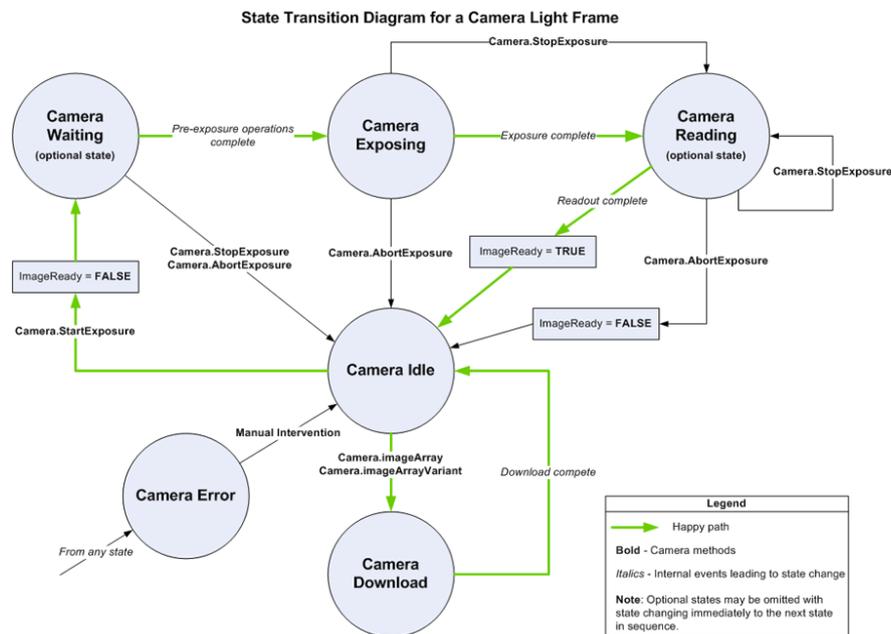


Figure 6.4: The ASCOM camera state diagram for starting/stopping an exposure, and reading the resulting image array if available. “Happy path” indicates the normal path of operation without error. Credit: ASCOM Initiative.

(< 1 s), there was a single banding structure that was more visible on the right of the frame with shorter frame times (< 50 ms longer than the exposure time) and was seen to disappear at frame times 100 ms longer than the exposure time.

Once an image was received by MaxIm DL, about 5 s would be spent by the connected camera control layer on ACP adding metadata to the FITS headers. However, with the internally timed triggered mode the 1280SciCam was operating under, the next frame would already be exposing whilst ACP was adding the previous frame’s metadata. Figure 6.5 shows an example time sequence of acquiring an image with ACP. Thus, despite this inefficiency provided by ACP, there would be minimal deadtime lost with this mode of operation providing the exposure time was longer than approximately 7 s. Consequently however, ACP would be unable to capture all the images the 1280SciCam would produce if we chose to expose below 7 s, which is unfortunate given the highest signal-to-noise would be attained on with shorter exposures.

A second issue derived from ACP and its communication with MaxIm DL. Here, if MaxIm DL had received an image less than 1 s after ACP had begun its request, ACP would fail to acknowledge its receipt and stall the night’s observation run. To resolve this, a minimum 1.5 s would have to pass before the ImageReady variable was assigned true after the Camera.StartExposure command was called (see Figure 6.4).

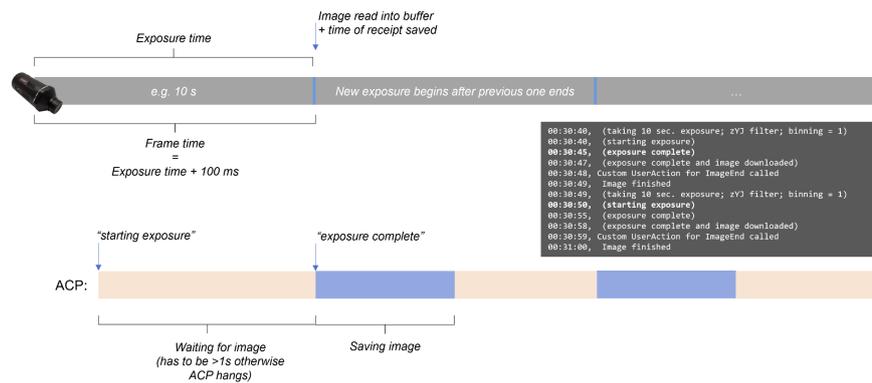


Figure 6.5: Example image acquisition timeline with the 1280SciCam operating with an internally timed trigger (top timeline) and ACP receiving the captured image (bottom timeline). The top timeline shows an example exposure time for 10 s (in grey) and the additional 100 ms time (in blue) from the set frame time. The bottom timeline shows ACP starting a request for an exposure, then waiting for an image to arrive (in pale pink). Once received, it spends 5 s saving the image (in blue). It then gets in sync with the 1280SciCam on receiving images. Inset example ACP log of a series of image acquisition in dark grey box.

6.1.3.2 Guiding software

For precise time-series photometry, guiding is essential for minimising systematic noise derived from pixel-to-pixel variations. These variations, for example, can stem from differences in pixel sensitivity and from errors in flat fielding. Guiding similarly permits one to maintain the target and the same comparison stars in the field of view throughout a series of observations, if the same reference frame is used.

SPECULOOS uses a python based guiding software called DONUTS (McCormac et al. 2013). It has enabled sub-pixel guiding accuracy for our observations when used with the existing CCD based instrumentation. However, with a CMOS based infrared detector, the presence of hot pixels (see Section 5.7) will introduce false identification of stars, which will negatively influence DONUTS’ guiding algorithm. To resolve this, I worked with the main author of DONUTS, James McCormac, over the course of both missions to enable DONUT’s guiding algorithm to function in the presence of hot pixels.^d

We initially intended to use a bad pixel map to mask out the hot pixels. However, this approach was not successful because of the random time variable hot pixels which were not encapsulated in the produced bad pixel map (which was produced from a series of darks). Instead, and likewise to avoid the need to update the bad pixel map, we chose to run a median 4x4 filter across the frame before feeding it into DONUT’s guiding algorithm. This approach

^dThe adapted version of DONUTS for SPIRIT can be found here https://github.com/ppp-one/DONUTS_ACP.

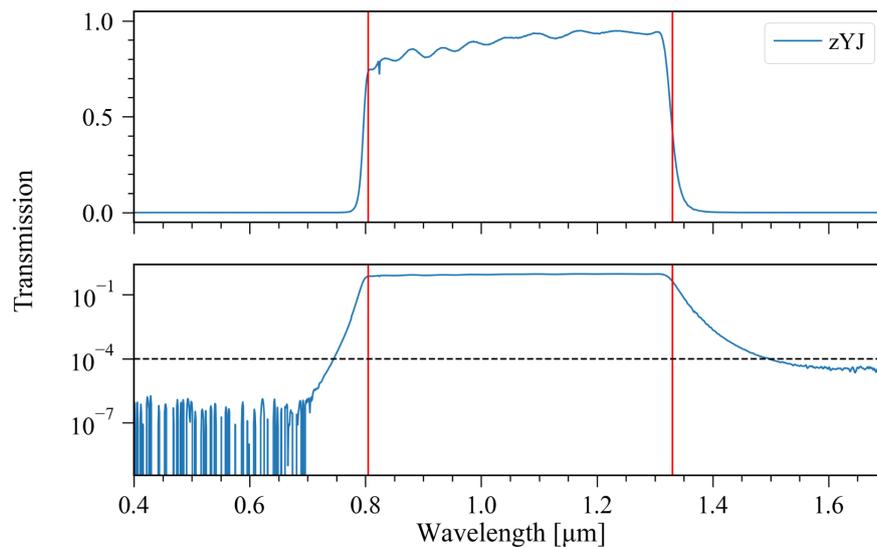


Figure 6.6: Transmission data (in blue) of the accepted zYJ filter set from Brinell Vision, with linear (top) and log-scale (bottom). Red vertical lines show the requested half power points from the specifications. In the log-scale plot, a dashed black line shows the requested OD4 blocking specification. The sharp dip at 0.82 μm was assumed to be due to instrument anomalies as opposed to a real feature.

has demonstrated to maintain the field aligned on a sub-pixel level, as illustrated in Section 6.2.2.

6.1.4 Filters

As per the filter procurement document in Appendix B, we requested 3 filter sets from Brinell Vision, titled zYJ, J, and H-long. In the procurement, Brinell Vision presented the highest theoretical mean transmission curves across the requested ranges over the other companies asked to be part of the tender. We received the first filter sets of zYJ, J, and H-long, in early 2021. Unfortunately, there were quality control issues which led us to reject the batch. In particular, there were visual surface defects in all the filter sets and some of the measured parameters were out of agreed specifications. Consequently, after some deliberation, we requested a re-make of all the filter sets, with a focus tailored towards remaking the zYJ filter set first given it was to be our primary filter.

In June 2021, a revised version of the zYJ filter set was made. In this set, the surface quality had fewer imperfections than the previous set to an extent deemed acceptable. It similarly had better out-of-band blocking in the visible domain by ~ 2 orders of magnitude due to the removal of a small feature at 670 nm, from $\sim \text{OD}3$ to $\sim \text{OD}6$, assessed using the transmission data provided by the manufacturer. From the transmission curve of the revised filter, shown in Figure 6.6, I could assess the extent the new filter set had met the requested specifications,

Table 6.1: Assessed specifications points from the manufacturer’s transmission curve for the accepted zYJ filter set. The roll off slope was assessed with the original definition, and also with a revised definition (under the row with an asterisk) since the transmission curve did not reach the 80% on the blue side. HPP corresponds to the Half-Power Point, the wavelengths where the transmission is 50% of the peak. See Appendix B.

| Parameter | Specification | Value |
|----------------------------|--|-----------------------|
| Average total transmission | >90% transmission, have a ripple of less than $\pm 5\%$ | 90.1% Ripple < 5% |
| Roll off slope | < 2.5% $(\lambda(80\%) - \lambda(5\%) /\lambda(5\%)) \times 100$ | 1: 7.12% 2: 2.36% |
| Roll off slope* | < 2.5% $(\lambda(70\%) - \lambda(5\%) /\lambda(5\%)) \times 100$ | 1: 1.95% 2: 2.17% |
| Out-of-band transmission | > OD4, starting < 70 nm from each HPP till 0.3 μm and 1.85 μm | 1: 52 nm 2: 167 nm |
| σ_{HPP-1} | ± 7 nm (propagated from ± 10 nm FWHM in quote) | -8 nm |
| σ_{HPP-2} | ± 7 nm (propagated from ± 10 nm FWHM in quote) | -1 nm |
| σ_{FWHM} | ± 10 nm | +7 nm |
| σ_{Centre} | ± 5 nm | -5 nm |

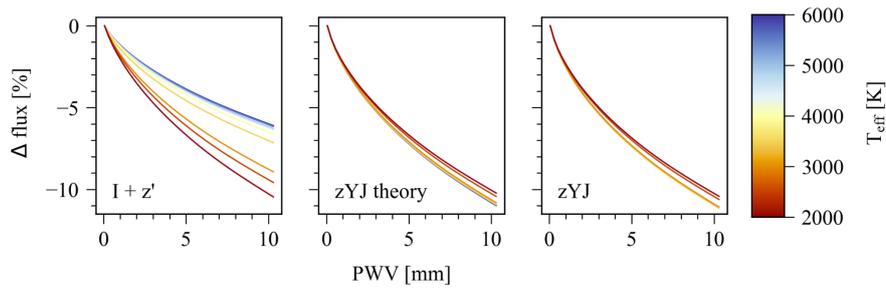


Figure 6.7: The change in flux as a function of PWV for different-temperature stars (from 6000 K to 2000 K in steps of 500 K) as modelled through three observational bandpasses, $I+z'$, zYJ theory, and the real zYJ , with respect to a 0.05 mm PWV atmosphere at airmass 1, using the methods described in Chapter 3.

quantified in Table 6.1. From Table 6.1 and Figure 6.6, we can see that there were some instances where the measured transmission curve did not meet the specifications. Despite these discrepancies, the expected precipitable water vapour modelled performance was similar to the theoretical performance, shown in Figure 6.7.

In Figure 6.7 we can see the PWV-flux behaviour for SPECULOOS’ standard filter, $I+z'$, when coupled with a deeply-depleted CCD. The spread in Δflux for different effective temperatures of stellar models can give rise to variations in observed differential flux when PWV changes during an observation, as described in Chapter 3. Figure 6.7 also shows the PWV-flux behaviour for a theoretical zYJ filter, modelled as part of Figure 4.6 (box filter with cut-on and cut-off defined by the half-power-point in the requested specification), in addition to the

manufactured zYJ filter by Brinell Vision coupled with the 1280SciCam’s quantum efficiency at $-60\text{ }^{\circ}\text{C}$ (shown in Figure 5.2). The spread of Δflux at a PWV value of 10 mm for zYJ bandpass is significantly lower than that presented by the $I+z'$ bandpass, with the modelled zYJ filter showing a similar behaviour to the manufactured filter by Brinell Vision. Nonetheless, there is some deviation for the coolest stars evaluated, but not near to an extent where a false transit feature could be induced. However, should we want to fully evade the effects of PWV, in future we could consider the YJ filter, since it had a smaller spread, or instead J_{MKO} , but at the cost of lost photometric precision. The true PWV suppressing nature of the zYJ filter is evaluated later in this chapter.

At the time of writing, Brinell Vision has failed to provide the revised J and H-long filters due to a failure with their spectrometer to measure the produced filters transmission curves. The final “filter” implemented was a “dark” filter, which was needed since the camera did not have a shutter. It was composed of a 3 mm thick aluminium plate.

6.1.5 Setbacks

The pandemic undoubtedly had a major impact on the progression of delivering first light. The delayed procurement of core equipment from manufacturers and travel restrictions did not permit the amount of work one would have liked within the expected timeline of a PhD. Nonetheless, the results in this chapter should give sufficient indication of the performance of SPIRIT and avenues for future improvements.

Whilst the pandemic caused the largest setback, there were a handful of events which affected the timeline of getting SPIRIT on sky. In summary, technical downtime of the 1280SciCam, electromagnetic interference (EMI) debugging, software timing issues (described in Section 6.1.3), bad weather, and finally long-term telescope downtime hindered the progression of getting results.

The first technical downtime of the 1280SciCam occurred during the first mission, where the camera died during its first mounting on the telescope. After a remote session with the manufacturer the following day, a power cable within the camera was found to be loose. The cable unfortunately detached itself again at the end of the first mission. A replacement cable was thus installed during the second mission.

A longer delay was caused by EMIs, which affected the communication between the camera and the computer. Here, the computer would blue screen or freeze after interference, where restarting the camera and computer then became necessary. Initially, it was thought to be related to some Windows 7 operating system bug, where the system would occasionally blue screen during an observation run. After changing computers and changing to a Windows 10 based system, the computer would instead freeze under similar conditions. It was later discovered,

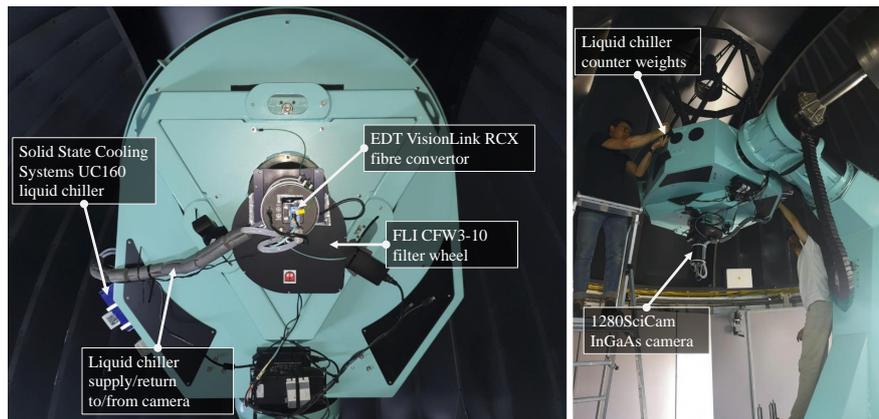


Figure 6.8: Left: A view of the back of the OTA, with the attached equipment which make up SPIRIT, mounted in their final position. Right: Process of balancing the declination axis, to counteract the additional ~ 4 kg mass from the liquid chiller. The right ascension axis was similarly balanced by shifting the counterweight ~ 1 cm away from the axis.

after a few days of debugging during the second mission, that grounding the fibre convertor on the camera end (although it was connected to the camera which was grounded) would resolve the issue, with the EMI deriving from dome rotation.

The longest setback was caused by telescope downtime, where the hour axis faced a torque error, which occurred during an observation run on the last night of the second mission while the observatory was operating in remote mode. This error took 2 months to resolve with the involvement of ESO engineers, who found a cable which unluckily became loose within the telescope's control cabinet.

Lastly, a warning to future infrared instrumentalists, the shipping process of sending equipment to Chile could have potentially induced a far greater setback than the above. Since the 1280SciCam had U.S. origin, it was subjected to export restrictions due to its potential military applications. Fortunately, the UK customs agency helped in providing the relevant documentation, as did the manufacturer with getting the relevant documentation from the U.S. Commerce Department, which permitted the equipment to arrive in Chile in time for the first mission. In future, it's recommended to have this documentation ready months in advance.

6.1.6 Final installed system

Figure 6.8 shows SPIRIT mounted on the back of the telescope's OTA. Here, we can see the cooling system described in Section 6.1.2, with the liquid tubing connected to the back of the 1280SciCam. The tubing was wrapped with segmented insulating foam, which was more to permit securing the tubing to the OTA with cable ties without imposing damage to the plastic

tubing than to insulate heat. It was also to minimise the travel distance of any leaks from the chiller along the tubing towards the camera. The supply and return lines shared the same insulating foam, thus insulated more from ambient heat. Nonetheless, the implemented setup was sufficient to permit the FPA to reach $-60\text{ }^{\circ}\text{C}$ under the ambient conditions experienced in Paranal. The system could be made more efficient by separating the supply and return lines, and insulating them independently. In the future we would like to remove SPIRIT's liquid dependence by introducing a larger heatsink and fan combination to the cold finger within the camera build.

On the camera body were 3D printed guides to guide the power, communication, and liquid tubing back towards the OTA. Since the camera's ports were at the back of its body, it meant that the length of the 1280SciCam was effectively extended by the minimum turning radius of the liquid tubing and cables connected. We were lucky that this additional distance did not intercept with the mount when the OTA was pointing from beneath the hour axis, where we were a few centimetres from intercepting. Had interception occurred, we would have had to introduce a meridian flip to the telescope control software, or more extremely, modify the camera build to reposition the camera interface board. At each mating face that formed the optical train at the back of the OTA was taped to minimise the extent of dust migrating into the optical path.

In the event of instrument failure, the old power and communication hardware for the CCD based camera was kept on the OTA. This was decided to permit an easier exchange of instrumentation should a long term failure occur with the new equipment. SPIRIT utilised the existing optical fibre routed through the mount to the OTA, which was extended to the fibre converter located on the back of the camera. The fibre converter was grounded to the OTA, as shown in Figure 6.8 and described in Section 6.1.5. The filter wheel communication was done via USB, which was connected to a powered USB hub. Cable management was less elegant than desired, but the use of 3D printed guides and cable ties enabled secure mounting. Since the other SPECULOOS telescopes are identical in construction and share the same optical setup, it would be possible to replicate the system, both in hardware and software without much difficulty other than the potential lead times on some of the core components.

6.2 Early results and observations

In this section I detail early observational results and look into any discrepancies with the models developed in Chapter 4. I begin with a quantification of the thermal background flux experienced by the detector, and then progress into on sky observations accompanied by simultaneous observations with a second telescope at SPECULOOS' Southern Observatory. I

then conclude with a discussion on the limitations and future improvements for SPIRIT.

6.2.1 Quantifying the true dark current

In the lab, I had measured the detector's background (dark current + thermal) flux to be $\sim 260 \text{ e}^-/\text{s}/\text{pixel}$ at room temperature, when the detector was shielded by a plastic and aluminium cover. The exact proportion contributed by the dark current or thermal background was unknown at the time. The thermal flux was likely to be different with SPIRIT mounted, given the different geometry, materials, and temperatures which were to be in front of the detector. To quantify this, while the telescope had technical downtime, a long series of 60 s darks were taken over a period of 3 days. Over this time, the dome temperature varied from 13 to 21 °C, a range which was not possible to simulate in the lab. This large temperature range permitted better calibration of determining the dark current and thermal contributions to the detector's measured background flux.

Modelling the thermal emissions from the exact geometry and respective emissivities in front of the detector was deemed unnecessarily complex. I instead assumed the overall thermal flux could be modelled as if emitted from a single blackbody. One can then integrate Planck's law over SPIRIT's wavelength range to calculate the change in flux as a function of temperature. Correlating the modelled value to the recorded detector's background can permit us to estimate the true dark current produced by the detector.

A temperature sensor located at the edge of the dome was used to measure the temperature variation over the course of the experiment. There was a clear relationship between the background flux recorded and the temperature recorded. The nature of this relationship improved when a 45 minute delay was assigned to the temperature sensor values, where it was assumed that SPIRIT and the telescope took approximately 45 minutes to reach thermal equilibrium with a change in dome temperature.

Thus, continuing with the previous assumptions, we can write the observed background flux to be a function of environment temperature, T_{env} , which is composed of thermal flux and dark current, DC :

$$Background(T_{env}) = Thermal(T_{env}) + DC. \quad (6.1)$$

Taking the derivative with respect to environment temperature leaves us with:

$$\frac{dBackground}{dT_{env}} = \frac{dThermal}{dT_{env}} = A \frac{dThermal_{est}}{dT_{env}}, \quad (6.2)$$

Where the dark current derivative is 0 given its independence to environment temperature since the FPA remains at a fixed temperature of $-60 \text{ }^\circ\text{C}$. The true thermal flux is assumed to be

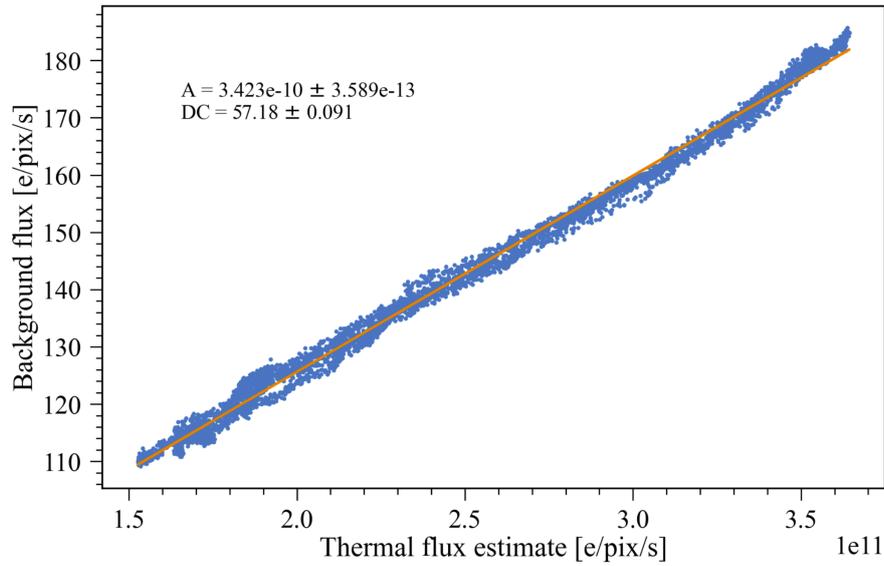


Figure 6.9: The change in background flux from 60 s darks over 3 days (taking the median of the frame) compared to the estimated thermal flux (from Equation 6.3) experienced by SPIRIT when mounted on the telescope during telescope downtime over a 13 to 21 °C temperature range. Raw data points in blue, and line of best fit in orange with the gradient value, A , and intercept, DC , shown with errors from the fit.

proportional, by a constant A , to an estimate, $Thermal_{est}$, given by an integrated Planck's law over SPIRIT's wavelength range:

$$Thermal_{est}(T_{env}) = \int_{\lambda_0}^{\lambda_1} \frac{2c/\lambda^4}{e^{hc/\lambda kT_{env}} - 1} d\lambda, \quad (6.3)$$

Where wavelength, λ , Planck's constant, h , speed of light in a vacuum, c , and Boltzmann constant, k , are used. Equation 6.2 can be rearranged for the constant of proportionality

$$A = \frac{dBackground}{dThermal_{est}}. \quad (6.4)$$

This constant of proportionality can be found as the gradient of the observed background flux against estimated thermal flux, shown in Figure 6.9. Likewise, the intercept will give us the dark current of the detector as suggested by Equation 6.1.

With the constant of proportionality, A , calculated from Figure 6.9, we can show how the thermal flux varied as a function of time in Figure 6.10. Here, the dark current was calculated as the residual of the background flux minus the calculated thermal flux using the constant of proportionality, A . The dark current was seen to be relatively constant, ranging between 50 and 60 $e^-/s/pixel$. This range is likely to be associated with the inaccuracies in the SPIRIT/telescope temperature estimate from the dome temperature sensor, rather than the

actual dark current varying by that amount. Adding a systematic ± 5 °C, for example, to the temperature estimate, and recalculating A , gives an intercept value range for the dark current between $54 - 60$ $e^-/s/pixel$. Nonetheless, it suggests that the true dark current is ~ 60 $e^-/s/pixel$ at a FPA temperature of -60 °C, which is similar to the value of 68 $e^-/s/pixel$ attained by [Birch et al. \(2022\)](#) who had a -10.6 °C lens cap in front of the detector.

The FPA temperature was seen to remain at the set temperature of -60 °C with some instances of deviation of a few 0.1 °C. These deviations are unexplained, but perhaps caused by air getting into the circulated liquid by the mixing caused by the return line which is sometimes above the liquid level. At the position the telescope was locked at during its technical down time, it's difficult to confirm whether or not the return line was above the liquid level. The first event is visible in the background flux, as is the second event, but to a lesser extent given the lower temperature fluctuation range. The effect on the background flux and connected quantum efficiency is minor, and so it's not a cause for concern at the moment. A side note, the bias was similarly recorded over a 24 hour period, and was not seen to change as a function of time/temperature.

The varying background flux during the night will have an effect on the data reduction. Presently, darks for calibration are taken at the end of a night when the temperature of the equipment is likely to be at its coolest, before being rapidly heated up by the rising sun. As a consequence, the calculated sky background from each observation will be different, likely higher, from reality. Normally, the dark exposures will have finished exposing before the dome has been rapidly heated. The data shown in Figure 6.10 was taken when the dome was shut. However, when the dome is open, the temperature will be likely to fluctuate depending on the weather conditions, as well as the relative pointing direction with respect to the wind direction. This result suggests we would be better to take dark calibrations during an observation run as opposed to at the end in order to have more accurate photometry. It similarly presents the argument that we should regulate the temperature of the optics in front of the detector, perhaps to the extent of near-nullifying the thermal flux.

Since we now have an understanding of the thermal flux as a function of temperature, and similarly wavelength range, we can estimate the temperature needed to significantly reduce the thermal background, but similarly understand what effect the cut off wavelength of a detector will have on the recorded thermal flux. Figure 6.11 shows this calculation for a range of cut offs which cover a standard Si based CCD to a HgCdTe 1.75 μm cut off detector. Thus, to reduce the thermal flux of SPIRIT, which has a cut off of 1.62 μm , to for example < 1 $e^-/s/pixel$, it would require an environment temperature of -18 °C. An HgCdTe based detector in comparison, would require -36 °C. However, since HgCdTe detectors typically come in pitches 10 , 15 , and 18 μm , the temperature to reach < 1 $e^-/s/pixel$ thermal flux will be in the range of -41 to -33 °C.

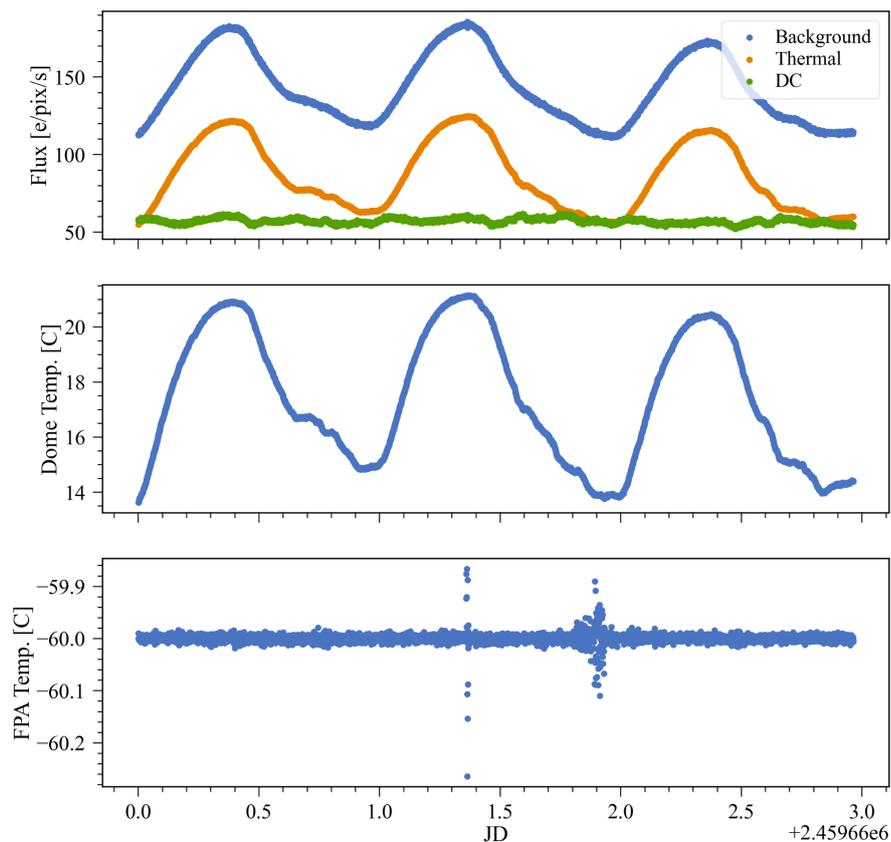


Figure 6.10: Top: Background flux measured from 60 s darks over 3 days (taking the median of the frame), in blue. The constituents, from Equation 6.3, which make up the background flux, thermal flux in orange, and dark current (DC) in green. Middle: Dome temperature from a sensor near the edge of the dome, offset by 45 minutes to account for the time for SPIRIT/telescope needed to thermally equilibrate. Bottom: FPA temperature, which was set to run at -60°C .

Without cooling the instrument’s environment, during a nightly temperature of 14°C (median temperature at night over the course of a year), the thermal flux for SPIRIT is approximately equivalent to its dark current at $57\text{ e}^{-}/\text{s}/\text{pixel}$. In comparison, a HgCdTe based detector of the same pitch would see $493\text{ e}^{-}/\text{s}/\text{pixel}$ at a 14°C environment temperature.

More significantly, the dark current on a theoretical $12\text{ }\mu\text{m}$ pitch HgCdTe based detector at -60°C FPA temperatures would be $\sim 4000\text{ e}^{-}/\text{s}/\text{pixel}$, 2 orders of magnitude higher than the InGaAs based detector. It would take approximately a -90°C operating temperature to match the dark currents we are currently experiencing with the 1280SciCam (Beletic et al. 2008). These temperatures are feasible in theory to set up without requiring regular maintenance, but not presently commercially available off-the-shelf with HgCdTe. Should the 1280SciCam be able to cool to -90°C , the dark current would be expected to reach $\sim 3\text{ e}^{-}/\text{s}/\text{pixel}$.

It is currently unclear how the thermal background will change once the camera is exposing

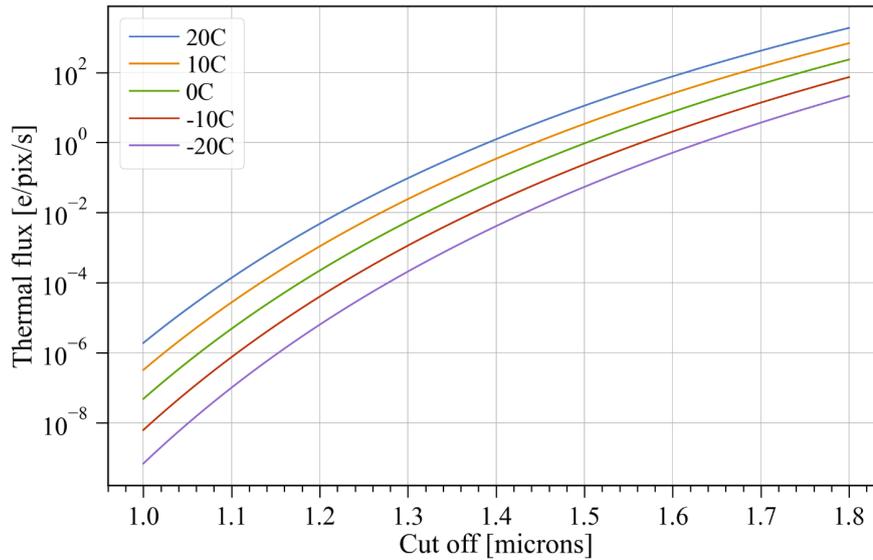


Figure 6.11: Expected thermal flux at different quantum efficiency cut offs and temperatures received by an instrument mounted in the same position as SPIRIT and same 12 μm pixel pitch.

through a filter. The results in the following sections suggest the additional thermal background through the zYJ filter are negligible. With a cut off of 1.33 μm , the zYJ filter under the same calculations as above for 14 $^{\circ}\text{C}$ environmental temperature, the thermal flux is expected to be near $\sim 0.1 \text{ e}^-/\text{s}/\text{pixel}$ through the filter plus the previously calculated thermal flux under the same geometry of the calculation. This is assuming the different emissivity of the filter, in comparison to the “dark” filter used for the above calculations, does not significantly change the thermal calculation. However, given the filter is exposed to the remaining optics, the thermal contribution may be larger. According to [Holzlöhner et al. \(2018\)](#), thermal radiation in a telescope beyond the instrument that arrives at the detector primarily stems from mirror imperfections and dust on the mirror surface which both emits and scatters thermal radiation (where dust is likely to have a higher emissivity coefficient than the mirror coating). Similarly, if the telescope has not been appropriately baffled, then additional stray thermal flux will arrive at the detector. Fortunately, all of SPECULOOS telescopes have included baffles, but we have not actively engineered a cold stop as part of the optical path. [Simcoe et al. \(2019\)](#) similarly did not add a cold stop as part of their InGaAs camera build. This is to be a much greater issue when observing beyond the zYJ filter, which was not possible to assess during this work.

6.2.2 Cool star simultaneous observation

Near the end of the second mission, once all the hardware and software bugs related to SPIRIT had been resolved, a couple of simultaneous observations were made with other SPECULOOS telescopes on site. We observed one of the cooler targets from the SPECULOOS target list ([Sebastian et al. 2020a](#)), Sp1507-1627, a 1694 ± 104 K target at 7.39 ± 0.02 pc, 12.830 ± 0.027 J magnitude, to assess whether SPIRIT would indeed provide better photometric precision than the existing instrumentation for cooler targets.

We first observed the target simultaneously with SPIRIT's neighbouring telescope, called Ganymede, but its guiding and focus were not optimal, and thus deemed an unfair comparison. The following night we observed the same target with a different neighbouring telescope, Europa, which did not have guiding or focus issues. SPIRIT observed with 30 s exposures and Europa with 36 s. SPIRIT's exposure time was calculated to fill 70% of the well, assuming a seeing of 1", since it was unknown to what extent the seeing would change with the new instrument, discussed later in this section. The exposure time for Europa to fill 70% of the central pixel's well would have exceeded SPECULOOS's set cadence limit of 120 s. It was decided to observe at a lower target well fill to minimise the effect of red noise ([Pont et al. 2006](#)), without sacrificing significantly on binned precision (difference in precision estimated in the Section 6.2.3).

The aligned stacks are shown in Figure 6.12, which were processed with `Prose` ([Garcia et al. 2022](#)), which also processed the resulting light curves. In `Prose`, a bad pixel map was computed for SPIRIT from the darks and flats in the same manner as Section 5.7. An average of the surrounding pixels (excluding neighbouring bad pixels) around a bad pixel was taken to mask the region. This approach is preliminary, and there is scope for potential improvements. A dedicated pipeline for SPIRIT is intended to be developed, but for this work we are using the standard approach for reducing and differential photometry, with one difference being this method to deal with the bad pixels. One further exception to the reduction process from that utilised by SPECULOOS' dedicated pipeline ([Murray et al. 2020](#)) is with the creation of the master dark. Instead of estimating the master dark from a series of darks of increasing exposure time, a master dark of matching exposure time to the science frames was instead used. The reason for this difference stems from artefacts that were seen to appear with the former approach. In particular, some vertical banding and ROIC glow were more visible if the dark exposure times were not matched. For the master flat, it's presently impossible to produce a flat which matches the science exposure since the observatory can only perform sky flats, whose exposures range from 0.1 s to 15 s. As a consequence, these artefacts do appear from time to time, but less significant than before the dark exposure matching. A deeper investigation into

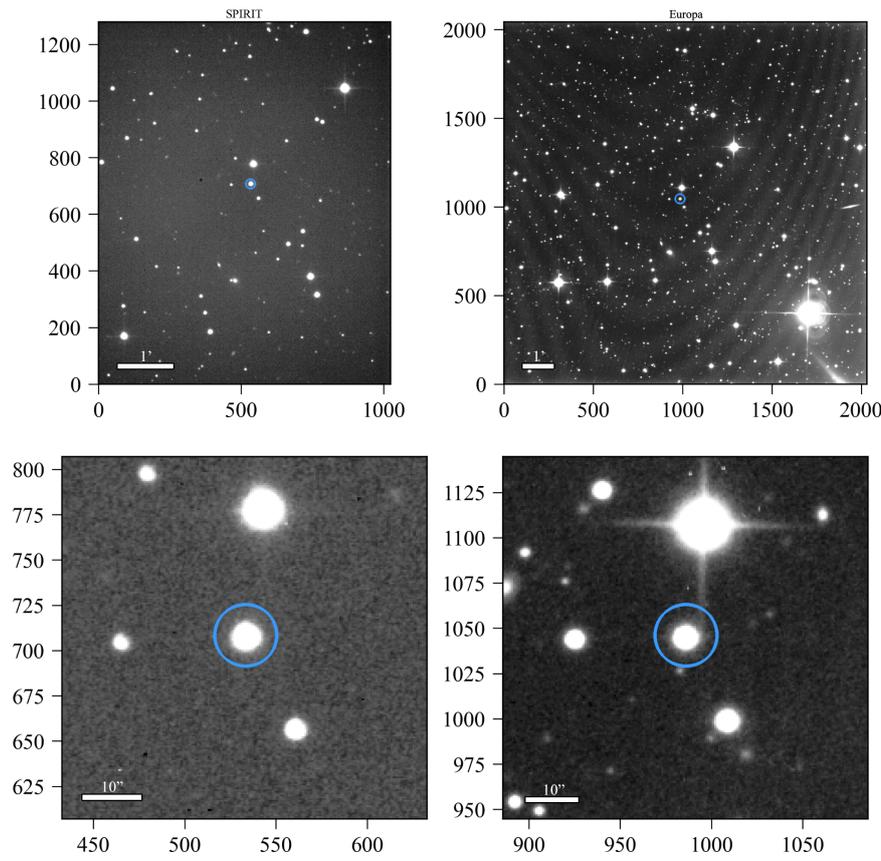


Figure 6.12: Reduced stacked images of the Sp1507-1627 field of view with target circled in blue. Top-left: SPIRIT’s full field of view. Bottom-left: SPIRIT’s cropped in field of view, with the target star centred. Top-right: Europa’s full field of view with the existing instrumentation. Bottom-right: Europa’s cropped in field of view, with the target star centred.

the artefact’s effects on photometry is needed. In future, the introduction of dome flats which can permit exposure matching is desired. The last difference to the reduction process is related to the background flux estimate. `Prose` estimates via an annulus as opposed to the sliding median method utilised in `SPECULOOS`’ dedicated pipeline via `CASUtools`^e. The different background estimation methods do not seem to impact the final photometric product.

From the aligned stack, shown in Figure 6.12, there were some notable visual differences. The first noticeable difference was the lack of fringing in SPIRIT’s reduced stack, which has likewise not been seen in other zYJ imaged fields by SPIRIT from other nights. [Birch et al. \(2022\)](#) similarly did not see fringing with their unit during lab tests when exposing with a narrow pass filter covering the edge of 1280SciCam’s quantum efficiency range. Fringing can cause issues with the differential photometry process, particularly when the guiding is poor.

^e<http://casu.ast.cam.ac.uk/surveys-projects/software-release>

The second difference in the stacks was a dark ring like feature, a bright feature in flats, $\sim 1'$ left from the centre of the frame. In the stack, the region was $\sim 5\%$ dimmer than the surrounding area. The cause of this feature is unknown but potentially due to a deposition of some additional photosensitive material which has led to different behavioural characteristics to the rest of the detector. It similarly may be due to imperfections on the detector's surface or fused-silica window causing light to image on the detector plane in this manner. This region should thus be currently ignored for photometry until further investigation is performed. The third difference was the number of fewer visible stars (and thus fewer comparison stars) in the same angular area. This was expected given the higher noise floor of SPIRIT, and therefore fewer dimmer stars are to be visible above this noise floor. Similarly, for the brightest stars in the field, they are often of much higher effective temperatures than the target star, and will therefore have more flux received in visible wavelengths than in the infrared.

The resulting light curves from the night's observation are shown in Figure 6.13, processed by `Prose`. This short comparison was enough to give an indication of SPIRIT's better photometric capabilities for cooler targets. Similarly, it was sufficient to verify the noise models developed in Chapter 4, which are discussed in detail in the next subsection. The conditions for the night were clear, humid ($\sim 60\%$), with a $\sim 75\%$ illuminated moon setting in the middle of the observation. A better median FWHM was achieved by Europa, $1.17''$ in comparison to $1.35''$ achieved with SPIRIT. Europa was known to have a better historical FWHM than the telescope SPIRIT is mounted on (Callisto), with a median observed FWHM of $1.18''$ vs $1.34''$ on Callisto over the course of its operating history since 2018, as computed by SPECULOOS' dedicated pipeline with the existing instrumentation and $I+z'$ filter.

The aperture which normally presents the best photometric precision is one which optimises the total number of pixels to encapsulate enough flux whilst minimising the summed noise which each pixel brings (described in Equation 4.1). However, this optimisation is time variable due to changing atmospheric conditions. Presently, it's standard practice to choose one aperture for the course of the observation rather than to vary the aperture and detrend, which would possibly lead to correlated noise. In the observations made, an aperture radius of 8.9 pixels for SPIRIT and 6.5 pixels for Europa were chosen. The binned signal-to-noise ratio naturally improves with smaller apertures, up to a point (Howell 2006), given the smaller number of pixels used, but with an increased risk of correlated noise. The structure remained similar for larger apertures. At the chosen apertures, `Prose` chose 2 comparison stars for SPIRIT, and 7 for Europa (majority of which were outside SPIRIT's field of view) using the algorithm described in Broeg et al. (2005). Choosing the same comparisons as SPIRIT for Europa's photometry (both of which were above 5000 K) yielded the same structure but with marginally larger noise than shown in Figure 6.13.

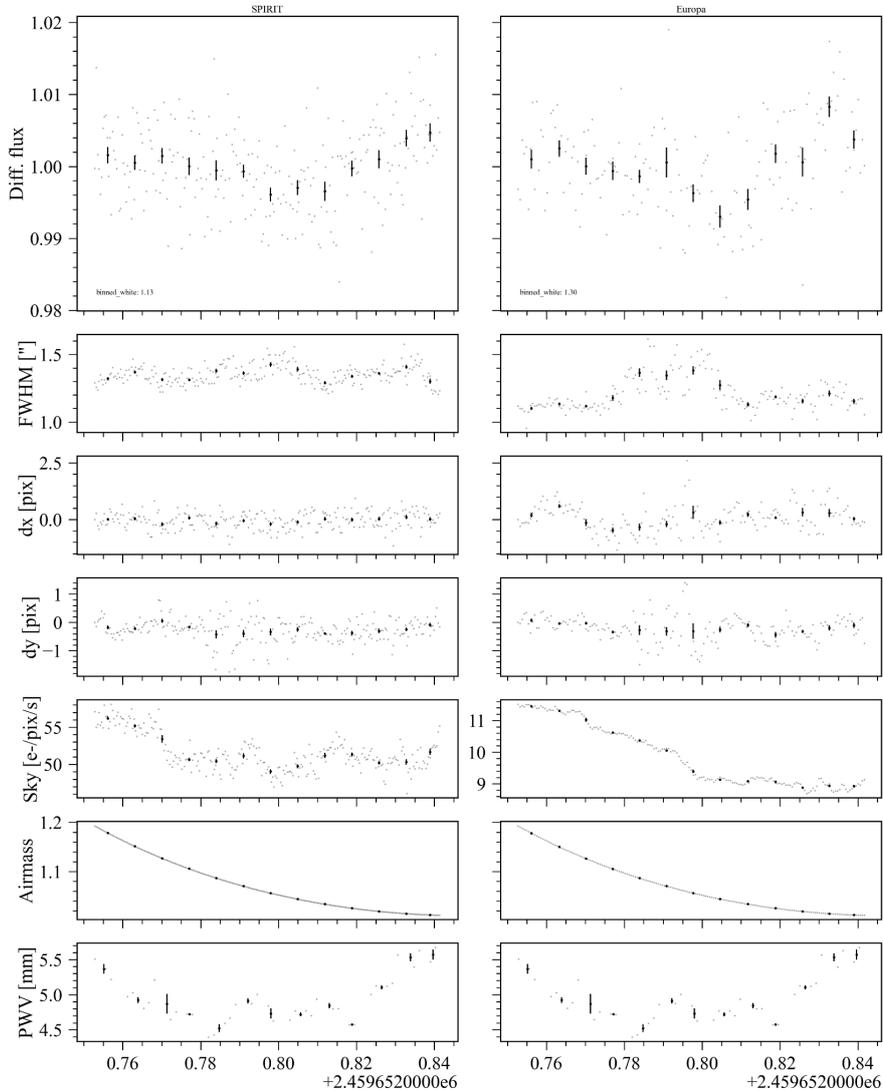


Figure 6.13: Top: 10 minute binned light curves (in black, raw data in grey) of Sp1507-1627 taken on the night of 2022-03-13 under an aperture radius of 8.9 pixels for SPIRIT and 6.5 pixels for Europa. SPIRIT observed with the zYJ filter (left) and the existing instrumentation on Europa with I+z (right). Observing related parameters, such as full-width-half-maximum (FWHM), image shift dx, dy, sky background, airmass, and PWV are similarly plotted below the light curves for reference across the same time period. The x-axis is Julian Date (JD). The PWV data for Paranal was retrieved from the LHATPRO query form for LHATPRO ID 2, with periodic spikes removed, assumed to be conescan measurements (described in Chapter 3).

There was a dip seen in both light curves, which is not believed to be a transit. If it were to be, Sp1507-1627 would be the coolest exoplanet host star after Trappist-1 (Gillon et al. 2017) found via the transit method. The moon was setting around the time the beginning of the dip was seen in both light curves, with the telescope pointing near zenith. The dip was less pronounced in SPIRIT's light curve suggesting a lower sensitivity to the effect. However, the return to normality could not be correlated with any background flux changes. There was a change in FWHM during the dip duration, but the change in FWHM structure did not match the inverse of the dip, neither was it seen in the comparison stars. The data was not PWV corrected as the models had not been extended to below 2000 K. If it was PWV induced, we would have expected an opposing shape to have been observed in SPIRIT's light curve if the flux-PWV relationships in Figure 6.7 are true for the zYJ filter. An alternative hypothesis is that it could be a real feature, but rather sourced from a stellar spot, whose contrast is lower in the infrared (discussed in detail with another example in Section 6.3.1). Further observations of this target are required to probe its origin.

We were expecting SPIRIT to have marginally better seeing than the existing instrumentation given its bandpass operating further in the infrared (Boyd 1978). However, it has continued to have higher seeing than the existing instrumentation on other nights. The immediate cause for this is unknown, and it is a cause for concern as the quality of the photometry is a function of seeing (c.f. Figure 4.13). If it is not optics/dome seeing (from the camera/liquid chiller expelling more hot air than the previous set up) induced, perhaps it is linked with PWV/relative humidity causing a blurring effect (Absil et al. 2022), where the zYJ band has larger water absorption features in its bandpass than the $I+z'$ filter. If this is the case, then we would possibly see the FWHM vary with PWV, although this relationship has not been seen with the existing instrumentation. Nonetheless, the preliminary conclusions from these observations are that SPIRIT has the potential to deliver better photometric quality than the existing instrumentation, but there remains the unknown extent of the effect of red noise and hot pixels affecting the photometry. This requires a more in depth global analysis before deriving any conclusions, which was not possible to perform in this work.

6.2.3 Noise models

To assess the models made in Chapter 4 with the camera parameters determined here and in Chapter 5, we can plot the median binned noise of the detected stars in the field with respect to their magnitude. Ideally, the binned noise should match the total noise contributions from the detector, target, and atmosphere, described in Equation 4.1.

In Figure 6.14, we see the 10 minutes binned precision against filter magnitude for the observations described in the previous sub-section. The binned precision was calculated from

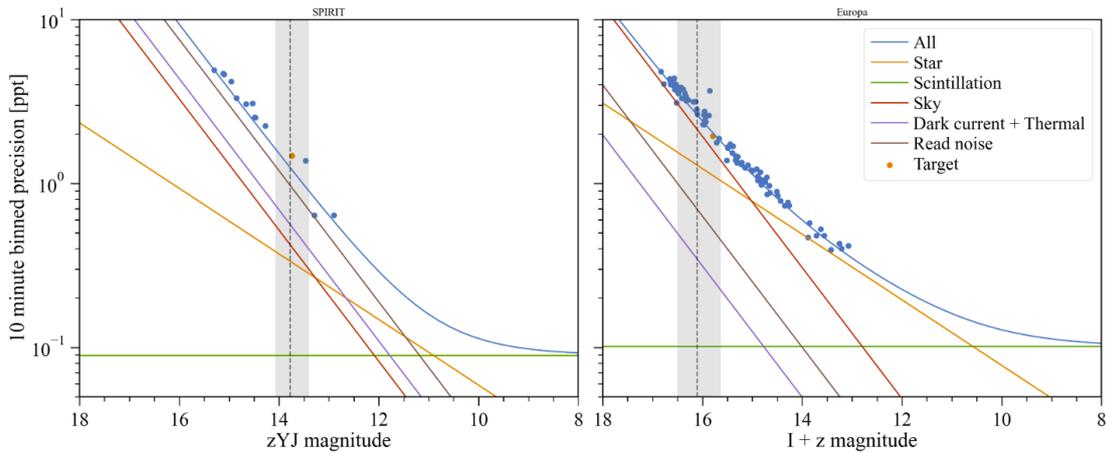


Figure 6.14: Comparing the photometric 10 minute binned precision of SPIRIT and Europa with the observations made in Section 6.2.2, with all the identified stars as blue points, and target star in orange, plotted against their respective magnitudes. Here, a larger set of aperture radii were used to encapsulate the majority of the target’s flux, 12.16 pixels for SPIRIT, 11.31 for Europa. Vertical dashed lines show the expected magnitudes for a 1694 K star at 7.4 pc using the models utilised as part of the modelling in Chapter 4, the shaded grey region shows the expected magnitude range for ± 100 K from the target temperature. The coloured lines show the respective contributions of the noise model, with the blue line representing the total model.

the differential light curves for all the identified stars in the field whose peaks were at least three times higher than the background flux and similarly did not saturate over the course of the observation. The median binned error over the course of the observation was used to quantify the star’s respective precision. The magnitude of the stars were calculated using an estimated Vega flux for the respective bandpasses, using a spectrum sourced from the Astrolib PySynphot’s spectrum library (Lim et al. 2015), integrated over the respective bandpasses and atmospheric conditions.

SPIRIT’s camera parameters were taken from Table 5.4 and from the darks taken the morning after the observations (dark current + thermal: $90 \text{ e}^-/\text{s}/\text{pixel}$). Europa’s camera parameters were taken from its manufacturer and from measurements on site (gain: $1.072 \text{ e}^-/\text{ADU}$, read noise: 5.8 e^-_{rms} , dark current: $0.2 \text{ e}^-/\text{s}/\text{pixel}$). The median sky background for SPIRIT was measured to be $51 \text{ e}^-/\text{s}/\text{pixel}$, $\sim 20\%$ lower than my models predicted. For Europa, the median sky background was $\sim 33\%$ lower than the models, suggesting that it was a better than average night for sky background. Earlier in the night, when observing a different target, the sky background was 3-4x higher for both telescopes due to a closer angular distance to the moon. The scintillation appeared to match with Equation 4.2, however, there were not enough bright objects that did not saturate to line up with the tail end of the model.

Notably fewer stars met the noise and saturation cut for SPIRIT’s field of view, with the

target star being the 4th brightest in the field, and the two brightest stars forming the artificial star as part of the differential photometry process. Europa in comparison had approximately 10x more stars that met the cutting criteria, with the target star being the 47th brightest star in the field, and the comparison stars being similarly the brightest in the field. The target star's flux matched closely to the expected model flux in SPIRIT's field, after the models were calibrated with the 2MASS catalogue (described in Section 4.3.2). The flux in the $I+z'$ band was greater than expected, suggesting that either the model spectra used were not as accurate, or that the temperature of the star was higher than estimated. Nonetheless, there appears to be a good agreement between the precision models and the observed data, with the target star in SPIRIT's data appearing above the model line possibly due to the use of only two comparison stars. However, there have been instances with other targets, where SPIRIT's precision has had significantly more noise, to the extent that where the read noise has effectively increased by $\sim 40\%$. The cause for this is yet unclear, whether it's to do with hot pixels landing on the target and comparison stars, or to do with CMOS architecture and particular exposure times being linked to higher read noises. In the case evaluated here, neither the target or comparison stars had bad pixels within them. The other stars in the field were not checked.

The observations made could have had greater precision if the exposure times were optimised for the night's observations. The peaks of the target's well did not reach the 70% well depth target for either of the telescopes. In Europa's case, even approaching the 120 s cadence limit would still not reach the 70% well depth at the seeing conditions usually observed by the telescope. Nonetheless, had its exposure time matched 120 s, the 10 minute binned precision would have likely improved by an estimated $\sim 10\%$. Similarly, if SPIRIT had reached 70% of its available well depth (by doubling its exposure), we would have seen a $\sim 15\%$ improvement in its 10 minute binned precision.

Inspecting Figure 6.14, we can see that Europa was limited by its sky background, whereas SPIRIT was limited by its read noise. This suggests that there is scope for future improvements with the next generations of this instrument. Prior to receiving the 1280SciCam we were expecting the read noise to be half of that received, in which case we would have been dark+thermal limited on moonless nights. On moon-full nights/observing close to the moon, it may be that we will be sky background limited. Taking what has been learnt from the observations above, and taking into account the potential for $\sim 40\%$ higher read noise in some instances, we can recreate Figure 4.10, in Figure 6.15, to gauge an understanding of where SPIRIT will present better photometric precision than the existing instrumentation.

In Figure 6.15, we have used the camera parameters from Table 5.4, with a fixed the dark current + thermal flux value of $200 \text{ e}^-/\text{s}/\text{pixel}$, which would be the highest one would expect to see it reach given the upper temperatures we have observed Paranal to reach at night (22°C). The

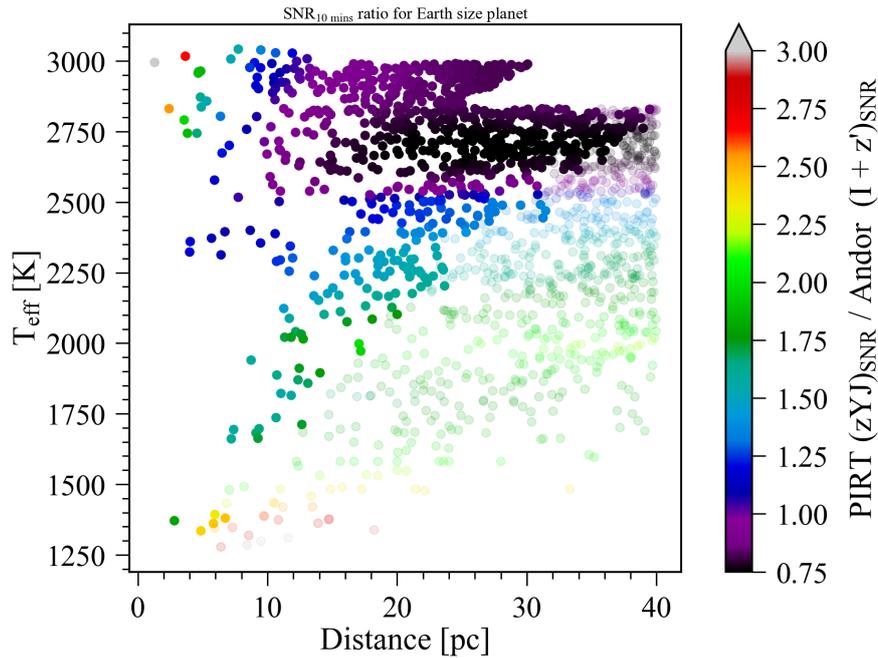


Figure 6.15: Recreation of Figure 4.10. SNR ratio between SPIRIT and existing instrumentation with the $I+z'$ filter for the entire SPECULOOS target list, plotted as a function of effective temperature against distance. Solid coloured points represent the cases where the $\text{SNR} > 3$ (for the metric defined as part of Figure 4.10) for the proposed NIR instrumentation.

camera's read noise was set at $130 e^-_{rms}$ to account for the instances where the field has higher read noise, for reasons presently unknown. The seeing was set to that observed historically with Callisto, 1.34". The sky background flux value was left at the model's output value of $64 e^-/s/\text{pixel}$ under median atmospheric conditions at zenith. The higher read noise places SPIRIT as read noise limited for all the targets in the SPECULOOS target list, and decreases the overall performance of the instrument despite the higher dark current + thermal previously modelled. Nonetheless, it still presents a better performance than the existing instrumentation for targets below 2550 K, as exemplified by the observations made in Section 6.2.2. The better seeing typically seen with Europa would marginally change Figure 6.15, but still maintains the 2550 K temperature point as where SPIRIT begins to outperform the existing instrumentation.

SPIRIT similarly still maintains an advantage for the brightest targets of the SPECULOOS target list, independent of their temperature. This is primarily because of its significantly faster read time, which minimises the amount of deadtime in a binned period. However, this benefit for the brightest targets will be reduced if one were to observe with the existing instrumentation using a different/narrower filter than $I+z'$. For example, the brightest star in the target list (3000 K, 1.3 pc) would near match (but not improve upon) SPIRIT's binned photometric

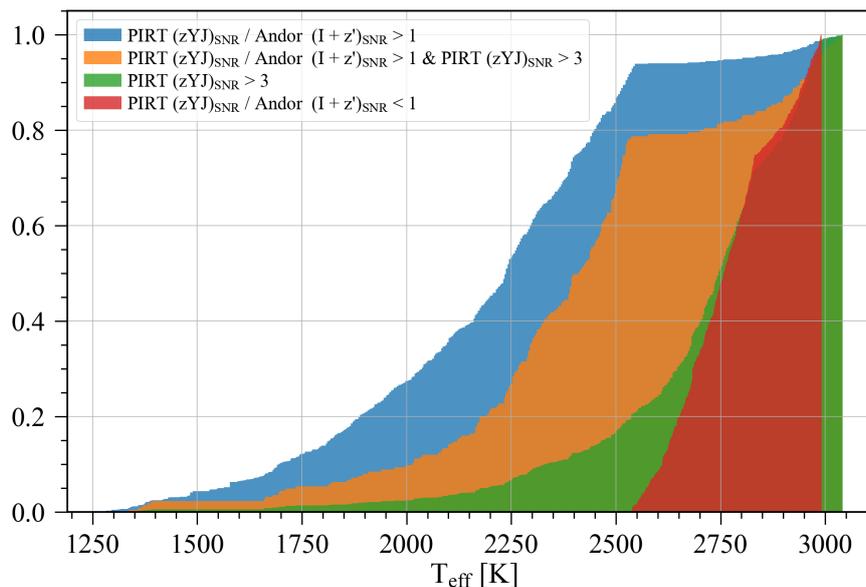


Figure 6.16: Cumulative histograms of different conditions evaluated on the modelled precision for SPIRIT and the existing instrumentation (see Figure 6.15) on the SPECULOOS target list, as a function of target effective temperature. Here, four conditions are evaluated. When the SNR ratio > 1 in blue (783 targets), SNR ratio > 1 and when SPIRIT's SNR > 3 (for the metric defined as part of Figure 4.10) in orange (225 targets), SPIRIT's SNR > 3 in green (909 targets), and finally when the SNR ratio < 1 in red (874 targets).

precision if instead it were to be observed with a r' filter. In $I+z'$, it would have a photometric precision ~ 4 times worse than SPIRIT's. Here, the existing instrumentation would increase its exposure time from 0.2 s ($I+z'$) to 4.8 s in r' , minimising the amount of deadtime in a binned period. However, one has to be careful with filter choice as it may introduce more observed variability, as illustrated in Section 6.3.

This leaves SPIRIT with 783 (from 1657) SPECULOOS targets that present an advantage to be observed over the existing instrumentation, 225 of this selection presenting an SNR > 3 (for the metric defined as part of Figure 4.10). The effective temperature distributions of the respective cuts are shown in Figure 6.16. Here, we can see the proportion of targets demonstrating a SPIRIT advantage, primarily for the coolest targets below 2550 K (94% of 783 targets), and a smaller fraction of benefit for the brightest targets above 2550 K. With 225 targets of potential interest to dedicate SPIRIT's time to (79% of which below 2550 K), a dedicated survey will be of interest. The order of survey priority should be evaluated in future work.

6.2.4 PWV insensitivity

The design of the zYJ filter (Section 4.2 and 6.1.4) was not only to increase the photometric precision of observations, but also to minimise the effect of PWV variability. The existing instrumentation with the $I+z'$ filter is highly sensitive to PWV changes during an observation, as discussed in depth in Chapter 3. It was thus desirable to mitigate those effects whilst maximising flux with a custom designed filter for SPIRIT. To test the effectiveness of the produced zYJ filter, we observed a known quiet target previously observed by SPECULOOS over a series of nights until a large PWV change occurred. The results of the observations are shown in Figure 6.17.

In Figure 6.17 we observed Sp1424-3514 over the course of a few days, a 2720 ± 104 K target at a distance of 19.1 ± 0.1 pc, with a J magnitude of 11.840 ± 0.022 . It was observed by both SPIRIT and Europa, at an exposure time of 12 s and 26 s respectively. The available comparison stars for both fields were all around 5000 K. During the observation run, we were lucky to observe a large PWV change, 5.6 mm to 2.1 mm, a Δ PWV of 3.5 mm, which induced a 0.98% change in Europa's light curve, and a 0.07% change in SPIRIT's as suggested by the modelled change. The induced changes in Europa's light curve was a factor of 15 higher than that seen by SPIRIT. The residual light curves show good visual agreement with each other, with a flare seen on the final night. This agreement is indicative that the PWV suppressing nature of the zYJ filter design is working as expected for this example target. However, to have greater confidence in its performance, further observations with cooler targets would permit the necessary statistical tests to assess the significance of zYJ's PWV suppressing capabilities. A cooler target would have been preferred for this test as it's expected to induce greater changes. Unfortunately, there wasn't a known quiet target that was cooler and observable when this test was performed. We did attempt to observe a cooler target at a later date over a period of a week, but the change of PWV was not as significant as the example assessed in this work to show the performance as well as this example.

This result has implications for the future of ground-based infrared observations which are time series orientated. If maximising photometric precision is the main objective of a survey, for example when searching for transits, then this work demonstrates that wider bandpasses beyond the standard filter sets (shown in Figure 4.4) can be considered, and still avoid PWV induced effects.

6.3 Example exoplanet photometry

The following subsections detail some of the early observations of transiting planets made with SPIRIT and the zYJ filter, in addition to observations made by other telescopes in other

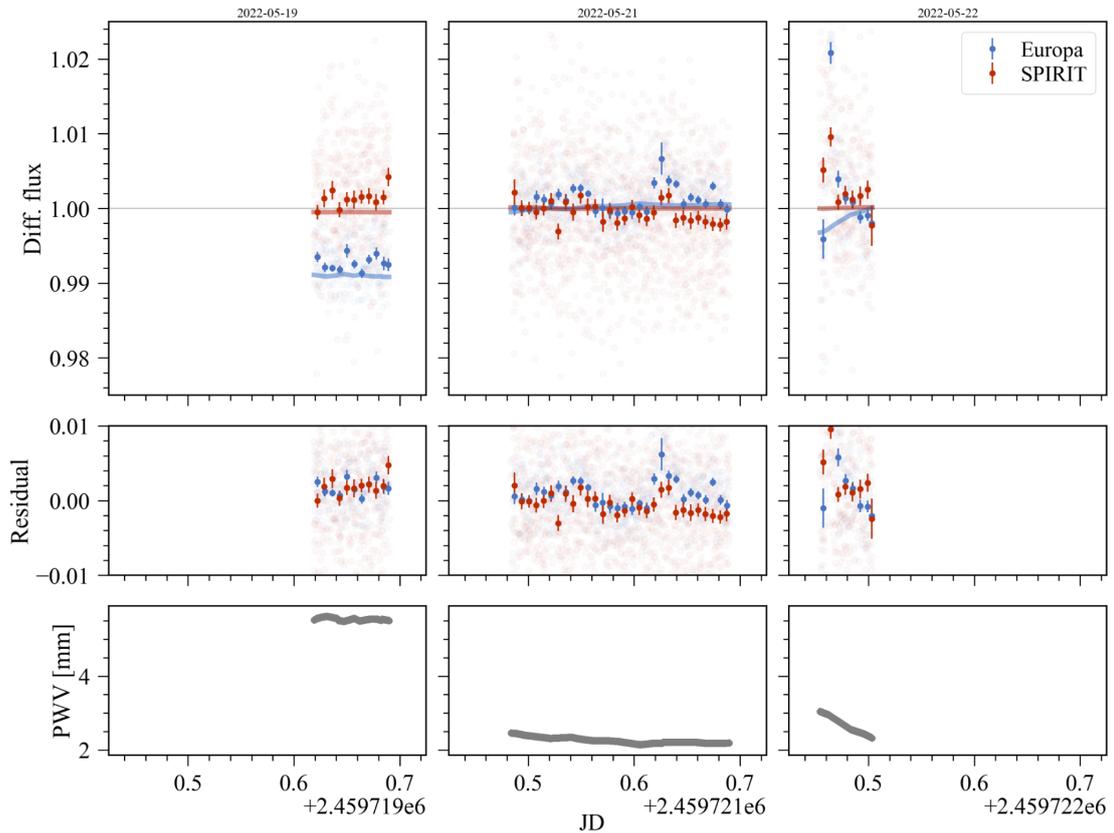


Figure 6.17: Globally normalised light curves of a 2720 K target observed by SPIRIT (zYJ) and Europa ($I+z'$). Top row: 10 minute binned differential light curves as solid points, and raw data as fainter points, without PWV correction. SPIRIT's observations in red, and Europa's in blue. Solid lines showing the expected trend induced by the PWV variability assuming a 5000 K comparison star, in the same respective colours as the observations. Middle row: The residual light curve after subtracting the expected trend induced by the PWV variability. Bottom row: PWV measurements made by the LHATPRO at zenith. Each x-axis has the same length to permit comparisons of feature changes over the same timescale. Missing data were due to either bad weather or technical downtime.

bandpasses for comparison.

6.3.1 TOI-3884.01

From observations made by TESS of TOI-3884 (Guerrero et al. 2021), a 3192 ± 157 K ($M4 \pm 0.5$) star, a super-Neptune candidate (TOI-3884.01) was found to be transiting with a period of 4.44 days and transit depth of 31 ± 17 ppt (values taken from Exoplanet Follow-up Observing Program (ExoFOP) database).^f From its transits, each lasting 1.4 hours, an interesting near-asymmetric feature was observed in transit. The recurring feature was suspected to be due to a large polar spot in view, which is partially obscured during TOI-3884.01's transit.

As stellar spots are cooler regions on the photosphere, it changes the observed stellar flux (Cracchiolo et al. 2020; Pont et al. 2013) when in view. Thus, stellar spots can shift the apparent transit depth, where visible spots will reduce the amount of flux captured by an observer. When spots are occulted by a transiting planet, an apparent increase in flux is observed in transit as it moves from occulting a non-spot to a spot region. This leads to issues with estimating the planetary radius. The effect however is chromatic, where the contrast between the spot and non-spot regions will change as a function of wavelength - permitting one to possibly constrain the transit depth with multiband photometry.

The contrast between a cooler spot region will become less apparent at longer wavelengths in the NIR domain, and at a stellar effective temperature of 3192 ± 157 K, the contrast increases steeply at visible wavelengths given how a blackbody spectra's peak transitions further into the infrared at cooler temperatures (c.f. Figure 4.2) in addition to the inherent lower flux. To test this hypothesis, multiband photometry was performed with a range of SPECULOOS and partnering facilities. Given SPIRIT's higher sensitivity later in the infrared over the other available instruments, it presented a unique opportunity to test the instrument and polar spot hypothesis. Table 6.2 details the series of multiband observations that were thus made.

There were three attempts to get photometry with SPIRIT. The first attempt (2022-05-15) unfortunately did not succeed. The wind that evening was hovering around SPIRIT's dome's wind limit, which prevented capturing the majority of the transit. Moreover, SPIRIT's narrow field of view, with the target centred, did not yield any suitable comparison stars to perform differential photometry with. The field only had one other visible star, which was oversaturated. The field was not previously vetted, unlike the fields in SPECULOOS' target list (Section 4.5.3). The other SSO telescope also observing TOI-3884 (in r') that night did not have its wind limit triggered.

For the next visible transit, an ingress event (2022-05-24), a new field of view was vetted with 2MASS catalogues and manually vetted prior to observing. There was only one possible

^fDate accessed 2022-07-03: <https://exofop.ipac.caltech.edu/tess/target.php?id=86263325>

Table 6.2: Range of ground-based multiband observations of TOI-3884.01's transit.

| Observatory | Bandpass | Date |
|-------------------------------|---------------|---------------|
| SSO | r' | 2022-05-15 |
| | z' (ingress) | 2022-05-24 |
| SSO - SPIRIT | zYJ (ingress) | 2022-05-24 |
| TRAPPIST South | I+z' | 2022-05-15 |
| | z' | 2022-05-25 |
| Las Cumbres Observatory-Tiede | g' | 2022-04-27 |
| | i' | 2022-05-06 |
| Las Cumbres Observatory-SAAO | g' | 2022-04-27 |
| TESS | 600 – 1000 nm | Sector 46, 49 |

field of view which could provide more than one comparison star. In this new field of view, there were only two comparison stars, one with approximately half the flux and the other with twice the flux of TOI-3884. To capture the flux of both comparison stars and to maintain an exposure time (> 7 s) which ACP could handle, I had to defocus to match ACP's lower exposure limit of 7 s so as to not saturate the brightest comparison star. Defocusing decreased the photometric performance of the observation given the increased number of pixels required to form an aperture. Nonetheless, it yielded a binned photometry sufficient to determine differences in the transit feature (Figure 6.18), with the comparison stars stable. The alternative option would have been to expose for half the duration whilst in focus, however ACP would have missed at least half the images, yielding a similar precision to what was attained as suggested by my models. The final attempt (2022-06-25) was unsuccessful due to bad weather conditions (clouds and rain). The star is no longer visible, we will have to wait till next year for the next opportunity to observe.

Figure 6.18 shows all the multiband photometry collected for TOI-3884.01. The in transit feature is visible in all passbands, displaying similar asymmetry with an inflection near mid-transit. The amplitude of the feature generally decreases with longer wavelengths as would be expected for a transit occurring over a stellar spot. The amplitude of the feature with SPIRIT's data was seen to be ~ 8 ppt, whereas the feature in g' was seen to be ~ 24 . The exact origin of this feature will be impossible to determine given the degeneracy of the problem. If a polar spot does indeed exist, the size of the spot, and the exact alignment of the polar axis with respect to the planet's orbital plane and the observer has numerous possible setups. Similarly, the rate of rotation, and circular symmetry of the spot will have an effect on the amplitude of the feature over time, and likewise any rate of stellar precession. Longer term monitoring of this target will therefore be necessary to constrain any of these variables. However, the unknown evolution of the size and shape of the potential spot will make this difficult. Radial velocity observations will help constrain the system further, but we do not currently have such measurements.

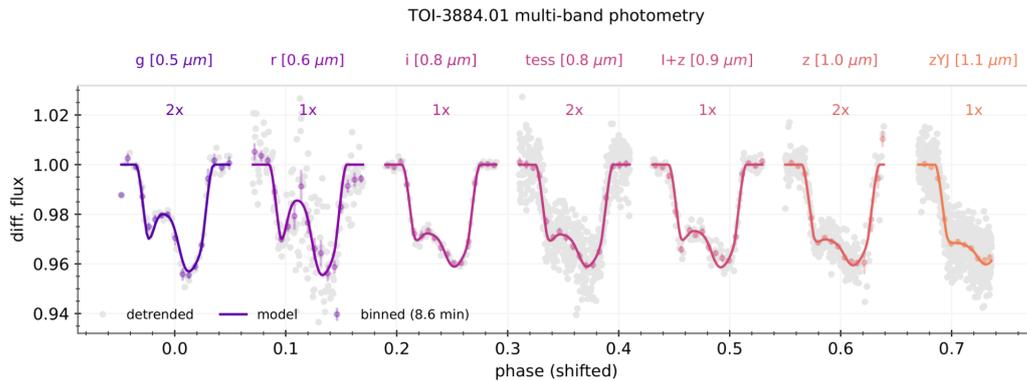


Figure 6.18: Range of multiband photometry of the TOI-3884.01 transit, in 7 passbands, with the ground-based photometry processed with `Prose` (Garcia et al. 2022). The label and centre of each passband placed above their respective data. A modelled scenario (detailed in Figure 6.19) is superimposed onto each passband’s data. Figure credit: Lionel Garcia

From the shape of the in-transit feature, we can see where the planet lies in the dimmest portion of the stellar surface from our line of sight along its orbital path (the peak of the feature), and then out of this feature (as suggested by the inflection point) to a brighter portion of the stellar surface (the trough of the feature where the planet then exits stellar disc). We can estimate the angular size of the portion of the spot the planet is transitioning through based on the time it took to reach the peak and inflection point of the in-transit feature. However, with the effect of limb darkening, the spot’s edge hardness, and the planet potentially having an atmosphere, the exact respective edge positions of the spot will have multiple solutions.

Considering the simplest polar spot scenario, with the polar axis in line with the orbital path of the planet, we can use a Python package called `starry` by Luger et al. (2019) to iterate through possible scenarios of spot position along the orbital path, size, and edge hardness. This aspect of the modelling was performed by Lionel Garcia and is shown in Figure 6.19, where he assumed a planet radius of $6.0 R_E$. As part of modelling, he also iterated through the various spot temperatures below the stellar effective temperature of 3192 K using the BT Settl models (Allard et al. 2012), which then quantifies the spot contrast expected at different wavelengths - where the spot contrast was defined as the fraction of flux from the spot’s centre divided by non-spot stellar flux per unit surface area. We can estimate the spot contrast by measuring the amount of light left at the peak and trough of the features, where the contrast can be calculated to be as $(1 - \text{peak}) / (1 - \text{trough})$. Figure 6.19 shows the resulting position, size, and modelled temperature of the polar spot. 200 K below the star’s effective temperature was found to be the closest temperature for the spot with a spot radius of 27.14 degrees. This work will go into

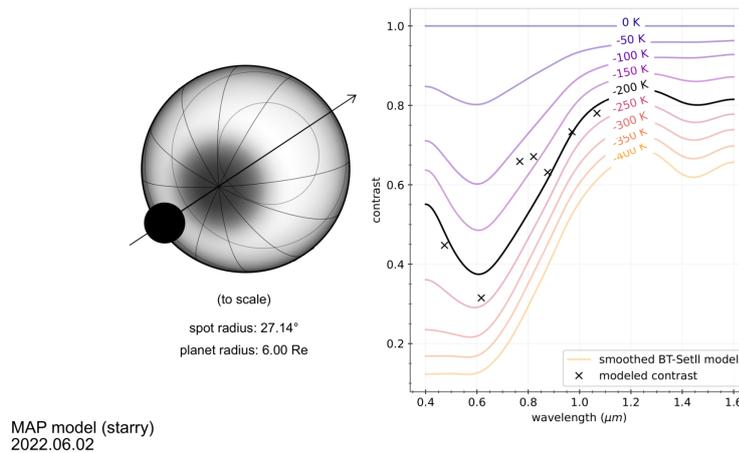


Figure 6.19: Left: Modelled polar spot position and size with respect to the orbital path of TOI-3884.01. Right: Spot contrast modelled with BT-Settl models as a function of wavelength at different temperatures below the stellar effective temperature of 3192 K. The points (x) show the measured spot contrast using the peak and troughs for the band passes in Figure 6.18, positioned at the centres of their wavelengths. Figure credit: Lionel Garcia.

much further depth in Garcia et al. (in prep.).

6.3.2 Trappist-1b

We have similarly observed other transiting planets with success^g, most notably Trappist-1b (Gillon et al. 2017) on the night beginning 2022-08-06. Again, observing simultaneously with Europa and the $I+z'$ filter, we were able to recover the transit in both light curves, shown in Figure 6.20.

The FWHM for Trappist-1 on SPIRIT's stack was 1.46", larger than that of Europa's (1.17") - even after focus checking. Thankfully however, the SPIRIT's FWHM was very stable over the course of the observation, permitting the use of a small aperture to maintain good precision to compare with Europa's data. Another issue with SPIRIT's observation was the availability of comparison stars. There were effectively only two good comparison stars in SPIRIT's field of view to give the precision shown in Figure 6.20, further emphasising the need for a broader field of view for SPIRIT.

There was good agreement between SPIRIT's light curve and the transit model produced by Batman (Kreidberg 2015). The transit models were produced using quadratic limb darkening coefficients (u_1, u_2) from Claret & Bloemen (2011), in J band observing a 2600 K star ($u_1 = 0.10, u_2 = 0.57$), and the stellar and planet parameters from Gillon et al. (2017). There was

^gsee Appendix C for an example reduced image and further light curves produced by SPIRIT

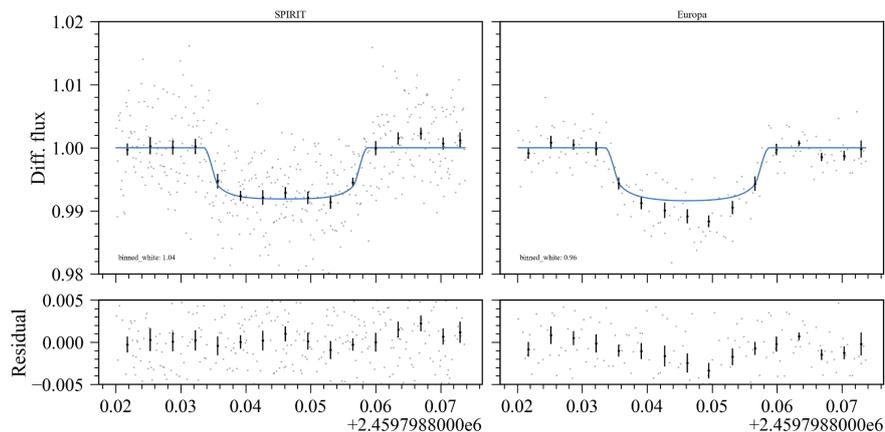


Figure 6.20: Simultaneous observations of Trappist-1b by SPIRIT (zYJ) and Europa ($I+z'$). Top row: 5 minute binned light curves (in black, raw data in grey), with a transit model from Batman (Kreidberg 2015). Both observations were not corrected for any PWV changes. Bottom row: Residual from transit model. The x-axis is Julian Date (JD).

less good agreement between Europa’s light curve and the transit model, which similarly used a quadratic limb darkening model derived from Claret & Bloemen (2011) by a SPECULOOS team member for the $I+z'$ band observing Trappist-1 ($u_1 = 0.24, u_2 = 0.48$). A limb darkening model specifically for the zYJ was not calculated as the J band model was deemed sufficient.

There was a deviation of -3.4 ppt in transit from Europa’s binned data from the 8.4 ppt transit depth. It was not correlated with any atmospheric, instrument, or comparison star systematics, suggesting perhaps it could be due to a faculae being obscured in transit. However, the lack of a similar feature in SPIRIT’s data does not give confidence to the idea. The zenith PWV data from the LHATPRO did not show any significant changes to suggest it was PWV induced. However, it may be that there was a line-of-sight induced PWV change that was not captured by the LHATPRO.

6.4 Conclusions and future improvements

In this work, a new infrared instrument (SPIRIT) tailored to the photometric applications of cool stars has successfully been integrated at SPECULOOS’ Southern Observatory. Beyond a successful technology demonstration of an InGaAs based detector as a cost-effective, robotic observatory friendly alternative to the traditionally used HgCdTe based detectors in astronomy. It has also permitted one to demonstrate the advantages of moving into the infrared for observing late type M, L type stars for SPECULOOS. From reducing the effect of stellar features on light curves in Section 6.3.1, to increasing the photometric precision of final data products in

Sections 6.2.2.

Likewise, this new instrument presented the opportunity to design a bandpass to minimise the effect of PWV on our light curves, which has significantly impacted the existing instrumentation (as discussed in depth in Chapter 3). An experimental filter design, described in Section 6.1.4, covered a much wider wavelength range than traditional filters. Observations made in Section 6.2.4 demonstrated its invariance to PWV changes. However, a global study on cooler stars is desired before drawing firm conclusions.

The models derived in Chapter 4 showed close agreement with the noise models presented in Section 6.2.3. However, there have been instances where the read noise of an observation has been suspected to have effectively increased by 40%, nearly 3x higher than the goal specifications in Appendix A. A further investigation is necessary to diagnose the issue. To improve future InGaAs based instrumentation, the read noise needs to be significantly reduced, for example by introducing Fowler sampling (Fowler & Gatley 1990). Nonetheless, the photometric precision of SPIRIT is still expected to be better than the Si based instrumentation for targets below 2550 K, allowing us to probe cooler into potential host star's of exoplanets within SPECULOOS' target list.

This success has not come without pain, as described in Section 6.1.5. Nonetheless, through the lessons learnt here, there is plenty of scope for improvement, from improving SPIRIT to future instruments. In the instance of improving SPIRIT, increasing the field of view would yield multiple benefits. Specifically, by giving an increased number of comparison stars for the differential photometry process, this would increase the photometric precision of the final light curve, but likewise avoid the issues initially faced in Section 6.3.1. It would similarly allow us to observe more fields, as calculated in Section 4.5. Moreover, the photometric precision would be expected to improve, as exemplified in Figure 4.12, given the lower number of pixels used within an aperture and shorter exposure times. However, before pursuing this avenue, the effect of bad pixels and flat-fielding errors should be evaluated before committing to an intervention on the optical train of the telescope.

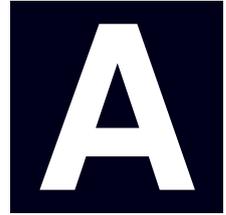
As seen in the observations shown in this chapter, the FWHM of SPIRIT was higher than a neighbouring telescope. An investigation into this is recommended, as lowering the spread of a star's flux on the detector would improve a light curve's photometric precision (see Figure 4.13), however similar to changing the field of view, the effect of bad pixels and flat-fielding errors should be evaluated prior to targeting FWHM issues. One further avenue to improve SPIRIT's photometric precision would be to consider cooling the optical region in front of the detector, as calculated in Section 6.2.1, to minimise the influence of the noise and variation induced by the thermal background flux.

The last suggestion to improve SPIRIT is to modify or replace the software control layer

ACP, which is limiting SPIRIT to observe targets above an exposure time of 7 s, as described in Section 6.1.3. The pipeline to then produce the light curves should likewise be optimised, taking perhaps the lessons learnt already in astronomy with HgCdTe based detectors.

As well as lower-cost ground based instruments, space missions with InGaAs detectors are also likely to be more common in the future, in part due to their significantly lower costs (see Table 4.1), but also due to the increasing accessibility of launching satellites. In recent years, space based InGaAs detectors have been demonstrated for Earth remote sensing applications (e.g. [Pack et al. \(2022\)](#)), however none have been used for astrophysical applications. For space-based astronomy, any infrared detector's photometric precision will increase given the lack of noise introduced by the Earth's atmosphere (specifically, noise from sky background, scintillation, PWV variability, and variable seeing conditions). In addition, a space-based instrument could consider using a wider bandpass to increase its photometric precision. In the context of observing cool stars for transits, it would be of benefit to consider a bandpass optimised to minimise the contrasting effects of stellar spots. In both scenarios, one would not have to worry about the effects of PWV variability or sky background noise as was done for SPIRIT. However, the susceptibility to radiation damage on InGaAs detectors and the additional engineering required for cooling and shielding from the Sun's thermal radiation for a desired low-cost space mission will need to be evaluated. Likewise, the minimum mirror size and cost for such a satellite and desired science goals will need to be compared to the cost of a ground-based equivalent.

APPENDIX



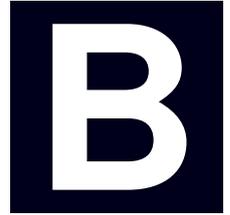
1280SciCAM PROCUREMENT DOCUMENT

An extract of the procurement document for the Princeton Infrared Technologies 2nd generation 1280SciCam showing the core technical requirements.

| Parameter | Unit | Values | Verification method | Notes |
|--|--------------------------------------|---------------------------------|---------------------|--|
| Array format | pixel | 1280 × 1024 | By design | |
| Pixel pitch | μm | 12 | By design | |
| Data output | bits | 14 | By design | |
| Readout rate | Hz | > 1 | Measured | Equivalent to < 1 s period between consecutive images. |
| Detector material | | InGaAs | By design | |
| Detector substrate | | InP – removed | By design | |
| Mean total system efficiency between 0.81 – 1.33 μm | % | > 55 | Measured | with default AR coatings, optimised at 1.3 μm. |
| Median dark current | e ⁻ /s/pixel | < 300 | Measured | PIRT will screen for lowest dark current in inventory |
| Median read noise | e _{rms} ⁻ /pixel | < 45 | Measured | |
| Well capacity (without bias) | e ⁻ /pixel | > 45000 | Calculated | |
| Inoperable pixels | % | < 0.5 | Measured | |
| Non-linearity | % | < 1 | Measured | Extra test – nonstandard measurement |
| Persistence after 1 s (after exposure filling 80 % well depth) | % | < 0.1 | Measured | Extra test – nonstandard measurement |
| SCA Flatness | μm | < 20 | By design | Assumed |
| Temperature stability | °C | ± 0.2 | Measured | Extra test – nonstandard measurement |
| Integration time lower limit | s | < 0.01 | Measured | |
| Integration time upper limit | s | > 150, or until all wells full. | Measured | |
| Read out mode | | Global. CDS . | By design | |
| Ambient operating temperature | °C | -8 to +25 | Measured | Above 15 °C ambient water cooling required to maintain -60 °C setpoint |
| Ambient operating humidity | % | 0 to 80 | By design | PIRT does not have equipment to test |

Table 1: Core technical requirements for the [PIRT](#) 1280SciCam 2nd generation.

APPENDIX



FILTER PROCUREMENT DOCUMENT

An extract of the procurement document for zYJ, J_{MKO}, H-long filters showing the core technical requirements.

1 INTRODUCTION

The following document presents the desired specifications of three NIR filters, intended for the ground-based exoplanet transit search, SPECULOOS. The following section outlines the technical requirements for one standard astronomy filters, ‘MKO J’ (Tokunaga and Simons (2003)), one ‘long pass H’ filter, and one custom filter, ‘z+Y+J’.

2 FILTER SPECIFICATIONS

Table. (1) details the base specification for all filters. Table. (2) details the transmission range of each filter, with the corresponding illustrative filter profiles in Figures. (1-3). Table. (3) details information about the telescope the filter will be mounted on. Application of AR coating(s) is expected, upon recommendation of the provider, to minimise internal reflections between the camera and filter, and to maximise transmission.

| Parameter | Goal specification |
|---------------------------------------|---|
| Average total transmission* | > 90% transmission, have a ripple of less than $\pm 5\%$ |
| Roll off slope [†] | < 2.5% |
| Wedge of substrate | < 0.5 arcmin |
| Angle of Incidence | 0 ± 5 degrees |
| Cone half angle | 3.6 degrees |
| Out-of-band transmission [‡] | > OD4, starting < 70 nm from each HPP till 0.3 μm and 1.85 μm |
| Flatness | < $\lambda/4$ at 0.63 μm peak to valley |
| Surface Quality | 40/20 over central 90% clear aperture |
| Material | Single substrate, non-radioactive |
| Thickness | 3.0 mm ± 0.1 mm |
| Size | 50.0 mm \times 50.0 mm +0.0/-0.4 mm |
| Ambient operating temperature | -8 °C to 25 °C |
| Ambient operating humidity | 0 % to 80 % |

Table 1: Base technical specifications.

*This includes the substrate. This applies to the long pass H filter only up to 1.7 μm .

[†]slope = $[\lambda(80\%) - \lambda(5\%)]/\lambda(5\%) \times 100$

[‡]1.85 μm limit does not apply to the long pass H filter.

Appendix B. Filter procurement document

| Name | Centre | HPP 1 | HPP 2 |
|-------------|--------|-------|-------|
| MKO J | 1.250 | 1.170 | 1.330 |
| Long pass H | - | 1.517 | - |
| z+Y+J | 1.068 | 0.805 | 1.330 |

Table 2: Filter centres and HPPs in μm , errors $< \pm 0.007 \mu\text{m}$ (equivalent to $\pm 0.010 \mu\text{m}$ on the FWHM).

| Parameter | Value |
|--------------|------------------------------------|
| Filter wheel | FLI CFW3-10, filters are unmounted |
| Focal ratio | F8 |

Table 3: Telescope parameters.

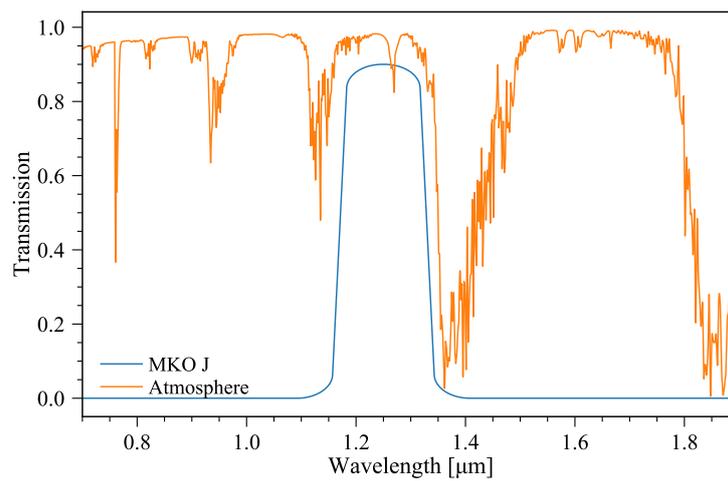


Figure 1: MKO J filter profile superimposed on the median atmospheric transmission at Paranal, Chile.

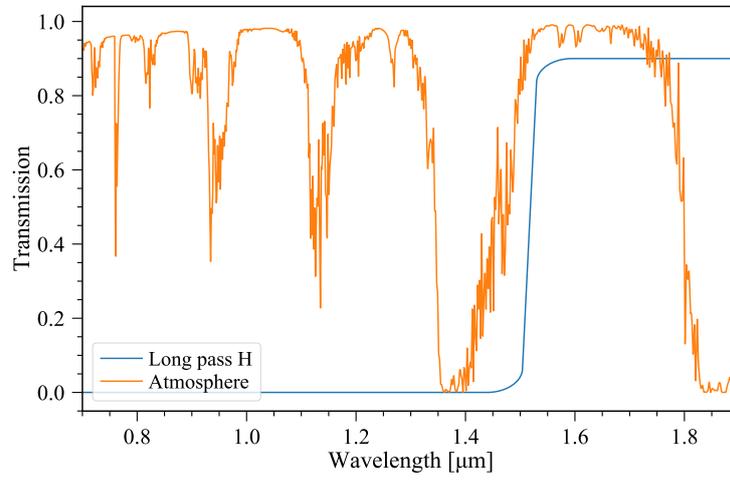


Figure 2: Long pass H filter profile superimposed on the median atmospheric transmission at Paranal, Chile.

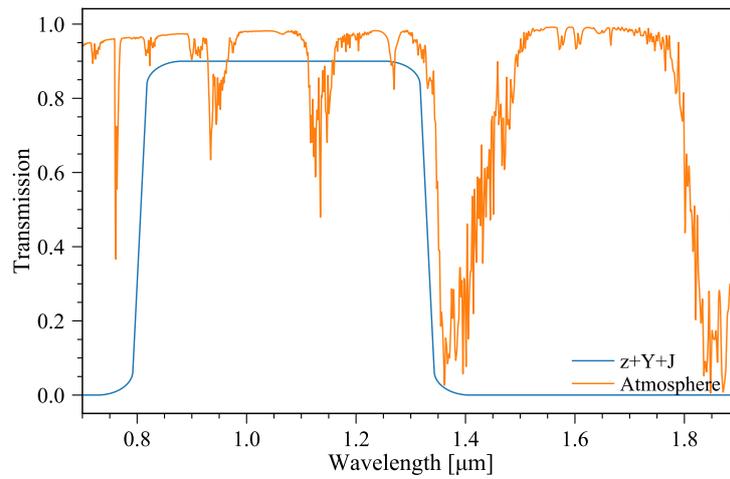
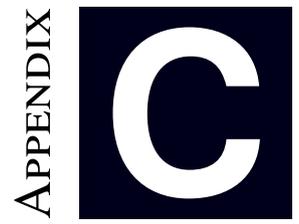


Figure 3: z+Y+J filter profile superimposed on the median atmospheric transmission at Paranal, Chile.



**EXAMPLE IMAGE AND LIGHT CURVES FROM
SPIRIT**

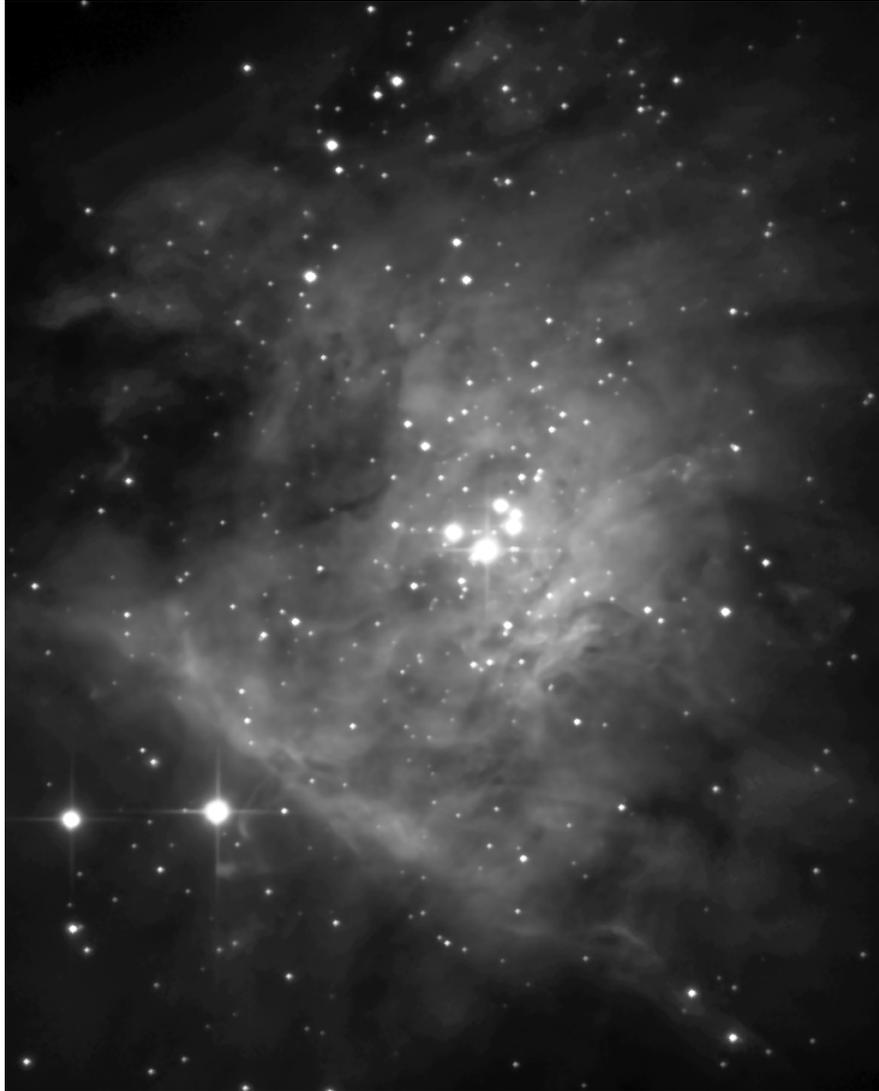


Figure C.1: A reduced image of the Trapezium Cluster taken by SPIRIT with the zYJ filter.

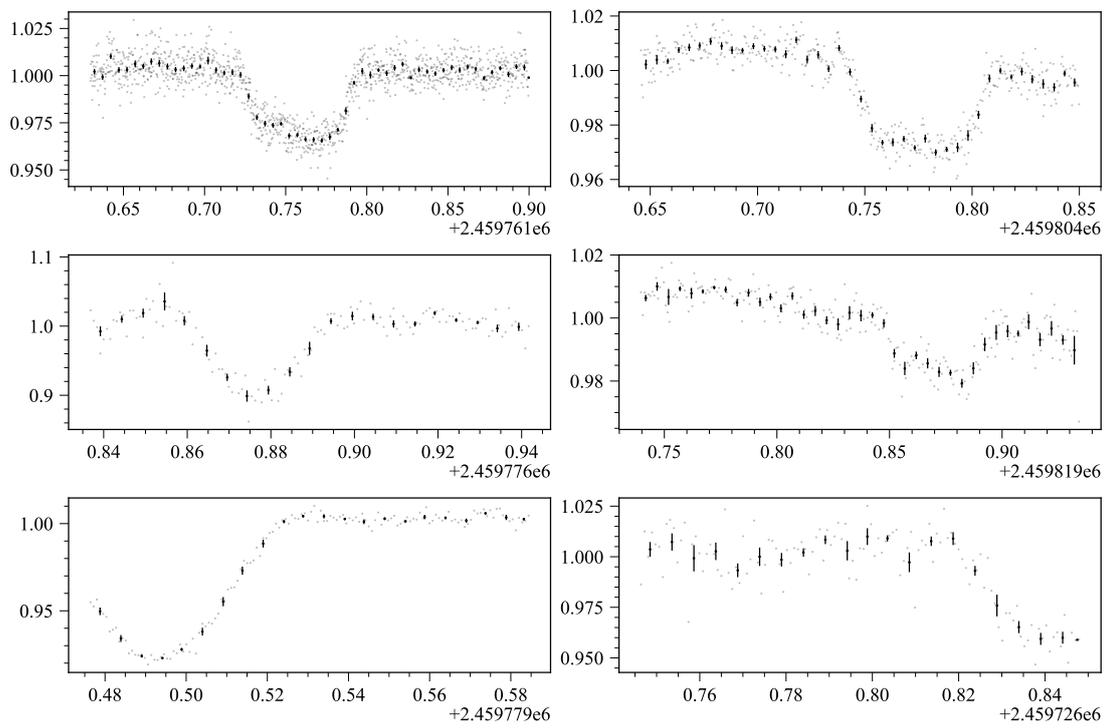


Figure C.2: A series of transits or occultations of different M dwarfs from TESS Objects of Interest catalog, observed by SPIRIT.

REFERENCES

- Absil O., Delacroix C., de Xivry G. O., Pathak P., Willson M., Berio P., van Boekel R., Matter A., Defrère D., Burtscher L., et al., 2022, in *Adaptive Optics Systems VIII*, pp 298–310
- Alibert Y., Benz W., 2017, *A&A*, **598**, L5
- Allard F., Homeier D., Freytag B., 2012, *Philosophical Transactions of the Royal Society of London Series A*, **370**, 2765
- Astier P., Antilogus P., Juramy C., Le Breton R., Le Guillou L., Sepulveda E., 2019, *Astronomy & Astrophysics*, **629**, A36
- Baker A. D., Blake C. H., Sliski D. H., 2017, *Publications of the Astronomical Society of the Pacific*, **129**, 085002
- Barclay T., Pepper J., Quintana E. V., 2018, *ApJS*, **239**, 2
- Bardalez Gagliuffi D. C., Burgasser A. J., Schmidt S. J., Theissen C., Gagné J., Gillon M., Sahlmann J., Faherty J. K., Gelino C., Cruz K. L., Skrzypek N.,Looper D., 2019, *The Astrophysical Journal*, **883**, 205
- Baumer M., Davis C. P., Roodman A., 2017, *Publications of the Astronomical Society of the Pacific*, **129**, 084502
- Beletic J. W., Blank R., Gulbransen D., Lee D., Loose M., Piquette E. C., Sprafke T., Tennant W. E., Zandian M., Zino J., 2008, in *High Energy, Optical, and Infrared Detectors for Astronomy III*, pp 161–174
- Benneke B., et al., 2019, *ApJ*, **887**, L14
- Berta-Thompson Z. K., et al., 2015, *Nature*, **527**, 204
- Bhan R., Dhar V., 2019, *Opto-Electronics Review*, **27**, 174
- Birch M., Soon J., Travouillon T., Mendel T., Taylor B., Tiong B. C. K., 2022, *Journal of Astronomical Telescopes, Instruments, and Systems*, **8**, 016001
- Blake C. H., Bloom J. S., Latham D. W., Szentgyorgyi A. H., Skrutskie M. F., Falco E. E., Starr D. S., 2008, *Publications of the Astronomical Society of the Pacific*, **120**, 860
- Bohndiek S. E., Blue A., Clark A. T., Prydderch M. L., Turchetta R., Royle G. J., Speller R. D., 2008, *IEEE Sensors Journal*, **8**, 1734
- Boyd R. W., 1978, *J. Opt. Soc. Am.*, **68**, 877

- Boyle W. S., Smith G. E., 1970, *Bell System Technical Journal*, 49, 587
- Bradley L., Sipocz B., Robitaille T., Tollerud E., Deil C., Vinícius Z., Barbary K., Günther H. M., Bostroem A., Droettboom M., et al., 2016, *Astrophysics Source Code Library*, pp ascl-1609
- Broeg C., Fernández M., Neuhäuser R., 2005, *Astronomische Nachrichten: Astronomical Notes*, 326, 134
- Burdanov A., Delrez L., Gillon M., Jehin E., 2018, *SPECULOOS Exoplanet Search and Its Prototype on TRAPPIST*. p. 130, doi:10.1007/978-3-319-55333-7_130
- Burdanov A. Y., de Wit J., Gillon M., Rebolo R., Sebastian D., Alonso R., Sohy S., Niraula P., Garcia L., Barkaoui K., et al., 2022, arXiv preprint arXiv:2209.09112
- Burke C. J., Bryson S. T., Mullally F., Rowe J. F., Christiansen J. L., Thompson S. E., Coughlin J. L., Haas M. R., Batalha N. M., Caldwell D. A., et al., 2014, *The Astrophysical Journal Supplement Series*, 210, 19
- Charbonneau D., Brown T. M., Latham D. W., Mayor M., 2000, *The Astrophysical Journal*, 529, L45
- Charbonneau D., Berta Z. K., Irwin J., Burke C. J., Nutzman P., Buchhave L. A., Lovis C., Bonfils X., Latham D. W., Udry S., et al., 2009, *Nature*, 462, 891
- Chuh T., Loose M., Gulbransen D. J., Anglin S. W., Beletic J., Piquette E. C., Garnett J. D., 2006, in *Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter*. p. 62652K
- Claret A., Bloemen S., 2011, *Astronomy & Astrophysics*, 529, A75
- Cohen M., Wheaton W. A., Megeath S., 2003, *The Astronomical Journal*, 126, 1090
- Coleman G. A. L., Leleu A., Alibert Y., Benz W., 2019, *A&A*, 631, A7
- Cracchiolo G., Micela G., Peres G., 2020, *Monthly Notices of the Royal Astronomical Society*, 501, 1733
- Crossfield I. J. M., et al., 2019, *ApJ*, 883, L16
- Dangi M., Saxena R. S., Niranjana V., 2022, *Microsystem Technologies*, pp 1–10
- Delrez L., et al., 2018, in *Proceedings of the SPIE*. p. 107001I (arXiv:1806.11205), doi:10.1117/12.2312475
- Delrez L., et al., 2022, *A&A*, 667, A59
- Demory B.-O., Gillon M., Madhusudhan N., Queloz D., 2016, *Monthly Notices of the Royal Astronomical Society*, 455, 2018
- Demory B. O., et al., 2020, *A&A*, 642, A49
- Dittmann J. A., et al., 2017, *Nature*, 544, 333
- Downing M., Baade D., Sinclair P., Deiries S., Christen F., 2006, in *High Energy, Optical, and Infrared Detectors for Astronomy II*. pp 76–86
- Ettenberg M. H., Nguyen H., Martin C. R., Lange M., 2018, in *Infrared Technology and*

- Applications XLIV. pp 22–28
- Fowler A., Gatley I., 1990, *The Astrophysical Journal*, 353, L33
- Gaia Collaboration et al., 2016, *Astronomy and Astrophysics*, 595, A1
- Gaia Collaboration et al., 2018, *A&A*, 616, A1
- Garcia-Mejia J., Charbonneau D., Fabricant D., Irwin J., Fata R., Zajac J. M., Doherty P. E., 2020, in *Ground-based and Airborne Telescopes VIII*. pp 1387–1406
- Garcia L. J., Timmermans M., Pozuelos F. J., Ducrot E., Gillon M., Delrez L., Wells R. D., Jehin E., 2022, *Monthly Notices of the Royal Astronomical Society*, 509, 4817
- Gaudi B. S., 2012, *Annual Review of Astronomy and Astrophysics*, 50, 2012
- Gibbs A., et al., 2020, *AJ*, 159, 169
- Gillon M., 2018, *Nature Astronomy*, 2, 344
- Gillon M., Jehin E., Magain P., Chantry V., Hutsemékers D., Manfroid J., Queloz D., Udry S., 2011, in *European Physical Journal Web of Conferences*. p. 06002 ([arXiv:1101.5807](https://arxiv.org/abs/1101.5807)), [doi:10.1051/epjconf/20101106002](https://doi.org/10.1051/epjconf/20101106002)
- Gillon M., Jehin E., Fumel A., Magain P., Queloz D., 2013, in *European Physical Journal Web of Conferences*. p. 03001, [doi:10.1051/epjconf/20134703001](https://doi.org/10.1051/epjconf/20134703001)
- Gillon M., et al., 2016, *Nature*, 533, 221
- Gillon M., et al., 2017, *Nature*, 542, 456
- Glass I. S., Glass I., 1999, *Handbook of infrared astronomy*. No. 1, Cambridge University Press
- Glendinning R., Beard S. M., Mountain C. M., Pettie D. G., Pickup D. A., Wade R., 1990, in *Instrumentation in Astronomy VII*. pp 34–43
- Google 2022, Google Maps Elevation API digital elevation model, <https://developers.google.com/maps/documentation/elevation/overview/>
- Guerrero N. M., Seager S., Huang C. X., Vanderburg A., Soto A. G., Mireles I., Hesse K., Fong W., Glidden A., Shporer A., et al., 2021, *The Astrophysical Journal Supplement Series*, 254, 39
- Günther M. N., Berardo D. A., Ducrot E., Murray C. A., Stassun K. G., Olah K., Bouma L., Rappaport S., Winn J. N., Feinstein A. D., et al., 2022, *The astronomical journal*, 163, 144
- Hardegree-Ullman K. K., Cushing M. C., Muirhead P. S., Christiansen J. L., 2019, *AJ*, 158, 75
- He M. Y., Triaud A. H., Gillon M., 2017, *Monthly Notices of the Royal Astronomical Society*, 464, 2687
- Hecht J. H., Ramsay Howat S. K., Walterscheid R. L., Isler J. R., 1995, *Geophysical research letters*, 22, 2873
- Heidt J., 2022, *Astronomy in the Near-Infrared-Observing Strategies and Data Reduction Techniques*. Springer Nature
- Hoffman A., Sessler T., Rosbeck J., Acton D., Ettenberg M., 2005, in *Infrared Technology and*

- Applications XXXI. pp 32–39
- Holzlohner R., Kosmalski J., Guisard S., 2018, in *Ground-based and Airborne Telescopes VII*. pp 146–160
- Hopkinson G. R., Goodman T. M., Prince S. R., 2004, *A Guide to the use and calibration of detector array equipment*
- Howard A. W., et al., 2012, *ApJS*, **201**, 15
- Howell S. B., 2006, *Handbook of CCD Astronomy*, 2 edn. Cambridge Observing Handbooks for Research Astronomers, Cambridge University Press, doi:10.1017/CBO9780511807909
- Irwin M. J., Lewis J., Hodgkin S., Bunclark P., Evans D., McMahon R., Emerson J. P., Stewart M., Beard S., 2004, in Quinn P. J., Bridger A., eds, *Proc. SPIE Vol. 5493, Optimizing Scientific Return for Astronomy through Information Technologies*. pp 411–422, doi:10.1117/12.551449
- Irwin J., Charbonneau D., Nutzman P., Falco E., 2008, *Proceedings of the International Astronomical Union*, **4**, 37
- Irwin J., Berta Z. K., Burke C. J., Charbonneau D., Nutzman P., West A. A., Falco E. E., 2011, *The Astrophysical Journal*, **727**, 56
- Janesick J. R., 2007.
- Jehin E., Gillon M., Queloz D., Magain P., Manfroid J., Chantry V., Lendl M., Hutsemékers D., Udry S., 2011, *The Messenger*, **145**, 2
- Jehin E., et al., 2018, *The Messenger*, **174**, 2
- Johnson H. L., 1966, *The Astrophysical Journal*, **143**, 187
- Jones A., Noll S., Kausch W., Szyszka C., Kimeswenger S., 2013, *Astronomy & Astrophysics*, **560**, A91
- Jun G., Zhongjian C., Wengao L., Wentao C., Lijiu J., 2004, in *Proceedings. 7th International Conference on Solid-State and Integrated Circuits Technology, 2004..* pp 1437–1440
- Kaltenegger L., Traub W. A., 2009, *The Astrophysical Journal*, **698**, 519
- Karpov S., Bajat A., Christov A., Prouza M., Beskin G., 2020, in *X-Ray, Optical, and Infrared Detectors for Astronomy IX*. pp 86–98
- Kerber F., Rose T., Chacón A., Cuevas O., Czekala H., Hanuschik R., Momany Y., Navarrete J., Querel R. R., Smette A., et al., 2012, in *Ground-based and Airborne Instrumentation for Astronomy IV*. p. 84463N
- Kerber F., Querel R., Neureiter B., Hanuschik R., 2017, in *ESO Calibration Workshop: The Second Generation VLT Instruments and Friends*. p. 22, doi:10.5281/zenodo.887263
- Kirkpatrick J. D., 2005, *ARA&A*, **43**, 195
- Kirkpatrick J. D., Henry T. J., Irwin M. J., 1997, *AJ*, **113**, 1421
- Koch D. G., Borucki W. J., Basri G., Batalha N. M., Brown T. M., Caldwell D., Christensen-

- Dalsgaard J., Cochran W. D., DeVore E., Dunham E. W., et al., 2010, *The Astrophysical Journal Letters*, 713, L79
- Kopparapu R. K., Ramirez R., Kasting J. F., Eymet V., Robinson T. D., Mahadevan S., Terrien R. C., Domagal-Goldman S., Meadows V., Deshpande R., 2013, *ApJ*, 765, 131
- Kreidberg L., 2015, *PASP*, 127, 1161
- Krishnamurthy A., Villasenor J., Seager S., Ricker G., Vanderspek R., 2019, *Acta Astronautica*, 160, 46
- Laidler V. G., Greenfield P., Busko I., Jedrzejewski R., 2008, in *Proceedings of the 7th Python in Science Conference*. pp 36–38
- Lang D., Hogg D. W., Mierle K., Blanton M., Roweis S., 2010, *AJ*, 139, 1782
- Lee J., Georgitzikis E., Li Y., Lin Z., Park J., Lieberman I., Cheyns D., Jayapala M., Lambrechts A., Thijs S., Stahl R., Malinowski P. E., 2020, in *2020 IEEE International Electron Devices Meeting (IEDM)*. pp 16.5.1–16.5.4, doi:10.1109/IEDM13553.2020.9372018
- Levski D., Wány M., Choubey B., 2021, *IEEE Transactions on Circuits and Systems II: Express Briefs*, 68, 102
- Li T., DePoy D., Kessler R., Burke D., Marshall J., Wise J., Rheault J.-P., Carona D., Boada S., Prochaska T., et al., 2012, in *Ground-based and Airborne Instrumentation for Astronomy IV*. p. 84462L
- Li T., DePoy D., Marshall J. L., Nagasawa D., Carona D., Boada S., 2014, in *Ground-based and Airborne Instrumentation for Astronomy V*. p. 91476Z
- Li D., Blake C. H., Nidever D., Halverson S. P., 2017, *Publications of the Astronomical Society of the Pacific*, 130, 014501
- Lienhard F., et al., 2020, *MNRAS*, 497, 3790
- Lim P., Diaz R., Laidler V., 2015, *Astrophysics Source Code Library ascl*, 1303
- Lissauer J. J., 2007, *ApJ*, 660, L149
- Liu B., Lambrechts M., Johansen A., Pascucci I., Henning T., 2020, *A&A*, 638, A88
- Luger R., et al., 2017, *Nature Astronomy*, 1
- Luger R., Agol E., Foreman-Mackey D., Fleming D. P., Lustig-Yaeger J., Deitrick R., 2019, *The astronomical journal*, 157, 64
- Lustig-Yaeger J., Meadows V. S., Lincowski A. P., 2019, *The Astronomical Journal*, 158, 27
- Macdonald E. J. R., Cowan N. B., 2019, arXiv e-prints, p. arXiv:1908.10873
- Madhusudhan N., 2019, *Annual Review of Astronomy and Astrophysics*, 57, 617
- Mann A. W., Feiden G. A., Gaidos E., Boyajian T., Von Braun K., 2015, *The Astrophysical Journal*, 804, 64
- Marín J. C., Pozo D., Curé M., 2015, *Astronomy & Astrophysics*, 573, A41
- Martin T., Brubaker R., Dixon P., Gagliardi M.-A., Sudol T., 2005, in *Infrared Technology and*

- Applications XXXI. pp 12–21
- Mayor M., Queloz D., 1995, *Nature*, 378, 355
- McCormac J., Pollacco D., Skillen I., Faedi F., Todd I., Watson C., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 548
- Ment K., et al., 2019, *AJ*, 157, 32
- Ment K., et al., 2020, arXiv e-prints, p. arXiv:2009.13623
- Miguel Y., Cridland A., Ormel C. W., Fortney J. J., Ida S., 2020, *MNRAS*, 491, 1998
- Montgomery R., Laughlin G., 2009, *Icarus*, 202, 1
- Morley C. V., Kreidberg L., Rustamkulov Z., Robinson T., Fortney J. J., 2017, *The Astrophysical Journal*, 850, 121
- Morris B. M., Tollerud E., Sipőcz B., Deil C., Douglas S. T., Berlanga Medina J., Vyhmeister K., Smith T. R., Littlefair S., Price-Whelan A. M., Gee W. T., Jeschke E., 2018, *AJ*, 155, 128
- Muirhead P. S., et al., 2012, *The Astrophysical Journal*, 747, 144
- Murray C., Delrez L., Pedersen P., Queloz D., Gillon M., Burdanov A., Ducrot E., Garcia L., Lienhard F., Demory B., et al., 2020, *Monthly Notices of the Royal Astronomical Society*, 495, 2446
- Murray C. A., Queloz D., Gillon M., Demory B.-O., Triaud A., de Wit J., Burdanov A., Chinchilla P., Delrez L., Dransfield G., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 513, 2615
- Niraula P., et al., 2020, *AJ*, 160, 172
- Noll S., Kausch W., Barden M., Jones A., Szyszka C., Kimeswenger S., Vinther J., 2012, *Astronomy & Astrophysics*, 543, A92
- Nutzman P., Charbonneau D., 2008, *Publications of the Astronomical Society of the Pacific*, 120, 317
- Osborn J., Föhring D., Dhillon V., Wilson R., 2015, *Monthly Notices of the Royal Astronomical Society*, 452, 1707
- Otarola A., De Breuck C., Travouillon T., Matsushita S., Nyman L.-Å., Wootten A., Radford S. J., Sarazin M., Kerber F., Pérez-Beaupuits J. P., 2019, *Publications of the Astronomical Society of the Pacific*, 131, 045001
- O'Brien S. M., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 509, 6111
- Pack D., Coffman C., Santiago J., Russell R., 2022, *Authorea Preprints*
- Pain B., Hancock B. R., 2003, in *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications IV*. pp 94–103
- Payne M. J., Lodato G., 2007, *MNRAS*, 381, 1597
- Pearsall T., Hopson R., 1978, *Journal of Electronic Materials*, 7, 133
- Pecaut M. J., Mamajek E. E., 2013, *ApJS*, 208, 9

- Pedersen P. P., Murray C., Queloz D., Gillon M., Demory B., Triaud A., de Wit J., Delrez L., Dransfield G., Ducrot E., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 518, 2661
- Perryman M., Hartman J., Bakos G. Á., Lindegren L., 2014, *The Astrophysical Journal*, 797, 14
- Phillips M. W., Tremblin P., Baraffe I., Chabrier G., Allard N. F., Spiegelman F., Goyal J. M., Drummond B., Hebrard E., 2020, *Astronomy & Astrophysics*, 637, A38
- Piquette E., Edwall D., Arnold H., Chen A., Auyeung J., 2008, *Journal of Electronic Materials*, 37, 1396
- Pont F., Zucker S., Queloz D., 2006, *Monthly Notices of the Royal Astronomical Society*, 373, 231
- Pont F., Sing D. K., Gibson N. P., Aigrain S., Henry G., Husnoo N., 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 2917
- Querel R. R., Kerber F., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, *Proc. SPIE Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V*. p. 914792 ([arXiv:1408.1753](https://arxiv.org/abs/1408.1753)), [doi:10.1117/12.2055500](https://doi.org/10.1117/12.2055500)
- Quintavalla M., Spagnol M., Abe L., Carbillet M., Aristidi E., Mocchi J., Muradore R., Bonora S., 2020, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 029004
- Raymond S., 2007, in *American Astronomical Society Meeting Abstracts #210*. p. 110.03
- Rest A., Hilbert B., Leisenring J. M., Misselt K., Rieke M., Robberto M., 2016, in *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*. pp 1820–1828
- Riccardi A., Puglisi A., Grani P., Briguglio R., Esposito S., Agapito G., Biliotti V., Bonaglia M., Carbonaro L., Xompero M., et al., 2022, in *Adaptive Optics Systems VIII*. pp 53–60
- Ricker G. R., Winn J. N., Vanderspek R., Latham D. W., Bakos G. Á., Bean J. L., Bert-Thompson Z. K., Brown T. M., Buchhave L., Butler N. R., et al., 2014, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Ricker G. R., et al., 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Sagear S. A., Skinner J. N., Muirhead P. S., 2019, arXiv e-prints, [p. arXiv:1912.04286](https://arxiv.org/abs/1912.04286)
- Sandrock S., Amestica R., Sarazin M., 1999, Doc no.: VLTMAN-ESO-17440-1773
- Schoonenberg D., Liu B., Ormel C. W., Dorn C., 2019, *A&A*, 627, A149
- Sebastian D., et al., 2020a, arXiv e-prints, p. arXiv:2011.02069
- Sebastian D., Pedersen P., Murray C., Ducrot E., Garcia L., Burdanov A., Pozuelos F., Delrez L., Wells R., Dransfield G., et al., 2020b, in *Ground-based and Airborne Telescopes VIII*. pp 431–446

- Sensirion 2009, Introduction to Humidity: Basic Principles on Physics of Water Vapor, https://www.soselectronic.com/a_info/resource/c/sensirion/Sensirion_Introduction_to_Relative_Humidity_V2.pdf
- Seshadri S., Cole D., Hancock B., Ringold P., Peay C., Wrigley C., Bonati M., Brown M., Schubnell M., Rahmer G., et al., 2006, in High Energy, Optical, and Infrared Detectors for Astronomy II. p. 62760S
- Simcoe R. A., Fűrész G., Sullivan P. W., Hellickson T., Malonis A., Kasliwal M., Shectman S. A., Kollmeier J. A., Moore A., 2019, *The Astronomical Journal*, 157, 46
- Simons D., Tokunaga A., 2002, *Publications of the Astronomical Society of the Pacific*, 114, 169
- Skrutskie M. F., et al., 2006, *The Astronomical Journal*, 131, 1163
- Smette A., Sana H., Noll S., Horst H., Kausch W., Kimeswenger S., Barden M., Szyszka C., Jones A., Gallenne A., et al., 2015, *Astronomy & Astrophysics*, 576, A77
- Smith R. M., Hale D., Wizinowich P., 2014, in High Energy, Optical, and Infrared Detectors for Astronomy VI. pp 372–387
- Strausbaugh R., Jackson R., Butler N., 2018, *Publications of the Astronomical Society of the Pacific*, 130, 095001
- Sullivan P. W., Croll B., Simcoe R. A., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 1021
- Sullivan P. W., Croll B., Simcoe R. A., 2014, in High Energy, Optical, and Infrared Detectors for Astronomy VI. p. 91541F
- Sullivan P. W., et al., 2015, *The Astrophysical Journal*, 809, 77
- Tamburo P., Muirhead P. S., 2019, *PASP*, 131, 114401
- Tamburo P., Muirhead P. S., McCarthy A. M., Hart M., Gracia D., Vos J. M., Gagliuffi D. C. B., Faherty J., Theissen C., Agol E., et al., 2022, arXiv preprint arXiv:2201.01794
- Traub W. A., Cutri R., 2008, in Fischer D., Rasio F. A., Thorsett S. E., Wolszczan A., eds, *Astronomical Society of the Pacific Conference Series Vol. 398, Extreme Solar Systems*. p. 475
- Traub W. A., Oppenheimer B. R., Seager S., 2010, *Direct imaging of exoplanets*. University of Arizona Press, Tucson
- Tsiaras A., Waldmann I. P., Tinetti G., Tennyson J., Yurchenko S. N., 2019, *Nature Astronomy*, 3, 1156
- Valdés E. M., Morris B., Demory B., 2021, *Astronomy & Astrophysics*, 649
- Vanderspek R., et al., 2019, *The Astrophysical Journal Letters*, 871, L24
- Ware R. H., Fulker D. W., Stein S. A., Anderson D. N., Avery S. K., Clark R. D., Droege-meier K. K., Kuettner J. P., Minster J. B., Sorooshian S., 2000, *Bulletin of the American*

- Meteorological Society, 81, 677
- Wheatley P. J., et al., 2018a, *MNRAS*, 475, 4476
- Wheatley P. J., West R. G., Goad M. R., Jenkins J. S., Pollacco D. L., Queloz D., Rauer H., Udry S., Watson C. A., Chazelas B., et al., 2018b, *Monthly Notices of the Royal Astronomical Society*, 475, 4476
- Willey R. L., 1992, *Publications of the Astronomical Society of the Pacific*, 104, 285
- Wilson R., Jenkins C., 1996, *Monthly Notices of the Royal Astronomical Society*, 278, 39
- de Wit J., Seager S., 2013, *Science*, 342, 1473